

Hierarchical Reasoning in Local Theory Extensions*

Viorica Sofronie-Stokkermans**

Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, Saarbrücken, Germany
sofronie@mpi-sb.mpg.de

Abstract. We show that for special types of extensions of a base theory, which we call *local*, efficient hierarchic reasoning is possible. We identify situations in which it is possible, for an extension \mathcal{T}_1 of a theory \mathcal{T}_0 , to express the decidability and complexity of the universal theory of \mathcal{T}_1 in terms of the decidability resp. complexity of suitable fragments of the theory \mathcal{T}_0 (universal or $\forall\exists$). These results apply to theories related to data types, but also to certain theories of functions from mathematics.

1 Introduction

Many problems in mathematics and computer science and, in particular, problems involving reasoning in and about complex systems, can be reduced to proving the satisfiability of conjunctions of literals modulo some background theory. This theory may be quite complex: it can for instance be the extension of a base theory with additional functions (free, monotone, or recursively defined) or a combination of theories. It is therefore extremely important to find methods for efficient reasoning in extensions and combinations of theories.

In this paper we address the problem of reasoning in extensions of theories. We show that for special types of theory extensions, which we call *local*, hierarchic reasoning in which a theorem prover for the base theory is used as a “black box” is possible. Many theories important for computer science or mathematics are local extensions of a base theory. Examples are theories of data structures, e.g. theories of lists (or arrays cf. [6]); but also theories of monotone functions or of functions satisfying the Lipschitz conditions at a given point. We identify situations where the decidability of the universal theory of an extension of a theory is a consequence of the decidability of a certain fragment of the base theory.

The notion of local extension of a theory which we introduce in this paper generalizes the notion of *locality of a theory* introduced by Givan and McAllester [7,8], and of locality of equational theories studied by Ganzinger [4]. For local theories, validity of ground Horn clauses can be checked in polynomial time.

* This work was partly supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS). See www.avacs.org for more information.

** Many of the results presented here are the result of joint work with Harald Ganzinger[†].

Similar ideas also occurred in algebra. To prove that the uniform word problem for lattices is decidable in polynomial time, Skolem (1920) used the following idea: replace the lattice operations \vee and \wedge by ternary relations r_\vee and r_\wedge , required to be functional, but not necessarily total. The lattice axioms were translated to a relational form, by flattening them and then replacing every atom of the form $x \vee y \approx z$ with $r_\vee(x, y, z)$ (similarly for \wedge -terms). Additional axioms were added, stating that equality is an equivalence and that the relations are compatible with equality and functional. This new presentation, consisting only of Horn, function-free clauses, can be used for deciding in polynomial time the uniform word problem for lattices. The correctness and completeness of the method relies on the fact that every partially-ordered set (where \vee and \wedge are partially defined) embeds into a lattice. The idea described above was extended by Burris [2] to quasi-varieties of algebras. He proved that if a quasi-variety axiomatized by a set \mathcal{K} of Horn clauses has the property that *every finite partial algebra which is a partial model of the axioms in \mathcal{K} can be extended to a total algebra model of \mathcal{K}* then the uniform word problem for \mathcal{K} is decidable in polynomial time.

In [4], Ganzinger established a link between proof theoretic and semantic concepts for polynomial time decidability of uniform word problems. He defined two notions of locality for equational Horn theories, and established relationships between these notions of locality and corresponding semantic conditions, referring to embeddability of partial algebras into total algebras.

Our paper continues this line of research. Its main contributions are the following. First, we generalize in several ways the notion of locality of an equational theory:

- We consider local extensions $\mathcal{T}_0 \subseteq \mathcal{T}_1$, where the base theory \mathcal{T}_0 can be arbitrary. If \mathcal{T}_0 is the empty theory the original notion of locality is recovered.
- In defining locality of a theory extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ by a set \mathcal{K} of formulae we allow \mathcal{K} to be an arbitrary set of clauses (not necessarily Horn).

Second, we relate the extended notions of locality we consider with semantic properties, involving embeddability of partial algebras into total algebras.

Third, we use these results for hierarchic reasoning in local theory extensions, and identify situations in which this allows us to express the complexity of the universal theory of the extension as a function of the complexity of appropriate fragments of the base theory. We also sketch a possibility of extending the results beyond universally quantified formulae.

Structure of the Paper: Section 2 contains basic notions and notations. In Section 3, embeddability conditions are introduced and illustrated by examples; in Section 4 notions of locality of an extension are defined. The main contributions of the paper are contained in Sections 5 and 6: In Section 5 we establish links between various notions of locality of a theory extension and semantic properties, involving embeddability of partial algebras into total algebras. This helps to identify cases in which suitable locality conditions for theory extensions hold.

In Section 6 we establish parameterized complexity results of the universal theory of the extension in terms of the complexity of fragments of the base theory. Section 7 sketches a possibility of going beyond the universal fragment.

1.1 Idea

We illustrate the idea of our approach. Let $\mathbb{R} \cup \mathbf{L}_f^{\lambda_1} \cup \mathbf{L}_g^{\lambda_2}$ be the extension of the theory \mathbb{R} of reals with function symbols f, g satisfying the following axioms:

$$(\mathbf{L}_f^{\lambda_1}) \quad |f(x) - f(c_0)| \leq \lambda_1 \cdot |x - c_0| \quad (\mathbf{L}_g^{\lambda_2}) \quad |g(x) - g(c_0)| \leq \lambda_2 \cdot |x - c_0|$$

where c_0, λ_1 , and λ_2 are constants and the free variable x is, in both cases, implicitly universally quantified. We want to prove:

$$\mathbb{R} \cup (\mathbf{L}_f^{\lambda_1}) \cup (\mathbf{L}_g^{\lambda_2}) \models \forall x (|f(x) + g(x) - (f(c_0) + g(c_0))| \leq (\lambda_1 + \lambda_2) \cdot |x - c_0|).$$

Standard theorem provers for first order logic cannot always be used in such situations, as these can usually handle only approximations of the theory of real numbers. Provers for reals do not know about additional functions. The Nelson-Oppen method for reasoning in combinations of theories cannot be used either. The method we propose reduces the task of proving the formula above to the problem of checking the satisfiability of a set of constraints over \mathbb{R} as follows:

Negate. Let $\mathcal{K} = (\mathbf{L}_f^{\lambda_1}) \cup (\mathbf{L}_g^{\lambda_2})$. Note that $\mathbb{R} \cup \mathcal{K} \models \forall x C(x)$ (where $C(x)$ is $(|f(x) + g(x) - (f(c_0) + g(c_0))| \leq (\lambda_1 + \lambda_2) \cdot |x - c_0|)$) if and only if $\mathbb{R} \cup \mathcal{K} \cup G$ is unsatisfiable, where $G = \{|f(c) + g(c) - (f(c_0) + g(c_0))| \not\leq (\lambda_1 + \lambda_2) \cdot |c - c_0|\}$ is the set of ground clauses obtained from $\neg \forall x C(x)$ by Skolemization.

Take Ground Instances of Extension Axioms. We will show that $\mathbb{R} \cup \mathcal{K} \cup G$ is satisfiable if and only if $\mathbb{R} \cup \mathcal{K}[G] \cup G$ has a partial model in which all terms in the set $\text{st}(\mathcal{K}, G)$ consisting of all ground subterms in \mathcal{K} or in G are defined (and hence $f(c_0), g(c_0), f(c), g(c)$ are defined). ($\mathcal{K}[G]$ denotes the set of all instances of \mathcal{K} in which the terms starting with f or g are in $\text{st}(\mathcal{K}, G)$.) We compute $\mathcal{K}[G] \cup G$ and flatten replacing the ground terms starting with f or g with new constants:

$$\begin{aligned} (\mathcal{K}[G] \cup G)_{\text{flat}} := & f(c) \approx d \wedge f(c_0) \approx d_0 \wedge g(c) \approx e \wedge g(c_0) \approx e_0 \wedge \\ & |d - d_0| \leq \lambda_1 \cdot |c - c_0| \wedge |d_0 - d_0| \leq \lambda_1 \cdot |c_0 - c_0| \wedge \\ & |e - e_0| \leq \lambda_2 \cdot |c - c_0| \wedge |e_0 - e_0| \leq \lambda_2 \cdot |c_0 - c_0| \wedge \\ & |(d + e) - (d_0 + e_0)| \leq (\lambda_1 + \lambda_2) \cdot |c - c_0| \end{aligned}$$

Relational Translation. We compute the relational translation of the clauses above, using instead of f and g the functional binary predicates r^f and r^g :

$$\begin{aligned} (\mathcal{K}[G] \cup G)^* := & r^f(c, d) \wedge r^f(c_0, d_0) \wedge r^g(c, e) \wedge r^g(c_0, e_0) \wedge \\ & |d - d_0| \leq \lambda_1 \cdot |c - c_0| \wedge |d_0 - d_0| \leq \lambda_1 \cdot |c_0 - c_0| \wedge \\ & |e - e_0| \leq \lambda_2 \cdot |c - c_0| \wedge |e_0 - e_0| \leq \lambda_2 \cdot |c_0 - c_0| \wedge \\ & |(d + e) - (d_0 + e_0)| \leq (\lambda_1 + \lambda_2) \cdot |c - c_0| \end{aligned}$$

$$\text{Fun} := x_1 \approx x_2 \wedge R(x_1, y_1) \wedge R(x_2, y_2) \rightarrow y_1 \approx y_2 \quad \text{for } R = r^f \text{ or } R = r^g$$

We will show that we only need to consider those instances Fun^* of Fun in which the r^f resp. r^g literals are the ground literals already occurring in $(\mathcal{K}[G] \cup G)^*$,

and that $\mathbb{R} \cup (\mathcal{K}[G] \cup G)^* \cup \text{Fun}^*$ has a (relational) model if and only if the following set of constraints in \mathbb{R} is satisfiable:

$$\{\lambda_1 > 0, \lambda_2 > 0, (c \approx c_0 \rightarrow d \approx d_0), (c \approx c_0 \rightarrow e \approx e_0), (|d - d_0| \leq \lambda_1 \cdot |c - c_0|), (|e - e_0| \leq \lambda_2 \cdot |d - c_0|), |(d + e) - (d_0 + e_0)| \not\leq (\lambda_1 + \lambda_2) \cdot |c - c_0|\}.$$

We proved the unsatisfiability of this set of non-linear constraints using the REDLOG demo [3] (we considered the disjunction over the cases $c \leq c_0$ and $c > c_0$ and used quantifier elimination).

2 Basic Notions and Notations

Local Theories. The notion of *local theory* was introduced in [7,8] by Givan and McAllester. A local theory is a set of Horn clauses \mathcal{K} such that, for any ground Horn clause C , $\mathcal{K} \models C$ only if already $\mathcal{K}[C] \models C$ (where $\mathcal{K}[C]$ is the set of instances of \mathcal{K} in which all terms are subterms of ground terms in either \mathcal{K} or C). In [4], Ganzinger defined *locality* and *stable locality* of equational Horn theories, and established relationships between these notions of locality and embeddability of partial algebras into total algebras.

Total and Partial Algebras. We now present some generalities on partial algebras. Further details on partial algebras can be found in [1].

A *partial Σ -algebra* is a structure $(A, \{f_A\}_{f \in \Sigma})$, where A is a non-empty set and for every $f \in \Sigma$ with arity n , f_A is a partial function from A^n to A . A (*total*) *Σ -algebra* is a partial Σ -algebra where all functions f_A are total. In what follows we usually denote with the same symbol both an algebra and its support.

The notion of evaluating a term t with respect to a variable assignment $\beta : X \rightarrow A$ for its variables in a partial algebra A is the same as for total algebras, except that this evaluation is undefined if $t = f(t_1, \dots, t_n)$ and either one of $\beta(t_i)$ is undefined, or else $(\beta(t_1), \dots, \beta(t_n))$ is not in the domain of f_A .

A total map $h : A \rightarrow B$ between partial Σ -algebras A and B is called a *weak Σ -homomorphism* if whenever $f_A(a_1, \dots, a_n)$ is defined in A , then also $f_B(h(a_1), \dots, h(a_n))$ is defined in B and $h(f_A(a_1, \dots, a_n)) = f_B(h(a_1), \dots, h(a_n))$. A partial algebra A weakly embeds into a (total) algebra B if there exists an injective weak Σ -homomorphism from A to B .

In what follows we will consider structures $(A, \{f_A\}_{f \in \Sigma}, \{P_A\}_{P \in \text{Pred}})$, where Pred is a set of predicate symbols and $(A, \{f_A\}_{f \in \Sigma})$ is a partial Σ -algebra. We will refer to this type of structures as *Π -algebras* (or *Π -models*), where $\Pi = (\Sigma, \text{Pred})$. We say that a partial Π -algebra A weakly embeds into a (total) Π -algebra B if there exists $i : A \rightarrow B$ which is an injective weak Σ -homomorphism from A to B , and an embedding with respect to Pred .

We define *Evans validity* in structures $(A, \{f_A\}_{f \in \Sigma}, \{P_A\}_{P \in \text{Pred}})$, where Pred is a set of predicate symbols and $(A, \{f_A\}_{f \in \Sigma})$ is a partial Σ -algebra. In what follows the symbol \approx standing for formal equality will be considered to be symmetric also syntactically, so $s \approx t$ denotes at the same time also $t \approx s$. Let $\beta : X \rightarrow A$.

- (1) $(A, \beta) \models t \approx s$ if and only if (a) $\beta(t)$ and $\beta(s)$ are both defined and equal; or (b) $\beta(s)$ is defined, $t = f(t_1, \dots, t_n)$ and $\beta(t_i)$ is undefined for at least one of the direct subterms of t ; or (c) both $\beta(s)$ and $\beta(t)$ are undefined.
- (2) $(A, \beta) \models t \not\approx s$ if and only if (a) $\beta(t)$ and $\beta(s)$ are both defined and different; or (b) at least one of $\beta(s)$ and $\beta(t)$ is undefined.
- (3) $(A, \beta) \models P(t_1, \dots, t_n)$ if and only if (a) $\beta(t_1), \dots, \beta(t_n)$ are all defined and $(\beta(t_1), \dots, \beta(t_n)) \in P_A$; or (b) at least one of $\beta(t_1), \dots, \beta(t_n)$ is undefined.
- (4) $(A, \beta) \models \neg P(t_1, \dots, t_n)$ if and only if (a) $\beta(t_1), \dots, \beta(t_n)$ are all defined and $(\beta(t_1), \dots, \beta(t_n)) \notin P_A$; or (b) at least one of $\beta(t_1), \dots, \beta(t_n)$ is undefined.

(A, β) satisfies a clause C (notation: $(A, \beta) \models C$) if $(A, \beta) \models L$ for at least one literal L in C . A satisfies C (notation: $A \models C$) if $(A, \beta) \models C$ for all assignments β . A satisfies a set of clauses \mathcal{K} (notation: $A \models \mathcal{K}$) if $A \models C$ for all $C \in \mathcal{K}$.

The notion of *weak validity* is obtained from Evans validity by replacing conditions (1)(b) and (c) in the definition of truth of equality atoms with condition

- (b') at least one of $\beta(s), \beta(t)$ is undefined.

Validity of non-equality literals is the same. The notion of weak validity extends to clauses and sets of clauses in the usual way. We use the notation: $(A, \beta) \models_w L$ for a literal L ; $(A, \beta) \models_w C$; $A \models_w C$ for a clause C , etc.

Example 1. Let A be a partial Σ -algebra, where $\Sigma = \{\text{car}/1, \text{nil}/0\}$. Assume that nil_A is defined and $\text{car}_A(\text{nil}_A)$ is not defined. Then $A \not\models \text{car}(\text{nil}) \approx \text{nil}$ (since $\text{car}_A(\text{nil})$ is undefined in A , but nil is defined in A); and $A \models \text{car}(\text{nil}) \not\approx \text{nil}$, $A \models_w \text{car}(\text{nil}) \approx \text{nil}$, $A \not\models_w \text{car}(\text{nil}) \not\approx \text{nil}$ (because one term is not defined in A).

Theory Extensions. In this paper we consider extensions of theories, in which the signature is extended by new *function symbols*. For the sake of simplicity we assume that the set of predicate symbols remains unchanged in the extension. A theory can be regarded as a set of formulae. Then extension with a set of formulae is set union. In what follows we regard theories as sets of formulae.¹

Let \mathcal{T}_0 be an arbitrary theory with signature $\Pi_0 = (\Sigma_0, \text{Pred})$, where the set of function symbols is Σ_0 . We consider extensions \mathcal{T}_1 of \mathcal{T}_0 with signature $\Pi = (\Sigma, \text{Pred})$, where the set of function symbols is $\Sigma = \Sigma_0 \cup \Sigma_1$. We assume that \mathcal{T}_1 is obtained from \mathcal{T}_0 by adding a set \mathcal{K} of (universally quantified) clauses.

A *partial model* of \mathcal{T}_1 with totally defined Σ_0 function symbols is a partial Π -algebra A where (i) the reduct $A|_{\Pi_0}$ of A to the signature Π_0 is a model of \mathcal{T}_0 (in the classical sense, i.e. all operations in Σ_0 are total); (ii) A satisfies (in the Evans sense) all clauses in \mathcal{K} .

A partial Π -algebra A is a *weak partial model* of \mathcal{T}_1 with totally defined Σ_0 function symbols if (i) $A|_{\Pi_0}$ is a (classical) model of \mathcal{T}_0 and (ii') A weakly satisfies all clauses in \mathcal{K} .

¹ If a theory \mathcal{T}_0 is regarded as a collection of models, then its extension with a set \mathcal{K} of formulae consists of all structures (in the extended signature) which are models of \mathcal{K} and whose reduct to the signature of \mathcal{T}_0 is in \mathcal{T}_0 . All the notions defined in this paper can easily be reformulated to a setting in which \mathcal{T}_0 is a collection of models.

In what follows, if the base theory \mathcal{T}_0 and its signature are clear from the context, we will refer to *partial models* of \mathcal{T}_1 , resp. *weak partial models* of \mathcal{T}_1 . We will denote by $\text{PMod}(\Sigma_1, \mathcal{T}_1)$ the class of all partial models of \mathcal{T}_1 in which the functions in Σ_1 are partial, and all other function symbols are total; and by $\text{PMod}_w(\Sigma_1, \mathcal{T}_1)$ the class of all weak partial models of \mathcal{T}_1 in which the Σ_1 functions are partial and all the other function symbols are total. We denote by $\text{PMod}^f(\Sigma_1, \mathcal{T}_1)$, resp. $\text{PMod}_w^f(\Sigma_1, \mathcal{T}_1)$ the class of all finite partial models (resp. weak partial models) of \mathcal{T}_1 , with total Σ_0 functions, and partial Σ_1 functions. $\text{Mod}(\mathcal{T}_1)$ denotes the class of all models of \mathcal{T}_1 in which all functions in $\Sigma_0 \cup \Sigma_1$ are totally defined. Note that $\text{Mod}(\mathcal{T}_1) \subseteq \text{PMod}(\Sigma_1, \mathcal{T}_1) \subseteq \text{PMod}_w(\Sigma_1, \mathcal{T}_1)$.

3 Embeddability

For theory extensions $\mathcal{T}_0 \subseteq \mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$, where \mathcal{K} is a set of clauses, and for the classes of partial algebras mentioned above we consider the following conditions:

- (Emb) Every $A \in \text{PMod}(\Sigma_1, \mathcal{T}_1)$ weakly embeds into a total model of \mathcal{T}_1 .
- (Emb_w) Every $A \in \text{PMod}_w(\Sigma_1, \mathcal{T}_1)$ weakly embeds into a total model of \mathcal{T}_1 .

Weaker conditions, which only refer to embeddability of *finite* partial models can also be defined. These will be denoted by (Emb^f), resp. (Emb_w^f). We also define stronger notions of embeddability, which we call *completability*:

- (Comp) Every $A \in \text{PMod}(\Sigma_1, \mathcal{T}_1)$ weakly embeds into a total model B of \mathcal{T}_1 such that $A|_{\Pi_0}$ and $B|_{\Pi_0}$ are isomorphic (or, more generally: elementarily equivalent).

(Comp^f), (Comp_w) and (Comp_w^f) are defined analogously.

Example 1. *We present several examples of theory extensions for which embedding conditions among those mentioned above hold.*

- (1) **Shallow extensions:** *Suppose that $\mathcal{T}_0 \subseteq \mathcal{T}_1$ is a shallow theory extension, i.e. $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$, where \mathcal{K} is a set consisting only of clauses in which partial function symbols occur only in equality atoms, only positively and only at the root of terms. Assume that all extension functions are declared partial. Then the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies the embeddability condition (Comp) [6]. Extensions with functions defined by tail recursions are shallow [6].*
- (2) **Extensions with free functions:** *Any extension of a theory \mathcal{T}_0 with a set of free function symbols satisfies condition (Comp_w).*
- (3) **Extensions with selector functions:** *Let \mathcal{T}_0 be a theory with signature $\Pi_0 = (\Sigma_0, \text{Pred})$, let $c \in \Sigma_0$ with arity n , and let $\Sigma_1 = \{s_1, \dots, s_n\}$ consist of n unary function symbols. Let $\mathcal{T}_1 = \mathcal{T}_0 \cup \text{Sel}$ (a theory with signature $\Pi = (\Sigma_0 \cup \Sigma_1, \text{Pred})$) be the extension of \mathcal{T}_0 with the set Sel of clauses below. Then the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies condition (Comp).*

If in addition \mathcal{T}_0 satisfies the (universally quantified) formula $\text{Inj}(c)$ (i.e. c is injective in \mathcal{T}_0) then the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies condition (Comp_w) .

$$\begin{aligned}
 (\text{Sel}) \quad & s_1(c(x_1, \dots, x_n)) \approx x_1 \\
 & \dots \\
 & s_n(c(x_1, \dots, x_n)) \approx x_n \\
 & x \approx c(x_1, \dots, x_n) \rightarrow c(s_1(x), \dots, s_n(x)) \approx x \\
 (\text{Inj}(c)) \quad & c(x_1, \dots, x_n) \approx c(y_1, \dots, y_n) \rightarrow \left(\bigwedge_{i=1}^n x_i \approx y_i \right)
 \end{aligned}$$

- (4) **Extensions with monotone functions:** Let \mathcal{T}_0 be one of the following theories: (1) \mathcal{P} (posets), (2) \mathcal{T} (totally-ordered sets), (3) \mathcal{DO} (dense totally-ordered sets), (4) \mathcal{S} (semilattices), (5) \mathcal{L} (lattices), (6) \mathcal{DL} (distributive lattices), (7) \mathcal{B} (Boolean algebras), (8) \mathbb{R} (theory of reals).

Let Mon_f be the monotonicity axiom:

$$(\text{Mon}_f) \quad \bigwedge_{i=1}^n x_i \leq y_i \rightarrow f(x_1, \dots, x_n) \leq f(y_1, \dots, y_n).$$

The extension $\mathcal{T}_0 \subseteq \mathcal{T}_0 \cup \text{Mon}_f$ satisfies condition (Emb_w) in the cases (1)–(5); satisfies condition (Comp_w^f) in the cases (6) and (7); and satisfies condition (Comp_w) in case (8).

- (5) **Lipschitz functions:** The extension $\mathbb{R} \subseteq \mathbb{R} \cup (\mathbf{L}_f^\lambda)$ of \mathbb{R} with a unary function which is λ -Lipschitz in a point x_0 (for $\lambda > 0$) satisfies condition (Comp_w) .

$$(\mathbf{L}_f^\lambda) \quad \forall x |f(x) - f(x_0)| \leq \lambda \cdot |x - x_0|$$

4 Local Theory Extensions

We now define two notions of locality of a theory extension which generalize the notion of local equational theory studied by Ganzinger in [4] and of locality of a theory in general [7,8].

Let \mathcal{K} be a set of clauses in the signature $\Pi = (\Sigma_0 \cup \Sigma_1, \text{Pred})$. In what follows, when we refer to sets G of ground clauses we assume that they are in the signature $\Pi^c = (\Sigma \cup \Sigma_c, \text{Pred})$, where $\Sigma = \Sigma_0 \cup \Sigma_1$, and Σ_c is a set of new constants.

If Ψ is a set of ground $\Sigma_0 \cup \Sigma_1 \cup \Sigma_c$ -terms, where Σ_c is a set of (new) constants, we denote by \mathcal{K}_Ψ the set of all instances of \mathcal{K} in which all terms starting with a Σ_1 -function symbol are ground terms in Ψ . We denote by \mathcal{K}^Ψ the set of all instances of \mathcal{K} in which all variables occurring below a Σ_1 -function symbol are instantiated with ground terms in the set $T_{\Sigma_0}(\Psi)$ of Σ_0 -terms generated by Ψ .

If G is a set of ground clauses and $\Psi = \text{st}(\mathcal{K}, G)$ is the set of ground subterms occurring in either \mathcal{K} or G then we write $\mathcal{K}[G] := \mathcal{K}_\Psi$, and $\mathcal{K}^{[G]} := \mathcal{K}^\Psi$.

We identify the following types of locality of a theory extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$, where $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$ with \mathcal{K} a set of (universally quantified) clauses:

- (Loc) For every set G of ground clauses $\mathcal{T}_1 \cup G \models \perp$ iff $\mathcal{T}_0 \cup \mathcal{K}[G] \cup G$ has no weak partial model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined.
- (SLoc) For every set G of ground clauses $\mathcal{T}_1 \cup G \models \perp$ iff $\mathcal{T}_0 \cup \mathcal{K}^{[G]} \cup G$ has no partial model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined.

Weaker notions (Loc^f), resp. (SLoc^f) can be defined if we require that the respective conditions hold only for *finite* sets G of ground clauses. An extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ is *local* (*stably local*) if it satisfies condition (Loc^f) (resp. (SLoc^f)). A local (stably local) theory [4] is a local (stably local) extension of the empty theory.

5 Locality and Embeddability

We establish links between the notions of locality and embeddability. This extends the results established for local equational theories in [4].

Let \mathcal{T}_0 be an arbitrary theory with signature $\Pi_0 = (\Sigma_0, \text{Pred})$. Let \mathcal{T}_1 be an extension of \mathcal{T}_0 by a set \mathcal{K} of clauses in signature $\Pi = (\Sigma_0 \cup \Sigma_1, \text{Pred})$. Under appropriate assumptions, locality implies embeddability. The converse, which is proved in this section, will be used to provide examples of local theory extensions.

5.1 Flattening of Goals

We first show that in the locality condition we can assume, w.l.o.g., that G consists only of flat and linear (resp. purified) clauses.

We say that a ground clause is Σ_1 -*flat* if only constants appear as arguments of function symbols in Σ_1 . A Σ_1 -flat ground clause is Σ_1 -*linear* if whenever a constant occurs in two terms in the clause starting with function symbols in Σ_1 , the two terms are identical, and if no term starting with a function symbol in Σ_1 contains two occurrences of the same constant.

Any set G of ground clauses in a signature Σ containing Σ_1 can be transformed into a set $G_{\text{fin}(\Sigma_1)}$ of ground clauses in which subterms starting with function symbols in Σ_1 are flat and linear. This can be done by introducing, in a bottom-up manner, new constants for subterms occurring below functions in Σ_1 , and adding the corresponding definitions to the set of clauses. A set G of ground clauses can be transformed into a *purified* set of clauses $G_{\text{sep}(\Sigma_1)}$ (i.e. the function symbols in Σ_1 are separated from the other symbols) by introducing, in a bottom-up manner, new constants c_t for subterms $t = f(g_1, \dots, g_n)$ with $f \in \Sigma_1$, g_i ground $\Sigma_0 \cup \Sigma_c$ -terms (where Σ_c is a set of constants which contains the constants introduced by flattening), together with corresponding definitions $c_t \approx t$. These transformations preserves satisfiability and unsatisfiability with respect to total algebras, and also with respect to partial algebras in which all ground subterms which are flattened are defined.

Lemma 1. *Let \mathcal{K} be a set of clauses containing only Σ_1 -flat ground subterms. Assume that for any set G of Σ_1 -flat and Σ_1 -linear (resp. purified, resp. flat,*

linear and purified) ground clauses, if $\mathcal{T}_0 \cup \mathcal{K} \cup G \models \perp$ then $\mathcal{T}_0 \cup \mathcal{K}[G] \cup G$ has no partial algebra model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined. Then the extension $\mathcal{T}_0 \subseteq \mathcal{T}_0 \cup \mathcal{K}$ satisfies condition (Loc).

5.2 Embeddability of Weak Partial Models Implies Locality

We prove that for extensions which are Σ_1 -flat and Σ_1 -linear embeddability of weak partial models into total models implies locality.

A non-ground clause is Σ_1 -flat if function symbols (including constants) do not occur as arguments of function symbols in Σ_1 . A Σ_1 -flat non-ground clause is called Σ_1 -linear if whenever a variable occurs in two terms in the clause which start with function symbols in Σ_1 , the two terms are identical, and if no term which starts with a function in Σ_1 contains two occurrences of the same variable.

Theorem 2. *Let \mathcal{K} be a set of clauses which are Σ_1 -flat and Σ_1 -linear, and let $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$. Then the following hold:*

- (1) *If the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Emb_w) then it satisfies (Loc).*
- (2) *Assume that \mathcal{T}_0 is a locally finite universal theory, and that \mathcal{K} contains only finitely many ground subterms. If the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Emb_w^f), then $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Loc^f).*

Proof: (1) Assume that $\mathcal{T}_0 \cup \mathcal{K}$ is not a local extension of \mathcal{T}_0 . Then there exists a set G of ground clauses (with additional constants) such that $\mathcal{T}_0 \cup \mathcal{K} \cup G \models \perp$ but $\mathcal{T}_0 \cup \mathcal{K}[G] \cup G$ has a weak partial model P in which all terms in $\text{st}(\mathcal{K}, G)$ are defined. By Lemma 1 we can assume w.l.o.g. that $G = G_0 \cup G_1$, where G_0 contains no function symbols in Σ_1 and G_1 consists of ground unit clauses of the form $f(c_1, \dots, c_n) \approx c$, where c_1, \dots, c_n, c are constants in $\Sigma_0 \cup \Sigma_c$ and $f \in \Sigma_1$.²

We construct another structure, A , having the same support as P , which inherits all relations in **Pred** and all maps in $\Sigma_0 \cup \Sigma_c$ from P , but on which the domains of definition of the Σ_1 -functions are restricted as follows: for every $f \in \Sigma_1$, $f_A(a_1, \dots, a_n)$ is defined if and only if there exist constants c^1, \dots, c^n such that $f(c^1, \dots, c^n)$ is in $\text{st}(\mathcal{K}, G)$ and $a^i = c_P^i$ for all $i \in \{1, \dots, n\}$. In this case we define $f_A(a_1, \dots, a_n) := f_P(c_P^1, \dots, c_P^n)$. The reduct of A to $(\Sigma_0 \cup \Sigma_c, \text{Pred})$ coincides with that of P . Thus, A is a model of $\mathcal{T}_0 \cup G_0$. By the way the operations in Σ_1 are defined in A it is clear that A satisfies G_1 , so A satisfies G .

To show that $A \models_w \mathcal{K}$ we use the fact that if D is a clause in \mathcal{K} and $\beta : X \rightarrow A$ is an assignment in which $\beta(t)$ is defined for every term t occurring in D , then (by the way Σ_1 -functions are defined in A) we can construct a substitution σ with $\sigma(D) \in \mathcal{K}[G]$ and $\beta \circ \sigma = \beta$. As $(P, \beta) \models_w \sigma(D)$ we can infer $(A, \beta) \models_w D$.

As $A \models_w \mathcal{K}$, A weakly embeds into a total algebra B satisfying $\mathcal{T}_0 \cup \mathcal{K}$. But then $B \models G$, so $B \models \mathcal{T}_0 \cup \mathcal{K} \cup G$, which is a contradiction.

² All results below hold if only purified goals are considered; flattening and linearity of goals is not absolutely necessary.

(2) Proof similar to (1), with the difference that we start with a finite set G of ground clauses, and as support for A we take $\{t_P \mid t \in T_{\Sigma_0}(\text{st}(\mathcal{K}, G))\}$; all operations and relations are defined as above. \mathcal{T}_0 is a universal theory, so $A_{|_{\Pi_0}}$ (a Π_0 -substructure of $P_{|_{\Pi_0}}$) is also a model of \mathcal{T}_0 . As $\text{st}(\mathcal{K}, G)$ is finite and \mathcal{T}_0 is locally finite, A is finite, so (Emb_w^f) is sufficient to find a total model B of $\mathcal{T}_0 \cup \mathcal{K} \cup G$. \square

Example 2. *The following theory extensions $\mathcal{T}_0 \subseteq \mathcal{T}_1$ are local:*

- (1) *Any extension \mathcal{T}_1 of a theory \mathcal{T}_0 with a set of free function symbols.*
- (2) *Any extension \mathcal{T}_1 of a theory \mathcal{T}_0 with a set of selectors $\{s_1, \dots, s_n\}$ for an n -ary function c which is injective in \mathcal{T}_0 .*
- (3) *The extension of any of the theories \mathcal{T}_0 in Example 1(4) with monotone functions: $\mathcal{T}_1 = \mathcal{T}_0 \cup \text{Mon}_f$, where Mon_f is the monotonicity axiom of the n -ary function f .*
- (4) *The extension of the theory of reals with a λ -Lipschitz function at x_0 .*

Proof: (Sketch) Locality follows from the embeddability properties in Example 1. Ad (3): To prove condition (Loc^f) when \mathcal{T}_0 is \mathcal{DL} or \mathcal{B} it suffices to show that (Emb_w^f) holds, because these theories are universal and locally finite. \square

5.3 Embeddability of Evans Partial Models Implies Stable Locality

We now show that, for an extension $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$ of a universal theory \mathcal{T}_0 , embeddability of Evans partial models into total models implies stable locality.

Theorem 3. *Let \mathcal{T}_0 be a universal theory and \mathcal{K} be a set of clauses. Then:*

- (1) *If the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Emb) then it satisfies (SLoc) .*
- (2) *Assume that \mathcal{T}_0 is a locally finite universal theory, and that \mathcal{K} contains only finitely many ground subterms. If the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Emb^f) , then $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (SLoc^f) .*

Proof: The proof is similar to that of Theorem 2. The first difference is in the construction of the partial model A of $\mathcal{T}_0 \cup \mathcal{K} \cup G$. Let $A = \{t_P \mid t \in T_{\Sigma_0}(\text{st}(\mathcal{K}, G))\}$. Define the functions and relations in Π_0 as in P . If $f \in \Sigma_1$ is an n -ary function and $t_P^1, \dots, t_P^n \in A$, then $f_A(t_P^1, \dots, t_P^n)$ is defined and equal to t_P if and only if $t_P = f(t^1, \dots, t^n)_P \in A$. $A_{|_{\Pi_0}}$ is a Π_0 -substructure of $P_{|_{\Pi_0}}$. As \mathcal{T}_0 is a universal theory, $A_{|_{\Pi_0}}$ is a total model of \mathcal{T}_0 . As all terms in $\text{st}(\mathcal{K}, G)$ are defined both in P and in A , and $P \models G$, A satisfies all clauses in G . To show that A satisfies \mathcal{K} note that every assignment $\beta : X \rightarrow A$ defines at least one substitution $\sigma : X \rightarrow T_{\Sigma_0}(\text{st}(\mathcal{K}, G))$ such that $(\sigma(t))_P = \beta(t)$. Then $\sigma(D)$ is an instance of D in $\mathcal{K}^{[G]}$, so $P \models \sigma(D)$, hence $(A, \beta) \models D$. It follows that A satisfies $\mathcal{T}_0 \cup \mathcal{K} \cup G$. The existence of a total model of $\mathcal{T}_0 \cup \mathcal{K} \cup G$ follows from (Emb) .

(2) If G is a finite set of clauses then the additional conditions guarantee that A is finite, so only embeddability of finite partial models is necessary in the proof. \square

Example 3. (1) A shallow extension of a universal theory is stably local.

(2) Let \mathcal{T}_0 be a universal theory with signature $\Pi_0 = (\Sigma_0, \text{Pred})$, and let $c \in \Sigma_0$ be a function symbol with arity n . Let $\Pi = (\Sigma_0 \cup \Sigma_1, \text{Pred})$, where $\Sigma_1 = \{s_1, \dots, s_n\}$. Let $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}_{\text{sel}}$ be the extension of \mathcal{T}_0 with the selector axioms for the unary functions s_1, \dots, s_n . Then $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies condition (SLoc).

6 Relational Encodings, Decidability and Complexity

The locality conditions we consider relate satisfiability in total models to satisfiability of certain ground instances with respect to partial models. We can replace reasoning about partially defined functions with reasoning about relations.

For the signature $\Pi = (\Sigma_0 \cup \Sigma_1, \text{Pred})$ let Π^* denote the signature $(\Sigma_0, \Sigma_1^* \cup \text{Pred})$, where every n -ary function symbol f in Σ_1 is replaced by an $(n+1)$ -ary relation symbol r^f . If A is a Π -algebra, its relational variant is the Π^* -structure A^* for which $r_A^f(a_1, \dots, a_n, a)$ if and only if $f_A(a_1, \dots, a_n)$ is defined and equal to a .

The idea of the relational translation is to replace each atom $f(c_1, \dots, c_n) \approx c$ with the $r^f(c_1, \dots, c_n, c)$.

We use the relational translation to establish relationships between the decidability resp. complexity of the universal clause theory of the extension and the decidability resp. complexity of a suitable fragment of the base theory.

6.1 Flattening and Relational Encoding

The locality conditions defined in Section 4 require that $\mathcal{T}_1 \cup G$ is satisfiable (where G is a set of ground clauses) if and only if $\mathcal{T}_0 \cup \mathcal{K}^*[G] \cup G$ has a (Evans, weak, finite) partial model with additional properties, where, depending on the notion of locality, $\mathcal{K}^*[G]$ is $\mathcal{K}[G]$ or $\mathcal{K}^{[G]}$. In these sets of clauses the function symbols in Σ_1 only occur at the root of ground terms. Therefore, they can be flattened as explained in Section 5.1. They can also be purified (i.e. the function symbols in Σ_1 are separated from the other symbols) by introducing, in a bottom-up manner, new constants c_t for subterms $t = f(g_1, \dots, g_n)$ with $f \in \Sigma_1$, g_i ground $\Sigma_0 \cup \Sigma_c$ -terms (where Σ_c is a set of constants which contains the constants introduced by flattening, resp. purification), together with corresponding definitions $c_t \approx t$. The set of clauses thus obtained has the form $\mathcal{K}_0 \cup G_0 \cup D$, where D is a set of ground unit clauses of the form $f(g_1, \dots, g_n) \approx c$, where $f \in \Sigma_1$, c is a constant, g_1, \dots, g_n are ground terms without function symbols in Σ_1 ; and \mathcal{K}_0 and G_0 are clauses without function symbols in Σ_1 . (If we flatten and then purify $\mathcal{K}^*[G] \cup G$ we ensure that D consists of ground unit clauses of the form $f(c_1, \dots, c_n) \approx c$, where $f \in \Sigma_1$, and c_1, \dots, c_n, c are constants.) These flattening and purification transformations preserve both satisfiability and unsatisfiability with respect to total algebras, and also with respect to partial algebras in which all ground subterms which are flattened are defined.

For the sake of simplicity in what follows we will always flatten and then purify G and $\mathcal{K}^*[G]$. All results also hold if only purification is applied.

Lemma 4. *Let \mathcal{K} be a set of clauses and G a ground clause, and let $\mathcal{K}_0 \cup G_0 \cup D$ be obtained from $\mathcal{K} * [G] \cup G$ by flattening and purification, as explained above. Then the following are equivalent:*

- (1) $\mathcal{T}_0 \cup \mathcal{K} * [G] \cup G$ has a partial model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined.
- (2) $\mathcal{T}_0 \cup \mathcal{K}_0 \cup G_0 \cup D$ has a partial model with all terms in $\text{st}(\mathcal{K}_0, G_0, D)$ defined.
- (3) $\mathcal{T}_0 \cup \mathcal{K}_0 \cup G_0 \cup \text{Fun}(D^*) \cup D^*$ has a relational model, where $D^* = \{r^f(c_1, \dots, c_n, c) \mid (f(c_1, \dots, c_n) \approx c) \in D\}$ and $\text{Fun}(D^*) = \{\bigwedge_{i=1}^n c_i \approx d_i \wedge r^f(c_1, \dots, c_n, c) \wedge r^f(d_1, \dots, d_n, d) \rightarrow c \approx d \mid f \in \Sigma_1, r^f(c_1, \dots, c_n, c), r^f(d_1, \dots, d_n, d) \in D^*\}$.
- (4) $\mathcal{T}_0 \cup \mathcal{K}_0 \cup G_0 \cup N_0$ has a (total) model, where $N_0 = \{\bigwedge_{i=1}^n c_i \approx d_i \rightarrow c = d \mid r^f(c_1, \dots, c_n, c), r^f(d_1, \dots, d_n, d) \in D^*\}$.

6.2 Decidability and Complexity

Let \mathcal{T}_0 be an arbitrary Π_0 -theory, where $\Pi_0 = (\Sigma_0, \text{Pred})$ and let $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$, where \mathcal{K} is a finite set of clauses in a signature $\Pi = (\Sigma_0 \cup \Sigma_1, \text{Pred})$.

Theorem 5. *Assume that the theory extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ either (1) satisfies condition (Loc^f) , or else (2) satisfies condition (SLoc^f) and \mathcal{T}_0 is locally finite. Then:*

- (a) *If all variables in the clauses in \mathcal{K} occur below some function symbol from Σ_1 and if the universal theory of \mathcal{T}_0 is decidable, then the universal theory of \mathcal{T}_1 is decidable.*
- (b) *If the $\forall\exists$ theory of \mathcal{T}_0 is decidable then the universal theory of \mathcal{T}_1 is decidable.*

Proof: It is sufficient to show that the universal clause theory of \mathcal{T}_1 is decidable. We present the proofs under hypotheses (1) and (2) in parallel.

Let C be a clause in the signature Π with variables x_1, \dots, x_n . Obviously, $\mathcal{T}_0 \cup \mathcal{K} \models \forall x_1 \dots x_n C$ if and only if $\mathcal{T}_0 \cup \mathcal{K} \cup G$ is unsatisfiable, where G is the set of ground unit clauses obtained from $\exists x_1 \dots x_n \neg C$ by Skolemization. By the locality assumption, the last statement is equivalent to saying that $\mathcal{T}_0 \cup \mathcal{K} * [G] \cup G$ has no (weak) partial model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined (where $\mathcal{K} * [G]$ is $\mathcal{K}[G]$ in the case of local extensions and $\mathcal{K}^{[G]}$ for stably local extensions). Let $\mathcal{K}_0 \cup G_0 \cup D$ be the flattened form of $\mathcal{K} * [G] \cup G$. By Lemma 4 we know that $\mathcal{T}_0 \cup \mathcal{K} * [G] \cup G$ has no (weak) partial model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined if and only if $\mathcal{T}_0 \cup \mathcal{K}_0 \cup G_0 \cup N_0$ has a total model, where $N_0 = \{\bigwedge_{i=1}^n c_i \approx d_i \rightarrow c \approx d \mid f(c_1, \dots, c_n) \approx c, f(d_1, \dots, d_n) \approx d \in D\}$.

Flattening and purification increase the size (i.e. the total number of symbols) of clauses only by a linear factor. So the size of \mathcal{K}_0 is linear in the size of $\mathcal{K} * [G]$ and the size of $G_0 \cup D$ is linear in the size of $G \cup \mathcal{K}$. N_0 contains at most $|D|^2$ clauses, so the number of clauses in N_0 is quadratic in the number of Σ_1 ground terms occurring in \mathcal{K} and G . The maximal length of the clauses in N_0 is $m + 1$, where m is the maximal arity of a function symbol in Σ_1 . The only difference between (1) and (2) is the number of clauses in $\mathcal{K} * [G]$.

- (1) For a local extension, $\mathcal{K} * [G] = \mathcal{K}[G]$. If \mathcal{K} is finite then $\mathcal{K}[G]$ has at most $n_c \cdot n_t \cdot |\text{st}(\mathcal{K}, G)|$ clauses, where n_c is the number of clauses in \mathcal{K} and n_t the maximal number of distinct Σ_1 terms in a clause in \mathcal{K} . Then $\mathcal{K}_0 \cup G_0 \cup N_0$ is finite, of size polynomial in the size of $\mathcal{K} \cup G$.

(2) For a stably local extension, $\mathcal{K} * [G] = \mathcal{K}^{[G]}$. If \mathcal{K} is finite and \mathcal{T}_0 is locally finite then there are only finitely many equivalence classes in $T_{\Sigma_0}(\text{st}(\mathcal{K}, G))$ with respect to equality modulo \mathcal{T}_0 (say $n_{\mathcal{K}, G}$). If we only choose the representatives for instantiation in $\mathcal{K}^{[G]}$, the resulting set of clauses is finite, of size polynomial in $n_{\mathcal{K}, G}$ and in the size of $\mathcal{K} \cup G$. Then $\mathcal{K}_0 \cup G_0 \cup N_0$ is finite.

The proof now continues for both local and stably local extensions:

(a) Assume that for every clause in \mathcal{K} , every variable occurs below a Σ_1 function. Then $\mathcal{K}[G]$ (and \mathcal{K}_0) consists only of ground clauses. If checking the satisfiability of (existentially quantified) conjunctions of clauses w.r.t. \mathcal{T}_0 is decidable, then the universal clause theory of \mathcal{T}_1 is decidable, and its complexity is determined by the complexity of satisfiability checking for sets of clauses in \mathcal{T}_0 and the size of $\mathcal{K}_0 \cup G_0 \cup N_0$. The problem of checking the satisfiability of conjunctions of clauses is decidable iff the universal theory of \mathcal{T}_0 is decidable: if $\{k_1, \dots, k_m\}$ is the set of all constants that occur in $\mathcal{K}_0 \cup G_0 \cup N_0$ the following are equivalent:

- (i) $\mathcal{T}_0 \cup \bigwedge_{C \in \mathcal{K}_0 \cup G_0 \cup N_0} C(k_1, \dots, k_m) \models \perp$.
- (ii) $\mathcal{T}_0 \cup \exists x_1, \dots, x_m \left(\bigwedge_{C \in \mathcal{K}_0 \cup G_0 \cup N_0} C(x_1, \dots, x_m) \right) \models \perp$.
- (iii) $\mathcal{T}_0 \models \forall x_1, \dots, x_m \left(\bigvee_{C \in \mathcal{K}_0 \cup G_0 \cup N_0} \neg C(x_1, \dots, x_m) \right)$.

(b) If some variables in clauses in \mathcal{K} do not occur below Σ_1 -function symbols then the clauses in \mathcal{K}_0 are not necessarily ground: they contain variables $\{y_1, \dots, y_k\}$, and constants in $\{c_1, \dots, c_n\}$. The following statements are equivalent:

- (i) $\mathcal{T}_0 \cup \left(\bigwedge_{C \in \mathcal{K}_0} \forall y_1 \dots \forall y_k C(c_1, \dots, c_n, y_1, \dots, y_k) \wedge \bigwedge_{C \in G_0 \cup N_0} C(c_1, \dots, c_n) \right) \models \perp$.
- (ii) $\mathcal{T}_0 \cup \exists x_1 \dots x_n \left(\bigwedge_{C \in \mathcal{K}_0} \forall y_1 \dots y_k C(x_1 \dots x_n, y_1 \dots y_k) \wedge \bigwedge_{C \in G_0 \cup N_0} C(x_1 \dots x_n) \right) \models \perp$.
- (iii) $\mathcal{T}_0 \models \forall x_1 \dots x_n \left(\bigvee_{C \in \mathcal{K}_0} \exists y_1 \dots y_k \neg C(x_1 \dots x_n, y_1 \dots y_k) \vee \bigvee_{C \in G_0 \cup N_0} \neg C(x_1 \dots x_n) \right)$.

If the $\forall \exists$ fragment of \mathcal{T}_0 is decidable then we can use this and the equivalence of (i) and (iii) to check whether $\mathcal{T}_0 \cup \mathcal{K}_0 \cup G_0 \cup N_0$ is satisfiable. The size of $\mathcal{K}_0 \cup G_0 \cup N_0$ and the complexity of the $\forall \exists$ fragment of \mathcal{T}_0 then determine the complexity of the universal fragment of \mathcal{T}_1 . \square

Corollary 6. *Let \mathcal{T}_0 be a theory for which the satisfiability of a set of ground clauses of size n can be checked in time at most $g(n)$, and let $\mathcal{T}_0 \subseteq \mathcal{T}_0 \cup \mathcal{K}$ be a local theory extension where in every clause in \mathcal{K} each variable occurs below some extension function. The validity of a set of clauses in the extension can be checked in time $g(c \cdot n^2)$, where c is a constant. This holds for:*

- (1) Extensions with free function symbols (alternative proof of results in [5, 10]).
- (2) Extensions with monotone functions (see also Example 4)
 - If \mathcal{T}_0 is the theory \mathcal{DL} (of distributive lattices) or \mathcal{B} (of Boolean algebras) the complexity of the universal clause theory of an extension of \mathcal{T}_0 with monotone functions is in co-NP.
 - If \mathcal{T}_0 is the theory \mathcal{L} (of lattices) or \mathcal{SL} (of semilattices) then the complexity of the universal clause theory of an extension of \mathcal{T}_0 with monotone functions is in co-NP, and that of the universal Horn theory of \mathcal{T}_1 is in PTIME.

- (3) Extensions of theories of injective constructors with selectors.
- (4) Extensions of \mathbb{R} with Lipschitz functions: *the universal clause theory is in EXPTIME (an example was already presented in Section 1.1).*

Example 4. Let \mathcal{T}_0 be a theory (with a binary predicate \leq), and \mathcal{T}_1 a local extension of \mathcal{T}_0 with two monotone functions f and g . Consider the following problem:

$$\mathcal{T}_0 \cup \text{Mon}_f \cup \text{Mon}_g \models \forall x, y, z, u, v (x \leq y \wedge f(y \vee z) \leq g(u \wedge v) \rightarrow f(x) \leq g(v))$$

The problem reduces to the problem of checking whether $\mathcal{T}_0 \cup \text{Mon}_f \cup \text{Mon}_g \cup G \models \perp$, where $G = c_0 \leq c_1 \wedge f(c_1 \vee c_2) \leq g(c_3 \wedge c_4) \wedge f(c_0) \not\leq g(c_4)$.

After flattening, using the locality of the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$, making the relational translation, and computing N_0 , we obtain the following set of clauses:

$c_0 \leq c_1$	$r_f(d_1, e_1)$	$d_1 = c_0 \rightarrow e_1 = e_3$	$d_1 \leq c_0 \rightarrow e_1 \leq e_3$
$d_1 = c_1 \vee c_2$	$r_f(c_0, e_3)$	$d_2 = c_4 \rightarrow e_2 = e_4$	$c_0 \leq d_1 \rightarrow e_3 \leq e_1$
$d_2 = c_3 \wedge c_4$	$r_g(c_4, e_4)$		$d_2 \leq c_4 \rightarrow e_2 \leq e_4$
$e_1 \leq e_2$	$r_g(d_2, e_2)$		$d_4 \leq d_2 \rightarrow e_4 \leq e_2$
$e_3 \not\leq e_4$			

(1) Assume \mathcal{T}_0 is \mathcal{DL} or \mathcal{B} . The universal clause theory of \mathcal{DL} (resp. \mathcal{B}) is the theory of the two element lattice (resp. two element Boolean algebra), so testing Boolean satisfiability is sufficient. (This is in NP.) We proved unsatisfiability using SPASS [11], but SAT solvers such as, e.g. CHAFF [9], can be used as well.

(2) If $\mathcal{T}_0 = \mathcal{L}$ we can reduce the problem above to the problem of checking the satisfiability of a set of ground Horn clauses (via the relational translation of Skolem described in the introduction). This can be checked in PTIME.

(3) If $\mathcal{T}_0 = \mathbb{R}$ we first need to explain what \vee and \wedge are. For this, we replace $d_1 = c_1 \vee c_2$ with $(c_1 \leq c_2 \rightarrow d_1 = c_2) \wedge (c_2 < c_1 \rightarrow d_1 = c_2)$ and similarly for $d_2 = c_3 \wedge c_4$. We proved unsatisfiability using the REDLOG demo [3].

We can therefore conclude that in all cases above:

$$\mathcal{T}_1 \models \forall x, y, z, u, v (x \leq y \wedge f(y \vee z) \leq g(u \wedge v) \rightarrow f(x) \leq g(v)).$$

7 Beyond the Universal Fragment

Analyzing the proof of Theorems 2 and 3 we notice that the embeddability conditions (Comp) and (Comp_w) imply, in fact, stronger locality conditions. Consider a theory extension $\mathcal{T}_0 \subseteq \mathcal{T}_0 \cup \mathcal{K}$ with a set \mathcal{K} of formulae of the form $\forall x_1 \dots x_n (\Phi(x_1, \dots, x_n) \vee C(x_1, \dots, x_n))$, where $\Phi(x_1, \dots, x_n)$ is an arbitrary first-order formula in the base signature Π_0 with free variables x_1, \dots, x_n , and $C(x_1, \dots, x_n)$ is a clause in the signature Π .

We can extend the notion of locality of an extension accordingly:

- (ELoc) For every formula $\Gamma = \Gamma_0 \cup G$, where Γ_0 is a Π_0 -sentence and G is a set of ground clauses, $\mathcal{T}_1 \cup \Gamma \models \perp$ iff $\mathcal{T}_0 \cup \mathcal{K}[\Gamma] \cup \Gamma$ has no weak partial model in which all terms in $\text{st}(\mathcal{K}, G)$ are defined.

A stable locality condition (ESLoc) can be defined similarly. The proofs of Theorems 2 and 3 can be adapted with minimal changes to prove a stronger result:

- Theorem 7.** (1) Assume all terms of \mathcal{K} starting with a Σ_1 function are flat and linear. If the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Comp_w) then it satisfies (ELoc) .
 (2) Assume that \mathcal{T}_0 is a universal theory. If the extension $\mathcal{T}_0 \subseteq \mathcal{T}_1$ satisfies (Comp) then it satisfies (ESLoc) .

Proof: (Idea) By (Comp) , the partial model and its total completion have supports whose Π_0 -reducts are isomorphic, hence elementarily equivalent. Therefore the (weak) embedding guaranteed by (Comp_w) resp. (Comp) preserves and reflects the truth of all first-order formulae in the base signature. \square

Further generalizations are possible (concerning both the form of the set of extension formulae, and the form of the goals). This is work in progress.

8 Conclusions

We introduced notions of locality for theory extensions and showed that for local theory extensions we may regard w.l.o.g. the extension functions as functional relations. Using a relational translation we identified situations where it is possible to express the decidability (complexity) of an extension \mathcal{T}_1 in terms of the decidability (complexity) of a fragment of the base theory \mathcal{T}_0 (universal or $\forall\exists$). These results apply to theories of data types and to some theories of functions from algebra or mathematical analysis.

There seem to exist relationships with results on combinations of non stably infinite theories [10]. The result on extensions of an arbitrary theory with free functions which we obtain as an example was discovered independently in a different context by Ganzinger [5] and Tinelli and Zarba [10]. However, here we go beyond analyzing mere combinations of theories: we look at proper extensions, in which the extension axioms contain functions from the base theory. In this paper we restrict ourselves to one-sorted theories. Similar results can be obtained in a many-sorted framework.

Acknowledgements. Many of the results presented here are the direct or indirect result of discussions and joint work with Harald Ganzinger during the last years. Some of the ideas on local extensions sketched in [6] are now presented in full detail and generalized. I thank Uwe Waldmann for discussions on partial algebras, and Carsten Ihlemann for commenting on an early draft of this paper. Many thanks to the referees for their helpful comments.

References

1. P. Burmeister. *A Model Theoretic Oriented Approach to Partial Algebras: Introduction to Theory and Application of Partial Algebras, Part I*, volume 31 of *Mathematical Research*. Akademie-Verlag, Berlin, 1986.
2. S. Burris. Polynomial time uniform word problems. *Mathematical Logic Quarterly*, 41:173–182, 1995.
3. A. Dolzmann and T. Sturm. Redlog: Computer algebra meets computer logic. *ACM SIGSAM Bulletin*, 31(2):2–9, 1997.

4. H. Ganzinger. Relating semantic and proof-theoretic concepts for polynomial time decidability of uniform word problems. In *Proc. 16th IEEE Symposium on Logic in Computer Science (LICS'01)*, pages 81–92. IEEE Computer Society Press, 2001.
5. H. Ganzinger. Shostak light. In A. Voronkov, editor, *Automated Deduction – CADE-18*, LNCS 2392, pages 332–346, Copenhagen, Denmark, 2002. Springer.
6. H. Ganzinger, V. Sofronie-Stokkermans, and U. Waldmann. Modular proof systems for partial functions with weak equality. In *Proc. International Joint Conference on Automated Reasoning (IJCAR'04)*, LNCS 3097, pages 168–182. Springer, 2004.
7. R. Givan and D. McAllester. New results on local inference relations. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR'92)*, pages 403–412. Morgan Kaufmann Press, 1992.
8. D. McAllester. Automatic recognition of tractability in inference relations. *Journal of the Association for Computing Machinery*, 40(2):284–303, 1993.
9. M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *Proceedings of the 39th Design Automation Conference (DAC 2001)*, pages 530–535. ACM, 2001.
10. C. Tinelli and C.G. Zarba. Combining non-stably infinite theories. *Journal of Automated Reasoning*, 2005. to appear.
11. Ch. Weidenbach, B. Gaede, and G. Rock. SPASS & FLOTTER Version 0.42. In M.A. McRobie and J.K. Slaney, editors, *Proceedings of CADE-13*, LNCS 1104, pages 141–145. Springer Verlag, 1996.