

Computing Constraint Sets for Differential Fields

Russell Miller¹

*Department of Mathematics, CUNY Queens College, 65-30 Kissena Blvd., Queens, New York 11367, USA
Ph.D. Programs in Mathematics & Computer Science, CUNY Graduate Center, 365 Fifth Avenue, New
York, New York 10016, USA*

Alexey Ovchinnikov²

*Department of Mathematics, CUNY Queens College, 65-30 Kissena Blvd., Queens, New York 11367, USA
Ph.D. Program in Mathematics, CUNY Graduate Center, 365 Fifth Avenue, New York, New York 10016,
USA*

Dmitry Trushin

*Einstein Institute of Mathematics, Edmond J. Safra Campus, Givat Ram, The Hebrew University of
Jerusalem, Jerusalem, 91904, Israel*

Abstract

Kronecker's Theorem and Rabin's Theorem are fundamental results about computable fields F and the decidability of the set of irreducible polynomials over F . We adapt these theorems to the setting of differential fields K , with constrained pairs of differential polynomials over K assuming the role of the irreducible polynomials. We prove that two of the three basic aspects of Kronecker's Theorem remain true here, and that the reducibility in one direction (but not the other) from Rabin's Theorem also continues to hold.

Key words: differential algebra, computable differential fields, differential closure

2010 MSC: primary 12H05, secondary 03D25, 13N10, 34M15

1. Introduction

Differential algebra is the study of differential equations from a purely algebraic standpoint. The differential equations studied use polynomials in a variable Y and its derivatives $\delta Y, \delta(\delta Y), \dots$, with coefficients from a specific field K which admits

Email addresses: Russell.Miller@qc.cuny.edu (Russell Miller), aovchinnikov@qc.cuny.edu (Alexey Ovchinnikov), trushindima@yandex.ru (Dmitry Trushin)

¹R. Miller was partially supported by the NSF grant DMS-1001306, by the *Infinity Project* of the Templeton Foundation at the Centre de Recerca Matemàtica (grant #13152), by the Isaac Newton Institute as a Visiting Fellow, by the European Science Foundation (grant #4610), and by grants #62632-00 40 and #63286-00 41 from the PSC-CUNY Research Award Program.

²A. Ovchinnikov was supported by the grants: NSF CCF-0952591 and PSC-CUNY #60001-40 41.

³The extended abstract [16] presented proofs of results in Section 6; those proofs are omitted here. Much of Section 1 appeared as an extended abstract in the unpublished proceedings of a weeklong workshop at the Mathematisches Forschungsinstitut Oberwolfach.

differentiation on its own elements via the operator δ . Such a field K is known as a *differential field*: it is simply a field with one or more additional unary functions δ on its elements, satisfying the usual properties of derivatives: $\delta(x + y) = (\delta x) + (\delta y)$ and $\delta(x \cdot y) = (x \cdot \delta y) + (y \cdot \delta x)$.

It is therefore natural to think of the field elements as functions, and standard examples include the field $\mathbb{Q}(X)$ of rational functions in one variable under differentiation $\frac{d}{dX}$, and the field $\mathbb{Q}(t, \delta t, \delta^2 t, \dots)$ with a *differential transcendental* t satisfying no differential equation over the ground field \mathbb{Q} . Additionally, every field becomes a differential field when we set $\delta x = 0$ for all x in the field; we call such a differential field a *constant field*, since an element whose derivative is 0 is commonly called a *constant*. We give a good deal of further background on computability and on differential fields in Sections 2, 3, and 4. First, though, without dwelling on formal definitions, we summarize the situation addressed in this article.

Although the natural examples are fields of functions, the treatment of differential fields regards the field elements merely as points. There are strong connections between differential algebra and algebraic geometry, with such notions as the ring $K\{Y\}$ of differential polynomials (namely the algebraic polynomial ring $K[Y, \delta Y, \delta^2 Y, \dots]$, with each $\delta^i Y$ treated as a separate variable), differential ideal, differential variety, and differential Galois group all being direct adaptations of the corresponding notions from field theory.

Characteristically, these concepts behave similarly in both areas, but the differential versions are often a bit more complicated. In terms of model theory, the theories \mathbf{ACF}_0 and \mathbf{DCF}_0 (of algebraically closed fields and differentially closed fields, respectively, of characteristic 0) are both complete and ω -stable with effective quantifier elimination, but \mathbf{ACF}_0 has Morley rank 1, whereas \mathbf{DCF}_0 has Morley rank ω . The higher Morley rank has elicited intense interest in \mathbf{DCF}_0 from model theorists.

Just as the algebraic closure \bar{F} of a field F (of characteristic 0) can be defined as the prime model of the theory $\mathbf{ACF}_0 \cup \Delta(F)$ (where $\Delta(F)$ is the atomic diagram of F), the differential closure \hat{K} of a differential field K is normally taken to be the prime model of $\mathbf{DCF}_0 \cup \Delta(K)$. This \hat{K} is unique up to isomorphism over K , but not always minimal: it is possible for \hat{K} to embed into itself over K (i.e., fixing K pointwise) with image a proper subset of itself. This has to do with the fact that some 1-types over K are realized infinitely often in \hat{K} , so that the image of the embedding can omit some of those realizations.

As a prime model, the differential closure realizes exactly those 1-types which are principal over K , i.e., generated by a single formula with parameters from K . It therefore omits the type of a differential transcendental over K , since this type is not principal, and so every element of \hat{K} satisfies some differential polynomial over K . On the other hand, the type of a *transcendental constant*, i.e., an element x with $\delta x = 0$ but not algebraic over K , is also non-principal and hence is also omitted, even though such an element would be “differentially algebraic” over K .

The goal of this article is to adapt the two fundamental theorems from computable field theory to computable differential fields. These two theorems, each used very frequently in work on computable fields, are the following.

Theorem 1.1 (Kronecker’s Theorem (1882); see [11] or [2], or Theorem 3.3 below).

1. *The field \mathbb{Q} has a splitting algorithm. That is, the set of irreducible polynomials in $\mathbb{Q}[X]$, commonly known as the splitting set of \mathbb{Q} , is decidable.*
2. *If a computable field F has a splitting algorithm, so does the field $F(x)$, for every element x algebraic over F (within a larger computable field).*
3. *If a computable field F has a splitting algorithm, then so does the field $F(t)$, for every element t transcendental over F .*

(The algorithms deciding irreducibility in Parts II and III are different, and no unifying algorithm exists.)

Theorem 1.2 (Rabin’s Theorem (1960); see [19], or Theorem 3.5 below).

1. *Every computable field F has a Rabin embedding, i.e., a computable field embedding $g : F \rightarrow E$ such that E is a computable, algebraically closed field which is algebraic over the image $g(F)$.*
2. *For every Rabin embedding g of F , the image $g(F)$ is Turing-equivalent to the splitting set S_F of F .*

For differential fields, the analogue of the first part of Rabin’s Theorem was proven in 1974 by Harrington [6], who showed that for every computable differential field K , there is a computable embedding g of K into a computable, differentially closed field L such that L is a differential closure of the image $g(K)$. Harrington’s proof used a different method from that of Rabin, and therefore did not address the question of the Turing degree of the image. Indeed, the first question to address, in attempting to adapt either of these theorems for differential fields, is the choice of an appropriate analogue for the splitting set S_F in the differential context.

Kronecker, like many others before and since, saw the question of reducibility of a polynomial in $F[X]$ as a natural question to ask, with applications in a broad range of areas. However, with twentieth century model theory, we can specify more exactly the reasons why it is important. Specifically, every irreducible polynomial $p(X) \in F[X]$ generates a principal type over the theory $\mathbf{ACF}_0 \cup \Delta(F)$, and every principal type is generated by a unique monic irreducible polynomial. (More exactly, the formula $p(X) = 0$ generates such a type.)

On the other hand, no reducible polynomial generates such a type (with the exception of powers $p(X)^n$ of irreducible polynomials, in which case $p(X)$ generates the same type). So the splitting set S_F gives us a list of generators of principal types, and every element of \overline{F} satisfies exactly one polynomial on the list. Moreover, since these generating formulas are quantifier-free, we can readily decide whether a given

element satisfies a given formula from the list or not. Thus, a decidable splitting set allows us to identify elements of \overline{F} very precisely, up to their orbit over F .

From model theory, we find that the set $\overline{T_K}$ of *constrained pairs* over a differential field K plays the same role for the differential closure. A pair (p, q) of differential polynomials from $K\{Y\}$ is *constrained* if

- p is monic and irreducible and of greater order than q (i.e., for some r , $p(Y)$ involves $\delta^r Y$ nontrivially while $q(Y) \in K[Y, \delta Y, \dots, \delta^{r-1} Y]$), and
- for every $x, y \in \hat{K}$, if $p(x) = p(y) = 0$ and $q(x) \neq 0 \neq q(y)$, then for every $h \in K\{Y\}$, either $h(x) = 0 = h(y)$ or $h(y) \neq 0 \neq h(x)$. This says that, if x and y both *satisfy* the pair (p, q) , then the differential fields $K\langle x \rangle$ and $K\langle y \rangle$ that they generate within \hat{K} must be isomorphic, via an isomorphism fixing K pointwise and mapping x to y .

This is sufficient to ensure that the formula $p(Y) = 0 \neq q(Y)$ generates a principal type over $\mathbf{DCF}_0 \cup \Delta K$, and conversely, every principal type is generated by such a formula with (p, q) a constrained pair.

For these reasons, model theorists have come to see constrained pairs as the appropriate differential analogue of irreducible polynomials. Computability theorists have reserved their judgment. They would go along with the model theorists if it could be shown that every other list of existential generators of the principal types over $\mathbf{DCF}_0 \cup \Delta K$ has complexity $\geq_T T_K$, or even if this could be shown at least for those computable differential fields appearing commonly in differential algebra; but this is not known. Essentially, the question turns on whether Kronecker's Theorem can be adapted and proven in the setting of computable differential fields K and constraint sets T_K .

With this background, we may state our results, first addressing Rabin's Theorem and then Kronecker's.

Theorem 1.3 (See Theorem 6.1 and Proposition 6.2). *For every embedding g of a computable differential field as described by Harrington in [6], the image $g(K)$ is Turing-computable from the set $\overline{T_K}$. So too is algebraic independence of finite tuples from \hat{K} , and also the function mapping each $x \in \hat{K}$ to its minimal differential polynomial over K . However, there do exist such embeddings g for which $\overline{T_K}$ has no Turing reduction to $g(K)$.*

Theorem 1.4 (See Theorem 8.6 and Theorem 9.6). *Let K be a computable non-constant differential field, and let z be an element of a larger computable differential field $L \supseteq K$ such that K is computably enumerable within L . (So $K\langle z \rangle$ is also c.e. within L , and thus has a computable presentation.)*

- *If z is constrained over K , then $\overline{T_{K\langle z \rangle}}$ is Turing-computable from $\overline{T_K}$.*
- *If z is differentially transcendental and \hat{K} is not algebraic over K , then $\overline{T_{K\langle z \rangle}}$ is Turing-computable from $\overline{T_K}$.*

So the last two parts of Kronecker's Theorem hold in most cases. The first part remains open: it is unknown whether the set $\overline{T_{\mathbb{Q}}}$ of constrained pairs over the constant differential field \mathbb{Q} is decidable. We regard this as the most important question currently open in this area of study. A positive answer would likely give us a much better intuition about the structure of various simple differentially closed fields, well beyond any current understanding. It would also be desirable to make the failure of the second part of Rabin's Theorem more precise, by finding sets which are always equivalent to the Rabin image $g(K)$, and by finding sets which are always equivalent to $\overline{T_K}$.

Acknowledgments

We are grateful to P. Cassidy, D. Marker, M. Singer, W. Sit, and the referee for their helpful suggestions.

2. Background in Computability

We recall here the concepts from computability theory which will be essential to our work on differential fields. Computable functions are defined in [25], and indeed, several very different definitions give rise to the same class of functions. Functions on the set ω of nonnegative integers are usually identified with their graphs in ω^2 , and we then code ω^2 into ω , so that the graph corresponds to a subset of ω . Conversely, for our purposes, a subset A of ω may be identified with its characteristic function χ_A , and we say that A is *computable* (or *decidable*) if the function χ_A is computable.

The partial computable functions (those for which the computation procedure halts on certain inputs from ω , but not necessarily on all of them) can be enumerated effectively, and are usually denoted as ϕ_0, ϕ_1, \dots , with the index e coding the program for computing $\phi_e(x)$ on $x \in \omega$. The domains of these functions constitute the *computably enumerable sets*, and we write W_e for the domain of ϕ_e . These are precisely the sets which are definable by Σ_1^0 formulas, i.e., sets of the form

$$\{x \in \omega : \exists y_1 \cdots \exists y_m (x, y_1, \dots, y_m) \in R\},$$

where $m \in \omega$ is arbitrary and R may be any computable subset of ω^{m+1} . We usually write " $\phi_e(x) \downarrow = y$ " to indicate that the computation of ϕ_e on input x halts and outputs y , and so $\phi_e(x) \downarrow$ iff $x \in W_e$; otherwise we write $\phi_e(x) \uparrow$. Also, if the computation halts within s steps, we write $\phi_{e,s}(x) \downarrow$. The set $W_{e,s}$ is the domain of $\phi_{e,s}$, so $W_e = \cup_s W_{e,s}$. Every set $W_{e,s}$ is computable (although the union W_e may not be), and we take it as a convention of our computations that only numbers $\leq s$ lie in $W_{e,s}$.

More generally, we define the Σ_n^0 formulas by induction on n . The Σ_0^0 formulas are those formulas with free variables x_1, \dots, x_m which define computable subsets of ω^m (for any $m \in \omega$). A Π_n^0 formula is the negation of a Σ_n^0 formula. Therefore,

the Π_0^0 formulas are exactly the Σ_0^0 formulas, but for $n > 0$ this is no longer true: for instance a Π_1^0 formula is universal, in the same sense that a Σ_1^0 formula is existential. A Σ_{n+1}^0 formula in the variable x is a formula of the form

$$\exists y_1 \cdots \exists y_m R(x, y_1, \dots, y_m),$$

where R is a Π_n^0 formula. Thus the subscript counts the number of quantifier alternations. (We sometimes omit the superscript 0, which refers to the fact that we quantify only over natural numbers, not over sets of naturals, or sets of sets of naturals, etc.)

Turing reducibility and 1-reducibility are ways of comparing the complexity of subsets $A, B \subseteq \omega$. We refer the reader to [25] for the definition of Turing reducibility, and write $A \leq_T B$ to denote that under this reducibility, A is no more complex than B . 1-Reducibility is more simply defined.

Definition 2.1. A set A is 1-reducible to a set B , written $A \leq_1 B$, if there exists a computable injective function f , whose domain is all of ω , such that

$$(\forall x)[x \in A \iff f(x) \in B].$$

It is well known that, for every $n \in \omega$, there exists a set S which is Σ_{n+1}^0 -complete: S itself is Σ_{n+1}^0 , and every Σ_{n+1}^0 set T has $T \leq_1 S$. Indeed, the Halting Problem, written here as \emptyset' , is Σ_1^0 -complete. The set \emptyset'' is the halting problem relative to \emptyset' , and is Σ_2^0 -complete, and one iterates this *jump* operation, always taking the halting problem relative to the previous set, to get the Σ_{n+1}^0 -complete set \emptyset'''''' , or $\emptyset^{(n+1)}$. Likewise, the complement of a Σ_{n+1}^0 -complete set S is Π_{n+1}^0 -complete. This is regarded as an exact assessment of the complexity of S ; among other things, Σ_{n+1}^0 -completeness ensures that S is *not* Π_{n+1}^0 , nor Σ_n^0 .

Note that the class of Σ_0^0 sets and the class of Π_0^0 sets coincide: these are the computable sets. For $n > 0$, a set which is both Σ_n^0 and Π_n^0 is said to be Δ_n^0 . The Δ_{n+1}^0 sets are exactly those which are Turing-reducible to a Σ_n^0 -complete oracle set. As a canonical Σ_n^0 -complete set, we usually use $\emptyset^{(n)}$, the n -th jump of the empty set, as defined in [25].

Turing reducibility \leq_T is a partial pre-order on the power set $\mathcal{P}(\omega)$. We define $A \equiv_T B$, saying that A and B are *Turing-equivalent*, if $A \leq_T B$ and $B \leq_T A$. The equivalence classes under this relation form the *Turing degrees*, and are partially (but not linearly) ordered by \leq_T . In fact, they form an upper semi-lattice under \leq_T , with least element $\mathbf{0}$, the degree of the computable sets, but no greatest element. One often speaks of a set A as being *computable in a Turing degree* \mathbf{d} , meaning that for some (equivalently, for every) $B \in \mathbf{d}$ we have $A \leq_T B$.

3. Background on Fields

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

Definition 3.1. 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 :

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

Fröhlich and Shepherdson were the first to consider computable algebraically closed fields, in [4]. 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 3.2. 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,

$$R_F = \{p(X) \in F[X] : (\exists a \in F) p(a) = 0\}$$

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

If S_F is computable, F is said to have a *splitting algorithm*.

With F computable, R_F and S_F must both be computably enumerable, being defined by existential conditions. Theorem 3.5 will show them to be Turing-equivalent. For most computable fields one meets, both are computable; the first steps in this direction were taken by Kronecker in 1882.

Theorem 3.3 (Kronecker's Theorem; see [11]).

- (i) \mathbb{Q} has a splitting algorithm.
- (ii) Let L be a c.e. subfield of a computable field F . If L has a splitting algorithm, then for every $x \in F$ algebraic over L , $L(x)$ also has a splitting algorithm. The specific decision procedure for $S_{L(x)}$ can be determined from that for S_L and from the minimal polynomial of x over L .
- (iii) Let L be a c.e. subfield of a computable field F . If L has a splitting algorithm, then for any $x \in F$ transcendental over L , $L(x)$ also has a splitting algorithm, which can be determined just from that for L , given that x is transcendental.

More generally, for any c.e. subfield L of a computable field F and any $x \in F$, the splitting set of $L(x)$ is Turing-equivalent to the splitting set for L , via reductions uniform in x and in the minimal polynomial of x over L (or in the knowledge that x is transcendental).

The algorithms for algebraic and transcendental extensions are different, so it is essential to know whether x is algebraic. If it is, then from S_L one can determine the minimal polynomial of x . This yields the following.

Lemma 3.4. *For every computable field F algebraic over its prime subfield P , there is a computable function which accepts as input any finite tuple $\vec{x} = \langle x_1, \dots, x_n \rangle$ of elements of F and outputs an algorithm for computing the splitting set for the subfield $P[\vec{x}]$ of F . (We therefore say that the splitting set of $P[\vec{x}]$ is computable uniformly in \vec{x} .)*

Proof. Clearly there are splitting algorithms for all finite fields, just by checking all possible factorizations. (So in fact there is a single algorithm which works in all positive characteristics.) In characteristic 0, one can readily compute the unique isomorphism onto the prime subfield P of F from the computable presentation of \mathbb{Q} for which Kronecker's splitting algorithm works, and this computable isomorphism allows us to compute the splitting set of P .

The lemma then follows by induction on the size of the tuple $\vec{x} = \langle x_1, \dots, x_n \rangle$, using part (ii) of Theorem 3.3. Since our F is algebraic over P , we may simply search for a polynomial $p(X)$ with root x_n and coefficients in $P[x_0, \dots, x_{n-1}]$, and then factor it, using the splitting algorithm for $P[x_0, \dots, x_{n-1}]$ (by inductive hypothesis), until we have found the minimal polynomial of x_n over $P[x_0, \dots, x_{n-1}]$. \square

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 [14, Lemma 7] for a simple example. Fröhlich and Shepherdson showed that R_F is computable if and only if 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 his proof there.) More recent works [15, 26] 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 3.5 (Rabin's Theorem; see [19]). *For every computable field F , there exist an algebraically closed computable field E and a computable field homomorphism $g : F \rightarrow E$ 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 $g : F \rightarrow E$ 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.

When we begin to consider polynomials in several variables, the connection between R_F and S_F breaks down.

Theorem 3.6 (see [3]). *Suppose that F is a computable field. Then for every n , the set of irreducible polynomials in the ring $F[X_0, \dots, X_n]$ is computable in an oracle for the splitting set S_F , via a Turing reduction uniform in n .*

Hence, in a computable differential field K , it is decidable in S_K whether a differential polynomial $p(Y) \in K\{Y\}$, viewed as an algebraic polynomial over K in $Y, \delta Y, \delta^2 Y, \dots$, is irreducible. (Differential polynomials are described in the next section.)

For a proof, see [3, §19]. In contrast, the decidability of the existence of rational solutions to arbitrary polynomials in $\mathbb{Q}[X_1, X_2, \dots]$ remains an open question: this is *Hilbert's Tenth Problem for \mathbb{Q}* , the subject of much study. At present, it is not clear that these questions impinge on single-variable problems for differential fields, but since a differential polynomial can be viewed as an algebraic polynomial in several variables, it is not implausible that a connection might exist.

We now turn to questions of algebraic dependence in fields. Predictably, these issues are closely tied to transcendence bases.

Definition 3.7. For a computable field F with computably enumerable subfield E , the *algebraic dependence set* $A_{F/E}$ is the set of all finite tuples of F algebraically dependent over E :

$$A_{F/E} = \{(x_1, \dots, x_n) \in F^{<\omega} : (\exists \text{ nonzero } p \in E[X_1, \dots, X_n]) p(\vec{x}) = 0\}.$$

Likewise, for any computable F -vector space V (including computable field extensions of F), the *linear dependence set* is

$$L_V = \left\{ \{v_1, \dots, v_n\} \subseteq V : (\exists \langle c_1, \dots, c_n \rangle \in F^n - \{0, \dots, 0\}) \sum_{i \leq n} a_i v_i = 0 \right\}.$$

Below, when considering a differential field K within its differential closure \hat{K} , we will often want to consider this set for \hat{K} over K , and we will write

$$D_K = A_{\hat{K}/K}$$

for the set of all finite subsets of \hat{K} algebraically dependent over K . The following lemma is considered in more depth in [13].

Lemma 3.8. *For every computable field F and every computably enumerable subfield E , there is a transcendence basis $B_{F/E}$ for F over E computable using the set $A_{F/E}$ as an oracle. Conversely, for every transcendence basis B for F over E , we have $A_{F/E} \leq_T B$.*

Proof. We define the *canonical transcendence basis* $B_{F/E}$ for F over E as $\cup_s B_s$, where $B_0 = \emptyset$ and

$$B_{s+1} = \begin{cases} B_s \cup \{x_s\}, & \text{if this set is algebraically independent over } E, \\ B_s, & \text{if not.} \end{cases}$$

Clearly $B_{F/E} \leq_T A_{F/E}$ (and so, by the next paragraph, $B_{F/E} \equiv_T A_{F/E}$).

If F has finite transcendence degree over E , then every transcendence basis is computable, with no oracle at all. So we assume the transcendence degree to be infinite. If B is a transcendence basis and

$$S = \{x_1, \dots, x_n\} \subseteq F,$$

then S is algebraically independent iff there exist an $r \geq n$, an r -element subset $B_0 = \{b_1, \dots, b_r\} \subseteq B$, and elements $y_{n+1}, \dots, y_r \in F$ such that every element in each of the sets B_0 and

$$\{x_1, \dots, x_n, y_{n+1}, \dots, y_r\}$$

is algebraic over the other set. This condition is Σ_1^0 relative to B . Of course, algebraic dependence of S is Σ_1^0 (without any oracle), so membership of S in $A_{F/E}$ is decidable from B . \square

An exactly analogous proof likewise holds for bases of vector spaces.

Lemma 3.9. *For every computable field F and every computable F -vector space V (such as a computable field extension of F), the relation L_V of linear dependence is computable relative to any basis for V over F . Conversely, this relation itself computes a basis.*

4. 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 4.1. A *differential field* is a field K with one or more additional unary functions δ_i satisfying the following two axioms for all $x, y \in K$:

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

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 $F(X_1, \dots, X_n)$ 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 .

A differential field K gives rise to a ring $K\{Y\}$ of *differential polynomials*. This can be described as

$$K[Y, \delta Y, \delta^2 Y, \dots],$$

the ring of algebraic polynomials in the infinitely many variables shown. However, for any differential polynomial p and any single $y \in K$, it makes sense to speak of $p(y)$, by which we mean the element

$$p(y, \delta y, \delta^2 y, \dots) \in K$$

calculated using δ and the field operations of K . Likewise, in any differential field extension L of K , $p(x)$ will be an element of L for every $x \in L$.

We can similarly discuss the derivative of a polynomial, bearing in mind that the coefficients are elements of K , not necessarily constants, and may require the Leibniz Rule. For instance, if

$$p(Y) = a(\delta Y)^2 + bY(\delta^2 Y)$$

with $a, b \in K$, then

$$(\delta p)(Y) = (\delta a)(\delta Y)^2 + 2a(\delta Y)(\delta^2 Y) + (\delta b)Y(\delta^2 Y) + b(\delta Y)(\delta^2 Y) + bY(\delta^3 Y).$$

The *order* of $p \in K\{Y\}$ is the greatest $r \geq 0$ such that $\delta^r Y$ appears nontrivially in $p(Y)$. So, in the example above, $p(Y)$ has order 2 and $\delta p(Y)$ has order 3. The algebraic polynomials (in $K[Y]$) of positive degree in Y have order 0. By convention the zero polynomial has order $-\infty$, and all other algebraic polynomials of degree 0 have order -1 .

Just as polynomials in $F[X]$ are ranked by their degree, differential polynomials in $K\{Y\}$ are ranked as well. First, if $p, q \in K\{Y\}$ and $\text{ord}(p) < \text{ord}(q)$, then p has lower rank than q . Second, if $\text{ord}(p) = \text{ord}(q) = r$ but $\delta^r Y$ has lesser degree in p than in q , then p has lower rank than q . This is not the entire notion of rank, but it is as much as we need in this paper: rank is the lexicographic order on the pair $(\text{ord}(p), \deg_{\delta^{\text{ord}(p)} Y}(p))$, hence of order type ω^2 .

The notion of a computable differential field extends that of a computable field in the natural way.

Definition 4.2. A *computable differential field* is a computable field with one or more differential operators δ as in Definition 4.1, each of which is likewise given by some Turing-computable function h with $\delta(x_i) = x_{h(i)}$ (where $\{x_0, x_1, \dots\}$ is again the set of elements of the differential field, as in Definition 3.1).

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 $\text{ord}(p) > \text{ord}(q)$, there exists some $y \in K$ with $p(y) = 0 \neq q(y)$. (By convention, all nonzero constant polynomials have order -1 . Blum's axioms therefore include formulas saying that K is algebraically closed, by taking $q = 1$ and arbitrary nonconstant $p \in K[Y]$.)

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 $z \in L$ is *constrained over K* if z satisfies some *constrained pair* over K , as defined here. The terminology of “principal types” reflects the model theory behind the definition. In fact, the type generated is principal over the theory \mathbf{DCF}_0 .

Definition 4.3. Let K be a differential field. A pair (p, q) of differential polynomials in $K\{Y\}$ *generates a principal type* if, 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 $x_0 \cong_K x_1$. The pair (p, q) is *constrained* for K if p is monic and algebraically irreducible over K , with $\text{ord}(q) < \text{ord}(p)$, and (p, q) generates a principal type. Elements x in an extension of K with $p(x) = 0 \neq q(x)$ are said to *satisfy the constrained pair* (p, q) . We let

$$\begin{aligned} T_K &= \{(p, q) \in (K\{Y\})^2 : (p, q) \text{ is not a constraint}\} \\ &= \left\{ (p, q) : \left(\exists x, y \in \hat{K} \right) (\exists h \in K\{Y\}) [h(x) = 0 \neq h(y) \ \& \ x, y \text{ satisfy } (p, q)] \right\}. \end{aligned}$$

(The second of these equivalent definitions should logically follow Theorem 4.4, where we define the differential closure \hat{K} of K , and Proposition 4.5, which establishes the equivalence.) Thus $\overline{T_K}$ is the *constraint set* for K . If T_K is computable, we say that K has a *constraint algorithm*.

The broad intention is to quantify over all x and y in all differential field extensions of K . However, since the definition considers only those x_0 and x_1 satisfying

the constrained pair, it turns out to be equivalent to quantify over all x and y in a differential closure \hat{K} of K . Once we state Harrington's Theorem below, the quantification will just be over elements of ω , and so the second definition of T_K above uses a Σ_1^0 formula, provided that K is computable. (The first definition can also be seen directly to be Σ_1^0 : the model-theoretic notion of generating a principal type is Π_1^0 , since the theory $\mathbf{DCF}_0 \cup \Delta(K)$ described below is complete and decidable.) This was our reason for defining T_K to be the *complement* of the constraint set: we thus parallel the notation R_F and S_F for the corresponding sets for fields, in that all are existential definitions. (For this purpose we eschew the symbol C_K , which is already widely used to denote the constant subfield of K .)

Definition 4.3 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 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_1^0 sets, given that K and F are both computable, and both are the negations of the properties we need to produce isomorphic extensions.

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 r of any nonzero element p of $K\{Y\}$ for which x is a zero, and that

$$\{x, \delta x, \dots, \delta^{r-1}x\}$$

forms a transcendence basis for $K\langle x \rangle$ as a field extension of K . The unique irreducible p of smallest order is called the *minimal differential polynomial of x* , and its order r is called the *order of x* . (Fact 5.1 and Definition 5.2 will elaborate on this. For more general results, see [9, II.12, Theorem 6(d)], and also [17, Lemma 6.12] for ordinary differential fields.) 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 \bar{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 .

Model-theoretically, the two situations are closely analogous: the algebraic closure \bar{F} of F (of characteristic 0) is the prime model of the theory $\mathbf{ACF}_0 \cup \Delta(F)$, given by extending the language to include constants for all elements of F and adjoining to \mathbf{ACF}_0 the atomic diagram $\Delta(F)$. Likewise, the differential closure \hat{K} of

K is the prime model of $\mathbf{DCF}_0 \cup \Delta(K)$.

As with fields, the differential closure \hat{K} of K turns out to be unique up to isomorphism over K . On the other hand, \hat{K} is generally not *minimal*: there exist differential field embeddings of \hat{K} into itself over K whose images are proper differential subfields of \hat{K} . This provides a first contrast between \mathbf{DCF}_0 and \mathbf{ACF}_0 , since the corresponding statement about algebraic closures is false.

Theorem 4.4 ([10, 22, 24]). *For every differential field K of characteristic 0, there exists a differential field extension $\hat{K} \supseteq K$ which satisfies the axiom set \mathbf{DCF}_0 and has the property that every element of \hat{K} satisfies some constrained pair from $\overline{T_K}$. We refer to \hat{K} as the differential closure of K , since it can be shown to be unique up to isomorphisms which fix K pointwise.*

Our next fact was already cited in Definition 4.3 and is standard.

Proposition 4.5. *Let x and y lie in any two differential field extensions of K . Then $x \cong_K y$ iff, for all $h \in K\{Y\}$, we have $h(x) = 0$ iff $h(y) = 0$.*

With this much information in hand, we can now state the parallel to the first half of Theorem 3.5.

Theorem 4.6 ([6, Corollary 3]). *For every computable differential field K , there exists a differentially closed computable differential field L and a computable differential field homomorphism $g : K \rightarrow L$ 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 4.6 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 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 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 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 $\overline{T_K}$, from Definition 4.3, which is analogous in several ways to the splitting set and will be the focus of our attempts to adapt Kronecker's Theorem (Theorem 3.3 above) to differential fields.

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

Theorem 4.7 ([8, 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 \hat{F} with $F \subset E$, and E is generated (as a differential field over F) by finitely many elements, then there is a single element of E that generates all of E as a differential field over F .*

Kolchin extended this theorem to partial differential fields with m derivations: the generalized condition there is the existence of m elements whose Jacobian is nonzero. He offered counterexamples in the case where δ is the zero derivation on F , but in the counterexamples, the generators of E are constants which are algebraically independent over F and therefore do not lie in \hat{F} , although they are differentially algebraic over F . (At that time, the definition of differential closure was not yet well established.) It remains open whether the theorem as stated here holds for constant differential fields as well.

One sees readily that Theorem 4.7 can be carried out effectively: given an enumeration of F within E , along with the finitely many generators of E over F , it is easy to find a single generator as described in the theorem, simply by enumerating the elements generated over F by each single $x \in E$ until some such x is seen to generate all of the finitely many given generators of E . Of course, without Theorem 4.7, one would not be sure whether this process would ever halt.

5. Constrained Pairs

For fields, the usefulness of the set S_F of reducible polynomials in S_F is that the irreducible polynomials in $F[X]$ exactly define the isolated 1-types over $\mathbf{ACF}_0 \cup \Delta(F)$. That is, if x and y lie in field extensions of F and satisfy the same irreducible $p(X) \in F[X]$, then they generate isomorphic subfields: $F(x) \cong F(y)$, via an isomorphism fixing F pointwise and mapping x to y . (Reducible polynomials in $F[X]$ fail to have the same property, except for the special case of a power of an irreducible polynomial, which can be recognized effectively using the formal derivative and the Euclidean algorithm.) Moreover, satisfaction of a particular irreducible polynomial $p(X)$ by a particular x is decidable, in a computable field extending F : the formula $p(X) = 0$ generating the 1-type is quantifier-free, and the field operations are computable. Of course, since $\mathbf{ACF}_0 \cup \Delta(F)$ has effective quantifier elimination, we could start with any generating formula and find a quantifier-free generating formula.

The important point is that we have a list of formulas, each of which generates a principal type, and such that every principal type is generated by a formula on the list. Theorem 3.3 says that this list can be given effectively for the field \mathbb{Q} , and that the effectiveness carries over to finitely generated computable field extensions. This is the theorem which we wish to adapt for differential fields K , with the constraint

set $\overline{T_K}$ in place of the set $\overline{S_F}$ of irreducible polynomials in $F[X]$. After some further preliminaries, we will prove that the natural adaptations of parts (ii) and (iii) of Theorem 3.3 do indeed carry over to differential fields. Whether part (i) can likewise be adapted remains an open question.

Fact 5.1. *Let K be a differential field. Then for each $x \in \hat{K}$, there is exactly one $p \in K\{Y\}$ such that x satisfies a constrained pair of the form $(p, q) \in \overline{T_K}$. (Recall that by definition p is required to be monic and irreducible.) Moreover, $\text{ord}(p)$ 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 $\delta^{\text{ord}(p)}Y$ in any polynomial in $K\{Y\}$ of order $\text{ord}(p)$ with root x .

Proof. Since \hat{K} is constrained over K , each $x \in \hat{K}$ satisfies at least one constrained pair $(p, q) \in \overline{T_K}$. 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 $y \in \hat{K}$ 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 $K\langle x \rangle \not\cong K\langle y \rangle$. This would contradict Definition 4.3. Hence r is the least order of any nonzero differential polynomial with root x .

It follows from minimality of r that $\{x, \delta x, \dots, \delta^{r-1}x\}$ is algebraically independent over K . Next we claim that the minimal polynomial of $\delta^r x$ over $K(x, \dots, \delta^{r-1}x)$ must equal

$$p(x, \delta x, \dots, \delta^{r-1}x, Y).$$

Clearly the minimal polynomial must divide $p(x, \delta x, \dots, \delta^{r-1}x, Y)$, and if it were a proper factor, then by Blum's axioms for \mathbf{DCF}_0 , the quotient of $p(x, \delta x, \dots, \delta^{r-1}x, Y)$ by the minimal polynomial would have a root y with $q(y) \neq 0$ which would not be a root of the minimal polynomial, so that $K\langle x \rangle \not\cong K\langle y \rangle$, contradicting Definition 4.3. (We used here the fact that with $p(Y, \delta Y, \dots, \delta^r Y)$ irreducible in $K[Y, \delta Y, \dots, \delta^r Y]$,

$$p(x, \delta x, \dots, \delta^{r-1}x, Y)$$

cannot be a power of the minimal polynomial.) Thus p is the *minimal differential polynomial* of x over K . \square

Definition 5.2. If $K \subseteq L$ is an extension of differential fields, then for each $x \in L$, we define $\text{ord}_K(x) = \text{ord}(p)$, where (p, q) is any constrained pair in $\overline{T_K}$ satisfied by x . Notice that in the differential closure of K , every element has a well-defined finite order over K , by Fact 5.1. (A more general definition of order for elements of

differential field extensions sets $\text{ord}_K(x) \leq \omega$ be the transcendence degree of $K\langle x \rangle$ as a field extension of K .)

In fact, the irreducibility of $p(Y)$ is barely necessary in Definition 4.3. The condition that $K\langle x \rangle \cong K\langle y \rangle$ for all x and y satisfying the constrained pair 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 Fact 5.1 would no longer hold.

Proposition 5.3. *Let $p, q, \tilde{q} \in K\{Y\}$, with (p, q) in the constraint set $\overline{T_K}$ and $\text{ord}(\tilde{q}) < \text{ord}(p)$. Then $(p, \tilde{q}) \in \overline{T_K}$ iff, for every $x \in \hat{K}$, x satisfies (p, q) iff x satisfies (p, \tilde{q}) .*

Proof. First suppose $(p, \tilde{q}) \in \overline{T_K}$. If x satisfies (p, q) , then p is the minimal differential polynomial of x over K (by Fact 5.1), and so $\tilde{q}(x) \neq 0$. Likewise, if x satisfies (p, \tilde{q}) , then it satisfies (p, q) as well. Conversely, if the second condition holds, then every x and y satisfying (p, \tilde{q}) also satisfy (p, q) , hence have $x \cong_K y$, putting $(p, \tilde{q}) \in \overline{T_K}$. \square

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 Proposition 5.3 is nontrivial, and the constrained pair satisfied by an $x \in \hat{K}$ is not unique, although its first component is unique. The proposition shows that the first component essentially determines the constrained pair: two constrained pairs with the same first component define the same set. On the other hand, not all monic irreducible differential polynomials p can be the first component of a constrained pair; the polynomial $p(Y) = \delta Y$ is a simple counterexample. In Section 7 we will address the question of *constrainability*: for which $p \in K\{Y\}$ does there exist some $q \in K\{Y\}$ with $(p, q) \in \overline{T_K}$?

6. Decidability in the Constraint Set and Rabin Image

6.1. Decidability in the Constraint Set

The two theorems in this section were proven in [16], a preliminary report on the work in this article. They address the adaptation of Rabin's Theorem to the context of differential fields. The proofs are straightforward, and we do not repeat them here.

Theorem 6.1 ([16, Theorem 10]). *Let K be any computable differential field, and $g : K \rightarrow \hat{K}$ a (differential) Rabin embedding of K . Then all of the following are computable in an oracle for the constraint set $\overline{T_K}$: the splitting set S_K , the Rabin image $g(K)$, and the order function ord_K on \hat{K} . If additionally our derivation is nontrivial, then the set D_K of finite subsets of \hat{K} algebraically dependent over $g(K)$ is also computable in a $\overline{T_K}$ -oracle.*

In particular, the Rabin image $g(K)$ is computable in a T_K -oracle. This means that $(g(K) \cap C_{\hat{K}})$ is a T_K -computable subfield of the constant field $C_{\hat{K}}$, which in turn is a computable subfield of \hat{K} . Indeed, the restriction of g to C_K is a Rabin embedding of the computable field C_K into its algebraic closure $C_{\hat{K}}$, in the sense of Theorem 3.5, 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 4.6 gives a Rabin embedding g of this differential field K into a computable presentation of \hat{K} . Theorem 3.5 shows that $g(K) = g(C_K)$ is noncomputable within the computable subfield $C_{\hat{K}}$, and therefore must be noncomputable within \hat{K} itself. Finally, Theorem 6.1 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 constrained pairs are exactly those pairs of the form $(Y - a, b)$ with $a, b \in K$ and $b \neq 0$. (Such a pair is satisfied by exactly one $x \in K$, hence by exactly one $x \in \hat{K} = g(K)$, using the identity function as the Rabin embedding g . Thus it trivially satisfies Definition 4.3.)

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 $(p, q) \in \overline{T_K}$? We address this question in Section 7. The comments in the proof of Theorem 6.1 in [16] make it clear that $p(Y) = \delta Y$ is an example of a differential polynomial which is not constrainable.

6.2. Decidability in the Rabin Image

Rabin's Theorem for fields, stated above as Theorem 3.5, 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 6.1 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 \hat{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 $C_{\hat{K}}$. By additivity, the solutions in \hat{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 \hat{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 6.2. [16, Proposition 11] *In a computable differential field K whose field C_K of constants is algebraically closed, the full linear root set FR_K :*

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

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

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 algebraically independent constants.

7. The Constrainability Set

We now address the question of whether a given differential polynomial $p(Y) \in K\{Y\}$ is *constrainable* within its computable differential field K . This question asks whether there exists some $q \in K\{Y\}$ such that (p, q) forms a constrained pair, i.e., such that $(p, q) \notin T_K$. Such a q is called a *constraint* on p .

Definition 7.1. For a differential field K , the set of *unconstrainable* differential polynomials is the set

$$U_K = \{p \in K\{Y\} : (\forall q \in K\{Y\})(p, q) \in T_K\}.$$

Any $p(Y) \notin U_K$ is said to be *constrainable*.

As the natural c.e. and co-c.e. sets for fields are all named in their existential forms (R_K , S_K , T_K , and D_K), we propose that related Σ_2^0 and Π_2^0 sets should always be named in their Π_2^0 (that is, $\forall\exists$) forms. As in the case of U_K , this will often involve a single \forall quantifier over an existential set, which makes a simple and natural definition. An alternative definition follows from our next proposition.

Fact 7.2. *A differential polynomial $p \in K\{Y\}$ is constrainable over K iff p is the minimal differential polynomial of some x in the differential closure \hat{K} of K .*

Proof. Every constrained pair (p, q) is satisfied by some $x \in \hat{K}$. Fact 5.1 shows that p must be the minimal differential polynomial of this x , and also shows that the minimal differential polynomial of each $x \in \hat{K}$ is the first component of the constrained pair satisfied by that x . \square

This alternative definition of constrainability is no simpler, however: it says that there exists an x satisfying p such that no polynomial in $K\{Y\}$ of lesser rank than p can have x as a zero. We now show that U_K cannot be defined by any formula simpler than Π_2^0 .

Theorem 7.3. *There exists a computable differential field K in which the set of constrainable polynomials $p \in K\{Y\}$ is Σ_2^0 -complete.*

Proof. This proof evolved out of a suggestion by David Marker. The key to the Π_2^0 -completeness is Theorem 6.2 from [12, p. 73], which introduces the differential polynomials

$$p_e(Y) = \delta Y - t_e(Y^3 - Y^2)$$

we will use, over a ground field $K_0 \cong \mathbb{Q}\langle t_0, t_1, \dots \rangle$ in which the elements t_e form a differentially independent set over \mathbb{Q} . (That is, no tuple $\langle t_0, \dots, t_n \rangle$ is a zero of any nonzero differential polynomial over \mathbb{Q} .) Now $p_e(Y)$ certainly will have zeroes in the differential closure of K_0 (although its only zeroes in K_0 itself are the trivial ones 0 and 1), but, imitating the proof of Corollary 6.3 in [12] with $f(Y) = Y^3 - Y^2$, we see that the set of its nonconstant zeroes is algebraically independent over \mathbb{Q} .

Moreover, in an extension K of K_0 such as we shall build, p_e is constrainable iff K contains only finitely many zeros of p_e . This condition is readily exploited to show Σ_2^0 -completeness of constrainability, which is to say, Π_2^0 -completeness of U_K , using requirements:

$$\mathcal{R}_e : p_e \text{ has infinitely many zeroes in } K \iff e \in \text{Inf}.$$

These will show that the Π_2^0 -complete set $\text{Inf} = \{e : |W_e| = \infty\}$ is 1-reducible to U_K . Since U_K is already known to be Π_2^0 , it will therefore be Π_2^0 -complete, and many computable model theorists could fill in the rest of these details immediately.

The differential field K is built as a c.e. subfield of the differential closure \hat{K}_0 , in stages, as a computably enumerable differential subfield of the (computable) differential closure of $\mathbb{Q}\langle t_0, t_1, \dots \rangle$. Each requirement \mathcal{R}_e will be eligible at infinitely many stages, at each of which it will have the opportunity to add more elements to K if it needs to.

At stage 0, we apply Harrington's Theorem 4.6 to (a computable presentation of) the differential field $\mathbb{Q}\langle t_0, t_1, \dots \rangle$, as described above, with the t_e all differentially independent over \mathbb{Q} . This yields a differential Rabin embedding f of $\mathbb{Q}\langle t_0, t_1, \dots \rangle$ into a computable presentation L of the differential closure of a c.e. subfield K_0 which is the image of $\mathbb{Q}\langle t_0, t_1, \dots \rangle$ under f . We may assume that the function $e \mapsto t_e$ in the original presentation of $\mathbb{Q}\langle t_0, t_1, \dots \rangle$ was computable, and thus that $e \mapsto f(t_e)$ is computable as well. From here on, we deal only with K_0 , writing

$$a_e = f(t_e) \in K_0$$

for the image in K_0 of each differential transcendental t_e . Within L , we can now name the nonconstant zeroes of the polynomials

$$p_e(Y) = \delta Y - a_e(Y^3 - Y^2).$$

There must be infinitely many such zeroes, since L is differentially closed and p_e has order 1, and each list

$$\{z_{e,0} < z_{e,1} < z_{e,2} < \dots\}$$

of all nonconstant zeroes of p_e in L is computable, uniformly in e .

We now move to stage $s + 1$. Here we use the convention that in our enumeration of all the c.e. sets W_e , for each s , there is a unique pair (e, x) for which $x \in W_{e,s+1} - W_{e,s}$ (that is, for which x enters W_e at stage $s + 1$). Of course, $W_{e,0} = \emptyset$ for every e . So, at the stage $s + 1$, we find the unique corresponding pair (e, x) , choose $i = |W_{e,s}|$, and set

$$K_{s+1} = K_s(z_{e,i}) \subset L.$$

Notice that since

$$\delta z_{e,i} = a_e(z_{e,i}^3 - z_{e,i}^2),$$

this K_{s+1} is closed under δ , hence is a differential subfield of L . Lemma 7.4 will show that $z_{e,i}$ is transcendental over K_s as a field (even though it is constrained over K_s as a differential field), making K_{s+1} a purely transcendental field extension.

This completes the construction, and the differential field K **is the union of the fields** K_s defined at each stage. Thus K is a c.e. subfield of L . (A computable presentation of K with domain ω may readily be found, since the infinite c.e. set K is the image of ω under some injective total computable function f . Just pull the differential field operations from K back to ω via f .)

Lemma 7.4. *The set $\{z_{e,i} : e, i \in \omega\}$ is algebraically independent over K_0 .*

Proof. Each $z_{e,i}$ by itself is transcendental over K_0 . A theorem of Rosenlicht which appears as [12, Theorem 6.2, p. 73], along with Example 2 there, shows that if $z = z_{e,i}$ and $\tilde{z} = z_{j,k}$ are algebraically dependent over K_0 , then

$$a_e z^2 = a_j \tilde{z}^2.$$

Applying δ to both sides and using the equations

$$\delta z = a_e(z^3 - z^2) \quad \text{and} \quad \delta \tilde{z} = a_j(\tilde{z}^3 - \tilde{z}^2),$$

we derive a second equation

$$z^2 \delta a_e + 2a_e^2(z^4 - z^3) = \tilde{z}^2 \delta a_j + 2a_j^2(\tilde{z}^4 - \tilde{z}^3).$$

Applying δ and substituting for δz and $\delta \tilde{z}$ again yields a third polynomial relation on $z, \tilde{z}, a_e, \delta a_e, \delta^2 a_e, a_j, \delta a_j,$ and $\delta^2 a_j$. But now these three distinct algebraic relations show that the field

$$\mathbb{Q}(z, \tilde{z}, a_e, \delta a_e, \delta^2 a_e, a_j, \delta a_j, \delta^2 a_j)$$

has transcendence degree at most 5 over \mathbb{Q} . Assuming that $e \neq j$, this contradicts the differential independence of a_e and a_j . If $e = j$, then the equation $a_e z^2 = a_j \tilde{z}^2$ actually showed that $\tilde{z} = \pm z$, and the only way for z and $-z$ both to be zeroes of p_e is for $z = 0 = -z$. Thus, every set $\{z_{e,i}, z_{j,k}\}$ of two distinct elements is algebraically independent over K_0 .

One then argues by induction. Let S be any subset of $\{z_{e,i} : e, i \in \omega\}$ of size $n + 1 \geq 3$. Form

$$S' = S - \{z_{e,i}, z_{j,k}\}$$

by deleting any two elements from S . Then each of $z_{e,i}$ and $z_{j,k}$ is a transcendental over $K_0(S')$, by inductive hypothesis, and the argument from the preceding paragraph, with $K_0(S')$ and $\mathbb{Q}(S')$ in place of K_0 and \mathbb{Q} , shows that $\{z_{e,i}, z_{j,k}\}$ is algebraically independent over $K_0(S')$, making S algebraically independent over K_0 . \square

Lemma 7.5. *In the construction above, the nonconstrainability set U_K is Π_2^0 -complete.*

Proof. If W_e is finite, then by our construction, there are only finitely many stages at which any zero $z_{e,i}$ of the polynomial $p_e(Y)$ is added to K . It follows from Lemma 7.4 that the only elements $z_{e,i}$ in K are the ones we specifically enumerated into K_{s+1} at some stage $s + 1$: Lemma 7.4 shows that no finite set of such elements $z_{e,i}$ generates any $z_{e',i'}$ as a field, except those already in the finite set, and we remarker in the proof that the field generated by this set is closed under differentiation. Therefore, an easy induction on s shows that, for every e and s , p_e has exactly $|W_{e,s}|$ nontrivial zeroes in K_s . But now, if $|W_e| = n$ is finite, then

$$(p_e(Y), Y(Y - 1)(Y - z_{e,0}) \cdot \dots \cdot (Y - z_{e,n-1})) \notin T_K,$$

and so $p_e \notin U_K$.

Conversely, if W_e is infinite, then $p_e(Y)$ has infinitely many zeroes in K , since a new one is added at each stage at which W_e receives a new element. But then there is no $q \in K\{Y\}$ with $(p_e, q) \notin T_K$: such a q would have to have order 0, since p_e has order 1, and hence q would have only finitely many zeroes in K . This would leave infinitely many $z_{e,i}$ in K satisfying (p_e, q) , yet clearly $z_{e,i} \not\equiv_K z_{e,j}$ for all $i \neq j$ with either $z_{e,i}$ or $z_{e,j}$ in K , since no isomorphism fixing K pointwise could map $z_{e,i}$ to

any element except itself. Thus

$$p_e \in U_K \iff e \in \text{Inf},$$

and the computable injective function $e \mapsto p_e$ is a 1-reduction from Inf to U_K , proving the lemma. (Recall that Definition 7.1 shows that U_K is a Π_2^0 set.) \square

This completes the proof of Theorem 7.3. \square

The usual examples of nonconstrainable polynomials are along the lines of $p(Y) = \delta Y$. Based on just this, one might wonder whether the nonconstrainable polynomials in $K\{Y\}$ are exactly those of the form $\delta\tilde{p}$ for some $\tilde{p} \in K\{Y\}$. In one direction, this holds true: polynomials $\delta\tilde{p}$ are nonconstrainable. To see this, note that for every constant $c \in K$ and every $q \in K\{Y\}$ of order $< \text{ord}(\delta\tilde{p})$ the polynomial $(\tilde{p}(Y) - c)$ would have order $> \text{ord}(q)$, so there would exist an $x_c \in \hat{K}$ satisfying

$$\tilde{p}(x_c) - c = 0 \neq q(x_c),$$

by Blum's axioms. Every such x_c would satisfy $(\delta\tilde{p}, q)$, yet for constants $c \neq d$, we would have

$$\tilde{p}(x_c) = c \neq d = \tilde{p}(x_d),$$

proving that $(\delta\tilde{p}, q)$ is not a constrained pair.

However, Theorem 7.3 shows that nonconstrainability is not always equivalent being the derivative of another polynomial.

Corollary 7.6. *There exist a computable differential field K and a differential polynomial $p \in K\{Y\}$ such that, for all $\tilde{p} \in K\{Y\}$, $p \neq \delta\tilde{p}$, yet p_0 is not constrainable in $K\{Y\}$.*

Proof. In the differential field K built in Theorem 7.3, nonconstrainability cannot be equivalent to being a derivative, since nonconstrainability cannot be defined by the Σ_1^0 condition

$$(\exists \tilde{p} \in K\{Y\}) \delta\tilde{p} = p.$$

We argued above that every $\delta\tilde{p}$ with $\tilde{p} \in K\{Y\}$ is nonconstrainable, so the opposite containment relation must fail. That is, there must exist some nonconstrainable p which is not a derivative. \square

In fact, we can say specifically that in the field K constructed in Theorem 7.3, those p_e which turned out to be nonconstrainable are specific instances of Corollary 7.6. If such a p_e were of the form $\delta\tilde{p}$, then \tilde{p} would have appeared in $K\{Y\}$ at some finite stage s , and thus p_e would have been nonconstrainable even in the differential subfield K_s generated by the elements enumerated by stage s . But this did not happen: each p_e was constrainable within each K_s , even though certain of them became nonconstrainable in the larger field K .

Finally, we note that although U_K is Π_2^0 -complete for this particular K , there may still be many other computable differential fields L for which U_L is not this complex. Theorem 7.3 shows that no definition of the constrainability set can be any simpler than Π_2^0 , but in specific cases a definition equivalent to constrainability might have lower complexity. In particular, for a constant differential field F , and especially for the constant field \mathbb{Q} itself, one suspects that the complexity may be lower. These questions remain open.

8. Extension by Constrained Elements

Since the constraint set T_K plays the same role for a differential field K that the splitting set S_F plays for an algebraic field F , it is natural to ask which results about S_F carry over to T_K . The most basic results about S_F are those of Theorem 3.3, given by Kronecker in [11]: that $S_{\mathbb{Q}}$ is computable, and that the splitting set $S_{F(x)}$ for a computable extension $F(x)$ of a computable field F is computable from an oracle for S_F .

Moreover, the process of computing $S_{F(x)}$ from S_F is uniform in F and in one single piece of information about x , namely, whether x is algebraic or transcendental over F . (It is often said that the process is uniform in the minimal polynomial of x over F , but if we know that x is algebraic, then we can find its minimal polynomial using the S_F oracle.)

It follows that every finitely generated computable field of characteristic 0 has computable splitting set, and this fact is extremely useful when one tries to build embeddings or isomorphisms among computable fields. The same holds in characteristic $p > 0$, since the splitting set of a finite field such as \mathbb{F}_p is obviously computable.

In this section, we investigate the analogue of this result for constrained differential field extensions. If K is a computable differential field, with computable extension $K\langle z \rangle$, must $T_{K\langle z \rangle} \leq_T T_K$? We will show that when z is constrained over K , the answer is always positive, uniformly in z , provided only that K is nonconstant. For constant K , the question remains open.

In the previous section we exploited the fact that if $K \subseteq L \subseteq \hat{K}$, then \hat{K} can fail to be (isomorphic over L to) a differential closure of L , although it must contain some subfield $\hat{L} \supseteq L$ which is a differential closure of L . The problem was that certain elements of \hat{K} turned out not to be constrainable over L . When we consider a finitely generated differential field extension $L = K\langle z_0, \dots, z_n \rangle \subseteq \hat{K}$, this is no longer possible.

Lemma 8.1 ([10, §4, Proposition 5]). *Let K be a differential field and L a subfield of any differential closure \hat{K} of K , with L finitely differentially generated as a differential field extension of K . Then \hat{K} must also be a differential closure of its differential subfield L .*

Lemma 8.1 shows that every $x \in \hat{K}$ satisfies a constrained pair over $K\langle z \rangle$, but does not specify the two polynomials which constitute that constrained pair. Below we will prove Theorem 8.6, stating that if we can recognize constrained pairs over K , then we can also recognize them over $K\langle z \rangle$. Once this is known, then given an oracle for T_K , we will be able to identify the constrained pair over $K\langle z \rangle$ satisfied by x , simply by searching for it. However, there is a good deal of work still to be done before we can prove Theorem 8.6.

Lemma 8.2. *For a computable differential field K with Rabin embedding $g : K \rightarrow \hat{K}$, and an arbitrary $z \in \hat{K}$, the splitting set $S_{K\langle z \rangle}$ is computable in an oracle for T_K , uniformly in z .*

Proof. With the oracle T_K , we can find a constrained pair $(p, q) \in \overline{T_K}$ satisfied by z . Then, with $r = \text{ord}(p)$, the differential field $K\langle z \rangle$ is the field

$$K(z, \delta z, \dots, \delta^{r-1}z)[\delta^r z],$$

and p gives the minimal polynomial of $\delta^r z$ over the purely transcendental field extension $K(z, \delta z, \dots, \delta^{r-1}z)$. With this information, Kronecker's Theorem allows us to compute the splitting set of $K\langle z \rangle$ from S_K . \square

The same holds when z is differentially transcendental over K : then $K\langle z \rangle$ is just the field $K(z, \delta z, \delta^2 z, \dots)$, with computable transcendence basis $\{\delta^i z : i \in \omega\}$, hence has a splitting algorithm. However, the uniformity of the result fails when we do not know whether z is constrained or not.

In what follows, for $p \in K\{Y\}$, $[p]$ is the least differential ideal in $K\{Y\}$ containing p . Let $r = \text{ord}(p)$. Recall that the initial I_p of p is defined as the leading coefficient when p is expressed as a polynomial in $\delta^r Y$ (say of degree d) with coefficients in $K(Y, \dots, \delta^{r-1}Y)$, so that $p = I_p(\delta^r Y)^d + \dots$. Let

$$h_p = I_p \cdot s_p, \quad \text{where } s_p = \frac{\partial p}{\partial(\delta^r Y)}.$$

For example, if $p = Y(\delta Y)^2 + \delta Y + 1$, then $h_p = Y \cdot (2Y\delta Y + 1)$. For $J \subset K\{Y\}$, we define the *colon ideal* $J : h_p^\infty$ by:

$$J : h_p^\infty = \{f \in K\{Y\} \mid h_p^n \cdot f \in J \text{ for some } n \geq 0\}.$$

It turns out that $[p] : h_p^\infty$ is a prime differential ideal if and only if $(p) : h_p^\infty$ is a prime ideal [9, III.8, Lemma 6 and IV.9, Lemma 2]. Moreover, if p is an irreducible polynomial, then (p) is a prime ideal. Therefore, in this case, $(p) : h_p^\infty$ is prime as well. Thus, if p is irreducible, then $[p] : h_p^\infty$ is a prime differential ideal.

Lemma 8.3. *For every computable differential field K and every differential polynomial $p \in K\{Y\}$, the relation of equivalence mod $[p] : h_p^\infty$ is computable in $K\{Y\}$,*

uniformly in K and p , and so there is a computable presentation of the differential ring $K\{Y\}/[p] : h_p^\infty$.

Proof. The basic point of the proof is that we have an algorithm, called differential pseudo-division, which takes an $f \in K\{Y\}$ as input and produces $g \in K\{Y\}$ of lower rank than p , with

$$f - g \in [p] : h_p^\infty,$$

(Recall that “lower rank” means that either $\text{ord}(g) < r = \text{ord}(p)$ or else $\text{ord}(g) = r$ and $\delta^r Y$ has strictly lesser degree in g than its degree d in p .) First, if $\text{ord}(f) = s > r$, then $\delta^{s-r} p$ has order s and is linear in $\delta^s Y$ with “leading coefficient” s_p . So, by subtracting an appropriate multiple of $\delta^s Y$ from $s_p \cdot f$, we get an f_0 with order $< s$ such that

$$(s_p \cdot f - f_0) \in [p].$$

Repeating this process produces an f_n of order $\leq r$. If its order is exactly r , then we reduce it by multiples of p until the degree of the variable $\delta^r Y$ in f_m is $< d$ multiplying by h_p whenever needed. A general algorithm is described in [9, I.9].

Of course, if S is the set of all $f \in K\{Y\}$ for which either $\text{ord}(f) < r$ or else $\text{ord}(f) = r$ and $\delta^r Y$ has degree $< d$ in f , then no two distinct elements of S can be equivalent mod $[p] : h_p^\infty$, since

$$[p] : h_p^\infty \cap S = \{0\}.$$

So two elements of $K\{Y\}$ are equivalent mod $[p] : h_p^\infty$ iff the above algorithm on those elements gives the same output for both. Moreover, S is computable and is the domain of our computable presentation of $K\{Y\}/[p] : h_p^\infty$, with addition, multiplication, and differentiation computed exactly as in $K\{Y\}$, with the result reduced modulo $[p] : h_p^\infty$ to an element of S in the sense of the algorithm described above. \square

Corollary 8.4. *For every computable differential field K , every Rabin embedding $g : K \rightarrow \hat{K}$, and every $z \in \hat{K}$, the differential subfield $g(K)\langle z \rangle$ is computably presentable, uniformly in the minimal differential polynomial p of z over K . (In turn, p can be computed from z using a D_K oracle.)*

Proof. $K\langle z \rangle$ is isomorphic to the differential field $\text{Quot}(K\{Y\}/[p] : h_p^\infty)$, with z mapping to the equivalence class of Y . In this case, the irreducible polynomial $p \in K\{Y\}$ forms a characteristic set of the prime differential ideal $P := \ker(K\{Y\} \rightarrow K\langle z \rangle)$ [9, I.10]. \square

Of course, this corollary is clear for other reasons, even without needing to know p . $g(K)\langle z \rangle$ is a c.e. differential subfield of \hat{K} , hence can easily be pulled back to a computable presentation. The interesting point is that one cannot readily use the subfield $g(K)\langle z \rangle$ to prove the uniformity over p in Lemma 8.3. It is easy to find

zeroes $z \in \hat{K}$ of the given p , but unless D_K is computable, one cannot decide which of these zeroes, if any, is generic for p – that is, which z have minimal differential polynomial p .

Moreover, if p were unconstrainable, then no z at all would be generic for p over K . So the proof of Lemma 8.3 is essential to the study of computable differential fields, particularly those for which D_K or U_K is noncomputable. (For computable fields F , the analogous procedure works for all irreducible $p \in F[X]$, so one only needs to be able to decide the splitting set S_F , which is always c.e., whereas in general U_K is only Σ_2^0 .)

We will need the following lemma as well.

Lemma 8.5. *Let \hat{K} be a computable differentially closed field. Then it is computable, given arbitrary m, n , and coefficients for differential polynomials f_0, \dots, f_m and g_0, \dots, g_n in $\hat{K}\{Y\}$, whether $V(f_0, \dots, f_m) \subseteq V(g_0, \dots, g_n)$, where these are the differential varieties in \hat{K} defined by these polynomials. Consequently, it is also uniformly computable whether $V(\bar{f}) = V(\bar{g})$, and whether containment or equality holds between the radical differential ideals $\sqrt{[\bar{f}]}$ and $\sqrt{[\bar{g}]}$.*

Proof. This follows from quantifier elimination for the theory \mathbf{DCF}_0 , but it can also be seen as a pleasing application of the Differential Nullstellensatz, which states that with \hat{K} differentially closed, $V(\bar{f}) \subseteq V(\bar{g})$ iff $\sqrt{[\bar{g}]} \subseteq \sqrt{[\bar{f}]}$. (The analogous result over algebraically closed fields is Hilbert’s Nullstellensatz.) So the statement about radical differential ideals will follow from that about varieties.

Moreover, the reversal between the two inclusions quickly yields the lemma, since $V(\bar{f}) \subseteq V(\bar{g})$ is a universal condition, saying that every common zero of all f_i is also a zero of every g_j , while $\sqrt{[\bar{g}]} \subseteq \sqrt{[\bar{f}]}$ is an existential condition: each of the (finitely many) g_i in \bar{g} lies in $\sqrt{[\bar{f}]}$ iff there exist k_j and a linear combination (over $\hat{K}\{Y\}$) of the f_i ’s and their derivatives such that the combination is equal to $g_j^{k_j}$. Being defined uniformly by both an existential and a universal condition, the containment is therefore computable. \square

We now come to the differential analogue of Kronecker’s Theorem, for finitely generated constrained extensions.

Theorem 8.6. *Let L be any computable ordinary differential field with nontrivial derivation δ , and K its image under any differential Rabin embedding of L into any \hat{K} . Then for every $z \in \hat{K}$, the constraint set $\overline{T_{K\langle z \rangle}}$ is computable in an oracle for T_K . Also, the constrainability set $\overline{U_{K\langle z \rangle}}$ is computably enumerable relative to T_K . Both the computation of $T_{K\langle z \rangle}$ and the enumeration of $\overline{U_{K\langle z \rangle}}$ are uniform in z and T_K .*

Proof. First, using the oracle T_K , search for a pair $(p_z, q_z) \in \overline{T_K}$ satisfied by z . Since $z \in \hat{K}$, we eventually find such a pair, in which p_z is the minimal differential

polynomial of z over K . Set $r_z = \text{ord}(p_z)$. Then $K\langle z \rangle$ has transcendence basis

$$B_0 = \{z, \delta z, \dots, \delta^{r_z-1} z\}$$

over K , and is generated by this basis along with $\delta^{r_z} z$, whose minimal polynomial over the transcendence basis is $p_z(z, \dots, \delta^{r_z-1} z, Y)$. With this information, we can compute the splitting set $S_{K\langle z \rangle}$ from S_K , which in turn is computable from our T_K -oracle, by Theorem 6.1.

Moreover, T_K can also compute D_K , the set of finite subsets of \hat{K} which are algebraically dependent over K . But a finite subset S of $\widehat{K\langle z \rangle} = \hat{K}$ is algebraically independent over $K\langle z \rangle$ iff $S \cap B_0 = \emptyset$ and $S \cup B_0 \notin D_K$. Thus

$$D_{K\langle z \rangle} \leq_T D_K \leq_T T_K,$$

again uniformly in z . (“Computing $D_{K\langle z \rangle}$ ” normally means deciding algebraic independence over $K\langle z \rangle$ for all finite subsets from some particular differential closure of $K\langle z \rangle$. In this case, by Lemma 8.1, the identity map from $K\langle z \rangle$ into \hat{K} is a differential Rabin embedding, and we have given a process for deciding algebraic independence over $K\langle z \rangle$ which works for all finite tuples from \hat{K} .)

Since $K\langle z \rangle$ is a computable differential field, $T_{K\langle z \rangle}$ is Σ_1^0 , uniformly in z (and even in K). So we need only show that its complement is Σ_1^0 relative to T_K . Briefly summarizing, the following argument says that there are two ways for (p, q) to lie in $T_{K\langle z \rangle}$: either p is nonconstrainable there, or else q is not the correct constraint for p . Constrainability is Σ_1 in an oracle for D_K , which we have. Assuming p is constrainable, we use the differential closure to find some x satisfying (p, q) , and then use Lemma 8.3 to compare q over $[p] : h_p^\infty$ with a constrained pair (p, \tilde{q}) over $K\langle z \rangle$, which x is known to satisfy. (Possibly (p, \tilde{q}) is not actually a constrained pair over $K\langle z \rangle$, but it is known to generate a principal type there, which is sufficient for our purpose.)

We claim that, for every $(p, q) \in (K\langle z \rangle\{Y\})^2$, the following process will halt iff $(p, q) \notin T_{K\langle z \rangle}$. Set $r = \text{ord}(p)$, and check first that $r > \text{ord}(q)$, and that p is monic and irreducible as an algebraic polynomial over $K\langle z \rangle$. (This uses $S_{K\langle z \rangle}$.) If so, then we search for an $x \in \hat{K}$ with the properties that

$$p(x) = 0 \neq q(x) \quad \text{and} \quad \{x, \delta x, \dots, \delta^{r-1} x\} \notin D_{K\langle z \rangle}.$$

Assuming $(p, q) \notin T_{K\langle z \rangle}$, such an x exists, so eventually we find one, and Fact 5.1 shows that p is the minimal differential polynomial of this x over $K\langle z \rangle$, hence is constrainable, by Fact 7.2. (If p were not constrainable, then no such x could exist, since we have already confirmed that p is irreducible.) Moreover, the differential

ideal $[p] : h_p^\infty$ in $K\langle z \rangle\{X\}$ is precisely the ideal

$$\{j \in K\langle z \rangle\{X\} : j(x) = 0\}.$$

Having found this x , we search for some u in the c.e. subfield $K\langle x, z \rangle$ of \hat{K} such that x and z both lie in $K\langle u \rangle$.

Theorem 4.7 shows that we do eventually find such a u in \hat{K} , since we assumed δ to be nontrivial, and indeed we find differential rational functions f, g, h with coefficients in K such that

$$f(x, z) = u, \quad g(u) = x, \quad \text{and} \quad h(u) = z.$$

We also use the oracle to find a pair $(p_u, q_u) \in \overline{T_K}$ satisfied by u , and another pair $(p_x, q_x) \in \overline{T_K}$ satisfied by x . Now

$$h(f(x, z)) - z = 0,$$

and if we choose $h_0, h_1 \in K\langle z \rangle\{X\}$ so that

$$\frac{h_0(X)}{h_1(X)} = h(f(X, z)) - z,$$

then $h_0(x) = 0 \neq h_1(x)$. Similarly, setting

$$\frac{g_0(X)}{g_1(X)} = g(f(X, z)) - X \quad \text{and} \quad \frac{p_0(X)}{p_1(X)} = p_u(f(X, z))$$

gives $g_0(x) = 0 \neq g_1(x)$ and $p_0(x) = 0 \neq p_1(x)$, with all these polynomials in $K\langle z \rangle\{X\}$.

Now we define $\tilde{q} \in K\langle z \rangle\{X\}$ by

$$\tilde{q}(X) = q_u(f(X, z)) \cdot g_1(X) \cdot h_1(X) \cdot p_1(X) \cdot h_p(X),$$

with h_p as on page 25 for our p . Suppose $\tilde{x} \in \hat{K}$ satisfies (p, \tilde{q}) . Now every $j \in K\langle z \rangle\{X\}$ with $j(x) = 0$ lies in $[p] : h_p^\infty$, hence must have $j(\tilde{x}) = 0$ (since $h_p(\tilde{x}) \neq 0$). We also know that x satisfies (p, \tilde{q}) , since h_p has strictly lesser rank than the minimal differential polynomial p of x , and thus must have $h_p(x) \neq 0$. We set

$$\tilde{u} = f(\tilde{x}, z),$$

so by our assumption about \tilde{x} , we know $q_u(\tilde{u}) \neq 0$. Also,

$$0 = p_u(u) = p_u(f(x, z)),$$

so p_0 lies in $[p] : h_p^\infty$, while $p_1(\tilde{x}) \neq 0$, and therefore

$$0 = p_u(f(\tilde{x}, z)) = p_u(\tilde{u})$$

as well. Therefore this \tilde{u} satisfies the constrained pair $(p_u, q_u) \in \overline{T_K}$, and so the map σ with $\sigma(u) = \tilde{u}$ which restricts to the identity on K is an isomorphism from $K\langle u \rangle$ onto $K\langle \tilde{u} \rangle$.

Now $0 = h_0(\tilde{x}) = g_0(\tilde{x})$, since these differential polynomials both have x as a zero, while we know

$$g_1(\tilde{x}) \cdot h_1(\tilde{x}) \neq 0.$$

It follows that

$$h(f(\tilde{x}, z)) - z = 0 = g(f(\tilde{x}, z)) - \tilde{x},$$

and so $h(\tilde{u}) = z$ and $g(\tilde{u}) = \tilde{x}$. But g and h have coefficients in K , so

$$\sigma(z) = \sigma(h(u)) = h(\sigma(u)) = h(\tilde{u}) = z \quad \text{and} \quad \sigma(x) = \sigma(g(u)) = g(\sigma(u)) = g(\tilde{u}) = \tilde{x}.$$

Thus this σ maps x to \tilde{x} , fixes z , and fixes K pointwise, so σ witnesses that $x \cong_{K\langle z \rangle} \tilde{x}$. We have thus shown that $(p(X), \tilde{q}(X))$ generates a principal type over $K\langle z \rangle$ and is satisfied by x .

Notice that (p, \tilde{q}) could fail to be a constrained pair, since \tilde{q} could have order $\geq \text{ord}(p)$, but it has the basic property of constrained pairs, which is to generate a principal type. (Recall that this means that every two elements of $\widehat{K\langle z \rangle}$ which both satisfy (p, q) are isomorphic over $K\langle z \rangle$.)

Now suppose $t \in V(p, q)$. Then $q(t) = 0 \neq q(x)$, so

$$x \not\cong_{K\langle z \rangle} t.$$

Since x satisfies (p, \tilde{q}) , t cannot satisfy it, and with $p(t) = 0$, this forces $\tilde{q}(t) = 0$. Thus

$$V(p, q) \subseteq V(p, \tilde{q}).$$

We now claim that

$$V(p, \tilde{q}) = V(p, q) \iff (p, q) \in \overline{T_{K\langle z \rangle}}.$$

For the forwards direction, suppose t and v both satisfy (p, q) . Then

$$t \notin V(p, q) = V(p, \tilde{q}),$$

but $p(t) = 0$, so $\tilde{q}(t) \neq 0$. Thus t satisfies (p, \tilde{q}) , and so does v , by the same

argument. Since (p, \tilde{q}) generates a principal type, we have $t \cong_{K\langle z \rangle} v$, which proves

$$(p, q) \in \overline{T_{K\langle z \rangle}}.$$

For the converse, suppose

$$V(p, \tilde{q}) \neq V(p, q).$$

Then there must exist some $w \in V(p, \tilde{q}) - V(p, q)$, since we saw above that $V(p, q) \subseteq V(p, \tilde{q})$. Thus $p(w) = 0$ and $q(w) \neq 0$, so w satisfies (p, q) , as does x . Yet

$$\tilde{q}(x) \neq 0 = \tilde{q}(w)$$

(since $w \in V(p, \tilde{q})$), and so $x \not\cong_{K\langle z \rangle} w$, which proves

$$(p, q) \notin \overline{T_{K\langle z \rangle}}.$$

So we use Lemma 8.5 to check whether $V(p, \tilde{q}) = V(p, q)$, and if so, we conclude that (p, q) is a constrained pair over $K\langle z \rangle$. Thus, being such a constrained pair is Σ_1^0 relative to T_K , and not being such a constrained pair is Σ_1^0 without any oracle. Finally, since a $T_{K\langle z \rangle}$ -oracle can enumerate $\overline{U_{K\langle z \rangle}}$, so can a T_K -oracle, now that we know $T_{K\langle z \rangle} \leq_T T_K$. This completes the proof of Theorem 8.6. \square

9. Extension by Differential Transcendentals

Lemma 9.1. *Suppose that $p \in K\{Y\}$ is algebraically irreducible and has order r . If an element $y \in \hat{K}$ has $p(y) = 0$, but p is not the minimal differential polynomial of y , then there exists some $h \in K\{Y\}$ of order $< r$ such that $h(y) = 0$ as well.*

Proof. This y must have a minimal differential polynomial h over K . If h has order $\geq r$, then by minimality it must have order r and degree $< d$, the degree of $\delta^r Y$ in $p(Y)$. But $h(y) = 0$ gives the minimal (algebraic) polynomial in

$$K(y, \delta y, \dots, \delta^{r-1} y)[\delta^r Y]$$

of $\delta^r y$ over $K(y, \delta y, \dots, \delta^{r-1} y)$. Thus h must divide p , contradicting the irreducibility of p . \square

Lemma 9.2. *Suppose that $p \in K\langle z \rangle\{Y\}$, where z is differentially transcendental over K . Suppose further that all coefficients of p lie within the field $K(z, \delta z, \dots, \delta^r z)$. If there are elements x, y in the differential closure of $K\langle z \rangle$ for which $p(x) = p(y) = 0$ but $x \not\cong_{K\langle z \rangle} y$, then there exists $g \in K(z, \dots, \delta^r z)\{Y\}$ of strictly lesser rank (in z) than p , such that*

$$\text{either } g(x) = 0 \neq g(y) \text{ or } g(y) = 0 \neq g(x).$$

If p is constrainable over $K\langle z \rangle$, and $g(x) = 0 \neq g(y)$, then there is also some \tilde{y} and some $\tilde{g} \in K(z, \dots, \delta^{r-1}z)\{Y\}$ such that $\tilde{g}(x) = p(\tilde{y}) = 0 \neq \tilde{g}(\tilde{y})$.

Proof. The proof of the main statement is by induction on the rank of p with respect to the variable z , where we view p as a polynomial in both Y and z , clearing denominators if needed. (Technically, this is a transfinite induction, since ranks of differential polynomials, even in our simple definition, form an order of type ω^2 .) The base case is trivial, since when $r = -1$, then the two elements x and y must satisfy $x \not\cong_K y$, and Proposition 4.5 then provides a polynomial $g \in K\{Y\}$ with $g(x) = 0 \neq g(y)$ (or vice versa).

For the inductive step, let p , x , and y be as described. Since $x \not\cong_{K\langle z \rangle} y$, Proposition 4.5 yields some $h \in K\langle z \rangle\{Y\}$ with

$$h(x) = 0 \neq h(y).$$

That proposition does not provide any *a priori* bound on the order of z in the coefficients of this h , but we may view $h(Y)$ as a differential polynomial in the two variables Y and z (clearing denominators in $h(Y)$ if needed) and reduce h modulo $[p] : h_p^\infty$ with respect to z , finding some g , whose rank with respect to z is less than the rank of $p(Y)$, and some exponent k such that $h_p^k \cdot h - g \in [p]$ (see [9, Section I.9]). Therefore,

$$h_p(x)^k \cdot h(x) - g(x) = 0,$$

forcing $g(x) = 0$. Now either g satisfies the lemma (if $g(y) \neq 0$), or we apply induction to g (if $g(y) = 0$), since g also has strictly lesser rank than p .

Now suppose that p is also constrainable. (Notice that, if $(p, q) \in \overline{T_{K\langle z \rangle}}$, then q has lower order than p in the variable Y , but not necessarily in z . So we cannot simply take \tilde{g} to equal q .) Assume that g has positive degree in $\delta^r z$, since otherwise we can set $\tilde{g} = g$. This degree must be less than that of p , since g has lower rank than p , and so we may apply the pseudo-division algorithm from Lemma 8.3 repeatedly, with respect to the degree of the variable $\delta^r z$ over coefficients from $K(z, \dots, \delta^{r-1}z)\{Y\}$. This yields

$$\begin{aligned} i_g(Y)^{m_g} \cdot p(Y) &= g(Y) \cdot d_1(Y) + q_1(Y) \\ i_{q_1}(Y)^{m_1} \cdot g(Y) &= q_1(Y) \cdot d_2(Y) + q_2(Y) \\ &\vdots \\ i_{q_n}(Y)^{m_n} \cdot q_{n-1}(Y) &= q_n(Y) \cdot d_{n+1}(Y) + q_{n+1}(Y), \end{aligned}$$

where, for all $i \leq n$, $q_{i+1}(Y)$ has strictly lower degree in $\delta^r z$ than $q_i(Y)$ has (including $i = 0$, with $q_0 = g$). The process ends when we reach an n for which $q_{n+1}(Y)$ has degree 0 in $\delta^r z$. Notice that if $q_{n+1}(Y)$ were the zero polynomial, then $q_n(Y)$ would

have positive degree in $\delta^r z$ and would divide

$$i_g(Y)^{m_g} \cdot i_{q_0}(Y)^{m_0} \cdot \dots \cdot i_{q_n}(Y)^{m_n} \cdot p(Y),$$

which is impossible, since

$$i_g(Y)^{m_g} \cdot i_{q_0}(Y)^{m_0} \cdot \dots \cdot i_{q_n}(Y)^{m_n}$$

is free of $\delta^r z$ and the constrainable polynomial p , even after its denominators were cleared, must be irreducible (as a polynomial in $Y, \delta Y, \delta^2 Y \dots$) and can have no nontrivial factor from $K\langle z \rangle$. Therefore,

$$q_{n+1}(Y) \in K(z, \dots, \delta^{r-1} z)\{Y\}$$

is nonzero. But by induction on i , we see that either $i_{q_i}(x) = 0$, or else $q_i(x) = 0$ for every x , since $p(x) = g(x) = 0$. In the first case, we take $\tilde{g} := i_g$, while in the second case, $\tilde{g} = q_{n+1}$ is the desired differential polynomial. In both cases, we get

$$\tilde{g}(x) = 0 = p(x),$$

while Fact 7.2 yields some \tilde{y} in the differential closure of $K\langle z \rangle$ with $p(\tilde{y}) = 0 \neq \tilde{g}(\tilde{y})$. \square

We state the next lemma in a specific form which will be useful for the results in this section.

Lemma 9.3. *For any differential polynomials $p, q, h \in K\{Y\}$, the set*

$$\left\{ x \in \hat{K} : p(x) = h(x) = 0 \neq q(x) \right\}$$

is empty iff $1 \in \sqrt{[p, h]} : q$.

Proof. If $p(x) = h(x) = 0 \neq q(x)$, then every $f \in \sqrt{[p, h]} : q$ has $f^n \cdot q^n \in [p, h]$ for some n , hence has $f(x) = 0$, precluding the constant function 1 from appearing in $\sqrt{[p, h]} : q$.

Conversely, if $1 \notin \sqrt{[p, h]} : q$, then, for every n , we have $q^n \notin [p, h]$, and the Differential Nullstellensatz yields an $x \in \hat{K}$ with $p(x) = h(x) = 0$ but $q(x) \neq 0$. \square

Corollary 9.4. *Suppose the pair (p, q) from $K\{Y\}$ has p monic and irreducible of order $> \text{ord}(q)$. Then the following are equivalent.*

1. $(p, q) \in \overline{T_K}$.
2. For all $h \in K\{Y\}$ of lesser rank than p , $1 \in \sqrt{[p, h]} : q$.
3. For all $g \in K\{Y\}$, we have either $1 \in \sqrt{[p, g]} : q$ or $g \in \sqrt{[p]} : q$.

Proof. If (p, q) is a constrained pair, then p is the minimal differential polynomial of some $y \in \hat{K}$, and so $q(y) \neq 0 \neq h(y)$ for every h of lesser rank than p . But now every x with $p(x) = 0 \neq q(x)$ has $x \cong_K y$, so $h(x) \neq 0$ for all such h , forcing

$$\left\{ x \in \hat{K} : p(x) = h(x) = 0 \neq q(x) \right\}$$

to be empty. By Lemma 9.3, it follows that $1 \in \sqrt{[p, h]} : q$ for all such h .

Conversely, suppose $1 \in \sqrt{[p, h]} : q$ for all $h \in K\{Y\}$ of lesser rank than p . Then, for every such h ,

$$\left\{ x \in \hat{K} : p(x) = h(x) = 0 \neq q(x) \right\}$$

is empty, by Lemma 9.3. Therefore, if $x, y \in \hat{K}$ both satisfy (p, q) , then $h(x) \neq 0 \neq h(y)$ for every such h , and Lemma 9.2 then shows that $x \cong_K y$. Thus, $(p, q) \notin T_K$, and so (1) \iff (2).

The equivalence of (2) and (3) follows because every g is equivalent modulo $\sqrt{[p]} : q$ to some h of lesser rank than p . If this h is zero, then $g \in \sqrt{[p]} : q$, while if not, then every x with $p(x) = 0 \neq q(x)$ has $g(x) = h(x) \neq 0$, and so $1 \in \sqrt{[p, g]} : q$ by Lemma 9.3. \square

Lemma 9.5. *If \hat{K} is not algebraic over K , then this extension is of infinite transcendence degree, and indeed \hat{K} contains elements of arbitrarily large order over K .*

Proof. The argument showing the first part of the statement is due to Michael Singer. For simplicity, assume that K is algebraically closed. When K and \hat{K} have the same field of constants, one can use [23, Corollary, p. 489] to show that the transcendence degree of \hat{K} over K is infinite. To see this, assume that this transcendence degree is finite, say n . Since K is algebraically closed, $n \geq 1$. Let y be in \hat{K} and transcendental over K . Let

$$v_i = y^i, \quad i = 1, \dots, n+1.$$

Since \hat{K} is differentially closed, let $u_i \in \hat{K}$ satisfy

$$\delta(u_i)/u_i = -\delta(v_i) \quad i = 1, \dots, n+1.$$

The conclusion of the corollary implies that there exist constants c_i such that $\sum_i c_i y^i$ is algebraic over K , which is a contradiction.

When \hat{K} has new constants, argue as follows. Let C be the constants of \hat{K} . If \hat{K} is not algebraic over KC , then there is $y \in \hat{K}$ transcendental over KC , and we can argue as above to get a contradiction. Therefore, we can assume \hat{K} is algebraic over KC . We can assume that K contains a non-constant x (otherwise there are lots of elements in \hat{K} that are algebraically independent over K , for example, non-zero

solutions of

$$\delta(Y) = z^n \cdot Y, \quad n \geq 0,$$

where $z \in \hat{K}$ with $\delta(z) = 1$ and that $\delta(x) = 1$ (replacing δ by $(1/\delta(x)) \cdot \delta$ if needed). Since K is algebraically closed, there is a constant $c \in \hat{K}$ not algebraic over K . Since \hat{K} is differentially closed, there is $y \in \hat{K}$ such that

$$\delta(y) = 1/(c + x).$$

Since y is algebraic over KC , taking traces, we see [18, Exercise 1.24] that there is $z \in KC$ such that

$$\delta(z) = 1/(c + x).$$

Since KC is an extension of $K(c)$ by constants, [20, Proposition 1.2] yields $w = p(c)/q(c) \in K(c)$ such that

$$\delta(w) = 1/(c + x).$$

Expanding w in partial fractions with respect to c , differentiating and comparing terms shows that this is impossible.

Now we consider the orders of elements of \hat{K} . If K is nonconstant, Theorem 4.7 then implies that \hat{K} contains individual elements of arbitrarily large order over K . In the case where K is a constant differential field, let $x \in \hat{K}$ be such that $x' = 1$; such an element exists in \hat{K} by Blum's axioms, and must be transcendental over K since an element algebraic over the constant field K would also be constant. Then, for all $a_1, \dots, a_n \in K$,

$$f(a_1, \dots, a_n) := x^n/n! + a_1 x^{n-1} + \dots + a_n$$

satisfies $f^{(n)} = 1$. Now there exist $a_1, \dots, a_n \in K$ such that $f(a_1, \dots, a_n)$ is not a solution of any polynomial differential equation F of order $n - 1$ or less. Indeed, for each $q < n$, the coefficient of x^0 in $f^{(q)}$ is equal to

$$q! \cdot a_{n-q}.$$

(This coefficient is often referred to as the “constant term” of $f^{(q)}$, but it would be confusing to call it that in this context.) Hence, for all q_1, \dots, q_r from $\{0, 1, \dots, n-1\}$ and all n_1, \dots, n_r , the coefficient of x^0 of

$$\prod_i (f^{(q_i)})^{n_i}$$

is equal to

$$\prod_i (q_i! \cdot a_{n-q_i})^{n_i}$$

Thus, the coefficient of x^0 in $F(f)$ is a polynomial in a_1, \dots, a_n , which is not identi-

cally zero. This implies the result as K is infinite. \square

In the proof of Theorem 9.6, we use the technique of characteristic sets from differential algebra. We will give a short introduction to it for the convenience of the reader unfamiliar with characteristic sets and differential rankings. Earlier, we introduced the ring of differential polynomials in one differential indeterminate, which we denoted by $K\{Y\}$ most often. Let us now have finitely many differential indeterminates $\{y_1, \dots, y_n\}$ (in the proof of Theorem 9.6, we will just need the case of $n = 2$). A *ranking* is a total order $>$ on the set of derivatives

$$D := \left\{ y_i^{(p)} \mid p \geq 0, 1 \leq i \leq n \right\}$$

satisfying the following conditions for all $p > 0$ and $u, v \in D$:

1. $u^{(p)} > u$, and
2. $u \geq v \implies u^{(p)} \geq v^{(p)}$.

Similarly to our previous introduced notation, $K\{y_1, \dots, y_n\} := K[D]$ (as commutative rings, and the δ -action is defined naturally) and, for $u = y_j^{(q)}$, $\text{ord } u := q$. If $f \in K\{y_1, \dots, y_n\} \setminus K$, then $\text{ord } f$ denotes the maximal order of the derivatives appearing effectively in f .

A ranking $>$ is called *orderly* if, for all $u, v \in D$, $\text{ord } u > \text{ord } v$ implies $u > v$. A ranking $>$ is called *elimination* with $y_{i_1} < \dots < y_{i_n}$ if

$$y_{i_j}^{(p)} < y_{i_{j+1}}^{(q)}, \quad p, q \geq 0, \quad 1 \leq j \leq n.$$

For example, in an orderly ranking with $y_1 < y_2$, we have: $y_1' > y_2$ and $y_1' < y_2'$. However, in an elimination ranking with $y_1 < y_2$, we have: $y_1' < y_2$ and $y_1' < y_2'$.

Let a ranking $>$ be fixed on S . The derivative $y_j^{(p)}$ of the highest rank appearing in $f \in K\{y_1, \dots, y_n\} \setminus K$ is called the *leader* of f , which we denote by u_f . Represent f as a univariate polynomial in u_f :

$$f = i_f u_f^d + a_1 u_f^{d-1} + \dots + a_d. \tag{1}$$

The monomial u_f^d is called the *rank* of f . Extend the ranking relation on derivatives to ranks: $u_1^{d_1} > u_2^{d_2}$ if either $u_1 > u_2$, or $u_1 = u_2$ and $d_1 > d_2$. As at the beginning of Section 8, the polynomial i_f is called the *initial* of f . Applying δ to f , we obtain

$$\delta f = \frac{\partial f}{\partial u_f} \delta u_f + \delta i_f u_f^d + \delta a_1 u_f^{d-1} + \dots + \delta a_d.$$

The leader of δf is δu_f , and the initial of δf is called the *separant* of f , which is denoted by s_f . Note that the initial of any proper derivative (i.e., a derivative of

order greater than 0) of f is equal to s_f . For example, if $K = \mathbb{Q}(x)$, $\delta(x) = 1$, the ranking is elimination with $y_1 > y_2$ and

$$f = (y_2 y_1' + 1) y_1''^2 + y_1'^2 y_2'^3 y_1'' + y_1 y_2^3 + x,$$

then $u_f = y_1''$, the rank of f is $y_1''^2$, $i_f = y_1' y_2 + 1$,

$$\begin{aligned} \delta f = & (2(y_2 y_1' + 1) y_1'' + y_1'^2 y_2'^3) y_1''' + \\ & + (y_2 y_1'' + y_2' y_1') y_1''^2 + (2 y_1' y_1'' y_2'^3 + 3 y_1'^2 y_2'^2 y_2'') y_1'' + y_1' y_2^3 + 3 y_1 y_2^2 y_2' + 1 \end{aligned}$$

and $s_f = 2(y_2 y_1' + 1) y_1'' + y_1'^2 y_2'^3$.

We say that $f \in K\{y_1, \dots, y_n\}$ is *partially reduced* w.r.t. $g \in K\{y_1, \dots, y_n\} \setminus K$ if no proper derivative of u_g appears in f . Also, f is said to be *algebraically reduced* w.r.t. g if $\deg_{u_g} f < \deg_{u_g} g$. Finally, f is called *reduced* w.r.t. g if f is partially and algebraically reduced w.r.t. g . Let $A \subset K\{y_1, \dots, y_n\} \setminus K$. For example, $y_1' + y_2$ is reduced w.r.t. $y_2' + y_1$ in an orderly ranking but is not partially reduced w.r.t. $y_2' + y_1$ in an elimination ranking with $y_1 > y_2$. Also, y_1 is partially reduced but not reduced w.r.t. y_1 . We say that A is *autoreduced* if each element of A is reduced w.r.t. all the others. So, in an orderly ranking, the set $y_1' + y_2, y_2' + y_1$ is autoreduced. However, it is not autoreduced in any elimination ranking.

Every autoreduced set is finite [9, Chapter I, Section 9]. For such sets we use the notation $A = A_1, \dots, A_p$ to specify the list of the elements of A arranged in order of increasing rank. We denote the sets of initials and separants of elements of A by i_A and s_A , respectively. Let $H_A = i_A \cup s_A$. For a finite set S in a commutative ring R , denote the smallest multiplicative set containing 1 and S by S^∞ . Let I be an ideal of R . The *colon ideal* $I : S^\infty$ is defined as

$$\{a \in R \mid \exists s \in S^\infty : sa \in I\}.$$

If R is a differential ring I is a differential ideal, then $I : S^\infty$ is also a differential ideal (see [9]).

Let $A = A_1, \dots, A_r$ and $B = B_1, \dots, B_s$ be autoreduced sets. We say the rank of A is lower than the rank of B if

- there exists $k \leq \min(r, s)$ such that $\text{rank } A_i = \text{rank } B_i$ for all $i, 1 \leq i < k$, and $\text{rank } A_k < \text{rank } B_k$,
- or if $r > s$ and $\text{rank } A_i = \text{rank } B_i$ for all $i, 1 \leq i \leq s$.

We say that $\text{rank } A = \text{rank } B$ if $r = s$ and $\text{rank } A_i = \text{rank } B_i$ for all $i, 1 \leq i \leq r$. For instance, let $A := y_1' + y_2, y_2' + y_1$ and $B := y_1' + y_2$. Then, in an orderly ranking with $y_1 < y_2$, the rank of A is lower than the rank of B because $\text{rank } A_1 = \text{rank } B_1$, but A has more elements than B does. For an orderly ranking with $y_2 < y_1$, we

write $A = y'_2 + y_1, y'_1 + y_2$, and the rank of A is lower than the rank of B because now $\text{rank } A_1 < \text{rank } B_1$.

For a differential ideal $I \subset K\{y_1, \dots, y_n\}$, its autoreduced subset of the least rank is called a *characteristic set* of I [9, p. 82]. If the differential ideal I is prime and C is its characteristic set, then $I = [C] : H_C^\infty$ [9, Lemma 2, Chapter IV, Section 9]. For example, the smallest radical differential ideal I of $K\{y\}$ containing $y'^2 + y$ is not prime, having two minimal differential prime ideals containing it, $P_1 = [y'^2 + y, 2y'' + 1]$ and $P_2 = [y]$. Moreover, $y'^2 + y$ and y are characteristic sets of P_1 and P_2 , respectively, and $P_1 = [y'^2 + y] : y'^\infty$. Also, $y'^2 + y$ is a characteristic set of I , but $[y'^2 + y] : y'^\infty$ strictly contains I .

Theorem 9.6. *Consider any computable ordinary nonconstant differential field K , and assume that \hat{K} is not an algebraic field extension of K . Let $K\langle z \rangle$ be a computable differential field extension generated by an element z which is differentially transcendental over K , presented so that K is a computably enumerable subset of $K\langle z \rangle$. Then the constraint set $\overline{T_{K\langle z \rangle}}$ is computable in an oracle for T_K . Also, the constrainability set $\overline{U_{K\langle z \rangle}}$ is computably enumerable relative to T_K . Both the computation of $T_{K\langle z \rangle}$ and the enumeration of $\overline{U_{K\langle z \rangle}}$ are uniform in z and T_K .*

Proof. It is only necessary to show that $T_{K\langle z \rangle}$ has both an existential definition and a universal definition, in which the quantifier-free parts are allowed to use the relation of membership in the oracle set T_K . Of course, $T_{K\langle z \rangle}$ is computably enumerable without any oracle; indeed, its existential definition comes straight from Definition 4.3, with no T_K -oracle required:

$$(p, q) \in T_{K\langle z \rangle} \iff (\exists x, y \in \widehat{K\langle z \rangle})(\exists h \in K\langle z \rangle\{Y\}) \\ [h(x) = 0 = p(x) \neq q(x) \ \& \ h(y) \neq 0 = p(y) \neq q(y)].$$

So we need only give a universal definition of $T_{K\langle z \rangle}$ (equivalently, an existential definition of $\overline{T_{K\langle z \rangle}}$) relative to T_K . For any $p, q \in K\langle z \rangle\{Y\}$, clear the denominators of the coefficients to form polynomials $p_0, q_0 \in K\{z, Y\}$. We claim that

$$(p, q) \in \overline{T_{K\langle z \rangle}} \iff (\exists f, g \in K\{Z\})(\exists \tilde{z} \in \hat{K}) [(f, g) \in \overline{T_K} \ \& \ f(\tilde{z}) = 0 \neq g(\tilde{z}) \\ \ \& \ (p_0(\tilde{z}, Y), q_0(\tilde{z}, Y)) \in \overline{T_{K\langle \tilde{z} \rangle}} \ \& \ p \in K(z, \delta z, \dots, \delta^{(\text{ord}(f) - \text{ord}_Y(p) - 1)/2} z)\{Y\}]$$

Once we have proven this equivalence, we will have a decision procedure for $T_{K\langle z \rangle}$ relative to T_K : given input (p, q) , search simultaneously for witnesses to either of the two existential statements above. (Notice that a T_K -oracle will decide membership in $T_{K\langle \tilde{z} \rangle}$, for every $\tilde{z} \in \hat{K}$, uniformly in \tilde{z} , according to Theorem 8.6.) Eventually, this procedure must find a witness for one or the other, and when it does, we have determined whether $(p, q) \in T_{K\langle z \rangle}$ or not.

For the forwards implication, let $(p, q) \in \overline{T_{K\langle z \rangle}}$. Then, for every $h_0 \in K\{z, Y\}$ of

rank (with respect to Y) less than the rank of p_0 , we know that

$$1 \in \sqrt{[p_0, h_0]} : q_0,$$

by Corollary 9.4. Therefore, no matter what f and g we choose from $K\{Z\}$, we will have

$$1 \in \sqrt{[p_0, h_0, f]} : (g \cdot q_0)$$

for every such h_0 , and hence 1 will also lie in the ideal

$$\sqrt{[p_0(\tilde{z}, Y), h_0(\tilde{z}, Y)]} : q_0(\tilde{z}, Y) \subset K\langle\tilde{z}\rangle\{Y\}$$

for each $\tilde{z} \in \hat{K}$ satisfying (f, g) . Set

$$\tilde{p}(Y) = p_0(\tilde{z}, Y) \quad \text{and} \quad \tilde{q}(Y) = q_0(\tilde{z}, Y).$$

Now every $\tilde{h} \in K\langle\tilde{z}\rangle\{Y\}$ of rank (with respect to Y) less than the rank of \tilde{p} has lower rank than p , since \tilde{p} has rank \leq the rank of p , and so viewing \tilde{h} as a polynomial in both \tilde{z} and Y and clearing denominators yields an h_0 as above and shows that

$$1 \in \sqrt{[\tilde{p}, \tilde{h}]} : \tilde{q}.$$

Then we apply Corollary 9.4 once more to see that (\tilde{p}, \tilde{q}) lies in $\overline{T_{K\langle z \rangle}}$, as required. By Lemma 9.5, there must exist a pair $(f, g) \in \overline{T_K}$ with $\text{ord}_z(f)$ large enough that

$$p \in K(z, \delta z, \dots, \delta^{(\text{ord}_z(f) - \text{ord}_Y(p) - 1)/2} z)\{Y\},$$

and this pair, along with any $\tilde{z} \in \hat{K}$ satisfying it, satisfies the existential condition given.

We prove the backwards implication by contraposition. Suppose that $(p, q) \in T_{K\langle z \rangle}$. Fix an r such that

$$p, q \in K(z, \delta z, \dots, \delta^r z)\{Y\}.$$

Then there exist elements x and y in the differential closure of $K\langle z \rangle$ which both satisfy (p, q) , but such that some $h \in K\{z, Y\}$ has

$$h(z, x) = 0 \neq h(z, y). \tag{2}$$

Let $h_0(z, Y) \in K\{z, Y\}$ be the result of multiplying h by the least common multiple of all denominators of its coefficients, as with p_0 and q_0 . By Lemma 9.2, we may assume that h_0 is of rank lower than that of p_0 (under the orderly ranking with

$z < Y$), that is,

$$\text{ord}_z(h_0) + \text{ord}_Y(h_0) \leq \text{ord}_z(p_0) + \text{ord}_Y(p_0) =: b \quad (3)$$

holds. It follows from (2) that

$$1 \notin \sqrt{[p, h]} : q \quad \text{and} \quad h \notin \sqrt{[p]} : q \quad (4)$$

implying that

$$\sqrt{[p_0, h_0]} : q_0 \cap K\{z\} = \{0\}. \quad (5)$$

It follows from (5) that there exists a minimal differential prime component of

$$\sqrt{[p_0, h_0]} : q_0$$

such that one of its characteristic sets C with respect to the elimination ranking with $z < Y$ has the form $\{C_1\}$ for some irreducible $C_1 \in K\{z, Y\} \setminus K\{z\}$. Indeed, suppose it had two elements in it, $B_1 < B_2$. We know that a characteristic set is autoreduced, and that $z < Y$ and the ranking is elimination. So, the leaders of B_1 and B_2 must be derivatives of z and Y , respectively. Again, since the ranking is elimination and $z < Y$, B_1 cannot depend on Y , which would contradict (5). It now follows from [5, Proposition 14] and (3) that

$$\begin{aligned} \text{ord}_z(C_1) &\leq \text{ord}_z(C_1) + \text{ord}_Y(C_1) \leq \\ &\leq \max(\text{ord}_z(h_0), \text{ord}_z(q_0), \text{ord}_z(p_0)) + \max(\text{ord}_Y(h_0), \text{ord}_Y(q_0), \text{ord}_Y(p_0)) \leq 2b. \end{aligned}$$

Moreover, for every irreducible $C_0 \in K\{z\}$ with $\text{ord}_z(C_0) > 2b$, the set $C := \{C_0, C_1\}$ is autoreduced and, by the Rosenfeld Lemma (see [7, Theorem 4.8]),

$$1 \in [C] : H_C^\infty \iff 1 \in (C) : H_C^\infty \quad (6)$$

with the latter statement being impossible because of the choice of C_0 and C_1 (see [21, pages 88–90]). Fix any $(f, g) \in \overline{T_K}$ with f of order $\geq 2b + 1$, and any $\tilde{z} \in \hat{K}$ satisfying (f, g) , and let \tilde{p} , \tilde{q} , and \tilde{h} be the result of replacing z by \tilde{z} in p_0 , q_0 , and h_0 , respectively. Now (4) and (6) imply that

$$1 \notin \sqrt{[\tilde{p}, \tilde{h}]} : \tilde{q} \subset \hat{K}\{Y\}. \quad (7)$$

Therefore, Corollary 9.4 shows that $(\tilde{p}, \tilde{q}) \in T_{K\langle \tilde{z} \rangle}$. This completes the proof of the backwards direction, and shows that our existential definition of $\overline{T_{K\langle z \rangle}}$ was correct. \square

References

- [1] L. Blum, Differentially closed fields: a model theoretic tour, in: H. Bass, P. Cassidy, J. Kovacic (Eds.), Contributions to Algebra (collection of papers dedicated to Ellis Kolchin), Academic Press, New York, 1977, pp. 37–61.
- [2] H. Edwards, Galois Theory, Springer-Verlag, New York, 1984.
- [3] M. Fried, M. Jarden, Field Arithmetic, Springer-Verlag, Berlin, 2008. URL: <http://dx.doi.org/10.1007/978-3-540-77270-5>.
- [4] A. Fröhlich, J. Shepherdson, Effective procedures in field theory, Philosophical Transactions of the Royal Society of London. Series A. Mathematical and Physical Sciences 248 (1956) 407–432. URL: <http://dx.doi.org/10.1098/rsta.1956.0003>.
- [5] O. Golubitsky, M. Kondratieva, M. Moreno Maza, A. Ovchinnikov, A bound for the Rosenfeld-Gröbner algorithm, Journal of Symbolic Computation 43 (2008) 582–610. URL: <http://dx.doi.org/10.1016/j.jsc.2007.12.002>.
- [6] L. Harrington, Recursively presentable prime models, The Journal of Symbolic Logic 39 (1974) 305–309. URL: <http://www.jstor.org/stable/2272643>.
- [7] E. Hubert, Notes on triangular sets and triangulation-decomposition algorithms II: Differential systems, in: Symbolic and Numerical Scientific Computing 2001, pp. 40–87. URL: http://dx.doi.org/10.1007/3-540-45084-X_2.
- [8] E. Kolchin, Extensions of differential fields, I, Annals of Mathematics 43 (1942) 724–729. URL: <http://www.jstor.org/stable/1968962>.
- [9] E. Kolchin, Differential Algebra and Algebraic Groups, Academic Press, New York, 1973.
- [10] E. Kolchin, Constrained extensions of differential fields, Advances in Mathematics 12 (1974) 141–170. URL: [http://dx.doi.org/10.1016/S0001-8708\(74\)80001-0](http://dx.doi.org/10.1016/S0001-8708(74)80001-0).
- [11] L. Kronecker, Grundzüge einer arithmetischen theorie der algebraischen größen, Journal für die reine und angewandte Mathematik 92 (1882) 1–122.
- [12] D. Marker, M. Messmer, A. Pillay, Model Theory of Fields, volume 5 of *ASL Lecture Notes in Logic*, 2nd ed., A.K. Peters, Ltd., Wellesley, MA, 2006.
- [13] C. McCoy, R. Miller, Independence and spanning sets for computable free groups and fields, 2014. In preparation.
- [14] R. Miller, Computable fields and Galois theory, Notices of the American Mathematical Society 55 (2008) 798–807. URL: <http://www.ams.org/notices/200807/tx080700798p.pdf>.
- [15] R. Miller, Is it harder to factor a polynomial or to find a root?, Transactions of the American Mathematical Society 362 (2010) 5261–5281. URL: <http://dx.doi.org/10.1090/S0002-9947-2010-04918-9>.
- [16] R. Miller, A. Ovchinnikov, Adapting Rabin’s theorem for differential fields, in: B. Löwe, D. Normann, I. Soskov, A. Soskova (Eds.), Models of Computation in Context. 7th

- Conference on Computability in Europe, CiE 2011, volume 6735 of LNCS, Springer-Verlag, 2011, pp. 211–220. URL: <http://dx.doi.org/10.1007/978-3-642-21875-0>.
- [17] B. Poizat, *A course in model theory*, Springer, New York, 2000. URL: <http://dx.doi.org/10.1007/978-1-4419-8622-1>.
- [18] M. van der Put, M.F. Singer, *Galois theory of linear differential equations*, Springer, Berlin, 2003. URL: <http://dx.doi.org/10.1007/978-3-642-55750-7>.
- [19] M. Rabin, Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* 95 (1960) 341–360. URL: <http://www.jstor.org/stable/1993295>.
- [20] R. Risch, The problem of integration in finite terms, *Transactions of the American Mathematical Society* 139 (1969) 167–189. URL: <http://dx.doi.org/10.1090/S0002-9947-1969-0237477-8>.
- [21] J. Ritt, *Differential Algebra*, American Mathematical Society, New York, 1950.
- [22] A. Robinson, On the concept of a differentially closed field, *Bull. Res. Council Israel Sect. F* 8F (1959) 113–128.
- [23] M. Rosenlicht, On Liouville’s theory of elementary functions, *Pacific Journal of Mathematics* 65 (1976) 485–492. URL: <http://projecteuclid.org/euclid.pjm/1102866806>.
- [24] S. Shelah, Uniqueness and characterization of prime models over sets for totally transcendental first-order theories, *Journal of Symbolic Logic* 37 (1972) 107–113. URL: <http://www.jstor.org/stable/2272553>.
- [25] R. Soare, *Recursively Enumerable Sets and Degrees*, Springer-Verlag, New York, 1987.
- [26] R. Steiner, Computable fields and weak truth-table reducibility, in: F. Ferreira, B. Löwe, E. Mayordomo, L. Gomes (Eds.), *Programs, Proofs, Processes – Sixth Conference on Computability in Europe, CiE 2010*, volume 6158 of LNCS, Springer-Verlag, 2010, pp. 394–405. URL: http://dx.doi.org/10.1007/978-3-642-13962-8_44.