Algorithmic Model Theory
SS 2016

Prof. Dr. Erich Grädel and Dr. Wied Pakusa

Mathematische Grundlagen der Informatik
RWTH Aachen

This work is licensed under:
http://creativecommons.org/licenses/by-nc-nd/3.0/de/
Dieses Werk ist lizenziert unter:
http://creativecommons.org/licenses/by-nc-nd/3.0/de/

© 2019 Mathematische Grundlagen der Informatik, RWTH Aachen.
http://www.logic.rwth-aachen.de

Contents

2 Descriptive Complexity

In this chapter we study the relationship between logical definability and computational complexity on finite structures. In contrast to the theory of computational complexity we do not measure resources as time and space required to decide a property but the logical resources needed to define it. The ultimate goal is to characterize the complexity classes known from computational complexity theory by means of logic.

We first define what it means for a logic to capture a complexity class. One of the main results is due to Fagin, stating that existential second order logic captures NP. At this point it is still unknown whether there exists a logic capturing PTIME on all finite structures. However, a deeper analysis of the proof of Fagin's Theorem shows that SO-HORN logic captures PTIME on all *ordered* finite structures.

2.1 Logics Capturing Complexity Classes

To measure the complexity of a property of finite τ-structures, (for instance, graph) we have to represent the structures by words over a finite alphabet Σ, so that they can serve as inputs for Turing machines. For graphs, a natural choice is to take an adjacency matrix, and write it, row after row, as binary string. Notice that one and the same graph can have many different adjacency matrices, and thus many different encodings. Moreover, it is an important open problem to decide efficiently (i.e. in polynomial time) whether two different matrices represent the same graph, up to isomorphism. The choice of an adjacency matrix means to fix an enumeration of the vertices, and thus an *ordering* of the graph. The same is true for encoding finite structures of any fixed finite vocabulary τ: to define an encoding it is necessary to fix an ordering on the universe.

By $\mathrm{Ord}(\tau)$ we denote the class of all finite structures $(\mathfrak{A}, <)$, where $\mathfrak{A}$ is a τ-structure and $<$ is a linear order on its universe. For any structure $\mathfrak{A} \in \mathrm{Ord}(\tau)$ with universe of size n, and for any fixed k, we can identify A^k with the set $\{0, 1, \ldots, n^k - 1\}$. This is done by associating each k-tuple $\overline{a}$ with its rank in the lexicographic ordering induced by $<$ on A^k. When we talk about the $\overline{a}$-th element, we understand it in this sense.

Definition 2.1. An *encoding* is a function mapping ordered structures to words. An encoding $\mathrm{code}(\cdot) : \mathrm{Ord}(\tau) \to \Sigma^*$ is good if it identifies isomorphic structures, is polynomially bounded, first-order definable and allows to compute the values of atomic statements efficiently. Formally, the following abstract conditions must be satisfied.

- $\mathrm{code}(\mathfrak{A}, <) = \mathrm{code}(\mathfrak{B}, <)$ iff $(\mathfrak{A}, <) \cong (\mathfrak{B}, <)$.
- There is a fixed polynomial p such that $|\mathrm{code}(\mathfrak{A}, <)| \leq p(|A|)$ for all $(\mathfrak{A}, <) \in \mathrm{Ord}(\tau)$.
- For all $k \in \mathbb{N}$ and all $\sigma \in \Sigma$ there exists a first-order formula $\beta_\sigma(x_1, \ldots, x_k)$ of vocabulary $\tau \cup \{<\}$ so that for all $(\mathfrak{A}, <)$ and all $\overline{a} \in A^k$ it holds that

$$(\mathfrak{A}, <) \models \beta_\sigma(\overline{a}) \Leftrightarrow \text{ the } \overline{a}\text{-th symbol of } \mathrm{code}(\mathfrak{A}, <) \text{ is } \sigma.$$

- Given $\mathrm{code}(\mathfrak{A}, <)$ a relation symbol R of τ and a tuple $\overline{a}$ one can efficiently decide whether $\mathfrak{A} \models R\overline{a}$.

The meaning of "efficiently" in the last condition may depend on the context, here we understand it is as evaluated in linear time and logarithmic space.

Example 2.2. Let $\mathfrak{A} = (A, R_1, \ldots, R_m)$ be a structure with a linear order $<$ on A. Let $|A| = n$ and let s_i be the arity of R_i. Let ℓ be the maximal arity of $R_1, \ldots, R_m$. For each relation we define

$$\chi(R_j) = w_0 \ldots w_{n^{s_j}-1} 0^{n^\ell - n^{s_j}} \in \{0,1\}^{n^\ell},$$

where $w_i = 1$ if the i-th element of A^{s_j} is in R_j. Now

$$\mathrm{code}(\mathfrak{A}, <) := 1^n 0^{n^\ell - n} \chi(R_1) \ldots \chi(R_m).$$

When we say that an algorithm decides a class $\mathcal{K}$ of finite τ-structures we actually mean that it decides

$$\text{code}(\mathcal{K}) = \{\text{code}(\mathfrak{A}, <) : \mathfrak{A} \in \mathcal{K}, < \text{ a linear order on } A\}.$$

Definition 2.3. A *model class* is a class $\mathcal{K}$ of structures of a fixed vocabulary τ that is closed under isomorphism, i.e. if $\mathfrak{A} \in \mathcal{K}$ and $\mathfrak{A} \cong \mathfrak{B}$, then $\mathfrak{B} \in \mathcal{K}$.

A *domain* is an isomorphism closed class $\mathcal{D}$ of structures where the vocabulary is not fixed. For a domain $\mathcal{D}$ and vocabulary τ, we write $\mathcal{D}(\tau)$ for the class of τ-structures in $\mathcal{D}$.

Definition 2.4. Let L be a logic, Comp a complexity class and $\mathcal{D}$ a domain of finite structures. L captures Comp on $\mathcal{D}$ if

(1) For every vocabulary τ and every (fixed) sentence $\psi \in L(\tau)$, the model-checking problem for ψ on $\mathcal{D}(\tau)$ is in Comp.
(2) For every vocabulary τ and any model class $\mathcal{K} \subseteq \mathcal{D}(\tau)$ whose membership problem is in Comp, there exists a sentence $\psi \in L(\tau)$ such that

$$\mathcal{K} = \{\mathfrak{A} \in \mathcal{D}(\tau) : \mathfrak{A} \models \psi\}.$$

Notice that first-order logic is very weak, in this sense. Indeed, for every fixed first-order sentence $\psi \in \text{FO}(\tau)$, it can be decided efficiently, with logarithmic space, whether a given finite τ-structure is a model for ψ. However, FO does *not* capture LOGSPACE, not even on ordered structures. Indeed, the reachability problem on undirected graphs can be solved in LOGSPACE, but it is not first-order expressible.

2.2 Fagin's Theorem

Existential second-order logic (Σ_1^1) is the fragment of second-order logic consisting of formulae of the form $\exists R_1 \ldots \exists R_m \varphi$ where $\varphi \in \text{FO}$ and $R_1, \ldots, R_m$ are relation symbols. As we will see in this chapter, the logic Σ_1^1 captures the complexity class NP on the domain of all finite structures.

Example 2.5. 3-Colourability of a graph $G = (V, E)$ is in NP and indeed there is a Σ_1^1-formula defining the class of graphs which possess a valid

3-colouring:

$$\begin{aligned} \exists R \exists B \exists Y \quad (& \quad \forall x (Rx \vee Bx \vee Yx) \\ \wedge & \quad \forall x \forall y (Exy \rightarrow \neg((Rx \wedge Ry) \vee (Bx \wedge By) \vee (Yx \wedge Yy)) \end{aligned}$$

Theorem 2.6 (Fagin)**.** Existential second-order logic captures NP on the domain of all finite structures.

Proof. The proof consists of two parts. First, let $\psi = \exists R_1 \ldots \exists R_m \varphi \in \Sigma_1^1$ be an existential second-order sentence. We show that it can be decided in non-deterministic polynomial time whether a given structure $\mathfrak{A}$ is a model of ψ.

In a first step, we guess relations $R_1, \ldots, R_m$ on A. Recall that relations can be identified with binary strings of length n^{s_i}, where s_i is the arity of R_i. Then we check whether $(\mathfrak{A}, R_1, \ldots, R_m) \models \varphi$ which can be done in LOGSPACE and hence in PTIME. Thus the computation consists of guessing a polynomial number of bits followed by a deterministic polynomial time computation, showing that the problem is in NP.

For the other direction, let $\mathcal{K}$ be an isomorphism-closed class of τ-structures and let M be a non-deterministic TM deciding code($\mathcal{K}$) in polynomial time. We construct a sentence $\psi \in \Sigma_1^1$ such that for all finite τ-structure $\mathfrak{A}$ it holds that

$$\mathfrak{A} \models \psi \Leftrightarrow M \text{ accepts code}(\mathfrak{A}, <) \text{ for any linear order } < \text{ on } A.$$

Let $M = (Q, \Sigma, q_0, F^+, F^-, \delta)$ with accepting and rejecting states F^+ and F^- and $\delta : (Q \times \Sigma) \rightarrow \mathcal{P}(Q \times \Sigma \times \{0, 1, -1\})$ which, given an input code($\mathfrak{A}, <$), decides in non-deterministic polynomial time whether $\mathfrak{A}$ belongs to $\mathcal{K}$ or not. We assume that all computations of M reach an accepting or rejecting state after precisely n^k steps ($n := |A|$).

We encode a computation of M on code($\mathfrak{A}, <$) by relations $\overline{X}$ and construct a first-order sentence $\varphi_M \in \text{FO}(\tau \cup \{<\} \cup \{\overline{X}\})$ such that for every linear order $<$ there exists $\overline{X}$ with $(\mathfrak{A}, <, \overline{X}) \models \varphi_M$ if and only if code($\mathfrak{A}, <$) $\in L(M)$. To this end we show that

- If $\overline{X}$ represents an accepting computation of M on code($\mathfrak{A}, <$) then $(\mathfrak{A}, <, \overline{X}) \models \varphi_M$.

- If $(\mathfrak{A}, <, \overline{X}) \models \varphi_M$ then $\overline{X}$ *contains* a representation of an accepting computation of M on $\mathrm{code}(\mathfrak{A}, <)$.

Accordingly the desired formula ψ is then obtained via existential second-order quantification

$$\psi := (\exists <)(\exists \overline{X})(\text{"} < \text{ is a linear order "} \wedge \varphi_M).$$

Details:

- We represent numbers up to n^k as tuples in A^k.
- For each state $q \in Q$ we introduce a predicate

$$X_q := \{\bar{t} \in A^k : \text{ at time } \bar{t} \text{ the TM } M \text{ is in state } q\}.$$

- For each symbol $\sigma \in \Sigma$ we define

$$Y_\sigma := \{(\bar{t}, \bar{a}) \in A^k \times A^k : \text{ at time } \bar{t} \text{ the cell } \bar{a} \text{ contains } \sigma\}.$$

- The head predicate is

$$Z := \{(\bar{t}, \bar{a}) \in A^k \times A^k : \text{ at time } \bar{t} \text{ the head of } M \text{ is at position } \bar{a}\}.$$

Now φ_M is the universal closure of START $\wedge$ COMPUTE $\wedge$ END.

$$\text{START} := X_{q_0}(\overline{0}) \wedge Z(\overline{0}, \overline{0}) \wedge \bigwedge_{\sigma \in \Sigma} (\beta_\sigma(\overline{x}) \rightarrow Y_\sigma(\overline{0}, \overline{x})).$$

Recall that β_σ states that the symbol at position $\overline{x}$ in $\mathrm{code}(\mathfrak{A}, <)$ is σ. The existence of the formulae β_σ is guaranteed by the fact that $\mathrm{code}(\cdot)$ is a good encoding. In what follows, we denote by $\overline{x} + 1$ and $\overline{x} - 1$ a first-order formula that defines the direct successor and predecessor of the tuple $\overline{x}$ (in the lexicographical ordering on tuples that is induced by the linear order $<$), respectively.

$$\text{COMPUTE} := \text{NOCHANGE} \wedge \text{CHANGE}.$$

$$\begin{aligned}\text{NOCHANGE} := \bigwedge_{\sigma\in\Sigma} (Y_\sigma(\bar{t},\bar{x}) \wedge Z(\bar{t},\bar{y}) \wedge \bar{y} \neq \bar{x} \\ \wedge\, \bar{t}' = \bar{t}+1 \to Y_\sigma(\bar{t}',\bar{x})).\end{aligned}$$

$$\text{CHANGE} := \bigwedge_{q\in Q,\sigma\in\Sigma} (\text{PRE}[q,\sigma] \to \bigvee_{(q',\sigma',m)\in\delta(q,\sigma)} \text{POST}[q',\sigma',m]),$$

where

$$\text{PRE}[q,\sigma] := X_q(\bar{t}) \wedge Z(\bar{t},\bar{x}) \wedge Y_\sigma(\bar{t},\bar{x}) \wedge \bar{t}' = \bar{t}+1,$$

$$\text{POST}[q',\sigma',m] := X_{q'}(\bar{t}') \wedge Y_{\sigma'}(\bar{t}',\bar{x}) \wedge \text{MOVE}_m[\bar{t}',\bar{x}],$$

and

$$\text{MOVE}_m[\bar{t}',\bar{x}] := \begin{cases} \exists\bar{y}(\bar{x}-1 = \bar{y} \wedge Z(\bar{t}',\bar{y})), & m=-1 \\ Z(\bar{t}',\bar{x}), & m=0 \\ \exists\bar{y}(\bar{x}+1 = \bar{y} \wedge Z(\bar{t}',\bar{y})), & m=1. \end{cases}$$

Finally, we let

$$\text{END} := \bigwedge_{q\in F^-} \neg X_q(\bar{t}).$$

It remains to show the following two claims.
Claim 1. If $\overline{X}$ represents an accepting computation of M on code$(\mathfrak{A},<)$ then $(\mathfrak{A},<,\overline{X}) \models \varphi_M$. This, however, follows immediately from the construction of φ_M.
Claim 2. If $(\mathfrak{A},<,\overline{X}) \models \varphi_M$, then $\overline{X}$ contains a representation of an accepting computation of M on code$(\mathfrak{A},<)$. We define

$$\text{CONF}[C,j] := X_q(\bar{j}) \wedge Z(\bar{j},\bar{p}) \wedge \bigwedge_{i=0}^{n^k-1} Y_{w_i}(\bar{j},\bar{i})$$

for configurations $C = (w_0 \dots w_{n^k-1}, q, p)$ (tape content $w_0 \dots w_{n^k-1}$, state q, head position p), i.e. the conjunction of the atomic statements that hold for C at time j. Let C_0 be the input configuration of M on $\text{code}(\mathfrak{A}, <)$. Since $(\mathfrak{A}, <, \overline{X}) \models \text{START}$ it follows that

$$(\mathfrak{A}, <, \overline{X}) \models \text{CONF}[C_0, 0].$$

Since $(\mathfrak{A}, <, \overline{X}) \models \text{COMPUTE}$ and $(\mathfrak{A}, <, \overline{X}) \models \text{CONF}[C_i, t]$, for some $C_i \vdash C_{i+1}$ it holds that $(\mathfrak{A}, <, \overline{X}) \models \text{CONF}[C_{i+1}, t+1]$.

Finally, no rejecting configuration can be encoded in $\overline{X}$ because $(\mathfrak{A}, <, \overline{X}) \models \text{END}$. Thus an accepting computation

$$C_0 \vdash C_1 \vdash \ldots \vdash C_{n^k-1}$$

of M on $\text{code}(\mathfrak{A}, <)$ exists, with $(\mathfrak{A}, <, \overline{X}) \models \text{CONF}[C_i, i]$ for all $i \leq n^k - 1$. This completes the proof of Fagin's Theorem. Q.E.D.

Theorem 2.7 (Cook, Levin)**.** SAT is NP-complete.

Proof. Obviously SAT $\in$ NP. We show that for any Σ^1_1-definable class $\mathcal{K}$ of finite structures the membership problem $\mathfrak{A} \in \mathcal{K}$ can be reduced to SAT. By Fagin's Theorem, there exists a first-order sentence ψ such that

$$\mathcal{K} = \{\mathfrak{A} \in \text{Fin}(\tau) : \mathfrak{A} \models \exists R_1 \ldots \exists R_m \psi\}.$$

Given $\mathfrak{A}$, construct a propositional formula $\psi_{\mathfrak{A}}$ as follows.

- replace $\exists x_i \varphi$ by $\bigvee_{a \in A} \varphi[x_i/a]$,
- replace $\forall x_i \varphi$ by $\bigwedge_{a \in A} \varphi[x_i/a]$,
- replace all closed τ-atoms $P\bar{a}$ in ψ with their truth values,
- replace all atoms $R\bar{a}$ with propositional variables $P_{R\bar{a}}$.

This is a polynomial transformation and it holds that

$$\mathfrak{A} \in \mathcal{K} \Leftrightarrow \mathfrak{A} \models \exists R_1 \ldots \exists R_m \psi \Leftrightarrow \psi_{\mathfrak{A}} \in \text{SAT}.$$

Q.E.D.

2.3 Second Order Horn Logic on Ordered Structures

The problem of whether there exists a logic capturing PTIME on all finite structures is still open. However, on *ordered* finite structures, there are several known logical characterizations of PTIME. The most famous result of this kind is the one is the Theorem by Immerman and Vardi which states that the least fixed-point logic LFP captures PTIME on the class of all ordered finite structures. We shall discuss this later. We here present a different characterization of PTIME, in terms of second-order Horn logic SO-HORN, which follows from a careful analysis of the proof of Fagin's Theorem. Indeed, the construction that we used in that proof is not the original one by Fagin, but an optimized version that has been tailored so that it can be adapted to a proof that SO-HORN captures PTIME on ordered structures.

Definition 2.8. *Second-order Horn logic*, denoted by SO-HORN, is the set of second-order sentences of the form

$$Q_1 R_1 \ldots Q_m R_m \forall y_1 \ldots \forall y_s \bigwedge_{i=1}^{t} C_i,$$

where $Q_i \in \{\exists, \forall\}$ and the C_i are Horn clauses, i.e. implications

$$\beta_1 \wedge \ldots \wedge \beta_m \rightarrow H,$$

where each β_j is either a positive atom $R_k\bar{z}$ or an FO-formula that does not contain $R_1, \ldots, R_m$. H is either a positive atom $R_j\bar{z}$ or the Boolean constant 0.

Σ_1^1-HORN denotes the existential fragment of SO-HORN, i.e. the set of SO-HORN sentences where all second-order quantifiers are existential.

Theorem 2.9. Every sentence $\psi \in$ SO-HORN is equivalent to a sentence $\psi' \in \Sigma_1^1$-HORN.

Proof. It suffices to prove the theorem for formulae of the form

$$\psi = \forall P \exists R_1 \ldots \exists R_m \forall \bar{z} \varphi,$$

where φ is a conjunction of Horn clauses and $m \geq 0$ (for $m = 0$, the formula has the form $\forall P \forall \overline{z} \varphi$). Indeed we can then eliminate universal quantifiers beginning with the inner most one by considering only the part starting with that universal quantifier.

Lemma 2.10. A formula $\exists \overline{R} \forall \overline{z} \varphi(P, \overline{R}) \in \Sigma^1_1$-HORN holds for all relations P on a structure $\mathfrak{A}$ if and only if it holds for those P that are false at at most one point.

Proof. Let k be the arity of P. For every k-tuple $\overline{a}$, let $P^{\overline{a}} = A^k - \{\overline{a}\}$, i.e. the relation that is false at $\overline{a}$ and true at all other points. By assumption, there exist $\overline{R}^{\overline{a}}$ such that

$$(\mathfrak{A}, P^{\overline{a}}, \overline{R}^{\overline{a}}) \models \forall \overline{z} \varphi.$$

Now consider any $P \neq A^k$ and let $R_i := \bigcap_{\overline{a} \notin P} R_i^{\overline{a}}$. We show that $(\mathfrak{A}, P, \overline{R}) \models \forall \overline{z} \varphi$ where $\overline{R}$ is the tuple consisting of all R_i.

Suppose that this is false, then there exists a relation $P \neq A^k$, a clause C of φ and an assignment $\rho : \{z_1, \ldots, z_s\} \to A$ such that $(\mathfrak{A}, P, \overline{R}) \models \neg C[\rho]$. We proceed to show that in this case there exists a tuple $\overline{a}$ such that $(\mathfrak{A}, P^{\overline{a}}, \overline{R}^{\overline{a}}) \models \neg C[\rho]$ and thus

$$(\mathfrak{A}, P^{\overline{a}}, \overline{R}^{\overline{a}}) \models \neg \forall \overline{z} \varphi$$

which contradicts the assumption.

- If the head of $C[\rho]$ is $P\overline{a}$, then take $\overline{a} = \overline{u} \notin P$.
- If the head of $C[\rho]$ is $R_i\overline{u}$, then choose $\overline{a} \notin P$ such that $\overline{u} \notin R_i^{\overline{a}}$, which exists because $\overline{u} \notin R_i$.
- If the head is 0, take an arbitrary $\overline{a} \notin P$.

The head of $C[\rho]$ is clearly false in $(\mathfrak{A}, P^{\overline{a}}, \overline{R}^{\overline{a}})$. $P\overline{a}$ does not occur in the body of $C[\rho]$, because $\overline{a} \notin P$ and all atoms in the body of $C[\rho]$ are true in $(\mathfrak{A}, P, \overline{R})$. All other atoms of the form P_i that might occur in the body of the clause remain true for $P^{\overline{a}}$. Moreover, every atom $R_i\overline{v}$ in the body remains true if R_i is replaced by $R_i^{\overline{a}}$ because $R_i \subseteq R_i^{\overline{a}}$. This implies $(\mathfrak{A}, P^{\overline{a}}, \overline{R}^{\overline{a}}) \models \neg C[\rho]$. Q.E.D.

Using the above lemma, the original formula $\psi = \forall P \exists R_1 \dots \exists R_m \forall \overline{z} \varphi$ is equivalent to

$$\exists \overline{R}\ \forall \overline{z} \varphi [P\overline{u}/\overline{u} = \overline{u}] \wedge \forall \overline{y}\ \exists \overline{R}\ \forall \overline{z} \varphi [P\overline{u}/\overline{u} \neq \overline{y}].$$

This formula can be converted again to Σ^1_1-HORN; in the second part we push the external first-order quantifiers inside while increasing the arity of quantified relations by $|\overline{y}|$ to compensate it, i.e. we get

$$\exists \overline{R'}\ \forall \overline{y}\, \overline{z} \varphi [P\overline{u}/\overline{u} \neq \overline{y}, R(\overline{x})/R'(\overline{x}, \overline{y})].$$

Q.E.D.

Theorem 2.11. If $\psi \in$ SO-HORN, then the set of finite models of ψ, $\mathrm{Mod}(\psi)$, is in PTIME.

Proof. Given $\psi' \in$ SO-HORN, transform it to an equivalent sentence $\psi = \exists R_1 \dots \exists R_m \forall \overline{z} \bigwedge_i C_i$ in Σ^1_1-HORN. Given a finite structure $\mathfrak{A}$ reduce the problem of whether $\mathfrak{A} \models \psi$ to HORNSAT (as in the proof of the Theorem of Cook and Levin).

- Omit quantifiers $\exists R_i$.
- Replace the universal quantifiers $\forall z_i \eta(z_i)$ by $\bigwedge_{a \in A} \eta[z_i/a]$.
- If there is a clause that is already made false by this interpretation, i.e. $C = 1 \wedge \dots \wedge 1 \rightarrow 0$, reject ψ. Else interpret atoms $R_i\overline{u}$ as propositional variables.

The resulting formula is a propositional Horn formula with length polynomially bounded in $|A|$ and which is satisfiable iff $\mathfrak{A} \models \psi$. The satisfiability problem HORNSAT can be solved in linear time. Q.E.D.

Theorem 2.12 (Grädel)**.** On ordered finite structures SO-HORN and Σ^1_1-HORN capture PTIME.

Proof. We analyze the formula φ_M constructed in the proof of Fagin's Theorem in the case of a deterministic TM M. Recall that φ_M is the universal closure of START $\wedge$ NOCHANGE $\wedge$ CHANGE $\wedge$ END. START, NOCHANGE and END are already in Horn form. CHANGE has the form

$$\bigwedge_{q\in Q,\sigma\in\Sigma} (\mathrm{PRE}[q,\sigma] \rightarrow \bigvee_{(q',\sigma',m)\in\delta(q,\sigma)} \mathrm{POST}[q',\sigma',m]).$$

For a deterministic M for each (q,σ) there is a unique $\delta(q,\sigma) = (q',\sigma',m)$. In this case $\mathrm{PRE}[q,\sigma] \rightarrow \mathrm{POST}[q',\sigma',m]$ can be replaced by the conjunction of the Horn clauses

- $\mathrm{PRE}[q,\sigma] \rightarrow X_{q'}(\bar{t}')$
- $\mathrm{PRE}[q,\sigma] \rightarrow Y_{\sigma'}(\bar{t}',\bar{x})$
- $\mathrm{PRE}[q,\sigma] \wedge \bar{y} = \bar{x} + m \rightarrow Z(\bar{t}',\bar{y})$.

Q.E.D.

Remark 2.13. The assumption that a linear order is explicitly available cannot be eliminated, since linear orderings are not definable by Horn formulae.