TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 00, Number 0, Pages 000–000
S 0002-9947(XX)0000-0

IS IT HARDER TO FACTOR A POLYNOMIAL OR TO FIND A ROOT?

RUSSELL MILLER

Abstract. For a computable field F, the *splitting set* S is the set of polynomials $p(X) \in F[X]$ which factor over F, and the *root set* R is the set of polynomials with roots in F. Work by Frohlich and Shepherdson essentially showed these two sets to be Turing-equivalent, surprising many mathematicians, since it is not obvious how to compute S from R. We apply other standard reducibilities from computability theory, along with a healthy dose of Galois theory, to compare the complexity of these two sets. We show, in contrast to the Turing equivalence, that for algebraic fields the root set has slightly higher complexity: both are computably enumerable, and computable algebraic fields always have $S \leq_1 R$, but it is possible to make $R \not\leq_m S$. So the root set may be viewed as being more difficult than the splitting set to compute.

1. Introduction

Let F be any field, and $p(X)$ any polynomial with coefficients in F. Two basic questions can immediately be asked. First, does $p(X)$ factor in the polynomial ring $F[X]$? (We ignore constant factors, of course.) Second, does $p(X)$ have a root in the field F?

Mathematicians often have conflicting instincts about which of these questions is easier. Plugging an element r into $p(X)$ to check whether $p(r) = 0$ seems easier than multiplying together two polynomials $f(X)$ and $g(X)$ to check whether $f \cdot g = p$, and so a blind search for roots will go faster than a blind search for factorizations. On this basis, the second question seems "easier." On the other hand, for most fields F, more polynomials will have factorizations than will have roots (aside from the trivial case of linear polynomials). A search for an arbitrary factorization has more possible answers than a search for a factorization in which one factor is linear, and so positive answers are more "easily" found for the first question. Of course, this reaction is soon tempered by the realization that therefore, negative answers are more easily found for the second question, blurring one's instincts about which question is easier.

In a field F, these questions are essentially the inductive steps in two larger processes: (1) factoring $p(X)$ into its irreducible factors in $F[X]$; and (2) finding

Received by the editors September 25, 2008.
2000 *Mathematics Subject Classification.* Primary 12E05, 03D45; Secondary 03C57, 12L05.
The author was partially supported by Grant # 13397 from the Templeton Foundation, and by Grants # PSCREG-38-967 and 61467-00 39 from The City University of New York PSC-CUNY Research Award Program. Kevin Keating provided excellent references and background in Galois theory, and the anonymous referee made several useful suggestions.

all roots of $p(X)$ in F. A full factorization of $p(X)$ in $F[X]$ clearly determines the roots of $p(X)$ in F, whereas knowledge of all the roots of $p(X)$ is often insufficient to determine all its irreducible factors. Thus, most mathematicians eventually agree that factorizability is the more difficult question: if one knows how to factor arbitrary polynomials, one can readily determine which polynomials have roots, whereas the converse seems false. Of course, to determine the existence of roots of $p(X)$, one must check not only whether $p(X)$ can be factored, but also whether its factors can be factored, and so on down to linear factors. So one may have to ask about factorizability many times, for different polynomials, in order to get a single answer about roots.

In computability theory, several reducibilities are widely used to compare the complexity of sets of natural numbers. Turing-reducibility is the best known of these, but 1-reducibility, m-reducibility, several types of truth-table reducibility, and assorted other methods are also well known and often applied. Some of these are strictly finer than others: for instance, 1-reducibility implies m-reducibility, which implies Turing reducibility, but both converse implications fail. We give exact definitions below for certain of these notions, and then use them to compare the problem of factoring an arbitrary polynomial $p(X)$ within a computable field F with the problem of finding a root of $p(X)$ in F. It has been known since the work of Frohlich and Shepherdson in [5] that these two problems are Turing-equivalent, which was already a surprise to many mathematicians: it means that, knowing which polynomials in $F[X]$ have roots, we can determine whether an arbitrary polynomial is reducible in $F[X]$ (and conversely, as described above). In this paper we review these results and then go further, showing that for algebraic fields the two problems are not always 1-equivalent, nor even m-equivalent. However, they are comparable under 1-reducibility, and the further surprise is that the reduction contradicts the mathematician's intuition: the factorization problem for a computable algebraic field is 1-reducible to the root problem, but the root problem may fail even to be m-reducible to the factorization problem.

Definitions are necessary before going further. We will need an assortment of results, both standard and advanced, about Galois theory and hilbertian fields, and these will be stated and considered in Section 2. Right now we give rigorous definitions for our basic object, a computable field; for our basic problems of whether polynomials split into factors and/or have roots; and for the reducibilities we will apply to these problems.

Definition 1.1. A *computable field* F consists of two computable functions f and g from $\omega \times \omega$ into ω, such that ω forms a field under these functions, with f as the addition and g as the multiplication. We may also refer to this F as a *computable presentation* of the isomorphism type of F.

When dealing with positive characteristic, we usually allow ω to be replaced by any finite subset of itself, so as to allow finite fields in our definition.

A function is *computable* if it can be computed by a Turing machine – which is to say, according to a finite program of instructions. As always in logic, ω denotes the set of natural numbers, beginning with 0. (Hence every computable field is countable.) Here f and g must be *total* functions, i.e. on every input $\langle m, n\rangle \in \omega^2$, their programs must eventually halt and give the correct outputs $m + n$ and $m \cdot n$ under the field operations. This is obviously necessary in order for us to say that f and g "compute" the field addition and multiplication; we remark it mainly

because computability theory also considers *partial computable functions*, which are also computed according to finite programs, but which allow the domain to be a subset of the natural numbers, since the program may not halt on every input. The *Halting Problem* is the question of whether a given program will halt on a given input; since this is (famously) unsolvable, it is more natural to consider all programs and allow "computable functions" to be partial. The term *strictly partial* refers to partial functions which are not total.

For computable fields, [16] gives a good survey of results, and [9] is useful as a basic introduction. For more general questions about computability theory, we suggest [15], the canonical reference, and also [14], which both give good explanations of the reducibilities we will use in this paper.

Definition 1.2. Let A and B be subsets of ω.

(a) A is *m-reducible to B*, written $A \leq_m B$, if there exists a total computable function h such that

$$n \in A \iff h(n) \in B$$

for every $n \in \omega$. (This h is called an *m-reduction* of A to B.)

(b) Likewise, A is *1-reducible to B*, written $A \leq_1 B$, if the computable function h in part (a) may be taken to be one-to-one.

(c) A is *Turing-reducible to B*, written $A \leq_T B$, if there exists an oracle Turing machine Φ, as described below, such that

$$\Phi^B(n) = \chi_A(n)$$

for every $n \in \omega$, where χ_A is the characteristic function of A.

We write $A \equiv_m B$ to indicate that both $A \leq_m B$ and $B \leq_m A$; likewise for $A \equiv_1 B$ and $A \equiv_T B$. In these cases A and B are *m-equivalent,* *1-equivalent*, and *Turing-equivalent*, respectively. A *Turing degree* is an equivalence class of sets under $\equiv_T$.

The oracle Turing machine Φ^B also runs according to a finite program, but is allowed to ask questions of the form "is m in B?" and to execute different instructions depending on whether the answer is yes or no. We refer to B as the *oracle*, and think of this as saying that if we were able to compute membership in B, then we could use that ability to compute membership in A as well. Thus A is no harder than B to compute, assuming that we are allowed to ask as many questions as we like about membership of different numbers in B. On the other hand, m-reducibility restricts us to a single question about membership of one particular number $f(n)$ in B, and requires that answer to be the correct answer about membership of n in A as well. (As an example, the reader could consider whether a given A need be either Turing-reducible or m-reducible to its complement $\omega - A$.)

Occasionally we will give a function $f : \omega \to \omega$, rather than a subset $B \subseteq \omega$, as an oracle. When this is the case, we mean the oracle to be the graph of f, viewed as a subset of $\omega \times \omega$, under a bijective computable coding between $\omega \times \omega$ and ω. A moment's reflection should make this seem reasonable.

Definition 1.3. The *splitting set* S_F for a field F is the set of reducible polynomials in $F[X]$, i.e. products of two nonconstant factors there. The *root set* R_F of F is $\{p(X) \in F[X] : (\exists a \in F)p(a) = 0\}$.

Both these sets are *computably enumerable* (or *c.e.*), being defined by existential formulas. Each may therefore be viewed as the range of a computable enumeration, that is, of a total computable function.

Throughout the literature on computable fields, the phrase "F has a splitting algorithm" is used to mean that F has a *computable* splitting set: the characteristic function of S_F is computable. In this paper we will be concerned with the Turing degree of the splitting set, not just with its computability, so we follow [11] and use the new term to avoid conflict with the existing one. Likewise, F has a *root algorithm* if its root set is computable.

Notice also that in the traditional terminology "splitting algorithm," and in our adaptation, a polynomial is said to split in $F[X]$ is it has any proper factorization there; it is *not* necessary for it to split into linear factors in $F[X]$. Nevertheless, we will also retain the traditional meaning of the term *splitting field*: the smallest field over which the polynomial splits into linear factors. Minimal fields over which the polynomial is reducible will be called *symmetric subfields* of the splitting field, with Lemma 2.12 as justification.

With S_F as an oracle, one can decompose any polynomial $p(X) \in F[X]$ into its irreducible components in $F[X]$: if $p(X)$ is not itself irreducible, then we simply search through pairs of elements of $F[X]$ until we find a factorization, and continue by induction on the degree of $p(X)$. Likewise, with R_F as an oracle, one can find all roots in F of an arbitrary $p(X)$: if R_F indicates that $p(X)$ has a root in F, search through F for such a root r, and then continue inductively on the polynomial $\frac{p(X)}{X-r} \in F[X]$.

2. Known Results on Computable Fields

Any discussion of computable fields of characteristic 0 should begin with the question of a splitting algorithm for $\mathbb{Q}$.

It is not obvious that $\mathbb{Q}$ must have a splitting algorithm, but Kronecker provided one. It works for every computable presentation of $\mathbb{Q}$, since $\mathbb{Q}$ is a computably categorical field. In fact, Kronecker showed that every finitely generated extension of $\mathbb{Q}$ has a splitting algorithm, using the following theorem. Since the original paper dates to 1882, the reader may prefer to see the more recent version in [1], or Lemmas 17.3 and 17.5 of [2]. Part (c) is an obvious relativization of the proofs there.

Theorem 2.1 (Kronecker [7])**.** (a) *$\mathbb{Q}$ has a splitting algorithm.*

(b) *Let L be a c.e. subfield of a computable field K. If L has a splitting algorithm, then for any $x \in K$ transcendental over L, $L(x)$ also has a splitting algorithm. When $x \in K$ is algebraic over L, again $L(x)$ has a splitting algorithm, which requires knowledge of the minimal polynomial of x over L.*

(c) *More generally, for any c.e. subfield L of a computable field K and any $x \in K$ transcendental over L, the splitting set of $L(x)$ is Turing-equivalent to the splitting set for L, via reductions uniform in x. Also, if $x \in K$ is algebraic over L, $L(x)$ and L have Turing-equivalent splitting sets, uniformly in x and the minimal polynomial of x over L.*

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

Lemma 2.2. *For every computable field F algebraic over its prime subfield P, there is a computable function which accepts as input any finite tuple $\vec{x} = \langle x_1, \ldots, x_n \rangle$ 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 $\vec{x} = \langle x_1, \ldots, x_n \rangle$, using part (b) of Theorem 2.1. Since our F is algebraic over P, we may simply search for a polynomial $p(X)$ with root x_n and coefficients in $P[x_0, \ldots, x_{n-1}]$, and then factor it, using the splitting algorithm for $P[x_0, \ldots, x_{n-1}]$ (by inductive hypothesis), until we have found the minimal polynomial of x_n over $P[x_0, \ldots, x_{n-1}]$. □

These splitting algorithms also allow us to compute the Galois groups of the corresponding fields. A proof appears in [11].

Lemma 2.3. *Let $\overline{\mathbb{Q}}$ be any computable presentation of the algebraic closure of the field of rational numbers. There is an algorithm which accepts any finite tuple $\langle x_0, \ldots, x_n \rangle$ of elements of $\overline{\mathbb{Q}}$ and computes the automorphism group G of the field $F = \mathbb{Q}[\vec{x}]$ – that is, the Galois group of F over $\mathbb{Q}$. Specifically, the algorithm computes both the cardinality and the characteristic function of $\{\langle y_0, \ldots, y_n \rangle \in F^{n+1} : (\exists \sigma \in G)(\forall i \leq n)\sigma(x_i) = y_i\}$.*

Therefore, if $\mathbb{Q} \subseteq E \subseteq F$ are finite field extensions within $\overline{\mathbb{Q}}$, we can compute $Gal(F/E)$ uniformly in finite generating sets for E and F over $\mathbb{Q}$, by computing $Gal(F/\mathbb{Q})$ and checking which of its elements fix every generator of E.

We will also require Rabin's Theorem. To begin with, we give his name to the type of field embedding he considered.

Definition 2.4. Let F and E be computable fields. A function $g : F \to E$ is a *Rabin embedding* if:

- g is a homomorphism of fields; and
- E is both algebraically closed and algebraic over the image of g; and
- g is a computable function.

Theorem 2.5 (Rabin [13])**.** *Let F be any computable field.*

1. *There exists a computable algebraically closed field $\overline{F}$ with a Rabin embedding of F into $\overline{F}$.*
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.*

The following result had been proven by Frohlich and Shepherdson in [5] four years before Rabin's work, which provided a much simpler proof. (In fact, [5] and

[13] both consider only the Turing degree $\mathbf{0}$, but their proofs relativize to other degrees.) It shows that under Turing reducibility, the splitting set and the root set of a field have the same degree of complexity, thus giving one answer to the basic question of this paper. Our work in Sections 3 and 4 will examine these sets under the finer notion of 1-reducibility – largely because this corollary contradicts the instincts of many mathematicians!

Corollary 2.6. *For any computable field F, the following are Turing equivalent:*

(i.) *the image $g(F)$ of F under any Rabin embedding g;*
(ii.) *the splitting set S_F of F;*
(iii.) *the root set R_F of F;*
(iv.) *the root function of F, i.e. the function with domain $F[X]$ which computes the number of distinct roots in F of any $p(X) \in F[X]$;*
(v.) *the root multiplicity function of F, i.e. the function with domain $F[X]$ which computes the number of roots in F, counted by multiplicity, of any $p(X) \in F[X]$.*

Proof. (i) and (ii) are Turing equivalent by Rabin's Theorem, the proof of which easily relativizes to the splitting set, or to the image $g(F)$, when either is not computable. Using an S_F-oracle, we may readily check whether any irreducible factor of a given $p(X)$ is linear, thereby computing R_F. From R_F, we may determine whether a given $p(X)$ has a root and, if so, find such a root $r \in F$ and repeat the process for $\frac{p(X)}{X-r}$ until there are no more roots, thereby computing the root function. The root function and the root multiplicity function are quickly seen to compute each other. It is possible to compute the splitting set from the root function using symmetric polynomials, as shown in [5], but we give a direct computation of $g(F)$ instead, based on Rabin's proof of his theorem. Given a Rabin embedding $g : F \hookrightarrow E$ and any $x \in E$, find any polynomial $p(X) \in F[X]$ such that $\overline{p}(x) = 0$, where $\overline{p} \in E[X]$ is the image of p under the map g on its coefficents. With a root function for F, we may find all the roots $r_0, \ldots, r_n$ of p in F. Then $x \in g(F)$ iff $(\exists i \leq n) x = g(r_i)$. □

Rabin's Theorem suggests that we may view a computable field as a computably enumerable subfield of its (computable) algebraic closure, using a computable isomorphism, namely the Rabin embedding. The converse is readily seen as well.

Lemma 2.7. *Let E be a computable field, and F any subfield of E which is computably enumerable (as a subset of the domain ω of E). Then F is computably isomorphic to a computable field F'.*

Proof. Fix a computable enumeration $\{x_0, x_1, \ldots\}$ of the subfield F. F' has domain ω, of course, and we define addition on F' by:

$$m + n = p \iff x_m + x_n = x_p \text{ in } E$$

and multiplication similarly. Clearly F' is a computable field, with a computable isomorphism onto F given by $n \mapsto x_n$. We sometimes speak of F' as the *pullback* of F to ω. In positive characteristic, if F happens to be a finite subfield, the same construction works, with the domain of F' now being $\{0, 1, \ldots, |F| - 1\}$. □

Lemma 2.8. *For any single computable presentation of the algebraic closure $\overline{\mathbb{Q}}$, there is an algorithm which accepts as input any finite subsets $\{x_0, \ldots, x_m\}$ and*

$\{y_0, \ldots, y_n\}$ of $\overline{\mathbb{Q}}$, and decides whether $\mathbb{Q}[\vec{x}] \subseteq \mathbb{Q}[\vec{y}]$ or not. (Consequently, there is also an algorithm for deciding equality of these subfields.)

Proof. By Theorems 2.1 and 2.5, $\mathbb{Q}[\vec{y}]$ is computable, uniformly in $\vec{y}$, so we simply check whether $(\forall i \leq m) x_i \in \mathbb{Q}[\vec{y}]$. □

In this paper we will be concerned only with algebraic fields. Since the second part of the following definition is not standard, we state it here:

Definition 2.9. If $F \subseteq E$ are fields, then E is *algebraic over* F if every $x \in E$ is algebraic over F, i.e. is a root of some $p(X) \in F[X]$. When F is the prime subfield of E, we simply call E an *algebraic field*.

Thus the algebraic fields are precisely the subfields of the algebraically closed fields $\overline{\mathbb{Q}}$ and $\overline{\mathbb{Z}/p\mathbb{Z}}$. Elements of $\overline{\mathbb{Q}}$ are traditionally called *algebraic numbers*, but, by a longstanding and widely used definition, an *algebraic number field* is a finite algebraic extension of $\mathbb{Q}$, not an infinite one. Thus $\overline{\mathbb{Q}}$ itself is a field of algebraic numbers, but not an algebraic number field. We reiterate here that for us, every algebraic extension of either $\mathbb{Q}$ or $\mathbb{Z}/p\mathbb{Z}$, whether finite or infinite, will be called an *algebraic field*.

The following are standard results in field theory; see for instance [6], p. 215, Thm. 4.2 and Lemma 4.14.

Lemma 2.10. *If $F \subseteq E \subseteq K$ are finite field extensions, then their indices satisfy $[K : F] = [K : E] \cdot [E : F]$. Hence if E_1 and E_2 are finite extensions of F within a larger field, and $[E_1 : F]$ is relatively prime to $[E_2 : F]$, then $E_1 \cap E_2 = F$.*

Lemma 2.11. *In a finite normal algebraic extension $F \subseteq L$, $|Gal(L/F)| = [L : F]$ and each root in L of an irreducible $p(X) \in F[X]$ can be mapped to each other root of $p(X)$ in L by an element of $Gal(L/F)$. In fact, $p(X)$ is irreducible in $F[X]$ iff the Galois group of the splitting field of $p(X)$ over F acts transitively on the roots of $p(X)$.*

Recall that the *elementary symmetric polynomials* in $\{X_1, \ldots, X_m\}$ over F are by definition the polynomials

$$s_k(X_1, \ldots, X_m) = \sum_{1 \leq i_1 < \cdots < i_k \leq m} X_{i_1} X_{i_2} \cdots X_{i_k} \text{ (for } 1 \leq k \leq m).$$

The *symmetric polynomials* are the elements of $F[s_1, \ldots, s_m]$; they include precisely those polynomials in $F[X_1, \ldots, X_m]$ invariant under permutations of the variables.

Lemma 2.12. *Let $p(X) \in F[X]$ be a polynomial over a field F, and let $F \subseteq E$ be a field extension. Let $\overline{E}$ be the algebraic closure of E, and A the set of roots of $p(X)$ in $\overline{E}$. Assume that every root of $p(X)$ is a simple root, i.e. of multiplicity 1. Then the following are equivalent.*

1. *$p(X)$ is reducible in $E[X]$.*
2. *There exists $I = \{x_1, \ldots, x_m\}$ with $\emptyset \subsetneq I \subsetneq A$ such that every symmetric polynomial $h \in F[X_1, \ldots, X_m]$ has $h(x_1, \ldots, x_m) \in E$.*
3. *There exists $I = \{x_1, \ldots, x_m\}$ with $\emptyset \subsetneq I \subsetneq A$ such that every elementary symmetric polynomial $h \in F[X_1, \ldots, X_m]$ has $h(x_1, \ldots, x_m) \in E$.*

Proof. The elementary symmetric polynomials in I are the coefficients (up to sign) of the polynomial $\Pi_{x \in I}(X - x)$, so if they all lie in E, then this is a proper factor

of $p(X)$ in $E[X]$. Thus (3) implies (1). Conversely, if $q(X) \in E[X]$ is any proper factor of $p(X)$, let I be the set of roots of $q(X)$ in $\overline{E}$. These are all simple roots of $p(X)$, hence of $q(X)$, so the coefficients of $q(X)$ are (up to sign) the elementary symmetric polynomials in this I, and they all lie in E. The equivalence of (2) and (3) is clear. □

Lemma 2.13. *With F, E, $p(X)$, and A as in Lemma 2.12, let $I \subseteq A$ and $J = A - I$. Then the elementary symmetric polynomials in I generate the same subfield L_I of $\overline{E}$ as do the elementary symmetric polynomials in J.*

Proof. The polynomial $q_I(X) = \Pi_{r \in I}(X - r)$ lies in $L_I[X]$, so $L_I[X]$ also contains the quotient $\frac{p(X)}{q_I(X)} = \Pi_{r \in J}(X - r)$, whose coefficients are the elementary symmetric polynomials in J. The reverse inclusion likewise holds. □

We will also need the following, which can be found in many sources, including Section 6.10 of [17].

Theorem 2.14 (Theorem of the Primitive Element)**.** *Every finite separable algebraic extension $F \subseteq E$ is a* simple *extension. That is, there exists $z \in E$ such that $E = F[z]$.*

Indeed, algorithms for finding this z were known to Kronecker; for a modern treatment, see Lemma 17.12 of [2].

The remaining theorem we will need from field theory requires nontrivial Galois-theoretic results, for which we recommend the book [18] by Völklein. The author is grateful to Kevin Keating for pointing out these results and explaining how they combine to yield the theorem.

Theorem 2.15. *Let E be any finite algebraic extension of $\mathbb{Q}$, and fix any positive integer d. Then there exists a polynomial $q(X) \in E[X]$ of degree d such that the splitting field of $q(X)$ over E has Galois group S_d, the symmetric group on the roots of $q(X)$.*

Proof. Keating showed how to prove this theorem from a series of results in [18], as follows. A field K is *hilbertian* if for every polynomial $f(X, Y)$ in $K[X, Y]$ which is irreducible over $K(X)$ as a polynomial in Y, there exist infinitely many $b \in K$ such that the polynomial $f(b, Y)$ is irreducible in $K[Y]$. (Corollary 1.8 in [18] gives two other equivalent conditions.) Hilbert's Irreducibility Theorem states that the rational numbers have this property.

Theorem 2.16 (Hilbert; Thm. 1.23 in [18])**.** *The field $\mathbb{Q}$ is hilbertian.*

It then follows that the E in our theorem is hilbertian:

Lemma 2.17 (Corollary 1.11 in [18])**.** *Every finitely generated extension (either algebraic or transcendental) of a hilbertian field is hilbertian.*

Lemma 2.18 (Example 1.17 in [18])**.** *For all fields K and all $n > 0$, the symmetric group S_d occurs as a Galois group over the rational function field $K(X_1, \ldots, X_n)$.*

Lemma 2.19 (Thm. 1.13 in [18])**.** *If K is hilbertian and a finite group G occurs as a Galois group over some $K(X_1, \ldots, X_n)$, then G also occurs as a Galois group over K.*

Since our field E is hilbertian, there must exist a Galois extension $E \subseteq L$ with Galois group $G = \mathrm{Gal}(L/E) \cong S_d$. Theorem 4.7 of [6] shows that L is the splitting field for some polynomial in $E[X]$, but our theorem requires that polynomial to have degree d. For each element $\sigma \in G$, let $\overline{\sigma}$ be the corresponding permutation of $\{1, \ldots, d\}$. Let $L_1 \subseteq L$ be the subfield fixed pointwise by the subgroup

$$G_1 = \{\sigma \in G : \overline{\sigma}(1) = 1\}$$

and let $r_1 \in L_1$ be a primitive generator of L_1 over E, as given by Theorem 2.14. Now for any $\sigma, \tau \in G$,

$$\sigma(r_1) = \tau(r_1) \iff \sigma^{-1} \circ \tau \in G_1 \iff \overline{\sigma}(1) = \overline{\tau}(1).$$

The polynomial $q(X) \in E[X]$ required by the theorem will be the minimal polynomial of r_1 over E. Since L is normal over E, Lemma 2.11 shows that the conjugates of r_1 over E are precisely the images $\sigma(r_1)$ with $\sigma \in G$, and since $|G/G_1| = d$, there are exactly d of them, say $r_1, \ldots, r_d$. Since E is separable, $q(X)$ must have degree d. A similar analysis for the other roots r_j shows the action of G on $\{r_1, \ldots, r_d\}$ to be precisely the action of S_d on $\{1, \ldots, d\}$. But each $\sigma \in G$ is determined by the action of $\overline{\sigma}$ on $\{1, \ldots, d\}$, so the roots $\{r_1, \ldots, r_d\}$ must generate all of L over E. Thus L is the splitting field of $q(X)$ over E, proving Theorem 2.15. □

3. **1**-Reducibility

In this section we give our positive result: uniform 1-reducibility of the splitting set of an algebraic field to its root set.

Theorem 3.1. *For every computable algebraic field F with splitting set S and root set R, in any characteristic, we have $S \leq_1 R$, via a computable function φ whose index is computable from the indices for addition and multiplication in F.*

Proof. Begin by fixing a Rabin embedding g of F into a computable field which we will call $\overline{F}$. For simplicity we will view the image $g(F)$ as F itself, a computably enumerable subfield within $\overline{F}$. Since g is computable, g^{-1} is also computable; g^{-1} is partial, of course, but we will only need to apply g^{-1} to elements already known to lie in the image of g.

Fix a computable enumeration $\{p_0(X), p_1(X), \ldots\}$ of the polynomial ring $F[X]$. We define $\varphi(p_e)$ for each $e = 0, 1, \ldots$ in order, according to the following computable process. Given e, first find an s such that all coefficients of $p_e(X)$ lie in $\{0, 1, \ldots, s\}$. We will write F_s for the computably enumerable subfield of F generated by $\{0, \ldots, s\}$. Thus F_s is a finite algebraic extension of the prime field P of $\overline{F}$. By Lemma 2.2, we have a splitting algorithm for F_s, uniformly in s, and so we can check whether $p_e(X)$ splits over F_s. If it does, then we immediately define $\varphi(p_e)$ to be the polynomial $(X - t)$, where t is the first element in the enumeration of F such that $(X - t) \notin \{\varphi(p_0), \ldots, \varphi(p_{e-1})\}$. In this case clearly $p_e \in S$ and $\varphi(p_e) \in R$.

Assuming that $p_e(X)$ is irreducible over F_s, therefore, let d be its degree, and find all the roots $r_1, \ldots, r_d$ of $p(X)$ in the algebraically closed field $\overline{F}$. (Being a finite algebraic extension of its prime field, F_s is perfect, so these roots are all simple and distinct.) Let G be the Galois group of the splitting field $K \subset \overline{F}$ of $p(X)$ over F_s, viewed as a group of permutations of the set $A = \{r_1, \ldots, r_d\}$. For each $I \subseteq A$, let $K_I \subseteq K$ be the subfield generated by the coefficients of the polynomial

$$\Pi_{r \in I}\ (X - r),$$

i.e. by the elementary symmetric polynomials in the elements of I. Lemma 2.2 yields splitting algorithms for K and for each K_I, uniformly in e and I, each of these is a computable subfield of $\overline{F}$.

By Theorem 2.14, there must exist a single element x_I for each $I \subseteq A$, such that $K_I = F_s[x_I]$. We search for such elements, and eventually, for each I, we find one and recognize it, since it must lie in K_I and generate the finite generating set of K_I. Using the splitting algorithm for F_s, we may find the minimal polynomial $q_I(X) \in F_s[X]$ of each x_I. Define $q(X)$ to be the product of these polynomials, for all nontrivial I:

$$q(X) = \Pi_{\emptyset \subsetneq I \subsetneq A}\ q_I(X)$$

Fix the least $k > 0$ such that $(q(X))^k \notin \{\varphi(p_0), \ldots, \varphi(p_{e-1})\}$, and set $\varphi(p_e) = (q(X))^k$. (Thus we ensure that φ remains one-to-one, by checking against all previous values of φ.)

Now we claim that the polynomial $\varphi(p_e)$ has a root in F iff $p_e(X)$ itself factors over F. First, suppose that $\varphi(p_e)$ has a root $x \in F$. Then for some nonempty $I \subsetneq A$ we have $q_I(x) = 0$. Now $x \in K$, since the splitting field K is normal over F_s, and so $F_s[x]$ is a subfield of K, and is isomorphic to $F_s[x_I]$ since $q_I(X)$ was irreducible over F_s. By Lemma 2.11, some ρ in the Galois group of q_I over F_s must have $\rho(x_I) = x$, and this ρ must extend to an element $\overline{\rho} \in G$ since K is normal. Since the elementary symmetric polynomials in I generate K_I, we see that all elementary symmetric polynomials in $\overline{\rho}(I) = \{\overline{\rho}(r) : r \in I\}$ lie in the image $\overline{\rho}(K_I)$ and generate this image. Therefore the coefficients of the polynomial

$$p_\rho(X) = \Pi_{r \in I}(X - \overline{\rho}(r))$$

all lie in $\overline{\rho}(K_I)$. But $\overline{\rho}(K_I)$ is generated by x over F_s, and $x \in F$ by assumption, so $\overline{\rho}(K_I) \subseteq F$, and thus $p_\rho(X) \in F[X]$. Since $\overline{\rho} \in G$ must map roots of p_e to roots of p_e, and since $\emptyset \subsetneq I \subsetneq A$, this $p_\rho(X)$ is a proper factor of $p_e(X)$ within $F[X]$.

Conversely, suppose that $p_e(X)$ has some proper factor in $F[X]$. Since $p_e(X) = \Pi_{r \in A}(X - r)$ in $\overline{F}[X]$, we may write this factor as

$$p_I(X) = \Pi_{r \in I}(X - r)$$

for some nonempty $I \subsetneq A$, and the coefficients of $p_I(X)$ are precisely the elementary symmetric polynomials in I. So all these coefficients lie in F. But the subfield K_I is generated by these polynomials, and since $x_I \in K_I$, we must have $x_I \in F$. Therefore $\varphi(p_e)$ has a root in F, completing the proof that φ is a 1-reduction from S to R.

Finally, the claim about computability of an index of φ follows from a careful reading of the above proof: the only information about F necessary for the construction of φ was the ability to add and multiply elements of F. With indices for the addition and multiplication functions in F, we can build the field $\overline{F}$ and the Rabin embedding g, enumerate $F[X]$, and perform all the steps required by the program we gave for computing $\varphi(p_e)$. Computability theorists therefore say that φ is given *uniformly* in the computable algebraic field F: from the field operations in F we can figure out the program for computing φ. If this notion is new to the reader, the beginning of the next section will help explain it. □

The foregoing proof generalizes to any computable extension of $\mathbb{Q}$ for which we have a computable transcendence basis.

Corollary 3.2. *Let F be a computable field of characteristic 0 with a computably enumerable transcendence basis. Then $S_F \leq_1 R_F$, and the reduction is uniform in F and an enumeration of the transcendence basis.*

Proof. Let $B = \{b_0, b_1, \ldots\}$ enumerate a transcendence basis. Rabin's Theorem still applies, and the image of B in $\overline{F}$ is a c.e. transcendence basis for $\overline{F}$. Take $F_0 = P(B)$, the (purely transcendental) extension of the prime subfield P by B. This F_0 is still c.e. within F, and has a splitting algorithm, since the coefficients of any $p(X) \in F_0[X]$ eventually all appear in some $P(b_0, \ldots, b_n)$, for which Theorem 2.1 provides a splitting algorithm. Let $F_s = F_0[0, \ldots, s]$, and proceed with the reduction described in the proof of Theorem 3.1. □

In fact, in a computable field, any c.e. transcendence basis is computable, since we can find the minimal polynomial of an arbitrary element x over the c.e. transcendence basis and use it to check whether x is in the basis.

4. **1**-Nonreducibility

As a warm-up for the principal negative result of this paper: we prove first that there is no uniform m-reduction of splitting sets to root sets. Here we will use the standard enumeration $\varphi_0, \varphi_1, \ldots$ of all partial computable functions by their programs; see for instance [15, I.3.1]. We claim that if ψ is a partial computable function, then there exist some a and b such that φ_a and φ_b define the addition and multiplication in some computable algebraic field F, yet $\psi(a, b)$ either diverges or converges to a value e such that φ_e is *not* an m-reduction from R_F to S_F. In fact, in Theorem 4.2 below we will build a single computable field F for which there is no such m-reduction at all, thereby generalizing Proposition 4.1 and showing the failure of the reverse reduction from Theorem 3.1. However, this proposition gives a useful introduction, in a simpler context, to the techniques we will use in Theorem 4.2.

Proposition 4.1. *For computable algebraic fields F, there is no m-reduction uniform in F from the root set R_F to the splitting set S_F.*

Proof. The key to this result is the Recursion Theorem, presented in most standard computability textbooks, including [15], where it appears as Theorem II.3.1. It allows us to take an arbitrary computable function ψ and run it on the indices a and b for addition and multiplication in the field that we construct below. The construction itself requires waiting for $\psi(a, b)$ to converge, and therefore appears circular, but the circularity is removed by application of the Recursion Theorem (or more precisely, by Smullyan's Double Recursion Theorem, item II.3.15 in [15]).

Fixing ψ, we define our field F as follows. First build a computable copy F_0 of the rationals themselves. We define the polynomial $q(X) = X^5 - X - 1 \in F_0[X]$. Now our programs φ_a and φ_b simply wait for $\psi(a, b)$ to converge. If it never converges, then we never add any more elements to F_0; in this case $F = F_0$. If it does converge, say to a value e, then we run the e-th program φ_e on input $q(X)$, and wait for it to converge. If this computation never converges, then again $F = F_0$. In both these cases ψ fails to produce the index for an m-reduction of R_F to S_F. (Technically, we actually build F by adding only finitely many elements of F_0 to F at a time; thus, if $\psi(a, b)$ never converges, the domain of F will be ω, whereas if $\psi(a, b)$ does eventually converge, we still have cofinitely many domain elements available on which to change our strategy as described below.)

If we find a stage s such that $\psi(a,b)\downarrow = e$ and also $\varphi_e(q(X))$ converges to some polynomial $p(X) \in F_0[X]$, all within s steps, then ψ thinks that φ_e is an m-reduction for F, and therefore thinks that $q(X) \in R_F$ iff $p(X) \in S_F$. If we left F equal to F_0, this might be the case, so when convergence occurs, we act to ensure that it is false.

Let $A = \{r_1, \ldots, r_5\}$ be the set of roots of $q(X)$ in a fixed computable presentation $\overline{\mathbb{Q}} \supset F_0$ of the algebraic closure of $\mathbb{Q}$. First we find these five roots and let K be the field they generate. This K has a splitting algorithm and is a computable subfield of $\overline{\mathbb{Q}}$, by Rabin's Theorem. It is shown in [17, Section 8.10] that $q(X)$ is irreducible over $\mathbb{Q}$ and that the Galois group $\mathrm{Gal}(K/F_0)$ is S_5, the symmetric group on A. (This is why we chose this q.)

We also have the polynomial $p(X) \in F_0[X]$ produced by φ_e. Let n be its degree, and use the splitting algorithm for F_0 to determine whether it is reducible over F_0. If so, then we leave $F = F_0$, since then $q(X) \notin R_F$ but $p(X) \in S_F$. Also, if $p(X)$ is a linear or constant polynomial, then it can never factor, so we set $F = F_0[r_1]$, thus putting $q(X)$ into R_F, while $p(X) \notin S_F$.

If $p(X)$ is irreducible over F_0 of degree > 1, then we find its n roots (all distinct) in $\overline{\mathbb{Q}}$. Let $B = \{x_1, \ldots, x_n\}$ be the set of these roots, and let $L = F_0[x_1, \ldots, x_n] \subset \overline{\mathbb{Q}}$ be the splitting field of $p(X)$. The subfield $F_0[x_1] \subset \overline{\mathbb{Q}}$ is computable, by Theorem 2.5, since Lemma 2.2 provides a splitting algorithm for it. So we check whether any r_i lies in $F_0[x_1]$. If not, then we adjoin x_1 in our construction of F, so that F is (a computable isomorphic copy of) $F_0[x_1]$. Thus $p(X)$ factors over F, having $(X - x_1)$ has a factor, yet $q(X) \notin R_F$, since no r_i lies in $\mathbb{Q}[x_1]$.

Finally, therefore, suppose that some r_i (say r_1, without loss of generality) lies in $F_0[x_1]$. Notice that since the splitting field L is normal over F_0, this forces $K \subseteq L$. Now we can find an $h(X) \in F_0[X]$ with $h(x_1) = r_1$. For every x_j we have some $\sigma \in \mathrm{Gal}(L/F_0)$ with $\sigma(x_1) = x_j$, since $p(X)$ is irreducible, and so

$$q(h(x_j)) = q(h(\sigma(x_1))) = \sigma(q(h(x_1)) = \sigma(q(r_1)) = \sigma(0) = 0,$$

forcing $h(x_j) = r_k$ for some k. Likewise, for every r_k, the element $\rho \in \mathrm{Gal}(K/F_0)$ interchanging r_1 with r_k extends by normality to $\overline{\rho} \in \mathrm{Gal}(L/F_0)$, so that $r_k = \overline{\rho}(r_1) = h(\overline{\rho}(x_1)) = h(x_j)$ for some j. We now add elements to F so that F is the subfield of L containing those field elements fixed by the subgroup $G_{12} \subset \mathrm{Gal}(L/F_0)$, where

$$G_{12} = \{\sigma \in \mathrm{Gal}(L/F_0) : \{\sigma(r_1), \sigma(r_2)\} = \{r_1, r_2\}\}.$$

This is the subgroup of automorphisms fixing $\{r_1, r_2\}$ setwise (but *not* necessarily pointwise; for instance, the $\tau \in \mathrm{Gal}(K/F_0)$ interchanging r_1 with r_2 does extend to an element of G_{12}). Let

$$I = \{x_j : h(x_j) = r_1 \text{ or } h(x_j) = r_2\}.$$

Then $x_1 \in I$, but $I \neq B$, since there is some j with $h(x_j) = r_3$, as remarked above. Moreover, I is fixed setwise by every element of G_{12}. Therefore, every symmetric polynomial $g(X_1, \ldots, X_{|I|})$ over F_0 has $g(I) \in F$, and so $p(X)$ factors in $F[X]$, by Lemma 2.12, with proper factor

$$\Pi_{x \in I}(X - x).$$

Thus we have $p(X) \in S_F$. On the other hand, since $\mathrm{Gal}(K/F_0) \cong S_5$, we have a $\tau \in \mathrm{Gal}(K/F_0)$ which interchanges r_1 with r_2 and has $\tau(r_3) = r_4$, $\tau(r_4) = r_5$, and

$\tau(r_5) = r_3$. By normality, this τ extends to a $\overline{\tau} \in G_{12}$, which shows that no r_k lies in F. Thus $q(X) \notin R_F$, and so once again φ_e fails to be an m-reduction from R_F to S_F. This completes our construction of F and proves Proposition 4.1. □

Now we wish to construct a single field F in which there is no m-reduction whatsoever from R_F to S_F. The difficulty here is that we cannot just repeat the construction above for all different partial computable functions φ_e. Obviously we can only use the same polynomial $q(X) = X^5 - X - 1$ once, to make one particular φ_e fail; but it is not hard to find other polynomials $q_e(X)$, indeed of arbitrary degree, whose splitting fields have symmetric Galois group, and to use one of them against each φ_e. The difficulty is that, as above, we may have to adjoin elements to $\mathbb{Q}$ to form F. In Proposition 4.1, if this happened, we found that either $[F : \mathbb{Q}] = n$, the degree of $p(X)$, if $F = \mathbb{Q}[x_1]$; or that $[F : \mathbb{Q}] = \binom{5}{2} = 10$ if we adjoined the fixed field of the group G. (In general, adjoining the fixed field of a subgroup G of the Galois group of a splitting field extends the ground field by a degree equal to the index of G in that Galois group.) When we adjoin elements to defeat one function φ_e, it may be that the polynomials $q_{e'}(X)$, with $e' \neq e$, no longer have symmetric Galois group over the new field. Worse yet, we might even have adjoined a root of some $q_{e'}$, rendering it useless in our effort to show that φ_e does not compute an m-reduction. ($\varphi_{e'}(q_{e'})$ could then converge to a polynomial which already factors over F, so that $q_{e'} \in R_F$ and $\varphi_{e'}(q_{e'}) \in S_F$.) The latter of these cases is less of a problem, since we have some control over it: we can define d_e, the degree of $q_e(X)$, to make $\binom{d_e}{2}$ have whatever value we wish. The former is more difficult: φ_e gets to choose n and the roots x_i, by choosing $\varphi_e(q_e)$, and so the degree of x_1 over F may be any value at all.

If we knew that all computations $\varphi_e(q_e)$ would eventually converge, then we might be able to carry out the strategy from Proposition 4.1 with φ_0 over $F_0 = \mathbb{Q}$ to build F_1, then repeat it with φ_1 over F_1 to build F_2, while being sure not to disturb the result for φ_0 on F_1, and so on. However, a further complication is that we have no way to determine in general whether the computation of φ_e on the polynomial q_e will even converge at all; this is (a subproblem of) the Halting Problem. To take care of all these difficulties, we appeal to the method of the *finite-injury priority construction*, introduced independently by Friedberg and Muchnik in [3] and [12] and described in Chapter VII of [15].

Theorem 4.2. *There exists a computable algebraic field F with splitting set $S = S_F$ and root set $R = R_F$, for which $R \not\leq_m S$ (and hence $R \not\leq_1 S$).*

Proof. We first give our construction of the computable field F, and then prove that $R \not\leq_m S$. Using an effective listing of all the computable partial functions:

$$\varphi_0, \varphi_1, \varphi_2, \ldots,$$

we will build F to satisfy for every e the requirement

$$\mathcal{R}_e : \ \varphi_e \text{ is not an } m\text{-reduction from } R \text{ to } S.$$

This will enable us to show that $R \not\leq_m S$: no total computable function succeeds in the role required for m-reducibility. Our numbering acts as a priority ranking on these requirements: satisfying $\mathcal{R}_0$ is our highest priority, satisfying $\mathcal{R}_1$ is our next highest, and so on. Occasionally two requirements will demand that we perform contradictory actions; when this happens, we follow the demand of the higher-priority requirement, and say that the lower-priority one has been *injured* by the

higher-priority one at this stage. Our construction will ensure that each requirement is injured at only finitely many stages, so that we will have cofinitely many stages during which to satisfy it.

(Of course, all total computable functions appear in our listing of the functions φ_e above. For arbitrary e and n, though, we have no way to know whether $\varphi_e(n)$ is defined or not; we can only run the e-th program on the input n and wait to see whether it ever halts. So our construction will allow for the possibility that $\varphi_e(n)$ never halts, but also must allow it infinitely much time to run, just in case it ever does halt. For instance, if $\varphi_2(q_2)$ runs for very many stages and then finally halts, then in order to satisfy $\mathcal{R}_2$ after that halt, we will likely have to injure many requirements $\mathcal{R}_e$ with $e > 2$, on whose behalf we had already acted. After satisfying $\mathcal{R}_2$, we will return our attention to those other $\mathcal{R}_e$ and start over from scratch to satisfy them.)

We enumerate F as a subfield of (a computable presentation of) the algebraic closure $\overline{\mathbb{Q}}$. By Lemma 2.7, there exists a computable field F' with a computable isomorphism from F' onto F. So any m-reduction of the root set of F' to its splitting set would yield such a reduction for F as well, by applying the computable isomorphism (and its inverse, which is also computable with domain F).

Our construction proceeds in stages. For each e, we will eventually choose a *witness polynomial* $q_e(X) \in \mathbb{Q}[X]$, and will give it to φ_e and wait (forever, if necessary) for $\varphi_e(q_e)$ to halt. If this computation ever does halt, then we will add elements to the field F to ensure that $q_e(X) \in R$ iff $\varphi_e(q_e) \notin S$. Thus we will satisfy the requirement $\mathcal{R}_e$. Of course, if $\varphi_e(q_e)$ fails to halt, then $\mathcal{R}_e$ is satisfied, since an m-reduction must be total. However, at certain stages we may have to change this witness polynomial used for $\mathcal{R}_e$, and so we write $q_{e,s}$ for the witness polynomial in use at stage s. We will ensure that it changes at only finitely many stages, so that the argument above does apply to q_e itself, which equals $q_{e,s}$ for all but finitely many s.

By each stage s we will have added finitely many elements to the prime field $F_0 = \mathbb{Q}$, and we will write F_s for the subfield they generate within $\overline{\mathbb{Q}}$. (We may imagine building the rest of the entire finitely-generated field F_s at the end of each stage s; the construction tacitly assumes this to have been done.) Each F_{s+1} will therefore be a finite algebraic extension of F_s, with $F = \cup_s F_s$, and since the extensions are finite, Lemma 2.2 yields splitting algorithms for each F_s, uniformly in s. Thus in fact each F_s is a computable subfield of $\overline{\mathbb{Q}}$.

To begin with, F_0 is just the field $\mathbb{Q}$ within $\overline{\mathbb{Q}}$. Every polynomial $q_{e,0}$ is undefined, and every requirement $\mathcal{R}_e$ is unsatisfied at this stage. For simplicity we set $d_{-1,0} = 3$; all $d_{e,0}$ with $e \in \omega$ are undefined.

At stage $s+1$, we assume that we have already constructed the field F_s to be a finite algebraic extension of F_0 within $\overline{\mathbb{Q}}$, so that we know a splitting algorithm for F_s. For each $e < s$ such that $q_{e,s}$ is defined and $\mathcal{R}_e$ is not currently satisfied, we check whether the computation of $\varphi_e(q_{e,s})$ halts within s steps.

If there is no $e < s$ for which this computation halts, then we consider the least e such that $q_{e,s}$ was not defined. For this e, we know by induction that $d_{e-1,s}$ is defined, and we use it to define a new witness polynomial $q_{e,s+1}(X) \in F_s[X]$, of degree $d_{e,s+1}$ as follows. Let M be the product of all the prime numbers $\leq d_{e-1,s}$. We set $d_{e,s+1} = 2M \cdot [F_s : \mathbb{Q}] - 1$, so that both $d_{e,s+1}$ and $\frac{d_{e,s+1}-1}{2} = M \cdot [F_s : \mathbb{Q}] - 1$ are relatively prime to $[F_s : \mathbb{Q}]$ and to all $d_{i,s}$ with $i < e$, and indeed to any number

less than any such $d_{i,s}$. Then by Theorem 2.15, there must exist some polynomial $q(X) \in F_s[X]$ of degree $d_{e,s+1}$ such that the Galois group of its splitting field over F_s is the symmetric group on the $d_{e,s+1}$ (distinct) roots of $q(X)$ in $\overline{\mathbb{Q}}$. Using Lemmas 2.2 and 2.3, we search until we find such a polynomial, and define it to be the witness polynomial $q_{e,s+1}(X)$. We leave $F_{s+1} = F_s$, and for all $i < e$ we set $q_{i,s+1} = q_{i,s}$ and $d_{i,s+1} = d_{i,s}$, so (by induction) every $d_{i,s+1}$ is divisible by the degree of the current witness polynomial $q_{i,s+1}(X)$ for $\mathcal{R}_i$. For $j > e$, $d_{j,s+1}$ and $q_{j,s+1}$ remain undefined. The requirements satisfied by stage $s+1$ are precisely those satisfied by stage s, since our action at this stage did not fulfill any more of the $\mathcal{R}_e$.

On the other hand, if there exists an $e < s$ for which $\varphi_e(q_{e,s})$ halts within s steps and $\mathcal{R}_e$ is not currently satisfied, then we fix the least such e and act to satisfy $\mathcal{R}_e$ at this stage. Immediately we make $q_{j,s+1}$ and $d_{j,s+1}$ undefined for all $j > e$, and say that these requirements $\mathcal{R}_j$ are all unsatisfied at stage $s+1$ and have been *injured* at this stage by the higher-priority requirement $\mathcal{R}_e$. For all $i < e$ we set $q_{e,s+1} = q_{e,s}$ and $d_{e,s+1} = d_{e,s}$; $\mathcal{R}_i$ is satisfied at stage $s+1$ iff it was satisfied at stage s.

Write $p(X)$ for the polynomial given by $\varphi_e(q_{e,s})$, and $x_1, \ldots, x_n$ for its roots in $\overline{\mathbb{Q}}$, and set $L = F_s[x_1, \ldots, x_n]$ to be the splitting field of $p(X)$ over F_s within $\overline{\mathbb{Q}}$. Likewise, let $r_1, \ldots, r_d \in \overline{\mathbb{Q}}$ be the roots of the polynomial $q(X) = q_{e,s}(X)$ (so $d = d_{e,s}$), and let $K = F_s[r_1, \ldots, r_d] \subseteq \overline{\mathbb{Q}}$ be its splitting field.

Define the intermediate fields $L_0, L_1, \ldots, L_{2^n-3}$ to be the *symmetric subfields* for $p(X)$ over F_s. That is, for each set I with $\emptyset \subsetneq I \subsetneq \{x_1, \ldots, x_n\}$, let the next L_i be the subfield of L generated by the elementary symmetric polynomials in I. We do not worry that this list includes some repetitions, but note that if $p(X)$ has degree ≤ 1, then there are no symmetric subfields. Now if $p(X)$ has only simple roots, then Lemma 2.12 shows that for fields E with $F_s \subseteq E \subseteq \overline{\mathbb{Q}}$, $p(X)$ factors in $E[X]$ iff E contains some L_i.

We act according to the following four cases.

(1) Check whether $p(X)$ is reducible over F_s, using the splitting algorithm for F_s. If it is, then $p(X) \in S$ and $q(X) \notin R_{F_s}$, so we simply set $F_{s+1} = F_s$, and keep $d_{e,s+1} = d_{e,s}$ and $q_{e,s+1} = q_{e,s}$ so as to ensure that $q(X)$ stays out of R.

(2) Otherwise, check whether, for all $i < 2^n - 2$, we have $L_i \not\subseteq F_s[r_1]$. (This includes the case in which $p(X)$ has degree ≤ 1.) If so, then set $F_{s+1} = F_s[r_1]$, so that $q(X) \in R$ but $p(X) \notin S_{F_{s+1}}$, by Lemma 2.12. In this case we set $q_{e,s+1} = q_{e,s}$ and define $d_{e,s+1}$ to be the least prime which is both $\geq d_{e,s}$ and $\geq |\mathrm{Gal}(L/F_s)| = [L : F_s]$. (So the degree of $q_{e,s+1}$ may not equal $d_{e,s+1}$.) Whenever a polynomial $q_{j,s'}$ with $j > e$ is defined at a stage $s' > s+1$, its degree $d_{j,s'}$ will be chosen to be relatively prime to all numbers $\leq d_{e,s+1}$, hence relatively prime both to $d_{e,s}$ and to $|\mathrm{Gal}(L/F_s)| = [L : F_s]$. This will ensure that if we subsequently adjoin a root of $q_{j,s'}$ to F to satisfy a lower-priority requirement $\mathcal{R}_j$, we cannot accidentally make $p(X)$ reducible, and so $p(X)$ will stay out of S.

(3) Otherwise, check whether there is an i such that $L_i \subsetneq F_s[r_1]$. If so, then for the least such i, set $F_{s+1} = L_i$; this will ensure that $p(X) \in S$. To preserve $q(X) \notin R$, we set $d_{e,s+1} = d_{e,s}$, with $q_{e,s+1} = q_{e,s}$.

(4) Otherwise, consider the subgroups G_{12}, G_{13}, and G_{23}, where

$$G_{jk} = \{\sigma \in \mathrm{Gal}(L/F_s) : \sigma \text{ fixes } \{r_j, r_k\} \text{ setwise}\}.$$

Lemma 4.4 below shows that the fixed field of at least one of these groups contains some symmetric subfield L_i for some i. Let F_{s+1} be that fixed field (for the least i, if we have a choice). Now we know that $p(X) \in S$, by Lemma 2.12. On the other hand, there is an element of $\mathrm{Gal}(L/F_s) \cong S_{d_{e,s}}$ interchanging r_j with r_k, by our choice of $q_{e,s}$ from an earlier stage and by Lemma 4.3 below. Hence r_j and r_k do not lie in the fixed field $L_i = F_{s+1}$. Also, $d_{e,s+1} > 3$, and so there is likewise an element of $\mathrm{Gal}(L/F_s)$ which permutes the set $(\{r_1, \ldots, r_{d_{e,s}}\} - \{r_j, r_k\})$ cyclically, with no fixed point. Thus no root of $q_{e,s}$ lies in F_{s+1}. We set $q_{e,s+1} = q_{e,s}$ and $d_{e,s+1} = d_{e,s}$, to ensure that this remains true at subsequent stages, so that $q_{e,s+1} \notin R_F$.

In each of these four cases, we now declare $\mathcal{R}_e$ satisfied. This completes stage $s+1$.

Our construction builds fields $F_s \subseteq F_{s+1} \subset \overline{\mathbb{Q}}$ for all s. The union of these is a computably enumerable subfield $F \subseteq \overline{\mathbb{Q}}$, which we may view as a computable field in its own right, using Lemma 2.7. It is clear (in light of Theorem 2.15) that the construction continues through all stages, without spending eternity at any single stage $s+1$. We claim that every requirement $\mathcal{R}_e$ is true for this F. For this, we need to prove that the relevant properties were preserved at every stage.

Lemma 4.3. *For every e and every stage s at which $F_{s+1} \neq F_s$, if $\mathcal{R}_e$ is the requirement satisfied at stage $s+1$, then $[F_{s+1} : F_s]$ is not divisible by any prime $\leq d_{e-1,s}$.*

Proof. We consider the four possible ways in which $\mathcal{R}_e$ may be satisfied at stage $s+1$, as listed on page 15, and argue by induction on s. Suppose that at stage $s+1$ we satisfied $\mathcal{R}_e$. Set $d = d_{e,s}$ and $q = q_{e,s}$.

Now r_1 is a root of the polynomial $q(X) \in F_s[X]$, which has degree dividing d (indeed equal to d, unless we used Case 2 to satisfy $\mathcal{R}_e$). Moreover, from the stage s_0 at which $q = q_{e,s_0}$ was chosen up until the current stage $s+1$, no requirement r_i with $i < e$ has acted, since such an action at such a stage s' would have caused $q_{e,s'}$ to become undefined. On the other hand, at stage s_0, all requirements $\mathcal{R}_j$ with $j > e$ had d_{j,s_0} undefined, and if $d_{j,s'}$ was subsequently defined at some stage s' with $s_0 < s' \leq s$, then it was chosen to be relatively prime to every prime $\leq d_{e,s'} = d$. By inductive hypothesis, therefore, we see that $d!$ is relatively prime to $[F_{s'+1} : F_{s'}]$ for all such s'. Therefore $q(X)$ remains irreducible over F_s, and indeed the Galois group of its splitting field over F_s is still the symmetric group on the roots $\{r_1, \ldots, r_d\}$, since the intersection of this splitting field with $F_{s'+1}$ is still F_s.

If we were in Case (1) at stage $s+1$, then $[F_{s+1} : F_s] = 1$. In Cases (2) and (3), F_{s+1} is chosen to be a subfield of $F_s[r_1]$. But by irreducibility of $q(X)$ over F_s, $[F_s[r_1] : F_s] = d$, relatively prime to all primes $\leq d_{e,s-1}$ (since $d = d_{e,s_0}$ was chosen thus), and so the intermediate field F_{s+1} must also have $[F_{s+1} : F_s]$ relatively prime to all those primes, by Lemma 2.10.

Finally, if we used Case (4) to satisfy $\mathcal{R}_e$, then we adjoined to F_s the fixed field of a subgroup G_{jk} of $\mathrm{Gal}(L/F_s)$. Notice that in this case, $p(X)$ is irreducible and not linear, since Cases (1) and (2) did not apply. Moreover, there is an i with $L_i = F_s[r_i]$: the failure of Case (2) gives us this i, and the failure of Case (3) shows the equality. So $r_i \in L_i \subseteq L$, and since L is normal over F_s, all of $r_1, \ldots, r_d$ lie in L. Let $N = \{\sigma \in \mathrm{Gal}(L/F_s) : (\forall i \leq d)\sigma(r_i) = r_i\}$; this is just

$\mathrm{Gal}(L/K)$. Now by normality, every $\rho \in \mathrm{Gal}(K/F_s)$ extends to exactly $|N|$-many distinct elements of $\mathrm{Gal}(L/F_s)$. But $\mathrm{Gal}(K/F_s) \cong S_d$ has $d!$-many elements, of which exactly $2 \cdot (d-2)!$-many fix the set $\{r_j, r_k\}$. So, out of the $|N| \cdot d!$ elements of $\mathrm{Gal}(L/F_s)$, exactly $|N| \cdot 2 \cdot (d-2)!$ lie in G_{jk}. Lemma 2.11 then yields

$$[F_{s+1} : F_s] = \frac{[L : F_s]}{[L : F_{s+1}]} = \frac{|\mathrm{Gal}(L/F_s)|}{|\mathrm{Gal}(L/F_{s+1})|} = \frac{|N| \cdot d!}{|N| \cdot 2 \cdot (d-2)!} = d \cdot \frac{d-1}{2}$$

and our choice of $d = d_{e,s_0}$ at stage s_0 made both d and $\frac{d-1}{2}$ relatively prime to all primes $\leq d_{e-1,s_0} = d_{e-1,s}$. This completes the induction. □

Lemma 4.4. *At stage $s+1$ in the above construction, if we reach Case (4) for the chosen value of e, then at least one of the given groups G_{12}, G_{13}, and G_{23} has fixed field containing some symmetric subfield L_i of $p(X)$ over F_s (and hence $p(X)$ factors over that fixed field).*

Proof. Recall that $G_{jk} = \{\sigma \in \mathrm{Gal}(L/F_s) : \sigma(r_1) \in \{r_j, r_k\}\}$, and define $G_j = \{\sigma \in \mathrm{Gal}(L/F_s) : \sigma(r_1) = r_j\}$. Since Cases (1), (2), and (3) all do not apply, we know that $F_s[r_1]$ itself is a symmetric subfield of L, and that $p(X)$, being irreducible, has distinct roots $x_1, \ldots, x_n$. For simplicity, reorder these roots so that $F_s[r_1]$ is the subfield of symmetric polynomials in the set $I = \{x_1, \ldots, x_d\}$, with $1 \leq d < n$. By Lemma 2.13, we may assume that $d \leq \frac{n}{2}$.

Now every elementary symmetric polynomial $s_1, \ldots, s_d$ in the elements of I lies in $F_s[r_1]$, so we have polynomials $h_i \in F_s[X]$ with $h_i(r_1) = s_i(I)$ for all $i \leq d$. (Since s_i is symmetric, we may simply write $s_i(I)$.) Sublemma 4.5 below, with $F_s[r_1]$ as E, will use these equations to show that x_1 must be a root of a polynomial $h(X)$ of degree d, with coefficients in $F_s[r_1]$. Therefore, if $\sigma \in \mathrm{Gal}(L/F_s)$, has $\sigma(r_1) = r_1$, there are only d possible values for $\sigma(x_1)$. (Those values must all satisfy $h(\sigma(x_1)) = \sigma(h(x_1)) = \sigma(0) = 0$, since such a σ fixes the coefficients of h.) Likewise, for any fixed value of $\sigma(r_1)$ (such as r_2 or r_3), there are at most d possible values for $\sigma(x_1)$: the roots of the image of the polynomial $h(X)$ after σ has acted on its coefficients. That is, the orbit $X_j = \{\sigma(x_1) : \sigma(r_1) = r_j\}$ of x_1 under G_j contains at most d elements, and the orbit $X_{jk} = \{\sigma(x_1) : \sigma \in G_{jk}\}$ of x_1 under G_{jk} contains at most $2d$ elements,

If $d < \frac{n}{2}$, this means that not every element of $\{x_1, \ldots, x_n\}$ lies in the orbit of x_1 under the action of G_{12} on this set. If $d = \frac{n}{2}$, then it is possible that $X_{12} = \{x_1, \ldots, x_n\}$. Assuming this to be the case, the nonempty subset X_3 of $\{x_1, \ldots, x_n\}$ intersects X_{12}, hence must intersect either X_1 or X_2. But each X_j contains at most d elements, so if $X_1 \cap X_3 \neq \emptyset$, then $X_{13} = X_1 \cup X_3$ contains $< n$ elements, and otherwise $X_2 \cap X_3 \neq \emptyset$ and X_{23} contains $< n$ elements.

So in every case we see that one of these three subgroups G_{jk} has more than one orbit in $\{x_1, \ldots, x_n\}$, since the orbit of x_1 contains $< n$ elements. But if J is a proper nonempty orbit under the action of G_{jk}, then every symmetric polynomial in J over F_s is fixed by every $\sigma \in G_{jk}$, because such a σ must map J bijectively onto itself. Lemma 2.12 then shows that $p(X)$ is reducible over the fixed field of G_{jk}.

It remains to prove the promised result on symmetric polynomials:

Sublemma 4.5. *For any d, any field E, and any system of d equations in the elementary symmetric polynomials*

$$a_1 = s_1(X_1, \ldots, X_d) \ \cdots \ a_d = s_d(X_1, \ldots, X_d)$$

with all $a_i \in E$, every coordinate b of every solution in E^d to this system satisfies $t(b) = 0$, where $t(Y) = Y^d - a_1 Y^{d-1} + a_2 Y^{d-2} - \cdots + (-1)^d a_d$.

Proof. Suppose $\vec{b} = (b_1, \ldots, b_d) \in E^d$ is a solution. Then each $a_i = s_i(\vec{b})$, and so

$$t(Y) = Y^d - s_1(\vec{b})Y^{d-1} + s_2(\vec{b})Y^{d-2} - \cdots + (-1)^d s_d(\vec{b}) = \Pi_{i=1}^d (Y - b_i).$$

Hence $t(b_i) = 0$ for every i, proving Sublemma 4.5 and also Lemma 4.4. □

□

Now we are ready to prove the following claim, by induction on e.

Lemma 4.6. *For every $e \in \omega$, there exists a stage s_1 such that:*

- *$\mathcal{R}_e$ is never injured after stage s_1; and*
- *$(\forall s \geq s_1) q_{e,s}(X) = q_{e,s_1}(X)$; and*
- *either $\varphi_e(q_{e,s_1})\uparrow$ or else $\mathcal{R}_e$ is satisfied at every stage $\geq s_1$.*

Moreover, if $\varphi_e(q_{e,s_1})\downarrow = p(X)$, then for all $s \geq s_1$,

$$q_{e,s_1}(X) \in R_{F_s} \iff p(X) \notin S_{F_s},$$

and so $q_e(X) = \lim_s q_{e,s}(X)$ exists and $(q_e(X) \in R_F \iff p(X) \notin S_F)$.

Proof. We assume inductively that the lemma holds for all $e' < e$. Hence there is a stage after which we never again act to satisfy any of the (finitely many!) requirements of higher priority than $\mathcal{R}_e$. At the greatest stage s at which any higher-priority requirement did act, $q_{e,s}$ became undefined, as did the witness polynomials for all requirements of lower priority than $\mathcal{R}_e$. At the end of that stage, only finitely many $q_{e',s}$ were defined at all, and within finitely many more stages $q_{e,s}$ will come to be defined as well. At the stage s when it is defined, its Galois group over F_s is the symmetric group on its roots, so clearly F_s contains no root of $q_{e,s}$. Moreover, from then on, $q_{e,s}$ will never again become undefined – that is, $\mathcal{R}_e$ will never again be *injured* – since no higher-priority requirement ever acts again. (This is the reason for calling this proof a *finite-injury* construction.) We may therefore write $q_e(X)$ for this final polynomial assigned to $\mathcal{R}_e$. Of course, we do not claim that we can compute whether a given polynomial assigned to $\mathcal{R}_e$ at some stage is the final one or not; but our proof only requires that a final one exist.

Now if the computation $\varphi_e(q_e)$ diverges, then $\mathcal{R}_e$ will never again act, so the induction goes through without a hitch. Moreover, in this situation, the requirement $\mathcal{R}_e$ does turn out to be true, since in this case φ_e is a strictly partial function, hence cannot be an m-reduction.

Assume, therefore, that $\varphi_e(q_e)$ does converge to some $p(X)$, and let s_1 be the least number $> s_0$ such that the computation converges within s_1 steps. This number will be the s_1 required by the Lemma. Now no $\varphi_{e'}(q_{e'})$ with $e' < e$ ever converges at any stage $> s_0$, so the construction dictates that at stage s_1 we will act to satisfy $\mathcal{R}_e$. We consider the four possible cases from the construction, and show that in each of them, our action at stage s_1 ensures that $\mathcal{R}_e$ will hold of the field F we build. In all four cases, the construction declared $\mathcal{R}_e$ to be satisfied at the end of stage s_1, and therefore never again acts on behalf of $\mathcal{R}_e$, so part of the Lemma is immediate for this e and s_1. What we need to show is that, for all $s \geq s_1$, $(q_{e,s_1} \in R_{F_s} \iff p \notin S_{F_s})$. This will be done by induction on $s \geq s_1$.

Suppose first that we are in Case (1) at stage s_1. Then $\varphi_e(q_e) = p(X) \in S_{F_{s_1}}$, and so $p(X) \in S_{F_s}$ for all subsequent s as well. On the other hand, we keep $d_{e,s}$

equal to d_e, the degree of $q_e(X)$, at stage $s = s_1$ and all subsequent stages, and we injure all lower-priority requirements at stage s_1. Whenever $d_{e+1,s}$ is defined again at a stage $s > s_1$, the construction chooses it to be relatively prime to all prime numbers $\leq d_e$, hence relatively prime to d_e as well, and the same holds by induction for all $d_{i,s}$ with $i > e$ chosen at any stage $s > s_1$. Lemma 4.3 then shows that for all $s \geq s_1$, $[F_{s+1} : F_s]$ is relatively prime to d_e. Therefore, by Lemma 2.10, no root of $q_e(X)$ can lie in $F_{s+1} - F_s$, and so $q_e(X) \notin R_{F_s}$ for all $s \geq s_1$.

Now suppose that we are in Case (2) at stage s_1. Then we took the opposite strategy: now $q_e \in R_{F_{s_1}}$, since the construction set $F_{s_1} = F_{s_1-1}[r_1]$ for a root r_1 of $q_e(X)$. On the other hand, since we are in Case (2), $p(X)$ is irreducible over F_{s_1-1} and $F_{s_1} = F_{s_1-1}[r_1]$ contains no symmetric subfield L_i for $p(X)$. Therefore $p(X)$ stays irreducible over F_{s_1}, by Lemma 2.12. For stages after s_1, we note that d_{e,s_1} is chosen relatively prime to all primes $\leq |\mathrm{Gal}(L/F_{s_1})|$. Therefore, by Lemma 4.3, for all $s \geq s_1$ we have $[F_{s+1} : F_s]$ relatively prime to $|\mathrm{Gal}(L/F_{s_1})| = [L : F_{s_1}]$. Therefore, no element of L ever enters F after stage s_1, and so $p(X)$ remains irreducible over each F_{s+1}, as required.

If Case (3) applies at stage s_1, then F_{s_1} is the symmetric subfield L_i chosen in that case, so $p(X) \in S_{F_{s_1}} \subseteq S_{F_s}$ for all $s \geq s_1$. By the choice of L_i, no root of $q_e(X)$ lies in F_{s_1}, and since $d_{e,s} = d_{e,s_1}$ is the degree of q_e for all $s \geq s_1$, Lemma 4.3 shows again that no root of q_e enters F_{s+1} at any subsequent stage. Thus $q_e \notin R_{F_s}$ for all those stages.

Case (4) is similar: we again chose F_{s_1} to contain a symmetric subfield L_i, so that $p(X) \in S_{F_{s_1}} \subseteq S_{F_s}$ for all $s \geq s_1$. Moreover, we again made sure that no root of $q_e(X)$ lies in F_{s_1}, as follows. At the stage s_0 when q_e was first defined, we chose it so that $\mathrm{Gal}(K_{s_0}/F_{s_0}) \cong S_{d_e}$ was the symmetric group on the roots $\{r_1, \ldots, r_{d_e}\}$ of q_e, where K_{s_0} was the splitting field of q_e over F_{s_0}. Lemma 4.3 shows that this has remained true at all subsequent stages up through stage $s_1 - 1$. So there is an element of the subgroup G_{jk} chosen in Case (4) which interchanges r_j with r_k and permutes the other roots of q_e cyclically. (We made sure to choose $d_e > 3$, so there are at least two more roots to permute!) Specifically, there is such an element in $\mathrm{Gal}(K_{s_1-1}/F_{s_1-1}) \cong S_{d_e}$, and by normality it extends to an element of $\mathrm{Gal}(L/F_{s_1-1})$. This shows that G_{jk} does not fix any single root of $q_e(X)$, and therefore F_{s_1}, the fixed field of G_{jk}, contains no such root. So $q_e \notin R_{F_{s_1}}$, and by Lemma 4.3 we keep $q_e \notin R_{F_s}$ for all $s \geq s_1$, completing the proof of Lemma 4.6. □

It now follows that F is a field, clearly computable and also clearly algebraic over $\mathbb{Q}$, such that no partial computable function φ_e can be an m-reduction from R_F to S_F. Theorem 4.2 is proven. □

5. Further Ideas

The most natural further reducibility to consider between R_F and S_F would be weak truth-table reducibility. By definition, for $A, B \subseteq \omega$, A is *weak truth-table reducible to* B, written $A \leq_{wtt} B$, if there exists an oracle Turing functional Φ_e and a total computable function f such that Φ_e^B computes the characteristic function of A and for all $x \in \omega$, the computation $\Phi_e^B(x)$ asks its oracle questions only about the membership in B of elements $< f(x)$. (So, for all x, $\Phi_e^{B\restriction f(x)}(x) \downarrow = \chi_A(x)$.) Details appear in [4] and in [15, V.2.16]. We conjecture that the proof of Theorem 4.2 can be adapted to build a computable algebraic field F such that $R_F \not\leq_{wtt} S_F$.

The idea is to adjust the requirements to say

$$\mathcal{R}_{e,i}: \text{If } \Phi_e^{S_F} \text{ and } \varphi_i \text{ are total and } \forall q(\Phi_e^{S_F}(q) = 1 \iff q \in R_F), \text{ then}$$
$$(\exists q \in F[X])(\exists y \geq \varphi_i(q))[\Phi_e^{S_F}(q) \text{ asks an oracle question about } y].$$

On the other hand, m-reducibility implies wtt-reducibility, and so Theorem 3.1 shows that $S_F \leq_{wtt} R_F$ for every computable algebraic field F. Therefore, if the above conjecture holds, wtt-reducibility also distinguishes the splitting set from the root set.

It would also be natural to ask how these reducibilities relate the splitting set and the root set to the image $g(F)$ of a computable algebraic field F under a Rabin embedding. That image too is naturally a computably enumerable set and closely tied to the splitting set and root set by Corollary 2.6, but there is no immediate reason to hope for m-reducibility or 1-reducibility among them. We leave these investigations for another time.

Finally, these questions can also be asked about computable fields which fail to be algebraic. Computable fields of finite transcendence degree over $\mathbb{Q}$ are covered by Corollary 3.2; the main difference is that some uniformity is lost, since one needs to assume knowledge of a finite transcendence basis for the field in order to compute the 1-reduction from S_F to R_F. We refer the reader to the final section of [11] for a discussion of these issues and their applicability to fields of positive characteristic (where separability becomes an issue). The proof of Theorem 4.2 could easily be adapted to make the field F have arbitrary finite transcendence degree over $\mathbb{Q}$, although extending a negative result in that way seems almost redundant. On the other hand, for positive characteristic, Theorem 4.2 appears more challenging, even in the algebraic case, since adapting our proof would require some version of Theorem 2.15 in that characteristic.

For computable fields of infinite transcendence degree, the situation is completely different. Rabin's Theorem still applies to such fields, but there is no obvious reason to expect the reducibility from Theorem 3.1 to carry over to that case, except in the specific case of a field of characteristic 0 with a computable (infinite) transcendence basis, which was covered in Corollary 3.2. In positive characteristic, separability issues again prevent us from applying the techniques used here, and moreover, Metakides and Nerode proved in [8] that a computable field, even of characteristic 0, can fail to have a computable transcendence basis. We offer no conjectures about the general case.

References

1. H.M. Edwards, *Galois Theory* (New York: Springer-Verlag, 1984).
2. M.D. Fried & M. Jarden, *Field Arithmetic* (Berlin: Springer-Verlag, 1986).
3. R.M. Friedberg, Two recursively enumerable sets of incomparable degrees of unsolvability, *Proc. Nat. Acad. Sci. (USA)* **43** (1957) 236–238.
4. R.M. Friedberg & H. Rogers, Jr., Reducibility and completeness for sets of integers, *Z. Math. Logik Grundlagen Math.* **5** (1959) 117-125.
5. A. Frohlich & J.C. Shepherdson, Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407-432.
6. N. Jacobson, *Basic Algebra I* (New York: W.H. Freeman & Co., 1985).
7. L. Kronecker, Grundzüge einer arithmetischen Theorie der algebraischen Größen, *J. f. Math.* **92** (1882), 1-122.
8. G. Metakides & A. Nerode, Effective content of field theory, *Annals of Mathematical Logic* **17** (1979), 289-320.

9. R.G. Miller, Computable fields and Galois theory, *Notices of the American Mathematical Society* **55** (August 2008) 7, 798-807.
10. R.G. Miller, Computability and differential fields: a tutorial, to appear in *Differential Algebra and Related Topics: Proceedings of the Second International Workshop*, eds. L. Guo & W. Sit, to appear. Also available at `qcpages.qc.cuny.edu/~rmiller/research.html`.
11. R.G. Miller, ***d***-Computable categoricity for algebraic fields, to appear.
12. A.A. Muchnik, On the unsolvability of the problem of reducibility in the theory of algorithms, *Dokl. Akad. Nauk SSSR, N.S.* **109** (1956) 194–197 (Russian).
13. M. Rabin, Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341-360.
14. H. Rogers, Jr., *Theory of Recursive Functions and Effective Computability* (New York: McGraw-Hill Book Co., 1967).
15. R.I. Soare, *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).
16. V. Stoltenberg-Hansen & J.V. Tucker, Computable Rings and Fields, in *Handbook of Computability Theory*, ed. E.R. Griffor (Amsterdam: Elsevier, 1999), 363-447.
17. B.L. van der Waerden, *Algebra*, volume I, trans. F. Blum & J.R. Schulenberger (New York: Springer-Verlag, 1970 hardcover, 2003 softcover).
18. H. Völklein, *Groups as Galois Groups: An Introduction* (Cambridge: Cambridge University Press, 1996).

Department of Mathematics, Queens College – C.U.N.Y., 65-30 Kissena Blvd., Flushing, New York 11367 U.S.A.; Ph.D. Programs in Computer Science & Mathematics, The Graduate Center of C.U.N.Y., 365 Fifth Avenue, New York, New York 10016 U.S.A.

E-mail address: `Russell.Miller@qc.cuny.edu`