Constructive Bounds from Ultraproducts and Noetherianity

Henry Towsner
joint work with William Simmons

University of Pennsylvania

February 19, 2016

Mathematics is full of theorems which have computable bounds.

For example:

Theorem

For every d, n there is an r so that if K is an algebraically closed field, $S \subseteq K[X_1, \ldots, X_n]$ is a set of polynomials of total degree $\leq d$, and $g \in (S)$ where g has total degree $\leq d$ then $g = \sum_i c_i f_i$ where the c_i have degree $\leq r$.

Mathematics is full of theorems which have computable bounds.

For example:

Theorem

For every d, n there is an r so that if K is an algebraically closed field, $S \subseteq K[X_1, \ldots, X_n]$ is a set of polynomials of total degree $\leq d$, and $g \in (S)$ where g has total degree $\leq d$ then $g = \sum_i c_i f_i$ where the c_i have degree $\leq r$.

Theorem

For every d, n there is an r so that if K is an algebraically closed field, $S \subseteq K[X_1, \ldots, X_n]$ is a set of polynomials of total degree $\leq d$, if $gh \in (S)$ implies $g \in (S)$ or $h \in (S)$ whenever $\deg(gh) \leq r$, then (S) is either prime or contains 1.

One often sees in the literature assertions that:

- A proof shows that computable bounds exist but does not say what they are, or
- A proof shows that an algorithm terminates but does not give a running time.

One often sees in the literature assertions that:

- A proof shows that computable bounds exist but does not say what they are, or
- A proof shows that an algorithm terminates but does not give a running time.

On the other hand, the proof-theoretic perspective on bounds says:

> *The existence of computable bounds is a syntactic property of a statement. In particular, there are no non-effective proofs that computable bounds exist.*

Suppose we have fixed a suitable language (like the language of rings or the language of differential rings) and a collection of structures $\mathcal{K}$ (like fields, or polynomial extensions of algebraically closed fields).

Definition

A $\Pi_2^{\mathbb{N}}$ *statement* is a statement of the form

$$\forall n_1, \ldots, n_k \ \forall K \in \mathcal{K} \ \exists b_1, \ldots, b_d \ \ K \models \phi_{n_1,\ldots,n_k,b_1,\ldots,b_d}$$

where each $\phi_{n_1,\ldots,n_k,b_1,\ldots,b_d}$ is a first-order sentence in the language.

Note that we're mixing the "internal" language—the language of fields—with an "external" language where we quantify over things like the natural numbers.

In this example:

> *For every d, n there is an r so that if K is an algebraically closed field, $S \subseteq K[X_1, \ldots, X_n]$ is a set of polynomials of total degree $\leq d$, and $g \in (S)$ where g has total degree $\leq d$ then $g = \sum_i c_i f_i$ where the c_i have degree $\leq r$.*

once d, n, r are fixed, we can write down a formula of first-order logic expressing that the statement holds in K.

Note that the language of rings allows us to quantify over arbitrary polynomials (by quantifying over elements of $K[X_1, \ldots, X_n]$) or over polynomials of bounded degree $\leq n$ (by writing the polynomial as a sum of monomials and quantifying over coefficients), but does not allow us to directly quantify over the degree.

Definition

The $\Pi_2^{\mathbb{N}}$ statement

$$\forall n_1, \ldots, n_k \; \forall K \in \mathcal{K} \; \exists b_1, \ldots, b_d \;\; K \models \phi_{n_1,\ldots,n_k,b_1,\ldots,b_d}$$

has (computable) bounds if for each $K \in \mathcal{K}$, there is a computable function B_K such that, for all $n_1, \ldots, n_k$, the witnesses $b_1, \ldots, b_d$ are bounded by $B_K(n_1, \ldots, n_k)$.

Definition

The $\Pi_2^{\mathbb{N}}$ statement

$$\forall n_1, \ldots, n_k \ \forall K \in \mathcal{K} \ \exists b_1, \ldots, b_d \ K \models \phi_{n_1, \ldots, n_k, b_1, \ldots, b_d}$$

has (computable) bounds if for each $K \in \mathcal{K}$, there is a computable function B_K such that, for all $n_1, \ldots, n_k$, the witnesses $b_1, \ldots, b_d$ are bounded by $B_K(n_1, \ldots, n_k)$.

The statement has *uniform (computable) bounds* if the bound B_K is just B and does not depend on the choice of K.

Some observations about $\Pi_2^{\mathbb{N}}$ statements:

- Under reasonable assumptions on $\mathcal{K}$, every provable $\Pi_2^{\mathbb{N}}$ statement has computable bounds.

Some observations about $\Pi_2^{\mathbb{N}}$ statements:

- Under reasonable assumptions on $\mathcal{K}$, every provable $\Pi_2^{\mathbb{N}}$ statement has computable bounds.
- If $\mathcal{K}$ is closed under ultraproducts then every provable $\Pi_2^{\mathbb{N}}$ statement has uniform bounds.

Some observations about $\Pi_2^{\mathbb{N}}$ statements:

- Under reasonable assumptions on $\mathcal{K}$, every provable $\Pi_2^{\mathbb{N}}$ statement has computable bounds.
- If $\mathcal{K}$ is closed under ultraproducts then every provable $\Pi_2^{\mathbb{N}}$ statement has uniform bounds.
- The proofs of these statements are constructive.

Some observations about $\Pi_2^{\mathbb{N}}$ statements:

- Under reasonable assumptions on $\mathcal{K}$, every provable $\Pi_2^{\mathbb{N}}$ statement has computable bounds.
- If $\mathcal{K}$ is closed under ultraproducts then every provable $\Pi_2^{\mathbb{N}}$ statement has uniform bounds.
- The proofs of these statements are constructive.

Theorem

Suppose we have a proof that a $\Pi_2^{\mathbb{N}}$ statement holds in some reasonable class of structures $\mathcal{K}$. Then we can extract explicit bounds from this proof.

Statements which are not $\Pi_2^{\mathbb{N}}$ typically do not have computable bounds (although there are exceptions).

For example, the Ritt problem essentially asks whether a certain $\Pi_3^{\mathbb{N}}$ statement has computable bounds.

These results come out of foundational proof theory techniques, going back to Gentzen.

Formally, if one takes a seemingly non-constructive proof of a $\Pi_2^{\mathbb{N}}$ statement and formalizes it in a theory (like Peano arithmetic, or the theories of reverse math), various results tell us how to translate the proof into a constructive one.

These results come out of foundational proof theory techniques, going back to Gentzen.

Formally, if one takes a seemingly non-constructive proof of a $\Pi_2^{\mathbb{N}}$ statement and formalizes it in a theory (like Peano arithmetic, or the theories of reverse math), various results tell us how to translate the proof into a constructive one.

The original proofs, via cut-elimination, are quite slow (the process of transforming a proof is itself often too slow to carry out in practice).

But newer proofs via the "functional interpretation" actually give a linear time transformation of a non-constructive proof into a constructive proof. The cost is that the new proof is in a theory of higher type functionals.

What sort of bounds can we hope for by this kind of extraction?

Definition

The *fast-growing hierarchy* of functions is given by transfinite recursion:

- $f_0(n) = n + 1$,
- $f_{\alpha+1}(n) = f_\alpha^n(n)$,
- $f_\lambda(n) = f_{\lambda[n]}(n)$ (where $\lambda[n]$ is a canonical sequence of ordinals approaching λ from below).

What sort of bounds can we hope for by this kind of extraction?

Definition

The *fast-growing hierarchy* of functions is given by transfinite recursion:

- $f_0(n) = n + 1$,
- $f_{\alpha+1}(n) = f_\alpha^n(n)$,
- $f_\lambda(n) = f_{\lambda[n]}(n)$ (where $\lambda[n]$ is a canonical sequence of ordinals approaching λ from below).

f_ω has rate of growth similar to the Ackermann function.

A reasonable theory T has a related ordinal $o(T)$ called its *proof-theoretic ordinal*. For example, the proof-theoretic ordinal of Peano arithmetic is $\epsilon_0 = \omega^{\omega^{\omega^{\cdots}}}$.

Theorem

Suppose we prove a $\Pi_2^{\mathbb{N}}$ statement in T. Then the statement is bounded by some f_α with $\omega^\alpha < o(T)$.

In practice, however, very few proofs achieve these worst-case bounds. Actual examples show that when a theorem is proven in Peano arithmetic, the bounds obtained by these methods are frequently around the Ackermann function.

Many examples actually end up substantially lower, on the order of multiple exponentials or towers of exponentials.

In practice, however, very few proofs achieve these worst-case bounds. Actual examples show that when a theorem is proven in Peano arithmetic, the bounds obtained by these methods are frequently around the Ackermann function.

Many examples actually end up substantially lower, on the order of multiple exponentials or towers of exponentials.

This isn't surprising: the bounds obtained by extraction are exactly those "hidden in the original proof", so they should often be similar to what one gets by writing down bounds from constructive but unoptimized proofs.

In practice, however, very few proofs achieve these worst-case bounds. Actual examples show that when a theorem is proven in Peano arithmetic, the bounds obtained by these methods are frequently around the Ackermann function.

Many examples actually end up substantially lower, on the order of multiple exponentials or towers of exponentials.

This isn't surprising: the bounds obtained by extraction are exactly those "hidden in the original proof", so they should often be similar to what one gets by writing down bounds from constructive but unoptimized proofs.

On the other hand, hand-optimization will essentially always get improvements. Often a different proof describes a fundamentally better algorithm, and gives a better order of magnitude.

If $\Pi_2^{\mathbb{N}}$ statements are essentially constructive, where do nonconstructive proofs of $\Pi_2^{\mathbb{N}}$ statements come from?

One common source is that they go through intermediate steps which fail to be $\Pi_2^{\mathbb{N}}$ and which are genuinely nonconstructive. In algebraic contexts, one standard example is Noetherianity.

If $\Pi_2^{\mathbb{N}}$ statements are essentially constructive, where do nonconstructive proofs of $\Pi_2^{\mathbb{N}}$ statements come from?

One common source is that they go through intermediate steps which fail to be $\Pi_2^{\mathbb{N}}$ and which are genuinely nonconstructive. In algebraic contexts, one standard example is Noetherianity.

Theorem

If $I_1 \subseteq I_2 \subseteq \cdots$ is an increasing sequence of ideals in $K[X_1, \ldots, X_n]$ then there is an N such that for every $m \geq N$, $I_m \subseteq I_N$.

This is $\Pi_3^{\mathbb{N}}$:

$$\forall \langle I_n \rangle_n \, \exists N \, \forall m \geq N \; I_m \subseteq I_N.$$

If $\Pi^{\mathbb{N}}_2$ statements are essentially constructive, where do nonconstructive proofs of $\Pi^{\mathbb{N}}_2$ statements come from?

One common source is that they go through intermediate steps which fail to be $\Pi^{\mathbb{N}}_2$ and which are genuinely nonconstructive. In algebraic contexts, one standard example is Noetherianity.

Theorem

If $I_1 \subseteq I_2 \subseteq \cdots$ is an increasing sequence of ideals in $K[X_1, \ldots, X_n]$ then there is an N such that for every $m \geq N$, $I_m \subseteq I_N$.

This is $\Pi^{\mathbb{N}}_3$:

$$\forall \langle I_n \rangle_n \ \exists N \ \forall m \geq N \ I_m \subseteq I_N.$$

This is genuinely non-constructive—if we had a function calculating N from (a computable description of) the sequence $\langle I_n \rangle$ then we could compute the halting problem.

The *functional interpretation* tells us that we should instead look at a notion of *local Noetherianity*:

Theorem

If $I_1 \subseteq I_2 \subseteq \cdots$ is a computable increasing sequence of ideals in $K[X_1, \ldots, X_n]$ so that I_i is finitely generated by polynomials of degree $\leq D(i)$ and $F : \mathbb{N} \to \mathbb{N}$ is a function then there is an N, computable from D and F and $\langle I_n \rangle$, such that $I_{F(N)} \subseteq I_N$.

The *functional interpretation* tells us that we should instead look at a notion of *local Noetherianity*:

Theorem

If $I_1 \subseteq I_2 \subseteq \cdots$ *is a computable increasing sequence of ideals in* $K[X_1, \ldots, X_n]$ *so that* I_i *is finitely generated by polynomials of degree* $\leq D(i)$ *and* $F : \mathbb{N} \to \mathbb{N}$ *is a function then there is an* N*, computable from* D *and* F *and* $\langle I_n \rangle$*, such that* $I_{F(N)} \subseteq I_N$*.*

Furthermore, any $\Pi_2^{\mathbb{N}}$ *statement which follows from Noetherianity also follows from local Noetherianity, and bounds on local Noetherianity suffice to give bounds on the* $\Pi_2^{\mathbb{N}}$ *statement.*

The *functional interpretation* tells us that we should instead look at a notion of *local Noetherianity*:

Theorem

If $I_1 \subseteq I_2 \subseteq \cdots$ is a computable increasing sequence of ideals in $K[X_1, \ldots, X_n]$ so that I_i is finitely generated by polynomials of degree $\leq D(i)$ and $F : \mathbb{N} \to \mathbb{N}$ is a function then there is an N, computable from D and F and $\langle I_n \rangle$, such that $I_{F(N)} \subseteq I_N$.

Furthermore, any $\Pi_2^{\mathbb{N}}$ statement which follows from Noetherianity also follows from local Noetherianity, and bounds on local Noetherianity suffice to give bounds on the $\Pi_2^{\mathbb{N}}$ statement.

It is not a coincidence that local Noetherianity is itself a $\Pi_2^{\mathbb{N}}$ sentence (where we allow the data to include the functions D and F).

For example, recall Buchberger's algorithm for computing a Gröbner basis: given the set of polynomials F in $K[X_1, \ldots, X_n]$,

1. Set $T_0 := S$.
2. Given T_i:
   - Calculate a polynomial $s_{f,g} = \frac{a}{L_f} f - \frac{a}{L_g} g$ for some $f, g \in T_i$ so that $s_{f,g} \notin (T_i)$.
   - Reduce $s_{f,g}$ relative to T_i.
   - Set $T_{i+1} = T_i \cup \{s_{f,g}\}$.
3. If T_{i+1} is a Gröbner basis, terminate. Otherwise repeat.

For example, recall Buchberger's algorithm for computing a Gröbner basis: given the set of polynomials F in $K[X_1, \ldots, X_n]$,

1. Set $T_0 := S$.
2. Given T_i:
   - Calculate a polynomial $s_{f,g} = \frac{a}{L_f} f - \frac{a}{L_g} g$ for some $f, g \in T_i$ so that $s_{f,g} \notin (T_i)$.
   - Reduce $s_{f,g}$ relative to T_i.
   - Set $T_{i+1} = T_i \cup \{s_{f,g}\}$.
3. If T_{i+1} is a Gröbner basis, terminate. Otherwise repeat.

The easy proof of termination goes through Hilbert's Basis Theorem: the T_i generate a strictly increasing sequence of ideals, so this process must terminte.

Bucherger's algorithm stops when $T_{i+1} \subseteq T_i$, so we can just take $F(m) = m + 1$. The sequence of ideals $\langle T_n \rangle$ is computable and the degrees are bounded.

So local Noetherianity gives us an actual bound M so that, for some $m \leq M$, $T_{m+1} \subseteq T_m$. This gives us a numeric bound on the running time of Buchberger's algorithm.

In fact, these bounds are quite poor relative to the actual bounds.

The bounds on local Noetherianity involve an Ackermannian number of iterations of the function F—that is, $M = O(F^{A(n)}(1))$ where:

- A is the Ackermann function, and
- n is the number of variables in the polynomial ring.

Another source of nonconstructive proofs is the use of ultraproducts. The main theorem that makes ultraproducts useful is the *transfer theorem*, which says

Theorem

Suppose K_i is a sequence of structures and that, for each $n_1, \ldots, n_d$, there exist $b_1, \ldots, b_k$ so that, for almost every i,

$$K_i \vDash \phi_{n_1,\ldots,n_d,b_1,\ldots,b_k}.$$

Then the ultraproduct K^* satisfies $\phi_{n_1,\ldots,n_d,b_1,\ldots,b_k}$.

Exactly what ultraproducts do is preserve uniform $\Pi_2^{\mathbb{N}}$ statements.

Whenever we prove a non-$\Pi_2^{\mathbb{N}}$ statement in an ultraproduct, the functional interpretation tells us how to reinterpret it as a $\Pi_2^{\mathbb{N}}$ statement about the original structures.

Whenever we prove a non-$\Pi_2^{\mathbb{N}}$ statement in an ultraproduct, the functional interpretation tells us how to reinterpret it as a $\Pi_2^{\mathbb{N}}$ statement about the original structures.

For example, here's a theorem from van den Dries and Schmidt:

Theorem

Let K_i be a sequence of fields, K^* the ultraproduct, and $K[X]^*_{int}$ the ultraproduct $(K_i[X])^*$. If I is an ideal of $K^*[X]$ then $\sqrt{IK[X]_{int}} \subseteq \sqrt{I} \cdot K[X]^*_{int}$.

That is, anything in the radical ideal generated by I in the larger ring $K[X]^*_{int}$ is a linear combination of things in the ideal $\sqrt{I}$ from the smaller ring.

At least when I is given by finitely many generators, this turns out to be the $\Pi_2^{\mathbb{N}}$ statement:

Theorem

For every d, d', n there is an r so that if $S \subseteq K[X]$, each S has total degree $\leq d$, $\deg(f) \leq d'$, and $f \in \sqrt{S}$ then $f^r \in (S)$.

van den Dries and Schmidt give a proof in the ultraproduct where the main ideas are:

- the prime decomposition of a radical ideal, and
- non-prime ideals have low degree witnesses to non-primality (i.e. if (S) is not prime then there is $fg \in (S)$ with $\deg(fg)$ bounded based on $\deg(S)$ so that $f, g \notin (S)$).

An effective proof: let d, d', n be given and suppose $S \subseteq K[X]$, each S has total degree $\leq d$, and $\deg(f) \leq d'$, and $f \in \sqrt{(S)}$.

We construct a tree of finitely generated ideals indexed by $0, 1$ sequences so that $\deg(S_\sigma) \leq D(\deg(S), |\sigma|)$. $S_{\langle\rangle} = S$.

An effective proof: let d, d', n be given and suppose $S \subseteq K[X]$, each S has total degree $\leq d$, and $\deg(f) \leq d'$, and $f \in \sqrt{(S)}$.

We construct a tree of finitely generated ideals indexed by $0, 1$ sequences so that $\deg(S_\sigma) \leq D(\deg(S), |\sigma|)$. $S_{\langle\rangle} = S$.

If (S_σ) is not prime, there is a witness $gh \in (S_\sigma)$ with $\deg(gh)$ a computable function of $\deg(S), |\sigma|$. Define $S_{\sigma^\frown\langle 0\rangle} = S_\sigma \cup \{g\}$ and $S_{\sigma^\frown\langle 1\rangle} = S_\sigma \cup \{h\}$. If (S_σ) is prime, σ is a leaf.

An effective proof: let d, d', n be given and suppose $S \subseteq K[X]$, each S has total degree $\leq d$, and $\deg(f) \leq d'$, and $f \in \sqrt{(S)}$.

We construct a tree of finitely generated ideals indexed by $0, 1$ sequences so that $\deg(S_\sigma) \leq D(\deg(S), |\sigma|)$. $S_{\langle\rangle} = S$.

If (S_σ) is not prime, there is a witness $gh \in (S_\sigma)$ with $\deg(gh)$ a computable function of $\deg(S), |\sigma|$. Define $S_{\sigma \frown \langle 0 \rangle} = S_\sigma \cup \{g\}$ and $S_{\sigma \frown \langle 1 \rangle} = S_\sigma \cup \{h\}$. If (S_σ) is prime, σ is a leaf.

Let M be the witness given by local Noetherianity applied to D and $F(i) = i + 1$. If any branch had length longer than M, we would contradict local Noetherianity. So there are at most 2^M leaves.

An effective proof: let d, d', n be given and suppose $S \subseteq K[X]$, each S has total degree $\leq d$, and $\deg(f) \leq d'$, and $f \in \sqrt{(S)}$.

We construct a tree of finitely generated ideals indexed by $0, 1$ sequences so that $\deg(S_\sigma) \leq D(\deg(S), |\sigma|)$. $S_{\langle\rangle} = S$.

If (S_σ) is not prime, there is a witness $gh \in (S_\sigma)$ with $\deg(gh)$ a computable function of $\deg(S), |\sigma|$. Define $S_{\sigma^\frown\langle 0\rangle} = S_\sigma \cup \{g\}$ and $S_{\sigma^\frown\langle 1\rangle} = S_\sigma \cup \{h\}$. If (S_σ) is prime, σ is a leaf.

Let M be the witness given by local Noetherianity applied to D and $F(i) = i + 1$. If any branch had length longer than M, we would contradict local Noetherianity. So there are at most 2^M leaves.

Since $f \in (S_\sigma)$ for every leaf σ, $f^{2^M} \in (S)$.

In fact, there is a simple proof using Gröbner bases and the Rabinowitsch trick which gives double exponential bounds.

In the sense of reverse mathematics, proving that a statement (proven using Hilbert's Basis Theorem) has Ackermannian lower bounds is essentially equivalent to showing that Hilbert's Basis Theorem is the only way to prove the statement.

More precisely:

Theorem

Over $\mathbf{RCA}_0$*, Hilbert's Basis Theorem is equivalent to the totality of the Ackermann function.*

Theorem

Suppose $K_1, \ldots, K_n, \ldots$ is a sequence of structures (possibly all the same) and K^* is their ultraproduct. Suppose some sentence σ is true in K^*, where σ is a sentence of first-order logic extended by quantifiers over the natural numbers.

Then there is a $\Pi_2^{\mathbb{N}}$ sentence

$$\forall n_1, \ldots, n_k \; \exists b_1, \ldots, b_d \; \sigma^{ND}_{n_1,\ldots,n_k,b_1,\ldots,b_d}$$

(where the n_i, b_j may include higher order functionals) so that, for each $n_1, \ldots, n_k$ there are $b_1, \ldots, b_d$, almost every K_i satisfies $\sigma^{ND}_{n_1,\ldots,n_k,b_1,\ldots,b_d}$.

In general, when we prove a theorem in an ultraproduct, the functional interpretation tells us to look for a corresponding theorem we can prove in the original structures.

Other than applications of Noetherianity, most results in the area of field theory and differential field theory actually stay entirely in the $\Pi_2^{\mathbb{N}}$ realm.

Some more striking examples have shown up in other areas.

The best studied example is the notion of *locally stable* or *metastable* convergence.

Definition

A sequence $\langle a_n \rangle$ *converges metastably* if for every $\epsilon > 0$ and every $F : \mathbb{N} \to \mathbb{N}$, there is an N so that for all $m \in [N, F(N)]$,

$$||a_N - a_m|| < \epsilon.$$

The best studied example is the notion of *locally stable* or *metastable* convergence.

Definition

A sequence $\langle a_n \rangle$ *converges metastably* if for every $\epsilon > 0$ and every $F : \mathbb{N} \to \mathbb{N}$, there is an N so that for all $m \in [N, F(N)]$,

$$||a_N - a_m|| < \epsilon.$$

Metastable convergence was introduced independently by Avigad-Gerhardy-Towsner and by Tao, and has been studied in the context of ergodic theory and fixed point theory in Banach spaces.

Some proofs appear to use ultraproducts because the corresponding proofs involve rather complicated statements whose finite version is unwieldy.

For example, a theorem from Banach space theory ("the James space is not locally unconditional") involves an intermediate step showing that certain limits can exchange:

Theorem

Under suitable assumptions (in a probability measure space),

$$\lim_n \lim_p \int f_n g_p \, d\mu = \lim_p \lim_n \int f_n g_p \, d\mu.$$

Depending on some assumptions about convergence, this could be interpreted several different ways, but the relevant one is

$$\forall \epsilon > 0 \; \forall n, p \; \exists m > n, q > p \; \forall k > m, r > q \; \exists l > m, s > r$$

$$|\int f_m g_s \, d\mu - \int f_l g_q \, d\mu| < \epsilon.$$

Depending on some assumptions about convergence, this could be interpreted several different ways, but the relevant one is

$$\forall\epsilon > 0\ \forall n, p\ \exists m > n, q > p\ \forall k > m, r > q\ \exists l > m, s > r$$

$$\left|\int f_m g_s\, d\mu - \int f_l g_q\, d\mu\right| < \epsilon.$$

Unfortunately, this is a $\Pi^{\mathbb{N}}_4$ statement. The functional interpretation tells us that the corresponding finite statement is:

Theorem

For every $\epsilon > 0$, every n, p, and all (suitably monotone) functions $\mathbf{K}, \mathbf{R}$, there are $m \geq n$ and $q \geq p$, and functions $\mathbf{L}$, $\mathbf{S}$ so that, setting

- $l = \mathbf{L}(\mathbf{K}(m, q, \mathbf{L}, \mathbf{S}), \mathbf{R}(m, q, \mathbf{L}, \mathbf{S}))$,
- $s = \mathbf{S}(\mathbf{K}(m, q, \mathbf{L}, \mathbf{S}), \mathbf{R}(m, q, \mathbf{L}, \mathbf{S}))$,

$$\left|\int f_m g_s\, d\mu - \int f_l g_q\, d\mu\right| < \epsilon.$$

Theorem (Gilmore-Robinson)

In characteristic 0, a field k is Hilbertian if and only if there is a $t \in k^{\mathcal{U}} \setminus k$ such that $\overline{k(t)} \cap k^{\mathcal{U}} = k(t)$.

This characterization of Hilbertianity says:

$\exists t \in k^{\mathcal{U}}$ *such that*

- $\forall^{ex} a \in k \;\; t \neq a$,
- $\forall^{ex} p \in k(t)[x] \; \forall c \in k^{\mathcal{U}}$
  $(p(c) = 0 \rightarrow \exists^{ex} u \in k(t) \;\; p(u) = 0)$.

We obtain the following characterization of Hilbertianity for countable fields:

There are functions

- $U : \mathcal{P}_{\mathrm{fin}}(k) \times \mathbb{N} \to \mathcal{P}_{\mathrm{fin}}(k)$, *and*
- $D : \mathcal{P}_{\mathrm{fin}}(k) \times \mathbb{N} \to \mathbb{N}$

such that for any finite sets $S, T \subseteq k$ *and any natural number* b, *there is a* $t \in k \setminus T$ *so that for each* $S_0 \subseteq S$ *and* $b_0 \leq b$, *whenever* $p \in k[x]$ *such that*

- *the degree of* p *is at most* b_0
- *each coefficient in* p *has the forum* $\sum_{i \leq b_0} a_i t^{c_i}$ *where* $a_i \in S_0$ *and* $|c_i| \leq b_0$,

then if p *has a root in* k, p *has a root of the form* $\sum_{i \leq D(S_0, b_0)} a_i t^{c_i}$ *where* $a_i \in U(S_0, b_0)$ *and* $c_i \leq D(S_0, b_0)$.

Ultraproducts transform complicated statements about the uniformity of different bounds into simple statement of the sort more common in mathematics.

In fact, this appears to be the entire role of ultraproducts in mathematics outside of model theory and set theory.

We work over a differential field K with a derivation δ.

Definition

$K\{X\}$ is the differential polynomial extension by infinitely many new variables $\delta^k X$ together with the obvious extension $\delta(\delta^k X) = \delta^{x+1} X$.

We work over a differential field K with a derivation δ.

Definition

$K\{X\}$ is the differential polynomial extension by infinitely many new variables $\delta^k X$ together with the obvious extension $\delta(\delta^k X) = \delta^{x+1} X$.

Theorem (Raudenbush)

If $I_1 \subseteq I_2 \subseteq \cdots$ is an increasing sequence of radical differential ideals then there is an N such that $I_{N+m} = I_N$ for all m.

Like Noetherianity in the algebraic setting, Ritt-Noetherianity is often used to prove that algorithms terminate, but doesn't explicitly give bounds.

Getting a local version is complicated by the fact that the best known bounds for checking membership in a radical ideal are Ackermannian using the differential nullstellensatz, when we are precisely interested in analying bounds on statements like the differential nullstellensatz.

We treat membership in a radical differential ideal as an existential statement.

Definition

We say $f \in_{\leq d} \{I\}$ if $f \in (\theta g \mid g \in I \text{ and } ord(\theta) \leq d)$.

We treat membership in a radical differential ideal as an existential statement.

Definition

We say $f \in_{\leq d} \{I\}$ if $f \in (\theta g \mid g \in I \text{ and } ord(\theta) \leq d)$.

Definition

The local Ritt-Noetherianity of $K\{X\}$ says that for any functions O, D, F there are M, L (computable from O, D, F) so that whenever $I_1 \subseteq I_2 \subseteq \cdots$ is an increasing sequence of radical differential ideals in $K\{X\}$ so that I_i is generated by elements of order $\leq O(i, L)$ and degree $\leq D(i, L)$, there is an $m \leq M$ so that each $h \in I_{F(m,L)}$ also has $h \in_{\leq L(O(m,L),D(m,L))} I_m$.