ENS Paris-Saclay

Informatique, Logique L3

Decidable Theories

Revision 1*

David Baelde

May 11th, 2020

Abstract

We recall some basic notions regarding logical theories, introduce the quantifier elimination technique for showing that a theory is decidable (or complete), and illustrate it on the theory of discrete orders[1].

1 Definitions

Two notions of theories are used in logic: a theory is sometimes a set of closed formulas, sometimes a set of closed formulas that is closed under deduction, i.e. a set T such that $T = \{\phi \mid \text{ closed and } T \models \phi\}$. We shall use the second style in this document. We can thus equivalently write $\phi \in T$ or $T \models \phi$; T is recursive or $\{\phi \mid \phi \text{ closed and } T \models \phi\}$ is recursive.

Definition 1.1. Let $\mathcal{F}, \mathcal{P}$ be a language. An *$\mathcal{F}, \mathcal{P}$-theory* is a set of closed $\mathcal{F}, \mathcal{P}$-formulas that is closed under deduction.

*Modifications made since the original version of the document are indicated in margins with their corresponding revision.

[1] This presentation is based on earlier notes by Hubert Comon.

Theories are often generated from axioms: starting from some set of axioms one takes the set of its closed logical consequences. Reasonable set of axioms are recursive, but the theories that they generate may not be recursive: it is the case with the axioms of arithmetic, but also with the empty set of axioms.

It is also possible to define a theory from a model, i.e. by taking all the closed formulas that are satisfied in that structure.

Exercise 1.2. Consider the canonical structures over $\mathbb{N}$ and $\mathbb{Z}$ for the language of arithmetic. Show that the theories that they generate are disjoint.

Definition 1.3. A theory T is *complete* when, for any closed formula ϕ, either $\phi \in T$ or $\neg\phi \in T$.

The theory of a structure $\mathcal{S}$ is always complete. But complete theories may have very different models — for an extremely simple example, take $\mathcal{F} = \emptyset$ and $\mathcal{P} = \emptyset$, and observe that the theory of any structure is the set of valid formulas, but there are still non-isomorphic models. Models of a complete theory are however equivalent in the following sense:

Definition 1.4. Two $\mathcal{F}, \mathcal{P}$-structures $\mathcal{S}$ and $\mathcal{S}'$ are *elementarily equivalent* when, for all closed formulas ϕ, we have $\mathcal{S} \models \phi$ iff $\mathcal{S}' \models \phi$.

In other words, $\mathcal{S}$ and $\mathcal{S}'$ are elementarily equivalent iff they have the same theory.

We will use one last notion, related to language extensions. We say that a language $\mathcal{F}', \mathcal{P}'$ *extends* $\mathcal{F}, \mathcal{P}$ when $\mathcal{F} \subseteq \mathcal{F}'$ and $\mathcal{P} \subseteq \mathcal{P}'$. In that case, for any $\mathcal{F}', \mathcal{P}'$-structure $\mathcal{S}$ we define $\mathcal{S}|_{\mathcal{F},\mathcal{P}}$ as the $\mathcal{F}, \mathcal{P}$-structure obtained from $\mathcal{S}$ by forgetting the interpretation of predicates and function symbols that are not in $\mathcal{F}, \mathcal{P}$. We obviously have $\mathcal{S} \models \phi$ iff $\mathcal{S}|_{\mathcal{F},\mathcal{P}} \models \phi$ for any closed $\mathcal{F}, \mathcal{P}$-formula ϕ.

Definition 1.5. Let T be an $\mathcal{F}, \mathcal{P}$-theory and T' an $\mathcal{F}', \mathcal{P}'$-theory for some extension $\mathcal{F}', \mathcal{P}'$ of $\mathcal{F}, \mathcal{P}$. We say that T' is a *conservative extension* of T when $T \subseteq T'$ and, for any model $\mathcal{S}$ of T, there exists a model $\mathcal{S}'$ of T' such that $\mathcal{S} = \mathcal{S}'|_{\mathcal{F},\mathcal{P}}$. rev. 1

Exercise 1.6. Assume that T' is a conservative extension of T. Show that, for any $\mathcal{F}, \mathcal{P}$-formula ϕ, we have $\phi \in T$ iff $\phi \in T'$. rev. 1

2 Quantifier elimination

The quantifier elimination technique is a method that allows to reduce important decision problems to their quantifier-free version.

Definition 2.1. We say that a theory T admits quantifier elimination when one can compute, given a quantifier-free formula ϕ, another quantifier-free formula ψ such that $T \models (\forall x.\phi) \Leftrightarrow \psi$.

Note that, in the above definition, the formulas ϕ and ψ are not necessarily closed. Free variables in a logical consequence are implicitly universally quantified: explicitly, we require that $T \models \forall \vec{y}.\ (\forall x.\phi) \Leftrightarrow \psi$ where $\vec{y} = \mathrm{fv}(\phi, \psi) \setminus \{x\}$.

Theorem 2.2. Let T be a theory that admits quantifier elimination. If the set of closed quantifier-free formulas of T is recursive, then T is recursive.

Proof. Starting from some closed formula ϕ, we compute a closed quantifier-free formula ψ such that $T \models \phi \Leftrightarrow \psi$. The process is iterative, starting from $\phi = \phi_0$ and computing ϕ_{i+1} from ϕ_i as follows as long as ϕ_i still contains quantifiers:

1. Consider a quantified subformula $\mathcal{Q}x.\phi'$ of ϕ_i of maximal depth, i.e. such that ϕ' is quantifier-free.

2. Since T admits quantifier elimination we can compute some quantifier-free ψ' such that $T \models (\mathcal{Q}x.\phi') \Leftrightarrow \psi'$. This is immediate if $\mathcal{Q} = \forall$, otherwise one simply has to apply quantifier elimination to $\neg(\mathcal{Q}x.\phi')$.

3. Let ϕ_{i+1} be ϕ_i where $\mathcal{Q}x.\phi'$ has been replaced by ψ'.

This process terminates as ϕ_{i+1} has one less quantifier than ϕ_i. Let k be the first index such that ϕ_k is quantifier-free, and let $\psi := \phi_k$. We have $T \models \phi_i \Leftrightarrow \phi_{i+1}$ for all $0 \leq i < k$, hence $T \models \phi \Leftrightarrow \psi$. Hence $T \models \phi$ iff $T \models \psi$, and the latter is decidable by hypothesis. □

Theorem 2.3. Let T be a theory that admits quantifier elimination and such that, for any closed quantifier-free formula ϕ, one has $\phi \in T$ or $\neg\phi \in T$. Then T is complete.

Exercise 2.4. Adapt the proof of theorem 2.2 to prove theorem 2.3.

For some $\mathcal{F}, \mathcal{P}$-theories T, eliminating quantifiers is not feasible without introducing new symbols in the logical language. Then the above results only yield the decidability (or completeness) of T seen as some $\mathcal{F}', \mathcal{P}'$-theory. But the desired result for the original language can still be obtained thanks to the following observation.

Proposition 2.5. Let T be a theory and T' be a conservative extension of T. If T' is recursive, then so is T. If T' is complete, then so is T.

3 Application to the theory of discrete orders

The theory D of discrete orders is the theory over $\mathcal{F} = \emptyset$ and $\mathcal{P} = \{\leq, =\}$ generated by the axioms of equality and the axioms of fig. 1.

$$
\begin{array}{rl}
\text{(Refl)} & \forall x.\ x \leq x \\
\text{(Trans)} & \forall x, y, z.\ (x \leq y \wedge y \leq z) \Rightarrow x \leq z \\
\text{(Anti)} & \forall x, y.\ (x \leq y \wedge y \leq x) \Rightarrow x = y \\
\text{(Total)} & \forall x, y.\ x \leq y \vee y \leq x \\
\\
\text{(Min)} & \exists x. \forall y.\ x \leq y \\
\text{(Succ)} & \forall x. \exists y.\ x \leq y \wedge x \neq y \wedge \forall z.\ (x \leq z \wedge z \neq x) \Rightarrow y \leq z \\
\text{(Pred)} & \forall x.\ (\forall y.\ x \leq y) \vee \exists y.\ y \leq x \wedge y \neq x \wedge \forall z.\ y \leq z \Rightarrow (z = y \vee x \leq z)
\end{array}
$$

Figure 1: Axioms of discrete orders

This theory has a canonical model over $\mathbb{N}$, but it admits much more complex models. Despite this, we will see that it is recursive and complete. The next exercises aim to get accustomed to the axioms, appreciate the complexity of discrete orders, and avoid abusive reasoning in the technical development that follows.

Exercise 3.1. Give a model of the theory of discrete orders in which the order is not well-founded.

Exercise 3.2. Show that the predecessor axiom (Pred) is not a consequence of the other axioms of the theory of discrete orders.

Exercise 3.3. Show that the totality axiom (Total) is not a consequence of the other axioms.

Unsurprisingly, it is more convenient to talk about discrete orders using a successor function symbol. We therefore define the theory D_1 in the language $\mathcal{F}', \mathcal{P}'$ extended with a constant symbol 0 and a successor symbol $s(\_)$, generated by the axioms of fig. 1 and:

(Def-0) $\forall x.\ x = 0 \Leftrightarrow \forall y.\ x \leq y$

(Def-s) $\forall x, y.\ x = s(y) \Leftrightarrow \big(y \leq x \wedge y \neq x \wedge \forall z.\ (y \leq z \wedge y \neq z) \Rightarrow x \leq z\big)$

As we shall see, the axioms added to D_1 force the values of the new function symbols. This leads to the following conservativity result.

Proposition 3.4. The theory D_1 is a conservative extension of D.

Proof. Given a model $\mathcal{S}$ of D, we define an $\mathcal{F}', \mathcal{P}'$-structure $\mathcal{S}'$ such that $\mathcal{S}'|_{\mathcal{F},\mathcal{P}} = \mathcal{S}$. We only have to define $0_{\mathcal{S}'}$ and $s_{\mathcal{S}'}$.

Since $\mathcal{S} \models (\mathrm{Min})$ there exists some element a in the domain of $\mathcal{S}$ such that $S, \{x \mapsto a\} \models \forall y.\ x \leq y$: we choose $0_{\mathcal{S}'} := a$. We can already check that $\mathcal{S}' \models$ (Def-0): it follows from the axioms of equality and the fact that, for any b such that $\mathcal{S}', \{x \mapsto b\} \models \forall y.\ x \leq y$ we also have $\mathcal{S}, \{x \mapsto b, y \mapsto a\} \models x \leq y$ and symmetrically, so by (Anti) we conclude $\mathcal{S}, \{x \mapsto b, y \mapsto a\} \models x = y$ and $\mathcal{S}, \{x \mapsto b\} \models x = 0$.

For the successor, axiom (Succ) gives us, for any a, a value b such that $\mathcal{S}, \{x \mapsto a, y \mapsto b\} \models x \leq y \wedge x \neq y \wedge \forall z.\ (x \leq z \wedge z \neq x) \Rightarrow y \leq z$. It is exactly what is required by (Def-s), so we can simply set $s(a) = b$. Again, there might be several values b satisfying the above formula but the precise choice does not matter up to equality, because another b' would be such that $b \leq_{\mathcal{S}} b' \leq_{\mathcal{S}} b$. □

Proposition 3.5. The following formulas belong to D_1:

(Zero) $\forall y.\ 0 \leq y$

(Succ-<) $\forall x.\ x \leq s(x) \wedge x \neq s(s)$

(Succ-Tot) $\forall x.\ x = 0 \vee \exists y.\ x = s(y)$

(Succ-0) $\forall x.\ s(x) \neq 0$

(Succ-Succ) $\forall x, y.\ x \leq y \Leftrightarrow s(x) \leq s(y)$

(Discrete) $\forall x, y.\ (x \leq y \wedge x \neq y) \Rightarrow s(x) \leq y$

Proof. We leave it as an exercise. It would typically be verified by selecting a sound proof system and deriving each formula from the axioms of D_1. □

Proposition 3.6 (Predecessor, $\mathrm{pred}^n(a)$)**.** Let $\mathcal{S}$ be a model of the theory D_1, of domain $D_{\mathcal{S}}$, such that $=_{\mathcal{S}}$ is the identity over $D_{\mathcal{S}}$. Assume that

$$\mathcal{S}, \sigma\{x \mapsto a\} \models s^n(u) \leq x$$

for some $n \in \mathbb{N}$, $a \in D_{\mathcal{S}}$ and term u such that $x \notin \mathrm{fv}(u)$. There exists a unique element $\mathrm{pred}^n(a) \in D_{\mathcal{S}}$ such that $a = s^n(\mathrm{pred}^n(a))$ and

$$\mathcal{S}, \sigma\{x \mapsto \mathrm{pred}^n(a)\} \models u \leq x.$$

Moreover, when $\mathrm{pred}^n(a)$ is defined, we have $\mathrm{pred}^n(s_{\mathcal{S}}(a)) = s_{\mathcal{S}}(\mathrm{pred}^n(a))$.

Proof. By induction on n using the fact that $\mathcal{S}$ satisfies (Succ-Succ) and (Succ-Tot) there exists $b \in D_{\mathcal{S}}$ such that $a = s^n_{\mathcal{S}}(b)$ and $\mathcal{S}, \sigma\{x \mapsto b\} \models u \leq x$. It is unique by antisymmetry and by assumption on $=_{\mathcal{S}}$. □

Theorem 3.7. The theory D_1 admits quantifier elimination.

Proof. Consider $\exists x.\phi$ where ϕ is quantifier-free. We construct an equivalent quantifier-free ψ by transforming ϕ as follows:

1. We replace any atom of the form $u = v$ by $u \leq v \wedge v \leq u$, which is logically equivalent by the axioms of equality and the antisymmetry axiom.

2. We put the formula in negative normal form and replace any literal of the form $\neg(u \leq v)$ by $s(v) \leq u$, which is again equivalent by totality and the axioms on successor. The obtained formula does not contain any negation. rev. 1

3. We repeatedly simplify occurrences of $s(u) \leq s(v)$ into $u \leq v$, which is equivalent by (Succ-Succ). We change any occurrence of $s(u) \leq 0$ into $\bot$ and $0 \leq u$ into $\top$, which is equivalent by (Succ-0) and (Zero).

4. We simplify literals of the form $x \leq s^n(x)$ into $\top$ and $s^n(x) \leq x$ with $n > 0$ into $\bot$, which is justified by (Succ-$<$).

At this point we obtain a formula ϕ' that contains only literals of the form $y \leq u$ or $u \leq y$ where y does not occur in u. As explained during the construction of ϕ', we have $D \models (\exists x.\phi) \Leftrightarrow (\exists x.\phi')$.

We assume wlog that ϕ' is a conjunction of literals, and we set out to compute an equivalent formula (under D) which does not feature the existentially quantified variable x anymore – we can indeed put ϕ' in disjunctive normal form and replace any of its disjuncts by an equivalent formula without x. We also assume wlog that $0 \leq x$ is part of ϕ'. Our formula ϕ' can be reorganized into $\phi_0 \wedge \phi^+ \wedge \phi^-$ where:

- $x \notin \mathrm{fv}(\phi_0)$;

- ϕ^+ is a conjunction of literals of the form $u \leq s^n(x)$;
- ϕ^- is a conjunction of literals of the form $s^m(x) \leq v$.

Intuitively, $\exists x.\phi'$ holds iff $s^m(u) \leq s^n(v)$ for all of the above literals. We thus define ψ as follows:

$$\psi \stackrel{def}{=} \phi_0 \wedge \bigwedge_{(u \leq s^n(x)) \in \phi^+, (s^m(x) \leq v) \in \phi^-} s^m(u) \leq s^n(v)$$

We have $D \models (\exists x.\phi') \Rightarrow \psi$ by (Succ-Succ) and transitivity. Let us verify the converse.

Consider a model $\mathcal{S}$ of domain $D_\mathcal{S}$ and an assignemnt σ such that $\mathcal{S}, \sigma \models \psi$. We need to extend it into some $\sigma\{x \mapsto a_x\}$ such that $\mathcal{S}, \sigma\{x \mapsto a_x\} \models \phi$. We can assume wlog that $=_\mathcal{S}$ is the identity relation over $D_\mathcal{S}$, thanks to the axiom of equality.

For each $s^n(x) \leq t$ in ϕ^-, we have $s^n(0) \leq t$ in ψ because $0 \leq x$ is in ϕ^+. Since $\mathcal{S}, \sigma \models \psi$, the predecessors $\mathrm{pred}^n(t)$ of these terms are well-defined, and so is the following definition:

$$a_x \stackrel{def}{=} \min_{s^n(x) \leq t \in \phi^-} \mathrm{pred}^n(\llbracket t \rrbracket_\sigma)$$

In other words, we propose to interpret x by the greatest possible value that satisfies the inequalities of ϕ^-. This is expressed as the minimum of a finite (totally ordered) set, which is thus reached: there exists $s^N(x) \leq T$ in ϕ^- such that $a_x = \mathrm{pred}^N(\llbracket T \rrbracket_\sigma)$. Let us verify that $\mathcal{S}, \sigma\{a \mapsto a_x\} \models \phi$:

- We obviously have $\mathcal{S}, \sigma\{a \mapsto a_x\} \models \phi_0$ because $\mathcal{S}, \sigma \models \phi_0$ and x does not occur in that formula.
- For each $u \leq s^n(x)$ in ϕ^+ we have $s^k(T) \geq s^N(u)$ in ψ and thus this inequality is satisfied by $\mathcal{S}, \sigma$, and we have $\mathrm{pred}^N(\llbracket s^k(T) \rrbracket_\sigma) \geq_\mathcal{S} \llbracket u \rrbracket_\sigma$. Thus $\llbracket s^k(x) \rrbracket_\sigma \geq_\mathcal{S} \llbracket u \rrbracket_\sigma$, i.e. $\mathcal{S}, \sigma \models s^k(x) \geq u$.
- For each $s^m(x) \leq v$ in ϕ^- we have $a_x \leq_\mathcal{S} \mathrm{pred}^m(\llbracket v \rrbracket_\sigma)$ by definition of a_x, and thus $\mathcal{S}, \sigma \models s^m(x) \leq v$ by definition of pred^m. □

Example 3.8. We give a few examples of the equivalent formulas computed using our quantifier elimination technique. Let u, v and v' be arbitrary terms.

The formula $\exists x.\ u \leq s(x)$ is equivalent to $\top$. The formula $\exists x.\ s(x) \leq u$ is equivalent to $s(0) \leq u$. The formula

$$\exists x.\ s^2(x) \leq v \wedge s(x) \leq v' \wedge u \leq s(x)$$

is equivalent to

$$s^2(0) \leq v \wedge s(0) \leq v' \wedge u \leq v' \wedge s(u) \leq v.$$

Theorem 3.9. The theories D and D_1 are recursive and complete.

Proof. Closed literals over $\mathcal{F}'$ are of the form $s^n(0) = s^m(0)$ and $s^n(0) \leq s^m(0)$. We verify easily that $D_1 \models s^n(0) = s^m(0)$ iff $n = m$, and similarly for $\leq$. Given a closed quantifier-free formula ϕ, we can thus replace all its literals by either $\top$ or $\bot$ (and this choice is computable) to obtain an equivalent boolean expression that we can simply evaluate to know whether $D_1 \models \phi$.

If it is not the case, then it will obviously hold for $\neg\phi$. Hence D_1 is recursive and complete by theorem 2.2 and theorem 2.3. By proposition 2.5, D is also recursive and complete. □

As a corollary of the completeness result and exercise 3.1, we can also deduce that no $\mathcal{F}, \mathcal{P}$-formula can express the well-foundedness of the order.