Microsoft Research
One Microsoft Way, Redmond WA 98052, USA
gurevich@microsoft.com

EMBEDDED FINITE MODELS

Leonid Libkin
University of Edinburgh & University of Toronto

Yuri asked me, the author (**A**) to meet his student Quisani (**Q**), who often appears in public just before a new issue of the *Bulletin* comes out, and for whom Yuri arranges meetings with computer science logicians.

As **Q** looked rather tired and suffering from a lack of sleep, I asked him what had caused it. He explained that in a recent meeting with Jan Van den Bussche, which was reported in this *Column* [38], he was given a chapter on embedded finite models from my book [29] as bedtime reading, but didn't find it very easy to start reading a 14-chapter book from chapter 13. So an email to Yuri followed, and a meeting with me was arranged. The following is my transcription of that meeting.

A. At the very least you're now familiar with the main definition of embedded finite models. Let's review it first.

Q. As I recall it, you start with an *infinite* model or structure, something like the real closed field $\mathfrak{R} = \langle \mathbb{R}, +, \cdot, 0, 1, < \rangle$, and then put a *finite* model on it, say, a finite graph whose nodes are real numbers.

A. That's right. Formally speaking, you have two vocabularies, say Ω for an infinite structure, and σ for a finite structure, and you look at (Ω, σ)-structures, where σ-relations are finite.

Bulletin of the EATCS no 90, pp. 134–149, October 2006

Q. So, for example, if I want to work with graphs whose nodes are real numbers, then Ω could be $(+, \cdot, 0, 1, <)$ and σ should have one binary relation $E(\cdot, \cdot)$ for the edges of my graphs.

A. Exactly. And you'll be working with logical formulae over both Ω and σ. So in first-order logic (which we abbreviate as FO), you can write a sentence

$$\exists a \exists b \forall x \forall y \; E(x, y) \rightarrow a \cdot x + b = y$$

saying that the graph lies on a line.

Q. Do you use special names to distinguish the Ω-structure and the σ-structure?

A. Yes, we usually refer to the Ω-structure as the *background* structure, and to finite σ-structures as *embedded finite models*. In our example, graphs are "embedded" into the real field $\Re$.

Q. Ok, I now remember the definition. But can you explain why anyone would study these objects?

A. Certainly. The initial motivation came from the field of database query languages.

Q. Yes, I heard from many people that databases provide much of the motivation for the development of finite model theory, but how do you come up with embedded finite models?

A. Simple. Do you remember what the main theoretical database query language is?

Q. Of course, it's relational calculus, which is just another name for FO.

A. Correct. For example, if you have a graph, you can ask for pairs of nodes (x, y) connected by a path of length 2 using the formula $\exists z \, (E(x, z) \wedge E(z, y))$, or for nodes x from which there is an edge to every other node: $\forall y \; E(x, y)$. And FO provides the basis of the most common real-life query language SQL.

Q. But we only store finite sets in databases, don't we?

A. Wait a minute. Much of database theory (say, as described in [1, 31]) concentrates on languages that operate with uninterpreted objects – in other words, it doesn't matter what those graph nodes are. But in real databases we operate with *interpreted* objects: say, numbers or strings. In fact, for every relation we put in a database, we must write a `create table` statement in SQL that specifies a type for each attribute: real, integer, Boolean, string, and so on.

Q. I think I see it: elements that we store in a database may come from an infinite set.

A. Not only that,but there are also some domain-specific operations, such as arithmetic operations for numerical domains, that we can use in queries.

Q. Can you give me an example?

A. Let's take a ternary relation $R(\cdot,\cdot,\cdot)$, whose tuples are interpreted as two city names and the distance between them. Then the query

$$\exists z, d_1, d_2\ (R(x, z, d_1) \wedge R(z, y, d_2) \wedge d_1 + d_2 < 100)$$

finds pairs of cities x and y so you can travel between them while visiting another city and the total traveled distance is less than 100.

Q. I remember now, Jan Van den Bussche [38] was talking about applications in Geographical Information Systems.

A. Yes, but this is not the only application. One can think of finite strings and various operations and relations on them, such as adding letters at either end of a string, or checking if one string is a prefix of another.

Q. I see. So, the background structure of vocabulary Ω provides information about the domain and operations on it, and the finite σ-structure is a "database" you put on the Ω-structure.

A. Yes. Note also that while Ω may contain function $(+,\cdot)$, relation $(<)$ and constant $(0, 1)$ symbols, we assume that σ has only relation symbols in it.

Q. This is a rather natural setting. Didn't database people study it to death during the early days of database theory?

A. Not really – they were not that interested in interpreted operations in query languages (although they are present in all real-life languages). Even more importantly from the relational databases point of view, when one writes queries in logical form, one normally assumes that a database is a finite structure with a finite universe. This suffices for many – but not all – database applications (a notable exception is constraint databases, to be discussed shortly). The formal setting of relational databases, however, assumes an *infinite* domain of possible values, albeit without any operations on it. So technically speaking, relational databases are often defined as finite structures embedded into an infinite structure of the empty vocabulary. So in this case, a logical formalism would be that of an infinite structure with a finite structure embedded into it, rather than just a "stand-alone" finite structure.

Q. And no one was curious whether these two settings are different?

A. Some people did. For example, Paris Kanellakis in his survey of relational databases in the Handbook of TCS [25] mentions this distinction. But by that time, it was known that infinite domains without operations don't add anything to the "everything-is-finite" relational model [2, 24].

Q. How can you state this formally?

A. We'll get to it soon – this is done via *collapse theorems*. But let's first talk about a new direction in database research brought to the fore the issues related to infinite domains and interpreted operations – *constraint databases*. They were introduced in 1990 [26], and a book about them appeared ten years later [28].

Q. Yes, I heard about constraint databases from [38]: they are used to represent infinite sets in databases, right?

A. Right. In fact the model of constraint databases is very similar to embedded finite models: all that changes is the interpretation of σ-relations. Now they are not just finite sets, but sets *definable* (in FO) in the background structure.

Q. And what can we represent in this setting?

A. Let's look at the real field again. An FO formula over $\mathfrak{R} = \langle \mathbb{R}, +, \cdot, 0, 1, < \rangle$ with, say, two free variables $\varphi(x, y)$ defines a subset of the the plane $\mathbb{R}^2$ of points that satisfy the formula. Do you remember what these sets are called?

Q. I think it has something to do with algebra. And somehow the name Tarski also comes to mind.

A. Right, they are *semi-algebraic* sets [13, 39]. And by Tarski's quantifier-elimination for the real field, each FO formula $\varphi(\bar{x})$ over $\mathfrak{R}$ is equivalent to a *quantifier-free* formula, that is, just to a Boolean combination of polynomial inequalities $p(\bar{x}) > 0$.

Q. And I presume you can represent a lot of useful information about, say, geography, using such polynomial constraints.

A. True. So now if your query language is FO over the real field and database predicates – interpreted as semi-algebraic sets – you can ask many queries about your geographical objects, which now are *finitely* represented in your database by means of a set of polynomial constraints. An example would be the "database lies on a line" query, which was our first example.

Q. What types of interesting queries can you write in this language?

A. You can test, for example, if a set is topologically open or closed, if it is bounded; for a trajectory $T(x, y, t)$ you can compute the speed at each time t; you can compute the boundary of a set, compare coordinates of specific points, and so on – many queries one needs to ask in GISs.

This language, by the way, is often called FO + Poly (for first-order with polynomial constraints). In many applications even simpler linear constraints are used [20]; the corresponding query language of firsr-order with linear constraints, and database relations definable with linear constraints, is called FO + Lin.

Q. But now I recall that certain things you can*not* ask in FO + Lin and FO + Poly – and that's why Jan suggested I read the embedded finite models chapter in [29].

A. And do you remember an example of a query that FO + Poly cannot express?

Q. I think it was topological connectivity, wasn't it? But then what does it have to do with finite models?

A. It turns out that many questions about expressiveness of FO + Poly over semi-algebraic sets can be reduced to questions about its expressiveness over finite sets. For example, topological connectivity and graph connectivity are very closely related.

Q. I believe I see why: we can embed any graph into $\mathbb{R}^3$ without self-intersections. So a graph is connected iff its embedding is topologically connected!

A. Exactly. There is one little detail: to reduce non-expressibility of topological connectivity to non-expressibility of graph connectivity you must show that the embedding itself is definable in FO + Poly, but this is easily done.

Q. This is a nice example, but it's quite ad hoc. Is there a general result that describes what problems can be reduced to the finite case?

A. Not really – although it would be nice to have such a result – but there are plenty of examples. For instance, Grumbach and Su [21] showed how inexpressibility of many topological properties in FO + Poly can be reduced to questions about embedded finite models. And many other results about constraint databases are obtained by reduction to the finite case [28]. So one can say that embedded finite models play the same role for constraint databases as usual finite models play in relational database theory.

Q. I think we had quite a detour since I asked you about a general result saying that embedded finite models behave just like the usual finite models.

A. You're absolutely right, let's get back to it. As I said, these results come in the form of *collapse theorems*. But before we state them, we need some notations. Let's use FO($\mathfrak{M}, \sigma$) to denote first-order logic over the background Ω-structure $\mathfrak{M}$ and relational vocabulary σ – remember that now σ-relations are finite. For example, FO + Poly is just another name for FO($\mathfrak{R}, \sigma$). Now what would you call the "standard" finite model-theoretic FO using this notation?

Q. Perhaps FO($\mathfrak{M}_\emptyset, \sigma$) where $\mathfrak{M}_\emptyset$ is a structure with an empty vocabulary?

A. Almost, but not quite. The issue, again, is the underlying domain. Let's say $\mathfrak{M}_\emptyset = \langle U, \emptyset \rangle$. If you write $\exists x \varphi(x)$, what does it mean?

Q. I guess it means that there is a witness a for $\varphi(x)$.

A. Correct, but where does this witness come from?

Q. It must come from the universe of $\mathfrak{M}_\emptyset$, that is, from U.

A. And now we have a little problem. When we work with *finite* models, $\exists x$ means that we can find a witness in the universe of the finite model, that is, somewhere in the σ-structure.

Q. Can you explain why this is a problem?

A. Sure. Let's say σ is the vocabulary $E(\cdot,\cdot)$ of graphs, and we want to say that a graph is reflexive. How would you express this in FO?

Q. I think I see what you want to say. I would like to write $\forall x\ E(x, x)$, but that would mean that $E(a, a)$ is true for all $a \in U$, and hence this sentence is false in all finite graphs embedded in $\mathfrak{M}_\emptyset$.

A. Precisely. So we introduce a new type of quantification that only refers to the σ-structure.

One calls the set of all elements of a finite σ-structure $\mathcal{A}$ its *active domain*, and denotes it by *adom*($\mathcal{A}$). And now we introduce active-domain quantification $\exists x \in$ *adom* $\varphi(x)$ and $\forall x \in$ *adom* $\varphi(x)$ with the meaning that there exists an element (or for all elements) a of *adom*($\mathcal{A}$), the formula $\varphi(a)$ is true.

Q. Does it make a logic more expressive?

A. No, it doesn't, because the active domain itelf is easily expressible in FO: say, for graphs by a formula $\exists y\ (E(x, y) \vee E(y, x))$. But then we can define an interesting fragment of the logic FO($\mathfrak{M}, \sigma$), namely its restriction in which all quantification is active-domain, that is, $\exists x \in$ *adom* φ or $\forall x \in$ *adom* φ. We shall denote it by $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma)$.

Q. I see – so now $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_\emptyset, \sigma)$ is the real finite-model theoretic FO over σ-structures, for which the background structure doesn't matter, and we somehow want to reduce FO($\mathfrak{M}, \sigma$) to $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_\emptyset, \sigma)$.

A. Almost - but for reasons that will become clear soon, we can't completely eliminate everything from the vocabulary of the background structure, and we need to keep a linear ordering in it. So with each $\mathfrak{M} = \langle U, \Omega\rangle$ we associate $\mathfrak{M}_< = \langle U, <\rangle$, where $<$ is an arbitrary linear ordering (if $\mathfrak{M}$ had it to start with, we'll keep that ordering), and we shall attempt to reduce questions about FO($\mathfrak{M}, \sigma$) to $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$.

Q. But what do we know about $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$?

A. Plenty, thanks to people working in finite model theory. This is just FO over σ-structures with a linear ordering on them. Many inexpressibility results in this setting are obtained by routine applications of Ehrenfeucht-Fraïssé games, and some heavy tools are available too: for example, the Grohe-Schwentick theorem [19] says that any property expressible in $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$ that does not depend on a particular linear ordering $<$ is local, i.e., determined by the isomorphism type of a small neighborhood of free variables of a formula, and Shelah's theorem [37], which says that even though $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$ does not have a 0-1 law, it has a very weak form of it, called the slow oscillation property.

Q. Ok, I am convinced we can use many facts about $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$ "by citation". But how do we go from $\mathrm{FO}(\mathfrak{M}, \sigma)$ to it?

A. We do it in two steps: first we try to show that $\mathrm{FO}(\mathfrak{M}, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma)$ – and this is called a *natural-active collapse*, because unrestricted quantification over $\mathfrak{M}$ is sometimes referred to as "natural" quantification. As the second step, we try to reduce $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma)$ to $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$.

Q. What do we do first?

A. Let's start with the second step, it's much easier.

Q. I don't see how it can be true that $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$. Say if $\mathfrak{M}$ is the real field and we write something like $\exists a \in adom \exists y \in adom\ E(x, y) \wedge (x + y \neq 1)$. How can we do this if only a linear ordering is available?

A. We cannot. But note that most queries over σ-structures that are of interest to us are queries such as graph connectivity, or cardinality comparisons, and they do not depend on which particular elements of $\mathfrak{M}$ that the active domain of a finite structure consists of. These queries are called *generic*.

Q. Can you define them formally?

A. Of course. Let's do it for Boolean (yes/no) queries. Such a query is just a class C of finite σ-structures $\mathcal{A}$ with $adom(\mathcal{A}) \subset U$. Now suppose $\mathcal{A} \in C$, and let $h : U \to U$ be a 1-1 partial map defined on $adom(\mathcal{A})$. The definition of a generic query Q says that then $h(\mathcal{A})$ must be in C too.

Q. Where $h(\mathcal{A})$ is simply $\mathcal{A}$ in which every $a \in adom(\mathcal{A})$ is replaced by $h(a)$?

A. Of course. Can you give me examples of generic and non-generic queries?

Q. I think I can – graph connectivity, evenness of cardinality are generic, but my earlier example – the existence of an edge (x, y) with $x + y \neq 1$ – is not.

A. Exactly. So our first "reduction" is often called an *active-generic collapse*: it says that every generic query expressible in $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma)$ is also expressible in $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$. That is, $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_<, \sigma)$ with respect to generic queries.

Q. This sounds like a strong result. And what conditions on $\mathfrak{M}$ do you need for it?

A. Here comes the good news – none whatsoever! This is true for all infinite $\mathfrak{M}$.

Q. That's wonderful! Is this hard to prove?

A. Not really. In fact two very similar proofs appeared almost at the same time [8, 33]. They used very similar ideas based on Ramsey's theorem.

Q. Ramsey's theorem? Isn't this about monochromatic cliques and other strange graph-theoretic constructions?

A. These are finite Ramsey theorems. Here we need the original result by Ramsey: if ordered n-tuples over an infinite set U are partitioned into $\ell \geq 2$ classes, then there is an infinite subset $U_0 \subseteq U$ such that all ordered n-tuples over U_0 belong to the same class of the partition.

So next we use this repeatedly to reduce every subformula involving symbols from Ω to a formula that only involves a linear ordering, and over some infinite subset is equivalent to the original one. For example, for $x + y \neq 1$ we can simply find an infinite set $U_0 \subseteq \mathbb{R}$ such that over it for all pairs (x, y) with $x < y$ we have $x + y \neq 1$. Then over U_0 we simply replace $(x + y \neq 1)$ with $x < y$ – and notice that we introduced an ordering!

Q. I think I see the idea now – you eliminate all symbols from Ω except an ordering and still have a formula equivalent to the original one on some infinite set, but by genericity you can assume that your finite structure comes from that set.

A. Exactly. So as you can see, it's a bit tedious but not hard at all. In fact the easiest proof of the active-generic collapse is simply by induction on the structure of a formula, and it is given in full detail in [10] and in Chap. 13 of my book [29].

Q. So far so good, we have the active-generic collapse for all structures. Is it the same for the natural-active collapse?

A. Far from it. Can you think of a simple counterexample?

Q. I think I can; what if we have an empty σ-structure? Then active-domain quantifiers make no sense and any $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma)$ formula is equivalent to a formula that has no quantifiers at all – but this cannot always be true.

A. Yes. In particular this means that every FO formula over $\mathfrak{M}$ is equivalent to a formula that has no quantifiers at all. Do you remember the name of this property?

Q. Of course, it's called quantifier-elimination. I even remember a few examples: $\langle \mathbb{Q}, < \rangle$, $\langle \mathbb{R}, +, \cdot, 0, 1, < \rangle$, or Presburger arithmetic $\langle \mathbb{N}, +, <, 0, 1 \rangle$ if you add all modulo comparisons $n = m(\mod k)$. So if $\mathfrak{M}$ has the natural-active collapse, it must have quantifier-elimination too.

A. Yes, but actually this is not the biggest problem. After all, quantifier-elimination is easy to achieve.

Q. How?

A. You take a structure $\mathfrak{M}$ and simply add a new k-ary predicate symbol P_φ for every formula $\varphi(x_1, \ldots, x_k)$ whose interpretation is $\{\bar{a} \in U^k \mid \mathfrak{M} \models \varphi(\bar{a})\}$. The new structure $\mathfrak{M}_{qe}$ is no different in terms of FO-definability, and it has quantifier-elimination.

Q. I see. And since the active-generic collapse applies to $\mathfrak{M}_{qe}$, it means that all we need to conclude that some generic queries – such as graph connectivity – are not definable in $\mathrm{FO}(\mathfrak{M}, \sigma)$ is to show that $\mathrm{FO}(\mathfrak{M}_{qe}, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_{qe}, \sigma)$.

A. You're absolutely right. In fact, there is even a special name for the statement that $\mathrm{FO}(\mathfrak{M}_{qe}, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_{qe}, \sigma)$: it's called a *restricted-quantifier collapse*.

Q. And it isn't true for all structures either?

A. No, and in fact some very familiar structures provide good counterexamples. Here is a hint: replace $\mathbb{R}$ by $\mathbb{N}$.

Q. I guess the best known structure on $\mathbb{N}$ is the standard arithmetic of addition and multiplication: $\mathfrak{N} = \langle \mathbb{N}, +, \cdot \rangle$. Are you saying that the restricted quantifier collapse fails for it, and we can have queries that are in $\mathrm{FO}(\mathfrak{N}, \sigma)$ but not in $\mathrm{FO}_{\mathrm{act}}(\mathfrak{N}_{qe}, \sigma)$?

A. That's right. Let's think of an example. What can you say about $\mathrm{FO}_{\mathrm{act}}(\mathfrak{N}_{qe}, \sigma)$?

Q. We have active-generic collapse for it, so I can't express queries such as 'is the cardinality of a structure even?'. So now I need to express it in $\mathrm{FO}(\mathfrak{N}, \sigma)$...

A. And if you remember some computability theory, you can tell me how.

Q. Of course – in $\mathfrak{N}$ I can code every finite structure by a natural number, and then I can express every computable property of natural numbers in FO. So of course I can say that the cardinality of a finite set is even. Now I see that we need to impose some conditions on the background structure.

A. Yes, and there's been quite a lot of work on identifying conditions that guarantee collapse: natural-active or restricted-quantifier. In fact, this work started with the simplest structure $\mathfrak{M}_\emptyset = \langle U, \emptyset \rangle$ with an empty vocabulary, and it was shown, by Hull and Su [24] to admit the natural-active collapse: $\mathrm{FO}(\mathfrak{M}_\emptyset, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_\emptyset, \sigma)$.

Q. How does one prove this?

A. We do it by induction on the formula, and the only case that requires work is that of an unrestricted existential quantifier: $\varphi(\bar{x}) = \exists y\ \psi(\bar{x}, y)$. This is equivalent to

$$\exists y \in adom\ \psi(\bar{x}, y)\ \vee\ \bigvee_{x_i \in \bar{x}} \psi(\bar{x}, x_i)\ \vee\ \exists y \notin adom\ \psi(\bar{x}, y).$$

So we need to take care of the last case. But then notice that since the vocabulary is empty, if there is one witness $y \notin adom$ for ψ, then every $y \notin adom$ is a witness for ψ. We thus modify ψ (which is, by the hypothesis, already an $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_\emptyset, \sigma)$) by carefully eliminating the variable y: for example, for each relation S in σ, we can safely replace $S(\ldots, y, \ldots)$ by *false*, since y does not belong to the active domain, and likewise we replace each comparison $y = z$, where z is a quantified variable, by *false* too, since all quantification in ψ is over the active domain. We thus get a formula that does not mention y and is equivalent to $\exists y \notin adom\ \psi(\bar{x}, y)$.

Q. I see. But this proof breaks the moment there is anything at all in the vocabulary.

A. Absolutely. And yet the result is true for the real field. Let's look at one example that we've seen already: all pairs (x, y) in a binary relation S lie on a line. That is, $\exists a \exists b \forall x \forall y\, (S(x, y) \rightarrow a \cdot x + b = y)$. There is an easy way to eliminate the unrestricted quantifiers $\exists a \exists b$. Can you try to say what it means for a set of points to lie on a line?

Q. Doesn't this happen iff every three points are collinear?

A. Exactly. So we can state this property as $\forall x_1, x_2, x_3, y_1, y_2, y_3 \in$ *adom* $(\bigwedge_{i=1}^{3} S(x_i, y_i) \rightarrow \alpha(\bar{x}, \bar{y}))$, where α states that (x_i, y_i), $i \leq 3$, are collinear. And it is easy to write α as a quantifier-free formula.

Q. This is a cute example but it's very ad hoc. And you're saying that we can do something similar with every $\mathrm{FO}(\Re, \sigma)$ query?

A. Yes. Let me tell you the history of this result. It was conjectured in 1990 [26] that some queries such as evenness and graph connectivity are not expressible in $\mathrm{FO}(\Re, \sigma)$, that is, $\mathrm{FO} + \mathrm{Poly}$. The suggested approach was to show the natural-active collapse for the real field $\Re$. This was first achieved in [9] by a non-constructive proof, and a constructive proof appeared in [10]. But a year before the proof of Benedikt and myself [9], Paredaens, Van Gucht and Van den Bussche [34] presented a nice constructive proof of the natural-active collapse for $\langle \mathbb{R}, +, -, 0, 1, < \rangle$ – that is, for the case of linear, rather than polynomial, constraints.

Q. Does multiplication make such a difference?

A. In retrospect, it doesn't. In fact, if you look at the proof of the natural-active collapse for $\Re$ in my book [29], it follows the ideas of Paredaens et al [34]. But the path to that proof wasn't straightforward. In fact, the result of [34] was first generalised quite a bit beyond the real field, as it was proved that every *o-minimal* structure has restricted-quantifier collapse [9, 10]. O-minimality is a central concept of contemporary model theory [35, 39]: it refers to ordered structures $\mathfrak{M} = \langle U, \Omega \rangle$ in which every definable subset of U is a finite union of intervals. Can you tell me why $\Re$ is an example of an o-minimal structure?

Q. I think I can: by Tarski's quantifier-elimination, every formula $\varphi(x)$ is equivalent to a Boolean combination of polynomial inequalities $p_i(x) > 0$, so if r and r' are two roots of polynomials p_i's such that no other root occurs between them, then the signs of all the p_i's on (r, r') don't change and hence the truth value of $\varphi(x)$ doesn't change on (r, r'). Are there other interesting examples of o-minimal structures?

A. There are, and perhaps the most celebrated of them is the "exponential field" – the expansion of $\Re$ with the function e^x. The o-minimality of the exponential field was proved by Wilkie [40].

So [9] proved the result for o-minimal structures, and its constructive version [10]

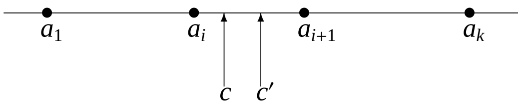

Figure 1: Illustration to the natural-active collapse for the linear case

did so for o-minimal structures again, assuming decidability of their theories. But when the proof was reworked specifically for the case of $\mathfrak{R}$, it looked remarkably similar to the proof for the case of linear constraints.

Q. Will you show this proof to me?

A. I think for this meeting it's better to understand the main idea of the proof for the linear case – after all, it's easier to deal with polynomials that can only have one root, rather than an arbitrary number of roots. So we shall work with $\langle \mathbb{R}, +, -, 0, 1, < \rangle$ as our background structure (and it is well-known to have quantifier-elimination). How do you think the proof will go?

Q. By induction?

A. Of course. So the only case that requires work is elimination of an unrestricted existential quantifier. Let's say we have $\varphi = \exists y \psi(y)$, where $\psi(x)$ is an $\mathrm{FO}_{\mathrm{act}}(\langle \mathbb{R}, +, -, 0, 1, < \rangle, \sigma)$ formula.

Q. Wait a minute, what happened to the free variables? Shouldn't you be looking at $\varphi(\overline{x}) = \exists y \psi(\overline{x}, y)$ to make your induction hypothesis general enough?

A. Of course, but free variables require some extra bookkeeping, and the main ideas can be already seen in the simple case. So let's understand the proof for that case, and you can fill in all the details later.

We assume that $\psi(y)$ is of the form $\exists x_1 \in adom \forall x_2 \in adom \ldots \alpha(\overline{x}, y)$, where α is a Boolean combination of atomic formulae $S(\cdot)$ for $S \in \sigma$ that don't use y (as $S(\cdot, y, \cdot)$ can be replaced by $\exists x' \in adom\ S(\cdot, x', \cdot) \wedge x' = y$), and linear constraints; we also assume that constraints involving y are rewritten as $y \{=, <\} \sum_{i=1}^{m} a_i \cdot x_i + b$.

Let $f_i(\overline{x})$, $1 \leq i \leq p$, enumerate all the functions that occur as right hand sides of linear constraints whose left-hand side is y, and let $f_0(x_1, \ldots, x_m) = x_1$. Now let $\mathcal{A}$ be a finite σ-structure, and let

$$A = \{ f_i(\overline{a}) \mid i = 1, \ldots, p, \ \overline{a} \in adom(\mathcal{A})^m \}.$$

Notice that $adom(\mathcal{A}) \subseteq A$. Assume that $A = \{a_1, \ldots, a_k\}$ with $a_1 < \ldots < a_k$.

Now look at the picture in Fig. 1: if $c \in (a_i, a_{i+1})$ satisfies ψ, then *every* $c' \in (a_i, a_{i+1})$ satisfies ψ because the truth values of $c, c' \{=, <\} f_i(\overline{a})$ are the same for all tuples $\overline{a}$ from the active domain, and all atomic formulae $S(\cdot, c, \cdot)$ and $S(\cdot, c', \cdot)$

are false, since $c, c' \notin adom(\mathcal{A})$.

Q. I see – so if we have a witness for $\psi(y)$ from an interval (a_i, a_{i+1}), the whole interval satisfies ψ. Thus, all we need now is to describe one potential witness from each interval.

A. Yes, and this is easy to do, in a way that is definable with linear constraints: for each interval (a_i, a_{i+1}) we take $(a_i + a_{i+1})/2$ as a witness, for $(-\infty, a_1)$ we take $a_1 - 1$ and for (a_k, ∞) we take $a_k + 1$. Thus, $\exists y \psi(y)$ is now equivalent to:

$$\exists \bar{u}, \bar{v} \in adom \Big(\bigvee_{i=0}^{p} \bigvee_{j=0}^{p} \psi([\frac{f_i(\bar{u}) + f_j(\bar{v})}{2} \;/\; y]) \vee$$
$$\bigvee_{i=0}^{p} (\psi([(f_i(\bar{u}) - 1) \;/\; y]) \vee \psi([(f_i(\bar{u}) + 1) \;/\; y])) \Big)$$

where $\psi([c/y])$ means that c is substituted for y in ψ. Thus, we replaced $\exists y$ with several active-domain quantifiers ($\exists \bar{u} \in adom \exists \bar{v} \in adom$) and a big disjunction over witnesses from the intervals generated by the set A.

Q. The definition of o-minimality you mentioned also talks about intervals...

A. A very good point. This proof is a special instance of a more general proof for o-minimal structures that uses the same ideas: if there is a witness, then a whole interval is a witness; the number of such intervals is finite; and one can choose specific witnesses from them. O-minimal structures $\mathfrak{M}$ have a remarkable "uniform bounds" property: for each formula $\varphi(x, \bar{y})$ there is a number ℓ such that the set $\{a \mid \mathfrak{M} \models \varphi(a, \bar{c})\}$ is composed of at most ℓ intervals, no matter how we choose $\bar{c}$. This is crucial in the proof as it gives us a finite disjunction of cases to check. In the case of the real field this uniform bounds property follows easily from the fundamental theorem of algebra, but in general this is a very nontrivial property [35, 39].

Q. So o-minimality is the best sufficient condition for collapse?

A. No, there are more conditions known now. They are quite model-theoretic in nature [4, 6, 17], and if you want ot learn about them, there are surveys [30, 7] you can check. And while there is no necessary and sufficient condition for collapse, the property that best describes it is finiteness of the VC (Vapnik-Chervonenkis) dimension.

Q. I remember this notion from computational learning theory [3]! It characterises concepts that are efficiently learnable. What does it have to do with embedded finite models?

A. This notion is used not only in learning, but also in model theory, where it is a very useful concept as was noticed by Shelah 35 years ago [36]. Now let's review the concept of VC dimension, shall we? You said that you know it.

Q. Yes, the VC dimension of a collection C of subsets of a set X is the maximum cardinality of a *shattered* finite set $F \subset X$ – if it exists, and if arbitrarily large sets can be shattered, then the VC dimension is infinite. And F is shattered if $\{F \cap Y \mid Y \in C\}$ is the powerset of X. And what does it mean in the language of an infinite structure $\mathfrak{M}$.

A. We say that $\mathfrak{M}$ has finite VC dimension if every definable family has finite VC dimension. And definable families are given by FO formulae $\varphi(\bar{x}, \bar{y})$ as follows: $\{\{\bar{a} \mid \mathfrak{M} \models \varphi(\bar{a}, \bar{b})\} \mid \bar{b} \in U^{|\bar{b}|}\}$.

Q. Can you give me some examples?

A. Yes: for example, all o-minimal structures [39], but also some unordered structures such as the field of complex numbers $\langle \mathbb{C}, +, \cdot \rangle$ [23].

Q. And in what sense is it close to characterising the collapse?

A. It is known that restricted-quantifier collapse ($\mathrm{FO}(\mathfrak{M}_{qe}, \sigma) = \mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_{qe}, \sigma)$) implies finiteness of VC dimension [11], and finiteness of VC dimension implies that $\mathrm{FO}(\mathfrak{M}_{qe}, \sigma)$ and $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_{qe}, \sigma)$, while not necessarily the same, define the same generic queries [4]. In particularly, this very strong result of [4] implies that over every structure of finite VC dimension, the set of generic queries in $\mathrm{FO}(\mathfrak{M}, \sigma)$ is the same as the set of queries definable in $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}_{<}, \sigma)$.

Q. You never said anything about the complexity of the collapse: how hard is it to convert an $\mathrm{FO}(\mathfrak{M}, \sigma)$ formula into an $\mathrm{FO}_{\mathrm{act}}(\mathfrak{M}, \sigma)$ formula?

A. Unfortunately not much is known about this, and complexity analyses may differ significantly for different structures, as such conversion algorithms need to make calls to quantifier-elimination procedures. One case though that was studied in detail is that of the real field and σ consisting of a single unary predicate. For this case Basu [5] developed special algorithms that also give the best known running time for quantifier-elimination for $\mathfrak{R}$.

Q. I think I have plenty of new information now ... I hadn't realised that there was a whole field within model theory developed when Jan Van den Bussche [38] made a passing remark about collapse theorems. It's quite nice to see this interplay between finite and infinite models.

A. Yes, but I don't want you to leave thinking that this is it for finite/infinite models interaction. There are *plenty* of other directions with very interesting results, techniques, and applications.

Q. Can you give me some examples?

A. Certainly. There are metafinite models of Grädel and Gurevich [18] which are finite models with some functions defined on their elements (or tuples of elements) whose range is in the universe of a fixed infinite structure. In logics over metafinite models, variables typically range over the finite part, so interplay is not as complete as in the case of embedded finite models; however, metafinite models make it easy to extend other logics typically studied in the finite model theory context.

There are various finite representations of infinite structures, like in the case of constraint databases. For example, in recursive structures, all predicate symbols are interpreted as recursive relations that are finitely representable by Turing machines. There are interesting connections between finite model theory and the behaviour of logics over recursive structures; a nice survey of this area was written by Harel [22]. As a special and more manageable case, we can consider structures in which all basic predicates (and thus by closure properties, all definable sets) are given by finite automata. These are automatic structures that have been studied rather actively in recent years [27, 15, 11, 12]. They have decidable theories – in fact, decision procedures use automata-theoretic techniques – and these structures found applications in verification and query languages. In particular, [11] looks at finite models embedded into automatic structures. In constraint satisfaction, logical studies of problems with infinite templates recently appeared [14], and those can be viewed as a special case of embedded finite models. In the field of verification people also have been looking at infinite graphs describing configurations of pushdown automata [32, 16]. These again are finitely represented infinite structures with decidable theories that have applications in software verification. So as you can see, there are many other interesting meetings that Yuri can arrange for you in the future, if you'd like to learn more about connections between the finite and the infinite in CS logic.

Q. I shall certainly think about it. And for now, thanks for your time today.

A. You're welcome.

References

[1] S. Abiteboul, R. Hull and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[2] A.K. Ailamazyan, M.M. Gilula, A.P. Stolboushkin and G.F. Shvarts. Reduction of a relational model with infinite domains to the finite-domain case. *Doklady Akademii Nauk SSSR*, 286 (1) (1986), 308–311. Translation in *Soviet Physics – Doklady*, 31 (1986), 11–13.

[3] M. Anthony and N. Biggs. *Computational Learning Theory.* Cambridge Univ. Press, 1992.

[4] J. Baldwin and M. Benedikt. Stability theory, permutations of indiscernibles, and embedded finite models. *Trans. Amer. Math. Soc.* 352 (2000), 4937–4969.

[5] S. Basu. New results on quantifier elimination over real closed fields and applications to constraint databases. *Journal of the ACM*, 46 (1999), 537–555.

[6] O. Belagradek, A. Stolboushkin, M. Taitslin. Extended order-generic queries. *Ann. Pure and Appl. Logic*, 97 (1999), 85–125.

[7] M. Benedikt. Generalizing finite model theory. In *Lecture Notes in Logic, vol. 24, Logic Colloquium 2003*, ASL, 2006.

[8] M. Benedikt, G. Dong, L. Libkin and L. Wong. Relational expressive power of constraint query languages. *Journal of the ACM* 45 (1998), 1–34.

[9] M. Benedikt and L. Libkin. On the structure of queries in constraint query languages. In *LICS'96*, pages 25–34.

[10] M. Benedikt and L. Libkin. Relational queries over interpreted structures. *Journal of the ACM*, 47 (2000), 644–680.

[11] M. Benedikt, L. Libkin, T. Schwentick, L. Segoufin. Definable relations and first-order query languages over strings. *Journal of the ACM* 50 (2003), 694–751.

[12] A. Blumensath, E. Grädel. Automatic structures. In *LICS 2000*, pages 51–62.

[13] J. Bochnak, M. Coste, M.-F. Roy. *Real Algebraic Geometry*. Springer Verlag, 1998.

[14] M. Bodirsky, J. Nešetřil. Constraint satisfaction with countable homogeneous templates. In *CSL 2003*, pages 44–57.

[15] V. Bruyère, G. Hansel, C. Michaux, R. Villemaire. Logic and p-recognizable sets of integers. *Bull. Belg. Math. Soc.* 1 (1994), 191–238.

[16] J. Esparza, A. Kucera, S. Schwoon. Model checking LTL with regular valuations for pushdown systems. *Inf.& Comput.* 186 (2003), 355–376.

[17] J. Flum and M. Ziegler. Pseudo-finite homogeneity and saturation. *J. Symbolic Logic* 64 (1999), 1689–1699.

[18] E. Grädel and Y. Gurevich. Metafinite model theory. *Inf. and Comput.*, 140 (1998), 26–81.

[19] M. Grohe, T. Schwentick. Locality of order-invariant first-order formulas. *ACM Trans. Comput. Log.* 1 (2000), 112–130.

[20] S. Grumbach, P. Rigaux, L. Segoufin. The DEDALE system for complex spatial queries. In *SIGMOD'98*, pages 213–224.

[21] S. Grumbach and J. Su. Queries with arithmetical constraints. *TCS* 173 (1997), 151–181.

[22] D. Harel. Towards a theory of recursive structures. In *MFCS'98*, pages 36–53.

[23] W. Hodges. *Model Theory*. Cambridge, 1993.

[24] R. Hull, J. Su. Domain independence and the relational calculus. *Acta Inform.* 31:513–524, 1994.

[25] P. Kanellakis. Elements of relational database theory. In *Handbook of TCS, Vol. B*, 1990, pages 1073–1156.

[26] P. Kanellakis, G. Kuper, and P. Revesz. Constraint query languages. *JCSS*, 51 (1995), 26–52. Extended abstract in *PODS'90*, pages 299–313.

[27] B. Khoussainov and A. Nerode. Automatic presentations of structures. In *Logic and Computational Complexity*, Springer LNCS vol. 960, 1994, pages 367–392.

[28] G. Kuper, L. Libkin and J. Paredaens, eds. *Constraint Databases*. Springer Verlag, 2000.

[29] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.

[30] L. Libkin. Embedded finite models and constraint databases. In E. Grädel et al., *Finite Model Theory and its Applications*, Springer, 2007, to appear.

[31] D. Maier. *The Theory of Relational Databases.* Computer Science Press, 1983

[32] D. Muller, P. Schupp. The theory of ends, pushdown automata, and second-order logic. *TCS* 37 (1985), 51–75.

[33] M. Otto and J. Van den Bussche. First-order queries on databases embedded in an infinite structure. *Inform. Proc. Letters* 60 (1996), 37–41.

[34] J. Paredaens, J. Van den Bussche, and D. Van Gucht. First-order queries on finite structures over the reals. *SIAM J. Comput.* 27 (1998), 1747–1763. Extended abstract in *LICS'95*.

[35] A. Pillay, C. Steinhorn. Definable sets in ordered structures. *Bull. of the AMS* 11 (1984), 159–162.

[36] S. Shelah. Stability, the f.c.p., and superstability. *Ann. of Math. Logic* 3 (1971), 271–362.

[37] S. Shelah. On the very weak 0-1 law for random graphs with orders. *J. Log. Comput.* 6 (1996), 137–159.

[38] J. Van den Bussche. First-order topological properties. *Bull. EATCS* 87 (2005), 155–164.

[39] L. van den Dries. *Tame Topology and O-Minimal Structures*. Cambridge, 1998.

[40] A.J. Wilkie. Model completeness results for expansions of the ordered field of real numbers by restricted Pfaffian functions and the exponential function. *J. Amer. Math. Soc.* 9 (1996), 1051–1094.