Generic Algebraic Fields

Russell Miller

Queens College & CUNY Graduate Center

Mathematisches Forschungsinstitut Oberwolfach
27 April 2021 (remote)

(Joint work with Kirsten Eisenträger, Caleb Springer, and Linda Westrick.)

HTP: Hilbert's Tenth Problem

Definition

For a ring R, *Hilbert's Tenth Problem for R* is the set

$$HTP(R) = \{f \in R[X_0, X_1, \ldots] : (\exists \vec{a} \in R^{<\omega})\ f(a_0, \ldots, a_n) = 0\}$$

of all polynomials (in several variables) with solutions in R.

So $HTP(R)$ is computably enumerable from the atomic diagram $\Delta(R)$.

HTP: Hilbert's Tenth Problem

Definition

For a ring R, *Hilbert's Tenth Problem for R* is the set

$$HTP(R) = \{f \in R[X_0, X_1, \ldots] : (\exists \vec{a} \in R^{<\omega})\ f(a_0, \ldots, a_n) = 0\}$$

of all polynomials (in several variables) with solutions in R.

So $HTP(R)$ is computably enumerable from the atomic diagram $\Delta(R)$.

Hilbert's original formulation in 1900 demanded a decision procedure for $HTP(\mathbb{Z})$.

Theorem (Matiyasevich-Davis-Putnam-Robinson, 1970)

$HTP(\mathbb{Z})$ is undecidable: indeed, $HTP(\mathbb{Z}) \equiv_1 \emptyset'$.

MDPR showed that $\emptyset'$ is *diophantine* in $\mathbb{Z}$, i.e., $\exists$-definable there.

Hilbert's Tenth Problem for $\mathbb{Q}$

Major Open Problem

The Turing degree of $HTP(\mathbb{Q})$ is unknown! All Σ_1 degrees are possible.

Hilbert's Tenth Problem for $\mathbb{Q}$

Major Open Problem

The Turing degree of $HTP(\mathbb{Q})$ is unknown! All Σ_1 degrees are possible.

If $\mathbb{Z}$ is existentially definable in the field $\mathbb{Q}$, then we would have $\mathrm{HTP}(\mathbb{Q}) \equiv \mathrm{HTP}(\mathbb{Z})$. At present the best-known definition of $\mathbb{Z}$ in $\mathbb{Q}$ is purely universal (Koenigsmann, 2016). The same situation applies in number fields, i.e., finite algebraic extensions of $\mathbb{Q}$.

Here we will go upwards from $\mathbb{Q}$, considering algebraic field extensions $E \supseteq \mathbb{Q}$ – or equivalently, subfields of the algebraic closure $\overline{\mathbb{Q}}$ – more generally.

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

$\sqrt{2}$

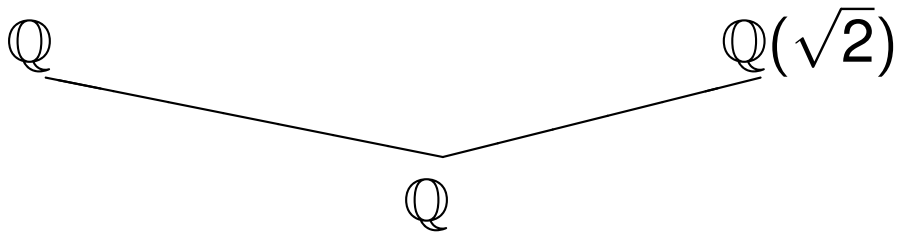

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

$\sqrt{3}$

$\sqrt{2}$

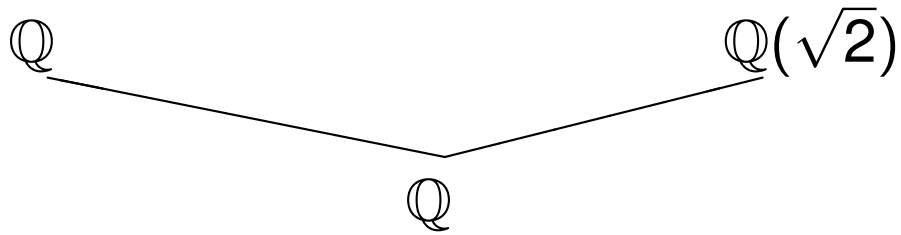

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

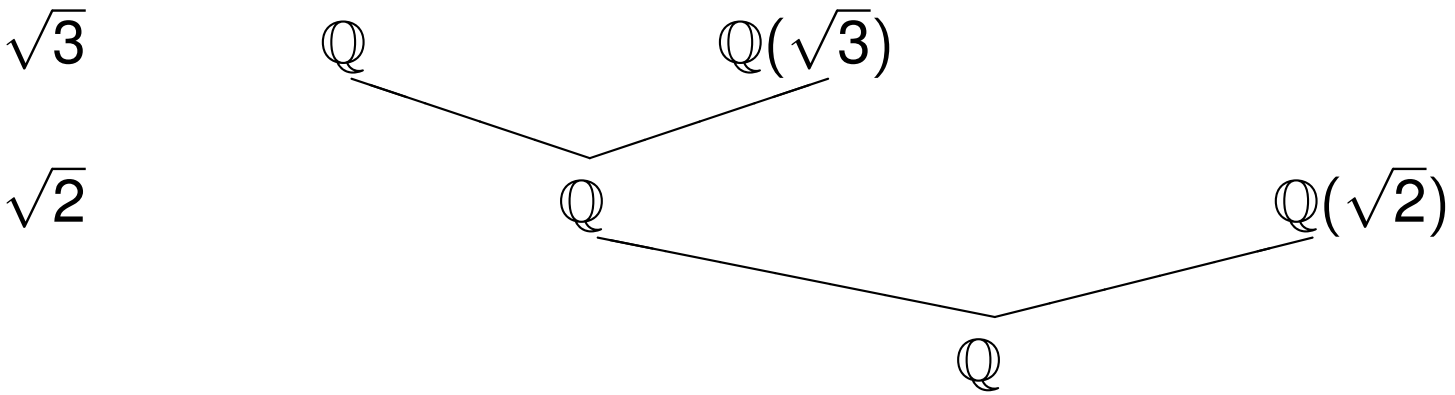

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

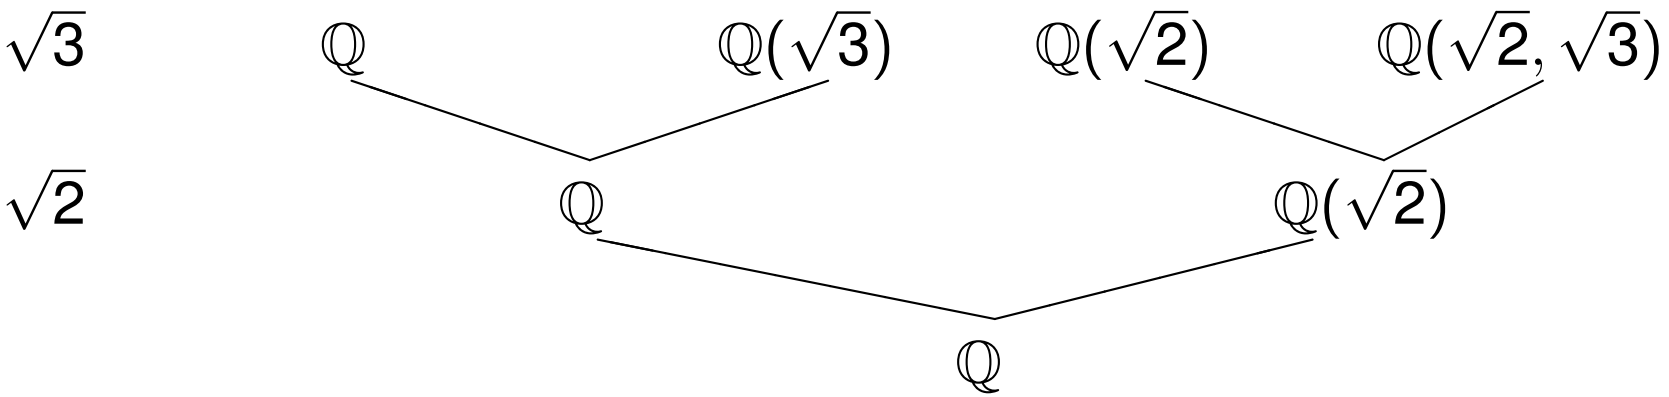

The space $\mathrm{Sub}(\overline{\mathbb{Q}})$ of all subfields of $\overline{\mathbb{Q}}$

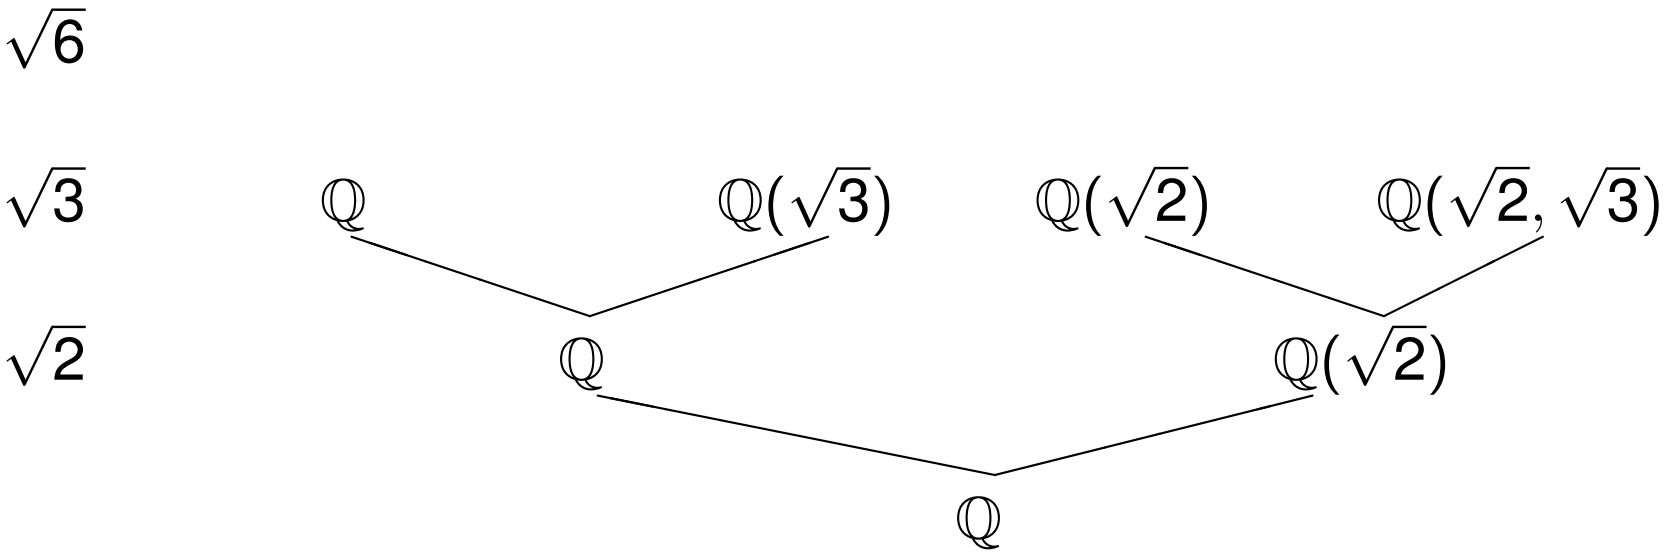

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

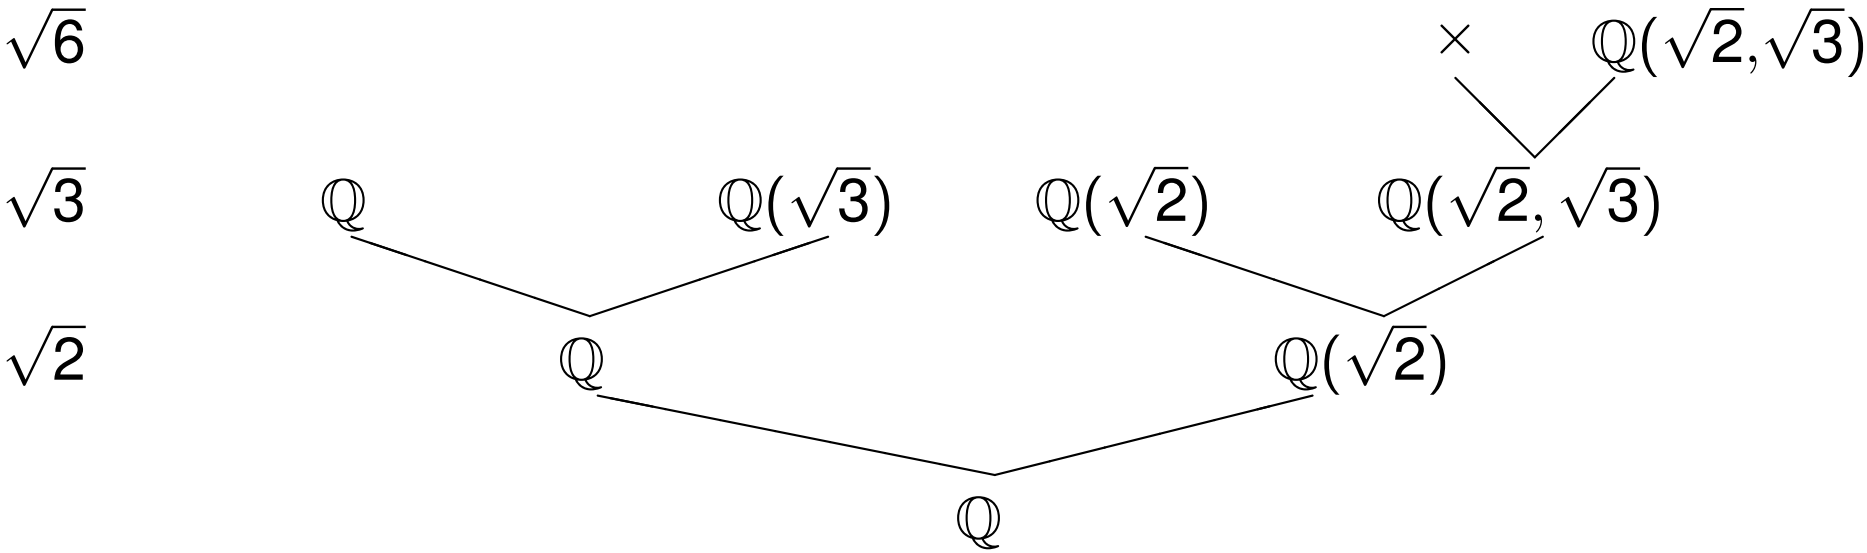

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

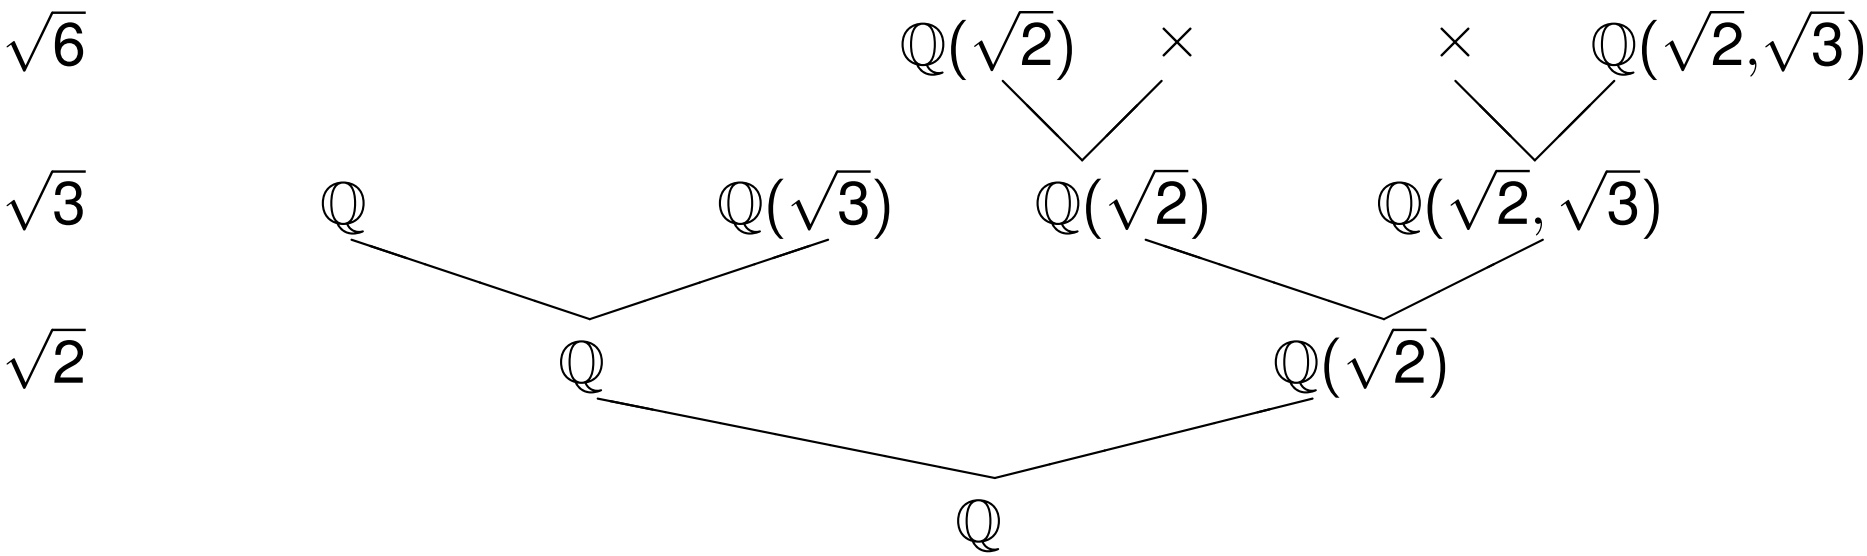

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

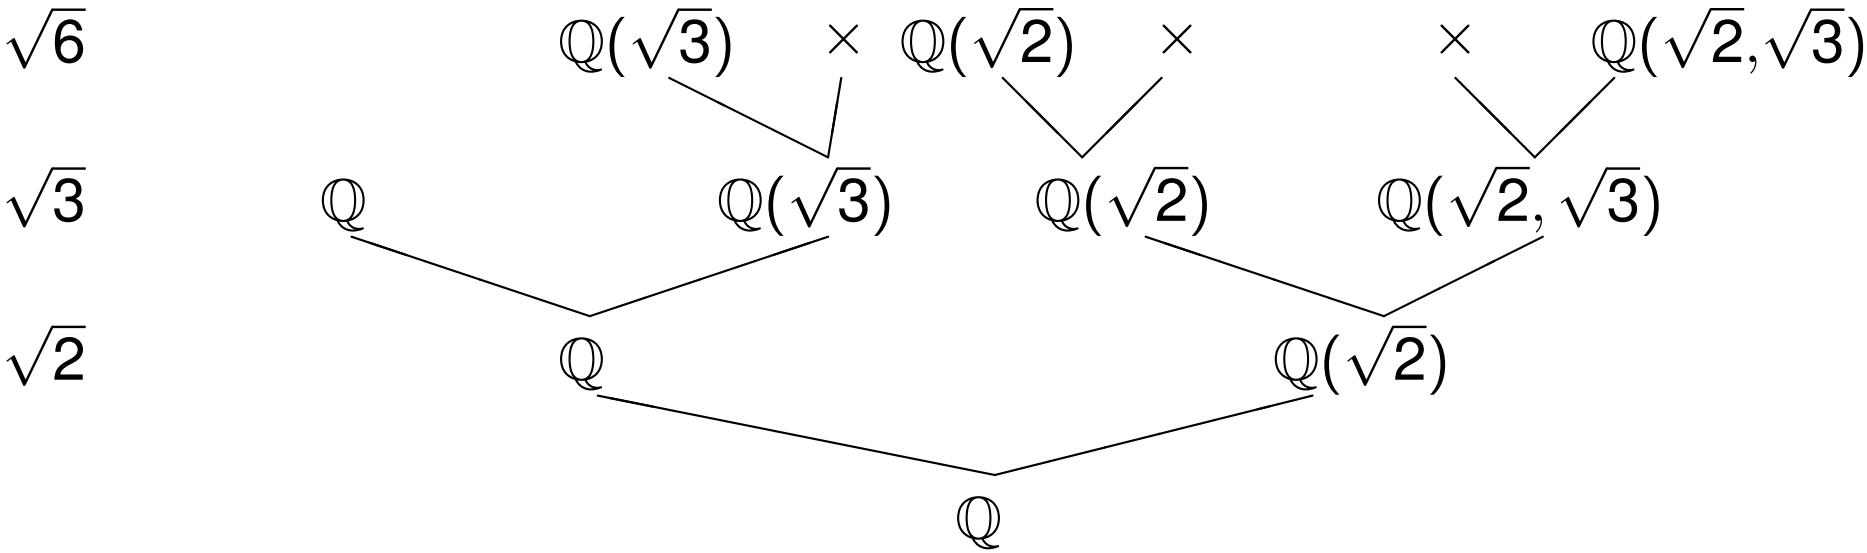

The space Sub($\overline{\mathbb{Q}}$) of all subfields of $\overline{\mathbb{Q}}$

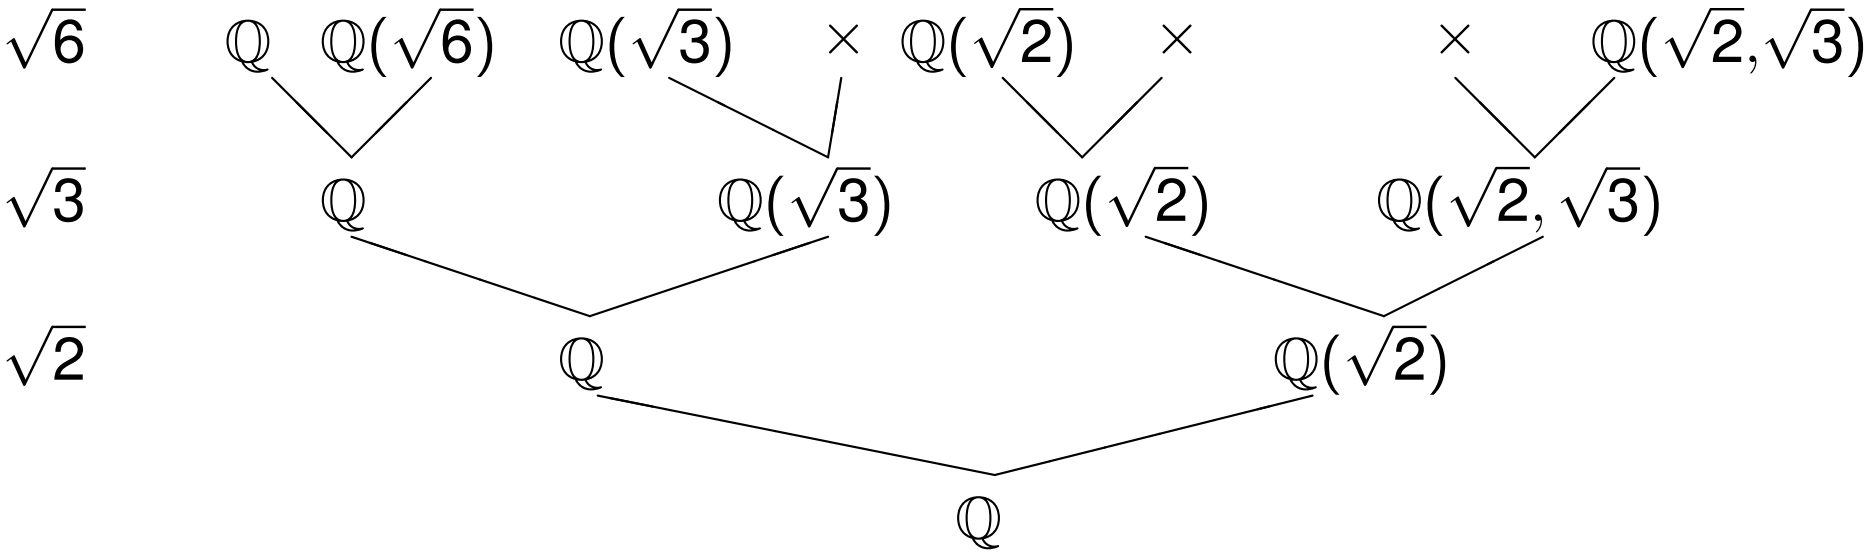

The space $\mathrm{Sub}(\overline{\mathbb{Q}})$ of all subfields of $\overline{\mathbb{Q}}$

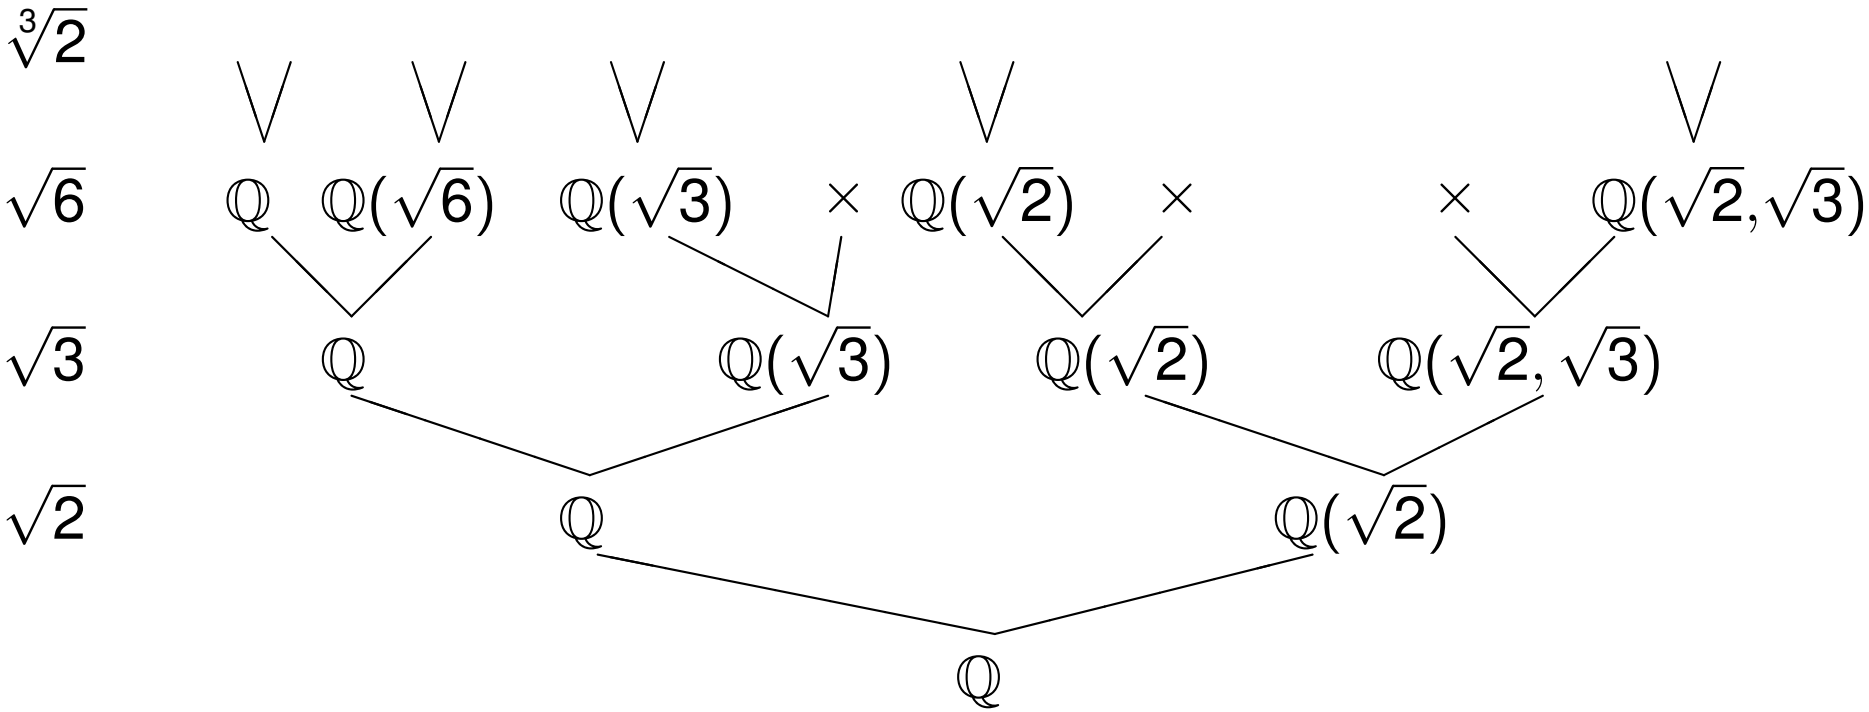

The nodes $\times$ are impossible: if we have ruled out $\sqrt{2}$, then E cannot contain both $\sqrt{3}$ and $\sqrt{6}$. But we still get a decidable subtree of $2^{<\omega}$, with no terminal nodes and no isolated paths. So the set of paths through it, which is the set of all subfields of $\overline{\mathbb{Q}}$, is homeomorphic to 2^{ω}.

Our question

Determining HTP($\mathbb{Q}$) is too hard.

Our question

Determining HTP($\mathbb{Q}$) is too hard.

Determining HTP($\overline{\mathbb{Q}}$) is too easy.

Our question

Determining HTP($\mathbb{Q}$) is too hard.

Determining HTP($\overline{\mathbb{Q}}$) is too easy.

We ask: what is the general situation for (presentations of) subfields F of $\overline{\mathbb{Q}}$?

"General situation" could refer to either measure theory, or Baire category. No canonical measure on this space **Sub**($\overline{\mathbb{Q}}$) is yet known, so we use Baire category, considering *comeager* subsets of the space to be "large." Since the set of all generic subsets of ω is comeager in Cantor space 2^ω, we are naturally led to examine generic sets.

Given a polynomial $f \in \overline{\mathbb{Q}}[Y_1, \ldots, Y_m]$, which generic fields contain a solution to $f = 0$? We address this mainly for $m > 1$, as for $m = 1$ there are only finitely many solutions in $\overline{\mathbb{Q}}$.

Topology of $\mathbf{Sub}(\overline{\mathbb{Q}})$

The clopen subsets of $\mathbf{Sub}(\overline{\mathbb{Q}})$ are those of the form:

$$\mathcal{U}_{\overline{a};\overline{b}} = \{E \in \mathbf{Sub}(\overline{\mathbb{Q}}) : \mathbb{Q}(\overline{a}) \subseteq E \ \& \ E \cap \overline{b} = \emptyset\}.$$

These form a basis for the topology on $\mathbf{Sub}(\overline{\mathbb{Q}})$, giving us a notion of density. Moreover, since this topological space is computable, the usual notions of genericity make sense.

A *weakly* 1*-generic field* is a field that lies in every dense c.e. open set, i.e., in every dense, computably enumerable union of these basis sets. For most of the work here, weak 1-genericity will suffice, so "generic" will mean "weakly-1-generic."

An *n-generic field* lies in every dense $\emptyset^{(n)}$-computable union of basis sets.

Forcing in Sub($\overline{\mathbb{Q}}$)

Each $(\overline{a}; \overline{b})$ from $\overline{\mathbb{Q}}^{<\omega}$ is a *condition*. For both existential and universal sentences φ, we say that $(\overline{a}; \overline{b}) \Vdash \varphi$ if $\{E : E \models \varphi\}$ is dense in $\mathcal{U}_{\overline{a},\overline{b}}$.

When φ is $\forall \vec{Y} \psi(\vec{Y})$, for instance $(\forall \vec{Y})\ f(\vec{Y}) \neq 0$,

$$(\overline{a}; \overline{b}) \Vdash \forall \vec{Y} \psi(\vec{Y}) \iff \text{all } E \in \mathcal{U}_{\overline{a};\overline{b}} \text{ satisfy } \forall \vec{Y} \psi(\vec{Y}).$$

When φ is $\exists \vec{Y} \psi(\vec{Y})$, for instance $(\exists \vec{Y})\ f(\vec{Y}) = 0$,

$$(\overline{a}; \overline{b}) \Vdash \exists \vec{Y} \psi(\vec{Y}) \iff \text{density holds.}$$

In each case, every generic $E \in \mathcal{U}_{\overline{a};\overline{b}}$ must realize the sentence forced.

Forcing in Sub($\overline{\mathbb{Q}}$)

Each $(\overline{a}; \overline{b})$ from $\overline{\mathbb{Q}}^{<\omega}$ is a *condition*. For both existential and universal sentences φ, we say that $(\overline{a}; \overline{b}) \Vdash \varphi$ if $\{E : E \models \varphi\}$ is dense in $\mathcal{U}_{\overline{a},\overline{b}}$.

When φ is $\forall \vec{Y} \psi(\vec{Y})$, for instance $(\forall \vec{Y})\ f(\vec{Y}) \neq 0$,

$$(\overline{a}; \overline{b}) \Vdash \forall \vec{Y} \psi(\vec{Y}) \iff \text{all } E \in \mathcal{U}_{\overline{a};\overline{b}} \text{ satisfy } \forall \vec{Y} \psi(\vec{Y}).$$

When φ is $\exists \vec{Y} \psi(\vec{Y})$, for instance $(\exists \vec{Y})\ f(\vec{Y}) = 0$,

$$(\overline{a}; \overline{b}) \Vdash \exists \vec{Y} \psi(\vec{Y}) \iff \text{density holds.}$$

In each case, every generic $E \in \mathcal{U}_{\overline{a};\overline{b}}$ must realize the sentence forced.

Intuition: working within $\overline{\mathbb{Q}}$ (with $m > 1$), one imagines that there should be some way to find a solution to $f(\vec{Y}) = 0$ in some field in $\mathcal{U}_{\overline{a};\overline{b}}$.

Density of solutions to $f = 0$ can fail!

Example where density fails: consider the sentence

$$(\exists Y_1 \exists Y_2)\ \ Y_1^2 - 2Y_2^2 = 0 \neq Y_1 Y_2.$$

Let $(\overline{a}; \overline{b})$ be $(\lambda; \sqrt{2})$, so $\mathcal{U}_{\overline{a};\overline{b}}$ contains all fields E that omit both $\pm\sqrt{2}$.

With the solution $(0, 0)$ ruled out, each nonzero solution $(y_1, y_2) \in \overline{\mathbb{Q}}^2$ has

$$\left(\frac{y_1}{y_2}\right)^2 = \frac{y_1^2}{y_2^2} = \frac{2y_2^2}{y_2^2} = 2.$$

So no field in $\mathcal{U}_{\overline{a},\overline{b}}$ can contain any such solution, and

$$(\overline{a}; \overline{b}) \Vdash (\forall Y_1, Y_2)\ \neg[Y_1^2 - 2Y_2^2 = 0 \neq Y_1 Y_2].$$

Things happen for a reason

Density of solutions to $f = 0$ can fail if every solution generates an element forbidden by the condition $(\overline{a}; \overline{b})$.
In our example, there was a reason: the formula $\frac{Y_1}{Y_2}$ produces the forbidden $\sqrt{2}$ from every nonzero solution to $f(Y_1, Y_2) = 0$.

If we find such a formula, we know that density of solutions fails. But is there always such a reason?

Things happen for a reason

Density of solutions to $f = 0$ can fail if every solution generates an element forbidden by the condition $(\overline{a}; \overline{b})$.
In our example, there was a reason: the formula $\frac{Y_1}{Y_2}$ produces the forbidden $\sqrt{2}$ from every nonzero solution to $f(Y_1, Y_2) = 0$.

If we find such a formula, we know that density of solutions fails. But is there always such a reason?

Theorem (cf. Stichtenoth, Cor. III.6.7)

A polynomial f, irreducible in $F[\vec{Y}]$, is absolutely irreducible $\iff$ F is algebraically closed within the function field of f over F.

Here the *function field* is the fraction field of the domain $F[\vec{Y}]/(f)$.
Absolutely irreducible means that f remains irreducible over $\overline{F}$.
In our example this fails: $Y_1^2 - 2Y_2^2 = (Y_1 + Y_2\sqrt{2})(Y_1 - Y_2\sqrt{2})$ in $\overline{\mathbb{Q}}$.

Refining the corollary

We want to consider the fields $F = \mathbb{Q}(\overline{a})$ and $K = F(\overline{b})$, for a condition $(\overline{a}; \overline{b})$.

Theorem

For K/F a finite Galois extension, and f irreducible in $F[\vec{Y}]$:
f remains irreducible in $K[\vec{Y}]$ iff the function field of f over F intersects K only in F.

Now, **Case 1:** if $f(Y_1, \ldots, Y_m)$ remains irreducible over K, then the Hilbert Irreducibility Theorem will give $y_1, \ldots, y_{m-1} \in \mathbb{Q}$ such that $f(y_1, \ldots, y_{m-1}, Y_m)$ stays irreducible in $K[Y_m]$. In this case, each root $y_m \in \overline{\mathbb{Q}}$ has this as its minimal polynomial over K, hence does not generate any element of K when adjoined to F. Thus $F(y_m)$ is a field in $\mathcal{U}_{\overline{a};\overline{b}}$ containing a solution to $f = 0$.

The reason appears

Case 2: if $f(Y_1, \ldots, Y_m)$ becomes reducible over K, then by the theorem, some element $\frac{p(\vec{Y})+(f)}{q(\vec{Y})+(f)}$ of the function field lies in K but outside F. Breaking into finitely many cases, we can assume that in fact this element is one of the forbidden b_i's.
So this is the "reason": every solution $f(\vec{y}) = 0$ will have $\frac{p(\vec{y})+0}{q(\vec{y})+0} = b_i$.

The reason appears

Case 2: if $f(Y_1, \ldots, Y_m)$ becomes reducible over K, then by the theorem, some element $\frac{p(\vec{Y})+(f)}{q(\vec{Y})+(f)}$ of the function field lies in K but outside F. Breaking into finitely many cases, we can assume that in fact this element is one of the forbidden b_i's.
So this is the "reason": every solution $f(\vec{y}) = 0$ will have $\frac{p(\vec{y})+0}{q(\vec{y})+0} = b_i$.

But then how did the solution $(0, 0)$ to $Y_1^2 = 2Y_2^2$ survive?

Full answer

Solutions $f(\vec{y}) = 0$ may still be possible in fields in $\mathcal{U}_{\bar{a};\bar{b}}$, but the only way for $\vec{y}$ to avoid generating b_i (over F) is to have $q(\vec{y}) = 0$.

Now in the function ring $F[\vec{Y}]/(f)$, we may choose $q(\vec{Y}) + (f)$ so that Y_m has smaller degree in q than in f: use the Euclidean algorithm modulo (f). This is progress, reducing the question of solvability of $f = 0$ (in fields in $\mathcal{U}_{\bar{a};\bar{b}}$) to that of solvability of a polynomial q of lesser multidegree.

Details: we get a remainder r with $\deg_{Y_m}(r) < \deg_{Y_m}(q) < \deg_{Y_m}(f)$ and $f(\vec{Y}) \cdot c(Y_1, \ldots, Y_{m-1}) = q(\vec{Y}) \cdot b(Y_1, \ldots, Y_{m-1}) + r(\vec{Y})$.
Now, for tuples $\vec{y}$ with $K \cap F(\vec{y}) = F$,

$$f(\vec{y}) = 0 \iff [q(\vec{y}) = r(\vec{y}) = 0 \neq c(\vec{y}) \text{ or } f(\vec{y}) = c(\vec{y}) = 0].$$

The final step is to well-order existential sentences in such a way that the right side has lower rank than the original $\exists \vec{Y}\, f(\vec{Y}) = 0$.

Well-ordering existential sentences

We rank an existential sentence of the form

$$(\exists Y_1 \cdots \exists Y_m)\ \ f_1(\vec{Y}) = \cdots = f_n(\vec{Y}) = 0 \neq g(\vec{Y})$$

as follows (assuming $g \notin \mathcal{V}(f_1, \ldots, f_n)$):

- first according to m, the dimension of the ambient space;
- next, according to the Krull dimension of the variety $\mathcal{V}(f_1, \ldots, f_n)$;
- and finally according to the multidegrees of $f_1, \ldots, f_m$ (arranged in nonincreasing order by multidegree).

If $\alpha_1 \vee \cdots \vee \alpha_j$ and $\beta_1 \vee \cdots \vee \beta_j$ are disjunctions of formulas of this form, put each in nonincreasing order under the above ranking. To compare them, compare α_1 to β_1; if these have the same rank, go on to α_2 and β_2, etc.

This is a well-order, and our sentence with q and r above does indeed have lower rank than the original sentence $\exists \vec{Y}\ f = 0$.

Conclusions

Conclusion (after further details!)

There is an effective procedure that decides, for all existential formulas φ in the language of fields and for all tuples $\overline{a}; \overline{b}$ from a fixed computable presentation of $\overline{\mathbb{Q}}$, whether $(\overline{a}; \overline{b}) \Vdash \varphi$, and also whether $(\overline{a}; \overline{b}) \Vdash \neg\varphi$.

So, if we are given the atomic diagram $\Delta(E)$ of a generic subfield of $\overline{\mathbb{Q}}$, then HTP(E) will be decidable from $\{(\overline{a}; \overline{b}) : E \in \mathcal{U}_{\overline{a};\overline{b}}\}$.

Conclusions

Conclusion (after further details!)

There is an effective procedure that decides, for all existential formulas φ in the language of fields and for all tuples $\overline{a}; \overline{b}$ from a fixed computable presentation of $\overline{\mathbb{Q}}$, whether $(\overline{a}; \overline{b}) \Vdash \varphi$, and also whether $(\overline{a}; \overline{b}) \Vdash \neg\varphi$.

So, if we are given the atomic diagram $\Delta(E)$ of a generic subfield of $\overline{\mathbb{Q}}$, then HTP(E) will be decidable from $\{(\overline{a}; \overline{b}) : E \in \mathcal{U}_{\overline{a};\overline{b}}\}$.

Theorem (EMSW)

For each generic subfield E, HTP$(E) \equiv_T E$, where E is presented as a subset of $\overline{\mathbb{Q}}$.

The procedures above assumed $m > 1$. For polynomials $f \in E[Y_1]$, simply find all the roots of f in $\overline{\mathbb{Q}}$, and ask the E-oracle whether each of them lies in E.

Rabin's Theorem

So, for all generic algebraic fields E, HTP(E) is decidable relative to its presentation *as a subfield of* $\overline{\mathbb{Q}}$. But what happens if we have a presentation of E as a freestanding field (i.e., if we have the atomic diagram $\Delta(E)$)? This requires the famous theorem of Rabin (1960).

The *index* of E is the set

$$I_E = \{\text{irreducible } h \in \mathbb{Z}[X] : h(X) \text{ has a root in } E\}.$$

This is an invariant of the isomorphism type of E, and distinct for distinct isomorphism types.

From $\Delta(E)$ itself we can uniformly compute an embedding ϵ of E into $\overline{\mathbb{Q}}$, and thus enumerate the tuples $\overline{a}$ with $\mathbb{Q}(\overline{a}) \subseteq \epsilon(E)$. But the ability to enumerate the tuples $\overline{b}$ disjoint from $\epsilon(E)$ is equivalent to knowing I_E as well: $\epsilon(E) \equiv_T \Delta(E) \oplus I_E$. This is the essence of Rabin's Theorem.

Rabin's Theorem

So, for all generic algebraic fields E, HTP(E) is decidable relative to its presentation *as a subfield of* $\overline{\mathbb{Q}}$. But what happens if we have a presentation of E as a freestanding field (i.e., if we have the atomic diagram $\Delta(E)$)? This requires the famous theorem of Rabin (1960).

The *index* of E is the set

$$I_E = \{\text{irreducible } h \in \mathbb{Z}[X] : h(X) \text{ has a root in } E\}.$$

This is an invariant of the isomorphism type of E, and distinct for distinct isomorphism types.

From $\Delta(E)$ itself we can uniformly compute an embedding ϵ of E into $\overline{\mathbb{Q}}$, and thus enumerate the tuples $\overline{a}$ with $\mathbb{Q}(\overline{a}) \subseteq \epsilon(E)$. But the ability to enumerate the tuples $\overline{b}$ disjoint from $\epsilon(E)$ is equivalent to knowing I_E as well: $\epsilon(E) \equiv_T \Delta(E) \oplus I_E$. This is the essence of Rabin's Theorem.

Indices of generic fields

For generic algebraic fields, the index I_E is *not* computable from a freestanding presentation $\Delta(E)$ in general, even nonuniformly.

Theorem (M, 2020)

Every generic algebraic field has a presentation $\Delta(E)$ such that $I_E \not\leq_T \Delta(E)$. However, all presentations satisfy $(I_E)' \leq_T (\Delta(E))'$: I_E is always *low relative to* $\Delta(E)$.

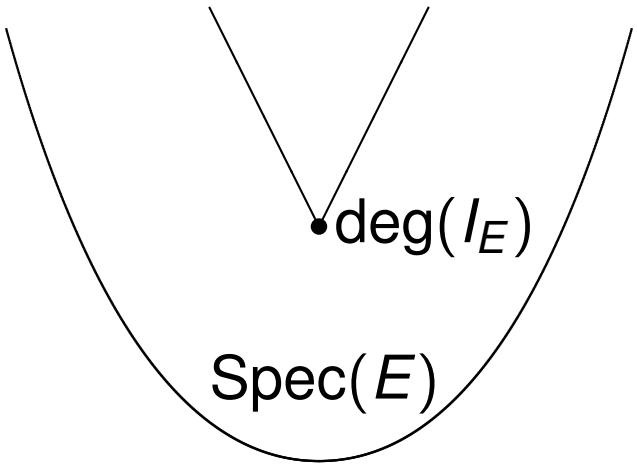

Freestanding presentations of generic fields

Theorem (EMSW)

Every generic algebraic field has a presentation $\Delta(E)$ for which $\mathrm{HTP}(E) \not\leq_T \Delta(E)$. However, all presentations of generic fields satisfy $\mathrm{HTP}(E) \equiv_T \mathrm{HTP}_1(E) \equiv_T I_E \oplus \Delta(E)$, and this degree is always low relative to $\Delta(E)$.

Here $\mathrm{HTP}_1(E)$ represents the restriction of $\mathrm{HTP}(E)$ to polynomials in a single variable. This set is also known as the *root set* R_E of E.

Addendum

We now have a procedure for deciding, for arbitrary f and any condition $(\overline{a}; \overline{b})$, whether $(\overline{a}; \overline{b})$ forces $f(\vec{Y}) = 0$ to have infinitely many solutions, and also whether it forces $f(\vec{Y}) = 0$ to have only finitely many solutions. The foregoing results are at the heart of this procedure. This means that for 2-generic subfields of $\overline{\mathbb{Q}}$, deciding whether polynomials have infinitely many solutions is exactly as hard as deciding whether they have solutions at all.

The question of forcing Σ^0_2 or Π^0_2 properties more generally remains open.