Finite variable logics

Ian Hodkinson*
Department of Computing
Imperial College
180 Queen's Gate, London SW7 2BZ, England.
Email: imh@doc.ic.ac.uk

Abstract

In this survey article we discuss some aspects of finite variable logics. We translate some well-known fixed-point logics into the infinitary logic $L^{\omega}_{\infty\omega}$, discussing complexity issues. We give a game characterisation of $L^{\omega}_{\infty\omega}$, and use it to derive results on Scott sentences. In this connection we consider definable linear orderings of types realised in finite structures. We then show that the Craig interpolation and Beth definability properties fail for $L^{\omega}_{\infty\omega}$. Finally we examine some connections of finite variable logic to temporal logic. Credits and references are given throughout.

1 Some extensions of first-order logic

Quisani: Hello. Who are you? I am Yuri's imaginary student, and I usually talk to him at this time.

Author: I'm afraid he may be a bit late. I am a computer scientist from London, England. I have some imaginary students myself, so maybe I can help. I was reading your earlier conversation on 0–1 laws [Gu3].

Quisani: I remember it. We examined the 0–1 law for logics such as $L^{\omega}_{\infty\omega}$, in which infinite conjunctions and disjunctions are allowed, as well as the first-order operations, but only finitely many variables can occur in any formula.

Author: Right: the top 'ω' represents this restriction on variables. The bottom '∞' means that subject to this restriction, the conjunction or disjunction of any set of formulas, even an uncountable one, is a formula. The bottom 'ω' is fixed and is not so important here, but for the record it means that only finite quantifier depth is permitted — unlike some other logics I could mention. So formally, any atomic L-formula is an $L^{\omega}_{\infty\omega}$-formula, if φ is an $L^{\omega}_{\infty\omega}$-formula and x a variable then $\neg\varphi$ and $\exists x\varphi$ are $L^{\omega}_{\infty\omega}$-formulas, and if Φ is a set of $L^{\omega}_{\infty\omega}$-formulas using only finitely many variables all told, then $\bigwedge\Phi$ and $\bigvee\Phi$ are $L^{\omega}_{\infty\omega}$-formulas.

Quisani: I like the way the restriction to finitely many variables in $L^{\omega}_{\infty\omega}$ is balanced by the possibility of infinite boolean operations; but for computing I guess the finite-variable restriction is the more important. Have you worked on finite variable logics?

*Supported by SERC Advanced Fellowship B/ITF/266.

Author: A little. Would you like to look into them now? We could discuss the way various fixed point logics are subsumed by $L^\omega_{\infty\omega}$ and so inherit its 0–1 law. Then there's an interesting game characterisation of $L^\omega_{\infty\omega}$. We could go on to look at the failure in $L^\omega_{\infty\omega}$ of properties such as Craig interpolation and Beth definability, and then maybe go over the use of finite variable logic in analysing the expressiveness of temporal logic.

Quisani: I think that might be interesting, yes.

Author: Very well. The study of finite variable logics goes back some years, but a lot of work has been done recently because of their connection with computer science. In the early 1980's people were searching for logics stronger than first-order logic for use on finite structures. First-order logic is weak in this regard.

Quisani: But I thought any finite structure could be defined up to isomorphism by a single first-order sentence. Doesn't that mean first-order logic is as powerful as could be?

Author: What you say is true, but first-order logic is not good at defining *classes* of finite structures. For instance, there is no first-order sentence defining the class of all finite structures (in any fixed signature or vocabulary) of even cardinality.

Quisani: Surely this is immediate from the compactness theorem for first-order logic. If arbitrarily large finite structures satisfy a first-order sentence, then some infinite ones do.

Author: True, but as computer scientists we are mainly interested in finite structures. So what I really meant was, there's no first-order sentence whose *finite* models are precisely those finite structures of even size.

Quisani: Ah yes: this is because of the 0–1 law for first-order logic! For any first-order sentence, the proportion of structures of size n satisfying it tends to 0 or 1 as $n \to \infty$. So no first-order sentence can express that its models have even size, as the property 'even size' does not have a limit.

Author: Your argument is not strictly correct, because there's only a 0–1 law for first-order logic if the signature is relational — without function or constant symbols. For example, if c is a constant and P a unary predicate, the proportion of finite structures of any given size that satisfy $P(c)$ is $1/2$.

Quisani: Well, at least I showed that there's no first-order sentence in a relational signature whose finite models are just the finite structures of even size. Hmm. Hang on, there's a much more direct way to see it. If there were a sentence σ with function symbols, relation symbols, anything — say $\forall x \exists y (P(f(g(x,x))) \to f(x) = g(x,y) \vee f(x) = x)$ — whose finite models are *all* the structures of its signature that have even size, then it can't matter at all how the non-logical symbols P, f, g are interpreted in a structure. σ must express 'even size' even if they give no help at all. P could be false everywhere, so we can replace it in σ by $\perp$. And f and g could just be constant functions, with any constant value at all! So $\rho = \exists z \forall x \exists y (\perp \to z = z \vee z = x)$ says that its models have even size, but in the empty signature! For any structure A in this signature, we have

$$A \models \rho \iff \langle A, \emptyset, (x \mapsto z), (\langle x, y\rangle \mapsto z)\rangle \models \sigma \text{ for some } z \in A, \iff |A| \text{ is even.}$$

So ρ defines the structures in the empty signature of even size, and we know this is impossible! So functions can't help in defining structures of even size.

Author: Right, good. We may see other ways that the 0–1 law can be used, later on. But to resume, papers including [AU], [T], [CH], [I2], [Gu1], [L], [Ko], [Va] and [TK] studied logics extending first-order logic by operations such as least fixed point (LFP), monotone fixed point (MFP), inflationary or iterative fixed point (IFP), and partial fixed point (PFP).

Quisani: I'm afraid you will have to remind me what these are.

Author: All right. Think of a first-order formula $\varphi(x_1, \ldots, x_n, P)$ of signature $L \cup \{P\}$, where P is an n-ary relation symbol not in L. If M is an L-structure, φ defines an operation $F_\varphi : M^n \to M^n$ by

$$F_\varphi(S) = \{(a_1, \ldots, a_n) \in M^n : M \models \varphi(a_1, \ldots, a_n, S)\}.$$

Quisani: Don't you mean $dom(M)^n$ here, rather than M^n?

Author: Strictly, yes, but I will use the same notation for a structure as for its domain or universe. This keeps it simpler. Anyway, we can iterate the function F_φ, starting from $S = \emptyset$.

Quisani: Ah yes: in LFP we add a new n-ary relation symbol, interpreted as the least fixed point of F_φ, for each such φ.

Author: Only for φ in which P occurs only positively. In this case we know that the least fixed point exists. For then, F_φ is monotonic, by which I mean $S \subseteq S' \Rightarrow F_\varphi(S) \subseteq F_\varphi(S')$, so the sequence $F_\varphi(\emptyset), F_\varphi^2(\emptyset), \ldots$ is increasing. Thus there must be an α such that $F_\varphi^\alpha(\emptyset) = F_\varphi^{\alpha+1}(\emptyset)$, and this set is the least fixed point of F_φ. In MFP we allow any φ such that F_φ is monotonic, though the set of such φ is undecidable so this logic is rarely taken seriously. For IFP, we allow any φ, and consider $G_\varphi(S) \stackrel{\text{def}}{=} S \cup F_\varphi(S)$. Then $G_\varphi(\emptyset), G_\varphi^2(\emptyset), \ldots$ is an increasing sequence and will contain a fixed point of G_φ, though this may not be its least fixed point. In each case we add a new relation symbol and interpret it as the fixed point.

Quisani: What about PFP?

Author: For PFP we just see if $F_\varphi(\emptyset), F_\varphi^2(\emptyset), \ldots$ reaches a fixed point. If it does, we interpret the new relation symbol as that fixed point; if not, we interpret it as the empty relation.

Quisani: So we have lots of new relation symbols, with specified interpretations. What happens then?

Author: Using this larger stock of relation symbols, we can form new formulas using the usual first-order formation rules. We can then get more new relation symbols from these by taking fixed points again, although here there are relevant normalisation results — see [I2] and [GuS1] for example.

Quisani: OK, I remember now. We get LFP $\subseteq$ MFP $\subseteq$ IFP $\subseteq$ PFP in expressive power.

Author: Right. Now it is easy to see that any sentence of LFP, MFP or IFP can be evaluated on a finite structure in time polynomial in the size of the structure.

Quisani: Right! For example, given $\varphi(x_1, \ldots, x_n, P)$, the sequence $F_\varphi(\emptyset), F_\varphi^2(\emptyset), \ldots$ is an increasing sequence of sets of n-tuples, so it must reach a fixed point in at most m^n steps on a structure of size m.

Author: That's it. But in around 1982 Immerman [I2] and Vardi [Va] independently showed that, in the presence of linear order, LFP is actually equivalent in expressive power on finite structures to the polynomial time queries, P. (This is when there's a special non-logical symbol '$<$' that is interpreted as a linear order on the structure, just as '$=$' is always equality.) The proof works by coding the run of a P-time Turing machine into the structure, using the linear order. We write a sentence of LFP saying that the Turing machine accepts.

Quisani: So in expressive power, LFP = MFP = IFP = P on linearly-ordered structures.

Author: Yes, good.

Quisani: What if there's no linear order around?

Author: Although Gurevich and Shelah [GuS1] showed that even then, LFP, MFP and IFP are equivalent in power on finite structures, in general they become less powerful than P. But in some situations — we will see this later — they do stay as strong as P.

Quisani: What about PFP?

Author: Because for arbitrary φ the sequence $F_\varphi(\emptyset), F_\varphi^2(\emptyset), \ldots$ may not be increasing, it may take more than P-time to see if it converges. But the calculation can still be done in space polynomial in the size of the structure. And Abiteboul and Vianu [AV1] showed that if there's a linear order, PFP = PSPACE in expressive power. They proved that PFP = first-order logic plus a WHILE operator, which [Va, remarks 2 and 4] (and see also [CH, theorem 6.3]) had earlier proved to equal PSPACE in the presence of linear order.

2 The logic $L^\omega_{\infty\omega}$

Quisani: How is all this related to finite variable logic?

Author: Yes, I'm coming to that. These logics, and also some others such as transitive closure and Datalog, are all translatable into $L^\omega_{\infty\omega}$. This was first proved (for LFP on a fixed structure) by Rubin [R], and it appeared in Barwise's [B]; then Kolaitis and Vardi observed that there's a translation that works on all finite structures.

Quisani: Let me think. Ah, I see how to do it. By substituting φ for P in $\varphi(x_1, \ldots, x_n, P)$ repeatedly, we can write a first-order L-formula $\varphi^m(x_1, \ldots, x_n, \emptyset)$ defining the set $F_\varphi^m(\emptyset)$ for each $m < \omega$. So on finite structures, the fixed point can be written in $L^\omega_{\infty\omega}$ as the countable disjunction of these?

Author: Yes, at least for LFP and MFP. For IFP, we replace φ by $\varphi \vee P$. For PFP, $F_\varphi(\emptyset), F_\varphi^2(\emptyset), \ldots$ may not have a fixed point, so there we use

$$\bigvee_{m<\omega} \varphi^m(\bar{x}, \emptyset) \wedge \forall \bar{x}(\varphi^m(\bar{x}, \emptyset) \leftrightarrow \varphi^{m+1}(\bar{x}, \emptyset)).$$

Also, you have to check that the φ^m can be written with a bounded number of variables.

Quisani: Well, inductively, if φ^m can be written with k variables, $x_1, \ldots, x_k$, then I think we can write φ^{m+1} with $x_1, \ldots, x_k$. We want to substitute $\varphi^m(x_1, \ldots, x_n, \emptyset)$ for each subformula $P(t_1, \ldots, t_n)$ in φ. Here, the t_i are terms. Ah, I see a problem! We can't just substitute $\varphi^m(t_1, \ldots, t_n, \emptyset)$: if the t_i use function symbols this may cause clashes of variables.

Author: You could start by replacing each such subformula $P(t_1, \ldots, t_n)$ of φ by $\exists v_1, \ldots, v_n(\bigwedge_{i\leq n} v_i = t_i \wedge P(v_1, \ldots, v_n))$, where the v_i are chosen not to clash with the t_j.

Quisani: Ah yes: so we can assume that all the atomic subformulas P of φ are of the form $P(x_{i_1}, \ldots, x_{i_n})$, where $x_{i_1}, \ldots, x_{i_n}$ are distinct. Now we can just permute the variables $x_1, \ldots, x_k$ used in $\varphi^m(x_1, \ldots, x_n, \emptyset)$ so that $x_{i_1}, \ldots, x_{i_n}$ become free. Then we can substitute it for $P(x_{i_1}, \ldots, x_{i_n})$ in φ.

Author: That's it.

Quisani: So all these logics are subsumed by $L^\omega_{\infty\omega}$ on finite structures. How does $L^\omega_{\infty\omega}$ relate to P and PSPACE?

Author: There are non-recursive classes of finite structures that can be defined in $L^\omega_{\infty\omega}$. For example:

Definition 2.1 1. If '$<$' is a binary relation symbol, let $lin(<)$ be the first-order sentence of $L^3_{\infty\omega}$ (the formulas of $L^\omega_{\infty\omega}$ that are written with just three variables) saying that $<$ is an irreflexive linear order.

2. For any $n < \omega$ let $len_{\geq n}(<)$ be the following (first-order) sentence of $L^2_{\infty\omega}$, the two-variable fragment of $L^\omega_{\infty\omega}$, with n quantifiers:

$$\exists x \exists y(x < y \wedge \exists x(y < x \wedge \ldots)) \cdots)$$

So $len_{\geq n}$ says that there is a $<$-chain of at least n elements.

3. Let $len_n(<) = len_{\geq n}(<) \wedge \neg len_{\geq n+1}(<)$.

Then for any set $S \subseteq \omega$, the $L^3_{\infty\omega}$-sentence

$$lin(<) \wedge \bigvee_{n\in S} len_n(<)$$

defines those linear orders whose size is in S.

Quisani: So if S is non-recursive, the class defined by this is not even recursively decidable, let alone in P or PSPACE.

Author: Exactly. So even $L^3_{\infty\omega}$ can express properties not in P. By using directed graphs instead of linear orders, we can even get away with $L^2_{\infty\omega}$ here.

Quisani: What about the converse? Is PSPACE 'contained' in $L^\omega_{\infty\omega}$?

Author: If there's a linear order, yes, as PSPACE is then equivalent to PFP in expressive power. But not in general: 'even' is in P, but not in $L^\omega_{\infty\omega}$ ...

Quisani: ... because there's a 0–1 law for $L^\omega_{\infty\omega}$.

Author: Right: Kolaitis and Vardi proved this [KV]. And it gives a 0–1 law for all the other logics. This had been proved for each in turn before. For instance, [BGK] proved the 0–1 law for IFP, and hence for LFP and MFP.

3 Games

Quisani: What else is known about $L^{\omega}_{\infty\omega}$ apart from 0–1 laws?

Author: There is a nice game-theoretic characterisation of $L^{\omega}_{\infty\omega}$, in the style of Ehrenfeucht–Fraïssé games and the related theorem of Karp [Kar] for the full $L_{\infty\omega}$. It comes from [B], where it's couched in terms of back-and-forth systems. [I1] and [P] give a 'game' version, which you can also read about in [KV] and [DLW]. Suppose we have structures M, N in the same signature, with disjoint domains, and two players, nowadays usually 'Spoiler' and 'Duplicator', play a game on them. Spoiler is male, and Duplicator, female. Besides being less sexist, this lets us use 'he' and 'she' in proofs, and in our heads, to distinguish the players.

Quisani: So how do Spoiler and Duplicator play here?

Author: For $L^{k}_{\infty\omega}$, the k-variable fragment of $L^{\omega}_{\infty\omega}$, we use what's called a *k-pebble game.* In fact there are $2k$ pebbles, 2 of each 'colour', $1, 2, \ldots, k$. At the start there may already be pebbles in play: some may be on the elements of an n-tuple in M, for some $n \leq k$, the other pebbles of the same colours being on an n-tuple in N. In each round, Spoiler picks up a pebble and places it on an element of one of the structures —

Quisani: — so if all pebbles are already in play, Spoiler will be *moving* a pebble?

Author: Yes, and then Duplicator puts the other pebble of the same colour on an element of the other structure. So at each stage of the game — at the start and after each round — the pebbles in M will be on some tuple $\bar{a}$ and those in N on a tuple $\bar{b}$; player Spoiler wins at that stage if $\bar{a}$ and $\bar{b}$ do *not* satisfy exactly the same atomic formulas.

Quisani: Let me get this clear. Say initially, before any moves, the pebbles coloured 3, 4 and 6 are in play. Suppose the two pebbles of colour 3 are on $a_3 \in M$ and $b_3 \in N$, and similarly for 4 and 6. Then Spoiler wins the game outright, before any moves are made, if there is an atomic formula $\varphi(x, y, z)$ such that the statement '$M \models \varphi(a_3, a_4, a_6) \iff N \models \varphi(b_3, b_4, b_6)$' is not true.

Author: Yes, exactly. Otherwise, the game goes on for another round. If the game goes on infinitely long, Duplicator wins.

Quisani: Isn't this what's called a 'closed game'?

Author: Yes, very good. The condition for Duplicator to win is that she never loses at any finite stage. So by a theorem of Gale and Stewart [GaS], our game is determined: in any situation, either Spoiler or Duplicator has a winning strategy.

Quisani: Now you will say that Duplicator has a winning strategy iff M and N agree on all sentences of $L^{k}_{\infty\omega}$.

Author: Yes, well spotted. We write $M \equiv^{k}_{\infty\omega} N$ in this case — when there are initially no pebbles in play. In the general case we have:

Theorem 3.1 (essentially Barwise) *In the k-pebble game beginning with pebbles on $\bar{a}$, $\bar{b}$, Duplicator has a winning strategy iff $M \models \varphi(\bar{a}) \iff N \models \varphi(\bar{b})$ for all formulas $\varphi \in L^{k}_{\infty\omega}$.*

Quisani: I have some questions about this theorem. What is a strategy, formally?

Author: It's just a set of rules telling a player what move to make in any position. We can allow non-deterministic strategies, so a player may be given more than one option. A strategy can be formalised by functions, though I don't think this helps much. Here, we could represent a strategy for Spoiler in an obvious way by a single function $\xi : M^k \times N^k \to \wp(\{1, \ldots, k\} \times (M \cup N))$. A useful, though not fully general, way of writing a strategy for Duplicator is as a 'back-and-forth system': a non-empty set $\mathcal{E}$ of partial functions $\theta : M \to N$ with $|\theta| \leq k$, $\mathcal{E}$ being closed under restrictions, and such that for any $\theta \in \mathcal{E}$ with $|\theta| < k$ and any elements $a \in M$, $b \in N$ there are $\theta', \theta'' \supseteq \theta$ with $a \in dom(\theta')$ and $b \in rng(\theta'')$. I'll leave you to work out how this formalises the notion of a strategy for Duplicator! As a hint, $\mathcal{E}$ is a winning strategy for Duplicator iff all its maps preserve atomic and negated-atomic formulas. As I said, [B] and others use back-and-forth systems instead of games.

Quisani: From this, it seems that the moves prescribed by a strategy depend only on the current position. Are strategies allowed to remember previous moves?

Author: No, only what is on the 'board' at the time. But this is no disadvantage to either player, as the notions of winning and losing don't depend on the history either. In the ordinary Ehrenfeucht–Fraïssé game for first-order logic [E], it is the same, but there we never move pebbles, so the entire history is always visible anyway.

Quisani: It also seems that if M and N are finite, there'll be only finitely many strategies.

Author: True.

Quisani: But the theorem is true for infinite structures too?

Author: Oh, yes. And the signature can be infinite, and can have function symbols, unlike in the Ehrenfeucht–Fraïssé game for first-order logic.

Quisani: Must $\bar{a}$ and $\bar{b}$ be k-tuples?

Author: They can be n-tuples for any $n \leq k$, including $n = 0$. But it simplifies the notation in the proof a bit if we stick to k-tuples. Let's prove Theorem 3.1, for it is a critical result and the proof is revealing.

Quisani: I see you want me to try to prove it, as usual! Well, the '$\Rightarrow$' direction seems easy enough. We show it by induction on φ. Assume that Duplicator has a winning strategy. Obviously, '$M \models \varphi(\bar{a})$ iff $N \models \varphi(\bar{b})$' holds for atomic φ, as otherwise Spoiler wins immediately, before any moves are made. The negation and infinitary conjunction and disjunction cases are clear by the inductive hypothesis. If now $M \models \exists x_i \varphi(a_1, \ldots, a_k)$, then $M \models \varphi(a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_k)$ for some $a' \in M$. If Spoiler moves a pebble of colour i to a', then Duplicator must have a response $b' \in N$, using her winning strategy. So she has a winning strategy in the game that starts with pebbles on $a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_k$ and $b_1, \ldots, b_{i-1}, b', b_{i+1}, \ldots, b_k$: namely, 'continue with the strategy already in progress'. Inductively, $N \models \varphi(b_1, \ldots, b_{i-1}, b', b_{i+1}, \ldots, b_k)$, so $N \models \exists x_i \varphi(b_1, \ldots, b_k)$, as required.

Author: Good; and the case where $N \models \exists x_i \varphi(b_1, \ldots, b_k)$ is similar. Now try the converse; it is more interesting.

Quisani: That sounds worrying. Let's see. Suppose Duplicator has no winning strategy. If she has already lost at the start, so that the pebbles are initially on tuples $\bar{a}$ and $\bar{b}$ satisfying different atomic formulas, then already the right-hand condition fails, even for atomic φ. Or maybe Spoiler can force a win after one move; then the starting position is better for her, but not much.

Author: So why don't you try to rank the positions by how good they are for her?

Quisani: OK, so a position $\bar{a}, \bar{b}$ has rank ≥ 0 iff $\bar{a}$ and $\bar{b}$ satisfy the same atomic formulas, and it has rank $\geq \alpha + 1$ iff (i) it has rank $\geq \alpha$, and (ii), for every move Spoiler makes, there is a move that Duplicator can make in response to create a new position $\bar{a}', \bar{b}'$ of rank $\geq \alpha$.

Author: Good. What about the case where α is a limit ordinal?

Quisani: Can this happen? If the structures are finite, there are only finitely many positions.

Author: In this theorem we are not assuming them to be finite. And even if they are, a position could still have rank $\geq \alpha$ for all finite α.

Quisani: Well, let's just say that $\bar{a}, \bar{b}$ has rank $\geq \delta$ for a limit ordinal δ iff it has rank $\geq \alpha$ for all $\alpha < \delta$.

Author: Good; and a position has rank α if it has rank $\geq \alpha$ but not $\geq \alpha + 1$, and rank -1 if it doesn't have rank ≥ 0. Now can you characterise the positions of rank $\geq \alpha$ by formulas?

Quisani: Well, the position $\bar{a}, \bar{b}$ has rank at least 0 iff $\bar{a}, \bar{b}$ agree on the set $\mathcal{F}_0$ of all atomic formulas written with variables from $x_1, \ldots, x_k$. Ah, I see: let's define a set $\mathcal{F}_\alpha$ so that the same holds for rank at least α.

Author: Good move!

Quisani: Obviously we'll put $\mathcal{F}_\delta = \bigcup_{\alpha<\delta} \mathcal{F}_\alpha$ for limit δ. For $\alpha+1$, we want every possible move of Spoiler to have a response from Duplicator leaving the resulting tuples agreeing on $\mathcal{F}_\alpha$-formulas. That is, if Spoiler has a move in one structure that produces a tuple satisfying exactly a particular set Φ of formulas in $\mathcal{F}_\alpha$, then Duplicator should have a similar move in the other structure; and this must hold for each set $\Phi \subseteq \mathcal{F}_\alpha$. So let's define

$$\mathcal{F}_{\alpha+1} = \{\exists x_i(\bigwedge \Phi \wedge \neg \bigvee (\mathcal{F}_\alpha \setminus \Phi)) : \Phi \subseteq \mathcal{F}_\alpha, i \leq k\}.$$

(Here, I wrote $\mathcal{F}_\alpha \setminus \Phi$ for the set of all formulas in $\mathcal{F}_\alpha$ but not in Φ.) Yes, this seems to work. We have

$$rank(\bar{a}, \bar{b}) \geq \alpha + 1 \quad \textit{iff} \quad \big(M \models \varphi(\bar{a}) \iff N \models \varphi(\bar{b}) \text{ for all } \varphi \in \mathcal{F}_{\alpha+1}\big),$$

for all $\alpha, \bar{a} \in M, \bar{b} \in N$.

Author: Excellent, but you should also include $\mathcal{F}_\alpha$ in $\mathcal{F}_{\alpha+1}$ — you forgot about part (i) of your definition of having rank $\geq \alpha + 1$. So the correct definition is

$$\mathcal{F}_{\alpha+1} = \mathcal{F}_\alpha \cup \{\exists x_i(\bigwedge \Phi \wedge \neg \bigvee (\mathcal{F}_\alpha \setminus \Phi)) : \Phi \subseteq \mathcal{F}_\alpha, i \leq k\}.$$

Can you finish the proof now?

Quisani: I think so. If $\bar{a}, \bar{b}$ agree on all $L^k_{\infty\omega}$-formulas, then they'll agree on all formulas of each $\mathcal{F}_\alpha$, as we used only k variables in the $\mathcal{F}_\alpha$. So the position $\bar{a}, \bar{b}$ has rank $\geq \alpha$ for all ordinals α —

Author: — say it has rank ∞ —

Quisani: — right, and Duplicator has the strategy 'keep the rank at ∞'. She can do this because if the current position has rank ∞, then for anything that Spoiler does, and any α, she has a response leaving a position of rank $\geq \alpha$ — ah, but does she have a response to give a single position of rank ∞?

Author: Yes. She only has a *set* of responses in any position, not a proper class. So by the axiom of replacement in set theory, the ranks of resulting positions form a set. If she can respond with rank better than any α, then she must have a response of rank ∞.

Quisani: Right: otherwise there'd be a strict upper bound α to the ranks of all resulting positions, so Duplicator couldn't have a response leaving a position of rank $\geq \alpha$. I think Theorem 3.1 is proved now. I can also see that the game is determined, without having to quote the Gale–Stewart theorem. If the starting position has rank ∞, Duplicator has the winning strategy 'keep the rank at ∞'. If not, then Spoiler has the strategy 'reduce the rank'. There is no infinite decreasing sequence of ordinals, so after finitely many rounds the rank will hit zero, and then Spoiler can win after one more move. So this strategy is winning for Spoiler.

Author: This is practically a proof of the Gale–Stewart theorem, anyway; but well done! You're good at games.

Quisani: Well, I remember an earlier discussion with my friend on games [Gu2]. But isn't the axiom of replacement a bit heavy? Surely it cannot be needed in computer science?

Author: It is not needed if the structures are finite. Then all ranks of positions are finite or ∞.

Quisani: Right: this is because there are only finitely many possible positions.

Author: Yes.

Quisani: It is interesting to compare all this with the first-order case. There, an Ehrenfeucht–Fraïssé game of length n characterises the formulas of quantifier depth at most n. Yet here, we have a game of length ω to characterise the formulas with infinite boolean operations, not infinite quantifier depth.

Author: It is striking, I agree. But the boolean operations become trivial in a game context, as your proof shows. And in infinitary logics like $L^k_{\infty\omega}$, the quantifier depth can be unbounded even in a single formula — think of a conjunction of formulas of arbitrarily large quantifier depth — and so the game must go on indefinitely long. [DLW] have more information here.

Quisani: I see.

4 Scott heights and sentences

Author: Do you also see that your proof of the theorem gives more? We can actually extract a single $L^\omega_{\infty\omega}$-formula specifying the ways the game can go, starting with a given tuple on one side.

Quisani: I don't see this. But so far, we considered a pair of tuples in two different structures that 'agreed' on a set $\mathcal{F}_\alpha$ of formulas. Now it sounds like you want to list explicitly which $\mathcal{F}_\alpha$-formulas are true for a single tuple, in one structure.

Author: Yes, I do. Let's fix k as in the theorem. For a given α, if $\bar{a}$ is a k-tuple in a structure M, let the α-type of $\bar{a}$ be:

$$tp_\alpha(\bar{a}) \stackrel{\text{def}}{=} \{\varphi \in \mathcal{F}_\alpha : M \models \varphi(\bar{a})\}.$$

So a position $\bar{a}, \bar{b}$ is of rank $\geq \alpha$ iff $\bar{a}$ and $\bar{b}$ have the same α-type. And in a fixed structure M, the relation $\sim_\alpha$ of having the same α-type is an equivalence relation on M^k.

Quisani: And as α increases, it gets finer: the $\sim_\alpha$-classes get smaller, or at least, no bigger. This must stop at some point.

Author: Right, there must be an ordinal α (with $|\alpha| \leq |M|^k$) such that in M, any two k-tuples having the same α-type have the same $(\alpha+1)$-type. The least such α is called the *Scott height* of M.

Quisani: Does it stop there? Maybe two tuples in M could have the same α- and the same $(\alpha+1)$-type, but different $(\alpha+2)$-types?

Author: It does stop. What does $\sim_\alpha$ really mean?

Quisani: It came from the games, really. Let's play the k-pebble game on two copies of M, not different structures M, N this time. If $\bar{a}, \bar{b}$ are tuples in M, then $\bar{a} \sim_\alpha \bar{b}$ means that the position $\bar{a}, \bar{b}$ has rank $\geq \alpha$.

Author: So if $\sim_\alpha$ is the same as $\sim_{\alpha+1}$ in M —

Quisani: — then *any* position of rank $\geq \alpha$ is also of rank $\geq \alpha+1$. Ah! But now, if a position has rank $\geq \alpha+1$ then by definition, for any move of Spoiler, Duplicator has a response to leave a position of rank $\geq \alpha$, *which must also have rank* $\geq \alpha+1$. So the original position must have had rank $\geq \alpha+2$. And so on: by induction on β we can show that any position of rank $\geq \alpha$ has rank $\geq \beta$ for all $\beta \geq \alpha$. So it has rank ∞, and all the $\sim_\beta$ for $\beta \geq \alpha$ are the same on M.

Author: You've got it.

Quisani: Can we say in $L^k_{\infty\omega}$ that α is the Scott height?

Author: Try it. It is not so hard.

Quisani: Well, we can define $\bar{x} \sim_\alpha \bar{y}$ by $\bigwedge_{\varphi \in \mathcal{F}_\alpha}(\varphi(\bar{x}) \leftrightarrow \varphi(\bar{y}))$. So the sentence $\forall \bar{x}\bar{y}(\bar{x} \sim_\alpha \bar{y} \rightarrow \bar{x} \sim_{\alpha+1} \bar{y})$ says that the Scott height is at most α.

Author: But you used $2k$ variables here: k for $\bar{x}$ and k for $\bar{y}$. Can you say it in $L^k_{\infty\omega}$?

Quisani: Let me see ... well, α is at least the Scott height if $\sim_\alpha = \sim_\alpha +1$, which happens when for every $\varphi \in \mathcal{F}_{\alpha+1}$, the set of tuples in M satisfying φ is a union of $\sim_\alpha$-classes. But the $\sim_\alpha$-class of a tuple $\bar{a}$ is definable in M by $\bigwedge tp_\alpha(\bar{a}) \wedge \neg \bigvee(\mathcal{F}_\alpha \setminus tp_\alpha(\bar{a}))$. Now take the disjunction of these formulas over all $\bar{a} \in M^k$ satisfying φ. This is equivalent to φ in M.

Author: Right — so if α is the Scott height or greater, each $\varphi \in \mathcal{F}_{\alpha+1}$ must be equivalent in M to

$$\varphi' = \bigvee_{\Phi \in \mathcal{S}} \Big(\bigwedge \Phi \wedge \neg \bigvee(\mathcal{F}_\alpha \setminus \Phi)\Big),$$

where $\mathcal{S} = \{tp_\alpha(\bar{a}) : \bar{a} \in M^k, M \models \varphi(\bar{a})\}$.

Quisani: And the $L^k_{\infty\omega}$-sentence

$$\eta_M \quad = \bigwedge_{\varphi(\bar{x}) \in \mathcal{F}_{\alpha+1}} \forall \bar{x}(\varphi(\bar{x}) \leftrightarrow \varphi'(\bar{x}))$$

says that the Scott height of M is at most α.

Author: Well, yes and no. It is true that if $N \models \eta_M$ then the α-type of a k-tuple in N determines its $(\alpha+1)$-type, yes?

Quisani: 'Determines' in what sense?

Author: In the usual sense that any two k-tuples in N of the same α-type will also have the same $(\alpha+1)$-type.

Quisani: Yes, this is true. But N has Scott height at most α iff this holds. I don't see the problem.

Author: We agree that if $N \models \eta_M$ then α-types determine $(\alpha+1)$-types in N, so that N has Scott height at most α. But the converse fails, as η_M says much more than this.

Quisani: Ah, now I see what you're getting at, I think. For each $\varphi \in \mathcal{F}_{\alpha+1}$, η_M specifies a particular infinitary boolean combination φ' of $\mathcal{F}_\alpha$-formulas that's equivalent to φ. So we can work out from η_M exactly *which* $(\alpha+1)$-type any k-tuple in a model of η_M has, once we know its α-type.

Author: Right. A structure N could have Scott height $\leq \alpha$ but not be a model of η_M, if a different formula φ' is equivalent to φ in N.

Quisani: Can you give an example?

Author: The obvious one would be where we have quantifier elimination. If every formula is equivalent in M to a quantifier-free one, then M has Scott height 0. If $k=2$, take M to be the complete graph on two vertices, and N to be its complement, the graph with two vertices and no edges.

These both have quantifier elimination, and Scott height 0. But if E is the edge relation, the formula $\varphi(x_1) = \exists x_2 E(x_1, x_2)$ is equivalent in M to truth and in N to falsity.

Quisani: Then η_M includes a conjunct equivalent to $\forall x_1(\varphi(x_1) \leftrightarrow x_1 = x_1)$, and this is false in N. So $N \models \neg\eta_M$.

Author: Good. I think we've cleared that one up. If we wanted an $L^\omega_{\infty\omega}$-sentence (in a fixed signature) really saying that the Scott height of an L-structure is at most α, we could use the disjunction $\bigvee_M \eta_M$ taken over all L-structures M of Scott height at most α; this is essentially well-formed as $\mathcal{F}_{\alpha+1}$ is a set.

Quisani: Yes, I see. It's then easy to say that the Scott height is exactly α.

Author: Right. But suppose that N does satisfy η_M. What does it mean if tuples $\bar{a} \in M^k$, $\bar{b} \in N^k$ have the same α-type?

Quisani: Er — it means that Duplicator can win the game on M, N starting from $\bar{a}, \bar{b}$?

Author: Yes: if $N \models \eta_M$ then the $(\alpha+1)$-types are determined by the α-types in N in the same way as they are in M. So if $\bar{a} \in M$, $\bar{b} \in N$ have the same α-type, as they do, they'll have the same $(\alpha+1)$-type.

Quisani: That is, for *any* $\bar{a} \in M$, $\bar{b} \in N$, if the position $\bar{a}, \bar{b}$ has rank $\geq \alpha$ then it has rank $\geq \alpha + 1$.

Author: Yes.

Quisani: So now it's just like when we were only looking at M, a minute ago. If the position $\bar{a}, \bar{b}$ has rank $\geq \alpha$ then it must have rank $\geq \alpha+1$, and so $\geq \alpha+2$, and so on. So it has rank ∞.

Author: Good.

Quisani: Now we have it. If $N \models \eta_M$ and tuples $\bar{a} \in M^k$, $\bar{b} \in N^k$ have the same α-type, then the position $\bar{a}, \bar{b}$ has rank $\geq \alpha$, so rank ∞, and Duplicator can win from it.

Author: And we can do this for n-tuples for any $n \leq k$. So by Theorem 3.1, we have:

Theorem 4.1 *Suppose that M, N are structures, M has Scott height at most α, $N \models \eta_M$, and $\bar{a} \in M$, $\bar{b} \in N$ are n-tuples of the same α-type, where $n \leq k$. Then Duplicator has a winning strategy in the k-pebble game starting with pebbles on $\bar{a}, \bar{b}$, and so $\bar{a}$ and $\bar{b}$ agree on all $L^k_{\infty\omega}$-formulas.*

Quisani: Yes, I believe it.

Author: In the case of $n = 0$, suppose Φ is the set of all sentences in $\mathcal{F}_\alpha$ that are true in M, that is, the α-type of the empty tuple, and let Φ' be all the rest — the $\mathcal{F}_\alpha$-sentences false in M. Then the single $L^k_{\infty\omega}$-sentence $\sigma = \eta_M \wedge \bigwedge \Phi \wedge \neg \bigvee \Phi'$ characterises M up to $L^k_{\infty\omega}$-equivalence. We have $N \models \sigma$ iff $M \equiv^k_{\infty\omega} N$. Such a sentence σ is called a *Scott sentence* of M.

Quisani: I have heard of Scott sentences before.

Author: Yes, you can read [Sco] to see. There are other, slightly different ways of writing Scott sentences, without using our $\mathcal{F}_\alpha$, and you can read [DLW] (for example) for details.

Quisani: I suppose we don't see Scott sentences in first-order logic because the necessary infinitary boolean operations are not allowed.

Author: Yes, but note that if the signature is finite and without function symbols then all formulas in the $\mathcal{F}_n$ for $n < \omega$ are first-order. To prove this, it is enough to show that $\mathcal{F}_n$ is finite for each n. We can do this by a trivial induction on n. The case $n = 0$ needs our hypothesis on the signature. So if M (and its signature) is finite, then the Scott sentence σ is first-order.

Quisani: Wow! So if M, N are structures in a finite signature that agree on all first-order sentences using k variables, and M is finite, then M and N agree on all $L^k_{\infty\omega}$-sentences.

5 Defining the types

Author: Do you think the relations $\bar{x} \sim_\alpha \bar{y}$ are themselves definable?

Quisani: Of course: I defined them a minute ago, in $L^{2k}_{\infty\omega}$, as you pointed out.

Author: Yes, but is there a more direct definition: by induction on α, perhaps?

Quisani: Ah, possibly. Let me see: from first principles, we can define $\sim_0$ explicitly, by $\bigwedge_{\varphi \in \mathcal{F}_0}(\varphi(\bar{x}) \leftrightarrow \varphi(\bar{y}))$. Then inductively, $\bar{a} \sim_{\alpha+1} \bar{b}$ holds iff (i) $\bar{a} \sim_\alpha \bar{b}$, and (ii) for every $\bar{a}'$ differing from $\bar{a}$ in one place, there's a $\bar{b}'$ differing from $\bar{b}$ in the same place and such that $\bar{a}' \sim_\alpha \bar{b}'$, and vice versa. Hey, this is like an inductive definition in LFP! We have a formula

$$\varepsilon(\bar{x}, \bar{y}, \sim) \quad = \quad \bar{x} \sim_0 \bar{y} \wedge \bar{x} \sim \bar{y} \wedge \bigwedge_{i \leq k} \Big(\forall x_i \exists y_i (\bar{x} \sim \bar{y}) \wedge \forall y_i \exists x_i (\bar{x} \sim \bar{y})\Big),$$

where $\bar{x} = x_1, \ldots, x_k$ and $\bar{y} = y_1, \ldots, y_k$. This is positive in $\sim$. Can we take a fixed point?

Author: Yes, except that as α increases the relations $\sim_\alpha$ shrink, rather than grow — the classes get smaller, remember? In fact, the least fixed point of your ε is $\emptyset$.

Quisani: Whoops! What's gone wrong?

Author: The way we want to use ε is not usual in LFP. We want to start with $\sim\ = M^{2k}$ and then iterate, and not with $\sim\ = \emptyset$ as usual. That is, we want to take the greatest fixed point of F_ε, not its least fixed point.

Quisani: Is this greatest fixed point expressible in LFP?

Author: Yes: if $\psi(\bar{x}, R)$ is positive in R then its greatest fixed point is obtainable, in the obvious notation, as $\neg\mathrm{LFP}(\neg\psi(\bar{x}, \neg R))$. This is an easy exercise. Note that $\neg\psi(\bar{x}, \neg R)$ is also positive in R, so this is well-formed.

Quisani: Right. I remember I've seen that trick before. I always find it funny how we get three negations when we do that.

Author: So to sum up, if the signature is finite and without function symbols we can get a formula $E(\bar{x}, \bar{y})$ of LFP saying that '$\bar{x} \sim_\alpha \bar{y}$ for all α'.

Quisani: Why do we need to restrict the signature?

Author: So that $\bar{x} \sim_0 \bar{y}$, and hence ε, are first-order.

Quisani: Does E work on infinite structures?

Author: Yes. The number of iterations before a fixed point is reached is the Scott height, which may be infinite; but once a fixed point is arrived at, it is what we want.

Quisani: Yes, I see.

6 Ordering the types

Author: We can say even more, still. Abiteboul and Vianu [AV2, section 3.2] showed that the E-classes can be linearly ordered. There's a single LFP formula $L(\bar{x}, \bar{y})$ defining a linear pre-order on M^k for any finite structure M, so that $L(\bar{a}, \bar{b}) \wedge L(\bar{b}, \bar{a})$ holds iff $E(\bar{a}, \bar{b})$. The formula L does not depend on M.

Quisani: That's amazing. You mean there's a LFP-definable linear order in any finite structure?

Author: A linear pre-order, yes.

Quisani: How is it done?

Author: [DLW] give an argument for it, which in a nutshell goes like this. Assume again that the signature is finite and without function symbols. As M is finite, $E = \sim_n$ for some $n < \omega$, so it's enough to order the $\sim_n$-classes for each finite n. We do this by induction on n. First, the finitely many $\sim_0$-classes are linearly ordered arbitrarily. Assume inductively that we've ordered the $\sim_n$-classes, as $p_1, \ldots, p_r$ say. We want to linearly order the $\sim_{n+1}$-classes. Remember that $\sim_{n+1}$ refines $\sim_n$, so each $\sim_{n+1}$-class lies in a unique $\sim_n$-class. If p, q are $\sim_{n+1}$-classes contained in *distinct* $\sim_n$-classes p_i, p_j, say, then we let $p < q$ iff $i < j$. So it remains to order the $\sim_{n+1}$-classes contained in each single $\sim_n$-class.

So fix a $\sim_n$-class, p_ℓ. We associate with each k-tuple $\bar{a} \in p_\ell$ a string of zeros and ones of length kr, as follows. Place our k pebbles on $\bar{a}$. Then define the $(r(i-1)+j)$th digit of the string to be 1 iff Spoiler (say) can move the ith pebble so as to give a new k-tuple $\bar{a}' \in p_j$.

Quisani: So for example, the digits $\ell, r+\ell, \dots, r(k-1)+\ell$ will always be 1, because Spoiler can choose not to move a pebble at all, in which case $\bar{a}' = \bar{a} \in p_\ell$.

Author: That's right. Now it follows from the definitions that the strings associated with k-tuples $\bar{a}, \bar{b} \in p_\ell$ are the same iff $\bar{a} \sim_{n+1} \bar{b}$. So the lexicographical order on strings induces a linear order on the $\sim_{n+1}$-classes in p_ℓ, and this completes the induction.

Quisani: Wow! And this order is LFP-definable?

Author: Yes. It's easiest to find a suitable IFP-formula first, and then use the theorem of [GuS1] that IFP = LFP on finite structures. To show the idea, let's write a formula $\varphi(\bar{x}, \bar{y}, <)$ whose greatest fixed point is the (irreflexive) linear pre-order we want. As before, $\bar{x} = x_1, \dots, x_k$ and $\bar{y} = y_1, \dots, y_k$ are k-tuples of distinct variables. First, the base case. Enumerate the maximal consistent conjunctions of atomic and negated atomic formulas of $L^k_{\infty\omega}$, as $\alpha_1(\bar{x}), \dots, \alpha_r(\bar{x})$, say. There are finitely many of them because of our assumption on the signature. Then φ will have a conjunct of the form

$$\alpha(\bar{x}, \bar{y}) = \bigvee_{i<j\le r} \alpha_i(\bar{x}) \wedge \alpha_j(\bar{y}).$$

Quisani: This is to order the $\sim_0$-classes.

Author: Right. Now for the iterative step. Let $\bar{x} \sim \bar{y}$ abbreviate the formula $\bar{x} \not< \bar{y} \wedge \bar{y} \not< \bar{x}$. At the nth stage of the iteration, this says $\bar{x} \sim_n \bar{y}$. We'll have $\bar{x} < \bar{y}$ at the $(n+1)$th stage if (1) $\bar{x} < \bar{y}$ already at the nth stage, or else (2) $\bar{x} \sim_n \bar{y}$ and for some $i \le k$ and some $\sim_n$-class c represented by $\bar{z}$, say, (2a) Spoiler can move y_i in such a way as to put $\bar{y}$ in c but can't move x_i to put $\bar{x}$ in there, while (2b) all lexicographically previous possibilities (smaller i, or same i and $<$-smaller $\bar{z}$) are the same for $\bar{x}$ and $\bar{y}$.

Quisani: Right! So now you just have to say this in logic.

Author: Yes... how about this:

$$\overbrace{(\bar{x} < \bar{y})}^{(1)} \vee \overbrace{(\bar{x} \sim \bar{y} \wedge \bigvee_{i\le k} \exists \bar{z} \psi_i)}^{(2)},$$

where ψ_i is

$$\begin{array}{ll} (2a) & \exists y_i(\bar{y} \sim \bar{z}) \wedge \neg \exists x_i(\bar{x} \sim \bar{z}) \\ (2b) & \wedge \bigwedge_{j<i} (\forall x_j \exists y_j(\bar{x} \sim \bar{y}) \wedge \forall y_j \exists x_j(\bar{x} \sim \bar{y})) \\ & \wedge \forall \bar{z}' < \bar{z}(\exists x_i(\bar{x} \sim \bar{z}') \leftrightarrow \exists y_i(\bar{y} \sim \bar{z}')) \end{array}$$

Here, $\bar{z}$ and $\bar{z}'$ are k-tuples of new variables. The final formula $\varphi(\bar{x}, \bar{y}, <)$ is thus

$$\alpha(\bar{x}, \bar{y}) \wedge \Big(\bar{x} < \bar{y} \vee (\bar{x} \sim \bar{y} \wedge \bigvee_{i\le k} \exists \bar{z} \psi_i)\Big).$$

It doesn't depend on the structure M.

Quisani: I believe it. Very nice. The predicate $<$ isn't positive in φ, but there is a greatest fixed point, which is the pre-order we want, and we can express this in IFP.

Author: Yes. But remember that the order we get is not on the points of a structure, but on its E-classes. In some structures, E is not very interesting. On structures of size $\ge k$ in the empty signature, for example, $tp_\alpha(\bar{a})$ really only says which elements of the tuple $\bar{a}$ are equal and which are not, however big α is.

Quisani: So all k-tuples of distinct elements are E-equivalent, for example. The Scott height is 0. This is not much use.

Author: On the other hand, sometimes the E-classes are closely related to the structure — for distinct points a, b, we have $\neg E(a, \ldots, a, b, \ldots, b)$. [DLW] calls such structures *k-rigid*.

Quisani: Doesn't this mean the structure itself can be linearly ordered?

Author: Yes, the LFP-formula $L(x, \ldots, x, y, \ldots, y)$ defines a linear order on any finite k-rigid structure, and so we get LFP = P on such structures. This generalises the result of [I2,Va] that I mentioned before.

Quisani: Right, you said we'd see some situations in which the fixed-point logics stay as strong as P. I thought you'd forget.

Author: So did I. But even for arbitrary structures, something can still be salvaged. The E-classes are really all that $L^k_{\infty\omega}$ can know about a structure. It has no finer resolution. So any $L^k_{\infty\omega}$-formula can be 'rephrased' to talk about the classes, which are linearly ordered. In effect, as far as $L^k_{\infty\omega}$ is concerned, we can take all structures to be linearly ordered! The same goes for fragments of $L^\omega_{\infty\omega}$ that can express the order, such as LFP and PFP. And we know linearly-ordered structures are nice in some ways —

Quisani: — yes, we have LFP = P and PFP = PSPACE on finite linearly-ordered structures.

Author: Right. This allows us to prove that LFP = PFP on finite structures iff P = PSPACE, a result of Abiteboul and Vianu [AV2].

Quisani: How?

Author: Do you see which is the easy direction?

Quisani: — I think '$\Rightarrow$'. If LFP is as expressive as PFP, then this is certainly true on linearly ordered structures, where we know LFP = P and PFP = PSPACE. So P = PSPACE for ordered structures. But P and PSPACE are by definition sets of sets of character strings, and a string can be regarded as an ordered structure. So P = PSPACE.

Author: Good. Now the converse: assume P = PSPACE, and take a sentence $\sigma \in$ PFP.

Quisani: Right, it dawns on me what to do ... we can rephrase σ, giving a sentence σ' of PFP talking about the linear ordering of the E-classes. So σ' is essentially equivalent to σ, but is evaluated in a linearly-ordered structure. We know that on linearly-ordered structures, PFP = PSPACE and P = LFP. But PSPACE = P, so there's a LFP-sentence σ'' equivalent to σ' on the E-classes. Now translate σ'' back into a sentence about the original structure.

Author: Well done. This result is interesting because it relates complexity to logic without a given linear order, and also because LFP and PFP do not always themselves correspond to P and PSPACE. Anyway, we've had a good run with the games. Much of the work we have covered is taken from [DLW], which I strongly recommend. I also recommend a coffee.

Quisani: Good idea.

7 The Craig interpolation property

Quisani [still drinking coffee]: Can we do anything else with all these results?

Author [could not find any coffee]: We can do quite a lot with them. For example, we can use them to show that, as in first-order logic on finite structures [Gu1], and unlike in classical first-order logic, the Craig interpolation and Beth definability properties fail for $L^{\omega}_{\infty\omega}$.

Quisani: As computer scientists, we are mainly interested in finite structures. Are you saying these properties fail on finite structures, or generally?

Author: They fail whoever you are — whether you restrict the semantics to finite structures or not.

Quisani: I know first-order logic has the interpolation property on all structures but not on finite ones [Gu1]. So is interpolation for finite structures stronger? Does it imply interpolation generally?

Author: I can't see why. Let's look at the definition. The Craig interpolation property says that if A, B are sentences in signatures L_A, L_B, and $\models A \to B$, then there's a sentence C in the common signature $L_A \cap L_B$ which is an 'interpolant': we have $\models A \to C$ and $\models C \to B$. By $\models X$ I mean that every structure in the signature of X is a model of X — and here we have a choice, as you said. We may restrict to finite structures X, or we may allow any structure. We get two different interpolation properties.

Quisani: So OK, let's write $\models$ for full validity and $\models_f$ for validity on finite structures. Assume we have interpolation for finite structures: so $\models_f A \to B \Rightarrow \models_f A \to C \wedge C \to B$ for some C. Now for arbitrary ones, if $\models A \to B$, then clearly $\models_f A \to B$, so there's a C with $\models_f A \to C \wedge C \to B$... ah, but how do we know that $\models A \to C \wedge C \to B$?

Author: Exactly.

Quisani: So are there logics with the interpolation property for finite structures but not for arbitrary structures?

Author: I can't think of one. Anyway, to show that the Craig interpolation property fails for $L^{\omega}_{\infty\omega}$ in either semantics, consider the $L^{3}_{\infty\omega}$-sentences

$$A = lin(<) \wedge \bigvee_{n \text{ even}} len_n(<), \qquad B = lin(\prec) \to \bigvee_{n \text{ even}} len_n(\prec).$$

Here, $<$ and $\prec$ are different binary relation symbols, and lin and len_n are as we defined them at the beginning (definition 2.1). Now, do you see that $\models A \to B$?

Quisani: Yes: any model of A is a $<$-linear order of finite even length, so any other linear order $\prec$ on it must also be of even length.

Author: And an interpolant C would be an $L^{\omega}_{\infty\omega}$-sentence in the empty signature, just using '='. It follows that the models of C would be exactly the finite structures of even size.

Quisani: Why?

Author: Well, if $M \models C$, choose a linear order '$\prec$' on M. We still have $(M, \prec) \models C$, as C doesn't mention '$\prec$'. But $\models C \to B$, so $(M, \prec) \models B$. So $(M, \prec) \models \bigvee_{n \text{ even}} len_n(\prec)$, and so M must be finite, of even size.

Quisani: Ah, I get the idea. For the converse, if M has even size then choose a linear order '$<$' on M. So $(M,<) \models A$. But $\models A \to C$, so $(M,<) \models C$, and as C doesn't involve '$<$' we get $M \models C$. Very good.

Author: So could the interpolant C exist?

Quisani: No. We know the property 'this structure has even size' can't be expressed in $L^{\omega}_{\infty\omega}$, as it has a 0–1 law.

Author: Good; or we can use Theorem 3.1. Clearly Duplicator has a winning strategy in the k-pebble game between any two structures of size at least k, if the signature is empty. So by Theorem 3.1, no sentence of '$\emptyset^k_{\infty\omega}$' will tell apart two structures of size $\geq k$.

Quisani: And the same idea works if we restrict to finite structures?

Author: Yes: the proof only involved finite structures. Come to think of it, it seems to me that a similar proof might show that the logic $L_{\infty\omega} + W$ has the interpolation property for finite structures but not for arbitrary ones, where W is a generalised quantifier saying that a binary relation well-orders the universe. You were asking for an example of this —

Quisani: — I'll try it. But something else is puzzling me now. If C can't express 'even size', how can A and B?

Author: Well, A says that its models are *linear orders* of even size. The limiting probability of this is 0, as almost all structures are not linear orders. For B, the probability is 1, for the same reason.

Quisani: Yes, I see now. It's obvious really.

8 The Beth definability property

Quisani: Now what about Beth? What is this property? I'm not sure I remember.

Author: Take a signature L and an n-ary relation symbol $R \notin L$. Let T be an $L \cup \{R\}$-theory in a logic. We say that T *implicitly defines* R over L if for any structures $M, N \models T$ such that the L-reducts $M\lceil L$, $N\lceil L$ of M and N are the same, we have $M = N$.

Quisani: I suppose an $L \cup \{R\}$-theory is a set of sentences of the logic in the signature $L\cup\{R\}$? And taking the L-reduct of a structure just means forgetting the interpretations of symbols not in L?

Author: That's right.

Quisani: What do you mean by $M\lceil L$, $N\lceil L$ being the same?

Author: I mean they're equal, identical, $=$. If you like, you can read the whole definition as saying that any L-structure has at most one expansion to a model of T.

Quisani: OK.

Author: Now we say that T *explicitly defines* R over L if there is an L-formula $\varphi(x_1, \ldots, x_n)$, an 'explicit definition', such that $M \models \forall x_1, \ldots, x_n(R(x_1, \ldots, x_n) \leftrightarrow \varphi(x_1, \ldots, x_n))$ for all $L \cup \{R\}$-structures $M \models T$.

Quisani: We call φ an explicit definition of R because R and φ are equivalent in any model of T even though φ doesn't mention R?

Author: Right. Then we say that the logic has the *Beth property* if any $L \cup \{R\}$-theory T that implicitly defines R over L also explicitly defines it.

Quisani: OK. And this property fails for $L^\omega_{\infty\omega}$?

Author: Yes, it does, both for finite structures (insert the word 'finite' before 'structure' throughout the definitions) and for arbitrary ones.

Quisani: I guess the reason is the same as for Craig. Perhaps we could define a nullary relation P and a unary relation Q implicitly, in terms of a given linear order $<$, by saying that P is true iff $<$ is a *finite* linear order (we use an infinite disjunction to say this), the $<$-first and last elements satisfy Q, and exactly one of each pair of $<$-adjacent elements satisfies Q. I think we can code two relations into a single one. Then P is true iff the structure is a linear order of odd size. I guess it's not explicitly definable, for the same reasons as before.

Author: Unfortunately, $lin(<) \wedge \bigvee_{n \text{ odd}} len_n(<)$ defines P explicitly (see definition 2.1). Your approach works in the first-order finite-model case, and [Gu1] used it, but I don't think it works here.

Quisani: Well, maybe Beth survives if the implicit definition is first-order but the explicit one may be infinitary.

Author: No. We'll show:

Theorem 8.1 *For all $k \geq 3$, there's a* first-order *theory T_k in $L^3_{\infty\omega}$ that implicitly defines a linear order which can't be explicitly defined in $L^k_{\infty\omega}$.*

Hence the Beth property also fails for L^k, the k-variable fragment of first-order logic.

Quisani: Yes, but you're only claiming that the Beth property fails for $L^k_{\infty\omega}$ for each $k \geq 3$. Is that enough to show that it fails for $L^\omega_{\infty\omega}$?

Author: Remind me at the end to mention that. For now, we'll fix a k and find our T_k. Here we use the 0–1 law in a deeper way. This part of the argument is due to the algebraic logicians Andréka, Düntsch and Németi [ADN]. They were only looking at the first-order case, but it's not hard to do $L^\omega_{\infty\omega}$ too, and this was done in [HKL]. The algebraic logic community have worked on finite variable logics for many years, and they've built a considerable body of further work, for instance on the 'weak Beth property'.

Quisani: The what?

Author: Ask me later! But I'd just like to say that they have studied logics related to k-variable logics but with different semantics, such as 'relativised semantics'. These logics often do have nice properties like Craig interpolation, Beth definability, decidability, and others, even on finite structures. Recent references for some of this work include [ABN, Ve2].

Anyway, let's prove theorem 8.1. We use graphs. Do you remember the extension axioms in graphs?

Quisani: Yes, my friend was explaining them when we discussed 0–1 laws [Gu3]. The k-extension axiom says that given any two disjoint sets X, Y of nodes of combined size less than k, there's a node outside X and Y that's joined by edges to every node of X but to no node of Y. This can be written as a first-order sentence with k variables.

Author: Good, you remembered. These axioms all have limiting probability 1. Now suppose we have graphs G, H satisfying the k-extension axiom, and k-tuples $\bar{a} \in G$, $\bar{b} \in H$ satisfying the same quantifier-free type. If Spoiler, Duplicator play the k-pebble game on G, H starting from the position $\bar{a}, \bar{b}$, then —

Quisani: — then I think Duplicator has a winning strategy. If Spoiler moves pebble i to a new point $a' \in G$, say, then we can chop the set of points in G covered by the remaining pebbles into two disjoint sets X and Y, where X gets the points connected to a' by edges, and Y gets the rest. We then translate these sets to the corresponding sets $X', Y' \subseteq H$, using the map $\bar{a} \mapsto \bar{b}$.

Author: Are X' and Y' disjoint?

Quisani: Yes: the map preserves all quantifier-free formulas. And obviously, the combined size of X' and Y' is at most $k-1$. So we can use the k-extension axiom in H, and find $b' \in H \setminus (X' \cup Y')$ connected to all of X' and none of Y'. Duplicator can respond with this b', and it'll leave new tuples $\bar{a}', \bar{b}'$ satisfying the same quantifier-free formulas again, just like $\bar{a}, \bar{b}$. So nothing is lost, and Duplicator can continue in the next round in the same way. If Spoiler moved in H instead, Duplicator uses the k-extension axiom in G. She never loses, so she wins.

Author: Very nice. It is important that 'nothing is lost': it is really saying that the rank of the new position is the same as the old, so must be ∞. You've shown in effect that any graph satisfying the k-extension axiom has Scott height 0.

Quisani: Yes, I see what you mean. But I thought I showed much more. I essentially determined η_G for $L^k_{\infty\omega}$: it is the k-extension axiom!

Author: — yes, I forgot! Well done.

Quisani: So to summarise, by Theorem 4.1 we get:

Theorem 8.2 *If graphs G, H satisfy the k-extension axiom, then whenever k-tuples $\bar{a} \in G$, $\bar{b} \in H$ satisfy the same quantifier-free formulas, they also satisfy the same $L^k_{\infty\omega}$-formulas.*

Author: OK, well put. Now let's leave all this for a moment. Fix a finite graph G, and suppose '$<$' linearly orders its nodes in some arbitrary way. So we can write the nodes of G as $\{g_1, \ldots, g_r\}$, with $g_1 < g_2 < \cdots < g_r$. Consider the following first-order $L^3_{\infty\omega}$-sentence δ_G:

$$\begin{aligned} & lin(<) \wedge len_r(<) \\ \wedge & \bigwedge_{\substack{1 \le i,j \le r \\ G \models E(g_i, g_j)}} \forall xy(\text{`}x, y \text{ are the } i\text{th}, j\text{th elements along } <, \text{resp.'} \to E(x,y)) \\ \wedge & \bigwedge_{\substack{1 \le i,j \le r \\ G \not\models E(g_i, g_j)}} \forall xy(\text{`}x, y \text{ are the } i\text{th}, j\text{th elements along } <, \text{resp.'} \to \neg E(x,y)). \end{aligned}$$

Here, the formula $lin(<) \wedge len_r(<)$ says that '$<$' is an r-element linear order (see definition 2.1). Let $\nu_n(x) = \exists y(y < x \wedge \exists x(x < y \wedge \cdots)) \cdots)$ with n quantifiers. Then the clause 'x is the ith element along $<$' abbreviates a formula of the form $\nu_i(x) \wedge \neg \nu_{i+1}(x)$, which uses only two variables. And E is the binary relation symbol used for the graph edge relation.

Essentially, δ_G describes the diagram of G.

Quisani: How's that?

Author: Well, what does a model of δ_G look like?

Quisani: Let $M \models \delta_G$. Then $<$ linearly orders $dom(M)$ as $\{m_1, \ldots, m_r\}$, say. And for each $i, j \leq r$ we have $M \models E(m_i, m_j)$ if $G \models E(g_i, g_j)$, and $M \models \neg E(m_i, m_j)$ otherwise. Ah, I see: δ_G forces M to be exactly the same as G. In fact, the map $g_i \mapsto m_i$ is a graph isomorphism from G to M.

Author: Good: the graph part (reduct) of any model of δ_G is isomorphic to G. So can we conclude that any two models of δ_G with the same graph reduct are actually the same?

Quisani: Yes.

Author: No! δ_G doesn't just specify the graph structure of its models. It also says what the graph looks like under a certain linear order. It specifies that in the given enumeration '$<$' of the nodes, *the first and ninth* elements (say) have a graph edge between them, etc. This certainly constrains '$<$'. But why should it determine it completely?

Quisani: Hmm — it may not do. For example, if G has no edges at all, *any* ordering fits. So then there are $r!$ models of δ_G, all with the same graph reduct, whereas what we really want is a G with a *unique* enumeration that fits the pattern described by δ_G.

Author: Yes. Now suppose there are two different enumerations (orderings) $(G, <), (G, \prec)$ of G, both of which fit the pattern of δ_G, so that G looks the same in both enumerations — the first and ninth elements in both are joined by a graph edge, etc. This simply means that $(G, <) \cong (G, \prec)$. That is, there's an automorphism of the *graph* G that takes $<$ to $\prec$.

Quisani: An automorphism is just an edge-preserving permutation of the nodes of G?

Author: Yes. Now what if G is *rigid:* without non-trivial automorphisms?

Quisani: Then this can't happen: there can't be two different enumerations of G fitting the pattern coded in δ_G. If $(G, <), (G, \prec) \models \delta_G$ then $< = \prec$. So δ_G implicitly defines $<$!

Author: Yes, exactly. It's a nice argument, isn't it?

Quisani: But can $<$ be explicitly defined in $L^k_{\infty\omega}$ over δ_G?

Author: No, not if the k-extension axiom holds in G. Do you see why?

Quisani: Well, we showed that any graph satisfying the k-extension axiom has quantifier-elimination in $L^k_{\infty\omega}$. Right, then. Take a model M of δ_G. If $<$ had an explicit definition $\varphi(x, y)$ over δ_G, then we'd have $M \models \forall xy(x < y \leftrightarrow \varphi(x, y))$. But the graph reduct of M is isomorphic to G, so it satisfies the k-extension axiom. φ is a formula of graphs, so by theorem 8.2 it's equivalent in G to a quantifier-free graph formula $\psi(x, y)$ of graphs — and this poor ψ has to define a linear order on G! But ψ will be symmetric, because graphs are. And $<$ isn't symmetric. Contradiction!

Author: Good.

Quisani: So if G is rigid, then $<$ is implicitly defined by δ_G; and if it satisfies the k-extension axiom then $<$ can't be explicitly defined over δ_G.

Author: Well put.

Quisani: So to show that the Beth property fails for $L^k_{\infty\omega}$, all we need to do now is find a finite rigid graph $G = G_k$ satisfying the k-extension axiom. Can this be done for all k?

Author: Yes: each k-extension axiom has limiting probability 1, and so does 'rigidity', as Erdös and Renyi [ER] proved. So there are many finite graphs with both properties. [ADN] construct such a graph explicitly, and it can be given other properties too.

Quisani: A very nice argument. But you've only covered the case $L^k_{\infty\omega}$. What about $L^\omega_{\infty\omega}$?

Author: The Beth property fails for $L^\omega_{\infty\omega}$ too. Choose a G_k as before, for each k. There is an infinite supply of each, so we can assume that the sizes of the G_k are strictly increasing: $|G_1| < |G_2| < \cdots$. Choose a linear order on each G_k, write δ_{G_k} as before using this order, and let $\delta = \bigvee_{k\geq 1} \delta_{G_k}$. This is an $L^3_{\infty\omega}$-sentence that implicitly defines a linear order $<$ in every G_k.

Quisani: Why is that? Oh, I think I see. Let M_1, M_2 be models of δ with the same graph parts. Then $M_1 \models \delta_{G_{k_1}}$ and $M_2 \models \delta_{G_{k_2}}$ for some k_1, k_2. So the graphs of M_1, M_2 are isomorphic to G_{k_1}, G_{k_2}. So G_{k_1} and G_{k_2} must be isomorphic. This means they must have the same size, at least. But you chose the G_k all different sizes, so $k_1 = k_2$. Now $M_1, M_2 \models \delta_{G_{k_1}}$, and we assumed M_1, M_2 have the same graph parts. So as before, $M_1 = M_2$.

Author: Very good — so δ implicitly defines $<$. Now, if $<$ were explicitly definable in $L^\omega_{\infty\omega}$, then it would be definable in some $L^k_{\infty\omega}$. But we know $L^k_{\infty\omega}$ can't explicitly define a linear order on G_k.

Quisani: Right! And the proof shows that the Beth property fails for $L^\omega_{\infty\omega}$ even if we restrict the semantics to finite structures.

9 The weak Beth property

Quisani: You also asked me to ask you about the weak Beth property.

Author: Oh yes. Well, the implicit definitions of the ordinary Beth property — the one we were just using — have the form $\exists_{\leq 1} R\varphi(R)$, where R is a relation symbol, and $\exists_{\leq 1} R$ is a second-order quantifier, meaning that there is *at most one* value of R making $\varphi(R)$ true.

Quisani: But we had no second-order quantifiers before... ah, I see what you mean! We said that φ defines R implicitly if whenever $M, N \models \varphi(R)$ and the reducts M', N' of M, N got by forgetting R are the same, then M, N agree on R as well. So there's *at most one* value of R in M' that will make $\varphi(R)$ true.

Author: Yes. That's it. But suppose we restrict to stronger implicit definitions, of the form $\exists_{=1} R\varphi(R)$. Here we insist there's *at least one* suitable R, as well.

Quisani: You mean that with stronger implicit definitions, there's more chance of finding an explicit definition?

Author: Yes. We define a logic to have the *weak Beth property* if any $L \cup \{R\}$-theory T that implicitly defines R over L also explicitly defines it (over L), so long as

- any model of the L-reduct of T (throw out all non-L-sentences from T) expands to a (necessarily unique) model of T.

Quisani: The ordinary Beth property implies the weak one?

Author: Of course.

Quisani: We already saw $L^{\omega}_{\infty\omega}$ doesn't have the ordinary Beth property. Was $\delta = \bigvee_{k\geq 1}\delta_{G_k}$ already a *strong* implicit definition of $<$? Then we'd have shown weak Beth fails too.

Author: No, I'm afraid it wasn't. The 'reduct' of the theory $T = \{\delta\}$ to graphs is empty; and it's not true that any model of $\emptyset$ in the signature of graphs expands to a model of δ.

Quisani: OK, so let's put some subsidiary information in T. I know: let $T = \{\delta\}\cup\{\neg\Delta_G : G$ a finite graph, $\forall k(G \ncong G_k)\}$, where Δ_G is a first-order sentence saying its models are isomorphic to G.

Author: Then any *finite* model of the graph reduct of T expands to a model of δ. But what about infinite ones?

Quisani: Yes, you are right. At least this T does seem to show weak Beth fails for $L^{\omega}_{\infty\omega}$ when we restrict the semantics to finite structures. But in the general case, I can't see how to add $L^{\omega}_{\infty\omega}$-sentences to the graph reduct of T to force its models to be finite. Besides, my additions to T were rather artificial. OK, so does $L^{\omega}_{\infty\omega}$ have the weak Beth property?

Author: No, it doesn't! This can be seen using a recent example of Gurevich and Shelah [GuS2]. They built finite structures called *multipedes.* A multipede looks a bit like this:

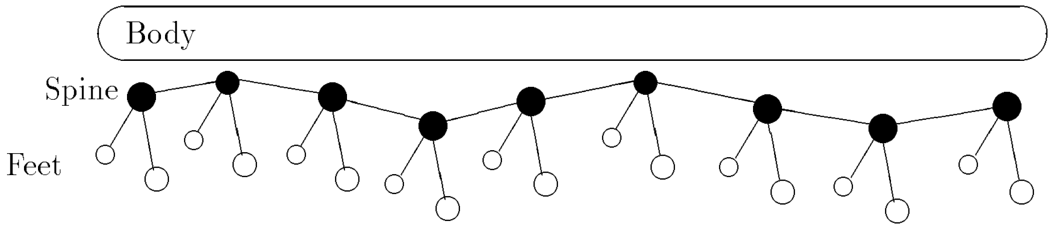

Quisani: [shudders]

Author: Their essential attributes are:

- The class $\mathcal{M}$ of finite 'odd' multipedes is definable by a first-order sentence μ, using (I think) four variables.
- Models of μ are in effect three-sorted structures. In the picture, the three sorts are the box (the 'body'), the black dots ('spine'), and the white dots ('feet'). All three sorts are needed, for technical reasons, but we'll only talk about the spine here. For biological reasons I should maybe follow [HKL] and call the spine the 'segments', but then it might get confused with the body, so I won't.
- μ forces the spine to be linearly ordered, by '$\prec$', say.
- If the spine of a model of μ is finite, so is the whole model.

The two main properties are:

- $\mathcal{M}$ is not k-rigid for any k (see section 6) — for each k, there's an odd multipede in $\mathcal{M}$ with two distinct points (actually the two feet of any black dot) that satisfy the same $L^{k}_{\infty\omega}$-formulas. This used a rather fine probabilistic argument. So, as with δ_k in §8, there is no $L^{k}_{\infty\omega}$-formula that explicitly defines a linear ordering on every multipede in $\mathcal{M}$.

- Every multipede in $\mathcal{M}$ is rigid: and in fact, [DHK] showed there's a first-order sentence $\lambda(<)$ that implicitly defines a linear order $<$ on every multipede in $\mathcal{M}$.

Quisani: I begin to like these multipedes a little. Can you explain how they are constructed?

Author: Let's wait till your friend arrives. You could ask him. It will be a very stimulating and valuable conversation.

Quisani: Ah well...but at least it seems clear that $<$ is strongly implicitly defined by $\{\mu, \lambda\}$, but can't be explicitly defined in $L^\omega_{\infty\omega}$, for the same reason as before!

Author: Does it? How do we know that *any* model of μ expands to a model of λ?

Quisani: But I thought you said μ defined $\mathcal{M}$. Surely this means any model of μ is a multipede in $\mathcal{M}$, so λ implicitly defines a linear order on it.

Author: No. Any *finite* model of μ is a multipede in $\mathcal{M}$. We don't know about infinite models: maybe λ doesn't implicitly define a linear order in them. In fact it can't always, as first-order logic has the Beth property!

Quisani: So we need somehow to restrict to finite models of μ. This is where I got stuck before. How are things better with multipedes?

Author: We can use the $\prec$-linearly-ordered spine. If this is finite, so is the multipede. And we can say in $L^2_{\infty\omega}$ that a linear order is finite. So let $\mu^+ = \mu \wedge \bigvee_{n<\omega} len_n(\prec) \in L^4_{\infty\omega}$ (see definition 2.1 for len). The models of μ^+ are exactly the finite odd multipedes. So any model whatever of μ^+ expands uniquely to a model of $\lambda(<)$. This means that $<$ is strongly implicitly definable in $L^\omega_{\infty\omega}$. But it is not explicitly definable, as $\mathcal{M}$ is not k-rigid for any finite k. Hence $L^\omega_{\infty\omega}$ fails the weak Beth property, and this is true even restricting to finite structures — which you knew already.

Quisani: Got you. Can't you use the same idea to show that the k-variable fragment L^k of first-order logic also fails the weak Beth property? At least, for any k large enough to write λ, μ in L^k. How many variables does λ need, by the way?

Author: I think five.

Quisani: Uhuh. Anyway, let k be any number large enough to write λ, μ. Find an odd multipede K in $\mathcal{M}$ that's not k-rigid. Suppose its spine has $< n$ elements. Consider the L^k-sentence

$$\mu \wedge \neg len_{\geq n}(\prec),$$

where $len_{\geq n}(\prec) \in L^2$ from definition 2.1 says there's a $\prec$-chain of at least n elements. The only models of this are odd multipedes, namely those with a spine of size $< n$. The problem with infinite models of μ disappears. So $\lambda(<)$ should strongly implicitly define $<$ on models of $\mu \wedge \neg len_{\geq n}(\prec)$. But $<$ is not explicitly definable in $L^k_{\infty\omega}$, never mind L^k, as $K \models \mu \wedge \neg len_{\geq n}(\prec)$ and K is not k-rigid.

Author: Yes, very nice. I believe the Budapest group have another (unpublished) example to show that L^k fails the weak Beth property. Shall we have another break?

Quisani: Yes please.

10 Temporal logic

Quisani [drinking coffee again]: What is your own interest in finite variable logics?

Author: I came to them from temporal logic. There, we model time by a *flow of time* — an irreflexive transitive partial order $(T, <)$, where T is the set of time points, and '$<$' the earlier–later relation on time points. Often we take the natural or real numbers, for example. We also have propositional letters or 'atoms', $p, q, r, \ldots$, to represent time-dependent facts, such as 'it is raining'. Each atom is given a truth value (true or false) at each time point.

Quisani: This assignment of truth values to atoms is not part of the original flow-of-time structure?

Author: That's right. Atoms are best regarded as propositional variables, akin to the variables in first-order formulas.

Quisani: But atoms don't get assigned to elements of a structure, do they?

Author: No, they don't. In temporal logic, each atom is assigned a subset of T, namely the set of time points where it's true.

Quisani: Then what?

Author: We then want to talk about the temporal behaviour of the atoms. To do this, we write temporal formulas, such as Fp and $S(p, q)$. They, too, get truth values at each time point, derived from the atoms' values at other points. The formula Fp is true at t iff p is true at some point $u \in T$ with $u > t$; and $S(p, q)$ is true at t iff there is $u < t$ where p is true, and moreover, q is true at all points strictly between u and t. We read Fp as 'in the future, p', and $S(p, q)$ as 'since p, q', for obvious reasons.

Quisani: And you can form more complex formulas like $F(p \wedge S(q, r))$, by substitution?

Author: Yes. In general, a temporal logic will specify a (finite) set of *connectives,* such as F (unary) and S (binary), and we form formulas by:

1. any atom $p, q, r, \ldots$ is a formula;
2. if A, B are formulas then so are $\neg A$ and $A \wedge B$;
3. if $\sharp$ is an n-ary connective in the given set of connectives, and $A_1, \ldots, A_n$ are formulas, then $\sharp(A_1, \ldots, A_n)$ is a formula.

Quisani: Temporal logic is of increasing importance for computer science, is it not?

Author: Yes, computer scientists and software engineers use it to write specifications with. So they need expressive temporal connectives. The study of expressive power of connectives is therefore important, and it connects with finite variable logics.

Quisani: This surprises me. Temporal logic seems a way of avoiding variables altogether!

Author: Yes, the 'pure' temporal logicians do sometimes say that sneaking variables in by the back door is inimical to true temporal logic. But we are interested in comparing its power with first-order logic, so we try to translate temporal formulas into first-order ones. For example, if we associate with each atom p a unary relation symbol P, then the atomic temporal formula p can be translated as the atomic first-order formula $P(x)$.

Quisani: The translation has a single free variable because the truth value of an atom depends on single time points?

Author: Yes. For the same reason, each translation of a temporal formula should be a first-order formula with a single free variable. For instance, Fp can be translated as $\exists y(y > x \wedge P(y))$, and $S(p,q)$ as $\exists y(y < x \wedge P(y) \wedge \forall z(y < z < x \rightarrow Q(z)))$. Each translation has x as the only free variable.

Quisani: A flood of variables! What about more complex temporal formulas?

Author: We can translate them too, by induction. Each connective comes with a *table,* explaining what it means. The table is in effect the first-order translation of a temporal formula consisting of the connective applied to distinct atoms. So an n-ary connective $\sharp$ has a table $\tau_\sharp(x, P_1, \ldots, P_n)$, which will be, by definition, the translation of $\sharp(p_1, \ldots, p_n)$.

Quisani: Right, I get the idea. For example, the table $\tau_S(x, P_1, P_2)$ of S will be $\exists y(y < x \wedge P_1(y) \wedge \forall z(y < z < x \rightarrow P_2(z)))$, the same formula we had a moment ago.

Author: Good. Now we can translate all formulas, by induction:

1. an atom p translates to $P(x)$;
2. if formulas A, B translate to $\psi_A(x)$ and $\psi_B(x)$ respectively, then $\neg A$ and $A \wedge B$ translate to $\neg\psi_A(x)$ and $\psi_A(x) \wedge \psi_B(x)$ respectively;
3. if $A_1, \ldots, A_n$ translate to $\psi_{A_1}(x), \ldots, \psi_{A_n}(x)$ respectively, and the table of the n-ary connective $\sharp$ is $\tau_\sharp(x, P_1, \ldots, P_n)$, then the formula $A = \sharp(A_1, \ldots, A_n)$ translates to $\tau_\sharp(x, P_1/\psi_{A_1}, \ldots, P_n/\psi_{A_n})$.

In this last formula, we substitute $\psi_{A_i}(v)$ for each atomic subformula $P_i(v)$ in $\tau_\sharp$, changing the variables if necessary to avoid clashes.

Quisani: This reminds me of something earlier ... but go on.

Author: So we obtain for each temporal formula $A(p_1, \ldots, p_n)$ a translation $\psi_A(x, P_1, \ldots, P_n)$, such that for any $(T, <)$, if the set of time points in T when the atom p_i is true is S_i, then A is true at t iff $(T, <) \models \psi_A(t, S_1, \ldots, S_n)$.

Quisani: What does it mean when you write $A(p_1, \ldots, p_n)$?

Author: Just that the atoms occurring in A are among $p_1, \ldots, p_n$.

Quisani: It seems to me that there's a possibility of translation only because every connective has a first-order table. Isn't this restrictive?

Author: Yes; second-order connectives and fixed-point extensions of classical temporal logic are considered too, especially in concurrency; the flow of time in these cases is usually the natural numbers. But first-order temporal logic is often quite adequate in power. For example, I believe Hans Kamp has conjectured that every English tense construction is expressible in first-order logic. You can also gain expressive power by using many-dimensional temporal logic [§13].

11 Expressive completeness

Author: Even in first-order temporal logic, there is the interesting question of how much of first-order logic can be obtained by translating temporal formulas.

Quisani: Ah! You are asking whether every first-order formula $\varphi(x, P_1, \ldots, P_n)$ is the translation of a temporal formula.

Author: Roughly, yes. For a given class $\mathcal{K}$ of flows of time, and a given set Ξ of connectives, the big question is: is it true that for every first-order formula $\varphi(x, P_1, \ldots, P_n)$ there is a temporal formula $A(p_1, \ldots, p_n)$, written with the connectives from Ξ, whose translation $\psi_A(x, P_1, \ldots, P_n)$ is $\mathcal{K}$-equivalent to φ?

Quisani: What is $\mathcal{K}$-equivalence?

Author: Formulas $\varphi(\bar{x}, P_1, \ldots, P_n), \psi(\bar{x}, P_1, \ldots, P_n)$ are said to be $\mathcal{K}$-equivalent if

$$(T, <) \models \forall \bar{x}(\varphi(\bar{x}, S_1, \ldots, S_n) \leftrightarrow \psi(\bar{x}, S_1, \ldots, S_n))$$

for all $(T, <) \in \mathcal{K}$ and all $S_1, \ldots, S_n \subseteq T$. I write $\varphi(\bar{x}, P_1, \ldots, P_n)$ to indicate that every free variable of φ is in the tuple $\bar{x}$; there is no requirement that every variable in $\bar{x}$ occurs free in φ.

Quisani: OK. So if you make $\mathcal{K}$ smaller or Ξ larger, the answer to your big question is more likely to be 'yes'.

Author: Right, the answer depends both on $\mathcal{K}$ and Ξ. If it is 'yes', we say that Ξ is *expressively complete* over $\mathcal{K}$.

Quisani: The word 'complete' is almost exhausted, I think.

Author: Yes, it is a bad choice, and it annoys people who work with second-order temporal logics, but we are probably stuck with it now.

Quisani: I guess we can trivially make a Ξ that's expressively complete over any class $\mathcal{K}$. For each formula $\varphi(x, P_1, \ldots, P_n)$ we add an n-ary connective $\sharp_\varphi$ with table φ.

Author: Yes, that's true; but we're mainly interested in *finite* sets Ξ.

Quisani: Is a finite set of connectives ever expressively complete over a class? It seems unlikely.

Author: As a trivial example, if $\mathcal{K}$ consists of a single one-point flow of time, so $\mathcal{K} = \{(\{t\}, \emptyset)\}$, then essentially we have propositional logic. We know from classical theory that the two connectives $\wedge$ and $\neg$ can express all other boolean connectives. So here, $\Xi = \emptyset$ is expressively complete over $\mathcal{K}$.

Quisani: We can take $\Xi = \emptyset$ because $\wedge$ and $\neg$ are always available in temporal logic?

Author: Yes, it's usual to include them by default — as we did.

Quisani: OK, then, are there any non-trivial examples of expressively complete connectives? It still seems unlikely to me.

Author: Kamp proved in [Kam] that S plus its 'mirror image', U, standing for 'until' and defined by replacing $<$ by $>$ in the table of S, are together expressively complete over the class of all *Dedekind complete linear orders* (such as the natural and real numbers, but not the rationals). This was the pioneering theorem in the area. [GPSS], [Gab2] (for the natural numbers) and [GHR2] have more recent proofs.

Quisani: Yet more completeness! Are there any other results like that?

Author: A few. In [GPSS] it was mentioned that adding two more connectives, U' and S' (they didn't give the semantics of these connectives) gives a set that's expressively complete over the class of all linear orders. [St] proves this, though I think the first full published proof will be in [GHR1] or [GHR2]; there are two different proofs in the latter. [GHR1,2] have other expressive completeness results for certain linear flows. There are also results for certain trees [Sch], and for many-dimensional logics [Ve1].

Quisani: How are these results proved?

Author: The proofs are always quite difficult. They're sometimes similar to eliminating quantifiers in classical logic, and you can view expressive completeness as a kind of quantifier-elimination result. Gabbay devised an interesting technique known as *separation* to prove expressive completeness results, which has also had practical applications in executable temporal logic. You can read about it in [Gab2].

12 H-dimension

Author: But what is interesting for us now, I think, is the more general question of when a given class $\mathcal{K}$ admits a finite expressively complete set Ξ of connectives.

Quisani: You could equally ask for which classes $\mathcal{K}$ is a given Ξ expressively complete.

Author: A fair point; I don't know if anyone's thought about that. I can, however, explain the connection of expressive completeness to finite variable logics. The memory that stirred in you a moment ago — what was it?

Quisani: When you were substituting the formulas ψ_{A_i} for P_i in $\tau_\sharp$ to obtain the translation of $\sharp(A_1, \ldots, A_n)$, it reminded me of LFP in §2, when we substituted φ^m for P in $\varphi(\bar{x}, P)$ to obtain φ^{m+1}.

Author: Yes, it's similar. Here, too, we can bound the number of variables needed, just as before. If for each $\sharp \in \Xi$, the table $\tau_\sharp$ uses only k variables, then by careful substitution we can write every translation ψ_A with k variables, for any temporal formula A written with the connectives from Ξ. Since for finite Ξ there will always be such a k, it follows that:

Theorem 12.1 (Gabbay, [Gab1]) *If the class $\mathcal{K}$ of flows of time admits a finite set of expressively complete connectives, then there is a finite k such that every first-order formula $\varphi(x, P_1, \ldots, P_n)$ with a single free variable is equivalent over $\mathcal{K}$ to some first-order formula $\varphi^*(x, P_1, \ldots, P_n)$ written with only k variables.*

Author: If the conclusion of the theorem holds, we say that the class $\mathcal{K}$ has *finite H-dimension;* the least such k is called the H-dimension of $\mathcal{K}$.

Quisani: Ah, so this is the finite-variable connection! [pause] What does the 'H' stand for?

Author: 'Henkin' — Leon Henkin worked on the proof theory of finite variable logics.

Quisani: It seems to me that we can use this theorem in two ways. First, we can use it to show that certain classes have finite H-dimension. For example, the class $\mathcal{L}$ of all linear orders has finite H-dimension, because U, S, U', S' are expressively complete for it.

Author: Yes; in fact, an inspection of their tables shows that they only need three variables, so that the H-dimension of $\mathcal{L}$ is at most 3. A game argument, given in [IK], will show that 3 is best possible, so that H-dim$(\mathcal{L}) = 3$. However, I must say that it's a very indirect method to use the expressive completeness of U, S, U' and S', a hard result, to prove that $\mathcal{L}$ has H-dimension 3, when another game argument gives this directly. Anyway, what was your second application?

Quisani: To use the contrapositive of the theorem to prove negative results. Are there classes that are known to have infinite H-dimension? If so, these would admit no finite expressively complete set of connectives.

Author: An excellent point. Indeed there are such classes; Gabbay showed in [Gab1] that the class of all flows of time is an example. Another is the class of trees with unbounded branching factor.

Quisani: How do you prove this?

Author: One way is to use a game, essentially that of [I1] and developed by [IK].

Quisani: Is this the 'game argument' you mentioned a moment ago, in connection with showing that H-dim$(\mathcal{L}) = 3$?

Author: Yes. Let M, N be *temporal structures*, i.e., flows of time with added assignments to an arbitrary finite number $P_1, \ldots, P_r$ of unary relation symbols. Let $k, n < \omega$. We let $G(M, N, k, n)$ be the game played by Spoiler, Duplicator on M, N using k pebbles, that stops after n rounds.

Quisani: Apart from stopping after a finite time, is this the same as the Barwise game we had earlier [§3]?

Author: Yes. Here there are no pebbles in play at the start. We regard M, N as structures in the signature $\{<, P_1, \ldots, P_r\}$, and the rules are as before. The winning condition for Spoiler is also as before: at some stage, the atomic formulas of this signature satisfied by the pebbles in M and in N should not be the same. Similarly, Duplicator wins if she survives each round: if the pebbles in M and N do always satisfy the same atomic formulas. We get the following theorem.

Theorem 12.2 (Immerman & Kozen, [IK]) *Suppose that $\mathcal{K}$ is a first-order-definable class of flows of time, and let $k < \omega$. Then $\mathcal{K}$ has H-dimension at most k iff for all temporal structures M, N whose underlying flows of time are in $\mathcal{K}$, if Duplicator has a winning strategy for $G(M, N, k, n)$ for all $n < \omega$ then she has a winning strategy in $G(M, N, n, n)$ for all $n < \omega$.*

Quisani: What does it mean to say that $\mathcal{K}$ is first-order-definable?

Author: Just that it is the class of all models of some first-order theory of signature $\{<\}$.

Quisani: I can see the '$\Rightarrow$'-direction of the theorem, I think. Thinking of the ordinary Ehrenfeucht–Fraïssé game, I'd guess that M and N agree on all first-order sentences of $L^k_{\infty\omega}$, i.e., on L^k-sentences, iff Duplicator has a winning strategy in $G(M, N, k, n)$ for all n. So the second condition of the theorem seems to be saying that if M and N agree on L^k then they agree on L^ω. But if $\mathcal{K}$ has H-dimension k then the L^k-theory should determine the L^ω-theory, so this follows.

Author: More or less, yes. The other direction uses first-order compactness, but apart from this the proof is similar to those in ordinary Ehrenfeucht–Fraïssé games. I'll not go into it now. But in fact you will not be able to prove theorem 12.2 correctly, because it's only true if you use Gabbay's original definition of H-dimension k, and not the one in theorem 12.1!

Quisani: Aagh! Why didn't you say so?

Author: I didn't want to complicate things too much. These matters get quite technical.

1. According to Gabbay [Gab1, Definition 1.3], a flow of time $(T, <)$ is said to have H-dimension k if k is the smallest number such that every formula $\varphi(x_1, \ldots, x_m, P_1, \ldots, P_n)$ written with $x_1, \ldots, x_m$ as free variable letters and any number of bound variable letters can be equivalently rewritten over $(T, <)$ using at most k bound variable letters and the same free $x_1, \ldots, x_m$.

2. A class $\mathcal{K}$ of flows of time can be said to have H-dimension k if every *sentence* $\sigma(P_1, \ldots, P_n)$ is equivalent over $\mathcal{K}$ to one written with only k variables, k being least such that this is possible. (For $\mathcal{K} = \{(T, <)\}$ this is equivalent to (1).)

3. The definition in theorem 12.1 above essentially replaces 'sentence' in (2) by 'formula with one free variable'.

Conditions (2) and (3) are not quite equivalent. Their differences are gone over in [HS].

Quisani: What is this paper [HS]?

Author: It is joint work with András Simon. We carefully compared several notions like these. Here's another well-known one: $\mathcal{K}$ is said to have the *k-variable property* if every formula $\varphi(x_1, \ldots, x_k, P_1, \ldots, P_n)$ is equivalent over $\mathcal{K}$ to a formula $\varphi^*(x_1, \ldots, x_k, P_1, \ldots, P_n)$ written with only k variables. This notion seems very similar to having H-dimension at most k, but in [HS] we constructed a class $\mathcal{K}$ that has H-dimension 3, and has a finite set of expressively complete connectives, but does not have the k-variable property for any finite k.

Quisani: Ouch! I see one has to be careful with definitions here.

Author: Yes, the exact form can make quite a difference. Actually, Flum [F] had related work already. He defines a first-order theory T in signature L to be *k-bounded* if (essentially) every first-order L-formula is T-equivalent to one where at most k distinct variables are bound in any branch of its formation tree.

Quisani: Is this related to the k-variable property?

Author: There is a conection, yes, though it seems not to be a completely tight one. If L is relational, of arity $< k$, then it can be shown (as in [HS, §3.2]) that T is k-bounded iff for all $n \geq k$, every formula $\varphi(x_1, \ldots, x_n)$ is T-equivalent to a formula $\varphi^*(x_1, \ldots, x_n)$ written with only n variables. We could say that T has the 'non-monadic' n-variable property for all $n \geq k$.

Flum also gives an example (suggested by Ziegler, and in some ways similar to that in [HS]) of a theory T that is not k-bounded for any k, but such that any formula can be equivalently rewritten over T using only one bound variable.

Quisani: You were going to use the game to prove that some classes don't have finite H-dimension.

Author: OK, let's see how to use theorem 12.2 to prove that the class $\mathcal{K}$ of all flows of time has infinite H-dimension. $\mathcal{K}$ is first-order-definable, by the axioms for irreflexive partial orders, so the theorem applies. Note that any *set* is essentially in $\mathcal{K}$: we interpret $<$ as $\emptyset$. We saw earlier [§7] that the k-pebble game cannot tell apart sets (structures in the empty signature) of size $\geq k$. Thus, if M, N are sets in $\mathcal{K}$ of size k and $k+1$ respectively, and we assign $\emptyset$ to all predicates $P_1, \ldots, P_r$ in each of them, then Duplicator has a winning strategy in $G(M, N, k, n)$ for all n.

Quisani: Right. But she can't win $G(M, N, k+1, k+1)$, as Spoiler can put the $k+1$ pebbles on different elements of N, and she has no response in M. So the second half of the theorem fails, for this k, and so $\mathcal{K}$ has H-dimension $> k$.

Author: Good; and this holds for all k, so $\mathcal{K}$ has infinite H-dimension. Hence by theorem 12.1, it admits no finite expressively complete set of connectives.

13 Many-dimensional temporal logic

Quisani: What about the converse of theorem 12.1? If $\mathcal{K}$ has finite H-dimension, is there necessarily a finite set of connectives that's expressively complete for it?

Author: Yes, if we allow temporal logics of the kind that I've already mentioned in passing, where formulas are evaluated at pairs of points ('intervals'), or in general, m-tuples of points. They are called *m-dimensional logics.* The atoms are still usually evaluated at single points, however.

Quisani: How can they be? Atoms are temporal formulas, too! Either the logic is m-dimensional or it isn't.

Author: Well, we cheat, by evaluating an atom at m points but requiring that its truth value only depends on the first of these points.

Quisani: But apart from this restriction on the assignments to the atoms, m-dimensional temporal logic is completely analogous to the original kind?

Author: Yes, completely. The tables of connectives $\sharp$ of m-dimensional temporal logic are of the form $\tau_\sharp(x_1, \ldots, x_m, P_1, \ldots, P_n)$, where $P_1, \ldots, P_n$ are m-ary relation symbols. The translations $\psi_A(x_1, \ldots, x_m, P_1, \ldots, P_n)$ of the formulas A of such a logic are constructed as before, by induction and substitution. But because of our restriction on assignments to atoms, we can arrange that they still have *unary* relation symbols $P_1, \ldots, P_n$ for the atoms. This is because, once we have formed the 'standard' translation of a formula, as before, we can replace its m-ary relations $P^*(v_1, \ldots, v_m)$, corresponding to atoms, by unary relations $P(v_1)$.

Quisani: Right: it's only v_1 that matters. But I think I would like to see an example.

Author: OK, here's an interesting example. Let's define the following set Ξ_k of connectives for a k-dimensional temporal logic:

- M_θ, for all maps $\theta : \{1, 2, \ldots, k\} \to \{1, 2, \ldots, k\}$;
- E;
- B;
- $\langle i \rangle$, for all $1 \le i \le k$.

M_θ and $\langle i \rangle$ are unary connectives, taking a single argument, and E and B are nullary (no arguments).

Quisani: How can you tell from this that the logic is k-dimensional?

Author: You can't. You have to look at the arity of the relations in the tables of the connectives to work out the dimension of the logic.

Quisani: What are the tables, then?

Author: Each connective $\sharp$ above has table $\tau_\sharp$ as follows:

- $\tau_{M_\theta}(x_1, \ldots, x_k, Q)$ is $\exists y_1 \ldots \exists y_k [Q(y_1, \ldots, y_k) \wedge \bigwedge_{i \le k} y_i = x_{\theta(i)}]$;
- $\tau_E(x_1, \ldots, x_k)$ is $x_1 = x_2$;
- $\tau_B(x_1, \ldots, x_k)$ is $x_1 > x_2$;

- $\tau_{\langle i \rangle}(x_1, \ldots, x_k, Q)$ is $\exists x_i Q(x_1, \ldots, x_k)$.

Quisani: So what would a sample formula look like?

Author: Well, if $k = 3$, consider the formula

$$A = \langle 2 \rangle [B \wedge M_{(1,2)} p \wedge \neg \langle 3 \rangle \neg (M_{(2,3)} B \wedge M_{(1,3)} B \rightarrow M_{(1,3)} q)].$$

(Here, I used the usual notation for permutations of the set $\{1, 2, 3\}$. For example, $(1, 3)$ denotes the permutation that swaps 1 and 3, leaving 2 fixed.)

Quisani: OK, and what about its translation?

Author: Slightly simplified, the initial translation is

$$\exists x_2 (x_1 > x_2 \wedge P^*(x_2, x_1, x_3) \wedge \forall x_3 (x_1 > x_3 \wedge x_3 > x_2 \rightarrow Q^*(x_3, x_2, x_1))).$$

Now we can replace the ternary P^*, Q^* by unary P, Q, and get the final translation

$$\psi_A = \exists x_2 (x_1 > x_2 \wedge P(x_2) \wedge \forall x_3 (x_1 > x_3 \wedge x_3 > x_2 \rightarrow Q(x_3))).$$

Quisani: Oooh — that's the table of $S(p, q)$! You seem to have built up the formula A almost by copying its intended translation!

Author: That's a shrewd observation. I did; and it can be done in general. Fix any $k \geq 2$. Take any formula $\varphi(x_1, \ldots, x_k, P_1, \ldots, P_n)$, where $P_1, \ldots, P_n$ are unary predicates. If φ is written with only the variables $x_1, \ldots, x_k$, then it can be 'back-translated' to a temporal formula $A(p_1, \ldots, p_n)$ using the connectives of Ξ_k. We define the map $\varphi \mapsto A$ by induction on φ:

- $P(x_i) \mapsto M_{(1,i)} p$;
- $(x_i = x_j) \mapsto M_\theta E$, where $\theta : \{1, \ldots, k\} \rightarrow \{1, \ldots, k\}$ is any map such that $\theta(1) = i$ and $\theta(2) = j$;
- $(x_i > x_j) \mapsto M_\theta B$, θ as above;
- if $\varphi_1 \mapsto A_1$ and $\varphi_2 \mapsto A_2$, then $\neg \varphi_1 \mapsto \neg A_1$, $\varphi_1 \wedge \varphi_2 \mapsto A_1 \wedge A_2$, and $\exists x_i \varphi_1 \mapsto \langle i \rangle A_1$ for $i = 1, \ldots, k$.

Then if $\varphi(x_1, \ldots, x_k, P_1, \ldots, P_n) \mapsto A$, we have

$$\vdash \forall x_1, \ldots, x_k (\varphi(x_1, \ldots, x_k, P_1, \ldots, P_n) \leftrightarrow \psi_A(x_1, \ldots, x_k, P_1, \ldots, P_n)).$$

Quisani: So the connectives of Ξ_k just mimic the first-order formula formation rules on k variables.

Author: Yes, and because of this they can mimic the construction of any first-order formula written with only k variables.

Quisani: This sounds like expressive completeness again.

Author: Right. So suppose a class $\mathcal{K}$ has finite H-dimension. How can we employ our logic?

Quisani: Well, suppose $\mathcal{K}$ has H-dimension k. This means that any first-order formula $\varphi(x_1, P_1, \ldots, P_n)$ can be equivalently rewritten over $\mathcal{K}$ using only variables $x_1, \ldots, x_k$. Ah, I see! For any such formula we can find a temporal formula A of Ξ_k whose translation is logically equivalent to it. So, combining the two parts, for any first-order formula $\varphi(x_1, P_1, \ldots, P_n)$ there is a temporal formula A of Ξ_k such that $\varphi(x_1, P_1, \ldots, P_n)$ and $\psi_A(x_1, \ldots, x_k, P_1, \ldots, P_n)$ are $\mathcal{K}$-equivalent.

Author: Excellent. This is a weak form of expressive completeness. Every $\varphi(x, P_1, \ldots, P_n)$ has a temporal equivalent, but in a k-dimensional temporal logic. You've shown that if we allow these many-dimensional logics, then we do get an equivalence in theorem 12.1. More exactly, we obtain:

Theorem 13.1 (Gabbay, [Gab1]) *If $\mathcal{K}$ is a class of flows of time, then $\mathcal{K}$ admits a finite set of expressively complete many-dimensional connectives iff it has finite H-dimension.*

Quisani: I see I proved the '$\Leftarrow$'-direction of theorem 13.1. How is the other direction proved?

Author: Just as in theorem 12.1. You might like to read [Ve1] or [GHR2] for more on many-dimensional logics; a great deal of work has been done.

14 Between finite H-dimension and expressive completeness

Quisani: You didn't really answer my question before. Does the converse of theorem 12.1 hold if we stick to one-dimensional temporal logic?

Author: No. Gabbay asked this in [Gab1]. It would be good to find a stronger condition than 'finite H-dimension' that's equivalent to the existence of a finite expressively complete set of one-dimensional connectives for a class of flows of time; but I showed in [Hodk] that finite H-dimension alone does not imply this.

Quisani: You found a counter-example?

Author: Yes. Shall we have a look at it?

Quisani: — OK, but I must be going soon.

Author: It uses *circles.* A circle is a structure $(C, <)$ satisfying the following axioms:

Trichotomy: For all $x, y \in C$, exactly one of $x < y, x = y, x > y$ holds.

Local linearity: For all $x \in C$, $\{y \in C : y > x\}$ and $\{y \in C : y < x\}$ (i.e., the future and past of x) are linearly ordered by $<$.

Circularity: $\forall xy(x < y \to \exists z(y < z \wedge z < x))$.

Quisani: Can you give an example of a circle?

Author: If C is the set of days of the week, and we let $x \leq y$ if and only if y is at most three days ahead of x, and $x > y$ otherwise, then $(C, <)$ is a circle of size seven.

Quisani: Hmm. I think there'll be a circle C_n of any finite size $n \geq 3$. Just arrange the n points equally-spaced around a real circle, and define $x < y$ iff it's quicker to go anticlockwise from x to y than to go clockwise.

Author: What if n is even? For example, with four points you'll get two pairs of opposite points, (N, S) and (E,W), say. How are opposite points such as N and S ordered?

Quisani: Ah. Well, let's only consider an odd number of points.

Author: That's better! In fact, the class $\mathcal{C} = \{C_n : 3 \leq n < \omega, n \text{ odd}\}$ is the counterexample we use.

Quisani: Wait a minute! Circles are not transitive orders, are they?

Author: No, non-trivial ones never are, because of the 'circularity' axiom.

Quisani: So $\mathcal{C}$ can't be a counterexample to '$\Leftarrow$' of theorem 12.1 — it's not a class of flows of time!

Author: Don't worry about this. It's easy to convert a circle into a flow of time, but the argument is a little technical and is otherwise not important. You can read it up in [Hodk] if you want. In any case, circles are intuitively quite natural 'flows of time'. Besides, you may have noticed that nowhere have we yet used the fact that we're dealing with flows of time.

Quisani: So all the definitions, of 'expressive completeness' and so on, can be made for general binary relations?

Author: Yes, even for general structures, and sometimes it's essential to do that. So it is legitimate to ask whether the converse of theorem 12.1 holds for structures such as circles.

Quisani: And it doesn't?

Author: No. We can show that (i) $\mathcal{C}$ has H-dimension 3 but (ii) $\mathcal{C}$ can't have any finite expressively complete set of (one-dimensional) connectives. For the first part you can start by using theorem 12.2 to show that the class of all circles has H-dimension ≤ 3; I'll leave it as an exercise for you. Of course, your approach, by counting variables in tables of connectives, isn't going to work here —

Quisani: — because $\mathcal{C}$ has no finite expressively complete set of connectives!

Author: Quite. We prove it — part (ii), that is — by contradiction. Fix a finite set Ξ of connectives, and suppose for contradiction that it's expressively complete for the class $\mathcal{C}$ of all circles of odd size. We can assume that $\top, \bot \in \Xi$; $\top$ is a nullary connective with table $x = x$, and $\bot$ is similar, with table $x \neq x$.

Quisani: Adding connectives to Ξ preserves expressive completeness, so it's OK to assume this.

Author: Yes. Now let A be a formula using the connectives of Ξ, without atoms. E.g., A might be $\sharp(\top, \top, \bot, \bot)$. We don't need to know the truth values of atoms, because A hasn't got any. If $C \in \mathcal{C}$, then the truth values of A at any two points $t, u \in C$ is the same.

Quisani: Right; C is completely symmetric so A can't distinguish between two points.

Author: Good: so we agree that any formula without atoms is true throughout C, or false throughout it. Now there are only finitely many formulas of the form $B = \sharp(B_1, \ldots, B_n)$, for $\sharp \in \Xi$ and $B_1, \ldots, B_n \in \{\top, \bot\}$.

Quisani: This is because Ξ is finite?

Author: Right. But $\mathcal{C}$ is infinite, so we can choose distinct circles $C, C' \in \mathcal{C}$ such that any such B is true in C iff it's true in C'.

Quisani: You are saying that C and C' are of different finite sizes, but they agree on all temporal formulas without atoms and of 'depth' 1.

Author: Yes, exactly. In fact, it follows that they agree on all formulas A without atoms, regardless of depth.

Quisani: Let me try to prove this — by induction on A. If $A \in \{\top, \bot\}$, it's clear, and the case of the boolean connectives is simple. So suppose $A = \sharp(A_1, \ldots, A_n)$ for $\sharp \in \Xi$ and A_i without atoms. Inductively, the two circles C and C' agree on the A_i. Replace each A_i in A by $\top$, if A_i is true in C and C', and by $\bot$ if it's false. We get a formula A', something like $\sharp(\top, \bot, \ldots, \bot)$, and it has the same truth value as A did in C, and in C'.

Author: Good, so far.

Quisani: Yes! — by choice of C and C', they agree on A', because it has depth 1! So they agree on A too. QED!

Author: OK. Now suppose that C has size n, and C' size $> n$. Let A be a formula of Ξ without atoms, whose translation $\psi_A(x_0)$ is $\mathcal{C}$-equivalent to

$$\nu(x_0) \stackrel{\text{def}}{=} \exists x_1, \ldots, x_n \bigwedge_{0 \leq i < j \leq n} x_i \neq x_j,$$

saying that there are $> n$ elements. We know there is such an A, because Ξ is expressively complete for $\mathcal{C}$.

Quisani: How do you know that A has no atoms?

Author: It's true that expressive completeness doesn't guarantee that A comes without atoms, but we can simply substitute $\top$ (say) for any atoms it does have. Clearly, the resulting formula will also be equivalent over $\mathcal{C}$ to ν.

Quisani: Right: now it's clear what to do. We know that A is true in C', so as it has no atoms, it'll be true in C too, and so C should have size $> n$. But it doesn't! This is the contradiction, and proves (ii).

Author: Well done!

Quisani: I'm tired now. I must go. Yuri never came, did he?

Author: No.

Quisani: Never mind, it was quite interesting. Thank you.

Acknowledgments

I warmly thank Yuri Gurevich for inviting me to write this, and for his excellent strategic advice and tactical suggestions. I am also very grateful to Marcelo Finger, Robin Hirsch and Mark Reynolds for computing and literary criticism, and Hajnal Andréka and István Németi for their permission to discuss the material on Beth's theorem in §8 prior to its formal publication. For this version, the Beth parts have been revised and expanded to take account of more recent views, and other minor alterations have been made. Finally, I thank the readers of the first version who sent me corrections, now I hope incorporated in the text.

References

AV1 S. Abiteboul, V. Vianu, *Datalog extensions for database queries and updates,* J. Comput. System Sci. 43 (1991) 62–124.

AV2 S. Abiteboul, V. Vianu, *Generic computation and its complexity,* in: Proc. 23rd ACM STOC, 1991.

AU A. Aho, J. Ullman, *Universality of data retrieval languages,* in: Proc. 6th ACM Symposium on Principles of Programming Languages, 1979, pp. 110–120.

ABN H. Andréka, J. van Benthem, I. Németi. Modal languages and bounded fragments of predicate logic. J. Philosophical Logic, to appear. Available at `http://www.fwi.uva.nl/research/illc/wwwreports.html`, report ML-96-03.

ADN H. Andréka, I. Düntsch, I. Németi, *Expressibility of properties of relations,* J. Symbolic Logic 60 (1995) 970–991.

B J. Barwise, *On Moschovakis closure ordinals,* J. Symbolic Logic 42 (1977) 292–296.

BGK A. Blass, Y. Gurevich, D. Kozen, *A zero-one law for logic with a fixed point operator,* Information and Control 67 (1985) 75–90.

CH A. Chandra, D. Harel, *Structure and complexity of relational queries,* J. Comput. System Sci. 25 (1982) 99–128.

DHK A. Dawar, L. Hella, Ph. Kolaitis, *implicit definability and infinitary logic in finite model theory,* Proc. LICS 1995, Springer LNCS 944 (1995) 624–635.

DLW A. Dawar, S. Lindell, S. Weinstein, *Infinitary logic and inductive definability over finite structures,* Information and Computation 119 (1995) 160–175.

E A. Ehrenfeucht, *An application of games to the completeness problem for formalized theories,* Fund. Math. 49 (1961) 128–141.

ER P. Erdös, A. Renyi, *On the evolution of random graphs,* Publ. Math. Inst. Hungar. Acad. Sci. 5 (1960) 17–61.

F J. Flum, *On bounded theories,* in: Proc. Computer Science Logic 91, Berne, eds. E. Boerger, G. Jaeger, H. Kleine Buening, M.M. Richter, Springer LNCS 626, pp. 111–118.

Gab1 D. M. Gabbay, *Expressive functional completeness in tense logic,* in: Aspects of Philosophical Logic, ed. U. Monnich, Reidel, Dordrecht, 1981, 91–117.

Gab2 D. M. Gabbay, *The declarative past and imperative future,* in: proceedings, Colloquium on Temporal Logic and Specification, ed. B. Banieqbal et al., Springer Lecture Notes in Computer Science 398, 1989.

GHR1 D. M. Gabbay, I. M. Hodkinson, M. A. Reynolds, *Temporal expressive completeness in the presence of gaps,* in: Logic Colloquium 1990, ed. J. Oikkonen, J. Väänänen, Lecture Notes in Logic 2, Springer-Verlag, 1993, pp. 89–121.

GHR2 D. M. Gabbay, I. M. Hodkinson, M. A. Reynolds, *Temporal logic, Volume 1,* Oxford University Press, 1994.

GPSS D. M. Gabbay, A. Pnueli, S. Shelah, J. Stavi, *On the temporal analysis of fairness,* 7th ACM Symposium on Principles of Programming Languages, Las Vegas, 1980, 163–173.

GaS D. Gale, F. M. Stewart, *Infinite games with perfect information,* in: Contributions to the theory of games, II, Ann. Math. Studies 28 (1953) pp. 245–266.

Gu1 Y. Gurevich, *Toward logic tailored for computational complexity,* in: Computation and Proof Theory, ed. M. M. Richter et al., Springer LNM 1104 (1984) pp. 175–216.

Gu2 Y. Gurevich, *Infinite Games,* Bull. EATCS, June 1989, 93–100.

Gu3 Y. Gurevich, *Zero-one laws,* Bull. EATCS, Feb. 1992, 90–106.

GuS1 Y. Gurevich, S. Shelah, *Fixed-point extensions of first-order logic,* Annals Pure Appl. Logic 32 (1986) 265–280.

GuS2 Y. Gurevich, S. Shelah, *On finite rigid structures,* J. Symbolic Logic 62 (1996) 549–562.

HKL L. Hella, Ph. Kolaitis, K. Luosto, *How to define a linear order on finite models,* In Proceedings, Ninth Annual IEEE Symposium on Logic in Computer Science, pages 40–49, Paris, France, 4–7 July 1994. IEEE Computer Society Press.

Hodk I. M. Hodkinson, *Finite H-dimension does not imply expressive completeness,* J. Philosophical Logic 23 (1994) 535–573.

HS I. Hodkinson, A. Simon, *The k-variable property is stronger than H-dimension k,* J. Philosophical Logic, to appear.

I1 N. Immerman, *Upper and lower bounds for first-order expressibility,* J. Comput. System Sci. 25 (1982) 76–98.

I2 N. Immerman, *Relational queries computable in polynomial time,* in: Proc. 14th ACM STOC (1982) pp. 147–152. Revised version in Information and Control 68 (1986) 86–104.

IK N. Immerman, D. Kozen, *Definability with bounded number of bound variables,* Proceedings IEEE 1987, 236–244.

Kam J.A.W. Kamp, *Tense Logic and the Theory of Linear Order,* Ph.D. thesis, University of California, 1968.

Kar C. R. Karp, *Finite-quantifier equivalence,* in: The theory of models, Proc. 1963 Internat. Symposium at Berkeley, ed. J.W. Addison et al., North-Holland, Amsterdam, 1965, pp. 407–412.

KV P. Kolaitis, M. Vardi, *Infinitary logics and 0–1 laws,* Information and Computation 98 (1992) 258–294.

Ko D. Kozen, *Results on the propositional μ-calculus,* in: Proc. 9th ICALP (1982) pp. 348–369.

L A. B. Livchak, *The relational model for systems of automatic testing,* Automatic documntation and math. linguistics 4 (1982) 17–19.

P B. Poizat, *Deux ou trois choses que je sais de L_n,* J. Symbolic Logic 47 (1982) 641–658.

R A. Rubin, Ph.D. thesis, CalTech, 1975.

Sch B-H Schlingloff, *Expressive completeness of temporal logic over trees,* preprint, 1990.

Sco D. Scott, *Logic with denumerably long formulas and finite strings of quantifiers,* in: *The Theory of Models,* ed. J.W. Addison, L. Henkin, A. Tarski, North-Holland, 1965, pp. 329–341.

St J. Stavi, *Functional completeness over the rationals,* manuscript, Bar-Ilan University, Ramat-Gan, Israel, 1979.

T V. A. Talanov, *Asymptotic solvability of logical formulas,* in: Combinatorial-algebraic methods in applied mathematics, pp. 118–126.

TK V. A. Talanov, V. V. Knyazev, *The asymptotic truth of infinite formulas,* in: Proc. All-Union Seminar on discrete and applied mathematics and its applications, pp. 56–61.

Va M. Vardi, *The complexity of relational query languages,* in: Proc. 14th ACM STOC, 1982, pp. 137–146.

Ve1 Y. Venema, *Expressiveness and completeness of an interval tense logic,* Notre Dame J. Formal Logic 31 (1990) 529–547.

Ve2 Y. Venema, *A modal logic for quantification and substitution,* Bull. IGPL 2 (1994) 31–43. Available from `http://theory.doc.ic.ac.uk:80/theory/forum/igpl/Bulletin/V2-1/` and `http://www.mpi-sb.mpg.de/igpl/Bulletin/V2-1/`