Computable Procedures for Fields *

Russell Miller[1,2]

[1] Queens College, 65-30 Kissena Blvd., Queens, NY 11367, U.S.A.
[2] C.U.N.Y. Graduate Center, 365 Fifth Avenue, New York, NY 10016 U.S.A.
Russell.Miller@qc.cuny.edu

Abstract. This tutorial will introduce listeners to many questions that can be asked about computable processes on fields, and will present the answers that are known, sometimes with proofs. This is not original work. The questions in greatest focus here include decision procedures for the existence of roots of polynomials in specific fields, for the irreducibility of polynomials over those fields, and for transcendence of specific elements over the prime subfield. Several of these questions are related to the construction of algebraic closures, making Rabin's Theorem prominent.

Keywords: computability, computable structure theory, factorization, field, Hilbert's Tenth Problem, irreducibility, polynomials, Rabin's Theorem, root set, splitting set.

1 Introduction

The classic problems of factoring polynomials and finding their roots arose long before any formal notion of decidability existed. In Greece, geometric problems led to it, such as finding side lengths in right triangles: the Greeks knew that $\sqrt{2}$ was irrational, which is to say, that $X^2 - 2$ does not factor over $\mathbb{Q}$. In India, the sixth-century mathematician Brahmagupta investigated integer solutions to what came to be known in the West as the Pell equations $X^2 - dY^2 = 1$. (Not only did Pell trail Brahmagupta by a millenium in this study, but he was also not even the first on his own continent to consider these equations, being preceded in this by Fermat.) In 1900, the tenth of the problems posed by Hilbert for the new century was to find a method of determining whether an arbitrary diophantine equation $f = 0$, with $f \in \mathbb{Z}[X_1, X_2, \ldots]$, has a solution in integers.

Algorithms for answering various of these questions had been discovered over the centuries, of course, often independently in different cultures. With Alan Turing's 1936 definition of an algorithm – using what came to be known as a *Turing machine* – questions of decidability could be studied more rigorously. Not only could one ask whether classical algorithms really could be implemented on a Turing machine – for the most part, the answers were affirmative, although some of these algorithms require prohibitive time and memory resources – but in certain cases, one could now prove that no algorithm at all could succeed.

* The author was partially supported by Grant #581896 from the Simons Foundation and by the City University of New York PSC-CUNY Research Award Program.

Hilbert appears not to have anticipated the possibility that his Tenth Problem would be resolved in this way, but indeed, in 1970, Matiyasevich [15] completed work by Davis, Putnam, and Robinson [2], proving that no Turing machine at all can accomplish the task demanded by Hilbert.

In this tutorial we will examine both decidability and undecidability results within these topics. We will focus on fields of characteristic 0, and on results in pure computability, rather than considering results from theoretical computer science and other disciplines. In doing so, we omit a significant and intriguing body of knowledge, but paradoxically, the time and space constraints on this abstract and tutorial preclude consideration of time and space constraints on the algorithms.

Apart from Section 6, we will restrict ourselves to countable (infinite) fields, the most natural subjects for our questions. A *presentation* of such a field is a first-order structure F with domain $\{x_0, x_1, \ldots\}$, in the signature of rings, that is isomorphic to F. (One can use $\mathbb{N}$ itself as the domain, but when studying fields this would create much confusion.) The *atomic diagram* $\Delta(F)$ of F essentially consists of the addition and multiplication tables for F, coded using a Gödel coding so that we may regard $\Delta(F)$ as a subset of $\mathbb{N}$. If $\Delta(F)$ is decidable, then F is a *computable presentation* of the field. A single field will have many presentations; some may be computable, but not all, and many fields have no computable presentation at all. In order to consider all fields, we often give ourselves the set $\Delta(F)$ as an oracle. The most basic countable field is the field $\mathbb{Q}$ of rational numbers, and we fix a single computable presentation of it to be used hereafter. (Often we will conflate the isomorphism type, the presentation, and the atomic diagram of a field, when it seems safe to do so.)

The work described here is not original in this article, although certain standard facts may go uncited. Historically important sources on computable fields includes work by van der Waerden [27], Fröhlich and Shepherdson [8], Rabin [23], Ershov [6], Metakides and Nerode [16], Fried and Jarden [7], and Stoltenberg-Hansen and Tucker [26], while [17] gives a helpful basic introduction to these topics.

2 Rabin's Theorem

The natural starting point is the simplest case: polynomials in a single variable X, over $\mathbb{Q}$. Systematic algorithms factoring polynomials in the polynomial ring $\mathbb{Q}[X]$ and finding their roots date back at least as far as Kronecker [14], who began with $\mathbb{Z}[X]$ and moved on to $\mathbb{Q}[X]$ and then to field extensions of $\mathbb{Q}$. It is worthwhile to examine his work: a good modern presentation appears in Edwards's *Galois Theory* [5, §§55-57]. In fact, determining whether an $f \in \mathbb{Z}[X]$ has a root x in $\mathbb{Z}$ is fairly trivial: x itself, if it exists, must divide the constant coefficient c_0, as it divides every other term in $0 = c_0 + c_1 x + \cdots + c_d x^d = f(x)$. If $c_0 = 0$, then 0 is a root; otherwise this leaves only finitely many possible values $x_1, \ldots, x_n$ for x, each of which can be tested by computing $f(x_i)$. A similar procedure applies when $f \in \mathbb{Q}[X]$: after one clears the denominators in

the coefficients, each prime power that divides the denominator of the possible root must also divide the leading coefficient. The more challenging problem is to determine reducibility: which $f \in \mathbb{Z}[X]$ can be factored there? Kronecker's algorithm for answering this question also appears in [5], and extends readily to $\mathbb{Q}[X]$.

We choose here to present a more general result: the theorem of Michael Rabin from 1960 [23, Thms. 7 & 8] that relates these questions to constructions of the algebraic closure of a given field. The version given here includes a fairly trivial extension to fields that have no computable presentation, as such fields are omitted from Rabin's own statement of the theorem.

Theorem 1 (Rabin's Theorem [23]). *There exist Turing functionals Φ and Ψ, such that, for every presentation F of any countable field,*

- *$\Phi^{\Delta(F)}$ computes $\Delta(K)$ for a presentation K of some algebraically closed field;*
- *and $\Psi^{\Delta(F)}$ computes an embedding $i : F \to K$ such that K is algebraic over the image $i(F)$.*

Thus K may be regarded as the algebraic closure of (the isomorphic image of) F, being both algebraically closed and algebraic over that image. Moreover, the following sets, each computably enumerable relative to $\Delta(F)$, are all Turing-equivalent (relative to $\Delta(F)$):

- *The image $i(F)$ of F, as a subset of the domain of K;*
- *the image $j(F)$ of F within any computable presentation K_0 of K, for an arbitrary F-computable embedding $j : F \to K_0$ with K_0 algebraic over $j(F)$;*
- *the* root set $R_F = \{f \in F[X] : (\exists x \in F)\ f(x) = 0\}$ *of F;*
- *the* splitting set $S_F = \{f \in F[X] : (\exists$ *nonconstant* $g, h \in F[X])\ f = g \cdot h\}$.

The Turing-equivalence of R_F and S_F may be surprising. It is quickly seen that $R_F \leq_T S_F$ (relative to $\Delta(F)$: this really means $R_F \leq_T S_F \oplus \Delta(F)$). Indeed, with an S_F-oracle, we can determine whether a given f factors over F, and if so, we can find a factorization (using $\Delta(F)$) and repeat the question for each factor until we have found the irreducible factors of f in $F[X]$. Then $f \in R_F$ just if one of these factors is linear. The reverse reduction is not so clear. However, it is soon seen that $S_F \leq_T i(F)$ (relative to $\Delta(F)$, again), since for $f \in F[X]$, we can factor the image of f in $K[X]$ into linear factors in $K[X]$, and the products of these linear factors yield all possible factorizations of f in $K[X]$. Then $f \in S_F$ just if one of those (finitely many) factorizations in $K[X]$ has all its coefficients in $i(F)$. To complete the equivalence, one reduces $i(F)$ to R_F: for any $x \in K$, we can find some $g \in F[X]$ with $(i \circ g)(x) = 0$, as K is algebraic over $i(F)$. Using R_F, we can determine whether g has roots in F – and if so, how many roots, by finding a root $a \in F$ and repeating the process for $\frac{g(X)}{X-a}$. Then compute $i(a)$ for each root a of g in F: x lies in $i(F)$ just if it is equal to one of these $i(a)$.

Rabin's Theorem is a classic example of computable structure theory. It reveals exactly how much one needs to know about F in order to construct the algebraic closure of F "around" F, with F as a decidable subfield. The pleasing Turing-equivalence of R_F and S_F is a byproduct. Those readers who still feel

that S_F is somehow more difficult to compute than R_F will find an affirmation of their intuition in [19, 25], where it is shown that (for computable fields F algebraic over $\mathbb{Q}$) R_F is always 1-reducible to S_F, uniformly in $\Delta(F)$, whereas S_F can fail to be 1-reducible to R_F, or even bounded-Turing-reducible to R_F.

Useful corollaries of Rabin's Theorem include several theorems first proven by Kronecker. For example, Kronecker gave an algorithm for deciding $S_{\mathbb{Q}}$, but this now follows directly from Rabin's Theorem and Kronecker's algorithm for $R_{\mathbb{Q}}$. (Irreducibility in $\mathbb{Z}[X]$ is also now quickly seen to be decidable.) Edwards [5] gives Kronecker's actual algorithm, which enables him to construct algebraic closures from an intuitionistic point of view. Rabin's proof of his theorem may be seen as nonconstructive in certain respects: it essentially assumes the existence of the algebraic closure and builds a computable presentation of that closure, rather than constructing the algebraic closure directly.

Furthermore, if $E = F(a)$ is an algebraic field extension of F, then by Rabin's Theorem $S_E \leq_T S_F$, uniformly relative to $\Delta(E)$ and the embedding of F into E. This follows by giving an algorithm for deciding membership (of each $x \in K$) in the image of $F(a)$ from membership in $i(F)$, once the embedding $i : F \to K$ is extended to $F(a)$. Thus all number fields have decidable splitting sets and root sets. In turn, this allows one to compute the Galois group G of a finite algebraic extension E/F, viewed as a set of automorphisms of E: one can determine the order n of G and name its elements $g_1, \ldots, g_n$ so that $g_m(x)$ is computable uniformly from $m \leq n$ and $x \in E$. All that is needed is an S_F-oracle and the minimal polynomial of a primitive generator of E over F.

3 Polynomials in Several Variables

The reader will notice that, while we sketched a proof of the Turing-equivalence claims of Rabin's Theorem above, we never addressed the initial claim of the theorem: the uniform method of producing an algebraic closure K of the given field F and of situating F inside K, via an embedding $i : F \to K$, so as to view K accurately as the algebraic closure of F. This claim is not that difficult to prove when F is an *algebraic field*, by which we mean an algebraic extension of its prime subfield (which here is always $\mathbb{Q}$; if we considered characteristics $p > 0$, it would be the p-element subfield $\mathbb{F}_p$). However, the proof is significantly more difficult for non-algebraic fields. For those with no computable transcendence basis over $\mathbb{Q}$, even a non-uniform construction of K and i requires real work. We content ourselves here with referring the reader to the original paper [23].

However, another algorithm of Kronecker is worthy of notice here. We mentioned above that his method of deciding the splitting set of an algebraic extension $F(a)$, given the splitting set for F, was superseded by Rabin's Theorem (apart from intuitionistic considerations). Kronecker showed the same for a purely transcendental extension $F(t)$ of F, and this does not follow from Rabin's Theorem. Here $F(t)$ can be presented (given $\Delta(F)$) as the set of all rational functions in one variable t over F, i.e., quotients of polynomials in $F[t]$. Of course,

$F(t)$ does not sit inside the algebraic closure of F, so we appeal instead to Kronecker [14].

Kronecker found a trick for deciding whether a polynomial $f \in F(t)[X]$ factors there. (Once again, the modern source [5, §59] expounds his method well.) First he argued that we can clear out the denominators of the rational functions in $F(t)$ serving as coefficients of f, reducing the problem to the situation where f can be viewed as an element of $F[t][X]$, or equivalently, a polynomial in two variables in $F[T, X]$. If n is the degree of T in f, then any factorization of f in $F[T, X]$ produces a factorization of $f(T, T^{n+1})$ in $F[T]$. Using the splitting set of F as an oracle, we find all (finitely many) factorizations of $f(T, T^{n+1})$ in $F[T]$, and check whether any of them arises from a factorization of $f(T, X)$, thus deciding $S_{F(t)}$. Moreover, Rabin's Theorem, with $F(t)$ as the given field, shows that $R_{F(t)} \equiv_T S_{F(t)}$, so $R_{F(t)} \equiv_T S_F \equiv_T R_F$ as well.

This result allows us to move beyond single-variable polynomials when considering irreducibility.

Proposition 1. *Irreducibility of polynomials in $F[X_0, X_1, X_2, \ldots]$ is decidable by a uniform procedure using the splitting set S_F of F (and the atomic diagram $\Delta(F)$, if F is not computable) as an oracle.*

Proof. Applying Kronecker's trick recursively, we derive procedures for deciding irreducibility in $R_n = F[X_0, \ldots, X_n]$ for each n, uniformly in n. So, given $f \in F[X_0, X_1, \ldots]$, we simply find an n with $f \in R_n$ and apply the algorithm for that R_n. Of course, if f factors in $F[X_0, X_1, \ldots]$ at all, both factors must lie in this R_n, so the algorithm for R_n gives the correct answer. □

Proposition 1 reveals a significant distinction between the single-variable and multi-variable situations. In $F[X]$, the questions of irreducibility and having a root are Turing-equivalent (relative to the atomic diagram $\Delta(F)$), by Rabin's Theorem: for example, with $F = \mathbb{Q}$, both $R_{\mathbb{Q}}$ and $S_{\mathbb{Q}}$ are decidable. However, in the multivariable situation, this fails. Irreducibility of polynomials in $\mathbb{Q}[X_0, X_1, \ldots]$ is decidable (and for F in general, it remains Turing-equivalent to S_F, so we do not even bother to give a separate name to the multivariable problem). However, the question of whether an $f \in \mathbb{Q}[X_0, X_1, \ldots]$ has a solution in $\mathbb{Q}$ poses a huge open problem. We refer to it as *Hilbert's Tenth Problem*, generalizing the original question posed by Hilbert.

Definition 1. *For a field F (or more generally a ring),* Hilbert's Tenth Problem for F *is the set*

$$\mathrm{HTP}(F) = \bigcup_n \{f \in F[X_0, \ldots, X_n] : (\exists (x_0, \ldots, x_n) \in F^{n+1})\ f(x_0, \ldots, x_n) = 0\}.$$

So R_F is just the single-variable case of $\mathrm{HTP}(F)$. Of course $R_F \leq_T \mathrm{HTP}(F)$, indeed via a 1-reduction, but the converse in general is false. Indeed, the decidability of $\mathrm{HTP}(\mathbb{Q})$ itself is unknown: this set is computably enumerable, but there is no proof yet whether its Turing degree is the computable degree $\mathbf{0}$, or the degree $\mathbf{0}'$ of the Halting Problem – or conceivably even a different c.e. degree in between these two! Of course, Hilbert's original Tenth Problem was to give

an algorithm deciding HTP($\mathbb{Z}$), which is now known from [15] to be undecidable, having degree $\mathbf{0}'$. It also remains unknown whether there is any existential formula defining the set $\mathbb{Z}$ within the field $\mathbb{Q}$: if such a definition exists, then HTP($\mathbb{Q}$) would have degree $\mathbf{0}'$ too, as membership questions about HTP($\mathbb{Z}$) could then be reduced to membership questions about HTP($\mathbb{Q}$) using that definition. Julia Robinson [24] created the first definition of $\mathbb{Z}$ in $\mathbb{Q}$, in 1949, by a $\forall\exists\forall\exists$-formula, thus showing that the theory of the field $\mathbb{Q}$ is undecidable. Within the past decade, Koenigsmann [13] gave a definition of $\mathbb{Z}$ in $\mathbb{Q}$ by a purely universal formula, but there are reasons to doubt whether an existential definition exists.

Although the situation of HTP($\mathbb{Q}$) remains unresolved, one certainly can build fields F (with $\Delta(F)$ computable) for which HTP(F) $\not\leq_T R_F$. Thus R_F can be strictly easier than HTP(F) under Turing reducibility. In Section 5 we will mention some further results concerning this question.

4 Transcendence Bases

For fields in general, the most basic question about an element is whether it is algebraic or transcendental. These questions can be asked relative to any subfield, but for us they will always refer to transcendence over the prime subfield: either $\mathbb{Q}$ or $\mathbb{F}_p$, depending on the characteristic. In every presentation F of any countable field, the prime subfield is always computably enumerable relative to $\Delta(F)$.

The prime subfield may be undecidable relative to $\Delta(F)$, but this can only occur if F has characteristic 0 and contains transcendental elements. For an algebraic x, Kronecker's decision procedure for $S_{\mathbb{Q}}$ allows us to find the minimal polynomial of x over $\mathbb{Q}$ and thus decide whether $x \in \mathbb{Q}$. Furthermore, it almost defines x within F, as only finitely many other elements of f can have the same minimal polynomial. (These other elements are the $\mathbb{Q}$-*conjugates* of x in F.)

A *transcendence basis* B for F (over $\mathbb{Q}$) is a way of extending this situation to all of F. By definition, B is a maximal subset of F algebraically independent over $\mathbb{Q}$, and so every $x \in F$ has a minimal polynomial over the subfield $\mathbb{Q}(B)$, which identifies x (relative to B) up to finitely many conjugates, in the same way that the minimal polynomial of an element of an algebraic field identifies that element. If we can enumerate a particular transcendence basis B, therefore, we are largely back in the comfortable situation of an algebraic field.

In [16], Metakides and Nerode provided the first example of a computable field with no computable transcendence basis. This result is sharp, as there is always a transcendence basis B that is co-c.e. relative to $\Delta(F)$: it consists of each element x_i in the domain of F that is independent over $\mathbb{Q}(x_0, \ldots, x_{i-1})$. The proof involves ensuring that F has infinite transcendence degree, but using a priority construction to guarantee that every infinite c.e. subset W_e of the domain contains an algebraic element. This is not difficult: the key is that the type of a transcendental x in a field is a nonprincipal type, i.e., not generated by any single formula, and therefore certainly not generated by any single existential formula. It follows that, no matter what finite amount of $\Delta(F)$ has been defined

so far, it will always be consistent with that amount for x to be algebraic. So, when it comes time to make some element of W_e algebraic, it is always possible to do so.

The further complication in transcendental fields is that there is no canonical transcendence basis. The prime subfield of F is always c.e. relative to $\Delta(F)$, by a uniform enumeration procedure, and is rigid, so elements in each copy of an algebraic field F can be identified, up to conjugacy, by their minimal polynomials. However, even in fields such as the purely transcendental extension $K = \mathbb{Q}(t_0, t_1, \ldots)$ of $\mathbb{Q}$ (which is just the field of all rational functions over $\mathbb{Q}$ in the variables t_i), this property no longer holds. It is quickly seen that there are presentations of this K in which no generating set of transcendentals is c.e., and therefore, even the natural candidate $\{t_0, t_1, \ldots\}$ for a canonical transcendence basis does not succeed in the way one requires. Indeed, this is a significant open question.

Question 1 (Melnikov & Miller). For the field $K = \mathbb{Q}(t_0, t_1, \ldots)$, find the lowest complexity level $\mathcal{S}$ such that every presentation F of K is generated by some transcendence basis of complexity $\mathcal{S}$.

It is known that Π^1_1 is such a complexity level, and that Π^0_1 is the least candidate for $\mathcal{S}$, but this leaves a wide spread of possibilities. This question is closely related to the categoricity spectrum of the field $\mathbb{Q}(t_0, t_1, \ldots)$; see [9] for basic definitions and [18] for some results involving fields.

One might hope that, for each computable field K, there might at least exist a computable copy F of K with a computable transcendence basis. However, this hope was dashed by Kalimullin, Schoutens, and the author in [11].

Theorem 2 (Corollary 3 from [11]). *For every Turing degree $\boldsymbol{c} \leq \mathbf{0}'$, there exists a computable field K such that, in every computable copy $F \cong K$, every transcendence basis for F has degree $\geq \boldsymbol{c}$. If $\boldsymbol{c}$ is itself a c.e. degree, then we can also ensure that every computable copy actually has a basis of degree $\boldsymbol{c}$.*

Oddly, the results here were based largely on work in [22] that produced a computable field K, of infinite transcendence degree, such that every copy F of K has a transcendence basis computable from $\Delta(F)$. (In particular, every computable copy has a computable transcendence basis.) The argument there used existential formulas to define the elements of one particular transcendence basis: the basis elements were those x such that, for some y in the field, (x, y) formed a nontrivial solution to a Fermat polynomial $X^p + Y^p = 1$. This technique of "tagging" basis elements by adjoining roots of polynomials over those elements was subsequently extended by Poonen, Schoutens, Shlapentokh, and the author in [21]. Discussion of that work is beyond the scope of this tutorial, but we provide a short version of the relevant theorem here. Roughly, it states that fields in general are just as complex as any other class of structures. Graphs, groups, and partial orders are also maximally complex, whereas linear orders and trees (for example) are not.

Theorem 3 (Theorem 1.8 from [21]). *For every countable first-order structure $\mathcal{M}$ in a finite signature, there exists a countable field K with the same computable-structure-theoretic properties as $\mathcal{M}$.*

To give a more precise, though incomplete, list of the properties preserved: K has the same Turing degree spectrum as $\mathcal{M}$, the same categoricity spectrum as $\mathcal{M}$, the same computable dimension as $\mathcal{M}$, and the same automorphism spectrum as $\mathcal{M}$. Moreover, for every relation R on $\mathcal{M}$, there is a relation on K with the same degree spectrum. (All of these properties are described in [21], and most in [10]. Some of them require $\mathcal{M}$ to be a computable structure; if it is, then K can also be taken to be computable.) Indeed, even properties unknown when Theorem 3 was proven have turned out to carry over from $\mathcal{M}$ to K, such as the degree of categoricity on a cone, defined in [1]. The theorem in general holds for countable structures in computable signatures, not just finite signatures, with the exception of certain simple but pathological structures known as *automorphically trivial* structures; see [12] for those details.

5 Algebraic Fields

Algebraic fields are fields in which every element is algebraic, i.e., is the root of some polynomial over the prime subfield, which in this section will always be $\mathbb{Q}$. The class $\mathfrak{A}$ of such fields is very far from satisfying Theorem 3: procedures involving fields in $\mathfrak{A}$ are in general much closer to computable, because each element of such a field can be effectively identified up to conjugacy over $\mathbb{Q}$, thanks to the decidability of $S_{\mathbb{Q}}$. The means that, for two presentations of fields in $\mathfrak{A}$, the property of being isomorphic is far simpler than for fields in general. (Theorem 3 shows the isomorphism relation to be Π^1_1-complete for fields in general.)

Theorem 4. *Two fields $E, F \in \mathfrak{A}$ are isomorphic just if*

$$\{f \in \mathbb{Q}[X] : f \text{ has a root in } E\} = \{f \in \mathbb{Q}[X] : f \text{ has a root in } F\}.$$

Similarly, the elements $x_0, x_1 \in F$ lie in the same orbit under automorphisms of F just if they have the same minimal polynomial over $\mathbb{Q}$ and

$$\{f \in \mathbb{Q}[X,Y] : f(x_0, Y) \in R_F\} = \{f \in \mathbb{Q}[X,Y] : f(x_1, Y) \in R_F\}.$$

Theorem 4 suggests that $\{f \in \mathbb{Q}[X] : f \text{ has a root in } F\}$ can serve as an index for the isomorphism type of F, for each $F \in \mathfrak{A}$. This is the foundation of work in [20] and [3, 4], which uses these indices to place a topology on the space $\mathfrak{A}$ of all algebraic field extensions of $\mathbb{Q}$. The same topology has been discovered by various field theorists independently over the years: it is sometimes known as the *étale topology*, or seen as the Vietoris topology on the space of all closed subgroups of $\mathrm{Aut}(\overline{\mathbb{Q}})$. Each of these is the same topology on the space of all subfields of $\overline{\mathbb{Q}}$; one mods out by the relation of isomorphism and imposes the quotient topology in order to topologize the space of isomorphism types in $\mathfrak{A}$. Both the

étale topology and the quotient modulo isomorphism have the pleasing property of being homeomorphic to the set $2^{\mathbb{N}}$ under the usual Cantor topology, and this allows one to use elements of Cantor space as indices for the isomorphism types. Each index essentially specifies $\{f \in \mathbb{Q}[X] : f \text{ has a root in } F\}$, as described above, although a certain amount of coding is necessary. This creates an effective classification of $\mathfrak{A}$ by the elements of $2^{\mathbb{N}}$.

Theorem 5 (see [20]). *There exist Turing functionals Φ and Ψ such that, for all presentations E and F of fields in $\mathfrak{A}$ and all $S \in 2^{\mathbb{N}}$:*

- $\Phi^{E \oplus R_E} \in 2^{\mathbb{N}}$, *with* $\Phi^{E \oplus R_E} = \Phi^{F \oplus R_F}$ *if and only if* $E \cong F$; *and*
- Ψ^S *computes* $\Delta(F) \oplus R_F$ *for some presentation* F *of a field in* $\mathfrak{A}$; *and*
- $\Phi^{(\Psi^S)} = S$.

Eisenträger, Springer, Westrick, and the author have recently exploited this homeomorphism, using the Baire-category property of co-meagerness in $2^{\mathbb{N}}$ to prove the following.

Theorem 6 (see [4]). *In $\mathfrak{A}$ under the topology described above, the (isomorphism types of) fields satisfying all of the following properties form a comeager set.*

- *in some presentation F of the field, $R_F \not\leq_T \Delta(F)$; but*
- *in every presentation F of the field, $(R_F)' \leq_T (\Delta(F))'$; and*
- *in every presentation F of the field, $R_F \oplus \Delta(F) \equiv_T \mathrm{HTP}(F) \oplus \Delta(F)$.*

Thus the "generic" situation for algebraic extensions of $\mathbb{Q}$ is that the root set is noncomputable but always low relative to the atomic diagram. Moreover, the question of solvability of polynomial equations in several variables is "generically" only as hard as the same question for polynomials in a single variable (i.e., the root set), hence also low but generally noncomputable relative to $\Delta(F)$.

6 The Field $\mathbb{R}$

Abstractly, it is natural to consider the problem of whether a polynomial $f \in \mathbb{R}[X_0, \ldots, X_n]$ has a real solution $\boldsymbol{x} \in \mathbb{R}^n$ with $f(\boldsymbol{x}) = 0$. In practice, since $\mathbb{R}$ is uncountable, the techniques used here are entirely different from those for countable fields, and we content ourselves with a brief summary.

It has been known since the work of Tarski that the theory of the field $\mathbb{R}$ is decidable. From this one directly infers a decision procedure for the question of whether an $f \in \mathbb{Q}[X_0, \ldots, X_n]$ has a solution in $\mathbb{R}^n$. Indeed, we can describe it succinctly: if we find $\boldsymbol{x}$ and $\boldsymbol{y}$ in the dense subset $\mathbb{Q}^n$ with $f(\boldsymbol{x}) < 0 < f(\boldsymbol{y})$, then the Intermediate Value Theorem guarantees a solution of f in $\mathbb{R}^n$; whereas, if no such pair $(\boldsymbol{x}, \boldsymbol{y})$ exists, then by a theorem of Artin f is either a sum of squares of polynomials in $\mathbb{Q}[X_0, \ldots, X_n]$ or else the negation of a sum of such squares. For a sum of squares, f will have a solution only if the absolute minimum value of f is 0, and so basic calculus yields the endgame.

When the polynomial f is allowed to have arbitrary real coefficients, one must first explain how those coefficients are to be presented. The usual procedure, in computable analysis, is to use *fast-converging Cauchy sequences* $\langle q_n \rangle_{n \in \mathbb{N}}$ of rational numbers, with the limit $c \in \mathbb{R}$ of the sequence satisfying $|c - q_n| < 2^{-n}$ for all n. This is best viewed as an approximation of c by open intervals $(q_n - 2^{-n}, q_n + 2^{-n})$, all containing c, whose lengths decrease effectively to 0. Of course, a single c will have many such representations, including noncomputable ones. The book [28] is a standard source for computable analysis.

Over countable fields, the only important aspect of the root set is determining whether a root exists: if it does, one can simply search through the field until a root is found. Over $\mathbb{R}^n$, this is no longer applicable, so this problem bifurcates: the first problem is to decide the existence of a solution, and if one exists, the second problem is to produce a solution. The first of these is undecidable, even for $n = 1$, and the proof is fairly quick. Suppose Φ were a Turing functional that, when given an oracle containing $(d+1)$ Cauchy sequences converging fast to real numbers $c_0, \ldots, c_d$, outputs either "yes" if $\sum c_i X^i = 0$ has a solution in $\mathbb{R}$, or "no" if it has no solution. Run Φ on the monomial X^2, given by constant Cauchy sequences $(1, 1, 1, \ldots)$, $(0, 0, 0, \ldots)$ and $(0, 0, 0, \ldots)$ to represent $1X^2 + 0X^1 + 0X^0$. Φ must output "yes" after examining the first u terms of each sequence, for some finite "use" $u \in \mathbb{N}$. But then, if we run it again and replace the coefficient in the X^0 term by $(0, 0, \ldots, 0, 2^{-(u+1)}, 2^{-(u+1)}, \ldots)$ with u initial 0's, it will give the same output "yes," which will be incorrect: the polynomial is now $X^2 + \frac{1}{2^{u+1}}$, which has no root in $\mathbb{R}$.

The second problem is also undecidable, and again tangency is the culprit. For example, the polynomial $f(X) = X^4 - 2X^2 + 1$ has real roots ± 1, but an arbitrarily small nonzero linear coefficient c can make either of them disappear: for $c > 0$, $X^4 - 2X^2 + cX + 1$ has only negative roots, while when $c < 0$ it has only positive roots. This allows us to use a strategy similar to the above: wait for a functional Φ to compute its first approximation q_0 to a root of $f(X)$, and then perturb the linear coefficient just slightly, making it either positive (if $q_0 \geq 0$) or negative (otherwise).

References

1. B.F. Csima & M. Harrison-Trainor; Degrees of categoricity on a cone via eta-systems, *Journal of Symbolic Logic* **82** (2017) 1, 325–346.
2. M. Davis, H. Putnam, & J. Robinson; The decision problem for exponential diophantine equations, *Annals of Mathematics* **74** (1961) 3, 425–436.
3. K. Eisenträger, R. Miller, C. Springer, & L. Westrick; A topological approach to undefinability in algebraic fields, submitted for publication.
4. K. Eisenträger, R. Miller, C. Springer, & L. Westrick; Genericity and forcing for algebraic fields, in preparation.
5. H.M. Edwards; *Galois Theory* (New York: Springer-Verlag, 1984).
6. Yu.L. Ershov; Theorie der Numerierungen, *Zeits. Math. Logik Grund. Math.* **23** (1977), 289-371.
7. M.D. Fried & M. Jarden; *Field Arithmetic* (Berlin: Springer-Verlag, 1986).

8. A. Frohlich & J.C. Shepherdson; Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407–432.
9. E. Fokina, I. Kalimullin, & R. Miller; Degrees of categoricity of computable structures, *Archive for Mathematical Logic*, **49** (2010), 51–67.
10. D.R. Hirschfeldt, B. Khoussainov, R.A. Shore, & A.M. Slinko; Degree spectra and computable dimensions in algebraic structures, *Annals of Pure and Applied Logic* **115** (2002), 71-113.
11. I. Kalimullin, R. Miller, & H. Schoutens; Degree spectra for transcendence in fields, in *Computing with Foresight and Industry: 15th Conference on Computability in Europe, CiE 2019*, eds. F. Manea, B. Martin, D. Paulusma, & G. Primiero, *Lecture Notes in Computer Science* **11558** (Berlin: Springer-Verlag, 2019), 205–216.
12. J.F. Knight; Degrees coded in jumps of orderings, *Journal of Symbolic Logic* **51** (1986), 1034-1042.
13. J. Koenigsmann; Defining $\mathbb{Z}$ in $\mathbb{Q}$. *Annals of Mathematics* **183** (2016) 1, 73–93.
14. L. Kronecker; Grundzüge einer arithmetischen Theorie der algebraischen Größen, *J. f. Math.* **92** (1882), 1-122.
15. Yu.V. Matiyasevich; The Diophantineness of enumerable sets. *Dokl. Akad. Nauk SSSR* **191** (1970), 279–282.
16. G. Metakides & A. Nerode; Effective content of field theory, *Annals of Mathematical Logic* **17** (1979), 289-320.
17. R. Miller; Computable fields and Galois theory, *Notices of the American Mathematical Society* **55** (August 2008) 7, 798–807.
18. R. Miller; $\boldsymbol{d}$-Computable categoricity for algebraic fields, *The Journal of Symbolic Logic* **74** (2009) 4, 1325–1351.
19. R. Miller; Is it easier to factor a polynomial or to find a root? *Transactions of the American Mathematical Society*, **362** (2010) 10, 5261–5281.
20. R. Miller; Isomorphism and classification for countable structures, *Computability* **8** (2019) 2, 99–117.
21. R. Miller, B. Poonen, H. Schoutens, & A. Shlapentokh; A computable functor from graphs to fields, *Journal of Symbolic Logic* **83** (2018) 1, 326–348.
22. R. Miller & H. Schoutens: Computably categorical fields via Fermat's Last Theorem, *Computability* **2** (2013) 51–65.
23. M. Rabin; Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341-360.
24. J. Robinson; Definability and decision problems in arithmetic. *Journal of Symbolic Logic* **14** (1949), 98–114.
25. R.M. Steiner; Computable fields and the bounded Turing reduction, *Annals of Pure and Applied Logic* **163** (2012), 730–742.
26. V. Stoltenberg-Hansen & J.V. Tucker; Computable Rings and Fields, in *Handbook of Computability Theory*, ed. E.R. Griffor (Amsterdam: Elsevier, 1999), 363-447.
27. B.L. van der Waerden: *Algebra*, volume I, trans. F Blum & J.R. Schulenberger (New York: Springer-Verlag, 1970 hardcover, 2003 softcover).
28. K. Weihrauch; *Computable Analysis: An Introduction* (Berlin: Springer, 2000).