Understanding Algebro-Geometric Quantifier Elimination: Part I, Algebraically Closed Fields of Characteristic Zero via Muchnik

Grant Olney Passmore*
grant.passmore@cl.cam.ac.uk
15 JJ Thomson Ave., University of Cambridge, CB3 0FD, UK

1 Overview: $\mathbf{ACF}_0$ and a Series of Notes

In this short expository note, we present a self-contained proof that algebraically closed fields of characteristic zero admit elimination of quantifiers over the elementary language of rings. We do this by fleshing out a method due to Muchnik[1] (and possibly Cohen[2]). This note is the first part of a series on algebro-geometric quantifier elimination: the next note shall use much of the machinery developed herein as the basis of a presentation of quantifier elimination over real closed fields, using again the method of Muchnik. Further notes in this series shall be written on quantifier elimination techniques based upon Gröbner bases, cylindrical algebraic decomposition and its variants, virtual term substitution, connected component sampling and the computation of semialgebraic Betti numbers.

2 Algebraically Closed Fields

2.1 Axiomatisation of ACF

We now present an axiomatisation of the elementary theory of algebraically closed fields. We shall then prove that this theory, once extended by fixing a characteristic (in our case, zero), admits elimination of quantifiers. This result is due, using predominantly syntactic methods, to Tarski [Tar48], though it was subsequently recast in model-theoretic terms by Abraham Robinson in his 1949 Ph.D. thesis [Rob49]. Many approaches to this result have since been developed. We have chosen to present a method due to Muchnik and discuss why we made this choice below.

[*] Most of this note has been extracted from **Section 2.1** of **Chapter 2** (Mathematical Preliminaries) of my University of Edinburgh PhD dissertation. I am very thankful to Dr. Paul B. Jackson for his careful reading and many suggestions. If you wish to cite this note, please cite my PhD dissertation (*Combined Decision Procedures for Nonlinear Arithmetics, Real and Complex*) instead.

[1] Unfortunately, it seems Muchnik never published his result. Instead, it was communicated to his students and colleagues and then appeared in two publications in Russian [Sem86, SV00] in which it was attributed to him. We have learned the method from two excellent English reconstructions of Muchnik's approach [Sch04, MO02]. For pedagogical reasons, we have presented the method in substantially more detail than the sources from which we learned it.

[2] Personal discussion with Alexander Shen.

Definition 2.1 (Axiomatisation of **ACF**)**.** Let $\mathcal{L}_{\mathcal{R}}$ be the first-order language with constants $\{0, 1\}$, relation symbol $\{=\}$, function symbols $\{+, -, *\}$, and logical symbols $\{\wedge, \vee, \neg, \forall, \exists\}$. $\mathcal{L}_{\mathcal{R}}$ is called the *language of rings.* Given an $\mathcal{L}_R$ term t, we use t^n as shorthand for $t * t * \ldots * t$ and $(t \neq 0)$ as shorthand for $\neg(t = 0)$. As no ambiguity will arise, we use $=$ as both object-theoretic and meta-theoretic equality. **ACF**, the elementary theory of algebraically closed fields, is then the $\mathcal{L}_{\mathcal{R}}$-theory defined as:

$$\mathbf{ACF} = \mathbf{F} \bigcup \mathbf{UPR},$$

where

1. $\mathbf{F}$ is an axiomatisation of the elementary $\mathcal{L}_{\mathcal{R}}$-theory of *fields*,
2. $\mathbf{U} \;=\; \{\forall a_0, \ldots, a_{n-1} \exists x (a_0 + a_1 * x + a_2 * x^2 + \ldots + a_n * x^n = 0) \mid n \in \mathbb{N} \text{ s.t. } n \geq 1\}.$

Note that $\mathbf{U}$, the collection of $\mathcal{L}_{\mathcal{R}}$-sentences stating that every univariate non-constant polynomial with coefficients in the field has a root in the field, is a countably infinite first-order axiom scheme. Note also that in the case of $\mathbf{F}$, field properties are readily expressed in $\mathcal{L}_{\mathcal{R}}$ by eliminating multiplicative inverses in favour of their defining multiplicative property (e.g., x^{-1} is replaced with a fresh variable y and the constraint $x * y = 1$ is conjoined with y quantified as is contextually appropriate). This is done so that every function symbol in $\mathcal{L}_{\mathcal{R}}$ denotes a total function.

As it stands, **ACF** is not a complete theory. This is because the *characteristic* of the field is not specified. For instance, in the absence of a specified characteristic, the ground sentence $(1 + 1 = 0)$ cannot be decided. If we specify a characteristic to obtain a theory $\mathbf{ACF}_p$ of algebraically closed fields of characteristic p, then $\mathbf{ACF}_p$ is complete, decidable, and admits elimination of quantifiers. As our interest in algebraically closed fields is chiefly motivated by making decisions over the complex numbers, we henceforth deal only with algebraically closed fields of characteristic *zero.*

2.2 $\mathbf{ACF}_0$ Admits Quantifier Elimination

We shall now prove the important theorem on quantifier elimination which will lead to the decidability of $\mathbf{ACF}_0$. Geometrically, it is essentially the theorem of Chevalley stating projections of constructible sets are themselves constructible, though proved effectively by presenting an algorithm for obtaining explicit descriptions of such projections as constructible sets. When proving this general quantifier elimination result about $\mathbf{ACF}_0$, we will often reason concretely over the complex numbers. Each time this is done, however, the reader should observe that the properties of $\mathbb{C}^n$ actually used in fact hold over *every* algebraically closed field of characteristic zero, and so our reasoning carries over to the theory $\mathbf{ACF}_0$ as a whole.

We prove this theorem by presenting the complex specialisation of a real quantifier elimination procedure due originally to Muchnik. This procedure is very elementary and is of limited practical interest. But, its simple nature makes it pedagogically superior to other more advanced methods, and it provides a vehicle for introducing many important concepts. It also has the advantage that much of the algebraic machinery we define in the context of this $\mathbf{ACF}_0$ result will be reusable in the **RCF** case in the next note where we present the Muchnik procedure in its full **RCF** form.

The result we will prove is as follows.

Theorem 2.1 (**ACF**$_0$ Quantifier Elimination)**.** *The theory of algebraically closed fields of characteristic zero admits effective elimination of quantifiers.*

We prove this by induction by showing how to eliminate a single existential quantifier from a formula with parameters. First, we introduce some algebraic machinery.

2.2.1 Complex Root Diagrams

Definition 2.2 (Labeled Row of Roots)**.** Let $p \in \mathbb{Z}[x]$. A finite binary sequence indexed by complex numbers $\alpha \in \{0,1\}^C$ s.t. $C \subset \mathbb{C}$ and $|C| < \omega$ is a *labeled row of roots* for p iff

- $p \neq 0 \Longrightarrow$ $[(\exists \zeta \in C \text{ s.t. } \alpha(\zeta) = 1) \ \wedge \ (|\{\zeta \in C \mid \alpha(\zeta) = 0\}| = |\{\zeta \in \mathbb{C} \mid p(\zeta) = 0\}|)]$,
- $\forall \zeta \in C \ (\alpha(\zeta) = 0 \iff p(\zeta) = 0)$.

We write $\mathcal{R}_{\mathbb{C}}(\alpha, p)$ to mean that α is a labeled row of roots for p.

Observation 2.1. If α is a labeled row of roots for $p \neq 0$, then α contains precisely as many 0's as there are **distinct** roots of p.

Observation 2.2. If α is a labeled row of roots for $p = 0$, then α is a row of 0's.

Definition 2.3 (Root Diagram)**.** If $P = \{p_1, \ldots, p_k\}$ is a set of polynomials in $\mathbb{Z}[x]$ then a *root diagram* for P is a labeled binary matrix M with columns labeled by members of $C \subset \mathbb{C}$ s.t. $|C| < \omega$ and rows labeled $p_1, \ldots, p_k$ s.t.

- $\forall p_i \in P \ (\mathcal{R}_{\mathbb{C}}(M(p_i), p_i))$,
- $\exists \zeta \in C \ \forall p_i \in P \ (M(p_i, \zeta) = 0 \iff p_i = 0)$,

where $M(p_i)$ is the row labeled by p_i. We call columns ζ witnessing the second property above *anti-solutions*, as they correspond to sample points within regions of $\mathbb{C}$ in which no non-zero $p_i \in P$ vanishes.

Observation 2.3. If M is a root diagram for P and M' is obtained from M by some combination of

- permuting the columns of M,
- adding or removing some (but not all) anti-solution columns from M,
- choosing a different label for the anti-solution column of M,

then M' is still a root diagram for P.

Lemma 2.1. *Up to the modifications described in* ***Observation 2.3****, a root diagram M for P is uniquely determined.*

Proof. This is immediate, as every root diagram for P must contain a *minimal core* consisting of columns labeled by every root of each $p \in P$, together with at least one anti-solution column for P. It is clear that none of these columns may be removed from M while maintaining its status as a root diagram for P, though a different label may be used for the anti-solution column. Thus, given two distinct root diagrams M, M' for P, M and M' may only differ by the operations given in **Observation 2.3**. □

Given this relative uniqueness, we will now write $\mathcal{D}(P)$ ("**the** root diagram for P") to mean a canonically chosen root diagram for P.

2.2.2 Muchnik Sets and Sequences

We now present the concepts of Muchnik sets and sequences which we will use to compute root diagrams for sets of polynomials.

Let $\mathbb{A}$ be a unique factorisation domain (UFD) and let $p \in \mathbb{A}[x]$ s.t.

$$(p = 0) \quad \vee \quad (p = \sum_{j=0}^{d} c_j x^j \quad \wedge \quad c_d \neq 0).$$

Definition 2.4 (Polynomial Degree)**.**

$$deg(p) = \begin{cases} d & \text{if } p \neq 0, \\ -\infty & \text{if } p = 0. \end{cases}$$

Definition 2.5 (Polynomial Tail)**.**

$$\tau(p) = \begin{cases} \sum_{j=0}^{d-1} c_j x^j & \text{if } p \neq 0, \\ 0 & \text{if } p = 0. \end{cases}$$

We now face a problem: we will need to perform division upon pairs of polynomials in $\mathbb{A}[x]$, where $\mathbb{A}$ is some non-Euclidean UFD such as $\mathbb{Z}[y_1, y_2]$. Recall that if $\mathbb{A}$ is a UFD, then $\mathbb{A}[x]$ is as well. We will thus make use of *polynomial pseudo-division* as the unique *pseudo-remainder* it computes for pairs of polynomials will be sufficient for inductively obtaining root diagrams.

Let $q \in \mathbb{A}[x]$ s.t.

$$q = \sum_{j=0}^{e} b_j x^j \quad \wedge \quad q \neq 0 \quad \wedge \quad deg(q) = e \leq d.$$

Definition 2.6 (Polynomial Pseudo-remainder)**.** Given p, q as specified above, *polynomial pseudo-division* of p by q will compute unique $h, r \in \mathbb{A}[x]$ s.t.

$$b_e^{d-e+1} p = hq + r \quad \wedge \quad deg(r) < e.$$

We refer to $r = rem(p, q)$ as the *pseudo-remainder*[3] of p by q.

Let $\mathbb{M} \neq \emptyset \subset \mathbb{A}[x]$ s.t. $|\mathbb{M}| < \omega$.

Definition 2.7 (Muchnik Set)**.** We say $\mathbb{M}$ is a *Muchnik set* iff

1. $p \in \mathbb{M} \implies \tau(p) \in \mathbb{M}$,
2. $p \in \mathbb{M} \implies \frac{\partial p}{\partial x} \in \mathbb{M}$,
3. $p, q \neq 0 \in \mathbb{M} \wedge deg(q) \leq deg(p) \implies rem(p, q) \in \mathbb{M}$.

Observation 2.4. If $\mathbb{M}$ is Muchnik, then $0 \in \mathbb{M}$.

Definition 2.8 (Muchnik Closure)**.** Given $\mathbb{M} \subset \mathbb{A}[x]$, let $\mathbb{M}^*$ be the smallest Muchnik set containing $\mathbb{M}$. We say $\mathbb{M}^*$ is the *Muchnik closure* of $\mathbb{M}$.

Observation 2.5. If $\mathbb{M} \subset \mathbb{A}$ s.t. $|\mathbb{M}| < \omega$ and $0 \in \mathbb{M}$ then $\mathbb{M}$ is Muchnik.

Definition 2.9 (Constant Fragment)**.** A *constant* is a polynomial $p \in \mathbb{A}$. If $\mathbb{M}$ is Muchnik then let $\mathbb{M}_0$ be $\mathbb{M} \cap \mathbb{A}$ – the *constant fragment* of $\mathbb{M}$ – which is also Muchnik.

Lemma 2.2 (Finiteness)**.** *Let $\mathbb{M} \subset \mathbb{A}[x]$ s.t. $|\mathbb{M}| < \omega$. Then $|\mathbb{M}^*| < \omega$.*

Proof. Immediate as each of the three operations placing polynomials into $\mathbb{M}^*$ are *strictly* degree reducing. □

All Muchnik sets we encounter will be finite and we henceforth omit the explicit assumption. We now introduce the notion of a *Muchnik sequence* and prove its important substructural property.

Definition 2.10 (Muchnik Sequence)**.** Let $\mathbb{M}$ be Muchnik. Then any sequence $\beta \in \mathbb{M}^{|\mathbb{M}|}$ is a *Muchnik sequence* iff

$$\forall 1 \leq i < |\mathbb{M}| \ (deg(\beta(i)) \leq deg(\beta(i+1)).$$

Observe that as $deg(0) = -\infty$, a Muchnik sequence always begins with 0.

The following lemma will be important for inductively extending Muchnik sequences.

Lemma 2.3 (Muchnik Subsequence)**.** *Let β be a Muchnik sequence. Then, any non-empty initial segment β' of β is a Muchnik sequence.*

Proof. By the definition of Muchnik, we must show β' is closed under the operations of *tail*, *partial differentiation* and *pseudo-remainder*. But this is immediate by the fact that these three closure operations are strictly degree reducing, and in a Muchnik sequence β (and hence any initial segment), the polynomials must be ordered so that $(deg(\beta(i)) \leq deg(\beta(i+1))$. □

[3] Knuth gives an excellent presentation of polynomial pseudo-division on pp 425-428 of [Knu97]. Sufficient background on UFDs is given on pp 422-424. To reiterate the elementary nature of the quantifier elimination procedure we are presenting, though, let us note that computing pseudo-remainders is conceptually very simple: In fact, over a UFD such as our $\mathbb{A}[x]$ above, one can compute the pseudo-remainder of p by q by first multiplying p by b_e^{d-e+1} and then performing standard polynomial division (that is, the division algorithm for polynomials over a field) between the product $(b_e^{d-e+1}p)$ and q.

2.2.3 Elimination of a Single Existential Quantifier

With the Muchnik machinery in hand, let us now discuss our strategy for eliminating an existential quantifier. Given a set of polynomials $P = \{p_1, \ldots, p_k\} \subset \mathbb{Z}[y_1, \ldots, y_n][x]$ (e.g., $\mathbb{A} = \mathbb{Z}[\vec{y}]$ is our ambient ring of coefficients for polynomials in x), let $P(\vec{c}) = \{p_1(\vec{c}), \ldots, p_k(\vec{c})\} \subset \mathbb{Z}[x]$ for any $\vec{c} \in \mathbb{C}^n$. Recall that P^* is the Muchnik closure of P with $P^*_0 = (P^* \cap \mathbb{Z}[\vec{y}])$ its subset of constants w.r.t. x. We will present an algorithm for computing all possible root diagrams for P^*. This will be done in such a way that each root diagram for P^* will be derived (and uniquely determined) from a root diagram for P^*_0. By viewing P^* as a collection of univariate polynomials in x, the set of all root diagrams for P^* arises by considering the specialisations of $\mathbb{A} = \mathbb{Z}[\vec{y}]$ to points $\vec{c} \in \mathbb{C}^n$. Let us introduce the notion of an *extended root diagram* to formalise this process.

Definition 2.11 (Extended Root Diagram)**.** Let $P = \{p_1, \ldots, p_k\}$ be a set of polynomials in $\mathbb{Z}[y_1, \ldots, y_n][x]$. Then, an *extended root diagram* for P w.r.t. $\vec{c} \in \mathbb{C}^n$ is a labeled binary matrix $M \in \{0,1\}^{P \times C}$ with $C \subset \mathbb{C}$ and $|C| < \omega$ s.t.

- $\forall p_i \in P(\mathcal{R}_{\mathbb{C}}(M(p_i), p_i(\vec{c})))$,
- $\exists \zeta \in C \; \forall p_i \in P \; (M(p_i, \zeta) = 0 \iff p_i(\vec{c}) = 0)$,

where $M(p_i)$ is identified with a function in $\{0,1\}^C$ in the obvious way.

Intuitively, if M is an extended root diagram for P w.r.t. $\vec{c}$, then M is in principle a normal root diagram for $P(\vec{c}) \subset \mathbb{Z}[x]$, but constructed so that row the labels of M hold a record of the polynomials $p_i \in \mathbb{Z}[\vec{y}][x]$ from which the $p_i(\vec{c}) \in \mathbb{Z}[x]$ were derived. Thus, one may see such an extended root diagram as what happens when one computes a root diagram for a specialisation of P to $P(\vec{c})$ and then *forgets* the specialisation of the row labels.

It is easy to see that the same uniqueness properties that hold for root diagrams (à la **Lemma 2.1**) also hold for extended root diagrams.

Finally, it turns out that all of the information needed to perform quantifier elimination can actually be obtained from a variant of extended root diagrams in which neither the columns nor rows carry explicit labels.

Definition 2.12 (Unlabeled Extended Root Diagram)**.** Let $P = \{p_1, \ldots, p_k\}$ be a set of polynomials in $\mathbb{Z}[y_1, \ldots, y_n][x]$. Then, an *unlabeled extended root diagram* for P w.r.t. $\vec{c} \in \mathbb{C}^n$ is an $(k \times m)$ binary matrix M obtained from an extended root diagram M' for P w.r.t. $\vec{c}$ by forgetting the row and column labels of M'. That is, M is simply the underlying binary matrix of M'. Given that P is ordered, we will use $M(p_i)$ and "the row corresponding to p_i" to mean the ith row of M, even though M is formally simply a matrix (without explicit polynomial row labels). If M is a matrix consisting of a single column, then we will use $M(p_i)$ to mean the value of the single entry in the row corresponding to p_i.

From now on, when we say "**the** unlabeled extended root diagram for P w.r.t. $\vec{c}$," we will mean a *canonically chosen* unlabeled extended root diagram for P w.r.t. $\vec{c}$. We will write $\mathcal{D}^*(P, \vec{c})$ to mean the unlabeled extended root diagram for P w.r.t. $\vec{c}$. Similarly, if β is a Muchnik sequence of polynomials in $\mathbb{Z}[\vec{y}][x]$, then $\mathcal{D}^*(\beta, \vec{c})$ will be the unlabeled extended root diagram for the underlying set of β w.r.t. $\vec{c}$.

Let

$$\mathbb{D} = \{\mathcal{D}^*(P^*, \vec{c}) \mid \vec{c} \in \mathbb{C}^n\},$$

and

$$\mathbb{D}_0 = \{\mathcal{D}^*(P_0^*, \vec{c}) \mid \vec{c} \in \mathbb{C}^n\}.$$

The key observations are:

1. Both sets $\mathbb{D}$ and $\mathbb{D}_0$ are *finite* (and every member of $\mathbb{D}_0$ consists of a *single-column* binary matrix), and
2. Given any $\vec{c} \in \mathbb{C}^n$, the unlabeled extended root diagram for P^* w.r.t. $\vec{c}$ (i.e., $\mathcal{D}^*(P^*, \vec{c}) \in \mathbb{D}$) may be obtained from the unlabeled extended root diagram for P_0^* w.r.t. $\vec{c}$ (i.e., $\mathcal{D}^*(P_0^*, \vec{c}) \in \mathbb{D}_0$).

This derivation of $\mathcal{D}^*(P^*, \vec{c})$ from $\mathcal{D}^*(P_0^*, \vec{c})$, which we call *diagram lifting*, will be done by an algorithm $\mathcal{A}_{\mathbb{C}}$ with the following universal property:

$$\forall \vec{c} \in \mathbb{C}^n (\mathcal{A}_{\mathbb{C}}(\mathcal{D}^*(P_0^*, \vec{c})) = \mathcal{D}^*(P^*, \vec{c})).$$

Let us now see how this machinery can be applied.

Consider a quantifier-free formula

$$\varphi(\vec{y}, x) = (\bigwedge_{i=1}^{k} (p_i \ \sigma_{p_i} \ 0)) \qquad \text{with} \qquad (\sigma_{p_i} \in \{=, \neq\}).$$

Let $Z(\sigma_{p_i})$ hold iff σ_{p_i} is '='. Say that the unlabeled extended root diagram $\mathfrak{C}$ for P^* is φ-*compatible* iff there exists a column j in $\mathfrak{C}$ s.t.

$$\forall 1 \leq i \leq k (\mathfrak{C}(p_i, j) = 0 \iff Z(\sigma_{p_i})).$$

Given any $\vec{c} \in \mathbb{C}^n$, we will then have

$$\exists x(\varphi(\vec{c}, x)) \quad \iff \quad \mathcal{A}_{\mathbb{C}}(\mathcal{D}^*(P_0^*, \vec{c})) \text{ is } \varphi\text{-compatible}.$$

Let k_0 be s.t. $P_0^* = \{p_1, \ldots, p_{k_0}\}$. Observe that $|\mathbb{D}_0| \leq 2^{k_0}$. That is, there are at most 2^{k_0} possible unlabeled extended root diagrams which could arise in the process of specialising P_0^* to any point in $\mathbb{C}^n$. Let $\mathcal{M}_0$ be the set of all $(k_0 \times 1)$ binary matrices. Observe that $\mathbb{D}_0 \subseteq \mathcal{M}_0$. Then, conditions $\psi(\vec{y})$ upon $\vec{y}$ s.t.

$$\forall \vec{y}(\psi(\vec{y}) \iff \exists x(\varphi(\vec{y}, x)))$$

are given by

$$\psi(\vec{y}) = \bigvee_{d_0 \in \mathcal{Q}_\varphi} \mathcal{E}_{\mathcal{L}_R}(d_0),$$

where

$$\mathcal{Q}_\varphi = \{d_0 \in \mathcal{M}_0 \mid \mathcal{A}_\mathbb{C}(d_0) \text{ is } \varphi - \text{compatible}\}$$

and

$$\mathcal{E}_{\mathcal{L}_R}(d_0) = \bigwedge_{q \in P_0^*} (q \ \odot_{d_0(q)} \ 0)$$

s.t.

$$\odot_{d_0(q)} = \begin{cases} \text{`='} & \text{if } d_0(q) = 0, \\ \text{`}\neq\text{'} & \text{if } d_0(q) = 1. \end{cases}$$

Now, there is one aspect of the above derivation of $\psi(\vec{y})$ which is counterintuitive. Naively, one would expect $\mathcal{Q}_\varphi$ to be defined as the set S as follows:

$$\mathcal{S} = \{d_0 \in \mathbb{D}_0 \mid \mathcal{A}_\mathbb{C}(d_0) \text{ is } \varphi - \text{compatible}\}.$$

The issue with this definition is that in practice, we will not a priori know if a given binary matrix $d_0 \in \mathcal{M}_0$ is actually the unlabeled extended root diagram for P_0^* w.r.t. any $\vec{c} \in \mathbb{C}^n$. That is, given some $d_0 \in \mathcal{M}_0$, we will not know in advance if it is a member of $\mathbb{D}_0$ or not. Thankfully, it will will not actually matter what our lifting algorithm gives as the value of $\mathcal{A}_\mathbb{C}(d_0)$ when $d_0 \in (\mathcal{M}_0 \setminus \mathbb{D}_0)$. Let us see why this is so.

Lemma 2.4. *Let $d_0 \in (\mathcal{M}_0 \setminus \mathbb{D}_0)$. Then, based upon our construction of $\psi(\vec{y})$ above, it does not matter which $(k \times m)$ binary matrix our lifting algorithm constructs as the value $\mathcal{A}_\mathbb{C}(d_0)$.*

Proof. Assume we are eliminating $\exists x$ from $\exists x\varphi(\vec{y}, x)$ to obtain a quantifier-free equivalent formula $\psi(\vec{y})$ as above. Consider $d_0 \in (\mathcal{M}_0 \setminus \mathbb{D}_0)$. That is, d_0 is a $k_0 \times 1$ binary matrix that is not realisable as the unlabeled extended root diagram for P_0^* w.r.t. any $\vec{c} \in \mathbb{C}^n$. Now, let us apply $\mathcal{A}_\mathbb{C}$ to lift d_0 and obtain an $(k \times m)$ binary matrix $d = \mathcal{A}_\mathbb{C}(d_0)$. We have two cases: either d is φ-compatible, or it is not. If d is not φ-compatible, then d_0 will not contribute at all to our construction of $\psi(\vec{y})$ and so the value of d does not matter. On the other hand, assume that d is φ-compatible. Then, d will contribute to our construction of $\psi(\vec{y})$ in the following way: $\mathcal{E}_{\mathcal{L}_R}(d_0)$ will be present as a disjunct in $\psi(\vec{y})$. But, since d_0 is not realisable as the unlabeled extended root diagram of P_0^* w.r.t. any $\vec{c} \in \mathbb{C}^n$, this means that

$$\langle \mathbb{C}, +, -, *, 0, 1 \rangle \models \forall \vec{y}(\mathcal{E}_{\mathcal{L}_R}(d_0) \iff 0 = 1).$$

Thus, $\mathcal{E}_{\mathcal{L}_R}(d_0)$ is only contributing a contradictory conjunction as a disjunct in our formula $\psi(\vec{y})$, which means $\psi(\vec{y})$ is logically equivalent to $\psi(\vec{y})$ with $\mathcal{E}_{\mathcal{L}_R}(d_0)$ removed. So, it is indeed the case that if d_0 is not realisable as the unlabeled extended root diagram of P_0^* w.r.t. any $\vec{c} \in \mathbb{C}^n$, then it does not matter which $(k \times m)$ binary matrix our lifting algorithm constructs as the value $\mathcal{A}_{\mathbb{C}}(d_0)$. □

This fact permits us a simple approach to constructing $\psi(\vec{y})$: We will generate *all* 2^{k_0} possible $(k_0 \times 1)$ binary matrices as *candidate* unlabeled extended root diagrams for P_0^*, lift each of them, and construct $\psi(\vec{y})$ as a disjunction of conjuncts corresponding to the φ-compatible lifted candidates.

Thus, once we have exhibited an algorithm $\mathcal{A}_{\mathbb{C}}$ for *diagram lifting*, we will have proved the following theorem establishing, by induction, that $\mathbf{ACF}_0$ admits elimination of quantifiers.

Theorem 2.2 (Projective Closure of Definability)**.** *Given any quantifier-free $\mathcal{L}_R$-formula $\varphi(\vec{y}, x)$ there exists a quantifier-free $\mathcal{L}_R$-formula $\psi(\vec{y})$ s.t.*

$$\boldsymbol{ACF}_0 \models \forall \vec{y} \ (\exists x \varphi(\vec{y}, x) \iff \psi(\vec{y}))\,.$$

Moreover, $\psi(\vec{y})$ is effectively computable from $\varphi(x, \vec{y})$.

Let us now finish the proof by constructing such an $\mathcal{A}_{\mathbb{C}}$. Again, let $P \subset \mathbb{Z}[\vec{y}][x]$ with P^* its Muchnik closure. Recall that P_0^* is the collection of constants in P^* w.r.t. x. Let $\beta = \langle p_1, \ldots, p_k \rangle$ be a Muchnik sequence for P^* and $\beta_0 = \langle p_1, \ldots, p_{k_0} \rangle$ a Muchnik sequence for P_0^*. We will construct an algorithm $\mathcal{A}_{\mathbb{C}}$ which will map a $(k_0 \times 1)$ binary matrix d_0 to a $(k \times m)$ binary matrix d s.t. *if* d_0 is the unlabeled extended root diagram for β_0 w.r.t. $\vec{c} \in \mathbb{C}^n$, *then* $\mathcal{A}_{\mathbb{C}}(d_0)$ will be the unlabeled extended root diagram for β w.r.t. $\vec{c}$. By **Lemma 2.4**, it will not matter what $\mathcal{A}_{\mathbb{C}}$ returns if its input is not realisable as the unlabeled extended root diagram for β_0 w.r.t. any $\vec{c} \in \mathbb{C}^n$. Recall that by **Lemma 2.3**, every subsequence $\beta \downarrow$ of β is Muchnik. Given an unlabeled extended root diagram for β_0 w.r.t. $\vec{c}$, we will use this property of Muchnik sequences to build an extended root diagram for β w.r.t. $\vec{c}$ inductively, by building one for each of its subsequences $\beta \downarrow$.

Lemma 2.5 ($\mathbf{ACF}_0$ Single-Step Diagram Lifting Algorithm)**.** *Let $\beta \in \mathbb{Z}[\vec{y}][x]^k$ be a Muchnik sequence. Let $\mathfrak{C}$ be a $(k \times m)$ binary matrix (a* candidate *unlabeled extended root diagram for β). Let $p \in \mathbb{Z}[\vec{y}][x]$ be a non-constant polynomial w.r.t. x s.t. $\beta^+ = \langle \beta(1), \ldots, \beta(k), p \rangle$ is Muchnik. Then, there is an algorithm $\mathcal{A}_{\mathbb{C}}$ sending $\mathfrak{C} \mapsto \mathfrak{C}^+$ s.t.* if $\mathfrak{C}$ *is the unlabeled extended root diagram for β w.r.t. $\vec{c} \in \mathbb{C}^n$,* then $\mathfrak{C}^+$ *is the unlabeled extended root diagram for β^+ w.r.t. the same $\vec{c}$. That is,*

$$\mathfrak{C} = \mathcal{D}^*(\beta, \vec{c}) \implies \mathfrak{C}^+ = \mathcal{D}^*(\beta^+, \vec{c}).$$

*As established by **Lemma 2.3**, if $\mathfrak{C}$ is in fact not an unlabeled extended root diagram for β w.r.t. any $\vec{c} \in \mathbb{C}^n$, then it is of no consequence which $(k + 1 \times m')$ binary matrix this algorithm returns. In certain cases, this algorithm may be able to "short-circuit" its processing by recognising that the* *candidate $\mathfrak{C}$ is not the unlabeled extended root diagram for β w.r.t. any $\vec{c} \in \mathbb{C}^n$. In these cases, the algorithm will return a special value $\perp$ to signify this.*

In summary, letting $\mathcal{M}(\beta)$ denote the collection of candidate extended root diagrams for β, we will present an algorithm $\mathcal{A}_{\mathbb{C}}$ with the following properties:

$$\mathcal{A}_{\mathbb{C}} \ : \ \mathfrak{M}(\beta) \to \mathfrak{M}(\beta^+)$$

$$\forall \vec{c} \in \mathbb{C}^n \left(\mathcal{A}_{\mathbb{C}}(\mathcal{D}^*(\beta, \vec{c})) = \mathcal{D}^*(\beta^+, \vec{c}) \right)$$

$$M \in \mathfrak{M}(\beta) \ \wedge \ \mathcal{A}_{\mathbb{C}}(M) = \perp \quad \Longrightarrow \quad \neg \exists \vec{c} \in \mathbb{C}^n (M = \mathcal{D}^*(\beta, \vec{c})).$$

Proof. Let $deg(p) = d$ and $\alpha \in \mathbb{Z}[y_1, \ldots, y_n]$ be the highest degree coefficient of p (both w.r.t. x). Recall that as Muchnik sets are closed under partial differentiation, $d!\alpha$ appears in β and thus corresponds to a row in $\mathfrak{C}$. Let r be this row. If r is not a constant row, then $\mathfrak{C}$ cannot be an unlabeled extended root diagram for β, so we return $\perp$. Otherwise, we have two cases:

[Case I: $r = \vec{0}$] In this case, the root conditions for p are equivalent to those for $0 * x^d + \tau(p) = \tau(p) \in \mathbb{Z}[\vec{y}][x]$. But, note that $deg(\tau(p)) < d$ w.r.t. x. Thus, by definition of Muchnik sequence, we have that the row of roots for $\tau(p)$ already exists in $\mathfrak{C}$. Hence we may simply copy this row of roots as the row for p and we are done.

[Case II: $r = \vec{1}$] If $\mathfrak{C} = \mathcal{D}^*(\beta, \vec{c})$ for $\vec{c} \in \mathbb{C}^n$, then $r = \vec{1}$ yields that $\alpha(\vec{c}) \neq 0$. To extend $\mathfrak{C}$ to $\mathfrak{C}^+$ by taking into account p, we must meet the following requirements:

- Any root of p not already represented by a column of $\mathfrak{C}$ must be represented by a column of $\mathfrak{C}^+$ (columns must be added for these roots),
- The nullity of every polynomial represented by a row of $\mathfrak{C}$ at each new root of p must be determined,
- The nullity of p at all points represented by columns of $\mathfrak{C}$ must be determined (p will be 0 in every column of $\mathfrak{C}^+$ which is not present in $\mathfrak{C}$, as these are roots of p),
- The existence of an *anti-solution* column must be maintained.

Let ζ be a column of $\mathfrak{C}$. We have two cases:

[Case II.a: ζ is not an anti-solution column] So, the column ζ contains a 0 which does not come from a 0-row. Let $q \in B$ be of minimal degree ($deg(q) = e$) s.t. $\mathfrak{C}(q, \zeta) = 0$ and $\mathfrak{C}(q) \neq \vec{0}$. If $\mathfrak{C}$ is an extended root diagram for β w.r.t. $\vec{c} \in \mathbb{C}^n$, then this means that $q(\vec{c}, \zeta) = 0$ and $q(\vec{c}) \in \mathbb{Z}[x]$ is not identically zero. Let $\gamma \in \mathbb{Z}[\vec{y}]$ be the highest degree coefficient of q s.t. $q = \gamma x^e + \tau(q)$. Observe by definition of Muchnik sequence that $e!\gamma$ corresponds to a row in $\mathfrak{C}$. If $\mathfrak{C}(e!\gamma)$ is not a constant row, then $\mathfrak{C}$ cannot be an unlabeled extended root diagram for β and we return $\perp$. Thus we assume $\mathfrak{C}(e!\gamma)$ is a constant row.

Let us now observe that if $\mathfrak{C}(e!\gamma) = \vec{0}$, then $\mathfrak{C}$ cannot be an unlabeled extended root diagram for β. If $\mathfrak{C}(e!\gamma) = \vec{0}$, then we have that the root conditions of q are equivalent to those of $0 + \tau(q) = \tau(q) \in \mathbb{Z}[\vec{y}][x]$. So, $\tau(q)(\vec{c}, \zeta) = 0$. But then by assumption that q was of minimal degree with $q(\vec{c}, \zeta) = 0$ and $\mathfrak{C}(q) \neq \vec{0}$, we have that $\mathfrak{C}(\tau(q))$ must be a 0-row. But, then $\mathfrak{C}(q)$ would be a 0-row as well, which is a contradiction. So, if $\mathfrak{C}(e!\gamma) = \vec{0}$ then $\mathfrak{C}$ cannot be an unlabeled extended root diagram for β and we return $\perp$. Thus we assume $\mathfrak{C}(e!\gamma) = \vec{1}$.

Let $r = rem(p, q)$ be the pseudo-remainder of p by q. So, $\gamma^{d-e+1}p = hq + r$ for some $h, r \in \mathbb{Z}[\vec{y}][x]$ s.t. $deg(r) \leq e - 1$. As $\mathfrak{C}(q, \zeta) = 0$, if $\mathfrak{C}$ is an unlabeled extended root diagram for β w.r.t. $\vec{c} \in \mathbb{C}^n$, then $p(\vec{c}) = r(\vec{c})$. By definition of Muchnik sequence, $r \in B$ so r is the label of a row in $\mathfrak{C}$. Therefore we simply set $\mathfrak{C}^+(p, \zeta) = \mathfrak{C}(r, \zeta)$ and this case is complete. Observe that this process allows us to determine the nullity of p for every column of $\mathfrak{C}$ which is a root of some non-constant polynomial in B.

[Case II.b: ζ is an anti-solution column] So, the ζ only has 0's coming from 0-rows. We have two requirements left to meet:

- We must add columns to $\mathfrak{C}^+$ corresponding to the roots of p which are not represented by columns of $\mathfrak{C}$ and determine the nullity of each polynomial corresponding to a row of $\mathfrak{C}$ at these new roots of p,
- We must guarantee the existence of an anti-solution column for $\mathfrak{C}^+$.

As $p \in \mathbb{Z}[\vec{y}][x]$ is non-constant by assumption, we can extend ζ to be an anti-solution column of $\mathfrak{C}^+$ by simply setting $\mathfrak{C}^+(p, \zeta) = 1$. Thus, the requirements of both determining the nullity of p at every column label of $\mathfrak{C}$ and guaranteeing the existence of an anti-solution column for $\mathfrak{C}^+$ have been met.

It now remains to add columns to $\mathfrak{C}^+$ representing the roots of p not already represented by columns of $\mathfrak{C}$ and to determine the nullity of every polynomial represented by a row of $\mathfrak{C}$ at these new roots. Observe that any root of p not represented by a column of $\mathfrak{C}$ must have multiplicity 1, as otherwise it would be a root of $\frac{\partial p}{\partial x}$ and hence would be already represented by a column of $\mathfrak{C}$. By the Fundamental Theorem of Algebra, we may determine the number of new roots of p to add to $\mathfrak{C}^+$ as $(deg(p) - \#\kappa)$ where $\#\kappa$ is the number of roots of p already represented in $\mathfrak{C}$ counted with multiplicity. To determine $\#\kappa$, it will suffice to determine the multiplicity $m(\xi)$ of every root ξ of p appearing in $\mathfrak{C}$. To compute $m(\xi)$, we examine the successive derivatives $\frac{\partial p}{\partial x}, \frac{\partial^2 p}{\partial x^2}, \ldots$ and check the nullity of each derivative at ξ. By definition of Muchnik sequence, all such derivatives will correspond to rows of $\mathfrak{C}$, and thus this information may be computed from $\mathfrak{C}$. Then, $m(\xi) = j$ where j is the least power s.t. $\mathfrak{C}(\frac{\partial^j p}{\partial x^j}, \xi) = 1$. Thus, $\#\kappa = \sum_{\xi \in \chi} m(\xi)$ where χ is the collection of points represented by columns of $\mathfrak{C}$ s.t. $\mathfrak{C}(p, \xi) = 0$. Now, we add $(deg(p) - \#\kappa)$ new columns to $\mathfrak{C}^+$ with 0's in their bottom row (corresponding to p), 0's in their rows corresponding to 0-rows, and 1's in all other rows. As we have met our final requirements, this completes our proof. □

As it is easy to see all properties of $\mathbb{C}$ used in the above construction hold over every $\mathcal{F}$ s.t. $\mathcal{F} \models$ **ACF**$_0$, **Theorem 2.2** follows by induction along β. That is, we may always eliminate a single existential quantifier. From this result, **Theorem 2.1** follows by induction by placing $\mathcal{L}_{\mathcal{R}}$ formulae in prenex normal form and successively eliminating the innermost existential quantifier until no quantified variables remain.

References

[Knu97] Donald E. Knuth. *The Art of Computer Programming, Volume 2 (3rd ed.): Seminumerical Algorithms*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1997.

[MO02] Christian Michaux and Adem Ozturk. Quantifier Elimination Following Muchnik. Universite de Mons-Hainaut Preprint Series (#10), April 2002.

[Rob49] Abraham Robinson. *The Metamathematics of Algebraic Systems*. PhD thesis, University of London, 1949.

[Sch04] Hans Schoutens. Muchnik's Proof of Tarski-Seidenberg. Short note, February 2004.

[Sem86] A. Semenov. Decision Procedures for Logical Theories. *Cybernetics and Computer Technology*, 2:134–146, 1986.

[SV00] A. Shen and N. K. Vereshchaigin. Languages and Calculi. *Moscow Center for Continuous Mathematical Education*, 2000.

[Tar48] Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. RAND Corporation, 1948.