Spectra of Algebraic Fields and Subfields

Andrey Frolov[1]⋆, Iskander Kalimullin[1]⋆⋆ and Russell Miller[2]⋆⋆⋆

[1] N.G. Chebotarev Research Institute of Mechanics and Mathematics
Kazan State University
Universitetskaya St. 17
Kazan 420008, Tatarstan Russia
Andrey.Frolov@ksu.ru
Iskander.Kalimullin@ksu.ru
[2] Queens College of CUNY
65-30 Kissena Blvd., Flushing, NY 11367 USA
and
The CUNY Graduate Center
365 Fifth Avenue, New York, NY 10016 USA
Russell.Miller@qc.cuny.edu
http://qcpages.qc.cuny.edu/~rmiller

Abstract. An algebraic field extension of $\mathbb{Q}$ or $\mathbb{Z}/(p)$ may be regarded either as a structure in its own right, or as a subfield of its algebraic closure $\overline{F}$ (either $\overline{\mathbb{Q}}$ or $\overline{\mathbb{Z}/(p)}$). We consider the Turing degree spectrum of F in both cases, as a structure and as a relation on $\overline{F}$, and characterize the sets of Turing degrees that are realized as such spectra. The results show a connection between enumerability in the structure F and computability when F is seen as a subfield of $\overline{F}$.

Key words: Computability, computable model theory, field, algebraic, spectrum.

1 Introduction

By definition, the *spectrum of a structure* $\mathfrak{A}$, written $\mathrm{Spec}(\mathfrak{A})$, is the set of all Turing degrees of structures $\mathfrak{B} \cong \mathfrak{A}$ with domain ω. The intention is to measure the inherent complexity of the isomorphism type of $\mathfrak{A}$, by describing the set of Turing degrees capable of computing a copy of $\mathfrak{A}$. We restrict the domain to ω in order to measure only the complexity of the functions and relations of $\mathfrak{A}$, without interference from an unusual choice of domain.

Likewise, the *spectrum of a relation* R on a computable structure $\mathfrak{M}$ is the set $\mathrm{DgSp}_{\mathfrak{M}}(R)$ of Turing degrees of all images of R under isomorphisms from $\mathfrak{M}$

⋆ The first two authors were partially supported by RFBR grants 05-01-00605 and 09-01-97010.
⋆⋆ The second author was partially supported by RF President grant MK-4314.2008.1.
⋆⋆⋆ The corresponding author was partially supported by Grant # 13397 from the Templeton Foundation, and by Grants # 69723-00 38 and 61467-00 39 from The City University of New York PSC-CUNY Research Award Program.

onto other computable structures $\mathfrak{B}$. (In general R will not lie in the signature of $\mathfrak{M}$, for otherwise its spectrum would contain only the degree $\mathbf{0}$.) This measures the complexity of the relation R, again by asking how simple or complex R can become under different presentations of the structure $\mathfrak{M}$. As before, restricting to computable presentations $\mathfrak{B}$ of $\mathfrak{M}$ keeps any extra complexity from creeping in when we select the underlying structure.

These two notions are compared at some length in [9]. Both are common in the study of computable model theory. In this paper we apply them to algebraic fields, by which we mean algebraic extensions of any of the prime fields $\mathbb{Q}$ and $\mathbb{Z}/(p)$ (for prime p). Theorem 2 will describe precisely the possible spectra of such fields, viewed as structures. Of course, the algebraic fields are exactly the subfields of the algebraic closures $\overline{\mathbb{Q}}$ and $\overline{\mathbb{Z}/(p)}$, so it is also natural to view such fields as unary relations on computable presentations of these algebraic closures. Theorem 4 will describe precisely the possible spectra of algebraic fields as relations in this sense. In Section 5, we will then compare the possible spectra of these fields in these two contexts and shed light on the relative usefulness of these two methods of presenting a field,

Notation and background for computability theory can be found in [19], and for fields in [10], [21], and many other sources. Many influential papers on computable fields have appeared over the years, among which we would cite in particular [7], [17], [4], [13], and [20]. For a basic introduction to the subject, [14] is useful. We will assume many standard field-theoretic results, but one is so important as to bear mention right now.

Theorem 1 (Kronecker [12]; or see [3]). *The fields $\mathbb{Q}$ and $\mathbb{Z}/(p)$, and all finite extensions of these fields, have splitting algorithms. (That is, the set of irreducible polynomials in $\mathbb{Q}[X]$ is computable, and likewise in $\mathbb{Z}/(p)[X]$.) Moreover, to determine a splitting algorithm for a computable field $E = F(x_1, \ldots, x_n)$, we need to know only a splitting algorithm for F, an enumeration of F within E, the atomic diagram of E, and (for each $i < n$) whether x_{i+1} is algebraic or transcendental over $F(x_1, \ldots, x_i)$.*

A splitting algorithm for E enables us to compute, for any $p(X) \in E[X]$, how many roots p has in E. With this information, we may then find all those roots in E (and know when we are finished!).

2 Spectra of Fields

The first part of the main theorem of this section follows by relativizing a theorem of Ershov [4, §12, Thm. 2]. Ershov's proof has only been published in German, and ours differs from his in our use of the Theorem of the Primitive Element. Moreover, the last two parts of Theorem 2 are (we believe) new. Appendix A includes the proof.

Theorem 2. *1. For every algebraic field F, there exists a set $V_F \in 2^\omega$ with*

$$\mathit{Spec}(F) = \{\boldsymbol{d} : V_F \text{ is c.e. in } \boldsymbol{d}\}.$$

2. *Conversely, for every $V \in 2^\omega$, there exists an algebraic field F of arbitrary characteristic such that*

$$Spec(F) = \{\boldsymbol{d} : V \text{ is c.e. in } \boldsymbol{d}\}.$$

3. *Finally, if Q is the prime field of any fixed characteristic c, then the set $\{V_F \upharpoonright n \ : \ F$ algebraic over Q & $n \in \omega\}$ forms a computable extendible subtree $T \subset 2^{<\omega}$ with no isolated paths, and every path through T is of the form V_F for some F.*

For reference in Section 4, we give the definition of V_F here, for arbitrary algebraic extensions F of any prime field Q. First we define two c.e. sets T_F and U_F, each of which depends only on the isomorphism type of F. Let $p_0(X), p_1(X), \ldots$ list the irreducible polynomials in $Q[X]$, and set $d_i = \deg(p_i(X))$ and $D_i = \sum_{j<i} d_j$. Then, for each i and each $k < d_i$, we define

$$D_i + k \in T_F \iff F \text{ contains at least } (k+1) \text{ distinct roots of } p_i(X).$$

So the first d_0 bits of T_F tell the exact number of roots of $p_0(X)$ in F, and the next d_1 bits tell the number of roots of $p_1(X)$, etc. Obviously T_F is computably enumerable in the degree of any field isomorphic to F.

To define U_F, we need to consider pairs of polynomials. Since we have a splitting algorithm for Q, it is computable whether a pair $\langle g_0(X), g_1(X,Y)\rangle \in (Q[X] \times Q[X,Y])$ satisfies both of the following conditions.

- $Q[X]/(g_0)$ is a field. This is equivalent to demanding that g_0 be nonconstant and irreducible in $Q[X]$.
- g_1, when viewed as a polynomial in Y, is irreducible in the polynomial ring $(Q[X]/(g_0))[Y]$. (The coefficients of Y in g_1 are really in $Q[X]$, of course; here we consider their images in $Q[X]/(g_0)$.) Equivalently, if x is a root of g_0, then $g_1(x,Y)$ is irreducible in the polynomial ring $Q(x)[Y]$.

So we may computably enumerate a list $\boldsymbol{G}_0, \boldsymbol{G}_1, \ldots$ of all pairs satisfying these properties, writing $\boldsymbol{G}_j = \langle g_{j0}, g_{j1}\rangle$. Now define

$$U_F = \{j : (\exists x, y \in F)[g_{j0}(x) = 0 = g_{j1}(x,y)]\}.$$

Again this set is c.e. in the degree of any field isomorphic to F.

In fact, $T_F \leq_1 U_F$, so it is not really necessary to take the join of T_F and U_F, but we consider it more perspicuous to do so. Define $V_F = T_F \oplus U_F$. The remainder of the proof of Theorem 2 appears in the Appendix A.

We regard Theorem 2 as a substantial step in classifying the possible spectra of models of standard theories of mathematics. We say this informally, because of course, the class of algebraic fields is not an EC class, nor even an EC_Δ class: there is no set T of first-order sentences for which the algebraic fields (even of a fixed characteristic) are precisely the models of T. So we must speak, non-rigorously, of standard classes $\mathcal{K}$ of mathematical structures. The goal, for a

given $\mathcal{K}$, is to provide a criterion ψ such that for all subsets $\mathcal{S}$ of the set of Turing degrees,

$$\psi \text{ holds of } \mathcal{S} \iff \mathcal{S} = \mathrm{Spec}(\mathfrak{M}) \text{ for some } \mathfrak{M} \in \mathcal{K}.$$

Ideally, ψ should use only set- and degree-theoretic properties: Turing reducibility, jumps, and so forth. In our case, ψ was the property that there exists a $V \subseteq \omega$ such that $\mathcal{S}$ is the set of degrees in which V is c.e.

Now there do exist classes $\mathcal{K}$, even EC-classes, for which such a ψ is known. For example, take θ to be the conjunction of the axioms for dense linear orders: the only possible spectrum $\mathcal{S}$ of a countable model of θ is the set of all Turing degrees. The one other class $\mathcal{K}$ for which a nontrivial criterion is known is the class of torsion-free abelian groups: Coles, Downey, and Slaman examined this class in [2], mainly with an eye to studying 1-degrees, and while they did not state it specifically, it is clear from their work that parts 1 and 2 of Theorem 2 would also hold with "algebraic field" replaced by "torsion-free abelian group." (These two classes are closely linked in computable model theory, and neither is an EC_Δ class.) A few other examples can be found, in addition to dense linear orders: the class of complete graphs, for example, or the class of finite models in a given signature, but for these classes the condition ψ is quite trivial. There is no ψ known for any of the following classes: graphs, trees, linear orders, Boolean algebras, abelian groups, p-groups, fields, or rational vector spaces. Moreover, for several pairs of these classes, the conditions ψ (while unknown as yet) must be nonequivalent. For each of these classes, we would regard the discovery of a condition ψ as an important result.

3 Consequences

Several known theorems combine nicely with Theorem 2. For example, we have a result on the jump degree of an algebraic field F. Recall that the *jump degree* of a countable structure $\mathfrak{M}$ is the least degree under $\leq_T$ (if any exists) in the set $\{\boldsymbol{d}' : \boldsymbol{d} \in \mathrm{Spec}(\mathfrak{M})\}$. Jump degrees are studied in [11] and several other papers. The useful result for us was proven by Coles, Downey, and Slaman.

Theorem 3 (Coles, Downey, Slaman; see [2]). *For all $A \subseteq \omega$ there is a set B such that (1) A is c.e. in B, and (2) every set C with A c.e. in C satisfies $B' \leq_T C'$.*

Applying this result along with Theorem 2 immediately gives jump degrees.

Corollary 1. *Every algebraic field has a jump degree.* □

In [18], Richter constructed a set A such that $\{\boldsymbol{d} : A \text{ is c.e. in } \boldsymbol{d}\}$ has no least Turing degree under $\leq_T$. Combining this with Theorem 2 yields a result already proven by other means by Calvert, Harizanov, and Shlapentokh:

Corollary 2 (Calvert, Harizanov, Shlapentokh; see [1]). *There exists an algebraic field whose spectrum contains no least Turing degree.* □

On the other hand, given any set S, the collection of degrees which can enumerate the join $S \oplus \overline{S}$ (where $\overline{S}$ is the complement of S) is the upper cone of Turing degrees above S. This allows us to reprove another result from the above paper.

Corollary 3 (Calvert, Harizanov, Shlapentokh; see [1]). *Every upper cone of Turing degrees forms the spectrum of some algebraic field.* □

4 Subfields of $\overline{\mathbb{Q}}$ and $\overline{\mathbb{Z}/(p)}$

Pick any prime field Q (either the rationals $\mathbb{Q}$, or $\mathbb{Z}/(p)$), and fix one computable presentation $\overline{Q}$ of the algebraic closure of Q. Since $\overline{Q}$ is computably categorical, we see that for any relation R on $\overline{Q}$, and any degree $\boldsymbol{d} \in \mathrm{DgSp}_{\overline{Q}}(R)$, there is a relation S on $\overline{Q}$ itself such that $(\overline{Q}, R) \cong (\overline{Q}, S)$ and $\boldsymbol{d} = \deg(S)$. (This appears in [9, Lemma 1.6]: if $(\overline{Q}, R) \cong (\mathfrak{A}, T)$ with $\boldsymbol{d} = \deg(T)$, take $S = g(T)$, where $g : \mathfrak{A} \to \overline{Q}$ is a computable isomorphism.)

It is well-known that any field isomorphism extends to an isomorphism of the algebraic closures of the fields; we express this in the following lemma.

Lemma 1. *For an algebraic field F, let f and g be any two embeddings of F into $\overline{Q}$. Then $(g \circ f^{-1})$ extends to an automorphism of $\overline{Q}$.* □

Consequently, we may speak without ambiguity of the degree spectrum of F, viewed as a relation on $\overline{Q}$; the exact choice of a subfield of $\overline{Q}$ isomorphic to F is irrelevant, since any two choices have the same degree spectrum.

When an algebraic field F is viewed as a structure, its spectrum is defined by a condition of computable enumerability. When the same field is viewed as a subfield of $\overline{Q}$, its spectrum as a relation will no longer involve enumerability: instead of simply waiting for a root of a polynomial $p(X) \in Q[X]$ to appear in F, we can find all the roots of $p(X)$ in $\overline{Q}$ and check immediately (using an oracle for the subfield F) whether any of those finitely many roots actually lies in F. This is easily seen in the case of normal extensions, which is now practically trivial.

Proposition 1. *Let F be a normal algebraic extension of the prime field Q, and fix a computable copy of $\overline{Q}$. Then F has exactly one homomorphic image in $\overline{Q}$, and $DgSp_{\overline{Q}}(F) = \{deg(T_F^*)\}$, where $T_F^* = \{i : \exists a \in F(p_i(a) = 0)\}$.*

Indeed, this proposition would apply not only to Q, but to any ground field E with a splitting algorithm, provided that we consider only embeddings of F into $\overline{E}$ which extend a given embedding $h : E \hookrightarrow \overline{E}$.

Proof. Fix any $g : F \hookrightarrow \overline{Q}$, and pick any $x \in \overline{Q}$. Either F contains no roots of the minimal polynomial $p(X)$ of x in $Q[X]$ (so $x \notin g(F)$), or, by normality, F contains $d = \deg(p)$ roots of $p(X)$. But $\overline{Q}$ also contains only d roots of $p(X)$, so all of them, including x, must then lie in $g(F)$.

This proves uniqueness of the image $g(F)$. Of course, there may still be many homomorphisms $g : F \hookrightarrow \overline{Q}$. Rabin proved in [17] that one of them (say h) must

be computable, and so $g(F) = h(F) \leq_T T_F^*$. Conversely, given any polynomial $p(X) \in Q[X]$, we can compute from a $g(F)$-oracle whether p has any roots in F, just by finding a root of p in $\overline{Q}$ and checking whether it lies in $g(F)$. (By normality, it suffices to check this for just one root of p.) So $T_F^* \leq_T g(F)$, and $\deg(T_F^*)$ is the unique degree in $\mathrm{DgSp}_{\overline{Q}}(F)$. □

If F is algebraic but not normal over Q, then F does have more than one possible homomorphic image in $\overline{Q}$. For us, a crucial distinction will be whether F is *almost normal* over Q.

Definition 1. *An algebraic field extension $L \subseteq F$ is* almost normal *if there is a finite field extension $L \subseteq E$ such that $E \subseteq F$ is a normal extension.*

Corollary 4. *If F is an almost normal field extension of the prime field Q, then $DgSp_{\overline{Q}}(F)$ is a singleton.*

Proof. Let E be a finite extension of Q over which F is normal. Then by the Proposition and the subsequent remark, the image $g(F)$ of any embedding $g : F \to \overline{P}$ is determined by $g(E)$, and specifically by the values of g on a finite set B generating E. Now for any two such embeddings g_0 and g_1, we need only finitely much information – namely $g_0 \restriction B$ and $g_1 \restriction B$ – to compute an automorphism of $\overline{Q}$ which maps $g_0(b)$ to $g_1(b)$ for each $b \in B$. This automorphism must then map $g_0(F)$ onto $g_1(F)$, and since it is computable, $g_0(F) \equiv_T g_1(F)$. □

Now we are ready to classify all spectra of algebraic subfields.

Theorem 4. *Let F be any algebraic field, with prime subfield Q, and let V_F be the set defined in Section 2. If F is almost normal over Q, then $DgSp_{\overline{Q}}(F) = \{deg(V_F)\}$. If not, then $DgSp_{\overline{Q}}(F) = \{\boldsymbol{d} : deg(V_F) \leq_T \boldsymbol{d}\}$, the upper cone of degrees above the degree of V_F.*

Proof. First we need a lemma.

Lemma 2. *For every algebraic field F, every subfield $\tilde{F} \subseteq \overline{Q}$ isomorphic to F computes V_F. Thus $DgSp_{\overline{Q}}(F)$ is contained in the upper cone above $deg(V_F)$.*

Proof. For any i, we can find all roots of $p_i(X)$ in $\overline{Q}$ and check how many of them lie in $\tilde{F}$. This computes T_F. Likewise, given any pair $\langle g_0(X), g_1(X,Y)\rangle$, we can find all roots of g_0 in $\overline{Q}$, check which ones (say $x_0, \ldots, x_k$) lie in $\tilde{F}$, and then check whether $\tilde{F}$ contains any roots of $g_1(x_j, Y)$, for each $j \leq k$. This computes U_F, so $V_F \leq_T \tilde{F}$. □

By Theorem 2, we may take F itself to be a presentation with $F \equiv_T V_F$. With a V_F-oracle, therefore, we may compute an embedding $g : F \to \overline{Q}$. But then the image $g(F)$ is also V_F-computable, since for any $x \in \overline{Q}$ we may find its minimal polynomial $p(X) \in Q[X]$, determine from V_F (really from T_F) the number of roots of $p(X)$ in F, find that many roots of $p(X)$ in F, and determine whether g maps any of them to x. Thus $\deg(g(F)) \in \mathrm{DgSp}_{\overline{Q}}(F)$ and $g(F) \leq_T V_F$, so by Lemma 2 $g(F) \equiv_T V_F$, putting $\deg(V_F) \in \mathrm{DgSp}_{\overline{Q}}(F)$.

If F is almost normal over Q, then Corollary 4 now proves the desired result. So we may assume that F is a non-normal extension of every finitely generated subfield of F. To complete our proof, we must show that in this case every degree $\boldsymbol{d}$ which computes V_F lies in $\mathrm{DgSp}_{\overline{Q}}(F)$.

So suppose that $D \in \boldsymbol{d}$ and $V_F \leq_T D$. The following construction is computable in a D-oracle. Let $F_0 = Q$ and let g_0 be the unique embedding of F_0 into $\overline{Q}$, and set $i_0 = -1$. We will build a computable increasing sequence $i_0 < i_1 < \cdots$ and an embedding $g = \cup_s g_s$ with $F_s = \mathrm{dom}(g_s) \subseteq F$ being the subfield of F generated by all roots in F of all polynomials $p_j(X)$ with $j \leq i_s$. Thus $\cup_s F_s$ will equal F. Moreover, the image $g_s(F_s)$ will be defined so as to code $D\restriction(s+1)$, with this coding being respected at subsequent stages.

It is important to note, in the following description of the procedure at stage $s+1$, that the first part of the procedure, creating the coding opportunity, requires only a V_F-oracle, along with knowledge of g_s. (Of course, our D-oracle computes V_F.) Given g_s and i_s and the finite algebraic extension $F_s = \mathrm{dom}(g_s) \supseteq Q$ within F, we search through all roots of $p_{i_s+1}(X)$ in F, then all roots of $p_{i_s+2}(X)$, etc., until we find an $x \in F$ which is a root of some $p_i(X)$ with $i > i_s$, such that the minimal polynomial $q(X) \in F_s[X]$ of x does not have $\deg(q)$-many roots in F. Eventually we must find such a root, since F is not almost normal over Q, and Sublemma 1 below shows that with our V_F-oracle, we can recognize the root when we find it. For the least $i > i_s$ for which $p_i(X)$ has such a root, we let $i_{s+1} = i$ and let $q_s(X)$ be the minimal polynomial in $F_s[X]$ of that root. This completes the first part of the procedure.

Now we use our full D-oracle to take the coding step. Let r_0 be the least root of $q_s(X)$ in $\overline{Q}$ (under the order $<$ on the domain ω of $\overline{Q}$). If $s \in D$, then define g'_s to extend g_s by mapping the least root of $q_s(X)$ in F to r_0. If $s \notin D$, then we wish to ensure that r_0 does not lie in the image $g_{s+1}(F_{s+1})$. Let $\overline{q}_s(X) \in \overline{Q}[X]$ be the image of $q_s(X)$ under the map g_s on its coefficients. Let F'_s be the subfield of F generated by F_s and all roots of $q_s(X)$ in F. Of course, all these generators of F'_s must be mapped to roots in $\overline{Q}$ of $\overline{q}_s$, of which there are only finitely many. So we check through these finitely many possible maps until we find a map g'_s, which extends to a field embedding of F'_s into $\overline{Q}$ and satisfies $r_0 \notin \mathrm{range}(g'_s)$. To see that such an g'_s must exist, let $K_s \subset \overline{Q}$ be the splitting field of $\overline{q}_s(X)$ over $g_s(F_s)$. By irreducibility of $q_s(X)$, the Galois group $\mathrm{Gal}(K_s/g_s(F_s))$ acts transitively on the roots of $\overline{q}_s$ in K_s. We know that there exists an embedding of F'_s into K_s, and that embedding must omit at least one root r_1 of $\overline{q}_s$ from its image, since F does not contain all $\deg(q_s)$-many possible roots of q_s. By transitivity, there is an element of $\mathrm{Gal}(K_s/g_s(F_s))$ mapping r_1 to r_0, and the composition of this element with the given embedding omits r_0 from its image. So when we search, we will find the desired extension g'_s.

The coding step is now finished. To complete stage $s+1$ (in both cases $s \in D$ and $s \notin D$), we now extend this g'_s to the remaining roots of $p_{i_{s+1}}$ in F and to all roots in F of all polynomials $p_j(X)$ with $i_s < j < i_{s+1}$. This can be done in any systematic way, and we let g_{s+1} be this extension, so g_{s+1} extends g_s to the subfield F_{s+1} generated by all roots of all $p_j(X)$ with $j \leq i_{s+1}$. This completes

stage $s+1$. Notice that if $s \notin D$, then r_0 remains outside the image of g_{s+1}, since all roots of $q_s(X)$ in F are mapped to elements $\neq r_0$. Indeed, the same reasoning shows that the embedding $\cup_s g_s$ constructed through all these stages will have r_0 in its image iff $s \in D$.

Now the union $g = \cup_s g_s$ is an embedding of F into $\overline{Q}$, computable in D. We define $\tilde{F} = g(F)$ to be its image, so that $\deg(\tilde{F}) \in \mathrm{DgSp}_{\overline{Q}}(F)$, and we claim that $\tilde{F} \equiv_T D$. First, the entire procedure is D-computable. To decide from a D-oracle whether an arbitrary $x \in \overline{Q}$ lies in $\tilde{F}$, just find the unique i for which $p_i(x) = 0$. Then use D to compute V_F and find all roots of $p_i(X)$ in F, and check whether g maps any of them to x. Thus $\tilde{F} \leq_T D$.

Conversely, our coding allows us to run the entire construction and thus compute g (and hence D) from an $\tilde{F}$-oracle, as follows. By Lemma 2, $V_F \leq_T \tilde{F}$, so given g_s, we can run the first part of stage $s+1$, defining $q_s(X)$ and setting up the coding. Then we find the least root r_0 of $q_s(X)$ in $\overline{Q}$, and check whether it lies in $\tilde{F}$. If not, then by the construction $s \notin D$, and with this knowledge we can run the remainder of the procedure for stage $s+1$ and compute g_{s+1}. Conversely, as argued at the end of the construction, r_0 can only lie in $\tilde{F}$ if it appeared there at stage $s+1$ under the steps followed when $s \in D$. So, if $r_0 \in \tilde{F}$, then we know that $s \in D$, and again this knowledge allows us to run the remainder of the procedure for stage $s+1$ and compute g_{s+1}. The real point is to compute D, of course, and it is now clear that we can decide by this process, for any s, whether $s \in D$ or not. So $D \leq_T \tilde{F}$, and thus $\boldsymbol{d} = \deg(D) \in \mathrm{DgSp}_{\overline{Q}}(F)$. (Computing g_{s+1} below $\tilde{F}$ is really the inductive step which allows our computation of D to continue through as many stages as necessary.)

It remains to prove the sublemma we used.

Sublemma 1 *With a V_F-oracle and our fixed presentation of F we can determine, uniformly in i and s, the irreducible factors $q(X)$ of $p_i(X)$ in $F_s[X]$, and check which ones have their full complement of $\deg(q)$-many roots in F.*

Proof. From a V_F-oracle we know a finite set generating F_s over Q, so we may find a primitive generator x of F_s and its minimal polynomial $g_0(X) \in Q[X]$. Also, we have a splitting algorithm for F_s, which allows us to find the irreducible factors of $p_i(X)$ in $F_s[X]$. If $q(X)$ is one of these, then we may compute its splitting field in $\overline{Q}$ over the image $g_s(F_s)$. We find a primitive generator $z \in \overline{Q}$ of this splitting field and determine the minimal polynomial $p_k(X) \in Q[X]$ of z. Then we use T_F (which is part of V_F) to find all roots $z_1, \ldots, z_n \in F$ of $p_k(X)$. However, an individual $Q[z_m] \subseteq F$ may only be conjugate to the splitting field of $q(X)$ over F_s, rather than being an actual splitting field. So, for each root z_m, we check whether the subfield $Q[z_m]$ of F contains all generators of F_s (hence contains the coefficients of $q(X)$) and also contains a full complement of roots of $q(X)$. This is easy: the splitting set of $Q[z_m]$ is computable in F, by Theorem 1. If there is an m for which this holds, then we have our answer; if not, then F cannot contain a full complement of roots of $q(X)$, since such a set of roots, along with the generators of F_s, would generate a subfield containing some root of $p_k(X)$.

In this computation, the use of g_s is not necessary; we could compute a splitting field of $q(X)$ over F_s without knowing g_s. Indeed, it would be sufficient to have as oracles T_F and a presentation of F. (From these we can find a finite generating set for F_s, using the definition of F_s in the construction. If F_s were replaced by an arbitrary subfield of F, we would need to know a finite set of generators for it.) Alternatively, a V_F-oracle would suffice, since from it we could compute both T_F and a presentation of F. □

This completes the proof of Theorem 4. □

5 Conclusions

The most obvious theorem relevant to the results of Sections 2 and 4 is Rabin's Theorem, from [17]; see also [14].

Theorem 5 (Rabin). *Let F be any computable field.*

1. *There exists a Rabin embedding g of F into a computable algebraically closed field $\overline{F}$. (This means that g is a computable field homomorphism and $\overline{F}$ is algebraic over the image $g(F)$.)*
2. *For every Rabin embedding g of F into any computable ACF E, the image of g is a computable subset of E iff F has a splitting algorithm.*

A relativized version states that the image $g(E)$ is Turing-equivalent to the *splitting set* of F, i.e. the set of reducible polynomials in $F[X]$. Details are available in [15].

We view Rabin's Theorem as saying that Σ_1 questions, such as the reducibility of polynomials over F, become computable if one is given F as a subfield of $\overline{F}$, rather than simply as a freestanding field. This phenomenon is specific to algebraic fields. For example, in the language of trees with an immediate-predecessor function, all computable trees of height $\leq \omega$ embed into the computable tree $\omega^{<\omega}$, yet existential questions about a computable tree T do not become computable when one is given the ability to compute a subtree of $\omega^{<\omega}$ isomorphic to T. Indeed, in a field of infinite transcendence degree over its prime subfield Q, being algebraic over Q is a Σ_1 property which need not become computable just because one can compute the image of the field in a computable algebraic closure. For fields, algebraicity is the key: one knows exactly how many roots $p(X) \in F[X]$ has (counted by multiplicity) in its algebraic closure $\overline{F}$, and so, if one can compute the image of F within $\overline{F}$, one can find them all and check how many lie in the image of F. We conjecture that similar results may hold for other algebraic structures (in the model-theoretic sense of *algebraic*), but that some further constraints on the structure are necessary, such as the ability to determine the maximum number of possible realizations of a type. (Local finiteness and finite axiomatizability of the theory, in a finite signature, would likely suffice for this.)

The relation of these matters to the current work is that in our results again, computable enumerability in fields as freestanding structures converts to computability when we view the fields as subfields of their algebraic closures. The

spectrum of the field F is defined by a Σ_1 condition: the ability of $\boldsymbol{d}$ to *enumerate* a given set V_F. The spectrum of F as a subfield of $\overline{F}$, on the other hand, is defined by the ability of $\boldsymbol{d}$ to *compute* V_F (and, for almost normal fields, the ability of V_F to compute $\boldsymbol{d}$). Of course, this makes it harder for a degree to present F as a subfield of a computable copy of $\overline{F}$ than it is to present F as a field; conversely, a presentation of F as subfield of $\overline{F}$ gives us more power than a mere presentation of F as a field.

References

1. W. Calvert, V. Harizanov & A. Shlapentokh; Turing degrees of isomorphism types of algebraic objects, *Journal of the London Math. Soc.* **73** (2007) 273-286.
2. R.J. Coles, R.G. Downey & T.A. Slaman; Every set has a least jump enumeration, *Journal of the London Mathematical Society* **62** (2000) 2 641-649.
3. H.M. Edwards; *Galois Theory* (New York: Springer-Verlag, 1984).
4. Yu.L. Ershov; Theorie der Numerierungen III, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **23** (1977) 4, 289-371.
5. Yu.L. Ershov & S.S. Goncharov; *Constructive Models* (New York: Kluwer Academic/Plenum Press, 2000).
6. M.D. Fried & M. Jarden, *Field Arithmetic* (Berlin: Springer-Verlag, 1986).
7. A. Frohlich & J.C. Shepherdson; Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407-432.
8. A. Frolov, I. Kalimullin, & R.G. Miller; Spectra of algebraic fields and subfields. Available at `qcpages.qc.cuny.edu/~rmiller/CiE09.pdf`, with Appendix A included.
9. V. Harizanov & R. Miller; Spectra of structures and relations, *Journal of Symbolic Logic* **72** (2007) 1, 324-347.
10. N. Jacobson; *Basic Algebra I* (New York: W.H. Freeman & Co., 1985).
11. J.F. Knight; Degrees coded in jumps of orderings, *Journal of Symbolic Logic* **51** (1986), 1034-1042.
12. L. Kronecker; Grundzüge einer arithmetischen Theorie der algebraischen Größen, *J. f. Math.* **92** (1882), 1-122.
13. G. Metakides & A. Nerode; Effective content of field theory, *Annals of Mathematical Logic* **17** (1979), 289-320.
14. R.G. Miller; Computable fields and Galois theory, *Notices of the AMS* **55** (August 2008) 7, 798-807.
15. R.G. Miller; Is it harder to factor a polynomial or to find a root?, to appear in the *Transactions of the American Mathematical Society*.
16. R.G. Miller; $\boldsymbol{d}$-Computable categoricity for algebraic fields, to appear in *The Journal of Symbolic Logic*.
17. M. Rabin; Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341-360.
18. L.J. Richter; Degrees of structures, *Journal of Symbolic Logic* **46** (1981), 723-731.
19. R.I. Soare; *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).
20. V. Stoltenberg-Hansen & J.V. Tucker; Computable Rings and Fields, in *Handbook of Computability Theory*, ed. E.R. Griffor (Amsterdam: Elsevier, 1999), 363-447.
21. B.L. van der Waerden; *Algebra*, volume I, trans. F. Blum & J.R. Schulenberger (New York: Springer-Verlag, 1970 hardcover, 2003 softcover).

Appendix A: Proof of Theorem 2

We first consider algebraic fields which are normal extensions of the prime field. The result for such fields is the same as the general result, but the proof is easier and introduces our technique. The set T^*_F defined here is a simpler version of the set T_F to be defined below for all algebraic fields.

Proposition 2. *For every algebraic field F which is normal over its prime subfield Q, there exists a set $T^*_F \subseteq \omega$, such that*

$$Spec(F) = \{\boldsymbol{d} : T^*_F \text{ is c.e. in } \boldsymbol{d}\}.$$

Conversely, for every $S \subseteq \omega$, there exists a normal algebraic field F of arbitrary characteristic such that

$$Spec(F) = \{\boldsymbol{d} : S \text{ is c.e. in } \boldsymbol{d}\}.$$

Finally, when we view T^*_F as a path through $2^{<\omega}$ and fix any characteristic c, the set $\{T^*_F \restriction n : n \in \omega \ \&\ F \text{ normal algebraic of characteristic } c\}$ forms a computable extendible subtree $T \subset 2^{<\omega}$ with no isolated paths, and every path through T is of the form T^*_F for some such F.

Proof. Since the prime field Q has a splitting algorithm, we may fix an enumeration $p_0(X), p_1(X), \ldots$ of the irreducible monic polynomials in $Q[X]$, and choose any normal algebraic extension field $F \supseteq Q$. Let T^*_F be the set $\{i : \exists a \in F(p_i(a) = 0)\}$. Notice that half of the first result is immediate: T^*_F is computably enumerable in any presentation of F, since we may enumerate the subfield of rational numbers in any such presentation.

For the reverse inclusion, let $\boldsymbol{d}$ be a degree such that T^*_F is c.e. in $\boldsymbol{d}$. We will construct a field $\tilde{F}$ isomorphic to F, using a $\boldsymbol{d}$-oracle. This will suffice, since it was shown by Knight in [11] that spectra of infinite fields are upwards closed under Turing reducibility. Indeed, the construction of $\tilde{F}$ is straightforward. Fix a computable copy $\overline{Q}$ of the algebraic closure of Q and a $\boldsymbol{d}$-computable enumeration of T^*_F, and start with $\tilde{F}_0 = Q$. At stage $s+1$, exactly one new number i appears in T^*_F, and we find all $\deg(p_i)$-many roots of $p_i(X)$ in $\overline{Q}$. By induction, $\tilde{F}_s$ is normal over Q, so either these roots all lie in $\tilde{F}_s$ already, in which case we set $\tilde{F}_{s+1} = \tilde{F}_s$; or else none of these roots is in $\tilde{F}_s$, in which case we build $\tilde{F}_{s+1}$ by adjoining all these roots to $\tilde{F}_s$. (We have root algorithms uniformly for all finite algebraic extensions of Q, hence for all $\tilde{F}_s$.) The field $\tilde{F} = \cup_s \tilde{F}_s$ is clearly $\boldsymbol{d}$-computable and normal over Q, hence has only one possible homomorphic image in $\overline{Q}$, as shown in Proposition 1. Likewise, F embeds into $\overline{Q}$ with that same image, giving an isomorphism from F onto $\tilde{F}$.

There exist nodes in $2^{<\omega}$ which lie on no path T^*_F, as F ranges over computable algebraic fields. For instance, the irreducible polynomial $p_i(X) = X^2 - 2X - 1$ has roots $(1 \pm \sqrt{2})$, which lie in F iff the irreducible polynomial $p_j(X) = X^2 - 2$ has roots in F. Thus no F could have $T^*_F(i) \neq T^*_F(j)$. To prove the final claim of Proposition 2, therefore, we consider the tree T it describes. (The second result in the proposition will be proven afterwards.)

T is the computable subtree of $2^{<\omega}$ defined as follows. The empty node λ lies in T, with $E_\lambda = Q$, and for each $\sigma \in T \cap 2^n$, we define the field E_σ to be the splitting field over Q of the product polynomial

$$\Pi\{p_i(X) : \sigma(i) = 1\}.$$

Proceeding recursively, for each $\sigma \in T \cap 2^n$, we define $\sigma\hat{}0 \in T$ iff $p_n(X)$ has no roots in E_σ; and $\sigma\hat{}1 \in T$ iff the splitting field of $p_n(X)$ over E_σ contains no roots of any $p_i(X)$ for which $\sigma(i) = 0$. The intuition is that E_σ is the minimal normal extension of Q with $T^*_{E_\sigma}\restriction n = \sigma$, and we keep $\sigma\hat{}i$ out of T iff there is no such normal extension. An easy induction on σ shows that this intuition holds, i.e. $T^*_{E_\sigma}\restriction n = \sigma$ for all $\sigma \in T$. Thus every $\sigma \in T$ is extendible.

Moreover, for any $\sigma \in T$, fix a prime q which is $> \deg(p_i)$ for all $i < |\sigma|$. and find a polynomial $p(X) \in Q[X]$ whose Galois group over Q is cyclic of order q. (Classical results show that such a polynomial exists.) Then $p(X)$ must have degree q, hence is not among $p_0(X), \ldots, p_{|\sigma|}(X)$, and its splitting field has prime degree q over Q, hence is generated by any root of $p(X)$. The extension E of E_σ by a single root of $p(X)$ also is normal of degree q over E_σ, since we chose q prime to $[E_\sigma : Q]$, and therefore $T^*_E\restriction |\sigma| = \sigma$, since no root of any $p_i(X)$ can generate an extension of degree q. But T^*_E branches from $T^*_{E_\sigma}$ at level k, where $p_k(X) = p(X)$. Thus there are no isolated paths.

The construction shows that for any path h through T, we have a field $F = \cup_n E_{h\restriction n} \subseteq \overline{Q}$ such that $h = T^*_F$. Conversely, the nodes of T are defined so that no path T^*_F could possibly contain any node not in T. So the set $\{T^*_F \in 2^\omega : F \text{ is a subfield of } \overline{Q}\}$ is precisely the set of paths through the subtree T of $2^{<\omega}$.

Finally, we prove the converse statement in the proposition. Given a set $S \subseteq \omega$, it is tempting simply to use S to define a path h through T: each time h reaches a branch point in T, just take the next value of S to determine which way h should go. Unfortunately, the path thus constructed does not lend itself to an easy proof of the result, because different paths through T have branch points at different levels, so that S would only be Δ_2 in an arbitrary copy of the field defined by p. (To see the dificulty, imagine the example in which $p_0(X) = X^2-2$, $p_1(X) = X^2 - 3$, and $p_2(X) = X^2 - 6$, and examine the first four levels of the corresponding tree T.)

Instead, we focus on specific polynomials from our list. Compute $i_0 < i_1 < \cdots$ such that the polynomials $p_{i_j}(X)$ are precisely the cyclic polynomials $X^{m-1} + X^{m-2} + \cdots + X + 1$, where m is a positive prime number. Define the path h by

$$h(n) = \begin{cases} S(j), \text{ if } n = i_j \\ 0, \quad \text{if } n \notin \{i_j : j \in \omega\} \ \&\ (p\restriction n)\hat{}0 \in T \\ 1, \quad \text{otherwise.} \end{cases}$$

This h is a path through T, because, for any field F generated by roots of any subset of $\{p_{i_j} : j \in \omega\}$ the existence of a root of p_{i_k} in F is independent of the set $\{j : j \neq k \ \&\ p_{i_j} \text{ has a root in } F\}$. (The splitting field of the polynomial $p_{i_j}(X)$ has degree m over Q, for the corresponding prime m, and so the splitting field of

any one p_{i_j} intersects the field generated by the splitting fields of all the others only in Q.)

Let F be the field extension of Q generated by the roots of the polynomials $p_{i_j}(X)$ for which $S(j) = 1$. It is clear that S is enumerable computably in any presentation of F, and that conversely, if S is $\boldsymbol{d}$-computably enumerable, then we can build a presentation of F computable in $\boldsymbol{d}$. This completes the proof of Proposition 2. □

The general case of algebraic fields is more complicated. We need to know not just whether an irreducible $p(X) \in Q[X]$ has roots in F, but how many roots, and how they interact with each other, and which irreducible polynomials they satisfy over elements added earlier to the field (since $p(X)$ may become reducible once we start adjoining other algebraic numbers to Q). The result, however, is the same.

Theorem 2. *For every algebraic field F, there exists a set $V_F \in 2^\omega$ with*

$$\mathit{Spec}(F) = \{\boldsymbol{d} : V_F \text{ is c.e. in } \boldsymbol{d}\}.$$

Conversely, for every $V \in 2^\omega$, there exists an algebraic field F of arbitrary characteristic such that

$$\mathit{Spec}(F) = \{\boldsymbol{d} : V \text{ is c.e. in } \boldsymbol{d}\}.$$

Finally, if P is the prime field of any fixed characteristic c, then the set $\{V_F \upharpoonright n \ : \ F \text{ algebraic over } P \ \&\ n \in \omega\}$ forms a computable extendible subtree $T \subset 2^{<\omega}$ with no isolated paths, and every path through T is of the form V_F for some F.

Proof. As in Proposition 2, our proof in characteristic 0 will work equally well in all positive characteristics. Also, Proposition 2 already proves the second of the three statements here. We repeat that the first statement is essentially just a relativization of a theorem of Ershov (see [4, §12, Thm. 2]).

Recall the definition of the c.e. set V_F, which was given in Section 2. First we define T_F, following the definition of T_F^* in Proposition 2, but allowing for the possibility of the field being non-normal. We use the same list $p_0(X), p_1(X), \ldots$ of irreducible polynomials in $Q[X]$ as before, and set $d_i = \deg(p_i(X))$ and $D_i = \sum_{j<i} d_j$. Then, for each i and each $k < d_i$, we define

$$D_i + k \in T_F \iff F \text{ contains at least } (k+1) \text{ distinct roots of } p_i(X).$$

So the first d_0 bits of T_F tell the exact number of roots of $p_0(X)$ in F, and the next d_1 bits tell the number of roots of $p_1(X)$, etc. Obviously T_F is computably enumerable in the degree of any field isomorphic to F.

To define U_F, we need to consider pairs of polynomials. Since we have a splitting algorithm for Q, it is computable whether a pair $\langle g_0(X), g_1(X,Y)\rangle \in (Q[X] \times Q[X,Y])$ satisfies both of the following conditions.

- $Q[X]/(g_0)$ is a field. This is equivalent to demanding that g_0 be nonconstant and irreducible in $Q[X]$.
- g_1, when viewed as a polynomial in Y, is irreducible in the polynomial ring $(Q[X]/(g_0))[Y]$. (The coefficients of Y in g_1 are really in $Q[X]$, of course; here we consider their images in $Q[X]/(g_0)$.) Equivalently, if x is a root of g_0, then $g_1(x, Y)$ is irreducible in the polynomial ring $Q(x)[Y]$.

So we may computably enumerate a list

$$\boldsymbol{G}_0, \boldsymbol{G}_1, \ldots$$

of all pairs satisfying these properties, writing $\boldsymbol{G}_j = \langle g_{j0}, g_{j1} \rangle$. Now define

$$U_F = \{j : (\exists x, y \in F)[g_{j0}(x) = 0 = g_{j1}(x, y)]\}.$$

Again this set is c.e. in the degree of any field isomorphic to F.

Now define $V_F = T_F \oplus U_F$. As remarked above, it is clear that $\mathrm{Spec}(F) \subseteq \{\boldsymbol{d} : V_F \text{ is c.e. in } \boldsymbol{d}\}$. To prove the reverse inclusion, suppose that $\boldsymbol{d}$ is a degree in which V_F is c.e. We will construct a field $\tilde{F} \cong F$ computably in $\boldsymbol{d}$. To begin, let $\tilde{F}_0 = Q$.

Given $\tilde{F}_s$, we continue enumerating T_F, using our $\boldsymbol{d}$-oracle. When the next new element t appears in T_F, it satisfies $D_i \leq t < D_{i+1}$ for some i, so $t = D_i + k$ for some $k < d_i$ and the enumeration shows that F contains at least $(k + 1)$ roots of $p_i(X)$. Since $\tilde{F}_s$ is (by induction) a finite extension of Q, we may determine exactly how many roots of $p_i(X)$ it already contains. If at least k distinct elements $y \in \tilde{F}_s$ satisfy $p_i(y) = 0$, then we do nothing. If not, then we search for a primitive generator x of $\tilde{F}_s$. This x must exist, by the Theorem of the Primitive Element, and we will eventually recognize one, by seeing that it generates all of the finitely many generators of $\tilde{F}_s$. Then we find the minimal polynomial $g_0(X)$ of this x in $Q[X]$. Now enumerate U_F until we find some $j \in U_F$ such that:

- $g_{j0} = g_0$, so that $Q[X]/(g_0) \cong \tilde{F}_s$; and
- there exists $h(Y) \in (Q[X]/(g_0))[Y]$ with $g_{j1}(Y) \cdot h(Y) = p_i(Y)$.

Let $\tilde{F}_{s+1} = (\tilde{F}_s[Y])/(g_{j1}(Y))$; that is, extend $\tilde{F}_s$ by adjoining a root of g_{j1}. This completes stage $s + 1$, and we set $\tilde{F} = \cup_s \tilde{F}_s$, so $\tilde{F}$ is indeed $\boldsymbol{d}$-computable.

We claim that in fact $\tilde{F} \cong F$; this will show that $\boldsymbol{d} \in \mathrm{Spec}(F)$ (since spectra of structures are closed upwards under $\leq_T$), and will thus complete the proof of the first part of Theorem 2. The main difficulty is to build an embedding f of $\tilde{F}$ into F. Of course, f need not be computable, and in general it will not be. We will define a sequence of embeddings $f_s : \tilde{F}_s \to F$, and construct f from this sequence.

To start, f_0 is the unique embedding of $\tilde{F}_0 = Q$ into F. Recall that $\tilde{F}_{s+1} = \tilde{F}_s[y_s]$ for some $y_s \in \tilde{F}_{s+1}$, with the minimal polynomial $g(Y)$ used in the construction of $\tilde{F}_{s+1}$. Recall that there was some $j \in U_F$ with $g = g_{j1}$ and with $Q[X]/(g_{j0}) \cong \tilde{F}_s$. Since $j \in U_F$, we know that F contains elements x and y satisfying $g_{j0}(x) = 0 = g_{j1}(x, y)$. So we map $\tilde{F}_s \cong Q[X]/(g_{j0})$ into F by sending X

to x, and then extend this map by sending y_s to y. This defines our embedding $f_{s+1} : \tilde{F}_{s+1} \to F$.

The reader is likely surprised that f_{s+1} was built without any use of f_s, but we are able to proceed nevertheless. $\tilde{F}$ is generated over Q by the elements $y_0, y_1, \ldots$ chosen above. There is only one possible definition of f on $\tilde{F}_0 = Q$, so we fix that definition, and let $s^0_n = n$ for all n.

Once we have defined f on $\tilde{F}_t$, we proceed recursively to $\tilde{F}_{t+1} = \tilde{F}_t[y_t]$. Since y_t is algebraic over Q, the set $\{f_{s^t_n}(y_0) : n \in \omega\}$ is finite: all its elements must be roots of the minimal polynomial of y_0. So the sequence $s^t_0 < s^t_1 < \cdots$ contains an infinite subsequence $s^{t+1}_0 < s^{t+1}_1 < \cdots$ of stages such that $f_{s^{t+1}_n}(y_0)$ is the same for all n. We define $f(y_t)$ to have this value, and extend f to all of $\tilde{F}_{t+1} = \tilde{F}_t[y_t]$. So, restricted to the domain $\tilde{F}_{t+1}$, f is equal to each $f_{s^{t+1}_n}$, allowing the recursion to proceed.

This defines f on all of $\tilde{F}$, and since its restriction to $\tilde{F}_t$ is a field embedding for every t, it is clear that f itself embeds $\tilde{F}$ into F. Now every element x of F has minimal polynomial $p_i(X)$ in $Q[X]$ for some i, since F is algebraic, and $p_i(X)$ has a certain number k of roots in F. Therefore $D_i + k - 1 \in T_F$, and so our construction ensured that $\tilde{F}$ also contains at least k distinct roots of $p_i(X)$. Since f is an embedding, these roots must map to distinct roots of $p_i(X)$ in F, and hence one of them maps to x. So the embedding f is onto, and $\tilde{F} \cong F$.

It remains to show the final clam of Theorem 2. As in the case of normal fields, certain finite binary sequences cannot be initial segments of V_F for any field F. However, we can compute which ones they are. Fix any $\sigma \in 2^{<\omega}$. Each i with $\sigma(i) = 1$ requires that one element (for even i, corresponding to T_F) or finitely many elements (for odd i, i.e. for U_F) be present in the field it describes, and if they are not already there, we adjoin them in every way possible. Thus we build finitely many minimal subfields $F^1_\sigma, \ldots, F^{j_\sigma}_\sigma$ of $\overline{Q}$ satisfying all the criteria required by these i. (Indeed, these criteria imply that each F^k_σ embeds into every other one, and minimality then shows that these subfields are all isomorphic to each other. So we may refer to the isomorphism type as F_σ.) Next we consider those i with $\sigma(i) = 0$. Since F_σ is a finite extension of Q, Theorem 1 allows us to compute its root set, and hence to decide whether F_σ satisfies the condition given by $\sigma(i) = 0$. (For the condition of not satisfying $\boldsymbol{G}_j = \langle g_{j0}, g_{j1} \rangle$, we may have to find all roots x of g_{j0} in F_σ and then check, for each of them, whether $g_{j1}(x, Y)$ has any roots in F_σ; but this is still computable.) The subtree T described in Theorem 2 contains precisely those σ such that F_σ satisfies these conditions for all i with $\sigma(i) = 0$. Thus T is computable.

It is clear that for all algebraic fields F, V_F is a path through T. For any $\sigma \in T$, the field F_σ itself satisfies $\sigma \subset V_{F_\sigma}$, and so σ is extendible. Moreover, T has no isolated paths, since it is easy to repeat the process from the proof of Proposition 2 and define a second field F' with $\sigma \subset V_{F'}$ and $V_{F'} \neq V_{F_\sigma}$.

Finally, consider any path h through T. For each n in turn, set $\sigma = h{\restriction}n$; we know $\sigma \in T$, so σ satisfies the conditions above for membership in T, allowing us to embed $F_{\sigma{\restriction}(n-1)}$ into F_σ. The union $F = \cup_n F_{h{\restriction}n}$ along all these embeddings is a field F with $V_F = h$, and this completes the proof of Theorem 2. $\square$