Adapting Rabin's Theorem for Differential Fields

Russell Miller[1⋆] and Alexey Ovchinnikov[2⋆⋆]

[1] Queens College of CUNY
and
The CUNY Graduate Center
365 Fifth Avenue, New York, NY 10016 USA
Russell.Miller@qc.cuny.edu
http://qcpages.qc.cuny.edu/~rmiller
[2] Queens College of CUNY
65-30 Kissena Blvd., Flushing, NY 11367 USA
Alexey.Ovchinnikov@qc.cuny.edu
http://qcpages.qc.cuny.edu/~aovchinnikov

Abstract. Harrington extended the first half of Rabin's Theorem to differential fields, proving that every computable differential field can be viewed as a computably enumerable subfield of a computable presentation of its differential closure. For fields F, the second half of Rabin's Theorem says that this subfield is Turing-equivalent to the set of irreducible polynomials in $F[X]$. We investigate possible extensions of this second half, asking both about the degree of the differential field K within its differential closure and about the degree of the set of constraints for K, which forms the closest analogue to the set of irreducible polynomials.

Key words: Computability, constraint, differential fields, Rabin's Theorem, recursion theory.

1 Introduction

Rabin's Theorem is fundamental to the study of computability theory and fields. Proven in [12] in 1960, it gives an effective construction of the algebraic closure of a computable field F around the field itself, and describes the exact conditions necessary for the original field to be decidable (as a subfield of the algebraic closure), namely the decidability of the set of reducible polynomials in $F[X]$. As the notion of algebraic closure is essential to modern field theory, the question it addresses is absolutely natural, and it answers that question convincingly.

⋆ The corresponding author was partially supported by Grant # DMS – 1001306 from the National Science Foundation and by grants numbered 62632-00 40 and 63286-00 41 from The City University of New York PSC-CUNY Research Award Program.

⋆⋆ The second author was partially supported by Grant # CCF – 0952591 from the National Science Foundation and by Grant # 60001-40 41 from The City University of New York PSC-CUNY Research Award Program. Both authors wish to acknowledge useful conversations with Dave Marker and Michael Singer.

The practice of closing a field algebraically foreshadowed the construction of the differential closure of a differential field, and the stream of results which flowed from this notion with the work of Kolchin and many others, beginning in the mid-twentieth century, emphasized its importance. Moreover, starting later in that century, differentially closed fields have become the focus of a great deal of work in model theory. The theory $\mathbf{DCF_0}$ of differentially closed fields of characteristic 0 is in many ways even more interesting to model theorists than the corresponding theory $\mathbf{ACF_0}$ for algebraically closed fields: both are ω-stable, but only $\mathbf{DCF_0}$ has infinite Morley rank. Today differentially closed fields are widely studied by both algebraists and logicians.

Therefore it is natural to attempt to replicate Rabin's Theorem in the context of computable differential fields and their differential closures. Harrington took a significant step in this direction in [5] in 1974, proving that every computable differential field does indeed have a computable differential closure, and can be enumerated inside that closure, just as Rabin showed can be done for a computable field inside its algebraic closure. All these results, and the terms used in them, are defined fully in the next section. However, Harrington's theorem mirrors only the first half of Rabin's Theorem: it remains to determine what conditions would be sufficient for – or better yet, equivalent to – decidability of the original differential field within its differential closure. With this abstract we begin these efforts, giving the current state of knowledge in Sections 4 and 5. Sections 2 and 3 sketch most of the necessary background. For further questions we suggest [13] for general background in computability, [4] for computable model theory, [2] for field arithmetic, [8, 11, 10] for introductions to computable fields, and [1, 7] for differential fields in the context of model theory.

2 Computable Differential Fields

Differential fields are a generalization of fields, in which the field elements are often viewed as functions. The elements are not treated as functions, but the differential operator(s) on them are modeled on the usual notion of differentiation of functions.

Definition 1 A *differential field* is a field K with one or more additional unary functions δ_i satisfying the following two axioms for all $x, y \in K$:

$$\delta_i(x+y) = \delta_i x + \delta_i y \qquad \delta_i(x \cdot y) = (x \cdot \delta_i y) + (y \cdot \delta_i x).$$

The *constants* of K are those x such that, for all i, $\delta_i x = 0$. They form a differential subfield C_K of K.

So every field can be made into a differential field by adjoining the zero operator $\delta x = 0$. For a more common example, consider the field $F(X_1, \ldots, X_n)$ of rational functions over a field F, with the partial derivative operators $\delta_i = \frac{\partial}{\partial X_i}$. We will be concerned only with *ordinary differential fields*, i.e. those with a single differential operator δ.

The next definitions arise from the standard notion of a computable structure. To avoid confusion, we use the domain $\{x_0, x_1, \ldots\}$ in place of ω.

Definition 2 A *computable field* F consists of a set $\{x_i : i \in I\}$, where I is an initial segment of ω, such that these elements form a field with the operations given by Turing-computable functions f and g:

$$x_i + x_j = x_{f(i,j)} \qquad x_i \cdot x_j = x_{g(i,j)}.$$

A *computable differential field* is a computable field with one or more differential operators δ as in Definition 1, each of which is likewise given by some Turing-computable function h with $\delta(x_i) = x_{h(i)}$.

Fröhlich and Shepherdson were the first to consider computable algebraically closed fields, in [3]. However, the definitive result on the effectiveness of algebraic closure is Rabin's Theorem. To state it, we need the natural notions of the root set and the splitting set.

Definition 3 Let F be any computable field. The *root set* R_F of F is the set of all polynomials in $F[X]$ having roots in F, and the *splitting set* S_F is the set of all polynomials in $F[X]$ which are reducible there. That is,

$$R_F = \{p(X) \in F[X] : (\exists a \in F)\ p(a) = 0\}$$
$$S_F = \{p(X) \in F[X] : (\exists \text{ nonconstant } p_0, p_1 \in F[X])\ p = p_0 \cdot p_1\}.$$

Both these sets are computably enumerable. They are computable whenever F is isomorphic to a prime field $\mathbb{Q}$ or $\mathbb{F}_p$, or to any finitely generated extension of these. At the other extreme, if F is algebraically closed, then clearly both R_F and S_F are computable. However, there are many computable fields F for which neither R_F nor S_F is computable; see the expository article [8, Lemma 7] for a simple example. Fröhlich and Shepherdson showed that R_F is computable iff S_F is, and Rabin's Theorem then related them both to a third natural c.e. set related to F, namely its image inside its algebraic closure. (Rabin's work actually ignored Turing degrees, and focused on S_F rather than R_F, but the theorem stated here follows readily from Rabin's proof.) More recent works [9, 14] have compared these three sets under stronger reducibilities, but here, following Rabin, we consider only Turing reducibility, denoted by $\leq_T$, and Turing equivalence $\equiv_T$.

Theorem 4 (Rabin's Theorem, in [12]) *For every computable field F, there exist an algebraically closed computable field E and a computable field homomorphism $g : F \to E$ such that E is algebraic over the image $g(F)$. Moreover, for every embedding g satisfying these conditions, the image $g(F)$ is Turing-equivalent to both the root set R_F and the splitting set S_F of the field F.*

We will refer to any embedding $g : F \to E$ satisfying the conditions from Rabin's Theorem as a *Rabin embedding* of F. Since this implicitly includes the presentation of E (which is required by the conditions to be algebraically closed), a Rabin embedding is essentially a presentation of the algebraic closure of F, with F as a specific, but perhaps undecidable, subfield.

As we shift to consideration of differential fields, we must first consider the analogy between algebraic closures of fields and differential closures of differential fields. The theory $\mathbf{DCF_0}$ of differentially closed fields K of characteristic 0 is a complete theory, and was axiomatized by Blum (see e.g. [1]) using the axioms for differential fields of characteristic 0, along with axioms stating that, for every pair of nonzero differential polynomials $p, q \in K\{Y\}$ with $\mathrm{ord}(p) > \mathrm{ord}(q)$, there exists some $y \in K$ with $p(y) = 0 \neq q(y)$. (The *differential polynomial ring* $K\{Y\}$ is the ring $K[Y, \delta Y, \delta^2 Y, \ldots]$ of all algebraic polynomials in Y and its derivatives. The *order* of $p \in K\{Y\}$ is the greatest $r \geq 0$ such that $\delta^r Y$ appears nontrivially in $p(Y)$. By convention the zero polynomial has order $-\infty$, and all other constant polynomials have order -1. Blum's axioms therefore include formulas saying that all nonconstant algebraic polynomials $p \in K[Y]$ have roots in $\overline{K}$, by taking $q = 1$.)

For a differential field K with extensions containing elements x_0 and x_1, we will write $x_0 \cong_K x_1$ to denote that $K\langle x_0\rangle \cong K\langle x_1\rangle$ via an isomorphism fixing K pointwise and sending x_0 to x_1. This is equivalent to the property that, for all $h \in K\{Y\}$, $h(x_0) = 0$ iff $h(x_1) = 0$; a model theorist would say that x_0 and x_1 realize the same atomic type over K. The same notation $x_0 \cong_F x_1$ could apply to elements of field extensions of a field F, for which the equivalent property would involve only algebraic polynomials $h \in K[Y]$.

Let $K \subseteq L$ be an extension of differential fields. An element $x \in L$ is *constrained over* K if x satisfies some *constraint* over K, as defined here.

Definition 5 Let K be a differential field. A *constraint* for K is a pair (p, q) of monic differential polynomials in $K\{Y\}$ with the properties that $\mathrm{ord}(p) > \mathrm{ord}(q)$, that p is irreducible as a polynomial in $K[Y, \delta Y, \ldots]$, and that for all differential field extensions L_0 and L_1 of K and all $x_i \in L_i$ such that $p(x_i) = 0 \neq q(x_i)$, we have $x_0 \cong_K x_1$. Such elements x_0 and x_1 are said to *satisfy the constraint* (p, q). We denote the complement by

$$T_K = \{(p, q) \in (K\{Y\})^2 : (p, q) \text{ is } \textit{not} \text{ a constraint}\},$$

and refer to $\overline{T_K}$ as the *constraint set* for K. If T_K is computable, we say that K has a *constraint algorithm*.

The notation T_K is intended to parallel the notation R_F and S_F. (Also, recall that C_K already denotes the constant subfield of K). Definition 5 parallels the definition of the splitting set S_F in function if not in form. For fields F, irreducible polynomials $p(X)$ have exactly the same property: if $p(x_0) = p(x_1) = 0$ (for x_0 and x_1 in any algebraic field extensions of F), then $x_0 \cong_F x_1$ (that is, $F(x_0) \cong F(x_1)$ via an F-isomorphism mapping x_0 to x_1). So T_K is indeed the analogue of S_F: both are Σ^0_1 sets, given that K and F are both computable, and both are the negations of the properties we need to produce isomorphic extensions. To see that T_K is Σ^0_1, note that $(p, q) \in \overline{T_K}$ iff all $x_0, x_1 \in \overline{K}$ and all $h \in K\{Y\}$ satisfy:

$$[p(x_0) = p(x_1) = 0 \ \&\ q(x_0) \neq 0 \neq q(x_1)] \implies (h(x_0) = 0 \iff h(x_1) = 0),$$

the latter condition (over all h) being equivalent to $K\langle x_0\rangle \cong K\langle x_1\rangle$.

If $x \in L$ is constrained over K by (p, q), then there exists a differential subfield of L, extending K and containing x, whose transcendence degree as a field extension of K is finite. Indeed, writing $K\langle x\rangle$ for the smallest differential subfield of L containing x and all of K, we see that the transcendence degree of $K\langle x\rangle$ over K is the smallest order of any nonzero element of $K\{Y\}$ with root x. This will be seen below to be exactly the order of p. The elements of L which are constrained over K turn out to form a differential field in their own right. If this subfield is all of L, then L itself is said to be a *constrained extension* of K.

An algebraic closure $\overline{F}$ of a field F is an algebraically closed field which extends F and is algebraic over it. Of course, one soon proves that this field is unique up to isomorphism over F (that is, up to isomorphisms which restrict to the identity on the common subfield F). On the other hand, each F has many algebraically closed extensions; the algebraic closure is just the smallest of them. Likewise, each differential field K has many differentially closed field extensions; a *differential closure* of K is such an extension which is constrained over K. As with fields, the differential closure of K turns out to be unique up to isomorphism over K. On the other hand, the differential closure $\overline{K}$ of K is generally not *minimal*: there exist differential field embeddings of $\overline{K}$ into itself over K whose images are proper differential subfields of $\overline{K}$. This provides a first contrast between $\mathbf{DCF_0}$ and $\mathbf{ACF}_0$, since the corresponding statement about algebraic closures is false.

With this much information in hand, we can now state the parallel to the first half of Theorem 4.

Theorem 6 (Harrington [5], Corollary 3) *For every computable differential field K, there exists a differentially closed computable differential field L and a computable differential field homomorphism $g : K \to L$ such that L is constrained over the image $g(K)$.*

For the sake of uniform terminology, we continue to refer to a computable function g as in Theorem 6 as a *Rabin embedding* for the differential field K.

Harrington actually proves the existence of a computable structure L which is the prime model of the theory T generated by $\mathbf{DCF_0}$ and the atomic diagram of K. Thus L is a computable structure in the language $\mathcal{L}'$ in which the language of differential fields is augmented by constants for each element of K. The embedding of K into L is accomplished by finding, for any given $x \in K$, the unique element $y \in L$ which is equal to the constant symbol for x. Clearly this is a computable process, since L is a computable $\mathcal{L}'$-structure, and so we have our embedding of K into L. Since L is the prime model of T, it must be constrained over K: otherwise it could not embed into the constrained closure, which is another model of T. So L satisfies the definition of the differential closure of K, modulo the computable embedding. The same holds in positive characteristic.

The root set and splitting set of a differential field K are still defined, of course, just as for any other field. However, with the differential operator δ now in the language, several other sets can be defined along the same lines and are

of potential use as we attempt to adapt Rabin's Theorem. The most important of these is the constraint set, which is analogous in several ways to the splitting set and will be discussed in the next section.

We will also need a version of the Theorem of the Primitive Element for differential fields. This was provided long since by Kolchin.

Theorem 7 (Kolchin; see [6], p. 728) *Assume that an ordinary differential field F contains an element x with $\delta x \neq 0$. If E is a differential subfield of the differential closure $\overline{F}$ and E is generated (as a differential field over F) by finitely many elements, then there is a single element of E which generates all of E as a differential field over F.* □

Kolchin gave counterexamples in the case where δ is the zero derivation on F, and also extended this theorem to partial differential fields with m derivations: the generalized condition there is the existence of m elements whose Jacobian is nonzero.

3 Constraints

Proposition 8 *Let K be a differential field. Then for each $x \in \overline{K}$, there is exactly one $p \in K\{Y\}$ such that x satisfies a constraint of the form $(p, q) \in \overline{T_K}$. Moreover, ord$(p)$ is least among the orders of all nonzero differential polynomials in the radical differential ideal $I_K(x)$ of x within $K\{Y\}$:*

$$I_K(x) = \{p \in K\{Y\} : p(x) = 0\},$$

and deg(p) is the least degree of $\delta^{ord(p)}Y$ in any polynomial in $K\{Y\}$ of order ord(p) with root x.

Proof. Since $\overline{K}$ is constrained over K, each $x \in \overline{K}$ satisfies at least one constraint $(p, q) \in \overline{T_K}$. Set $r = \text{ord}(p)$, and suppose there were a nonzero $\tilde{p}(Y) \in I_K(x)$ with $\text{ord}(\tilde{p}) < r$. By Blum's axioms for $\mathbf{DCF}_0$, there would exist $y \in \overline{K}$ with $p(y) = 0 \neq q(y) \cdot \tilde{p}(y)$, since the product $(q \cdot \tilde{p})$ has order $< r$. But then y also satisfies (p, q), yet $\tilde{p}(y) \neq 0 = \tilde{p}(x)$, so that $K\langle x\rangle \ncong K\langle y\rangle$. This would contradict Definition 5. Hence r is the least order of any nonzero differential polynomial with root x.

It follows from minimality of r that $\{x, \delta x, \ldots, \delta^{r-1}x\}$ is algebraically independent over K. The polynomial $p(x, \delta x, \ldots, \delta^{r-1}x, Y)$ must then be the minimal polynomial of $\delta^r x$ over the field $K(x, \ldots, \delta^{r-1}x)$, since $p(Y, \delta Y, \ldots, \delta^r Y)$ is irreducible in $K[Y, \delta Y, \ldots, \delta^r Y]$. This implies the claim in Proposition 8 about $\deg(p)$, and also shows that every constraint satisfied by x must have first component p. □

In fact, the irreducibility of $p(Y)$ is barely necessary in Definition 5. The condition that $K\langle x\rangle \cong K\langle y\rangle$ for all x, y satisfying the constraint shows that $p(Y)$ cannot factor as the product of two distinct differential polynomials. The only reason for requiring irreducibility of p is to rule out the possibility of p being

a perfect square, cube, etc. in $K\{Y\}$. If these were allowed, the uniqueness in Proposition 8 would no longer hold.

It is quickly seen that if $(p, q) \in \overline{T_K}$, then also $(p, q \cdot h) \in \overline{T_K}$ for every $h \in K\{Y\}$. So the constraint satisfied by an $x \in \overline{K}$ is not unique. However, Proposition 8 does show that the following definition makes sense.

Definition 9 If $K \subseteq L$ is an extension of differential fields, then for each $x \in L$, we define $\mathrm{ord}_K(x) = \mathrm{ord}(p)$, where (p, q) is the constraint in $\overline{T_K}$ satisfied by x. If no such constraint exists, then $\mathrm{ord}_K(x) = \infty$. Notice that in the constrained closure of K, every element has finite order over K.

4 Decidability in the Constraint Set

We now present our principal result thus far on constraint sets and Rabin embeddings for differential fields. Of course, we hope to establish more results, and perhaps the converse, in the near future.

Theorem 10 *Let K be any computable differential field with a single nonzero derivation δ, and $g : K \to \overline{K}$ a Rabin embedding of K. Then all of the following are computable in an oracle for the constraint set $\overline{T_K}$: the Rabin image $g(K)$, the set A of finite subsets of $\overline{K}$ algebraically independent over $g(K)$, and the order function ord_K on $\overline{K}$.*

Proof. First we show that ord_K is computable from a $\overline{T_K}$-oracle. Given any $x \in \overline{K}$, we use the oracle to find some $(p, q) \in \overline{T_K}$ satisfied by x. Since $\overline{K}$ is the constrained closure of K, such a constraint must exist, and when we find it, we know by Proposition 8 that $\mathrm{ord}_K(x) = \mathrm{ord}(p)$.

Next we show that $g(K) \leq_T \overline{T_K}$. Our procedure accepts an arbitrary $x \in \overline{K}$ as input, and searches through $\overline{T_K}$, using its oracle, for a constraint (p, q) with $p^g(x) = 0 \neq q^g(x)$. (Here p^g represents the polynomial in $\overline{K}\{Y\}$ whose coefficients are the images under g of the coefficients of p.) Since $\overline{K}$ is constrained over K, it must eventually find such a constraint (p, q), and it concludes that $x \in g(K)$ iff $\mathrm{ord}(p) = 0$ and $p(Y)$ is linear in Y. Of course, if $p(Y) = Y - b$ (hence is of order 0, when viewed as a differential polynomial), then $x = g(b) \in g(K)$. Conversely, if $x \in g(K)$, then $(Y - g^{-1}(x),\ 1)$ is readily seen to lie in $\overline{T_K}$, since it is satisfied by no element of $\overline{K}$ except x. Proposition 8 shows that our algorithm will find a constraint with first coordinate $(Y - g^{-1}(x))$, hence will conclude correctly that $x \in g(K)$.

It remains to show that $A \leq_T T_K$. Given a finite subset $S = \{b_1, \ldots, b_k\} \subseteq \overline{K}$, we decide whether $S \in A$ as follows. First, we search for a nonzero $h \in g(K)[X_1, \ldots, X_k]$ with $h(b_1, \ldots, b_k) = 0$. If we find such an h, we conclude that $s \notin A$. Simultaneously, we search for a constraint $(p, q) \in \overline{T_K}$ with $r = \mathrm{ord}(p) \geq k$, an $x \in \overline{K}$ satisfying (p, q), and elements $y_{k+1}, \ldots y_r$ of $\overline{K}$ such that:

- each of $x, \delta x, \ldots, \delta^{r-1}x$ is algebraic over the field generated over $g(K)$ by the set $S' = \{b_1, \ldots, b_k, y_{k+1}, \ldots, y_r\}$; and

– each element of S' is algebraic over the subfield $g(K)(x, \delta x, \ldots, \delta^{r-1}x)$.

Of course, being algebraic over a c.e. subfield of $\overline{K}$ is itself a Σ^0_1 property, so all of this involves a large search. If we do find all these items, we conclude that $s \in B$.

Now a polynomial h as described above exists iff S is algebraically dependent over $g(K)$. We must show that S is algebraically independent iff the second alternative holds. The backwards direction is quick: if we find all the required elements, then $\{x, \delta x, \ldots, \delta^{r-1}x\}$ is algebraically independent over $g(K)$ (since x satisfies the constraint (p, q) and $\mathrm{ord}(p) = r$) and algebraic over $g(K)(S')$, yet S' has only r elements itself, hence must be algebraically independent over $g(K)$. In particular, its subset S is algebraically independent, as required. It remains to show the forwards direction.

Now S generates a differential subfield F of $\overline{K}$, and by Theorem 7, since δ is assumed to be nonzero, F is generated as a differential field over $g(K)$ by a single element $x \in F$. Let (p, q) be any constraint satisfied by x. If S is indeed algebraically independent over $g(K)$, then F has finite transcendence degree $\geq k$ over $g(K)$, and since $\{x, \delta x, \ldots, \delta^{\mathrm{ord}(p)-1}x\}$ forms a transcendence basis for F (as a field) over $g(K)$, we know that $\mathrm{ord}(p) \geq k$. Moreover, S must extend to a transcendence basis $S \cup \{y_{k+1}, \ldots, y_{\mathrm{ord}(p)}\}$ of F over $g(K)$. This yields all the elements needed for the second alternative to hold. □

We note that the existence of a nonzero derivation δ was used only in the proof that $A \leq_T T_K$. In particular, the Rabin image $g(K)$ is computable in a T_K-oracle, regardless of the derivation. This means that $(g(K) \cap C_{\overline{K}})$ is a T_K-computable subfield of the constant field $C_{\overline{K}}$, which in turn is a computable subfield of $\overline{K}$. Indeed, the restriction of g to C_K is a Rabin embedding of the computable field C_K into its algebraic closure $C_{\overline{K}}$, in the sense of Theorem 4, the original theorem of Rabin for fields.

Therefore, if C is any computable field without a splitting algorithm, we can set $K = C$ to be a differential field with $C_K = C$ (by using the zero derivation). Theorem 6 gives a Rabin embedding g of this differential field K into a computable presentation of $\overline{K}$. Theorem 4 shows that $g(K) = g(C_K)$ is noncomputable within the computable subfield $C_{\overline{K}}$, and therefore must be noncomputable within $\overline{K}$ itself. Finally, Theorem 10 shows that the constraint set $\overline{T_K}$ of this differential field K was noncomputable.

So there do exist computable differential fields, even with the simplest possible derivation, for which the constraint set is noncomputable. In the opposite direction, it is certainly true that if K itself is already differentially closed, then its constraint set is computable, since the constraints are exactly those (p, q) with $p(Y)$ of order 0 and linear in Y. (Such a pair is satisfied by exactly one $x \in K$, hence by exactly one $x \in \overline{K} = g(K)$, using the identity function as the Rabin embedding g. Thus it trivially satisfies Definition 5.) We do not yet know any examples of computable differential fields which have computable constraint set, yet are not differentially closed. The decidability of the constraint set is a significant open problem for computable differential fields in general. So likewise

is the decidability of constrainability: for which $p \in K\{Y\}$ does there exist a q with $(p, q) \in \overline{T_K}$? The comments in the proof of Theorem 10 make it clear that $p(Y) = \delta Y$ is an example of a differential polynomial which is not constrainable.

5 Decidability in the Rabin Image

Rabin's Theorem for fields, stated above as Theorem 4, gave the Turing equivalence of the Rabin image $g(F)$ and the splitting set S_F. Our principal analogue of S_F for differential fields K is the set T_K, and Theorem 10 makes some headway in identifying sets, including $g(K)$ but not only that set, whose join is Turing-equivalent to T_K. It is also natural to ask about Rabin's Theorem from the other side: what set (or what join of sets) must be Turing-equivalent to $g(K)$? We now present one step towards an answer to that question, using the notion of a *linear* differential polynomial in $K\{Y\}$. Recall that "linear" here is used in exactly the sense of field theory: the polynomial has a constant term, and every other term is of the form $a\delta^i Y$, for some $a \in K$ and $i \in \omega$. If the constant term is 0, then the polynomial is *homogeneous* as well, every term having degree 1. The solutions in $\overline{K}$ of a homogeneous linear polynomial $p(Y)$ of order r are well known to form an r-dimensional vector space over the constant field $C_{\overline{K}}$. By additivity, the solutions in $\overline{K}$ to any linear polynomial of order r then form the translation of such a vector space by a single root x of $p(Y)$. Of course, not all of the solutions in $\overline{K}$ need lie in K: the solutions to $p(Y) = 0$ in K (if any exist) form the translation of a vector space over C_K of dimension $\leq r$

Proposition 11 *In a computable differential field K whose field C_K of constants is algebraically closed, the full linear root set FR_K:*

$$\{\textit{linear } p(Y) \in K\{Y\} : p(Y) = 0 \textit{ has solution space in } K \textit{ of dim ord}(p)\},$$

is computable from an oracle for the image $g(K)$ of K in any computable differential closure $\overline{K}$ of K under any Rabin embedding g. Moreover, the Turing reduction is uniform in indices for $\overline{K}$ and g.

Proof. To begin with, suppose that $x \in \overline{K}$ is a constant. Now x satisfies some constraint $(p, q) \in \overline{T_K}$, and with $\delta x = 0$, $p(Y)$ either is the polynomial δY or else has order 0. In the latter case, x is algebraic over K, and indeed turns out to be algebraic over the constant subfield C_K, hence lies in C_K. But $p(Y) = \delta Y$ is impossible: there is no $q(Y)$ such that $(\delta Y, q) \in \overline{T_K}$. (The infinitely many elements of C_K all are roots of δY, yet all satisfy distinct types over K, so in order for $(\delta Y, q)$ to lie in $\overline{T_K}$, all those elements would have to be roots of $q(Y)$, yet q must be nonzero with order < 1.) So every constant in $\overline{K}$ lies in the image $g(C_K)$. It follows that if a linear $p(Y) \in K\{Y\}$ lies in FR_K, then the entire solution space of $p(Y) = 0$ in $\overline{K}$ is contained within K.

Now, given as input a linear $p \in K\{Y\}$, say of order r and with constant term $z \in K$, write $p_0(Y) = p(Y) - z$, so $p_0(Y)$ is homogeneous. We search through $\overline{K}$ until we find:

- an $x \notin g(K)$ with $p(x) = 0$; or
- an $x \in g(K)$ with $p(x) = 0$ *and* $\mathrm{ord}(p)$-many solutions $x_0, \ldots, x_{r-1}$ in $g(K)$ to the homogeneous equation $p_0(Y) = 0$, such that $\{x_0, \ldots, x_{r-1}\}$ is linearly independent over $C_{\overline{K}}$.

Deciding linear independence over $C_{\overline{K}}$ is not difficult: one simply checks whether the Wronskian matrix $(\delta^i x_j)_{i,j<r}$ has determinant 0.

Recall that $C_{\overline{K}} = C_K$. In the former case, therefore, K cannot contain a space of solutions to $p(Y) = 0$ of dimension $\mathrm{ord}(p)$; while in the latter case, K must have a solution space of that dimension $\mathrm{ord}(p)$, i.e. a complete solution space. □

It would be of interest to try to extend this result to the case where C_K need not be algebraically closed, and/or to the situation involving the differential closure of an extension of K by finitely many (or possibly infinitely many) algebraically independent constants.

References

1. L. Blum; Differentially closed fields: a model theoretic tour, in *Contributions to Algebra*, eds. H. Bass, P. Cassidy, & J. Kovacic (New York: Academic Press, 1977).
2. M.D. Fried & M. Jarden, *Field Arithmetic* (Berlin: Springer-Verlag, 1986).
3. A. Frohlich & J.C. Shepherdson; Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407–432.
4. V.S. Harizanov; Pure computable model theory, *Handbook of Recursive Mathematics*, vol. 1 (Amsterdam: Elsevier, 1998), 3–114.
5. L. Harrington; Recursively presentable prime models, *The Journal of Symbolic Logic* **39** (1974) 2, 305–309.
6. E.R. Kolchin; Extensions of Differential Fields, I, *Annals of Mathematics* **43** (1942) 4, 724–729.
7. D. Marker; *Model Theory: An Introduction* (Springer, 2002).
8. R.G. Miller, Computable fields and Galois theory, *Notices of the American Mathematical Society* **55** (August 2008) 7, 798–807.
9. R.G. Miller; Is it harder to factor a polynomial or to find a root?, *Transactions of the American Mathematical Society*, **362** (2010) 10, 5261-5281.
10. R.G. Miller; An Introduction to Computable Model Theory on Groups and Fields, to appear in *Groups, Complexity and Cryptology*, May 2011.
11. R.G. Miller, Computability and differential fields: a tutorial, to appear in *Differential Algebra and Related Topics: Proceedings of the Second International Workshop*, eds. L. Guo & W. Sit (World Scientic, 2011), ISBN 978-981-283-371-6.
12. M. Rabin; Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341–360.
13. R.I. Soare; *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).
14. R.M. Steiner; Computable Fields and Weak Truth-Table Reducibility, *Programs, Proofs, Processes – Sixth Conference on Computability in Europe, CiE 2010*, eds. F. Ferreira, B. Löwe, E. Mayordomo, & L.M. Gomes, *LNCS* 6158 (Springer-Verlag, 2010), 394–405.