

Degree Spectra for Transcendence in Fields^{*}

Iskander Kalimullin¹, Russell Miller^{2,4}, and Hans Schoutens^{3,4}

¹ Kazan Federal University, Kremlyovskaya St. 18, 420008 Kazan, Russia
ikalimul@gmail.com

² Queens College, 65-30 Kissena Blvd., Queens, NY 11367, U.S.A.
Russell.Miller@qc.cuny.edu

³ New York City College of Technology, 300 Jay Street, Brooklyn, NY 11201 U.S.A.
hschoutens@citytech.cuny.edu

⁴ C.U.N.Y. Graduate Center, 365 Fifth Avenue, New York, NY 10016 U.S.A.

Abstract. We show that for both the unary relation of transcendence and the finitary relation of algebraic independence on a field, the degree spectra of these relations may consist of any single computably enumerable Turing degree, or of those c.e. degrees above an arbitrary fixed Δ_2^0 degree. In other cases, these spectra may be characterized by the ability to enumerate an arbitrary Σ_2^0 set. This is the first proof that a computable field can fail to have a computable copy with a computable transcendence basis.

Keywords: computability, computable structure theory, degree spectrum, field, transcendence basis

1 Introduction

It has been known since the work of Metakides and Nerode in [7] that a computable field need not have a computable transcendence basis. This result, readily established, is fundamental to the study of effectiveness for fields. Under the usual definition of computable structure, a computable field is simply a field whose domain is a computable subset of ω (usually just ω itself, the set of all nonnegative integers) and whose atomic diagram, in the language with addition and multiplication, is computable. The theorem of Metakides and Nerode shows that working with an arbitrary computable field will be difficult, as one cannot in general distinguish the algebraic elements of the field (relative to the prime subfield, either \mathbb{Q} or $\mathbb{Z}/(p)$) from those transcendental over the prime subfield.

We recently realized that the following very natural question had not been addressed: must every computable field be isomorphic to a computable field with

^{*} The research of the first author is supported by RSF Grant no. 18-11-00028; he is also funded by the Russian Ministry of Education and Science (project 1.451.2016/1.4) as a federal professor in mathematics. The second author was partially supported by Grant # 581896 from the Simons Foundation, and the second and third authors were both supported by grants from the City University of New York PSC-CUNY Research Award Program. The authors wish to acknowledge useful conversations with Dr. Kenneth Kramer.

a computable transcendence basis? It is well known that there need not exist a computable isomorphism between two isomorphic computable fields, and so it is plausible that the answer might be affirmative: two computable fields, one with a computable transcendence basis and the other without any such basis, can certainly be isomorphic. (We normally refer to isomorphic computable fields as *computable copies* of each other.) In case of an affirmative answer, one would be justified in always assuming a computable transcendence basis, as this would only require choosing a “nice” computable copy of the field in question.

Initially we were optimistic that the answer would indeed be affirmative, and even that a single Turing procedure might produce such a copy uniformly, using the original field’s atomic diagram as an oracle. (This could also then be extended to noncomputable fields.) However, conversations with our colleague Ken Kramer disabused us of that notion, and in fact we will demonstrate here that the answer is negative. For uniform procedures, the negative answer is proven in Section 2, which introduces and illustrates the use of algebraic curves of positive genus for this purpose. The remainder of the article shows that there is not even any nonuniform procedure: certain computable fields have no computable copy with any computable transcendence basis. Indeed, the spectrum of the transcendence relation on a field has many possible configurations, plenty of which do not include the degree $\mathbf{0}$. For each computably enumerable Turing degree \mathbf{c} , it is possible for transcendence to be intrinsically of degree \mathbf{c} , or for it intrinsically to compute \mathbf{c} , or for it intrinsically to enumerate a given Σ_2^0 set. The proofs here make substantial use of results on algebraic curves developed during earlier work by two of us in [11].

Transcendence bases are not in general definable, and a single field of infinite transcendence degree will have continuum-many different transcendence bases. This makes it difficult to define “the” Turing degree for transcendence bases. To address this, we use two $L_{\omega_1\omega}$ -definable relations on fields. The transcendence relation T , which is unary, holds of those elements not algebraic over the prime subfield Q of a field F :

$$x \in T \iff (\forall f \in Q[X]^*) f(x) \neq 0.$$

(Here $Q[X]^*$ is the set of nonzero polynomials over Q .) The algebraic independence relation I is a generalization of this to tuples of all arities n :

$$(x_1, \dots, x_n) \in I \iff (\forall f \in Q[X_1, \dots, X_n]^*) f(x_1, \dots, x_n) \neq 0.$$

A computable field of infinite transcendence degree will possess transcendence bases having each Turing degree $\geq_T \deg(I)$, but not of any other Turing degree: given any basis as an oracle, one can decide the independence relation on the field, and conversely, from an I -oracle, one can compute a transcendence basis for F . Thus the Turing degree of I can stand in for the set of Turing degrees of transcendence bases, as this set is the upper cone above $\deg(I)$. In turn, the unary relation T is always computable from I , although sometimes strictly below I under Turing reducibility. In the fields we consider here, we will always have $I \equiv_T T$. We remark the following useful property.

Lemma 1. *In a computable field F , for a Turing degree \mathbf{d} , every \mathbf{d} -computably enumerable transcendence basis B over the prime subfield Q is \mathbf{d} -computable.*

Proof. Given any $x \in F$, use a \mathbf{d} -oracle to list the elements b_0, b_1, \dots of B , and search for an n and a polynomial $f \in Q[X, Y_0, \dots, Y_n]^*$ with $f(x, b_0, b_1, \dots, b_n) = 0$ in F . This search must terminate, and $x \in B$ just if $x \in \{b_0, \dots, b_n\}$. \square

Our notation is generally standard. The articles [8,9] form good introductions to computable field theory, and myriad other articles have contributed to the area: [3,4,7,12,16] all have historical importance, while [5,10,11] describe related recent work in the discipline.

2 Curves of Positive Genus

Proposition 1. *Let $f(Y, Z) = 0$ define a curve over a field k of characteristic 0. If the genus of this curve is positive, then $f = 0$ has no solutions in the purely transcendental field extension $K = k(t_1, t_2, \dots)$ except those solutions within k itself. (We say that $f = 0$ has no nonconstant solutions in K .)*

Proof. A solution in K would lie within some subfield $k(t_1, \dots, t_n)$, so we prove by induction on n that no such subfield contains a nonconstant solution. For $n = 1$, we can view the extension $k(t_1)$ as an algebraic curve of genus 0 over k . If an extension $k(y, z)$ (where $f(y, z) = 0$) lies within $k(t_1)$, with $\{y, z\} \not\subseteq k$, then the Riemann-Hurwitz formula dictates that $f = 0$ must also have genus 0, contradicting the hypothesis of the proposition.

For the inductive step, suppose $y, z \in k(t_1, \dots, t_{n+1})$ satisfy $f(y, z) = 0$. We express $y = \frac{g_1}{h_1}$ and $z = \frac{g_2}{h_2}$ as rational functions of t_1, \dots, t_n over the field $k(t_{n+1})$. Of course, each of g_1, g_2, h_1, h_2 has finitely many nonzero coefficients in that field, and the pairs (g_1, h_1) and (g_2, h_2) may be taken to have no common factor. Having characteristic 0, k is infinite, so it must contain an element a such that, when t_{n+1} is replaced by a , all of these coefficients remain nonzero and no common factors are introduced. Substituting a for t_{n+1} in y and z yields a solution to $f = 0$ in $k(t_1, \dots, t_n)$. By inductive hypothesis this solution lies in k , meaning that the original y and z did not involve any of t_1, \dots, t_n (since no common factors arose to be cancelled when we made the substitution). But then y and z were a solution to $f = 0$ in $k(t_{n+1})$, hence must be a constant solution, according to the base case of the induction. \square

Corollary 1. *There is no uniform procedure for transforming a countable field into an isomorphic countable field that decides its own transcendence relation T .*

That is, there does not exist any Turing functional Ψ such that, for every atomic diagram F of a countable field with domain ω (in the signature with just $+$ and \cdot), Ψ^F computes the atomic diagram, in the larger signature with $+$, \cdot and T , of a structure with reduct F and for which T is the (unary) transcendence relation. The same holds with the ($<\omega$ -ary) algebraic independence relation I in place of T .

Proof. Suppose Ψ were such a functional. Fix an irreducible curve with affine equation $f(Y, Z) = 0$ of positive genus over \mathbb{Q} , and let F be a presentation of the field $\mathbb{Q}(y_1, y_2, y_3, \dots)(z_1)$, with $\{y_1, y_2, \dots\}$ algebraically independent over \mathbb{Q} and $f(y_1, z_1) = 0$. Then Ψ^F must compute the atomic diagram of an isomorphic field L with a transcendence relation T , say with isomorphism $h : F \rightarrow L$. Therefore T will hold of the 1-tuple $h(y_1)$ in the field L . Let σ be an initial segment of the atomic diagram of F such that Ψ^σ ensures that T holds of $h(y_1)$ and that $f(h(y_1), h(z_1)) = 0$ in L .

Now let E be a presentation of the field whose atomic diagram begins with σ . However, the atomic diagram of E (beyond σ) specifies that y_1 is in fact rational itself, in some way consistent with σ , and thus that z_1 is algebraic over \mathbb{Q} in E . (It may not be possible to make z_1 rational too, as $f = 0$ will have only finitely many solutions in \mathbb{Q} , by Faltings' proof of the Mordell Conjecture. However, there is no difficulty in making z_1 algebraic over \mathbb{Q} .) The rest of E is then generated by this portion and by elements y_2, y_3, \dots algebraically independent over \mathbb{Q} , just as in F .

Thus Ψ^E will build a field in which the domain element $h(y_1)$ is transcendental over \mathbb{Q} (being thus specified by Ψ^E) and $f(h(y_1), h(z_1)) = 0$. However, E is a purely transcendental extension of the field $k = \mathbb{Q}(z_1)$, which is algebraic over \mathbb{Q} . By Proposition 1, E does not contain any solution to $f(Y, Z) = 0$ outside of k , so every solution in E consists of elements algebraic over \mathbb{Q} . This ensures that E and the field with atomic diagram Ψ^E are not isomorphic as fields, proving the Corollary. (The result for the relation I follows directly.) \square

3 Background on Algebraic Curves

Corollary 1 proved that there is no uniform method of taking a computable field and producing a computable copy with a computable transcendence basis. Now we wish to show that a single computable field can entirely fail to have a computable copy with a computable transcendence basis. Indeed, we will establish far more specific results, with detailed descriptions of the possible degrees of transcendence bases in computable copies of the field. To do this, however, we need to work with infinitely many curves of positive genus at once, as a single curve will only allow our field to avoid being isomorphic to a single computable field with computable transcendence basis.

Fortunately, an appropriate collection of curves has already been built. We recall the following result from [11], as stated there.

Theorem 1 (Miller & Schoutens, Thm. 3.1 of [11]). *There exists a non-covering collection \mathcal{C} of curves with the effective Faltings property, over a computable ground field k .*

That is, $\mathcal{C} = \{f_0, f_1, \dots\}$ is an infinite set of smooth projective curves C_i with corresponding affine equations $f_i \in k[X, Y]$ such that:

- for each i , the function field $k(f_i)$ does not embed into $k(\mathcal{C} - \{f_i\})$; and

- the function $i \mapsto f_i(k)$ giving (a strong index for the finite set of) all solutions of $f_i(X, Y) = 0$ in k^2 is computable. (This is the *effective Faltings property*.)

In particular, the second item requires that each equation $f_i = 0$ should have only finitely many solutions by elements of k .

The specific example \mathcal{C} given in [11] is in fact a collection of Fermat curves $f_i = X^{q_i} + Y^{q_i} - 1$, for a fixed increasing computable sequence $q_0 < q_1 < \dots$ of odd prime numbers. By Fermat's Theorem, each has exactly two solutions in $k = \mathbb{Q}$, and the non-covering property for this \mathcal{C} is established in [11]. It is believed that many other computable sets of curves have the same property, but rather than pursuing that question here, we will use this same set \mathcal{C} . It should be borne in mind that not all odd primes belong to the sequence $\langle q_i \rangle_{i \in \omega}$. Indeed, this sequence is quite sparse within the primes: each element is the least prime $q_{i+1} > (4(q_i - 1)(q_i - 2))^2$, with q_0 chosen to be 5 (or any other odd prime except 3, which is ruled out because the genus of $(X^d + Y^d - 1)$ is $\frac{(d-2)(d-1)}{2}$ and we need genera > 1).

For the rest of this article we fix these curves C_0, C_1, \dots with affine equations f_0, f_1, \dots exactly as given here. The usefulness of Theorem 1 lies in the fact that it enables us to adjoin to a ground field k (such as \mathbb{Q} or $\overline{\mathbb{Q}}$) a transcendental element x and then an element y satisfying $f_i(x, y) = 0$ (so y is also transcendental) without creating any transcendental solutions to any other f_j in the new field $K = k(x)[y]/(f_i)$. Indeed, our k might already have been built this way, with pairs $(x_0, y_0), \dots, (x_{i-1}, y_{i-1})$ of transcendental solutions to f_0, \dots, f_{i-1} , say, and the new (x, y) will not generate any solutions to any of f_0, \dots, f_{i-1} that were not already in k . This allows us to work independently with the distinct polynomials f_i and their solutions, and avoids the need for priority arguments and the like.

It should be noted that $\mathbb{Q}(x_i)[y_i]/(f_i)$ actually contains eight solutions to $f_i(X, Y) = 0$. Two are the trivial solutions $(0, 1)$ and $(1, 0)$, which we can always recognize and ignore. Then the solution (x_i, y_i) generates $(\frac{-x_i}{y_i}, \frac{1}{y_i})$, $(\frac{-y_i}{x_i}, \frac{1}{x_i})$, and the transpositions of these three. Moreover, we will sometimes work over the algebraic closure $\overline{\mathbb{Q}}$ instead of \mathbb{Q} , and in this case we get $6q_i^2$ nontrivial solutions from (x_i, y_i) , since for a primitive q_i -th root θ of unity, $(x_i\theta^j, y_i\theta^k)$ will be another solution. (There will also be plenty of non-transcendental solutions within $\overline{\mathbb{Q}}$, of course.) However, it was shown by Leopoldt [6] and Tzermias [15] that no further transcendental solutions will exist; this result is also used in [11], appearing there as Theorem 4.4.

4 Examples of Degree Spectra

Our initial goal was to produce a computable field such that no computable copy of the field has a computable transcendence basis. In fact, we will give a much more specific answer to the question, using the well-established notion of the degree spectrum of a relation.

Definition 1. For a computable structure \mathcal{A} and a relation R on \mathcal{A} , the Turing degree spectrum of R on \mathcal{A} is the set of all Turing degrees of images of R in computable structures isomorphic to \mathcal{A} :

$$\text{DgSp}_{\mathcal{A}}(R) = \{\deg(g(R)) : g : \mathcal{A} \rightarrow \mathcal{B} \text{ is an isomorphism onto a computable } \mathcal{B}\}.$$

In many contexts this definition is restricted to n -ary relations R , but it applies equally well to finitary relations, i.e., those defined on all finite tuples from $(\mathcal{A})^{<\omega}$, of arbitrary length.

Notice first that both the (unary) transcendence relation T on a computable field and the (finitary) relation I of algebraic independence are definable in the field by computable infinitary Π_1^0 formulas. Therefore, in every computable field F , both T and I will be Π_1^0 sets, hence of c.e. Turing degree. This places an upper bound on the complexity of the degrees in $\text{DgSp}_F(T)$ and $\text{DgSp}_F(I)$, since all such degrees must be $\leq_T \mathbf{0}'$. Even below that bound, it also rules out a number of further candidates, namely those degrees that do not contain any c.e. set. (We call these *properly* Δ_2^0 degrees, meaning that they are Δ_2^0 but not Σ_1^0 .)

Our ultimate goal is to know the degrees of the transcendence bases in the various computable copies of F . Recall, however, that these are precisely the degrees $\geq_T \deg(I)$. Thus, once one knows the Turing degree of I in a particular computable copy of the field, one knows all the degrees of transcendence bases in that copy, and so we view $\text{DgSp}_F(I)$ as a reasonable answer to the question.

Our first example shows that, for a fixed computably enumerable Turing degree \mathbf{c} , the relations T and I can both be *intrinsically of degree* \mathbf{c} . This term was used in [2], in which Downey and Moses showed that the relation of adjacency in a computable linear order can be intrinsically of degree $\mathbf{0}'$. Subsequently, Downey, Lempp, and Wu showed in [1] that the only degrees \mathbf{c} for which the adjacency relation can be intrinsically of degree \mathbf{c} are $\mathbf{c} = \mathbf{0}'$ and (if the adjacency relation is finite) $\mathbf{c} = \mathbf{0}$. Therefore Theorem 2 distinguishes the situation for transcendence and for independence in fields from that of adjacency in linear orders.

Theorem 2. For each computably enumerable Turing degree \mathbf{c} , there exists a computable field F for which the spectrum of the transcendence relation T and of the independence relation I are both the singleton $\{\mathbf{c}\}$.

Proof. Fixing a computable enumeration $\langle C_s \rangle_{s \in \omega}$ of a c.e. set $C \in \mathbf{c}$, we can describe the isomorphism type of our field quickly. For every $i \notin C$, it will contain a transcendental element x_i (over the ground field \mathbb{Q}) and an additional element y_i with $f_i(x_i, y_i) = 0$. Moreover, these elements x_i will form a transcendence basis, as i ranges over \overline{C} . For each $i \in C$, the field will contain elements called x_i and y_i , again satisfying $f_i(x_i, y_i) = 0$, but this x_i will lie within \mathbb{Q} , making y_i algebraic over \mathbb{Q} . These x_i and y_i (for all $i \in \omega$) will generate the field. (The choice of which rational number equals x_i , for $i \in C$, will depend on the least s with $i \in C_s$.)

Next we give a computable presentation F of this field. At stage 0, F_0 consists of a finite substructure of the field \mathbb{Q} (with the operations viewed as relations,

so that it makes sense for \mathbb{Q} to have a finite substructure). At stage $s + 1$, we add elements x_s and y_s to F_s , along with as many new elements as are needed in order for the relational atomic diagram of F_{s+1} to specify that $f_s(x_s, y_s) = 0$ (but without making x_s itself algebraic over F_s). Then, for the least $i \leq s$ (if any) such that $i \in C_s$ and we have not yet acted on behalf of i , we add enough new elements to F_{s+1} and define the operations on them to make x_i lie within \mathbb{Q} (in a way consistent with F_s , of course: nothing in the atomic diagram should ever be redefined). This must be possible, since x_i has been treated as a transcendental up until this stage. Finally, we take another step to close F under the field operations, adding another element and extending the relations in F_{s+1} in a way consistent with the principle that the set

$$\{x_i : i \leq s \text{ \& we have not yet acted on behalf of } i\}$$

should form a transcendence basis for F_{s+1} . That is, we make sure not to create any algebraic relations involving these x_i 's, and all new elements added to the field are generated by $\{x_i, y_i : i \leq s\}$. This is the entire construction, and it is clear that it does ultimately build a computable field $F = \cup_s F_s$. Furthermore, $\{x_i : i \notin C\}$ will indeed be a transcendence basis for F , and every x_i with $i \in C$ will lie within \mathbb{Q} in the field F .

We now argue that for every computable field $E \cong F$, the transcendence relation T and the independence relation I on E have $T \equiv_T I \equiv_T C$. First, given a C -oracle, we enumerate a transcendence basis for F by collecting, for each $i \notin C$, the first coordinate of the first pair (x, y) that we find in E^2 for which $f_i(x, y) = 0 \neq xy$. (This pair is not unique, as mentioned in Section 3, but the six possible x -values are pairwise interalgebraic.) By the construction, this enumerates a transcendence basis B for F , hence computes one, by Lemma 1, and from B we can compute I and T .

To show that $C \leq_T T$, we claim that $i \in \overline{C}$ just if there exists a pair (x, y) of elements of E with $x \in T$ and $f_i(x, y) = 0$ in E . (Thus C is Π_1^T , as well as Σ_1 .) Indeed, for $i \notin C$, the isomorphic image in E of the elements (x_i, y_i) from F will be such a pair. For the converse, suppose $i \in C$. Then x_i and y_i were made algebraic at some stage in the construction of F , and by Theorem 1 with $k = \mathbb{Q}(y_j : j \in C) \subseteq F$ (which is the subfield containing all elements algebraic over \mathbb{Q}), the function field of the collection $\{f_j : j \notin C\}$ over k does not contain any nontrivial solution to $f_i = 0$. This function field is isomorphic to F itself, so we have $C \leq_T T$. \square

Theorem 2 answers the initial question posed above. Theorems 3 and 4 will provide further examples.

Corollary 2. *There exists a computable field F such that no computable field E isomorphic to F has a computable transcendence basis, nor even a computable transcendence relation.* \square

Our next example shows that the relations T and I can also *intrinsically compute* a c.e. degree \mathbf{c} , in the sense that the spectra of T and I can equal the

upper cone above \mathbf{c} (subject to the restriction that these spectra only contain c.e. degrees). Once again, this parallels a result of Downey and Moses in [2] for linear orders. In Corollary 3 below, we will generalize this result to all Δ_2^0 degrees \mathbf{c} , which is not known (to us) to be possible for adjacency on linear orders.

Theorem 3. *For each computably enumerable Turing degree \mathbf{c} , there exists a computable field F for which the spectrum of the transcendence relation T and of the independence relation I are both*

$$\text{DgSp}_F(T) = \text{DgSp}_F(I) = \{\mathbf{d} \geq \mathbf{c} : \mathbf{d} \text{ is a c.e. degree}\}.$$

Proof. Fix some c.e. set $C \in \mathbf{c}$, with a computable enumeration $\langle C_t \rangle_{t \in \omega}$ by finite nested sets C_t . The field F for this degree is the field

$$F = \overline{\mathbb{Q}}(x_k)[y_k]/(f_k(x_k, y_k)),$$

with k ranging over the set $\overline{C} \oplus \omega$. That is, F has the algebraic closure $\overline{\mathbb{Q}}$ as its ground field, and contains an algebraically independent set $\{x_{2i} : i \in \overline{C}\} \cup \{x_{2j+1} : j \in \omega\}$ of elements, along with corresponding elements y_{2i} and y_{2j+1} that “tag” the individual x -elements by forming solutions to $f_{2i} = 0$ or $f_{2j+1} = 0$. (The reason for the odd-indexed elements x_{2j+1} will become clear below: they will give us the upward closure we desire.)

To see that F has a computable presentation, start building a computable copy of $\overline{\mathbb{Q}}$, with only finitely many elements added at each stage. At stage $s + 1$, we add new elements x_s and y_s to the field, with $f_s(x_s, y_s) = 0$, and treat x_s as a transcendental over all previously existing field elements. For odd values $s = 2j + 1$, we simply continue at each subsequent stage to build the field, with x_s remaining transcendental. For even $s = 2i$, at each subsequent stage $t > s + 1$, we check whether $i \in C_t$. As long as $i \notin C_t$, we simply add to the field the next element generated by x_s , continuing to treat x_s as transcendental over the preceding elements. However, for the first t (if any) with $i \in C_t$, we switch strategies and make x_s a rational number, finding some way to do this that is consistent with the finite portion of the atomic diagram of F that has already been defined. Of course, this also makes y_s algebraic over \mathbb{Q} , though not rational. This enlarges our presentation of the ground field $\overline{\mathbb{Q}}$, of course, but since only finitely much of $\overline{\mathbb{Q}}$ had been built so far, it is easy to incorporate x_s and y_s into it and to continue building $\overline{\mathbb{Q}}$, including them, at each subsequent stage.

Now for any computable field $E \cong F$, with transcendence relation T , we can compute C from T . Indeed, by Theorem 1, a number i lies in \overline{C} if and only if E contains transcendental elements x and y such that $f_{2i}(x, y) = 0$, so \overline{C} is Σ_1^T , while C is Σ_1 . Thus $\text{DgSp}_F(T)$ contains only degrees above \mathbf{c} , and these must all be c.e. degrees, as the relation T is definable in F by a computable infinitary Π_1^0 formula. The same analysis applies to the independence relation I .

To prove the reverse inclusion, let \mathbf{d} be any c.e. degree that computes \mathbf{c} , and fix some c.e. set $D \in \mathbf{d}$ with computable enumeration $\langle D_s \rangle_{s \in \omega}$. We build a specific computable copy E of F in which $T \equiv_T I \equiv_T D$, by a process quite similar to the above construction of F itself. E includes a copy of $\overline{\mathbb{Q}}$, built slowly,

with only finitely many elements added at each stage. Once again, the even-indexed x_{2i} and y_{2i} are added at stage $2i$ and treated as transcendental until i enters C , at which point x_{2i} becomes rational. The odd-indexed elements x_{2j+1} and y_{2j+1} are added at stage $2j+1$ and treated as transcendental until (if ever) we reach a stage s with $j \in D_s$. If such a stage occurs, then this x_{2j+1} is made rational at that stage (in the same way as with x_{2i} if i enters C), and we adjoin to the field new elements x'_{2j+1} and y'_{2j+1} , again with x'_{2j+1} transcendental over all existing elements of F and with $f_{2j+1}(x'_{2j+1}, y'_{2j+1}) = 0$. These new elements will forever remain transcendental over the ground field $\overline{\mathbb{Q}}$, and the original x_{2j+1} and y_{2j+1} have now been “swallowed up” by $\overline{\mathbb{Q}}$. Thus the E built here is indeed isomorphic to F , and is a computable field. However, from the transcendence relation T on F , we can compute D , since $j \in \overline{D}$ if and only if the original x_{2j+1} lies in T . Conversely, from a D -oracle we can decide whether x_{2j+1} will ever be swallowed up by $\overline{\mathbb{Q}}$ or not, and also (since $C \leq_T D$) whether x_{2i} will remain transcendental in E or not. Thus $T \equiv_T D$, and so $\mathbf{d} \in \text{DgSp}_F(T)$.

The same argument also shows that $\mathbf{d} \in \text{DgSp}_F(I)$, since the elements x_s that stay transcendental forever form a transcendence basis for E , from which we can compute the independence relation. It should be remarked here, as in Section 3, that the first transcendental solution to $f_k = 0$ that one finds in E will only be one of the $6q_i^2$ such solutions, but in enumerating a transcendence basis, it is safe to choose the first coordinate of the first transcendental solution we find, and then to ignore all other solutions to the same f_k , as their coordinates are all either in $\overline{\mathbb{Q}}$ or interalgebraic with the coordinate we chose. In fact, since the automorphism group of E acts transitively on these solutions, there is nothing to distinguish one such choice from another. \square

Our next result suggests that many spectra of transcendence relations can be viewed as upper cones of enumeration degrees. To be clear, the spectrum is still a set of Turing degrees, by definition, but the defining property of the spectrum may be the ability to enumerate a particular set. (It remains true that only c.e. degrees may lie in $\text{DgSp}_F(T)$, although other Turing degrees may enumerate the same set. So the spectrum will never truly be an upper cone of e -degrees.)

Theorem 4. *Let S be any Σ_2^0 subset of ω . Then there exists a computable field F such that*

$$\text{DgSp}_F(T) = \text{DgSp}_F(I) = \{\text{c.e. degrees } \mathbf{d} : S \in \Sigma_1^{\mathbf{d}}\}.$$

That is, $\text{DgSp}_F(T)$ contains exactly those c.e. degrees that have the ability to enumerate S .

Proof. Since S is Σ_2^0 , there exists a computable total “chip function” $h : \omega \rightarrow \omega$ such that $S = \{n : h^{-1}(n) \text{ is finite}\}$. The field F we use for this set is the field

$$F = \overline{\mathbb{Q}}(x_k)[y_k]/(f_i(x_k, y_k)),$$

with k ranging over the set $S \oplus \omega$, much as in Theorem 3 but using the set S itself instead of its complement.

To give a computable presentation of F , we start building a copy of the field $\overline{\mathbb{Q}}(x_k)[y_k]/(f_k(x_k, y_k))$ with k ranging over all of ω , so that every x_k is initially treated as a transcendental. For odd $k = 2j + 1$, x_k stays transcendental throughout this construction. For even $k = 2i$, we write $x_{2i,0} = x_{2i}$ for the initial element described above. At each stage $s + 1$, we check whether $h(s) = i$. If not, then we keep $x_{2i,s+1} = x_{2i,s}$ and continue to treat it as a transcendental. If $h(s) = i$, however, then we suspect that i might not lie in S (since $h^{-1}(i)$ might turn out to be infinite). In this case we make the current $x_{2i,s}$ into a rational number, consistently with the finite portion of the atomic diagram of F built so far, and thus make $y_{2i,s}$ algebraic. We then adjoin new elements $x_{2i,s+1}$ and $y_{2i,s+1}$ to F , treating $x_{2i,s+1}$ as transcendental and setting $f_{2i}(x_{2i,s+1}, y_{2i,s+1}) = 0$. We continue building $\overline{\mathbb{Q}}$ as the ground field, now incorporating the old $x_{2i,s}$ and $y_{2i,s}$ into it, and continuing closing F itself under the field operations, but always adding only finitely many new elements at each stage. This completes the construction, and it is clear that $f_{2i}(X, Y) = 0$ will have a solution by transcendental elements in F just if $h^{-1}(i)$ is finite, which is to say, just if $i \in S$. The rest of the construction then makes it clear that the field we have built is a computable copy of the field F described above.

Given any computable field $E \cong F$, let T be the transcendence relation on E . Then, given a T -oracle, we may search in E for a solution to $f_{2i}(X, Y) = 0$ using transcendental elements x and y . If we find one, then by the definition of F we know that $i \in S$. Conversely, if $i \in S$, then such a solution exists, and we will eventually find it. Thus S is c.e. relative to the degree \mathbf{d} of T , as required.

Conversely, fix any c.e. degree \mathbf{d} such that S is $\Sigma_1^{\mathbf{d}}$, and fix a c.e. set $D \in \mathbf{d}$ and a computable enumeration of it. Also fix an index e such that $S = W_e^D = \text{dom}(\Phi_e^D)$; we will use this below to give a computable chip function for S , similar to that used in the original computation of F but specific to this D . To build a computable copy E of F whose transcendence relation T satisfies $T \equiv_T D$, we use the strategy from Theorem 3. The elements $x_{2i,0}$ and $y_{2i,0}$ are defined and initially treated as transcendentals. However, at each stage $s + 1$, the current $x_{2i,s}$ and $y_{2i,s}$ are made into algebraic elements and replaced by new elements $x_{2i,s+1}$ and $y_{2i,s+1}$ unless $\Phi_{e,s}^{D_s}(i) \downarrow$ with some use u such that $D_{s+1} \upharpoonright u = D_s \upharpoonright u$. This is our new chip function for S : if $i \in S = \text{dom}(\Phi_e^D)$, then there will be some s_0 such that we keep x_{2i,s_0} transcendental at all stages $\geq s_0$; whereas if $i \notin S$, then for every stage $s + 1$ at which $\Phi_{e,s}^{D_s}(i) \downarrow$ with a use u , there must be some $t > s$ with $D_t \upharpoonright u \neq D_s \upharpoonright u$, so that $x_{2i,s}$ will be made algebraic at stage $t + 1$ and replaced by a new $x_{2i,t+1}$.

We also revamp the construction for the odd-indexed elements x_{2j+1} and y_{2j+1} , using exactly the same process as in the proof of Theorem 3. If we ever reach a stage at which j enters D , then we turn x_{2j+1} into a rational number, consistently with the construction so far, and adjoin a new transcendental x'_{2j+1} and corresponding y'_{2j+1} with $f_{2j+1}(x'_{2j+1}, y'_{2j+1}) = 0$ in E . This completes the construction of E , which is clearly a computable field and isomorphic to F .

Now from an oracle for the transcendence relation T on E , we can determine whether x_{2j+1} is algebraic in E or not, thus deciding whether or not $j \in D$.

Thus $D \leq_T T$. For the reverse reduction, we claim that with a D -oracle we can enumerate a transcendence basis B for E , thus deciding the independence relation I on E , which in turn computes T . This will prove $I \equiv_T T \equiv_T D$ as required. The D -oracle allows us to decide, for each j , whether $j \in D$, from which we determine either that x_{2j+1} lies in B (if $j \notin D$) or that x'_{2j+1} does (if $j \in D$, in which case we identify x'_{2j+1} by waiting for a stage at which j has entered D). Next, for each i , we watch for a stage s at which $\Phi_{e,s}^{D_s}(i) \downarrow$ with a use u such that $D_s \upharpoonright u = D \upharpoonright u$. The D -oracle allows us to check this, and if we ever find such an s , then we enumerate $x_{2i,s+1}$ into our basis, since the computable enumeration of D will never again change below u . (This is where our argument would fail if \mathbf{d} were a properly Δ_2^0 degree, rather than a c.e. degree. With only a computable approximation to D , we could not be sure whether $D_s \upharpoonright u$ would ever again change, even knowing that $D_s \upharpoonright u = D \upharpoonright u$.) Thus we have enumerated exactly the set of elements x_k given when we first defined the isomorphism type of F above, and this set is a transcendence basis for E . \square

Corollary 3. *Let \mathbf{c} be any Δ_2^0 Turing degree. Then there exists a computable field F such that*

$$\text{DgSp}_F(T) = \text{DgSp}_F(I) = \{\text{c.e. degrees } \mathbf{d} : \mathbf{c} \leq_T \mathbf{d}\}.$$

In particular, both $\text{DgSp}_F(T)$ and $\text{DgSp}_F(I)$ can fail to contain a least degree.

The condition of not containing a least degree also holds for many of the spectra given in Theorem 4. In particular, if no set A' (with A c.e.) is 1-equivalent to S , then the spectrum has no least degree, as $\deg(D) \in \text{DgSp}_F(T)$ if and only if $S \leq_1 D'$. The proof uses the Sacks Jump Theorem (see [13], or [14, Thm. VII.3.1]), to avoid the upper cone above a hypothetical least degree.

Proof. Fix a set $C \in \mathbf{c}$, and apply Theorem 4 to the Σ_2^0 set $S = C \oplus \overline{C}$. The ability to enumerate S is exactly the ability to compute C , so the corollary follows. (To avoid having a least degree in the spectra, just choose a degree $\mathbf{c} \leq \mathbf{0}'$ that is not c.e. and has no least c.e. degree above it.) \square

Corollary 3 extends Theorem 3 to the Δ_2^0 degrees. The result can be viewed as an upper-cone result, but in a somewhat odd way. If the \mathbf{c} in the corollary is c.e., then the corollary merely repeats Theorem 3. If \mathbf{c} is Δ_2^0 but not c.e., then the degree spectrum is the restriction of the upper cone above \mathbf{c} to the c.e. degrees, and therefore does not contain the base degree \mathbf{c} itself, nor any other non-c.e. degree $\geq_T \mathbf{c}$.

References

1. R.G. Downey, S. Lempp, & G. Wu: On the complexity of the successivity relation in computable linear orderings, *Journal of Mathematical Logic* **10** (2010) 01n02, 83–99.

2. R.G. Downey & M.F. Moses: Recursive linear orders with incomplete successivities, *Transactions of the American Mathematical Society* **326** (1991) 2, 653–668.
3. Yu.L. Ershov: Theorie der Numerierungen, *Zeits. Math. Logik Grund. Math.* **23** (1977), 289–371.
4. A. Frohlich & J.C. Shepherdson: Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407–432.
5. A. Frolov, I. Kalimullin, & R. Miller: Spectra of algebraic fields and subfields, in *Mathematical Theory and Computational Practice: Fifth Conference on Computability in Europe, CiE 2009*, eds. K. Ambos-Spies, B. Löwe, & W. Merkle, *Lecture Notes in Computer Science* **5635** (Berlin: Springer-Verlag, 2009), 232–241.
6. H.-W. Leopoldt: Über die Automorphismengruppen des Fermatkörpers, *Journal of Number Theory* **56** (1996) 2, 256–282.
7. G. Metakides & A. Nerode: Effective content of field theory, *Annals of Mathematical Logic* **17** (1979), 279–320.
8. R. Miller: Computable fields and Galois theory, *Notices of the American Mathematical Society* **55** (August 2008) 7, 798–807.
9. R. Miller: An introduction to computable model theory on groups and fields, *Groups, Complexity and Cryptology* **3** (2011) 1, 25–46.
10. R. Miller, B. Poonen, H. Schoutens, & A. Shlapentokh: A computable functor from graphs to fields, *Journal of Symbolic Logic* **83** (2018) 1, 326–348.
11. R. Miller & H. Schoutens: Computably categorical fields via Fermat’s Last Theorem, *Computability* **2** (2013) 51–65.
12. M. Rabin: Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341–360.
13. G.E. Sacks: Recursive enumerability and the jump operator, *Transactions of the American Mathematical Society* **108** (1963), 223–239.
14. R.I. Soare: *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).
15. P. Tzermias: The group of automorphisms of the Fermat curve, *Journal of Number Theory* **53** (1995) 1, 173–178.
16. B.L. van der Waerden: *Algebra*, volume I, trans. F. Blum & J.R. Schulenberg (New York: Springer-Verlag, 1970 hardcover, 2003 softcover).