Minicourse: Lecture 1
Applying Topology to Spaces of Countable Structures

Russell Miller

Queens College & CUNY Graduate Center

DDC Program, Part I: Virtual Semester

Mathematical Sciences Research Institute
Berkeley, CA (remotely)
Autumn 2020

Plan of the Minicourse

Week 1: Specific example of subrings of $\mathbb{Q}$.
Online discussion: Thursday, Sept. 24, 11:00 PDT.

Week 2: Computability and continuity.
Online discussion: Thursday, Oct. 1, 11:00 PDT.

Week 3: Classifications of spaces of structures.
Online discussion: Thursday, Oct. 8, 11:00 PDT.

Week 4: The space of algebraic fields.
Online discussion: Thursday, Oct. 15, 11:00 PDT.

Week 5: Other related questions.
Online discussion: Thursday, Oct. 22, 11:00 PDT.

(Also watch Caleb Springer's MSRI Junior Seminar: Oct. 20, 09:00.)

The subrings of $\mathbb{Q}$

We begin with a natural class of structures: the subrings of $\mathbb{Q}$.
What are they?

The subrings of $\mathbb{Q}$

We begin with a natural class of structures: the subrings of $\mathbb{Q}$.
What are they?

Natural classification: subrings R of $\mathbb{Q}$ correspond bijectively to subsets $W \subseteq \mathbb{P}$ of the primes.

$$W \subseteq \mathbb{P} \mapsto \mathbb{Z}[W^{-1}] = \left\{ \frac{m}{n} \in \mathbb{Q} \ : \ \text{all } p \text{ dividing } n \text{ lie in } W \right\}.$$

$$R \subseteq \mathbb{Q} \mapsto \left\{ p \in \mathbb{P} \ : \ \frac{1}{p} \in R \right\} = \left\{ p \ : \ (\exists m, n)\ \frac{m}{np} \in R \ \&\ p \nmid m \right\}.$$

So the space of subrings of $\mathbb{Q}$ "looks like" the power set of $\mathbb{P}$.

Topology on the power set of $\mathbb{P}$

There is a natural topology, the *Cantor topology*, on the power set $\mathcal{P}(\mathbb{N})$ of $\mathbb{N}$, which transfers naturally to $\mathcal{P}(\mathbb{P})$. For a basis, we take the collection of all sets

$$\mathcal{U}_{Y,N} = \{W \subseteq \mathbb{P} : Y \subseteq W \;\&\; N \cap W = \emptyset\},$$

over all pairs (Y, N) of finite disjoint subsets of $\mathbb{P}$. So membership of W in $\mathcal{U}_{Y,N}$ is determined by a finite number of conditions on W.

Under the bijection between $\mathcal{P}(\mathbb{P})$ and $\{$subrings of $Q\}$,

$$\mathcal{U}_{Y,N} = \left\{ R \subseteq \mathbb{Q} \;:\; (\forall p \in Y)\, \frac{1}{p} \in R \;\&\; (\forall p \in N)\, \frac{1}{p} \notin R \right\}.$$

Open sets are unions of arbitrary collections of these $\mathcal{U}_{Y,N}$'s.

Usefulness of open sets

Fix any existential sentence φ, in the language of rings. It is well known that there is an equivalent (for all subrings!) sentence of the form

$$(\exists Y_1) \cdots (\exists Y_n)\, f(Y_1, \ldots, Y_n) = 0,$$

with $f \in \mathbb{Z}[Y_1, \ldots, Y_n]$. Then the set $\mathcal{A}_f$ of subrings R that satisfy φ is soon seen to be an open set.

Usefulness of open sets

Fix any existential sentence φ, in the language of rings. It is well known that there is an equivalent (for all subrings!) sentence of the form

$$(\exists Y_1) \cdots (\exists Y_n)\, f(Y_1, \ldots, Y_n) = 0,$$

with $f \in \mathbb{Z}[Y_1, \ldots, Y_n]$. Then the set $\mathcal{A}_f$ of subrings R that satisfy φ is soon seen to be an open set.

Reason: each solution $\vec{y}$ (in $\mathbb{Q}$) to $f = 0$ uses only finitely many primes in its denominators. If Y is this set of primes, then all rings in $\mathcal{U}_{Y,\emptyset}$ satisfy φ. So the class of all subrings realizing φ is a union of basic open sets.

Usefulness of open sets

Fix any existential sentence φ, in the language of rings. It is well known that there is an equivalent (for all subrings!) sentence of the form

$$(\exists Y_1) \cdots (\exists Y_n)\ f(Y_1, \ldots, Y_n) = 0,$$

with $f \in \mathbb{Z}[Y_1, \ldots, Y_n]$. Then the set $\mathcal{A}_f$ of subrings R that satisfy φ is soon seen to be an open set.

Reason: each solution $\vec{y}$ (in $\mathbb{Q}$) to $f = 0$ uses only finitely many primes in its denominators. If Y is this set of primes, then all rings in $\mathcal{U}_{Y,\emptyset}$ satisfy φ. So the class of all subrings realizing φ is a union of basic open sets.

What is unclear here is why we have the set N in the definition of $\mathcal{U}_{Y,N}$. Using $\mathcal{U}_{Y,\emptyset}$ would have worked just as well for these purposes.

Closed sets

Notice first that every $\mathcal{U}_{Y,N}$ is closed, as well as open.

Lemma

The clopen sets in our topology are exactly the finite unions of basic open sets $\mathcal{U}_{Y,N}$.

To see this, it is helpful to consider the primes one-by-one, in order. A set $W \subseteq \mathbb{P}$ is a path through the binary tree:

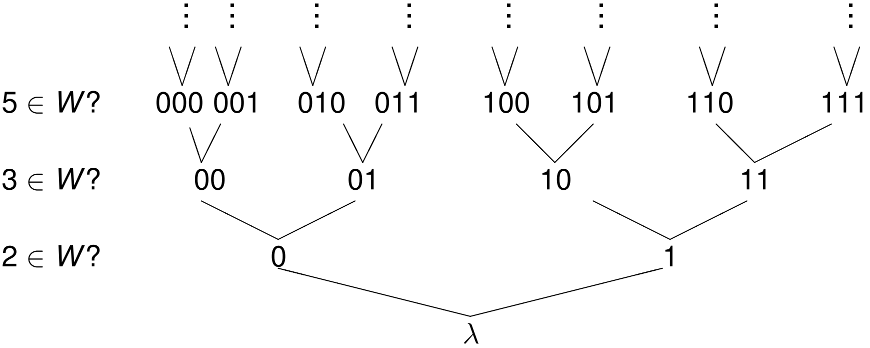

Clopen sets

Suppose that G (in green) and R (in red) are disjoint open sets (of paths through the tree). If there is a level at which they are divided up according to the nodes at that level, then each is a finite union of basic open sets $\mathcal{U}_{Y,N}$.

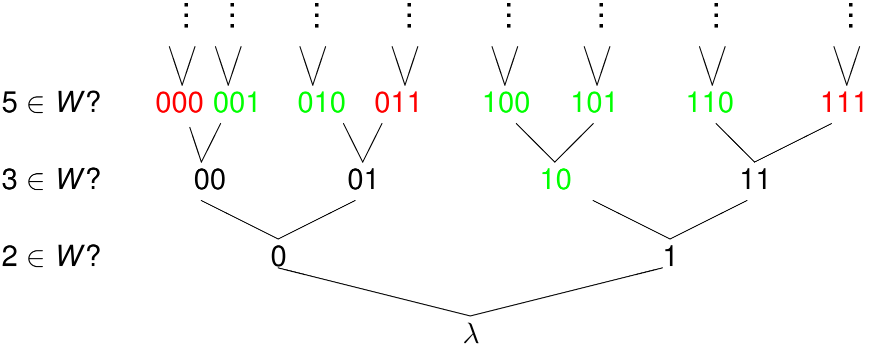

In this example, with revised notation, $R = \mathcal{U}_{000} \cup \mathcal{U}_{011} \cup \mathcal{U}_{111}$.

Non-clopen sets: König's Lemma!

If there is no level such as above, then infinitely many nodes are neither red nor green. Start at λ, and at each level, extend to a node such that infinitely many nodes above it are neither red nor green.

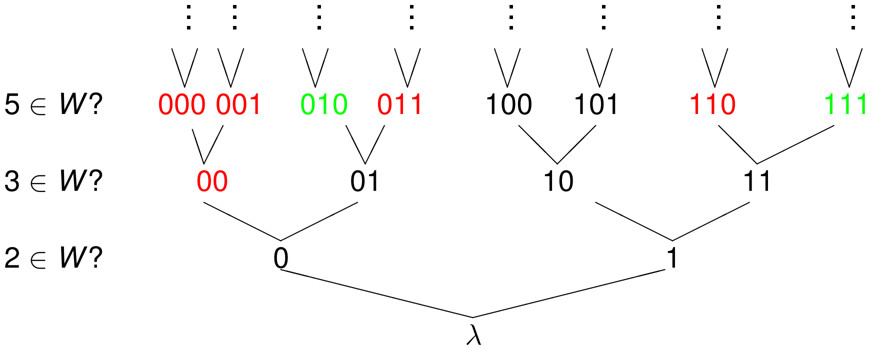

Non-clopen sets: König's Lemma!

If there is no level such as above, then infinitely many nodes are neither red nor green. Start at λ, and at each level, extend to a node such that infinitely many nodes above it are neither red nor green.

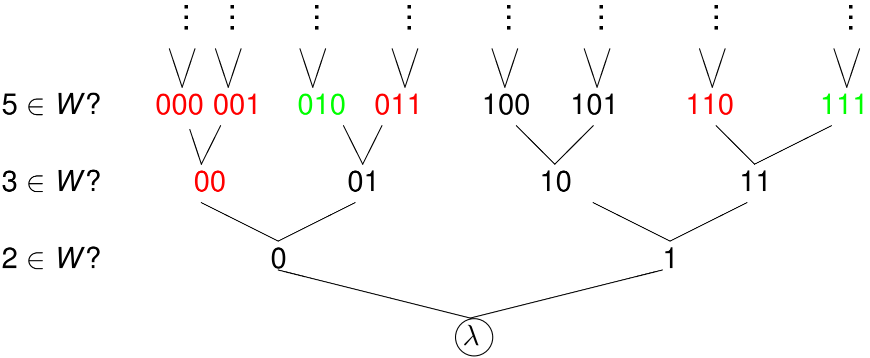

Non-clopen sets: König's Lemma!

If there is no level such as above, then infinitely many nodes are neither red nor green. Start at λ, and at each level, extend to a node such that infinitely many nodes above it are neither red nor green.

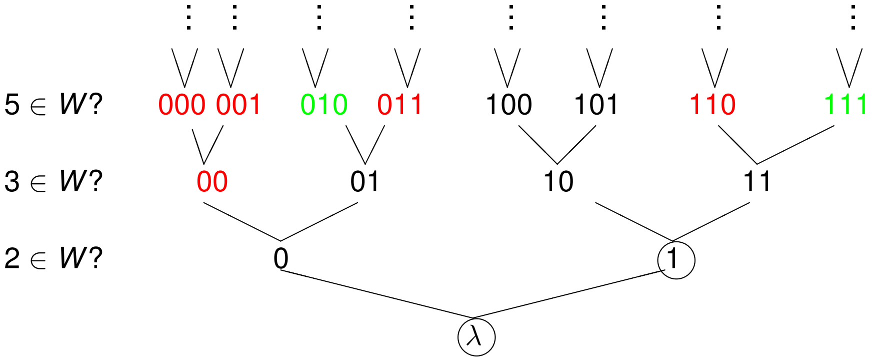

Non-clopen sets: König's Lemma!

If there is no level such as above, then infinitely many nodes are neither red nor green. Start at λ, and at each level, extend to a node such that infinitely many nodes above it are neither red nor green.

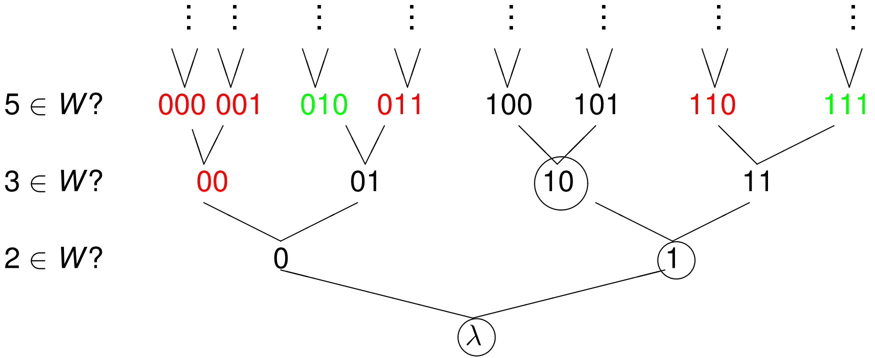

Non-clopen sets: König's Lemma!

If there is no level such as above, then infinitely many nodes are neither red nor green. Start at λ, and at each level, extend to a node such that infinitely many nodes above it are neither red nor green.

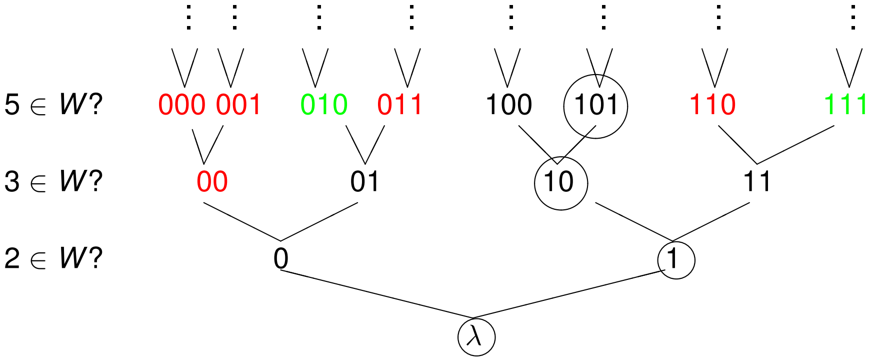

Non-clopen sets: König's Lemma!

If there is no level such as above, then infinitely many nodes are neither red nor green. Start at λ, and at each level, extend to a node such that infinitely many nodes above it are neither red nor green.

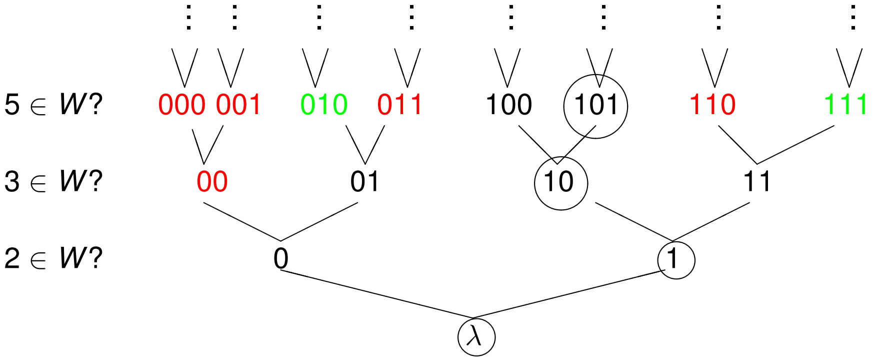

This defines a path $\notin G \cup R$. Thus G cannot be clopen.
Here the path is 101 . . ., meaning the subset $W = \{2, 5, \ldots\}$ of $\mathbb{P}$.

Polynomials

So the first question about polynomials: can they define non-clopen sets of subrings of $\mathbb{Q}$? One suspects so, and the answer is quickly seen to be positive.

Define $f(X, Y, \ldots) = (X^2 + Y^2 - 1)^2 + (\text{“ } X > 0 \text{ ”})^2 + (\text{“ } Y > 0 \text{ ”})^2$.

Solutions to $f = 0$ correspond to nonzero pairs $(\frac{a}{c}, \frac{b}{c})$ with $a^2 + b^2 = c^2$. Elementary number theory shows that $f = 0$ has solutions in exactly those subrings of $\mathbb{Q}$ in which some prime $p \equiv 1 \bmod 4$ is inverted. So the rings with solutions to $f = 0$ form an open but not clopen set $\mathcal{A}_f$.

The polynomials $X^2 + qY^2 - 1$ (modified so that $Y \neq 0$), with q prime, are similar examples, due to Ken Kramer. Here it is necessary and sufficient to invert a prime p for which $-q$ is a square modulo p.

Interior of the complement

An existential formula can fail to have solutions in an entire open set of rings. (Example: $(\exists X, Y, Z)\ (f(X,Y))^2 + (7Z-1)^2 = 0$ and $\mathcal{U}_{\emptyset,\{7\}}$.) That is, a set $\mathcal{U}_{Y,N}$ can lie within the complement of the open set $\mathcal{A}_f$. The *interior* $\mathcal{C}_f$ of the complement of $\mathcal{A}_f$ is the union of all such sets.

This is the first time that the set N in $\mathcal{U}_{Y,N}$ has mattered!

For a polynomial $f \in \mathbb{Z}[\vec{X}]$, here are the three relevant sets of rings:
$\mathcal{A}_f = \{R : f = 0 \text{ has a solution in } R\}$.
$\mathcal{C}_f = \text{Int}(\{R : f = 0 \text{ has no solution in } R\})$, the interior of the complement of $\mathcal{A}_f$.
$\mathcal{B}_f =$ complement of $(\mathcal{A}_f \cup \mathcal{C}_f)$, the topological *boundary* of $\mathcal{A}_f$.

Trying to enumerate $\mathcal{C}_f$

Given a polynomial f, we can computably enumerate all basic open sets $\mathcal{U}_{Y,N}$ within $\mathcal{A}_f = \{R : f = 0 \text{ has a solution in } R\}$. Enumerating the basic open sets that make up $\mathcal{C}_f$ seems much harder. But....

Trying to enumerate $\mathcal{C}_f$

Given a polynomial f, we can computably enumerate all basic open sets $\mathcal{U}_{Y,N}$ within $\mathcal{A}_f = \{R : f = 0 \text{ has a solution in } R\}$. Enumerating the basic open sets that make up $\mathcal{C}_f$ seems much harder. But....

Lemma (Shlapentokh, or Koenigsmann, following J. Robinson)

For each finite set $N \subseteq \mathbb{P}$, the semilocal subring $\mathbb{Z}[\overline{N}^{-1}]$ is diophantine in $\mathbb{Q}$, and its diophantine definition there is uniform in N.

The lemma gives computable maps $F_N : \mathbb{Z}[\vec{X}] \to \mathbb{Z}[\vec{X}]$ for all N, with

$$\mathcal{U}_{\emptyset,N} \subseteq \mathcal{C}_f \iff f \text{ has no solution in } \mathbb{Z}[\overline{N}^{-1}]$$
$$\iff F_N(f) = 0 \text{ has no solution in } \mathbb{Q}.$$

This means that, if we knew which polynomials have solutions in $\mathbb{Q}$, we would be able to enumerate $\mathcal{C}_f$ (by the same method for every f). Thus $\mathcal{C}_f$ is *HTP*($\mathbb{Q}$)*-computably enumerable*.

What about $\mathcal{B}_f$?

Recall: $\mathcal{A}_f$ is an open set. So it does not intersect its boundary $\mathcal{B}_f$: if $R \in \mathcal{B}_f$, then $f = 0$ has no solution in R. But also $R \notin \mathcal{C}_f$: there is no finitary reason why $f = 0$ has no solution in R. (Even if we know that R omits all of the first n primes, this does not rule out all possible solutions.) So, while R indeed contains no solution to $f = 0$, it "never loses hope." (This makes it hard to decide membership in $\mathcal{B}_f$!)

Sometimes $\mathcal{B}_f = \emptyset$. But for the $X^2 + Y^2 - 1$ example, $\mathcal{B}_f$ contains many rings: all those R in which no prime $\equiv 1 \bmod 4$ has an inverse. So this $\mathcal{B}_f$ has the cardinality of the continuum. We may still think this $\mathcal{B}_f$ is small, but the argument must be more subtle than mere counting: we need topology. We will appeal to both Lebesgue measure and Baire category, both of which apply naturally to Cantor space (namely, the power set of $\mathbb{P}$) and thus transfer readily to the space of all subrings of $\mathbb{Q}$.

Lebesgue measure

The Lebesgue measure of a set $\mathcal{U}_{Y,N}$ is defined to be $\frac{1}{2^{|Y \cup N|}}$. If you flip a coin independently for each prime p to decide whether $\frac{1}{p} \in R$, the odds are $2^{-|Y \cup N|}$ that your ring will lie in $\mathcal{U}_{Y,N}$.

This measure is extended to as many sets S of rings as possible (the *measurable sets*) by taking the infimum of the measures of countable covers of S by basic open sets.

Measure of the boundary set

In the $X^2 + Y^2 - 1$ example: to lie in $\mathcal{B}_f$, R must invert no primes $\equiv 1 \bmod 4$. Clearly this $\mathcal{B}_f$ has measure 0.

Open Question

Do all boundary sets $\mathcal{B}_f$ of polynomials $f \in \mathbb{Z}[\vec{X}]$ have measure 0?

This has proven to be a hard question! For a $\mathcal{B}_f$ of positive measure, one could try to build f having (for example) one solution using $\frac{1}{2}$ and $\frac{1}{3}$; another using $\frac{1}{5}$, $\frac{1}{7}$, and $\frac{1}{11}$; then another requiring the next four primes to be inverted, and so on. Is anything like this possible?

Measure of the boundary set

In the $X^2 + Y^2 - 1$ example: to lie in $\mathcal{B}_f$, R must invert no primes $\equiv 1 \bmod 4$. Clearly this $\mathcal{B}_f$ has measure 0.

Open Question

Do all boundary sets $\mathcal{B}_f$ of polynomials $f \in \mathbb{Z}[\vec{X}]$ have measure 0?

This has proven to be a hard question! For a $\mathcal{B}_f$ of positive measure, one could try to build f having (for example) one solution using $\frac{1}{2}$ and $\frac{1}{3}$; another using $\frac{1}{5}$, $\frac{1}{7}$, and $\frac{1}{11}$; then another requiring the next four primes to be inverted, and so on. Is anything like this possible?

Theorem

If $\mathbb{Z}$ has an existential definition in the field $\mathbb{Q}$, then there exist polynomials f with boundary sets of measure arbitrarily close to 1.

Baire category

Recall: a space has the *property of Baire* if no nonempty open set is meager, defined as follows.

A set S of rings is *nowhere dense* if, for every $\mathcal{U}_{Y,N}$, there exist disjoint sets $Y' \supseteq Y$ and $N' \supseteq N$ such that $S \cap \mathcal{U}_{Y',N'} = \emptyset$. (That is, for every $\mathcal{U}_{Y,N}$, S is not dense inside $\mathcal{U}_{Y,N}$.)

The union of a countable collection of nowhere dense sets can fail to be nowhere dense, but we still regard it as small. S is *meager* if S is a countable union of nowhere dense sets. The large sets are the *comeager* sets, the complements of meager sets.

The standard example is the usual topology on $\mathbb{R}$. But Cantor space also has the property of Baire, so we may use Baire category here.

Baire category and $\mathcal{B}_f$

Lemma

For every single polynomial $f \in \mathbb{Z}[\vec{X}]$, $\mathcal{B}_f$ is nowhere dense.

Proof: This is just the ordinary proof that boundaries of open sets (such as $\mathcal{A}_f$) are nowhere dense. Pick any $\mathcal{U}_{Y,N}$. If $\mathcal{A}_f \cap \mathcal{U}_{Y,N} = \emptyset$, then $\mathcal{U}_{Y,N}$, being open, is $\subseteq \mathcal{C}_f$, so $\mathcal{U}_{Y,N} \cap \mathcal{B}_f = \emptyset$. But if $\mathcal{A}_f \cap \mathcal{U}_{Y,N} \neq \emptyset$, then each R there lies within some $\mathcal{U}_{Y',N'} \subseteq \mathcal{A}_f \cap \mathcal{U}_{Y,N}$, just because this intersection is open. So $\mathcal{U}_{Y',N'} \cap \mathcal{B}_f = \emptyset$.

Baire category and $\mathcal{B}_f$

Lemma

For every single polynomial $f \in \mathbb{Z}[\vec{X}]$, $\mathcal{B}_f$ is nowhere dense.

Proof: This is just the ordinary proof that boundaries of open sets (such as $\mathcal{A}_f$) are nowhere dense. Pick any $\mathcal{U}_{Y,N}$. If $\mathcal{A}_f \cap \mathcal{U}_{Y,N} = \emptyset$, then $\mathcal{U}_{Y,N}$, being open, is $\subseteq \mathcal{C}_f$, so $\mathcal{U}_{Y,N} \cap \mathcal{B}_f = \emptyset$. But if $\mathcal{A}_f \cap \mathcal{U}_{Y,N} \neq \emptyset$, then each R there lies within some $\mathcal{U}_{Y',N'} \subseteq \mathcal{A}_f \cap \mathcal{U}_{Y,N}$, just because this intersection is open. So $\mathcal{U}_{Y',N'} \cap \mathcal{B}_f = \emptyset$.

Corollary

The countable union $\mathcal{B} = \cup_{f \in \mathbb{Z}[\vec{X}]} \mathcal{B}_f$ is meager.

So in Baire category, almost all rings lie outside *every* boundary set $\mathcal{B}_f$.

HTP-genericity: never on the boundary

Definition

A subring R of $\mathbb{Q}$ is *HTP-generic* if, for every $f \in \mathbb{Z}[\vec{X}]$, $R \notin \mathcal{B}_f$.

So the HTP-generic subrings form a comeager class. These are the rings where we expect it to be fairly easy to determine whether a polynomial has a solution.

HTP-genericity: never on the boundary

Definition

A subring R of $\mathbb{Q}$ is *HTP-generic* if, for every $f \in \mathbb{Z}[\vec{X}]$, $R \notin \mathcal{B}_f$.

So the HTP-generic subrings form a comeager class. These are the rings where we expect it to be fairly easy to determine whether a polynomial has a solution.

Definition

For a subring R of $\mathbb{Q}$, *Hilbert's Tenth Problem* is the set

$$\text{HTP}(R) = \{f \in \mathbb{Z}[\vec{X}] : f = 0 \text{ has a solution in } R\}.$$

Earlier we mentioned that $\mathcal{C}_f$ is always HTP($\mathbb{Q}$)-computably enumerable. However, the decidability of HTP($\mathbb{Q}$) is an open question.

HTP-genericity and computability theory

Julia Robinson's lemma showed that semilocal subrings $R \subseteq \mathbb{Q}$ all have HTP(R) exactly as hard as HTP($\mathbb{Q}$). All subrings R have HTP(R) $\geq_T$ HTP($\mathbb{Q}$), so these subrings have as simple HTP's as possible. The first use of HTP-genericity was to extend this result.

Theorem (Eisenträger-M.-Park-Shlapentokh, 2017)

There exist subrings $R \subseteq \mathbb{Q}$ such that infinitely many primes p have $\frac{1}{p} \notin R$, yet HTP(R) is Turing-equivalent to HTP($\mathbb{Q}$). Indeed, such rings can have computable presentations, and the set of primes inverted in R can have lower density 0.

The construction used a technique from computability theory called a *finite-injury construction*.
(For subrings of $\mathbb{Q}$, having a computable presentation essentially means that one can computably enumerate the elements of R.)

HTP for HTP-generic subrings

Proposition

For each HTP-generic subring R of $\mathbb{Q}$, $\text{HTP}(R) \equiv_T R \oplus \text{HTP}(\mathbb{Q})$.

The Turing-equivalence $\equiv_T$ here means two things. First, if you know which f have solutions in R, you (or a Turing machine) can decide which rational numbers lie in R itself, and also which g have solutions in $\mathbb{Q}$. Second, if you know these latter two items, then you can decide which f have solutions in R.

The Proposition shows that, if any HTP-generic ring R at all has $\text{HTP}(R) \not\leq_T R$, then $\text{HTP}(\mathbb{Q})$ is undecidable (as it gives R enough of a boost to compute $\text{HTP}(R)$).

Proving the Proposition

Proposition

For each HTP-generic subring R of $\mathbb{Q}$, $\text{HTP}(R) \equiv_T R \oplus \text{HTP}(\mathbb{Q})$.

Exercise: prove the first part (deciding R and $\text{HTP}(\mathbb{Q})$ from $\text{HTP}(R)$).

For the second part, knowing both R and $\text{HTP}(\mathbb{Q})$, and given any f, you can search for:

- a solution to $f = 0$ in R (placing $R \in \mathcal{A}_f$); and
- a finite set $N \subseteq \mathbb{P}$ such that $R \in \mathcal{U}_{\emptyset,N}$ and the polynomial $F_N(f)$ from Julia Robinson's lemma has no solution in $\mathbb{Q}$ (so $\mathcal{U}_{\emptyset,N} \subseteq \mathcal{C}_f$).

Since $R \notin \mathcal{B}_f$, one of these must exist, so you will eventually find it.

Recall: The lemma gives computable maps $F_N : \mathbb{Z}[\vec{X}] \to \mathbb{Z}[\vec{X}]$ with

$$\mathcal{U}_{\emptyset,N} \subseteq \mathcal{C}_f \iff f \notin \text{HTP}(\mathbb{Z}[\overline{N}^{-1}]) \iff F_N(f) = 0 \notin \text{HTP}(\mathbb{Q}).$$

HTP(R) $\equiv_T$ $R \oplus$ HTP($\mathbb{Q}$) for HTP-generic subrings

If any HTP-generic ring R has HTP(R) $\not\leq_T$ R, then HTP($\mathbb{Q}$) is undecidable (as it gives R enough of a boost to compute HTP(R)).

$\mathrm{HTP}(R) \equiv_T R \oplus \mathrm{HTP}(\mathbb{Q})$ for HTP-generic subrings

If any HTP-generic ring R has $\mathrm{HTP}(R) \not\leq_T R$, then $\mathrm{HTP}(\mathbb{Q})$ is undecidable (as it gives R enough of a boost to compute $\mathrm{HTP}(R)$).

Theorem

The following are equivalent, for every set C.

1. $\mathrm{HTP}(\mathbb{Q}) \geq_T C$.
2. A non-meager class of subrings R satisfy $\mathrm{HTP}(R) \geq_T C$.

$\implies$: clear. $\Leftarrow$: then a non-meager class of HTP-generic R have $R \oplus \mathrm{HTP}(Q) \geq_T C$. So some single Turing machine Φ computes χ_C from $R \oplus \mathrm{HTP}(Q)$ for a somewhere-dense set of R, say dense in $\mathcal{U}_\sigma$. Now whenever $\tau \supseteq \sigma$ and $\Phi^{\tau \oplus \mathrm{HTP}(\mathbb{Q})}(n)$ halts, we know it equals $\chi_C(n)$, because some $R \in \mathcal{U}_\tau$ computes χ_C this way.

HTP(R) $\equiv_T R \oplus$ HTP($\mathbb{Q}$) for HTP-generic subrings

If any HTP-generic ring R has HTP(R) $\not\leq_T R$, then HTP($\mathbb{Q}$) is undecidable (as it gives R enough of a boost to compute HTP(R)).

Theorem

The following are equivalent, for every set C.

1. HTP($\mathbb{Q}$) $\geq_T C$.
2. A non-meager class of subrings R satisfy HTP(R) $\geq_T C$.

$\Longrightarrow$: clear. $\Leftarrow$: then a non-meager class of HTP-generic R have $R \oplus$ HTP(Q) $\geq_T C$. So some single Turing machine Φ computes χ_C from $R \oplus$ HTP(Q) for a somewhere-dense set of R, say dense in $\mathcal{U}_\sigma$.
Now whenever $\tau \supseteq \sigma$ and $\Phi^{\tau \oplus \text{HTP}(\mathbb{Q})}(n)$ halts, we know it equals $\chi_C(n)$, because some $R \in \mathcal{U}_\tau$ computes χ_C this way.
So, with an *HTP*($\mathbb{Q}$)-oracle, we just search for such a τ, and when we find it, we have computed $\chi_C(n)$. Such a τ must exist, because $\mathcal{U}_\tau$ contains a ring from the somewhere-dense set.