

Local Theories and Efficient Partial Quantifier Elimination

by

Estifanos Sisay Getachew

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2025

© Estifanos Sisay Getachew 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Quantifier elimination is used in various automated reasoning tasks, including quantified SMT solving, exists/forall solving, program synthesis, model checking, and constrained Horn clause (CHC) solving. Complete quantifier elimination, however, is computationally intractable for many theories. The recent algorithm QEL shows a promising approach to approximate quantifier elimination, which has resulted in improvements in solver performance. QEL performs partial quantifier elimination with a completeness guarantee that depends on a certain semantic property of the given formula.

In this thesis, we study *local theories*, focusing on their proof theoretic and semantic characterization. We identify a subclass of local theories in which *partial quantifier elimination* can be performed efficiently. By considerably generalizing the previous approach, we present \mathcal{T} -QEL, a parametrized polynomial time algorithm that is relatively complete for this class of theories. The algorithm utilizes the proof theoretic characterization of the theories, which is based on *restricted derivations*. Finally, we prove for \mathcal{T} -QEL, soundness in general, and relative completeness with respect to the identified class of theories.

Acknowledgements

Thank you to everyone who supported me during my studies.

Table of Contents

Author's Declaration	ii
Abstract	iii
Acknowledgements	iv
List of Figures	vii
List of Algorithms	viii
1 Introduction	1
1.1 Summary of Contributions	3
1.2 Outline of Thesis	3
2 Background	5
2.1 Horn Theories and Direct Products	5
2.1.1 Preliminaries	5
2.1.2 McKinsey's lemma	7
2.2 Complete lattices and Fixed points	8
2.3 Quantifier Elimination	9
2.3.1 Presburger Arithmetic	10
2.3.2 Linear Real Arithmetic	14
2.4 Egraphs	15

3	Overview	16
3.1	Motivating Example	16
3.2	High level overview of \mathcal{T} -QEL	18
4	Locality: A Proof Theoretic Characterization	20
4.1	Derivations	21
4.2	Undecidability of Locality and Characterizing P	25
4.3	A Brief Interlude: Hilbert’s 24th problem	29
5	Locality: A Semantic Characterization	30
5.1	Basic Notions	31
5.2	Locality and Embeddability	33
5.2.1	Locality Implies Embeddability	33
5.2.2	Embeddability (with weak satisfaction) Implies Locality	34
5.3	Axiomatizability of Relational Substructures	36
5.4	Conclusion	37
6	Locally Ground Theories	39
6.1	Partial Orders	42
6.2	Recursively Defined Data Structures	44
7	Partial Quantifier Elimination	48
7.1	\mathcal{T} -QEL	49
7.2	Soundness and Relative Completeness	52
8	Conclusion and Related Work	56
8.1	Related Work	57
	References	58

List of Figures

3.1	An example egraph of a formula and a final transformed egraph representing the partition S_φ/\sim	18
4.1	Axioms to inference rules, where σ is a substitution.	21

List of Algorithms

1	Quantifier Elimination Procedure for \mathfrak{N}_{\equiv}	12
2	Fourier-Motzkin.	14
3	\mathcal{T} -QEL : An extension of QEL that utilizes theories.	50

Chapter 1

Introduction

Several automated reasoning tasks, including quantified SMT solving [5], exists/forall solving [13, 14], program synthesis [35], model checking [32, 21], and CHC solving [43] utilize quantifier elimination. Regardless, the intractability of complete quantifier elimination procedures, see e.g. [34, 24, 11], poses a challenge. As a result, solvers often use some form of an approximate technique. An interesting example of this, is the algorithm QEL proposed in [20], that resulted in significant performance improvements in Z3 [12] and the CHC solver Spacer [32].

The problem QEL solves can be formulated as a *partial quantifier elimination* problem. In particular, for any formula of the form $\exists \bar{x} \varphi(\bar{x})$, where $\varphi(\bar{x})$ is a conjunction of literals, QEL computes a quantifier free $\psi(\bar{y})$ whose free variables are among the x_i and its existential closure, $\exists \bar{y} \psi(\bar{y})$, is equivalent to $\exists \bar{x} \varphi(\bar{x})$. This can be seen as a relaxation of (complete) quantifier elimination, in which we allow some of the quantified variables to remain. An important property of QEL is that it is *relatively complete* for the theory of equality (with free functions), \mathcal{T}_{eq} ; that is, QEL guarantees to eliminate a variable x_i if it has an entailed ground definition — i.e., $\varphi(\bar{x}) \models x_i \approx s$ for some *arbitrary* ground term s .

Despite this, QEL only deals with formulas in the language of \mathcal{T}_{eq} , and consequently, its completeness guarantees are restricted to just \mathcal{T}_{eq} . We take a theoretical direction, and investigate for which other theories we can *efficiently* generalize QEL while maintaining *relative completeness* with respect to the theories. In particular, we explore theories in which we can do so in polynomial time. We use the notion of *locality* [22, 19] of a theory to identify theories in which we can perform partial quantifier elimination efficiently. Using the proof theoretic characterization of the theories, we present \mathcal{T} -QEL, a generalization of QEL.

The notion of *locality* in theories was first introduced by Givan and McAllester in [22, 23, 37]. They considered locality of theories in the context of First Order Logic (FOL) without equality. They gave a proof theoretic account of locality based on the concept of confining deductions to the subterms of an input formula. Ganzinger in [19] considers locality in the context of FOL with equality and gives a semantic characterization of local (equational) theories. We adopt some results due to Ganzinger [19], to lift the proof theoretic characterization of locality by Givan and McAllester to FOL with equality. We then base our algorithm \mathcal{T} -QEL on this *proof theoretic characterization*.

The reason we are interested in locality is that the uniform word problem for local theories is *polynomial time* decidable [19]. That is, given a ground Horn clause C and a local theory (locally) axiomatized by a finite set of Horn clauses Φ , we can decide entailment from the theory, $\Phi \models C$, in polynomial time. This is central to generalizing QEL to other theories in polynomial time.

Another central issue that arises, when trying to generalize QEL, is the problem of deciding whether in the theory \mathcal{T} , there is an entailed ground definition from a conjunction of literals φ . That is, given a free variable x of φ deciding whether for some *arbitrary* ground term s , $\varphi \models_{\mathcal{T}} x \approx s$. The problem is complicated by the fact that s is potentially an arbitrary term, not necessarily occurring in φ . Since we are working in a countable language, however, if \mathcal{T} has a decision procedure for its quantifier free fragment, there is the following (rather naive) semi-decision procedure: enumerate the ground terms s_0, s_1, \dots and for each s_i , check whether $\varphi \wedge x \not\approx s_i$ is \mathcal{T} -unsatisfiable, if so replace, in φ , x by s_i . Clearly, this is sound, but potentially non-terminating.

A way to deal with this potentially non-terminating behavior, is to reduce the search space for ground definitions to a finite one. Specifically, we introduce the finite notion of being a *constructively ground term* with respect to a theory \mathcal{T} and a conjunction of literals φ . We then give a definition for a subclass of local theories¹, which we call *locally ground theories*, based on this concept. It is worth noting that the constructively ground terms can be computed in polynomial time.

Interestingly, several important theories shown in the literature to exhibit *locality* properties are also *locally ground*. In particular, we show the theories listed below to be locally ground; for each of these theories, we include in our signature a countable set of free/uninterpreted function and constant symbols: **(i)** The theory of recursively defined data structures \mathcal{T}_{rd} ; **(ii)** the theory of partial orders \mathcal{T}_{po} ; **(iii)** the theory \mathcal{T}_{mo} , where some functions are axiomatized as monotone with respect to the partial order \leq ; and **(iv)** the theory of equality \mathcal{T}_{eq} .

¹Strictly speaking, we are considering an extended notion of *locality* from [19].

Finally, building on the ideas discussed above, we present the partial quantifier elimination algorithm \mathcal{T} -QEL. \mathcal{T} -QEL is a sound extension of QEL that utilizes the concepts of restricted derivations and constructively ground terms to solve the partial quantifier elimination problem in polynomial time for the locally ground theories. We prove, for \mathcal{T} -QEL, soundness in general and relative completeness for the locally ground theories.

Applications beyond locally ground theories. The algorithm \mathcal{T} -QEL can be used as a partial quantifier elimination algorithm for theories that might not necessarily be locally ground. Even for theories such as Linear Integer Arithmetic and Linear Real Arithmetic — where (complete) quantifier elimination is available — we can abstract away the semantics of the functions $+$ and \times as simply monotonic and free (uninterpreted) functions respectively and keep the semantics of \leq as a partial order. This abstraction allows us to use \mathcal{T} -QEL as a potentially efficient preprocessing step to partially eliminate quantified variables before performing the more expensive quantifier elimination algorithms of the theories.

1.1 Summary of Contributions

The contributions of this thesis can be summarized as follows.

- We identify a subclass of local theories, called *locally ground theories*, in which partial quantifier elimination can be done in polynomial time. This is done by reducing the search space for entailed ground definition, in a theory, to a finite one.
- We show several important theories to be locally ground.
- We give a *polynomial time algorithm* by lifting the proof theoretic characterization due to Givan and McAllester to FOL with equality.
- We prove for the algorithm \mathcal{T} -QEL *soundness* in general, and *relative completeness* for the locally ground theories.

1.2 Outline of Thesis

The rest of the thesis is organized as follows. In Chapter 2, we go over the basic notions and background for the thesis. In particular, we go over basic notions in first order logic (FOL),

Horn clauses and products of their models, as well as quantifier elimination procedures for Presburger and linear real arithmetic. Chapter 3 provides an overview of the intuition behind the *locally ground theories*. Moreover, the algorithm \mathcal{T} -QEL is demonstrated via an intuitive example. Chapter 4 looks at *local theories* using their proof theoretic characterization in the context of first order logic without equality. Finally, we show that the local theories give a characterization of the complexity class P.

Chapter 5 provides a semantic characterization of the *local theories* and discusses the uniform word problem. Polynomial time decidability of the uniform word problem for local theories is demonstrated. Chapter 6 introduces the *locally ground theories* using the idea of being a constructively ground term with respect to a given conjunction of literal and a theory. In Chapter 7, the proof theoretic characterization is used to provide the partial quantifier elimination algorithm \mathcal{T} -QEL. Proofs of soundness and relative completeness for the algorithm \mathcal{T} -QEL are presented.

Chapter 2

Background

We give in this chapter, a brief summary of the notational conventions and the necessary background for the rest of the thesis. We first review theories axiomatized by a set of Horn clauses, and the direct products of their models, with the goal of proving McKinsey’s lemma. Following that, we discuss the basics of quantifier elimination and give quantifier elimination procedures for two specific theories — Presburger arithmetic and linear real arithmetic. We discuss complexity issues regarding the procedures of these theories. The final section gives background on Egraphs.

2.1 Horn Theories and Direct Products

2.1.1 Preliminaries

Unless explicitly mentioned, we assume we are working in First order logic (FOL) with *equality*, and assume the usual concepts such as model, satisfaction, logical consequence, and theory for FOL with equality, as presented in e.g. [15, 18]. We will sometimes make use of FOL without *equality*, i.e., where equality is treated as just another predicate. In this case, we will explicitly mention it, and assume the corresponding concepts for FOL without equality. In particular, we will use \models_{neq} to denote logical consequence in FOL without equality. We briefly summarize the rest of notation used in the thesis.

For a given signature Σ , we take a Σ -structure \mathfrak{A} to be a tuple (A, I) where A is the universe, and I is a mapping that assigns each k -ary function and predicate symbol, f and P , a k -ary operation and relation, $f^{\mathfrak{A}}$ and $P^{\mathfrak{A}}$, on A . We use \mathcal{V} to denote the set of all

variables, and take a Σ -interpretation, \mathcal{I} , to be a tuple (\mathfrak{A}, α) of a Σ -structure \mathfrak{A} and a variable assignment $\alpha : \mathcal{V} \rightarrow A$. We denote by $t^{\mathcal{I}}$ the interpretation of the term t under \mathcal{I} . Throughout the thesis we assume bounded function and predicate arity. We use the notation $\text{FV}(\varphi)$ to denote the set of free variables of the formula φ . Given a conjunction of literals φ , we use Γ_φ for the set containing exactly the positive literals of φ .

For a given term t , we let S_t be the set consisting of the subterms of t , i.e., $S_t = \{x\}$ if t is some variable x , and $S_t = \{t\} \cup \bigcup S_{t_i}$ if t is $f(t_1, \dots, t_n)$. We extend the definition to S_φ and S_Φ for a Σ -formula φ and a set of Σ -formulas Φ , in the expected way. We denote by Σ^T the set of all terms generated from the signature Σ . Given a function $\lambda : \mathcal{P}(\Sigma^T) \rightarrow \mathcal{P}(\Sigma^T)$ we adopt the notation S_φ^λ to denote the set $\lambda(S_\varphi)$ assigned to S_φ .

Say that a subterm set Ψ is closed under the subterm relation, if, for all $t \in \Psi$ all the subterms of t are in Ψ . We say that a binary relation R on a subterm set Ψ is a *congruence relation* if it is an equivalence relation and whenever $t_i R s_i$ and $f(t_1, \dots, t_n), f(s_1, \dots, s_n) \in \Psi$ we have $f(t_1, \dots, t_n) R f(s_1, \dots, s_n)$.

We take a *theory* \mathcal{T} to be a set of Σ -sentences that is closed under logical consequence. Note that, for a given set of sentences (axioms) Φ , the consequences of Φ , which we denote by $\text{Cn } \Phi = \{\varphi \mid \Phi \models \varphi\}$, is a theory. We denote by $\text{Mod } \Phi$ the models of Φ . Say that φ is \mathcal{T} -satisfiable if there is a model of \mathcal{T} that is also a model of φ . We use the notation $\models_{\mathcal{T}}$ to denote logical consequence in the theory \mathcal{T} .

A *homomorphism* $h : A \rightarrow B$ from the structure \mathfrak{A} to \mathfrak{B} is a map that preserves the operations and relations. That is, if f and P are k -ary function and predicate symbols respectively, we have for every $a_1, \dots, a_k \in A$, $h(f^{\mathfrak{A}}(a_1, \dots, a_k)) = f^{\mathfrak{B}}(h(a_1), \dots, h(a_k))$, and $P^{\mathfrak{A}}(a_1, \dots, a_k)$ iff $P^{\mathfrak{B}}(h(a_1), \dots, h(a_k))$. We call h a *weak embedding* if it is injective, and an *isomorphism* if it is a bijection. We say \mathfrak{A} weakly embeds into \mathfrak{B} if there exists a homomorphism h from \mathfrak{A} to \mathfrak{B} , and \mathfrak{A} is isomorphic to \mathfrak{B} if h is an isomorphism. We say that \mathfrak{A} is a *substructure* of \mathfrak{B} , if $A \subseteq B$ and the identity function is a weak embedding.

We use the term basic Horn clause for formulas of the form $\bigwedge \psi_i \rightarrow \psi$, where ψ_i, ψ are atoms, to distinguish them from (universal) Horn clauses $\forall \bar{x} (\bigwedge \psi_i \rightarrow \psi)$. We will use *Horn formulas* when we want to refer to either of these without distinguishing them. We say a Horn formula is *superficial* if every term that appears in the head of the clause also appears in one of the antecedents — i.e., $S_\psi \subseteq S_{(\bigwedge \psi_i)}$ in the above formulas. We will call a theory \mathcal{T} a *Horn theory* if it is axiomatized by some set of Horn clauses, that is, $\mathcal{T} = \text{Cn } \Phi$ for some set of Horn clauses Φ . For convenience, we sometimes omit \mathcal{T} and simply refer to Φ as the theory, especially in Chapter 5.

2.1.2 McKinsey's lemma

Our goal in this subsection is to prove McKinsey's lemma which will be useful to us in most of the chapters in the thesis. The proof relies on the fact that Horn theories are closed under *direct products*. The concept of a direct product is a generalization of the usual notion of a product of algebraic structures, such as groups. We begin, then, by defining what a direct product of a family of structures is.

Let I be a non empty index set and $\{\mathfrak{A}_i\}_{i \in I}$ a family of Σ -structures, then the *direct product* $\prod_{i \in I} \mathfrak{A}_i$ is the Σ -structure \mathfrak{B} defined as follows. The universe of \mathfrak{B} is the cartesian product of each of the A_i , i.e., $B = \{a \mid a : I \rightarrow \bigcup_{i \in I} A_i \text{ s.t. } \forall i \in I, a(i) \in A_i\}$. For each k -ary function symbol $f \in \Sigma^F$, and $a_1, \dots, a_k \in B$, we let $f^{\mathfrak{B}}(a_1, \dots, a_k) = b$, such that for every $i \in I$, $b(i)$ is $f^{\mathfrak{A}_i}(a_1(i), \dots, a_k(i))$. Similarly for each k -ary predicate symbol P in Σ and $a_1, \dots, a_k \in B$, we let $(a_1, \dots, a_k) \in P^{\mathfrak{B}}$ iff for every $i \in I$, $(a_1(i), \dots, a_k(i)) \in P^{\mathfrak{A}_i}$.

Direct products suffice for our purposes, as we are only concerned with Horn theories, and the models of Horn theories are closed under products. However, this is not always the case, some classes of structures might not be closed under direct products. In such cases one might form a different type of product such as the ultraproduct or the reduced product, in which the class of structures are closed under. An example of this, is that every first order axiomatizable class of structures is closed under ultraproducts [26, 8].

We will need the following lemma, before proving McKinsey's lemma.

Lemma 2.1.1. *Let $\varphi(x_1, \dots, x_k)$ be a Horn formula, $\{\mathfrak{A}_i\}_{i \in I}$ a family of structures, and a_1, \dots, a_k be elements in $\prod_{i \in I} A_i$. Then,*

$$\prod_{i \in I} \mathfrak{A}_i \models \varphi(a_1, \dots, a_k) \text{ iff for each } i \in I, \mathfrak{A}_i \models \varphi(a_1(i), \dots, a_k(i)).$$

Proof. We proceed by induction on $\varphi(\bar{x})$. We may assume that each atomic formula is of flat, i.e., of the form $P(x_1, \dots, x_k)$, as otherwise we can transform the given Horn clause into the desired form (as for e.g $P(t_1, \dots, t_k)$ is equivalent with $\forall \bar{x}(x_1 \approx t_1 \wedge \dots \wedge x_k \approx t_k \rightarrow P(x_1, \dots, x_k))$ where the x_i are fresh variables). Then the conclusion for the base case, i.e., when φ is $P(t_1, \dots, t_k)$, follows from the definition of products. Assume φ is a basic Horn clause of the form $\bigwedge \psi_j \rightarrow \psi$ then we have

$$\begin{aligned} \prod_{i \in I} \mathfrak{A}_i \not\models \varphi(a_1, \dots, a_k) &\iff \prod_{i \in I} \mathfrak{A}_i \models \bigwedge \psi_j(a_1, \dots, a_k) \text{ and } \prod_{i \in I} \mathfrak{A}_i \not\models \psi(a_1, \dots, a_k) \\ &\stackrel{\text{IH}}{\iff} \forall i \in I, \mathfrak{A}_i \models \psi_j(a_1(i), \dots, a_k(i)) \text{ and } \mathfrak{A}_i \not\models \psi(a_1(i), \dots, a_k(i)) \\ &\iff \forall i \in I, \mathfrak{A}_i \not\models \varphi(a_1(i), \dots, a_k(i)) \end{aligned}$$

The conclusion for the universal case follows similarly from the inductive hypothesis. \square

Hence, it follows that the models of Horn formulas are preserved under direct products. We now prove the following variant of McKinsey's lemma for theories axiomatized by Horn clauses. Other similar formulations of McKinsey's lemma can be found in [26, 25].

Lemma 2.1.2. (*McKinsey's lemma*). *Let $\Phi(\bar{x})$ be a set consisting of universal and basic Horn clauses, and $\Psi(\bar{x})$ be a set of atoms. If $\Phi(\bar{x}) \models \bigvee \Psi(\bar{x})$ then $\Phi(\bar{x}) \models \psi$ for some $\psi \in \Psi(\bar{x})$.*

Proof. Assume to the contrary that for all $\psi \in \Psi(\bar{x})$, $\Phi(\bar{x}) \not\models \psi$. Then, for each ψ , let \mathfrak{A}_ψ be a structure s.t. $\mathfrak{A}_\psi \models (\bigwedge \Phi \wedge \neg\psi)(\bar{a}_\psi)$ for some $\bar{a}_\psi \in A_\psi^n$. Now let \mathfrak{B} be the direct product of the structures, i.e., $\mathfrak{B} = \prod_{\psi \in \Psi} \mathfrak{A}_\psi$, and $\bar{b} = (b_1, \dots, b_n)$ be a tuple in B such that $b_i(\psi) = \bar{a}_\psi(i)$. Then, as it follows from Lemma 2.1.1 that models of Horn formulas are preserved under direct products, we have $\mathfrak{B} \models \bigwedge \Phi(\bar{b})$, thus by our assumption that $\Phi(\bar{x}) \models \bigvee \Psi(\bar{x})$ we have $\mathfrak{B} \models \bigvee \Psi(\bar{b})$, and hence $\mathfrak{B} \models \psi(\bar{b})$ for some $\psi \in \Psi$. Again by Lemma 2.1.1, $\mathfrak{A}_\psi \models \psi(\bar{a}_\psi)$. Contradiction. \square

2.2 Complete lattices and Fixed points

A *partially ordered set (poset)* is a set S equipped with a partial order \leq on S . Let S be a poset, $X \subseteq S$, and a an element of S , we say that a is an *upper bound* of X if for all $b \in X$, we have $b \leq a$. We say the upper bound a of X is the *least upper bound* if for all upper bounds a' of X , we have $a \leq a'$. Similarly, $b \in S$ is a *lower bound* of X , if for all $c \in X$, $b \leq c$. A lower bound b of X is the *greatest lower bound* if for all lower bounds b' of X , $b' \leq b$.

Note that both the least upper bound and the greatest lower bound of X , when they exist, are unique. We will use $\text{lub}(X)$ and $\text{glb}(X)$ respectively to denote these elements.

A partially ordered set L is a *complete lattice* if for all subsets $X \subseteq L$, both the greatest lower bound of X , $\text{glb}(X)$, and the least upper bound of X , $\text{lub}(X)$, exist. An operation $h : L \rightarrow L$ on a complete lattice is *monotonic* if, for all $a_1, a_2 \in S$ we have $h(a_1) \leq h(a_2)$ whenever $a_1 \leq a_2$. An element a of L is a *least fixed point* of h , $\text{lfp}(h)$, if it is a fixed point of h , i.e., $h(a) = a$, and for all other fixed points a' , $a \leq a'$. Similarly, a is the *greatest fixed point* of h , $\text{gfp}(h)$, if for all other fixed points a' , $a' \leq a$.

Our interest in this section is the following weaker version of the Knaster-Tarski theorem, due to Tarski [51].

Theorem 2.2.1. (Knaster–Tarski [51]). Let L be a complete lattice and $h : L \rightarrow L$ a monotonic operation on L . Then, both the least and greatest fixed points of h exist, and are given by, $\text{lfp}(h) = \text{glb}(\{b \mid h(b) = b\})$, and similarly the greatest fixed point of h , $\text{gfp}(h) = \text{lub}(\{a \mid h(a) = a\})$.

Proof. Let $FP = \{b \mid h(b) = b\}$ and $FP' = \{b' \mid h(b') \leq b'\}$, then note that $\text{glb}(FP) = \text{glb}(FP')$. Let $g = \text{glb}(FP')$, then we have for all $b' \in FP'$, $g \leq b'$, by monotonicity of h , $h(g) \leq h(b') \leq b'$, hence $h(g)$ is a lower bound, it follows then $h(g) \leq g$ and $g \in FP'$. Similarly, by monotonicity $h(h(g)) \leq h(g)$, and hence $h(g) \in FP'$, and hence $g = h(g)$ and g is a fixed point of h . Finally, as $\text{glb}(FP) = \text{glb}(FP')$, g is the least fixed point of h . We can proceed in a similar way for the greatest fixed point. \square

2.3 Quantifier Elimination

In this section, we give an overview of the basics of quantifier elimination. Following that, we consider two concrete theories and discuss complexity issues related to their quantifier elimination procedures. A theory \mathcal{T} in some first order language \mathcal{L} is said to *admit quantifier elimination* if for every \mathcal{L} -formula $\varphi(\bar{x})$ there exists a computable, \mathcal{T} -equivalent quantifier free formula $\psi(\bar{x})$. Some theories that admit quantifier elimination include the theory \mathcal{T}_{dlo} of dense linear orders without endpoints, the theory of the structure $(\mathbb{N}, +, 0, 1, <, \{\equiv_k\}_{k \geq 2})$, i.e., Presburger arithmetic extended with congruence predicates¹, and the theory $Th(\mathbb{R}, 0, 1, +, \leq)$, i.e., Linear Real Arithmetic.

Note that when a theory \mathcal{T} admits quantifier elimination, deciding validity (and hence satisfiability) of all \mathcal{L} -formulas modulo \mathcal{T} reduces to that of deciding the validity (resp. satisfiability) of the quantifier free \mathcal{L} -formulas in \mathcal{T} . Hence, as, in all of the above cases the quantifier free fragment of the theories is known to be decidable, the whole theory is decidable.

Now, contrast this with the theory \mathcal{T}_{eq} of equality with free functions². A simple way to show that \mathcal{T}_{eq} does not admit quantifier elimination is as follows. We know that \mathcal{T}_{eq} , is in general undecidable, however, the quantifier free fragment of \mathcal{T}_{eq} has a decision procedure, e.g. the congruence closure algorithm. Hence, if \mathcal{T}_{eq} were to have a quantifier elimination procedure, one could reduce an arbitrary formula in \mathcal{T}_{eq} to a \mathcal{T}_{eq} -equivalent quantifier free formula and then effectively decide the quantifier free formula. Giving us a decision procedure for all \mathcal{T}_{eq} formulas.

¹One can equivalently extend Presburger arithmetic with divisibility predicates.

²This is just the theory axiomatized by \emptyset in FOL with equality.

Another way to show that \mathcal{T}_{eq} does not admit quantifier elimination is by noting that the models of \mathcal{T}_{eq} are not model complete. A theory is *model complete*, if for every models \mathfrak{A} and \mathfrak{B} of \mathcal{T} , with \mathfrak{A} a substructure of \mathfrak{B} , we have that \mathfrak{A} is an *elementary* substructure of \mathfrak{B} , i.e., \mathfrak{A} and \mathfrak{B} satisfy the same first order sentences. We know theories that admit quantifier elimination are *model complete*. However, in \mathcal{T}_{eq} , one can easily construct substructures that fail to be elementary substructures — just consider the sentence $\exists x f(x) \approx c$.

In what follows, we only consider formulas of the form $\exists x \varphi(x, \bar{y})$, where $\varphi(x, \bar{y})$ is a conjunction of literals. This is sufficient because, for any given theory, showing that one can eliminate a single existentially quantified variable from formulas of the above form guarantees that the theory admits full quantifier elimination.

Lemma 2.3.1. *Let \mathcal{T} be a theory and assume for every formula of the form $\exists x \varphi(x, \bar{y})$ where $\varphi(x, \bar{y})$ is a conjunction of literals, there exists a quantifier free formula $\psi(\bar{y})$ such that $\models_{\mathcal{T}} \varphi(x, \bar{y}) \leftrightarrow \psi(\bar{y})$. Then, \mathcal{T} admits quantifier elimination.*

Proof. We can proceed by induction on \mathcal{L} -formulas φ_0 . We may assume that φ_0 is in prenex normal form, $Q_1 x_1 \dots Q_k x_k \varphi'_0$ where $Q_k \in \{\forall, \exists\}$, for simplicity we assume $k \leq 1$, the general case can be shown by induction on k . If φ_0 is already a quantifier free formula then we can simply let ψ be φ_0 . In the case φ_0 is of the form $\exists x \varphi'_0$, note that we can transform φ'_0 into an equivalent DNF formula, $((\alpha_0^1 \wedge \dots \wedge \alpha_{k_1}^1) \vee \dots \vee (\alpha_0^\ell \wedge \dots \wedge \alpha_{k_\ell}^\ell))$. Hence, φ_0 is equivalent with $\exists x ((\alpha_0^1 \wedge \dots \wedge \alpha_{k_1}^1) \vee \dots \vee (\alpha_0^\ell \wedge \dots \wedge \alpha_{k_\ell}^\ell))$. Now, we can distribute the existential quantifier over disjunction, hence, φ_0 is equivalent with $(\exists x (\alpha_0^1 \wedge \dots \wedge \alpha_{k_1}^1)) \vee \dots \vee \exists x (\alpha_0^\ell \wedge \dots \wedge \alpha_{k_\ell}^\ell)$, by our assumption each $\exists x (\alpha_0^i \wedge \dots \wedge \alpha_{k_i}^i)$ is \mathcal{T} -equivalent with a quantifier free formula ψ_i , it follows then φ_0 is \mathcal{T} -equivalent with $\psi_1 \vee \dots \vee \psi_\ell$. The universal similarly follows by noting that $\forall x \varphi'_0$ is equivalent with $\neg \exists x \neg \varphi'_0$, and taking CNF instead of DNF. \square

We now consider two specific theories that admit quantifier elimination: Presburger arithmetic (extended with congruence predicates) and Linear Real Arithmetic. For each theory we give a quantifier elimination procedure and discuss complexity issues.

2.3.1 Presburger Arithmetic

Presburger arithmetic (PA) refers to the theory of the structure $\mathfrak{N}_+ = (\mathbb{N}, 0, 1, +, <)$, and by itself does not admit elimination of quantifiers. To see this, first, note that if a theory admits quantifier elimination then all the sets definable by formulas in the language of the theory are definable using just quantifier free formulas. Next, it is fairly easy to show that

the quantifier free formulas in the language of \mathfrak{N}_+ define either finite or co-finite sets — i.e., complements of finite sets. Now, take the formula stating that x is an even number, $\exists y x \approx y + y$. Note that the set defined by this formula is neither finite nor co-finite. Hence, can't be defined using a quantifier free formula.

However, it turns out, that the only thing we need to add to the language of PA, to have quantifier elimination, is congruence predicates or equivalently divisibility predicates. Hence, we extend the language of PA by adding for each $k \geq 2$ the congruence predicate symbols \equiv_k , and taking the theory of the structure $\mathfrak{N}_{\equiv} = (\mathbb{N}, 0, 1, +, <, \{\equiv_k\}_{k \geq 2})$, where each \equiv_k is the binary predicate of congruence modulo k . That is, $a \equiv_k b$ iff k divides $(a - b)$.

Theorem 2.3.2. The theory of the structure \mathfrak{N}_{\equiv} admits quantifier elimination.

We now give a quantifier elimination procedure for the theory of \mathfrak{N}_{\equiv} , based on Cooper's initial algorithm [9]. Recall, by Lemma 2.3.1, we only need to show that we can eliminate a single existentially quantified variable from a formula of the form $\exists x \varphi(x, \bar{y})$, where $\varphi(x, \bar{y})$ is a conjunction of literals. Let us denote by \mathbf{m} the m -fold sum of the term 1, i.e., $1 + \dots + 1$, representing the number m . Similarly, let us write mt for the m -fold sum of the term t .

Before presenting the procedure, we first establish the following lemma. It handles the simpler case in which all atoms are non-equalities and every coefficient of x is 1. We will rely on this result later.

Lemma 2.3.3. *We can compute, for any given formula $\varphi(x, \bar{y})$ of the form*

$$\exists x \left(\bigwedge_i u_i < x + u'_i \wedge \bigwedge_j v_j + x < v'_j \wedge \bigwedge_{\ell} w_{\ell} + x \equiv_{k_{\ell}} w'_{\ell} \right)$$

a quantifier free formula $\psi(\bar{y})$ that is equivalent (in \mathfrak{N}_{\equiv}) to $\varphi(x, \bar{y})$.

Proof. Let K be least common multiple of the k_{ℓ} , now note that for the congruence constraints, we can bound our search by making sure that we have considered one number from each of the congruence classes of $0, \dots, K - 1$; As for any $m \in \mathbb{N}$ we have, $\mathfrak{N}_{\equiv} \models \exists x \bigwedge_{\ell} w_{\ell} + x \equiv_{k_{\ell}} w'_{\ell}$ iff $\mathfrak{N}_{\equiv} \models \bigvee_{r \in \{0, \dots, K-1\}} \bigwedge_{\ell} w_{\ell} + \mathbf{m} + \mathbf{r} \equiv_{k_{\ell}} w'_{\ell}$. Now all that remains is handling the lower and upper bounds.

We construct ψ by considering two case, the first one when all the lower bounds are negative, i.e., $\bigwedge_i u_i < u'_i$, and hence trivially satisfied, and the second when there is at least one non trivial lower bound, i.e., $\bigvee_i u'_i \leq u_i$. Moreover, note that for a solution to exist, each of the upper bounds have to be greater than 0, i.e., $\bigwedge_j v_j < v'_j$. We let ψ be

the formula $\bigwedge_j v_j < v'_j \wedge (\bigwedge_i u_i < u'_i \rightarrow \psi_1) \wedge (\bigvee_i u'_i \leq u_i \rightarrow \psi_2)$ where ψ_1 and ψ_2 are as given below. Note, the difference being in the first case we start our search from 0 and in the second case from the maximum of the lower bounds $u_i - u'_i$.

$$\psi_1 := \bigvee_{r \in \{0, \dots, K-1\}} \left(\bigwedge_j v_j + \mathbf{r} < v'_j \wedge \bigwedge_\ell w_\ell + \mathbf{r} \equiv_{k_\ell} w'_\ell \right)$$

$$\psi_2 := \bigvee_{r \in \{0, \dots, K-1\}} \left(\left(\bigwedge_i \bigwedge_j v_j + u_i - u'_i + \mathbf{r} < v'_j \right) \wedge \left(\bigwedge_i \bigwedge_\ell w_\ell + u_i - u'_i + \mathbf{r} \equiv_{k_\ell} w'_\ell \right) \right)$$

Note that $-$ is not in our language however, it is easy to remove it by simply transposing the terms. \square

In what follows, we assume that atoms are of the form listed below. Note that this is not a restriction as any atom can be transformed into an equivalent atom having the desired form, by collecting like terms.

$$\begin{aligned} m_1 x + t &\approx u \\ m_1 x + t &< u \\ t &< m_1 x + u \\ m_1 x + t &\equiv_k u \end{aligned} \tag{2.1}$$

where t and u are x free terms. We now give the procedure for the general case $\exists x \varphi(x, \bar{y})$ where $\varphi(x, \bar{y})$ is a conjunction of literals.

Algorithm 1: Quantifier Elimination Procedure for \mathfrak{N}_\equiv .

Input: $\exists x \varphi(x, \bar{y})$ where $\varphi(x, \bar{y})$ is a conjunction of literals.

Output: a quantifier free ψ that is equivalent in \mathfrak{N}_\equiv to $\exists x \varphi(x, \bar{y})$.

Eliminate negation. Each negative literal can be eliminated by replacing it with an equivalent disjunction of atoms as follows. Replace the literal $t_1 \not\approx t_2$ by $t_1 < t_2 \vee t_2 < t_1$, the literal $t_1 \not< t_2$ by $t_2 \approx t_1 \vee t_2 < t_1$ and finally $t_1 \not\equiv_k t_2$ by $\bigvee_{r \in \{1, \dots, k-1\}} t_1 + \mathbf{r} \equiv_k t_2$. Note we can reorganize back into DNF and distribute the existential quantifier to obtain an equivalent formula of the form $\exists x \varphi_1 \vee \dots \vee \exists x \varphi_n$ where each φ_i is a conjunction of atoms of the form listed in 2.1. Hence, we can simply assume that $\varphi(x, \bar{y})$ is a conjunction of *atoms*. This will make the presentation succinct.

Uniformize coefficients. Next, let m_1, \dots, m_k be the coefficients of x in $\varphi(x, \bar{y})$ and let M be the least common multiple of the m_j . Multiply each atom by $\frac{M}{m_i}$, note that this will preserve equivalence. That is, we transform atoms of the following form

$$\begin{aligned} m_i x + t_i &\approx u_i, & m_i x + t < u_i, \\ t_i < m_i x + u_i, & & m_i x + t \equiv_{k_i} u_i \end{aligned}$$

to their corresponding equivalent forms

$$\begin{aligned} Mx + \frac{M}{m_i}t &\approx \frac{M}{m_i}u, & Mx + \frac{M}{m_i}t < \frac{M}{m_i}u \\ \frac{M}{m_i}t < Mx + \frac{M}{m_i}u, & & Mx + \frac{M}{m_i}t \equiv_{k \frac{M}{m_i}} \frac{M}{m_i}u \end{aligned}$$

Finally, we can replace Mx by a fresh variable z and assert that z is a multiple of M , i.e., add the atom $z \equiv_M 0$ to φ .

Eliminate equality. If at this point there is an equation of the form $z + u \approx t$ then we can simply eliminate z by replacing it with $t - u$ and adding the conjunct $t \leq u$. Again, we can transpose terms to account for $-$ not being in our language. Hence, at this point the formula has the following form, with lower and upper bounds of z , and congruence constraints:

$$\exists z \left(\bigwedge_i u_i < z + u'_i \wedge \bigwedge_j v_j + z < v'_j \wedge \bigwedge_\ell w_\ell + z \equiv_{k_\ell} w'_\ell \right)$$

Output an equivalent formula. Finally, if there are congruence constraints, then the formula is of the form described in Lemma 2.3.3, hence, we use Lemma 2.3.3 to compute an equivalent quantifier free formula. Otherwise, the formula simply asserts that there is a positive gap between the lower and upper bounds of z hence we output the following equivalent quantifier free formula:

$$\bigwedge_i \bigwedge_j u_i - u'_i + 1 < v'_j - v_j \wedge \bigwedge_j 0 < v'_j - v_j.$$

This completes the quantifier elimination procedure for the theory of \mathfrak{N}_{\equiv} . Presburger's initial quantifier elimination procedure had a non-elementary running time. The algorithm we presented above is based on Cooper's algorithm [9], which is itself an improved version of Presburger's algorithm. Cooper in [10] gave a new and more efficient version of his initial algorithm. The later improvements in [10] came by avoiding the (large amounts of) disjunctions introduced by the initial DNF transformation and the negation elimination step.

Moreover, it showed how to reduce the disjuncts introduced by Lemma 2.3.3. Following that, Oppen [41] showed that this new algorithm had a deterministic triply exponential running time.

Recently, the authors in [24] showed that a single block of existentially quantified variables can be eliminated in deterministic singly exponential time. This is rather interesting, as it directly contradicts an earlier claim by Weispfenning [56] asserting a doubly exponential lower bound for eliminating a single block of existential quantifiers. The authors in [24] show that this is indeed incorrect.

2.3.2 Linear Real Arithmetic

We now consider the theory of the reals with addition and order, for convenience we will add for each rational $c \in \mathbb{Q}$, a constant symbol c and a unary function symbol $c \cdot$ denoting multiplication by c , and take the theory of the structure $\mathfrak{R} = (\mathbb{R}, +, -, <, \{c, c \cdot\}_{c \in \mathbb{Q}},)$. Similar to the previous section, we show that this theory admits quantifier elimination by giving an explicit quantifier elimination procedure. Finally, we discuss complexity issues regarding quantifier elimination procedures for this theory. As in the previous section, it is sufficient to consider formulas of the form $\exists x \varphi(x, \bar{y})$, where $\varphi(x, \bar{y})$ is a conjunction of literals.

Algorithm 2: Fourier-Motzkin.

Isolate x : Note that any atom can be put into an equivalent atom of the form

$$x \approx t \quad x < u \quad u < x$$

where t, u are x free terms. Moreover, we can eliminate negation in a similar way as in the previous section. Hence, we may assume that φ is of the form

$$\bigwedge_i t_i < x \wedge \bigwedge_j x < u_j \wedge \bigwedge_k x \approx s_k$$

Eliminate x : If there is an equation $x \approx s$, then we can output $\varphi[x/s]$, otherwise we simply need to assert that all the lower bounds are less than that of the upper bounds:

$$\bigwedge_i \bigwedge_j t_i < u_j.$$

This completes our procedure. Note that this is a very simple but highly inefficient algorithm. A more efficient variant has been given in [38] reducing the doubly exponential

time to singly exponential time for systems of linear inequalities. However, for general formulas, i.e., those having arbitrary boolean structure, quantifier elimination in this theory has been shown to have a deterministic doubly exponential running time [55, 11].

2.4 Egraphs

An egraph is a well-known data structure to compactly represent a set of terms and an equivalence relation on those terms [40]. We assume that graphs have an ordered successor relation and use $n[i]$ to denote the i th successor (child) of a node n . We denote by $\text{deg}(n)$, the out degree of a node n , i.e., the number of edges leaving n .

Definition 2.4.1. (*Egraph [20]*) For a given signature Σ , an egraph is a tuple $G = \langle N, E, L, \text{root} \rangle$, where

- (a) $\langle N, E \rangle$ is a directed acyclic graph (possibly multigraph).
- (b) $L : N \rightarrow \Sigma^F \cup \mathcal{V}$ labels nodes by function and variable symbols s.t. nodes labelled by variables are leaves, and $\text{deg}(n) = k$ if n is labelled by a k -ary function symbol f .
- (c) $\text{root} : N \rightarrow N$ maps a node to its representative such that the relation $\rho_{\text{root}} := \{ (n, n') \mid \text{root}(n) = \text{root}(n') \}$ is closed under congruence w.r.t. root . That is, $(n, n') \in \rho_{\text{root}}$ whenever $L(n) = L(n')$, $\text{deg}(n) = \text{deg}(n') > 0$, and for $1 \leq i \leq \text{deg}(n)$, $(n[i], n'[i]) \in \rho_{\text{root}}$.

Given an egraph G , we let $\text{term}_G : N \rightarrow \Sigma^T$ be the function that maps nodes to their corresponding terms in the expected way. We assume the terms of different nodes are different (i.e., term is injective) and denote by n_t the node whose term is t . We denote by $\text{egraph}(\varphi)$ the egraph of φ built by the standard procedure given in [40], see Section 2 of [20] for more.

Chapter 3

Overview

We now discuss the intuition behind *locally ground theories* (Chapter 6) and \mathcal{T} -QEL, the algorithm presented in Chapter 7. Intuitively, \mathcal{T} -QEL is based on the following observation: for any formula of the form $\exists x \varphi$, if $\varphi \models_{\mathcal{T}} x \approx s$ for some *arbitrary* ground term s , then $\models_{\mathcal{T}} \exists x \varphi \leftrightarrow \varphi[x/s]$. Hence, effectively eliminating the quantified variable x . As discussed in the introduction, directly applying this observation to eliminate quantified variables might lead to non-terminating behavior, as there are infinitely many ground terms to check.

In the example below we illustrate via a sample formula and a theory, how the search space for entailed ground definitions can be reduced to a finite one. We further demonstrate how an egraph can be used to give a concrete partial quantifier elimination algorithm. Note that, our goal is to eliminate quantified variables in *polynomial time*, while maintaining *relative completeness* — that is, we want to eliminate every variable that has an entailed ground definition.

3.1 Motivating Example

Fix a signature $\Sigma_{po} = (\Sigma^F, \{\leq\})$, where Σ^F contains countably many free function and constant symbols, and let \mathcal{T}_{po} be the theory of partial orders in this signature. That is, \leq is reflexive, transitive, and antisymmetric. Now consider the formula $\exists \bar{x} \varphi(\bar{x})$, where

$$\varphi(\bar{x}) := x_1 \approx f(c) \wedge f(x_2) \approx g(x_1) \wedge x_3 \leq g(f(x_2)) \wedge g(f(x_2)) \leq x_3$$

and c, f, g are free constant and function symbols. It is straightforward to see that x_1 has an entailed ground definition, however, it is not so clear for x_2 and x_3 . In fact, note for

the term $g(g(f(c)))$, which is not a subterm of φ , we have $\varphi(\bar{x}) \models_{\mathcal{T}_{po}} x_3 \approx g(g(f(c)))$; this would be completely missed by QEL. It is essential that we detect such entailments to guarantee relative completeness. We proceed as follows:

Compute the partition S_φ/\sim of the subterms. Let \sim be the congruence relation defined on the subterm set of φ , S_φ , such that for all $t_1, t_2 \in S_\varphi$,

$$t_1 \sim t_2 \text{ iff } \varphi \models_{\mathcal{T}_{po}} t_1 \approx t_2.$$

Then, note that the quotient set

$$S_\varphi/\sim = \{ \{x_1, f(c)\}, \{c\}, \{x_2\}, \{f(x_2), g(x_1)\}, \{x_3, g(f(x_2))\} \}$$

represents terms which are equivalent under *every* model of φ and \mathcal{T}_{po} . Here is where the locality of the theories is crucial, to compute the partition in polynomial time. Let us use the notation $[t]$ to refer to the equivalence classes in S_φ/\sim , e.g. $[x_1] = \{x_1, f(c)\}$.


Construct a ground term. Next, we pick for each equivalence class in S_φ/\sim (if we can) a *representative* term that is either *already* ground or can be transformed via substitution of equivalent terms into an *equivalent* ground term. For the class $[x_1]$ we pick $f(c)$, for the class $[c]$ we pick c , and the class $[x_2]$ does not contain such a term. For the class $[f(x_2)]$, we note that $x_1 \sim f(c)$ and hence by functional congruence we know $\varphi \models_{\mathcal{T}_{po}} g(x_1) \approx g(f(c))$, thus we pick $g(x_1)$. With similar reasoning, we can see $\varphi \models_{\mathcal{T}_{po}} g(f(x_2)) \approx g(g(f(c)))$, hence we pick $g(f(x_2))$ for the class $[x_3]$. Interestingly, if a variable has *any* entailed ground definition, its equivalence class will contain such a *constructively* ground term.

Eliminate by substitution. Finally, we note that $\varphi(\bar{x})$ is equivalent (in \mathcal{T}_{po}) with the following formula:

$$f(c) \approx f(c) \wedge f(x_2) \approx g(f(c)) \wedge g(g(f(c))) \leq g(g(f(c))) \wedge g(g(f(c))) \leq g(g(f(c))).$$

Which we obtained by simply replacing each variable and representative with their equivalent ground formula we inferred in the above step. Hence, we output the formula $\exists x_2 f(x_2) \approx g(f(c)) \wedge g(g(f(c))) \leq g(g(f(c)))$ after removing redundant atoms.

Representation of S_φ/\sim on an egraph. Egraphs give us a compact way to represent congruence relations. For local theories such as \mathcal{T}_{po} , representing the partition S_φ/\sim amounts to saturating the egraph with implied atoms that are formed over the subterms of φ . Locality here allows us to consider only those atoms whose subterms are already in φ . We formalize this notion in Chapter 7. The next section provides a high level overview of the algorithm \mathcal{T} -QEL.



(a) Initial egraph of φ obtained by treating \leq as a free function symbol.

(b) The egraph obtained after instantiations of the axioms, representing the partition S_φ/\sim .¹

Figure 3.1: An example egraph of a formula and a final transformed egraph representing the partition S_φ/\sim .

3.2 High level overview of \mathcal{T} -QEL

The key insight behind \mathcal{T} -QEL is that given a formula in some locally ground theory, we can compactly represent on top of an egraph, the partition S_φ/\sim that is induced by \sim on the subterms of φ . After which, we can efficiently search through the equivalence classes and eliminate variables that have an entailed ground definition. As highlighted earlier, for the locally ground theories, searching for a *constructively ground term* in a variable’s class is sufficient.

A high level overview of the steps of the algorithm \mathcal{T} -QEL is given below. The idea, is to exploit the locality of \mathcal{T}_{p_0} to derive and represent *all* the implied equalities (over the subterms of φ) on the egraph. We will formally explore the notion of *locality* and *restricted derivations* in Chapters 4, 5 and 6. These concepts will help us confine our proof search to only those that are over the subterms of φ . For now we proceed informally.

Continuing with the example from the previous section, we now illustrate how the partition S_φ/\sim is represented on an egraph. Followed by how partial quantifier elimination is carried out.

Create the egraph. We first construct the egraph of φ using the standard procedure described in [40] and [20], this is shown in Figure 3.1a. Note that, at this stage we simply treat both function and predicate symbols uniformly as free functions. We indicate nodes in the same class, i.e., those having the same *root*, by drawing a red dashed arrow to a common node. A node with no outgoing red dashed arrow is its own *root*.

¹To avoid cluttering, some added predicate symbols during saturation that are not relevant to our discussion are omitted.

Saturate the egraph. Next we represent, i.e. add, to the egraph all the implied atoms that are over the subterms of φ . Representing simply means, that, when an equality $t \approx s$ is inferred we merge the equivalence classes of t and s , otherwise if a non equality atom, $P(t_1, \dots, t_k)$, is inferred we add a node labeled with P having as children the nodes n_{t_1}, \dots, n_{t_k} . In the case of the previous example, from $x_3 \leq g(f(x_2)) \wedge g(f(x_2)) \leq x_3$, we infer the equality $x_3 \approx g(f(x_2))$. Hence, we merge the equivalence classes of the nodes n_{x_3} and $n_{g(f(x_2))}$. This is shown in Figure 3.1b, by drawing a red dashed arrow to indicate the representative node.

Compute β a maximally ground representative function² and the set *core*. Now at this point, every implied equality between the terms is represented in the egraph. That is for every $t, s \in S_\varphi$ with $t \sim s$ we have merged the equivalence classes of their nodes — i.e., $t \sim s$ iff $root(n_t) = root(n_s)$. Moreover, we have for every $t \in S_\varphi$ that has an entailed ground definition, its class, $class(n_t)$, is ground as \mathcal{T}_{po} is locally ground. Intuitively, a maximally ground representative function picks a term that is either a *ground* or can be transformed into an equivalent ground term. The set *core* contains those terms that should be part of the final output. For both of these, we use QEL on the saturated egraph.

Substitute and extract an equivalent formula. Finally we use QEL, and the computed β and *core* to output the formula $\psi(\bar{y})$. This step requires extracting a formula from the egraph, which can result in unexpected behavior, if care is not taken. Extraction is thoroughly discussed in [20].

Chapter 4

Locality: A Proof Theoretic Characterization

In the current and the next chapter, we consider *local theories*. Local theories have a nice proof theoretic and semantic characterization that give tractability results for the *uniform word problem*. Our approach in this chapter, as opposed to the next, is syntactic, and has a proof theoretic flavor. We approach the semantic concept of a tractable uniform word problem through the proof theoretic notion of a tractable inference. This naturally gives rise to a proof theoretic algorithm for the uniform word problem. Following that, we discuss an interesting connection of local theories to two concepts: the characterization of the complexity class P, and the simplicity of proofs.

The core concept behind the proof theoretic characterization of local theories is that of *restricting* derivations to the subterms that occur within the given formula. We adopt the approach due to Givan and McAllester [22], and give a proof system for the local theories. We do, however, briefly mention in Section 4.3, a different approach by Negri and Plato [39]. In [39], they consider the concept of confining deductions to subterms, however, they do so by extending either the natural deduction or sequent calculus, with rules corresponding to the axioms of a (mathematical) theory.

Finally, we conclude the chapter with an interesting discussion connecting Hilbert's 24th problem to restricted derivations in local theories. Which on a high level is concerned with the notion of *simplicity* in proofs [39, 53, 52].

4.1 Derivations

We now introduce, in FOL without equality, a simple proof calculus for theories axiomatized by a finite set of Horn clauses. We prove for the calculus soundness and completeness with respect to semantics of Horn clauses. Finally, we give a definition for *restricted derivations* based on the calculus, and introduce the *local theories* in FOL without equality.

Given a finite set of atoms $\Gamma \cup \{\theta\}$, and Φ a finite set of Horn clauses, our goal is to construct a proof calculus such that θ is provable from $\Gamma \cup \Phi$ iff it is entailed by $\Gamma \cup \Phi$, i.e., $\Gamma \cup \Phi \models_{neq} \theta$. We can of course use any proof calculus for FOL without equality, however, as we will show here, for the case of Horn clauses, we can have a simpler calculus. The significance of this will become apparent when we give the algorithm in Chapter 7.

Definition 4.1.1. (*Derivation*). *Let $\Gamma \cup \{\theta\}$ be a set of atoms, then, a derivation of θ from Γ using the axioms Φ , is θ if $\theta \in \Gamma$ and otherwise is a sequence of atomic formulas $\theta_1, \dots, \theta_n$ s.t. θ_n is θ and for each θ_i there exists a Horn clause $\forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi) \in \Phi$ and a substitution σ such that θ_i is $\psi\sigma$ and each $\psi_j\sigma$ is either in Γ or appears earlier in the derivation.*

An intuitive way to look at the above definition is as follows. We convert each Horn clause $\forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi)$ in the axioms Φ to a corresponding inference rule as shown in Figure 4.1. Then, a derivation of θ from Γ in this calculus can be seen as simply θ when $\theta \in \Gamma$, and otherwise a sequence of atomic formulas $\theta_1, \dots, \theta_n$ where each θ_i is obtained via an application one of the inference rules obtained from the axioms, shown in Figure 4.1b and 4.1c.

$$\begin{array}{ccc}
 \frac{}{\theta} \quad \theta \in \Gamma & \frac{}{\psi\sigma} & \frac{\psi_1\sigma \quad \dots \quad \psi_k\sigma}{\psi\sigma} \\
 \text{(a) If } \theta \in \Gamma. & \text{(b) Case } k = 0. & \text{(c) Case } k > 0.
 \end{array}$$

Figure 4.1: Axioms to inference rules, where σ is a substitution.

We write $\Gamma \cup \Phi \vdash \theta$ if there is a derivation of θ from Γ using Φ . In what follows, we show that our calculus is sound and complete for Horn clause theories.

Theorem 4.1.2. (Soundness). *Let Φ be a finite set of Horn clauses and $\Gamma \cup \{\theta\}$ a set of atoms. Then, if $\Gamma \cup \Phi \vdash \theta$ then $\Gamma \cup \Phi \models_{neq} \theta$.*

Proof. This can be proven by a straightforward induction on the length of the derivation. Let $\theta_1, \dots, \theta_n$ be a derivation of θ from $\Gamma \cup \Phi$, then for the base case, either $\theta \in \Gamma$, in which case the conclusion follows immediately, or for some Horn clause $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi)$ in Φ , θ is $\psi\sigma$ and each of $\psi_i\sigma$ is in Γ . Hence, $\Gamma \cup \Phi \models_{neq} \theta$. The inductive step is a straightforward application of the inductive hypothesis. \square

We next show that the calculus is complete for Horn clause theories. Towards showing this, we introduce *Herbrand models* and construct the least Herbrand model for $\Gamma \cup \Phi$. Moreover, it is shown that the least Herbrand model is closely related to the *fixed point semantics* of logic programs.

Fix a signature Σ , and let \mathcal{B} be the set of all Σ -atoms, i.e., $\mathcal{B} = \{P(t_1, \dots, t_n) \mid t_i \text{ is a term in } \Sigma\}$. A Herbrand interpretation \mathcal{I} is an interpretation such that its universe $A \subseteq \Sigma^T$, and $t^{\mathcal{I}} = t$ for each term t . Since any two Herbrand interpretations differ only in the relations they assign to the predicate symbols, we can view them as being a subset of \mathcal{B} . That is \mathcal{I} as a subset of \mathcal{B} can be viewed as the set of atoms that are true in \mathcal{I} . We will use the notation \mathcal{I}^A to denote to the set of atoms true in \mathcal{I} .

For a given set of Horn clauses Φ , define the *immediate consequence* operator $T_\Phi : \mathcal{P}(\mathcal{B}) \rightarrow \mathcal{P}(\mathcal{B})$ to be the map such that for all $A \subseteq \mathcal{B}$,

$$T_\Phi(A) = \{ \psi\sigma \mid \forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi) \in \Phi, k \geq 0, \text{ and for each } i, \psi_i\sigma \in A \}.$$

We are now ready to define the *least Herbrand model* of $\Gamma \cup \Phi$, which will be central to the completeness theorem for the calculus. We simply denote least Herbrand model as \mathcal{H} . We construct \mathcal{H} in stages as follows, let

$$H_0 = \Gamma, \text{ and for } n > 0,$$

$$H_n = H_{n-1} \cup T_\Phi(H_{n-1}).$$

Finally, we let \mathcal{H} be the Herbrand interpretation for which the set of atoms true in \mathcal{H} , \mathcal{H}^A , is exactly the union of the above sets. Formally, for every predicate symbol P and t_1, \dots, t_k in its universe H , we let

$$(t_1, \dots, t_k) \in P^{\mathcal{H}} \text{ iff } P(t_1, \dots, t_k) \in \bigcup_{n < \omega} H_n.$$

Lemma 4.1.3. \mathcal{H} is a Herbrand model of $\Gamma \cup \Phi$. Moreover, it is the least Herbrand model, i.e., $\mathcal{H}^A = \bigcap_{\mathcal{I} \models \Gamma \cup \Phi} \mathcal{I}^A$, where \mathcal{I} is a Herbrand model of $\Gamma \cup \Phi$.

Proof. To see that \mathcal{H} is a Herbrand model of $\Gamma \cup \Phi$, note that $\Gamma = H_0$, hence $\mathcal{H} \models \Gamma$. Next, assume $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi)$ is in Φ , now assume for an arbitrary t_1, \dots, t_k in the universe H that $\mathcal{H}[t_1, \dots, t_k] \models (\psi_1 \wedge \dots \wedge \psi_k)$, then by definition each $\psi_i \sigma$, where $\sigma(x_i) = t_i$, is in some H_n , hence $\psi \sigma$ is in H_{n+1} . Therefore, $\mathcal{H}[t_1, \dots, t_k] \models \psi$, and as t_1, \dots, t_k were arbitrary, we have $\mathcal{H} \models \forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi)$.

We next show that \mathcal{H} is the least Herbrand model.

(\supseteq). This direction is straight forward as \mathcal{H} itself is a Herbrand model of $\Gamma \cup \Phi$. Hence, $\mathcal{H}^A \supseteq \bigcap_{\mathcal{I} \models \Gamma \cup \Phi} \mathcal{I}^A$.

(\subseteq). Let \mathcal{I} be an arbitrary Herbrand model of $\Gamma \cup \Phi$ and ψ be an atom in \mathcal{H}^A . Then by definition $\psi \in H_n$ for some $n \geq 0$, we proceed by induction on the least such n . For the base case, ψ in $H_0 = \Gamma$, hence ψ is true in \mathcal{I} , and thus $\psi \in \mathcal{I}^A$. Now assume $\psi \in H_n \setminus H_{n-1}$, for some $n > 0$, then there is a Horn clause $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi')$ in Φ , such that the atoms $\psi_i \sigma$ are in $H_{n-1} \subseteq \mathcal{H}^A$, and ψ is $\psi' \sigma$. Now by the inductive hypothesis, each $\psi_i \sigma$ is in \mathcal{I}^A , hence as $\mathcal{I} \models \Phi$, we have $\psi' \sigma \in \mathcal{I}^A$. \square

Another interesting way to look at the least Herbrand model \mathcal{H} is as the least fixed point of the *immediate consequence operator* T_Φ [36]. To facilitate discussion, we will stop making a distinction between a Herbrand interpretation \mathcal{I} and the set of atoms true in it \mathcal{I}^A . We will simply use \mathcal{I} as both the interpretation and the set of atoms true in \mathcal{I} . With this in mind, we can start viewing the set of all Herbrand interpretations as a complete lattice ordered by set inclusion \subseteq . We will now show that \mathcal{H} is indeed the least fixed point of T_Φ . We will need the following two lemmas.

Lemma 4.1.4. *An interpretation \mathcal{I} is a Herbrand model of $\Gamma \cup \Phi$ iff it is a fixed point of T_Φ , that is, $\mathcal{I} \models \Gamma \cup \Phi$ iff $T_\Phi(\mathcal{I}) = \mathcal{I}$.*

Proof. Both directions follow from the definition of T_Φ . \square

Lemma 4.1.5. *The immediate consequence operator T_Φ is monotonic on the complete lattice $\mathcal{P}(\mathcal{B})$, the set of all Herbrand interpretations.*

Proof. Let $\mathcal{I} \subseteq \mathcal{J} \subseteq \mathcal{B}$, then we need to show $T_\Phi(\mathcal{I}) \subseteq T_\Phi(\mathcal{J})$. Assume some ψ is in $T_\Phi(\mathcal{I})$, then either $\psi \in \Gamma$ in which case $\psi \in T_\Phi(\mathcal{J})$, or, there exists some Horn clause $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi')$ in Φ , such that, ψ is $\psi' \sigma$, and each $\psi_i \sigma$ are in $\mathcal{I} \subseteq \mathcal{J}$. Hence, $\psi' \sigma$ is in $T_\Phi(\mathcal{J})$. \square

Now, to see that \mathcal{H} is indeed the least fixed point, note that by Lemma 4.1.3 we have $\mathcal{H} = \bigcap_{\mathcal{I} \models \Gamma \cup \Phi} \mathcal{I}$. Hence, by Lemma 4.1.4, $\mathcal{H} = \inf \{ \mathcal{I} \mid T_\Phi(\mathcal{I}) = \mathcal{I} \}$. Finally, as T_Φ is monotonic (Lemma 4.1.5), by the Knaster-Tarski Theorem (Theorem 2.2.1), we have, $\text{lfp}(T_\Phi) = \inf \{ \mathcal{I} \mid T_\Phi(\mathcal{I}) = \mathcal{I} \} = \mathcal{H}$.

We now move on to prove the completeness theorem for the calculus. After seeing how we can construct the least Herbrand model \mathcal{H} inductively in stages, it is not surprising to see that the calculus we introduced is sufficient to capture logical consequence for Horn clause theories.

Theorem 4.1.6. (Completeness). Given a finite set of Horn clauses Φ and a set of atoms $\Gamma \cup \{ \theta \}$, if $\Gamma \cup \Phi \models_{neq} \theta$ then $\Gamma \cup \Phi \vdash \theta$.

Proof. Assume that $\Gamma \cup \Phi \models_{neq} \theta$, then θ is true in every Herbrand model of $\Gamma \cup \Phi$. In particular, it is true in \mathcal{H} , which is the least Herbrand model by Lemma 4.1.3, hence $\theta \in H_n \subseteq \mathcal{H}^A$ for some $n \geq 0$. We proceed by induction on the least such n . The base case is straightforward, as the length 1 derivation θ is a valid derivation. Now assume $\theta \in H_n \setminus H_{n-1}$ for some $n > 0$, then for some $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi)$, we have for some substitution σ , each $\psi_i \sigma \in H_{n-1}$ and θ is $\psi \sigma$. By the inductive hypothesis, for each $\psi_i \sigma$ we have $\Gamma \cup \Phi \vdash \psi_i \sigma$. Let $\theta_1^i \dots \theta_{m_i}^i$ be a derivation of $\psi_i \sigma$, and i_1, \dots, i_ℓ be indexes such that ψ_{i_j} is not in Γ . Then, by concatenating the derivations of each $\psi_{i_j} \sigma$, we have the valid derivation, $\theta_1^{i_1} \dots \theta_{m_{i_1}}^{i_1} \dots \theta_1^{i_\ell} \dots \theta_{m_{i_\ell}}^{i_\ell} \theta$, for θ . Hence, $\Gamma \cup \Phi \vdash \theta$. \square

Definition 4.1.7. (*Restricted Derivations in FOL without equality*) A restricted derivation of θ from Γ using the axioms Φ is a derivation, $\theta_1, \dots, \theta_k$, of θ such that for each θ_i we have $S_{\theta_i} \subseteq S_\Gamma \cup S_\theta$.

We write $\Gamma \cup \Phi \Vdash_{neq} \theta$ if there exists a restricted derivation of θ from Γ using Φ . We are now ready to introduce the local theories. A *local theory* in FOL without equality is a theory axiomatized by a finite set of Horn clauses Φ such that for every set of atoms $\Gamma \cup \{ \theta \}$, $\Gamma \cup \Phi \vdash \theta$ iff $\Gamma \cup \Phi \Vdash_{neq} \theta$. This immediately gives us polynomial time decidability of the *uniform word problem* for local theories.

Lemma 4.1.8. (*Tractability Lemma [22]*) Let \mathcal{T} be a local theory and $\Gamma \cup \{ \theta \}$ a finite set of atoms then we can decide $\Gamma \models_{\mathcal{T}} \theta$ (in FOL without equality) in polynomial time.

4.2 Undecidability of Locality and Characterizing P

We now turn our attention to two complexity issues related to locality. We first show that locality *as a property* of a given theory is in general undecidable. Hence, we can't uniformly determine whether or not a given theory is local. Following that, we give an interesting result which shows how to characterize the complexity class P using a set of local Horn clauses. We now state and prove the undecidability result.

Theorem 4.2.1 (Undecidability [22]). Given a finite set of Horn clauses Φ , the problem of determining whether $\mathcal{T} = \text{Cn } \Phi$ is a local theory is undecidable.

The proof for the above theorem proceeds by reducing the Halting problem for Turing machines to the problem of determining the locality of theories. We now introduce Turing machines and how computation in a Turing machine is carried out. Following that, we present a proof of the theorem.

Definition 4.2.2. We take a Turing machine to be a 5-tuple $(Q, \Omega, \delta, q_0, q_h)$, where:

- (i) Q is a finite set of states.
- (ii) Ω is a finite set denoting the alphabet, and contains the blank symbol \sqcup .
- (iii) $\delta : (Q \setminus \{q_h\}) \times \Omega \rightarrow Q \times \Omega \times \{L, R\}$ is the transition function.
- (iv) $q_0 \in Q$ is designated as a start state.
- (v) $q_h \in Q$ denotes the halting state.

Informally, we imagine a Turing machine M as having an infinite tape divided into cells, and equipped with a head to read and write to the tape. The machine M moves its head left and right, and computes by reading and writing to a single cell in the tape. Formally, we describe the computation of M as follows. A *configuration* of the machine M is a sequence of the form uqv where $u, v \in \Omega^*$ and $q \in Q$. This represents a snapshot of the tape content and the current state. We describe a single step computation in M using the binary predicate \rightarrow on configurations such that, $uqav \rightarrow ucq'v$ whenever $\delta(q, a) = (q', c, R)$ and $ubqav \rightarrow uq'bcv$ whenever $\delta(q, a) = (q', c, L)$. We handle the special cases where any of u, v, a, b are missing in the expected way; for instance we would add $uq \rightarrow ucq'$ whenever $\delta(q, \sqcup) = (q', c, R)$.

Then, a *computation* of M on the input string w is a possibly infinite sequence of configurations C_0, \dots, C_k, \dots such that,

- (i) C_0 is the *start* configuration, i.e., C_0 is q_0w , and
- (ii) For all $i \geq 0$, $C_i \rightarrow C_{i+1}$.

We say that M *halts* on the input string w iff there exists a finite computation C_0, \dots, C_k of M on w , where C_k is a halting configuration, i.e., C_k is uq_hv for some $u, v \in \Omega^*$. We now move on to the proof of Theorem 4.2.1.

The high level idea of the proof is that given the Turing machine M we construct a local set of Horn clauses ¹ Φ'_M such that M halts on the empty string iff there exists a term t such that $\Phi'_M \vdash H(t)$. We encode computations of M using terms, and the atom $H(t)$ intuitively says that t is a halting computation. From Φ'_M we then construct another finite set of Horn clauses Φ_M such that M halts on the empty string iff $\text{Cn } \Phi_M$ is local.

Let us fix a signature Σ , in which the constant and function symbols are the ones described below. For each $q \in Q$ and $a \in \Omega$, we include in our signature Σ , constant symbols q and a . We add the binary function symbols \cdot and $\#$ to our signature. Their purposes is to aid with representing string concatenation on Ω and sequence of configurations respectively. This way we can represent finite computations using terms, for instance the term² $q_0w\#\dots\#u_kq_kv_k$ would encode the computation $q_0w, \dots, u_kq_kv_k$. We assume the atom $P_{lst}(t_1\#\dots\#t_k, x_4)$ outputs the last configuration t_k as x_4 . It is straightforward to define P_{lst} using superficial Horn clauses.

Proof. (of Theorem 4.2.1). We first construct a finite set of *superficial* Horn clauses Φ_s such that M halts on the empty string iff there exists some term t s.t. $\{P_{in}(t)\} \cup \Phi_s \vdash H(t)$, for designated atoms P_{in} and H . We let $\Phi_s = \Psi_{sub} \cup \Psi_\delta \cup \Psi_{\delta^*} \cup \Psi_{halt}$. Where each set Ψ contains the superficial Horn clauses as outlined below.

Collect the subterms. The set Ψ_{sub} contains the following Horn clauses:

$$\forall x (P_{in}(x) \rightarrow P_{sub}(x)), \quad \forall \bar{x} (P_{sub}(f(t_1, \dots, t_k)) \rightarrow \bigwedge P_{sub}(t_i))$$

Encode δ and one step computations. Let Ψ_δ contain the following clauses:

$$\begin{aligned} \forall \bar{x} (P_{sub}(x_1qax_2) \wedge P_{sub}(x_1cq'x_2) \rightarrow P_\delta(x_1qax_2, x_1cq'x_2)), \quad \text{where } \delta(q, a) = (q', c, R) \\ \forall \bar{x} (P_{sub}(x_1bqax_2) \wedge P_{sub}(x_1q'bcx_2) \rightarrow P_\delta(x_1bqax_2, x_1q'bcx_2)), \quad \text{where } \delta(q, a) = (q', c, L) \end{aligned}$$

¹That is, $\text{Cn } \Phi$ is a local theory.

²We use infix notation to help with readability, i.e., we write $t_1\#t_2\#t_3$ for the term $\#(t_1, \#(t_2, t_3))$.

Encode all valid computations. We let Ψ_{δ^*} contain the following clauses.:

$$\begin{aligned} & \forall \bar{x} (P_{\delta}(x_1, x_2) \rightarrow P_{\delta^*}(x_1, x_2)) \\ & \forall \bar{x} (P_{sub}(x_3 \# x_4) \wedge P_{\delta}(x_1 q x_2, x_3) \wedge P_{\delta^*}(x_3, x_4) \rightarrow P_{\delta^*}(x_1 q x_2, x_3 \# x_4)) \end{aligned}$$

Verify halting computations. Finally, Ψ_{halt} contains the following clause:

$$\forall \bar{x} (P_{in}(q_0 \# x_1) \wedge P_{lst}(x_1, x_2 q_h x_3) \wedge P_{\delta^*}(q_0, x_1) \rightarrow H(q_0 \# x_1))$$

It is straightforward to show that M halts on the empty string iff there is some term t s.t. $\Phi_s \cup \{P_{in}(t)\} \vdash H(t)$. We now introduce a fresh predicate symbol P' of arity $n+1$ for each of the n -ary predicate symbol in Φ_s and we let Φ'_M be the the set consisting of the Horn clauses,

$$\forall x P'_{in}(x, x), \quad \forall x (H'(x, x) \rightarrow H(x)), \quad \forall x \forall \bar{x} (\bigwedge P'_i(x, t_1, \dots, t_{k_i}) \rightarrow P'(x, t_1, \dots, t_k))$$

Where x is a fresh variable, and $\forall \bar{x} (\bigwedge P'_i(t_1, \dots, t_{k_i}) \rightarrow P(t_1, \dots, t_k))$ is a clause in Φ_s . We can use a similar argument to Theorem 2 in [22] to see that Φ'_M is a local set of Horn clauses. Moreover, we can see that $\Phi_s \cup \{P_{in}(t)\} \vdash H(t)$ iff $\Phi'_M \vdash H(t)$. Hence M halts iff there exists a term t s.t. $\Phi'_M \vdash H(t)$. Finally, we let Φ_M be $\Phi'_M \cup \{\forall x H(x) \rightarrow P_{halt}\}$ for a fresh 0-ary predicate symbol P_{halt} . We now show that $\text{Cn } \Phi_M$ is local iff M does not halt on the empty string.

(\implies) We prove the contrapositive. Assume that M halts on the empty string, then for some term t , we have $\Phi'_M \vdash H(t)$, and hence $\Phi_M \vdash P_{halt}$. However, note that any derivation of P_{halt} must be of the form $\theta_1, \dots, H(t'), P_{halt}$. Hence, $\Phi_M \not\vdash P_{halt}$ and hence $\text{Cn } \Phi_M$ is not local.

(\impliedby) Conversely, assume that M doesn't halt. We must show for every finite set of atoms $\Gamma \cup \{\theta\}$, $\Gamma \cup \Phi_M \vdash \theta \implies \Gamma \cup \Phi_M \Vdash \theta$. We can assume w.l.o.g. that $P_{halt} \notin \Gamma$. Now, assume $\Gamma \cup \Phi_M \vdash \theta$, in the case θ is different from the atom P_{halt} we have,

$$\begin{aligned} \Gamma \cup \Phi_M \vdash \theta & \iff \Gamma \cup \Phi'_M \vdash \theta \\ & \iff \Gamma \cup \Phi'_M \Vdash \theta \\ & \iff \Gamma \cup \Phi_M \Vdash \theta. \end{aligned}$$

Otherwise, θ is P_{halt} and we have for some term t a restricted derivation $\theta_1, \dots, \theta_k, H(t)$ of $H(t)$ from $\Gamma \cup \Phi_M$. We claim $t \in S_{\Gamma}$. Now, as we have assumed M does not halt, Γ is

non empty. Moreover, we have either $H(t) \in \Gamma$ or for some θ_i in the derivation, the clause $\forall \bar{x} (\bigwedge \psi_j \rightarrow \psi)$ in Φ'_M was used such that θ_i is $\psi\sigma$ and some $\psi_j\sigma$ is in Γ . Now note that, by the way we have constructed Φ'_M , the term in first position of each θ_i in the derivation is the same. Hence, t occurs in some formula $\psi\sigma$ in Γ , and thus t is in S_Γ . We therefore, have the restricted derivation $\theta_1, \dots, H(t), P_{halt}$ for P_{halt} from $\Gamma \cup \Phi_M$. \square

We now move on to show how local theories characterize the complexity class P. In descriptive complexity [31, 30], one studies the complexity of computational problems via the complexity of how hard it is to express (define) them in *some* logic. There are several interesting results relating complexity classes to logics. We mention two of them here. The first one, is the seminal theorem of Fagin's [17] relating the class NP with the existential fragment of second order logic (SO \exists): A problem (a set of finite structures) \mathcal{K} is in the class NP iff there exists an existential second order (SO \exists) sentence φ that defines it. That is $\mathcal{K} \in NP$ iff $\mathcal{K} = \{ \mathfrak{A} \mid \mathfrak{A} \models \varphi \}$ for some SO \exists sentence φ .

The second result due to Immerman [29] and Vardi [54] relates in the same way, the class P with FOL plus a least fixed point operator. This time however, we assume there is a linear order on the domain of the finite structures. We conclude this section by stating a result from [22], with somewhat different flavor from the above, which characterizes the complexity class P using *local theories*. The proof makes use of the result of Immerman [29] and Vardi [54]. We refer the interested reader to Sections 3 and 4 of [22] for the full proof.

Theorem 4.2.3. Let Σ be a signature, and $R \subseteq (\Sigma^T)^k$ be a polynomial-time computable k -ary relation on the terms generated from Σ . Then, there exists a set of local Horn clauses, Φ , such that $(t_1, \dots, t_k) \in R$ iff $\Phi \vdash P_R(t_1, \dots, t_k)$, where P_R is a designated predicate symbol representing the k -ary relation R .

4.3 A Brief Interlude: Hilbert’s 24th problem

This section reflects on Hilbert’s 24th problem and its relevance to proof simplicity in local theories. It was interesting to see in [52, 53] that in addition to the famous 23 problems Hilbert posed in Paris, Hilbert had in mind one more problem that was not presented. The omitted 24th problem is concerned, on a high level, about finding simpler proofs and criteria for measuring simplicity. It was discovered by Rüdiger Thiele [52], in the notes of Hilbert, and reads as follows:

The 24th problem in my Paris lecture was to be: Criteria of simplicity, or proof of the greatest simplicity of certain proofs. Develop a theory of the method of proof in mathematics in general. Under a given set of conditions there can be but one simplest proof..

Hilbert, however, did not outline a clear criteria for determining the *simplicity* of proofs. There are of course criteria that seem natural to use. One can for instance use the length of proofs, or the total number or structure of terms that occur in a proof as a measure of simplicity [53]. In both of these cases, it is natural to ask whether in local theories, *restricted derivations* are the *simplest* proofs. Although a definitive answer – either affirming or refuting this – needs more work, there are a few properties of restricted derivations that seem to be in support of it.

First, note that logical axioms are not used in the proof system for local theories, the only axioms used are that of the theory itself. Second, provable formulas in local theories admit *polynomial length* proofs via restricted derivations. Third, every term that appears in the restricted derivation of a formula is a *subterm* already occurring in the formula itself. All of these seem to suggest that restricted derivations have simple structure.

We refer the interested reader to [39] for further discussions of other proof systems having the *subterm property* discussed above. In Parts I and II of [39], Negri and Plato analyze proofs in axiomatic theories by extending both the natural deduction and sequent calculus with rules corresponding to the axioms. They show, via normalization, for several axiomatic theories — such as quasi-orders, lattices, and groupoids — that the resulting extended system enjoys the subterm property.

Chapter 5

Locality: A Semantic Characterization

We now look at the notion of *locality* of theories in first order logic with equality, paying close attention to the polynomial time decidability of the uniform word problem. Our approach in this chapter is semantical in nature. After giving a semantic definition of *locality*, we explore a close relationship with two other concepts that lead to polynomial time decidability of the uniform word problem. The first of these, is Evans' [16] embeddability criteria, giving a sufficient condition for polynomial time decidability of the uniform word problem via embeddability of *finite partial algebras*. The second one is Burris' [6] certain *axiomatizability* criteria giving a similar result on polynomial time decidability.

We further give a definition for *stable locality*, a more general notion that is closely related to locality. We then show that Evans' [16] embeddability and Burris' [6] axiomatizability criteria lie in between locality and the more general notion of *stable locality*. While a weaker version of Evans' embeddability criteria is shown to coincide exactly with locality. This weaker version gives us an easy way of identifying local theories in first order logic with equality.

Perhaps a noteworthy point is that, even though the methods of Evans and Burris are essentially semantic, they have a close resemblance to the methods of the last chapter. All of the methods exploit structures induced by the linearly many terms of the input formula. In the case of Evans and Burris, these are algebraic structures, and in the case of the last chapter, these are deductive structures.

5.1 Basic Notions

In what follows, we assume a fixed but arbitrary algebraic signature $\Sigma = (\Sigma^F, \emptyset)$, where Σ^F contains countably many function and constant symbols. Our setting is first order logic with equality and similar to previous discussions we assume bounded function arity. Given a set of Horn clauses Φ and a ground clause ψ , we denote by $\Phi[S_\psi]$ the set consisting of instances of Φ whose subterms are all in S_ψ . That is, $\Phi[S_\psi] = \{\varphi\sigma \mid \forall \bar{x} \varphi \in \Phi \text{ and } S_{\varphi\sigma} \subseteq S_\psi\}$. Call a theory \mathcal{T} , *local* (in FOL with equality) [19] if it is axiomatized by some finite set of (universal) Horn clauses Φ ,¹ such that for any ground Horn clause ψ , $\Phi \models \psi$ iff $\Phi[S_\psi] \models \psi$. When it is more convenient we will omit mentioning \mathcal{T} and simply call Φ the theory.

A partial Σ -algebra \mathcal{A} is a Σ -structure $(A, \{f^{\mathcal{A}}\}_{f \in \Sigma^F})$, where for each k -ary function symbol $f \in \Sigma^F$, $f^{\mathcal{A}}$ is a partial function from A^k to A . For a given variable assignment $\alpha : \mathcal{V} \rightarrow A$, the valuation of a term t with respect to an interpretation $\mathcal{I} = (\mathcal{A}, \alpha)$ is the same as before except for the case $t = f(t_1, \dots, t_k)$, the interpretation of the term, $t^{\mathcal{I}}$, is *undefined* if either one of the $t_i^{\mathcal{I}}$ is *undefined* or $(t_1^{\mathcal{I}}, \dots, t_k^{\mathcal{I}})$ is not in the domain of $f^{\mathcal{A}}$, otherwise, $t^{\mathcal{I}}$ is $f^{\mathcal{A}}(t_1^{\mathcal{I}}, \dots, t_k^{\mathcal{I}})$. A (total) *algebra* is partial algebra where all the functions are total. A quasi-variety, \mathcal{K} , is a class of (total) algebras that is axiomatized by a finite set of Horn clauses, i.e., $\mathcal{K} = \text{Mod } \Phi$ for some finite set of Horn clauses Φ .

Following the presentation of Ganzinger [19], we now give a definition for Evans' notion of *satisfaction* of a clause in a partial algebra \mathcal{A} . Say that a partial algebra \mathcal{A} *satisfies* an equational clause $\forall \bar{x} (s_1 \approx t_1 \wedge \dots \wedge s_k \approx t_k \rightarrow s \approx t)$ iff for every assignment α , whenever all s_i and t_i are defined and equal in \mathcal{A} , i.e., $s_i^{(\mathcal{A}, \alpha)} = t_i^{(\mathcal{A}, \alpha)}$, we have:

- (i) If both $s^{(\mathcal{A}, \alpha)}$ and $t^{(\mathcal{A}, \alpha)}$ are defined then they are equal.
- (ii) If $t^{(\mathcal{A}, \alpha)}$ is defined, and s is a term of form $f(u_1, \dots, u_n)$ and each $u_i^{(\mathcal{A}, \alpha)}$ is defined then $s^{(\mathcal{A}, \alpha)}$ is defined.

We say that \mathcal{A} *weakly* satisfies a given clause whenever only condition (i) above is satisfied. For a given quasi-variety $\mathcal{K} = \text{Mod } \Phi$, we say that \mathcal{A} is a *partial \mathcal{K} -algebra* if \mathcal{A} *satisfies* each clause ψ in Φ . Let $h : A \rightarrow B$ be a mapping between two partial algebras \mathcal{A} and \mathcal{B} , then h is a *weak homomorphism* if whenever (a_1, \dots, a_k) are in the domain of $f^{\mathcal{A}}$, we have $(h(a_1), \dots, h(a_k)) \in \text{dom } f^{\mathcal{B}}$, and $h(f^{\mathcal{A}}(a_1, \dots, a_k)) = f^{\mathcal{B}}(h(a_1), \dots, h(a_k))$. Say that a partial \mathcal{K} -algebra \mathcal{A} *weakly embeds* into \mathcal{K} , if there exists an injective weak homomorphism from \mathcal{A} to some total algebra \mathcal{B} in \mathcal{K} .

¹For ease of presentation we assume no constants occur in the axioms.

Then the following result due to Evans gives us a criteria for the polynomial time decidability of the uniform word problem.

Theorem 5.1.1. (Evans [16], Burris [6]) Let $\mathcal{K} = \text{Mod } \Phi$ be a quasi-variety, then if every finite partial \mathcal{K} -algebra weakly embeds into \mathcal{K} , then the uniform word problem for $\text{Cn } \Phi$ is polynomial time decidable.

The proof can be found in Evans [16] and Burris [6], we outline the core idea behind the polynomial time decidability, which is that of encoding clauses from the language generated from Σ into a *relational* language generated by an associated relational signature Σ^* . Consider the following. We associate with $\Sigma = (\Sigma^F, \emptyset)$ a relational signature $\Sigma^* = (\emptyset, \{P_f\}_{f \in \Sigma^F})$. Now, note that given a Horn clause, ψ we can flatten it so that the each atom that has a functional term is of the form $f(x_1, \dots, x_k) \approx x$. Then, we let ψ^* be the Σ^* -formula obtained from ψ by replacing each atom of the form $f(x_1, \dots, x_k) \approx x$ by $P_f(x_1, \dots, x_k, x)$. If we treat equality in Σ^* as just another predicate symbol (i.e., as in FOL without equality) we get datalog, which gives us the the polynomial time decidability we are after.

Further, we can express, in datalog, that the equality predicate is an equivalence relation that is consistent with the relations. We can also state that each P_f is a *partial function*, using $\forall \bar{x}(P_f(x_1, \dots, x_k, x_{k+1}) \wedge P_f(x_1, \dots, x_k, x_{k+2}) \rightarrow x_{k+1} \approx x_{k+2})$. Now, given a set of Horn clauses Φ , and a ground Horn clause ψ , we proceed as follows.

Relational Encoding. We let Φ^* and ψ^* be the Σ^* formulas obtained from Φ and ψ by the relational encoding method described above. Furthermore, we let Φ_{eq}^* , be the set of Σ^* formulas stating that \approx is an equivalence relation consistent with the relations and that each P_f is a partial function.

Decide Datalog entailment. The relational encoding above gives us a datalog program $\Phi^* \cup \Phi_{eq}^*$ and a datalog clause ψ^* . It is well known that deciding entailment in datalog i.e., $\Phi^* \cup \Phi_{eq}^* \models_{neq} \psi^*$, is polynomial time decidable [54].

Equivalence to original entailment. Observe that deciding the relational encoding is equivalent to deciding the original entailment. As, for the case $\Phi^* \cup \Phi_{eq}^* \models_{neq} \psi^*$, it can be seen that $\Phi \models \psi$, by encoding Σ -algebras into Σ^* relational models that will preserve satisfaction of clauses. Otherwise, for the case $\Phi^* \cup \Phi_{eq}^* \not\models_{neq} \psi^*$, there would be a finite relational model, \mathcal{A}^* of $\Phi^* \cup \Phi_{eq}^*$ such that $\mathcal{A}^* \not\models \psi^*$. We can then extract a finite partial \mathcal{K} -algebra \mathcal{A} from \mathcal{A}^* . Thus, by our assumption that every *partial* \mathcal{K} -algebra embeds into a total model \mathcal{B} of Φ , we would have a total model of Φ such that $\mathcal{B} \not\models \psi$. Hence, $\Phi \not\models \psi$.

5.2 Locality and Embeddability

In what follows, we show a relationship between locality and Evans' weak embeddability criterion. In particular, we show that locality is subsumed by Evan's embeddability criterion. Moreover, we show when we relax the embeddability criterion by replacing (strong) satisfaction with weak satisfaction, embeddability implies locality, and hence in this case locality coincides exactly with embeddability.

Call a (universal) Horn clause φ *flat*, if every subterm $t \in S_\varphi$ occurring in it, has depth at most 1 and no constants occur as a subterm of a functional term. That is every term is either a variable, a constant, or a term of the form $f(x_1, \dots, x_k)$. Say that a flat universal Horn clause is *linear* if whenever a variable occurs in two terms then the two terms are identical, and no variable occurs more than once in a term. Similarly, we say that a ground clause is *flat* if every subterm occurring in it has depth at most 1. We call a flat ground clause *linear* if whenever a constant occurs in two terms then the two terms are identical, and no constant occurs more than once in a term.

Let us denote by Φ_{flin} and ψ_{flin} the flattened and linearized versions of the universal set of Horn clauses Φ and the ground Horn clause ψ . The lemma [19] below allows us to simply consider flattened and linearized versions of theories and ground clauses. We refer to [19] for the proof.

Lemma 5.2.1. ([19]) *Let Φ be a set universal Horn clauses then,*

- (i) *If Φ is a local theory, then so is Φ_{flin} .*
- (ii) *If for every flat and linear ground clause ψ , we have $\Phi \models \psi$ implies $\Phi[S_\psi] \models \psi$, then Φ is local.*

5.2.1 Locality Implies Embeddability

We now show that locality is subsumed by Evans' embeddability criterion. Note that this subsumption is proper, that is, Evan's embeddability criterion does not imply locality. We refer the interested reader to Sections 5 and 6 of [19] for a counterexample theory satisfying Evans' criterion but which fails to be a local theory. However, Evans' criterion does imply a different type of locality which we will define and discuss after proving the following theorem.

Theorem 5.2.2. (Ganzinger [19]) Let Φ be a local set of flat Horn clauses, and $\mathcal{K} = \text{Mod } \Phi$. Then, we have that every partial algebra that weakly satisfies Φ weakly embeds into \mathcal{K} .

Proof. Let \mathcal{A} be a partial algebra weakly satisfying Φ . Assume for the sake of contradiction that \mathcal{A} does not weakly embed into \mathcal{K} . Without loss of generality we can assume that every a in the domain A , is some c^A for some constant c . Let $\Gamma_{\mathcal{A}} = \{ f(c_1, \dots, c_k) = c \mid (c_1^A, \dots, c_k^A) \in \text{dom } f^A, \text{ and } f^A(c_1^A, \dots, c_k^A) = c^A \}$, then, note that for any arbitrary model \mathcal{B} of $\Phi \cup \Gamma_{\mathcal{A}}$, we have the map $h : A \rightarrow B$, s.t. $h(c^A) = c^B$, is a weak homomorphism. Hence, as by assumption \mathcal{A} does not embed in to \mathcal{K} , we have that h is not injective. Hence, for some distinct constants c_1 and c_2 , we have $c_1^A \neq c_2^A$ but $c_1^B = c_2^B$. As \mathcal{B} was arbitrary, we have $\Phi \cup \Gamma_{\mathcal{A}} \models \bigvee_{c_1^A \neq c_2^A} c_1 \approx c_2$. Hence, by McKinsey's lemma, $\Phi \cup \Gamma_{\mathcal{A}} \models c_1 \approx c_2$ for some distinct constants having $c_1^A \neq c_2^A$. By the compactness theorem for FOL, we have for some finite subset $\Gamma_{\mathcal{A}}^0$ of $\Gamma_{\mathcal{A}}$, $\Phi \cup \Gamma_{\mathcal{A}}^0 \models c_1 \approx c_2$.

Now, let ψ be $\bigwedge \Gamma_{\mathcal{A}}^0 \rightarrow c_1 \approx c_2$, and note that Φ entails ψ . Assume for the sake of contradiction that Φ indeed locally entails ψ , that is $\Phi[S_{\psi}] \models \psi$. Then, note that every term occurring in $\Phi[S_{\psi}]$ and in ψ is defined in \mathcal{A} . Hence, as \mathcal{A} satisfies each clause in $\Phi[S_{\psi}]$, we have that \mathcal{A} weakly satisfies $\bigwedge \Gamma_{\mathcal{A}}^0 \rightarrow c_1 \approx c_2$. But note that, each equation in $\Gamma_{\mathcal{A}}^0$ is true in \mathcal{A} , hence we have $c_1^A = c_2^A$. Contradiction. Therefore, Φ is not local. \square

Let us denote by $\Phi^{[S_{\psi}]}$ the set consisting of instances of Φ , using substitutions that send variables in Φ to subterms of ψ , i.e., substitutions σ , s.t. for all x , $\sigma(x) \in S_{\psi}$. Call a theory \mathcal{T} *stably local*, if it is axiomatized by a finite set of Horn clauses Φ , such that for every ground Horn clause ψ , $\Phi \models \psi$ iff $\Phi^{[S_{\psi}]} \models \psi$. Note that *locality* implies stable locality. Then, the following theorem shows that Evans' embeddability criterion coincides with *stable locality*.

Theorem 5.2.3. (Ganzinger [19]) Let Φ be a finite set of Horn clauses, and $\mathcal{K} = \text{Mod } \Phi$ a quasi-variety, then if every finite \mathcal{K} -algebra weakly embeds into \mathcal{K} , Φ is stably local.

5.2.2 Embeddability (with weak satisfaction) Implies Locality

If we weaken Evans' embeddability criterion by discarding condition (ii) in the definition of Evans' *satisfaction*, i.e. use weak satisfaction, we see that locality coincides exactly with embeddability.

Theorem 5.2.4. (Ganzinger [19]) Assume Φ is a set of flat and linear set of Horn clauses, and let $\mathcal{K} = \text{Mod } \Phi$. Then, if every partial algebra that weakly satisfies Φ embeds into \mathcal{K} , Φ is local.

Proof. Let ψ be an arbitrary flat and linear ground Horn clause, we show $\Phi[S_\psi] \models \psi$ whenever $\Phi \models \psi$. By Lemma 5.2.1, this is sufficient. Assume that Φ entails ψ and let \mathfrak{A} be a model of $\Phi[S_\psi]$. We construct, from \mathfrak{A} , a partial algebra \mathcal{A} that weakly satisfies Φ such that (i) each ground term in $S_\Phi \cup S_\psi$ is defined, and (ii) for any clause ψ' over the subterms in $S_\Phi \cup S_\psi$, we have $\mathfrak{A} \models \psi'$ iff \mathcal{A} weakly satisfies ψ' . Now, by assumption we know \mathcal{A} weakly embeds into some total model \mathfrak{B} of Φ , let h be the embedding. Let ψ be $\bigwedge s_i \approx t_i \rightarrow s \approx t$, and assume for the sake of contradiction, that \mathcal{A} does not weakly satisfy ψ . Then, as the terms s_i, t_i, t and s are all defined in \mathcal{A} , we have $s_i^{\mathcal{A}} = t_i^{\mathcal{A}}$ but $s^{\mathcal{A}} \neq t^{\mathcal{A}}$. Hence, as h is a homomorphism from \mathcal{A} to \mathfrak{B} , we have $s_i^{\mathfrak{B}} = t_i^{\mathfrak{B}}$, moreover, as h is injective, we have $s^{\mathfrak{B}} \neq t^{\mathfrak{B}}$, and it follows $\mathfrak{B} \not\models \psi$, contradicting our assumption $\Phi \models \psi$. Thus, \mathcal{A} weakly satisfies ψ and hence $\mathfrak{A} \models \psi$.

What remains is to construct \mathcal{A} , we proceed as follows.

- (i) Set the domain $A = \{ t^{\mathfrak{A}} \mid t \text{ is a ground term in } S_\psi \cup S_\Phi \}$.
- (ii) For each k -ary function symbol f , $f^{\mathcal{A}}(a_1, \dots, a_k) = f^{\mathfrak{A}}(c_1^{\mathfrak{A}}, \dots, c_k^{\mathfrak{A}})$ if there exists constants c_i , such that $c_i^{\mathfrak{A}} = a_i$ and $f(c_1, \dots, c_k)$ is in $S_\Phi \cup S_\psi$. Otherwise, $f^{\mathcal{A}}(a_1, \dots, a_k)$ is undefined.

Now note each ground term $t \in S_\Phi \cup S_\psi$ is defined in \mathcal{A} and equals $t^{\mathfrak{A}}$, hence \mathcal{A} weakly satisfies ψ iff $\mathfrak{A} \models \psi$. We now show that \mathcal{A} weakly satisfies Φ . Let $\alpha : \mathcal{V} \rightarrow A$ be an arbitrary assignment, and let $\forall \bar{x} \varphi(\bar{x})$ be a clause in Φ , where φ is $(\bigwedge u_i \approx v_i \rightarrow u \approx v)$. Assume that each u_i, v_i, u , and v are defined in (\mathcal{A}, α) , and that $u_i^{(\mathcal{A}, \alpha)} = v_i^{(\mathcal{A}, \alpha)}$. Now as Φ is a set of flat clauses, each of these terms are either a variable, constant, or of the form $f(x_1, \dots, x_k)$. Define the substitution σ as follows: If x is some x_i in a functional term t of the form $f(x_1, \dots, x_k)$, we have as $t^{(\mathcal{A}, \alpha)}$ is defined each $\alpha(x_i) = c_i^{\mathcal{A}}$ for some³ c_i s.t. $f(c_1, \dots, c_k) \in S_\Phi \cup S_\psi$, hence we let $\sigma(x) = c_i$, otherwise if a variable x doesn't occur under a functional term then $\sigma(x) = r$ for some ground term r s.t. $\alpha(x) = r^{\mathcal{A}}$. This is well defined as Φ is flat and linear. Now note for any term $t \in S_\varphi$, $(t\sigma)^{(\mathcal{A}, \alpha)} = t^{(\mathcal{A}, \alpha)}$. Moreover, as $\varphi\sigma$ is in $\Phi[S_\psi]$, \mathcal{A} weakly satisfies $\varphi\sigma$ and hence (\mathcal{A}, α) weakly satisfies φ . As α was arbitrary we have \mathcal{A} weakly satisfies $\forall \bar{x} \varphi$. \square

³If there are multiple c_i just choose one.

5.3 Axiomatizability of Relational Substructures

The final concept related to locality and polynomial time decidability of the uniform word problem, is that of Burris' [6] certain axiomatizability criterion. We show in this section that whenever a quasi-variety \mathcal{K} satisfies Burris' criterion we have for some stably local theory, Φ , $\mathcal{K} = \text{Mod } \Phi$, i.e., \mathcal{K} is axiomatized by Φ . Conversely, if \mathcal{K} is axiomatized by a (stably) local theory Φ then \mathcal{K} satisfies Burris' criterion.

Recall that for an algebraic signature $\Sigma = (\Sigma^F, \emptyset)$, we associate with it a relational signature $\Sigma^* = (\emptyset, \{P_f\}_{f \in \Sigma^F})$.⁴ For a given Σ -clause ψ , we denote by ψ^* the Σ^* -clause obtained from ψ by replacing all equations of the form $f(x_1, \dots, x_k) \approx x$ by $P_f(x_1, \dots, x_k, x)$. Similarly, we associate with a Σ -algebra \mathcal{A} the Σ^* -relational structure \mathcal{A}^* with the same domain, such that $f^{\mathcal{A}}(a_1, \dots, a_k) = a_{k+1}$ iff $P_f^{\mathcal{A}^*}(a_1, \dots, a_k, a_{k+1})$. For a given quasi-variety \mathcal{K} we let \mathcal{K}^* be the set $\{\mathcal{A}^* \mid \mathcal{A} \in \mathcal{K}\}$.

Let \mathcal{K} be a quasi-variety, then we let $S(\mathcal{K}^*)$ be the set consisting of the substructures of the structures in \mathcal{K}^* . We denote by $\bar{S}(\mathcal{K}^*)$ the set consisting of the structures that weakly embed (i.e., have an injective homomorphism) into some structure in \mathcal{K}^* . Note that $S(\mathcal{K}^*) \subseteq \bar{S}(\mathcal{K}^*)$. We now state Burris' result leading to polynomial time decidability of the uniform word problem.

Theorem 5.3.1. (Burris [6]) Let Σ be a signature and \mathcal{K} a Σ -quasi-variety. If there exists a finite set of Horn clauses Φ such that $S(\mathcal{K}^*) \subseteq \text{Mod } \Phi \subseteq \bar{S}(\mathcal{K}^*)$, then the uniform word problem for \mathcal{K} is polynomial time decidable.

We can easily apply the above theorem to show polynomial time decidability of the uniform word problem for the class of all Σ -algebras \mathcal{K} , by noting that (i) $S(\mathcal{K}^*) = \bar{S}(\mathcal{K}^*)$, and (ii) $S(\mathcal{K}^*)$ is axiomatized by a set of Horn clauses stating that the P_f are partial functions. Hence, by the above theorem gives us polynomial time decidability for the uniform word problem. This polynomial time decidability result for the class of all Σ -algebras was first given by Kozen [33].

We now show that quasi-varieties that satisfy Burris' criterion have a stably local axiomatization. Given a set of Σ^* -clauses Φ^* , we will simply use the notation Φ to denote the set of Σ -clauses obtained from the corresponding Σ^* -clauses Φ^* . That is, by replacing each atom of the form $P_f(x_1, \dots, x_k, x)$ in the Σ^* -clauses by the equation $f(x_1, \dots, x_k) \approx x$. Note that for any partial Σ -algebra \mathcal{A} and its corresponding Σ^* -algebra \mathcal{A}^* , \mathcal{A} strongly satisfies Φ iff \mathcal{A}^* satisfies Φ^* .

⁴We assume, for this section, that Σ is a finite signature.

Theorem 5.3.2. (Ganzinger [19]) Let \mathcal{K} be a quasi-variety and Φ^* a set of Σ^* -clauses satisfying Buriss' criterion, i.e., $S(\mathcal{K}^*) \subseteq \text{Mod } \Phi^* \subseteq \bar{S}(\mathcal{K}^*)$. Then, Φ is a stably local axiomatization of \mathcal{K} .

Proof. We first show that Φ is an axiomatization of \mathcal{K} , after which we show that Φ is indeed stably local.

($\mathcal{K} \subseteq \text{Mod } \Phi$). This follows easily by noting that $\mathcal{A} \in \mathcal{K}$ implies $\mathcal{A}^* \in S(\mathcal{K}^*) \subseteq \text{Mod } \Phi^*$. Hence, $\mathcal{A} \models \Phi$.

($\text{Mod } \Phi \subseteq \mathcal{K}$). Assume that an algebra $\mathcal{A} \models \Phi$, then we have $\mathcal{A}^* \models \Phi^*$, and hence $\mathcal{A}^* \in \text{Mod } \Phi^*$. By our assumption we have $\text{Mod } \Phi^* \subseteq \bar{S}(\mathcal{K}^*)$, hence \mathcal{A}^* weakly embeds into some $\mathcal{B}^* \in \mathcal{K}^*$. It follows then \mathcal{A} weakly embeds into some \mathcal{K} -algebra \mathcal{B} , and hence is isomorphic to a substructure of \mathcal{B} . Therefore, $\mathcal{A} \in \mathcal{K}$, as the set of universal sentences satisfied by \mathcal{B} is a subset of that of \mathcal{A} .

Next to see that Φ is stably local note that for any finite partial \mathcal{K} -algebra \mathcal{A} , by assumption \mathcal{A}^* weakly embeds into some structure in \mathcal{K}^* . Hence, \mathcal{A} weakly embeds into \mathcal{K} . By Theorem 5.2.3, Φ is stably local. \square

Conversely, one can show that a quasi-variety \mathcal{K} axiomatized by a stably local set of Horn clauses, Φ , satisfies Buriss' criterion, i.e., there exists a finite set of Horn clauses $\Phi_{\mathcal{K}}$ such that $S(\mathcal{K}^*) \subseteq \text{Mod } \Phi_{\mathcal{K}} \subseteq \bar{S}(\mathcal{K}^*)$. In fact, one can show this assuming just locality.

Theorem 5.3.3. (Ganzinger [19]) Let $\mathcal{K} = \text{Mod } \Phi$, for a local set of Horn clauses Φ , and $\Phi_{\mathcal{K}}$ defined as union of Φ^* and the set of clauses stating that the relations are partial functions, i.e., $\forall \bar{x}(P_f(x_1, \dots, x_k, x) \wedge P_f(x_1, \dots, x_k, y) \rightarrow x \approx y)$. Then we have $S(\mathcal{K}^*) \subseteq \text{Mod } \Phi_{\mathcal{K}} \subseteq \bar{S}(\mathcal{K}^*)$.

5.4 Conclusion

The key focus of this chapter has been the semantic characterization of *locality* along with the polynomial time decidability of the uniform word problem. We explored three different characterizations, all leading to polynomial time decidability of the uniform word problem. Moreover, we showed the approaches by Evans and Burris — based on embeddability criterion of partial algebras, and axiomatizability of relational substructures respectively — lie in-between locality and the more general notion of *stable locality*.

Beyond the polynomial time decidability results for local theories, we have at large refrained from discussing complexity issues regarding the uniform word problem. We conclude this chapter by providing a list of some quasi-varieties for which the uniform word problem is known to require superpolynomial time [6]— assuming $P \neq NP$. The uniform word problem for the following quasi-varieties is co-NP-hard: any non-trivial finitely generated variety of lattices [28], any finitely generated non-nilpotent variety of rings [7], any congruence distributive variety generated by a two-element algebra [4]. Moreover, the uniform word problem for commutative semigroups is in exponential space.

Chapter 6

Locally Ground Theories

In this chapter, we define a subclass of local theories in which the search space for entailed ground definitions — i.e., entailments of the form $\varphi(\bar{x}) \models_{\mathcal{T}} x_i \approx s$, where s is an arbitrary ground term — is reduced to a finite space. Our setting in this chapter is FOL with equality. In this chapter, we give a definition for the *locally ground theories* based on the notion of being a *constructively ground term* with respect to a conjunction of literal φ and a theory \mathcal{T} .

We adopt here the notion of a *restricted derivation* from Chapter 4, on which the definition of local theories (in FOL without equality) [22] is based. As we are working with theories in the context of FOL with equality, however, we consider the equality axioms in our definition. We use Φ_{eq} to denote the set of equality axioms (i.e., reflexivity, symmetry, transitivity, functional and relational congruence).

Recall that, a local theory in FOL with equality [19] is a theory that is axiomatized by some finite set of (universal) Horn clauses Φ such that for any basic Horn clause ψ , $\Phi \models \psi$ iff $\Phi[S_\psi] \models \psi$, where $\Phi[S_\psi]$ is the set consisting of instances of Φ whose subterms are all in S_ψ .^{1,2} We say Φ is a local axiomatization of \mathcal{T} . To motivate the definition of restricted derivations in FOL with equality, we restate, in the terminology of this paper, a result by Ganzinger [19].

¹Note that, only ground formulas are considered in [19] and [22]. It is straightforward to adopt the results to formulas with free variables, as is done here, by observing for any quantifier free $\varphi(\bar{x})$, we have $\models \varphi(\bar{x})$ iff $\models \varphi[x_i/c_{x_i}]$, where c_{x_i} are fresh constants.

²For ease of presentation, we assume no ground terms occur in the axioms Φ .

Theorem 6.0.1. ([19]). The set of Horn clauses Φ is a local axiomatization of \mathcal{T} in FOL with equality iff $\Phi \cup \Phi_{eq}$ is a local axiomatization of \mathcal{T} in FOL without equality. That is, for any basic Horn clause ψ the following are equivalent:

- (i) $\Phi \models \psi$ iff $\Phi[S_\psi] \models \psi$.
- (ii) $\Phi \cup \Phi_{eq} \models_{neq} \psi$ iff $(\Phi \cup \Phi_{eq})[S_\psi] \models_{neq} \psi$.

where \models_{neq} denotes logical consequence in FOL without equality.

This along with the characterization of local theories in FOL without equality given by Givan and McAllester in [22] motivates the proof theoretic definition we adopt below. We note here that, the definition of locality we give below subsumes the original notion of *locality* given in [22, 19], and Chapter 4 and 5; as we are restricting the terms that appear in the derivations to a (polynomially) *extended subterm set*. Note that the uniform word problem is still polynomial time decidable.

In particular, let $\lambda : \mathcal{P}(\Sigma^T) \rightarrow \mathcal{P}(\Sigma^T)$ be a map associating with each set S of terms a set $\lambda(S)$ of terms. We further require λ be computable on finite inputs S . For a given φ , recall we use the notation S_φ^λ to denote the set $\lambda(S_\varphi)$ assigned to S_φ . Call such a function λ , *proper*, if for each formula φ , (i) $S_\varphi \subseteq S_\varphi^\lambda$, (ii) the size of S_φ^λ is polynomial in $|S_\varphi|$ and (iii) S_φ^λ is closed under the subterm relation.

Definition 6.0.2. (*Restricted Derivation*). Let $\Gamma \cup \{\theta\}$ be a set of atoms, and λ a proper function, then, a *restricted derivation* of θ from Γ using the axioms Φ , is θ if $\theta \in \Gamma$ and otherwise is a sequence of atomic formulas $\theta_1, \dots, \theta_n$ s.t. θ_n is θ and we have for each θ_i (i) $S_{\theta_i} \subseteq S_\varphi^\lambda$, where $\varphi = \bigwedge \Gamma \rightarrow \theta$, and (ii) there exists a Horn clause $\forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_k \rightarrow \psi) \in \Phi \cup \Phi_{eq}$ and a substitution σ such that θ_i is $\psi\sigma$ and each $\psi_j\sigma$ is either in Γ or appears earlier in the derivation.

Write $\Gamma \cup \Phi \Vdash^\lambda \theta$ if there exists a restricted derivation of θ from Γ using Φ and λ . Then, we say a theory \mathcal{T} is a *local theory* if it can be axiomatized by a finite set of Horn clauses Φ , and there is a proper function λ , such that $\Gamma \cup \Phi \vdash \theta$ iff $\Gamma \cup \Phi \Vdash^\lambda \theta$. Where \vdash is the provability relation for some proof system that is sound and complete for FOL with equality — i.e., $\Gamma \vdash \theta$ iff $\Gamma \models \theta$ — as is done in [22]. It is clear to see that, if we fix λ to be the identity on $\mathcal{P}(\Sigma^T)$ the original notion of locality is recovered. In what follows we will simply write \Vdash and omit λ when it is clear from context.

We now introduce the finite notion of being a *constructively ground term* with respect to some given conjunction of literals and a theory. Constructively ground terms are terms which might not be ground themselves, but can be transformed into an equivalent ground term in polynomial time using the (extended) subterm set of φ .

More formally, let φ be a conjunction of literals, \mathcal{T} a theory, and λ a proper function. We let $\sim_{\mathcal{T},\lambda}^{\varphi}$ be the equivalence relation defined on S_{φ}^{λ} , such that $t \sim_{\mathcal{T},\lambda}^{\varphi} s$ iff $\varphi \models_{\mathcal{T}} t \approx s$. Note that $\sim_{\mathcal{T},\lambda}^{\varphi}$ is a congruence relation on S_{φ}^{λ} . We denote by $[t]_{\mathcal{T},\lambda}^{\varphi}$ the equivalence class of t induced by $\sim_{\mathcal{T},\lambda}^{\varphi}$. We simply write \sim and $[-]$, when φ , \mathcal{T} , and λ are clear from the context.

We now give a recursive definition for the constructively ground terms. A term $t \in S_{\varphi}^{\lambda}$ is *constructively ground* (*c-ground*) if either **(i)** t is a ground term, or, **(ii)** t is $f(t_1, \dots, t_n)$ and for each t_i , $[t_i] = [r_i]$ for some c-ground term r_i . We say that the equivalence class of t , $[t]$, is *ground* if it contains a c-ground term.

Definition 6.0.3. (*Locally Ground Theory*). *Call a theory \mathcal{T} locally ground, if there exists a proper function λ for \mathcal{T} such that **(i)** \mathcal{T} is a local theory, and **(ii)** for any \mathcal{T} -satisfiable conjunction of literals φ and t a subterm of φ , whenever $\varphi \models_{\mathcal{T}} t \approx s$ for an arbitrary ground term s , we have $[t]_{\mathcal{T},\lambda}^{\varphi}$ is ground.*

The additional condition in the definition of locally ground theories allows us to limit our attention to the c-ground terms. The following lemma for theories axiomatized by Horn clauses, follows from the variant of McKinsey's lemma we proved in Section 2.1. It will help us ignore the negative literals when considering the logical consequences of a (satisfiable) conjunction of literals.

Lemma 6.0.4. *Let \mathcal{T} be a theory axiomatized by a set of universal Horn clauses Φ and $\varphi(\bar{x})$ be a \mathcal{T} -satisfiable conjunction of literals, $\bigwedge \varphi_i \wedge \bigwedge \neg\psi_i$, where each φ_i and ψ_i is an atom. Then, if $\varphi(\bar{x}) \models_{\mathcal{T}} \psi(\bar{x})$ for some atom $\psi(\bar{x})$, we have $\bigwedge \varphi_i \models_{\mathcal{T}} \psi(\bar{x})$.*

Proof. Assume $\varphi(\bar{x}) \models_{\mathcal{T}} \psi(\bar{x})$ then $\Phi \cup \{\varphi_i\} \models \bigvee \psi_i \vee \psi(\bar{x})$. Assume towards a contradiction that $\Phi \cup \{\varphi_i\} \not\models \psi(\bar{x})$. Now, by McKinsey's lemma (Lemma 2.1.2) it must be the case then, that $\Phi \cup \{\varphi_i\} \models \psi_i$ for some ψ_i . It follows then, $\bigwedge \Phi \wedge \varphi(\bar{x}) \models \psi_i \wedge \neg\psi_i$. Contradicting our assumption that $\varphi(\bar{x})$ is \mathcal{T} -satisfiable. \square

6.1 Partial Orders

We take as our working signature $\Sigma_{po} = (\Sigma^F, \{\leq\})$ where Σ^F contains countably many function and constant symbols. Now take as axioms for \mathcal{T}_{po} , the set Φ_{po} , containing the following Horn clauses.

$$\begin{array}{lll}
 \forall x & x \leq x & \text{(Reflexivity)} \\
 \forall x \forall y \forall z & x \leq y \wedge y \leq z \rightarrow x \leq z & \text{(Transitivity)} \\
 \forall x \forall y & x \leq y \wedge y \leq x \rightarrow x \approx y & \text{(Antisymmetry)}
 \end{array}$$

The proof given below for \mathcal{T}_{po} can be adapted to show other theories to be locally ground. Specifically, it can be easily adapted for the theory \mathcal{T}_{eq} axiomatized by the empty set of axioms (i.e., EUF) as well as for the theory of *Recursively Defined Data Structures* [42] as shown in Section 6.2. Moreover, \mathcal{T}_{po} remains locally ground if we extend it by adding axioms expressing monotonicity of some functions.

In this section, we take λ_p to be the identity on $\mathcal{P}(\Sigma_{po}^T)$. Moreover, we simply write \sim and $[-]$ in place of $\sim_{\mathcal{T}_{po}, \lambda_p}^\varphi$ and $[t]_{\mathcal{T}_{po}, \lambda_p}^\varphi$, when there is no ambiguity.

Lemma 6.1.1. *For a given \mathcal{T}_{po} -satisfiable conjunction of literals φ , there exists a model, $\mathcal{I}_p = (\mathfrak{A}_p, \alpha_p)$, of \mathcal{T}_{po} such that $\mathcal{I}_p \models_{\mathcal{T}_{po}} \varphi$, and satisfies the following properties:*

- (i) *The universe $A_p = (S_\varphi / \sim) \cup \{\zeta\}$, where $\zeta \notin (S_\varphi / \sim)$.*
- (ii) *For every term t in S_φ , $t^{\mathcal{I}_p} = [t]$.*
- (iii) *For any ground term s , $s^{\mathcal{I}_p} \neq \zeta$ implies $s^{\mathcal{I}_p} = [r]$, for some c-ground term $r \in S_\varphi$.*

We defer the proof of Lemma 6.1.1 and proceed to show that \mathcal{T}_{po} is a locally ground theory. After which we present a proof of the lemma.

Theorem 6.1.2. *The theory of partial orders, \mathcal{T}_{po} , is a locally ground theory.*

Proof. The locality of \mathcal{T}_{po} follows from Theorem 2 in [45], by noting that every weak partial model of its axioms, Φ_{po} , weakly embeds into a total model of Φ_{po} . What remains to show is that \mathcal{T}_{po} satisfies condition (ii) in the definition of locally ground theories. Towards showing this, assume φ is a \mathcal{T}_{po} -satisfiable conjunction of literals, t a subterm of φ , and s an arbitrary ground term. Note that s is not necessarily in S_φ . Now, assume $\varphi \models_{\mathcal{T}_{po}} t \approx s$ for some arbitrary ground term s . By Lemma 6.1.1, \mathcal{I}_p is a model of φ and \mathcal{T}_{po} , hence $t^{\mathcal{I}_p} = s^{\mathcal{I}_p}$. By (ii) in Lemma 6.1.1, $t^{\mathcal{I}_p} = [t] = s^{\mathcal{I}_p}$ and hence $s^{\mathcal{I}_p} \neq \zeta$. Thus, by (iii) in lemma 6.1.1, $s^{\mathcal{I}_p} = [r]$ for some c-ground term r , and hence $[t]$ is ground. \square

Proof. (of Lemma 6.1.1) construct the model $\mathcal{I}_p = (\mathfrak{A}_p, \alpha_p)$, as follows.

- (i) Set the domain A_p to be $(S_\varphi/\sim) \cup \{\zeta\}$, where $\zeta \notin S_\varphi/\sim$.
- (ii) For each n -ary function symbol f and $a_1, \dots, a_n \in A_p$,

$$f^{\mathfrak{A}_p}(a_1, \dots, a_n) = [f(s_1, \dots, s_n)] \text{ if } \exists \bar{s}, f(s_1, \dots, s_n) \in S_\varphi \text{ and } \forall i, a_i = [s_i];$$
 otherwise $f^{\mathfrak{A}_p}(a_1, \dots, a_n) = \zeta$.
- (iii) For the binary predicate symbol \leq , (a) $\zeta \leq^{\mathfrak{A}_p} \zeta$, and (b) for any $t_1, t_2 \in S_\varphi$, $[t_1] \leq^{\mathfrak{A}_p} [t_2]$ iff $\Gamma_\varphi \cup \Phi_{p_0} \Vdash t_1 \leq t_2$.³

Finally, set $\mathcal{I}_p = (\mathfrak{A}_p, \alpha_p)$, where the variable assignment $\alpha_p : \mathcal{V} \rightarrow A_p$ is defined as $\alpha_p(x) = [x]$ if $x \in S_\varphi$, and otherwise $\alpha_p(x) = \zeta$.

We note here, as \sim is a congruence relation, that each n -ary function $f^{\mathfrak{A}_p}$ and the binary relation $\leq^{\mathfrak{A}_p}$ are well-defined. We proceed to show that \mathcal{I}_p satisfies conditions (ii) and (iii) outlined in Lemma 6.1.1. We can use induction on $t \in S_\varphi$ to see that for each subterm t of φ , $t^{\mathcal{I}_p} = [t]$, hence \mathcal{I}_p satisfies condition (ii).

To see that \mathcal{I}_p satisfies condition (iii), we proceed by induction on t . If $t = c$ for some constant symbol c , then either c is in S_φ and so by definition $c^{\mathcal{I}_p} = [c]$ or $c \notin S_\varphi$ and $c^{\mathcal{I}_p} = \zeta$. Otherwise, let t be the term $f(t_1, \dots, t_n)$ and assume $t^{\mathcal{I}_p} \neq \zeta$. Then, by definition, we have each $t_i^{\mathcal{I}_p} = [u_i]$ for some u_i in S_φ . Now, note that as each t_i is ground and each $t_i^{\mathcal{I}_p}$ is different from ζ , we have by the inductive hypothesis $[u_i] = [r_i]$ for some constructively ground term r_i . Moreover, by definition $t^{\mathcal{I}_p} = f^{\mathfrak{A}_p}([u_1], \dots, [u_n]) = [f(s_1, \dots, s_n)]$ for some s_i in $[u_i]$. Hence, $[s_i]$, which is equal to $[u_i]$, contains the c-ground term r_i . Therefore, the term $f(s_1, \dots, s_n)$ is a constructively ground term and the conclusion follows.

Finally, we show that \mathcal{I}_p is a model of both \mathcal{T}_{p_0} and φ . Note that, for any t_1, t_2 in S_φ , we have by Lemma 6.0.4 and the fact that Φ_{p_0} is a local axiomatization of \mathcal{T}_{p_0} , $[t_1] \leq^{\mathfrak{A}_p} [t_2]$ iff $\varphi \Vdash_{\mathcal{T}_{p_0}} t_1 \leq t_2$. Hence, $\leq^{\mathfrak{A}_p}$ is reflexive, transitive, and antisymmetric on A_p . Therefore, $\mathcal{I}_p \models \Phi_{p_0}$, and thus is a model of \mathcal{T}_{p_0} . Next, to see that \mathcal{I}_p satisfies φ , recall that φ is a \mathcal{T} -satisfiable conjunction of literals, φ_i , of the form $t_i \approx u_i$ or $r_i \leq s_i$ or their negations. Then, observe that $\mathcal{I}_p \models \varphi_i$, by noting $\varphi \Vdash_{\mathcal{T}_{p_0}} \varphi_i$, and that for each $t \in S_\varphi$, $t^{\mathcal{I}_p} = [t]$. Therefore, it follows $\mathcal{I}_p \models_{\mathcal{T}_{p_0}} \varphi$. \square

The construction is motivated by Shostak's decision procedure [44] for the quantifier free fragment of the theory of equality. However, unlike [44] we don't use the entire Herbrand universe, and simply interpret/map terms to their equivalence classes under \sim .

³Equivalently, $[t_1] \leq^{\mathfrak{A}_p} [t_2]$ iff $\varphi \Vdash_{\mathcal{T}_{p_0}} t_1 \leq t_2$.

6.2 Recursively Defined Data Structures

We now show that the theory of Recursively Defined Data Structures, \mathcal{T}_{rd} , [42] without the acyclicity axioms is locally ground. We take as our working signature $\Sigma_{rd} = (\Sigma^F \cup \{cr, sr_1, \dots, sr_k\}, \emptyset)$ where Σ^F contains countably many function and constant symbols, and cr and sr_i are intended to denote the k -ary constructor and unary selector functions respectively. Take as axioms for \mathcal{T}_{rd} , the set Φ_{rd} , containing the following Horn clauses:

$$\begin{aligned} \forall x \quad cr(sr_1(x), \dots, sr_k(x)) &\approx x && \text{(Construction)} \\ \forall \bar{x} \quad sr_i(cr(x_1, \dots, x_k)) &\approx x_i \quad \text{for } i = 1, \dots, k && \text{(Selection)} \end{aligned}$$

The approach we take in showing that \mathcal{T}_{rd} is a locally ground theory is similar to the approach we took in Section 6.1. Following Oppen's method in [42], we construct an infinite model of \mathcal{T}_{rd} . Let $\lambda_r : \mathcal{P}(\Sigma_{rd}^T) \rightarrow \mathcal{P}(\Sigma_{rd}^T)$ be the function such that for each $S \in \mathcal{P}(\Sigma_{rd}^T)$,

$$\lambda_r(S) = S \cup \bigcup_{t \in S} \{sr_i(t), cr(sr_1(t), \dots, sr_k(t))\}$$

Note that λ_r is a *proper* function. For the rest of the section, we simply write \sim and $[-]$ in place of $\sim_{\mathcal{T}_{rd}, \lambda_r}^\varphi$ and $[t]_{\mathcal{T}_{rd}, \lambda_r}^\varphi$, when there is no ambiguity. Recall that we use $S_\varphi^{\lambda_r}$ as a shorthand for $\lambda_r(S_\varphi)$.

Lemma 6.2.1. *For a given \mathcal{T}_{rd} -satisfiable conjunction of literals φ , there exists a model, $\mathcal{I}_r = (\mathfrak{A}_r, \alpha_r)$, of \mathcal{T}_{rd} such that $\mathcal{I}_r \models_{\mathcal{T}_{rd}} \varphi$, and satisfies the following properties:*

- (i) *The set $S_\varphi^{\lambda_r} / \sim$ is a subset of the universe A_r .*
- (ii) *For every term t in S_φ , $t^{\mathcal{I}_r} = [t]$.*
- (iii) *For any ground term s , $s^{\mathcal{I}_r} \in Z$. Where Z is a set such that each element of Z that is an equivalence class is ground — i.e., for all $[t] \in Z \cap S_\varphi^{\lambda_r} / \sim$, $[t]$ is ground.*

We defer the proof of Lemma 6.2.1 and proceed to show that \mathcal{T}_{rd} is a locally ground theory. After which we present a proof of the lemma.

Theorem 6.2.2. The theory of recursively defined data structures, \mathcal{T}_{rd} , is a locally ground theory.

Proof. It can be seen for any basic Horn clause ψ that $\Phi_{rd} \models \psi$ iff $\Phi_{rd}[S_\psi^{\lambda_r}] \models \psi$. Hence, \mathcal{T}_{rd} satisfies condition (i). What remains to show is that \mathcal{T}_{rd} satisfies condition (ii) in the definition of locally ground theories. Towards showing this, assume φ is a \mathcal{T}_{rd} -satisfiable conjunction of literals, $t \in S_\varphi$ a subterm of φ , and s an arbitrary ground term. Now, assume that $\varphi \models_{\mathcal{T}_{rd}} t \approx s$. Then, by Lemma 6.2.1, we have $\mathcal{I}_r \models_{\mathcal{T}_{rd}} t \approx s$, hence $t^{\mathcal{I}_r} = s^{\mathcal{I}_r}$. By (ii) in Lemma 6.2.1, $t^{\mathcal{I}_r} = [t] = s^{\mathcal{I}_r}$, moreover, by (iii) in Lemma 6.2.1, $s^{\mathcal{I}_r} \in Z$. Hence, by definition of Z , $[t]$ is ground. \square

Proof. (of Lemma 6.2.1). Construct the model, $\mathcal{I}_r = (\mathfrak{A}_r, \alpha_r)$, as follows. First define for each $n \geq 0$, the sets A_n , the partial functions cr_n and $sr_{i,n}$ for each $i \in \{1, \dots, k\}$.

(i) Let $A_0 = (S_\varphi^{\lambda_r} / \sim) \cup \{\zeta\}$, where $\zeta \notin S_\varphi^{\lambda_r} / \sim$.

(ii) Define the partial function cr_0 on A_0 as follows:

$$cr_0(a_1, \dots, a_k) = [cr(t_1, \dots, t_k)] \text{ if } \exists \bar{t}, cr(t_1, \dots, t_k) \in S_\varphi^{\lambda_r} \text{ s.t. } \forall i, a_i = [t_i];$$

otherwise $cr_0(a_1, \dots, a_k)$ is undefined.

(iii) For $i = 1, \dots, k$, define the partial functions $sr_{i,0}$ on A_0 as follows:

$$sr_{i,0}(a) = [sr_i(t)] \text{ if } \exists t, sr_i(t) \in S_\varphi^{\lambda_r} \text{ and } a = [t]; \text{ otherwise } sr_{i,0}(a) \text{ is undefined.}$$

We now keep extending the domains of cr_0 and $sr_{i,0}$. For each $n > 0$,

(i) Let $S_{n-1} = \{\zeta_1^a, \dots, \zeta_k^a \mid a \in (A_{n-1} \setminus \text{dom } sr_{i,n-1}) \text{ and } \zeta_i^a \notin (A_{n-1} \cup A_{n-1}^k)\}$

(ii) Let $A_n = A_{n-1} \cup (A_{n-1}^k \setminus \text{dom } cr_{n-1}) \cup S_{n-1}$.

(iii) Define cr_n to be a partial function on A_n that is an extension of cr_{n-1} , such that for each $\bar{b} = (b_1, \dots, b_k) \in A_n^k \setminus \text{dom } cr_{n-1}$:

$$cr_n(\bar{b}) = \bar{b} \text{ if } \bar{b} \in A_{n-1}^k \setminus \text{dom } cr_{n-1}; \text{ otherwise } cr_n(\bar{b}) = a \text{ if } \bar{b} = (\zeta_1^a, \dots, \zeta_k^a);$$

and finally, $cr_n(\bar{b})$ is undefined for every other case.

(iv) Similarly, for $i = 1, \dots, k$, define the partial function $sr_{i,n}$ on A_n , that is an extension of $sr_{i,n-1}$, such that, for all $a \in A_n \setminus \text{dom } sr_{i,n-1}$:

$$sr_{i,n}(a) = \zeta_i^a \text{ if } a \in A_{n-1} \setminus \text{dom } sr_{i,n-1}; \text{ otherwise } sr_{i,n}(a) = a_i \text{ if}$$

$$a = (a_1, \dots, a_k) \in A_{n-1}^k \setminus \text{dom } cr_{n-1}; \text{ and finally } sr_{i,n}(a) \text{ is undefined}$$

for every other case.

Note that both cr_n and $sr_{i,n}$ are total on A_{n-1} . We now construct the interpretation $\mathcal{I}_r = (\mathfrak{A}_r, \alpha_r)$ as follows.

- (i) Set the universe $A_r = \bigcup_{n < \omega} A_n$.
- (ii) For the function symbol, cr , and the k -tuple $\bar{a} = (a_1, \dots, a_k)$ in $(A_r)^k$,
 $cr^{\mathfrak{A}_r}(a_1, \dots, a_k) = cr_n(a_1, \dots, a_k)$, for the least index n , s.t. $\bar{a} \in \text{dom } cr_n$.
- (iii) Similarly, for each selector function symbol sr_i , and $a \in A_r$,
 $sr_i^{\mathfrak{A}_r}(a) = sr_{i,n}(a)$, for the least index n , such that $a \in \text{dom } sr_n$.
- (iv) For each k -ary function symbol $f \in \Sigma^F$ and a_1, \dots, a_k in A_r ,
 $f^{\mathfrak{A}_r}(a_1, \dots, a_k) = [f(s_1, \dots, s_k)]$ if $\exists \bar{s}, f(s_1, \dots, s_k) \in S_\varphi$ and $\forall i, a_i = [s_i]$;
and otherwise $f^{\mathfrak{A}_r}(a_1, \dots, a_k) = \zeta$.

Finally, define the variable assignment $\alpha_r : \mathcal{V} \rightarrow A_r$, as follows $\alpha_r(x) = [x]$ if $x \in S_\varphi^{\lambda_r}$ and $\alpha_r(x) = \zeta$ otherwise. We note that all the functions above are total on A_r . Moreover, as \sim is a congruence relation, the functions are well-defined. Now, to see that $\mathcal{I}_r \models_{\mathcal{T}_{rd}} \varphi$, we first observe, using induction similar to the proof given in [42] the following:

1. If $a \in \text{dom } sr_{i,n}$ then $\bar{b} = (sr_{1,n}(a), \dots, sr_{k,n}(a)) \in \text{dom } cr_n$, and $cr_n(\bar{b}) = a$.
2. If $\bar{a} = (a_1, \dots, a_k) \in \text{dom } cr_n$ then $sr_{i,n}(cr_n(a_1, \dots, a_k)) = a_i$.

Hence, the properties 1 and 2 above hold for the functions $cr^{\mathfrak{A}}$ and $sr_i^{\mathfrak{A}}$ and $\mathcal{I}_r \models \mathcal{T}_{rd}$. Furthermore, we can see by induction on each $t \in S_\varphi$ that $t^{\mathcal{I}_r} = [t]$, and hence for each positive literal $t_1 \approx t_2$ occurring in φ , $\mathcal{I}_r \models t_1 \approx t_2$. Additionally, as φ is \mathcal{T}_{rd} -satisfiable, we have for each negative literal $t_1 \not\approx t_2$ in φ , $[t_1] \neq [t_2]$ and hence $\mathcal{I}_r \models t_1 \not\approx t_2$. Thus, $\mathcal{I}_r \models_{\mathcal{T}_{rd}} \varphi$.

We now proceed to show that \mathcal{I}_r satisfies the conditions (i) – (iii) given in Lemma 6.2.1. Note $A_0 \subseteq A_r$, hence, \mathcal{I}_r satisfies condition (i). It is also straightforward to see by induction on $t \in S_\varphi$, that $t^{\mathcal{I}_r} = [t]$, hence \mathcal{I}_r satisfies condition (ii). To see \mathcal{I}_r satisfies condition (iii), let $Z = \bigcup_{n < \omega} Z_n$ where Z_n is defined as:

$$Z_0 = \{ \zeta \} \cup \{ [t] \in S_\varphi^{\lambda_r} / \sim \mid [t] \text{ is ground} \} \text{ and for } n > 0,$$

$$Z_n = Z_{n-1} \cup (Z_{n-1}^k \setminus \text{dom } cr_{n-1}) \cup \bigcup_{\zeta_i^a \in S_{n-1}} \{ \zeta_i^a \mid a \in Z_{n-1} \}$$

Note that for every $[t] \in Z \cap S_\varphi^{\lambda_r} / \sim$, $[t]$ is ground. Now, let s be a ground term, we show by induction on s that $s^{\mathcal{I}_r} \in Z$. It suffices to show $s^{\mathcal{I}_r} \in Z_n$ for some n . For the base case, if $s = c$ for some constant c , then $s^{\mathcal{I}_r}$ is $[c]$ if $c \in S_\varphi$ and ζ otherwise, in both cases $s^{\mathcal{I}_r} \in Z_0$. Otherwise, if $s = f(t_1, \dots, t_n)$ for some $f \in \Sigma^F$, the argument is the same as in Lemma 6.1.1. For the case $s = sr_i(t_1)$, let $a = t_1^{\mathcal{I}_r}$, then, we have by the inductive hypothesis, $a \in Z$. Let n be the least index such that $a \in \text{dom } sr_{i,n}$, we proceed by induction on n . First, note that for any $n \geq 0$, we have $Z \cap A_n \subseteq Z_n$. Hence, for the base case, $n = 0$, we have $a = [t]$ for some ground class $[t]$, and $s^{\mathcal{I}_r} = sr_{i,0}([t]) = [sr_i(t)]$. Hence, $s^{\mathcal{I}_r} \in Z_0$ as $sr_i(t)$ is c-ground. For $n > 0$, we have either $a \in A_{n-1} \setminus \text{dom } sr_{i,n-1}$ in which case $s^{\mathcal{I}_r} = sr_{i,n}(a) = \zeta_i^a$ and hence $s^{\mathcal{I}_r} \in Z_n$, or $a = (a_1, \dots, a_k) \in A_{n-1}^k \setminus \text{dom } cr_{n-1}$ and $sr_{i,n}(a) = a_i \in Z_{n-1}$.

Similarly, for the case $s = cr(t_1, \dots, t_k)$ we have by the induction hypothesis, $t_i^{\mathcal{I}_r} \in Z_{n_i}$. Let n be the least index such that $\bar{b} = (t_1^{\mathcal{I}_r}, \dots, t_k^{\mathcal{I}_r}) \in \text{dom } cr_n$, we proceed by induction on n . For $n = 0$, we can use a similar argument to the case $s = sr_i(t_1)$. For the inductive case, we have either $\bar{b} \in A_{n-1}^k \setminus \text{dom } cr_{n-1}$, in which case $s^{\mathcal{I}_r} = cr_n(\bar{b}) = \bar{b}$, and hence $s^{\mathcal{I}_r} \in Z_n$, or $\bar{b} = (\zeta_1^a, \dots, \zeta_k^a)$, and $s^{\mathcal{I}_r} = cr_n(\bar{b}) = a \in Z_{n-1}$. \square

Chapter 7

Partial Quantifier Elimination

In this chapter, we consider the *partial quantifier elimination* problem and give an efficient algorithm, \mathcal{T} -QEL, that is relatively complete for the locally ground theories. In particular, we consider the following problem: given a theory \mathcal{T} and a formula $\exists \bar{x} \varphi(\bar{x})$, where $\varphi(\bar{x})$ is a conjunction of literals, over the signature Σ , find a quantifier free Σ formula $\psi(\bar{y})$ such that: **(i)** $\models_{\mathcal{T}} \exists \bar{x} \varphi(\bar{x}) \leftrightarrow \exists \bar{y} \psi(\bar{y})$ **(ii)** $\text{FV}(\psi) \subseteq \text{FV}(\varphi)$ and **(iii)** for all $x \in \text{FV}(\varphi)$, if x has an entailed ground definition (i.e., $\varphi \models_{\mathcal{T}} x \approx s$ with s an arbitrary ground term), then $x \notin \text{FV}(\psi)$.

In essence, we have reduced the number of quantified variables that have entailed ground definitions from φ . We call ψ a *quantifier reduction* of $\exists \bar{x} \varphi(\bar{x})$ if it satisfies (i) and (ii). We say a procedure is *relatively complete* if for every input, it computes a quantifier free formula satisfying (iii). For the rest of the section, we fix a locally ground theory $\mathcal{T} = \text{Cn } \Phi$. Moreover, for simplicity of presentation, we take λ to be the identity on $\mathcal{P}(\Sigma^T)$.

Egraphs give us a compact way of representing congruence relations. In what follows, we describe a construction of a certain egraph, G^* , which we call *the completion* of the egraph of φ . The goal of the construction is to represent the partition S_{φ}/\sim induced by the congruence relation \sim on top of an egraph. For locally ground theories, this allows us to efficiently search the equivalence classes in S_{φ}/\sim for the constructively ground terms. Finally, we utilize QEL to extract an equivalent formula from the egraph G^* .

Intuitively, our construction saturates the egraph of φ with implied atoms. That is, we pick an instance of an axiom $\psi_1 \wedge \dots \wedge \psi_k \rightarrow \theta$, for which all ψ_i are represented (i.e., exist) in the egraph, after which we add θ to the egraph. We stop when no new atoms are added. Locality here allows us to restrict to instances of axioms whose subterms are already in φ .

We formalize this notion in the rest of the section, and use it to give proofs of soundness and relative completeness for our algorithm \mathcal{T} -QEL.

7.1 \mathcal{T} -QEL

We now present the details of the algorithm \mathcal{T} -QEL. In what follows, we refer to intermediate graphs arising from the construction of G^* , that might not satisfy the congruence condition (c) in the definition of egraphs, as *partial egraphs*. For simplicity of presentation, we follow [20] and allow labelling nodes with predicate symbols and introduce the fresh constant symbols \top and \perp .

Definition 7.1.1. (*Representability*). *Given a (partial) egraph, $G = (N, E, L, \text{root})$, say that a literal γ over the subterms of φ is representable in G if either (i) γ is $t_1 \approx t_2$, and $\text{root}(n_{t_1}) = \text{root}(n_{t_2})$, or (ii) γ is $t_1 \not\approx t_2$, and for some node w in N labelled with $\not\approx$, we have $w[i] = n_{t_i}$, or (iii) γ is $P(t_1, \dots, t_k)$ or $\neg P(t_1, \dots, t_k)$, and for some node w in N labelled with P , we have $w[i] = n_{t_i}$, and $\text{root}(w) = n_\top$ if γ is positive and $\text{root}(w) = n_\perp$ otherwise.*

Let $G_0 = \text{egraph}(\varphi)$, and for $k > 0$, form G_k from G_{k-1} by instantiating the axioms on the nodes of G_{k-1} . More formally, let $\psi \in \Phi \cup \Phi_{eq}$ be an axiom of the form $\forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_m \rightarrow P(t_1, \dots, t_n))$ where each ψ_i is an atom and P is different from the equality symbol, then let

$$\Pi_k^\psi = \{ \bar{u} \in N_{k-1}^n \mid \text{term}(u_j) = t_j\sigma \text{ and each } \psi_i\sigma \text{ is representable in } G_{k-1} \text{ for some } \sigma \}.$$

Then, let $\Pi_k = \bigcup \{ (\bar{u}, P) \mid \bar{u} \in \Pi_k^\psi \text{ for some } \psi \text{ of the above form} \}$ and form $G_k = (N_k, E_k, L_k, \text{root}_k)$ as follows:

- (a) $N_k = N_{k-1} \cup \{ w_{(\bar{u}, P)} \mid (\bar{u}, P) \in \Pi_k \}$
- (b) $L_k = L_{k-1} \cup \{ (w_{(\bar{u}, P)}, P) \mid w_{(\bar{u}, P)} \in N_k \setminus N_{k-1} \}$
- (c) $E_k = E_{k-1} \cup \{ (w_{(\bar{u}, P)}, u_i) \mid w_{(\bar{u}, P)} \in N_k \setminus N_{k-1} \text{ and } u_i \in \bar{u} \}$ and set the order of the children as $w_{\bar{u}}[i] = u_i$.

To complete the construction of G_k we now consider an equality that arises as a result of adding new nodes/atoms. Equality is handled by merging the equivalence classes of any

Algorithm 3 \mathcal{T} -QEL : An extension of QEL that utilizes theories.

Input: $\exists \bar{x} \varphi(\bar{x})$, where $\varphi(\bar{x})$ is a conjunction of literals; Φ a set of axioms.

Output: a quantifier reduction of $\exists \bar{x} \varphi(\bar{x})$ in the theory $\mathcal{T} := Cn \Phi$.

- 1: $G \leftarrow \text{egraph}(\varphi)$
 - 2: $G \rightsquigarrow_{\Phi} G^*$
 - 3: $\beta \leftarrow G^*.find_defs()$
 - 4: $core \leftarrow G^*.find_core(\beta)$
 - 5: **return** $G^*.to_formula(\beta, N^* \setminus core)$
-

two nodes that are inferred to be equal at this stage. We define Δ_k^ψ in a similar way we defined Π_k^ψ . Let $\psi \in \Phi \cup \Phi_{eq}$ be an axiom of the form $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_m \rightarrow t_1 \approx t_2)$, where each ψ_i is an atom. Then let

$$\Delta_k^\psi = \{ \bar{u} \in N_{k-1}^2 \mid \text{term}(u_j) = t_j \sigma \text{ and each } \psi_i \sigma \text{ is representable in } G_{k-1} \text{ for some } \sigma \}$$

Let $\Delta_k = \bigcup \Delta_k^\psi$ and let $(u_1, v_1), \dots, (u_r, v_r) \in \Delta_k$ be an enumeration of the elements of Δ_k . Now we iteratively merge the equivalence classes of each u_i with that of v_i , by setting the root of each u in the class of u_i to that of v_i . Let $root_k^0 = root_{k-1}$ and for $1 \leq i \leq r$ define $root_k^i : N_k \rightarrow N_k$, for all $u \in N_k$ as:

$$\begin{aligned} root_k^i(u) &= n_{\top} \text{ if } u \in N_k \setminus N_{k-1}, \text{ otherwise } root_k^i(u) = root_k^{i-1}(v_i) \text{ if} \\ &root_k^{i-1}(u) = root_k^{i-1}(u_i), \text{ and otherwise } root_k^i(u) = root_k^{i-1}(u). \end{aligned}$$

Finally, define $root_k = root_k^r$. This concludes the construction of G_k .

For a given (partial) egraph G , let $\mathcal{X}_G = \{ \gamma \mid \gamma \text{ is an atom representable in } G \}$. Then, let ℓ be the least index such that $\mathcal{X}_{G_\ell} = \mathcal{X}_{G_{\ell+1}}$, and let $G^* = G_\ell$. We use the notation $G \rightsquigarrow_{\Phi} G^*$ to denote the construction above. Below, we show that the maximum number of steps for the construction, is polynomial in $|S_\varphi|$. Hence, giving us an overall polynomial time algorithm.

Lemma 7.1.2. *The construction $G \rightsquigarrow_{\Phi} G^*$ takes at most $\mathcal{O}(|S_\varphi|^m)$ steps, where m is the maximum arity of the predicates occurring in the axioms and in φ .*

Proof. Note that there are at most $|\Phi| |S_\varphi|^m$ distinct number of atoms $P(t_1, \dots, t_n)$ formed over the subterms of φ where P occurs in the head of some clause. The construction above, at each step, only adds new nodes labelled by predicate symbols having as children the original nodes u_t for $t \in S_\varphi$. Hence, $|\mathcal{X}_G| \leq (|\Phi| + n) |S_\varphi|^m$ for all (partial) egraphs in the construction, where n is the number of distinct predicate symbols occurring in φ . Hence, in at most $k \leq (|\Phi| + n) |S_\varphi|^m$ steps we will have $\mathcal{X}_{G_k} = \mathcal{X}_{G_{k+1}}$ and the conclusion follows. \square

Our algorithm \mathcal{T} -QEL is presented in Algorithm 3. \mathcal{T} -QEL takes as input the axioms Φ of the theory and a formula $\exists \bar{x} \varphi(\bar{x})$. The full details of all the steps (except step 2) are discussed in [20]. Regardless, we will go ahead and provide a brief explanation of each of the steps.

In step 1, the egraph of φ is constructed using the standard procedure. Then, in step 2, we represent the partition S_φ/\sim of the subterms induced by \sim on the egraph. We note here that once S_φ/\sim has been represented, the newly added predicate symbols don't serve any purpose and can be removed from the egraph. In step 3, a function β that picks a representative node (term) from each equivalence class is computed. Step 4 identifies a subset of nodes (terms) that must be considered in the output. Finally, step 5 extracts a formula from the egraph using β and *core*.

We now proceed to show that the partition S_φ/\sim is represented on the egraph G^* . First, we prove the following Lemma.

Lemma 7.1.3. (*Representability lemma*). *Let \mathcal{T} be a locally ground theory axiomatized by Φ , φ a \mathcal{T} -satisfiable conjunction of literals, G^* the completion of the egraph(φ), and γ an atom such that $S_\gamma \subseteq S_\varphi$, then $\varphi \models_{\mathcal{T}} \gamma$ iff γ is representable in G^* .*

Proof. (\implies) To show representability in G^* we show that γ is representable in some G_k . The conclusion follows as for any $k \geq 0$, $\mathcal{X}_{G_k} \subseteq \mathcal{X}_{G_\ell}$. Now assume $\varphi \models_{\mathcal{T}} \gamma$, then by Lemma 6.0.4, we have $\bigwedge \varphi_i \models_{\mathcal{T}} \gamma$ hence for some proof calculus that is sound and complete for FOL(\approx) we have $\Gamma_\varphi \cup \Phi \vdash \gamma$. By definition of locality we have a restricted derivation, $\theta_1, \dots, \theta_d$, of γ . If γ is in Γ_φ then the conclusion follows trivially. For the case it does not, we proceed by induction on the length of the derivation. For $d = 1$, we have for some clause $\psi \in \Phi \cup \Phi_{eq}$ of the form $\forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_n \rightarrow \theta)$ and substitution σ , $\theta\sigma$ is γ and each $\psi_i\sigma$ is in Γ_φ . Hence, each $\psi_i\sigma$ is representable in G_0 by the way the *egraph*(φ) is constructed. Now, if γ is $P(t_1, \dots, t_n)$, different from an equality atom, then there exists nodes $\bar{u} \in \Pi_1^\psi$ and $w_{(\bar{u}, P)} \in N_1$ such that $w_{\bar{u}}$ is labelled with P and $w_{(\bar{u}, P)}[i] = n_{t_i}$. Hence, $P(t_1, \dots, t_n)$ is representable in G_1 . Otherwise, γ is an equality atom and there exists $(u_1, u_2) \in \Delta_1^\psi$ with $term(u_i) = t_i$. Let (u_1, u_2) be the j th element $(u_j, v_j) \in \Delta_k$ in the enumeration given during the construction of G_1 . Then we have merged the equivalence classes of u_j and v_j in G_1 . More precisely, we have $root_1^j(u_j) = root_1^j(v_j)$ and it can be shown by induction that for all $m \geq j$, $root_1^m(u_j) = root_1^m(v_j)$, hence $root_1(u_j) = root_1(v_j)$ and therefore $t_1 \approx t_2$ is representable in G_1 . In both cases we have shown γ is representable in G_1 . For the inductive step we have for some clause $\forall \bar{x} (\psi_1 \wedge \dots \wedge \psi_n \rightarrow \psi)$ and substitution σ , $\psi\sigma$ is γ , and either $\psi_i\sigma \in \Gamma_\varphi$ in which case $\psi_i\sigma$ is representable in G_0 or $\psi_i\sigma$ occurs earlier in the derivation. In the later case, we have by the inductive hypothesis $\psi_i\sigma$ is representable in

some G_{k_i} . Let $k = \max\{k_i\}$, then, as $\mathcal{X}_{G_{k_i}} \subseteq \mathcal{X}_{G_k}$, we have each $\psi_i\sigma$ is representable in G_k . It follows then $\psi\sigma$ which is γ is representable in G_{k+1} .

(\Leftarrow) Let $k \leq \ell$ be the least index s.t. γ is representable in G_k . We proceed by induction on k . For $k = 0$, note that γ is representable in $G_0 = \text{egraph}(\varphi)$ hence $\varphi \models_{\mathcal{T}} \gamma$. Now, for the inductive step, if γ is $P(t_1, \dots, t_n)$ that is different from an equality atom, then there are nodes $\bar{u} \in \Pi_k^\psi$ and $w_{(\bar{u}, P)} \in N_k \setminus N_{k-1}$ for some $\psi \in \Phi \cup \Phi_{eq}$ where ψ is of the form $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_n \rightarrow P(t'_1, \dots, t'_n))$ and $\text{term}(u_i) = t'_i\sigma = t_i$ for some substitution σ . Moreover, each $\psi_i\sigma$ is representable in G_{k-1} . By the inductive hypothesis we have $\varphi \models_{\mathcal{T}} \psi_i\sigma$ hence, $\varphi \models_{\mathcal{T}} P(t'_1, \dots, t'_n)\sigma$, and thus $\varphi \models_{\mathcal{T}} P(t_1, \dots, t_n)$ as required. Otherwise, γ is an equality atom, $t_1 \approx t_2$, and we have $(u_1, u_2) \in \Delta_k^\psi$ for some axiom ψ of the form $\forall \bar{x}(\psi_1 \wedge \dots \wedge \psi_n \rightarrow t'_1 \approx t'_2)$ s.t. $\text{term}(u_i) = t'_i\sigma$, (u_1, u_2) is the j th element, (u_j, v_j) , in the enumeration of Δ_k and $\text{root}_k^j(n_{t_1}) = \text{root}_k^j(u_j)$ and $\text{root}_k^j(n_{t_2}) = \text{root}_k^j(v_j)$. It can be shown by induction on j and the induction hypothesis that $\varphi \models_{\mathcal{T}} t'_i\sigma \approx t_i$. Furthermore, each $\psi_i\sigma$ is representable in G_{k-1} and hence by the inductive hypothesis $\varphi \models_{\mathcal{T}} \psi_i\sigma$. Therefore, $\varphi \models_{\mathcal{T}} (t'_1 \approx t'_2)\sigma$ and hence $\varphi \models_{\mathcal{T}} t_1 \approx t_2$ by transitivity. \square

Corollary 7.1.4. *Assume \mathcal{T} is a locally ground theory and φ a \mathcal{T} -satisfiable conjunction of literals. Then, for any subterms $t, s \in S_\varphi$, $t \sim_{\mathcal{T}, \lambda}^\varphi s$ iff $\text{root}^*(n_t) = \text{root}^*(n_s)$.*

It follows then the class of a node, $\rho_{\text{root}^*}(n_t)$, that emerges in G^* corresponds to the class of its term $[t]_{\mathcal{T}, \lambda}^\varphi$ that is induced by $\sim_{\mathcal{T}, \lambda}^\varphi$ on the subterm set of φ . This is essential as G^* faithfully represents the partition $S_\varphi / \sim_{\mathcal{T}, \lambda}^\varphi$, and hence allows for searching of the constructively ground terms.

7.2 Soundness and Relative Completeness

In this section, we prove for \mathcal{T} -QEL soundness in general and relative completeness for the locally ground theories. Given an egraph G , the class of a node $n \in N$, $\text{class}_G(n) := \rho_{\text{root}}(n)$, is the set of all nodes that are equivalent to n . We denote by $\text{children}(n)$ the set of nodes with an incoming edge from n . We now adopt two definitions from [20] below.

Definition 7.2.1. *(Constructively Ground Node [20]) Let G be an egraph and n be a node in G , then n is a constructively ground (c-ground) node if either (i) it is labelled by a constant symbol, or (ii) $\text{deg}(n) > 0$ and for each child $n[i]$ of n , there is a c-ground node in the class($n[i]$).*

Call an equivalence class of a node n , $\text{class}(n)$, *ground* if it contains a c-ground node.

Definition 7.2.2. (*Admissible Representative Function [20]*) Given an egraph $G = \langle N, E, L, \text{root} \rangle$, a representative function $\beta : N \rightarrow N$ is admissible for G if

- (a) β assigns unique representative per class.
- (b) $\rho_\beta = \rho_{\text{root}}$.
- (c) the graph G_β is acyclic where $G_\beta = \langle N, E_\beta \rangle$, and $E_\beta := \{ (n, \beta(c)) \mid c \in \text{children}(n), n \in N \}$.

Call a representative function $\beta : N \rightarrow N$ *maximally ground* if for every node $n \in N$, $\beta(n)$ is *c-ground* whenever $\text{class}(n)$ is ground. We assume for any representative function, β , whenever $\beta(n_\top) \neq \beta(n_\perp)$ it selects n_\top (resp. n_\perp) as their own representatives, i.e., $\beta(n_\top) = n_\top$ and $\beta(n_\perp) = n_\perp$.

As a consequence of Corollary 7.1.4 we have for any *c-ground term* $t \in S_\varphi$, the associated node $n_t \in N^*$ is a *c-ground node*. Another consequence of Corollary 7.1.4 is that G^* is indeed a valid egraph, that is the congruence condition (condition (c)) in the definition of egraphs is satisfied. Now let $\gamma = \bigwedge \gamma_i$ be a conjunction of literals such that $G^* = \text{egraph}(\gamma)$. Note that as φ is satisfiable, so is γ . Below we show using Lemma 7.1.3 that γ is equivalent with $\varphi^* = \varphi \wedge \bigwedge_{\theta \in D} \theta$, where D contains all the atoms over the subterm of φ obtained via restricted derivations, i.e., $D = \{ \theta \mid \Gamma_\varphi \cup \Phi \Vdash \theta \text{ and } S_\theta \subseteq S_\varphi \}$. Intuitively, the lemma below lets us view φ^* as the “completion” of φ and G^* as the egraph of φ^* .

Lemma 7.2.3. *Let γ and φ^* be as given above, then $\models \varphi^* \leftrightarrow \gamma$.*

Proof. We proceed in two steps, first we show $\models \varphi^* \rightarrow \gamma$ then we show $\models \gamma \rightarrow \varphi^*$.
i. ($\models \varphi^* \rightarrow \gamma$) We show $\models \varphi^* \rightarrow \gamma_i$. From which the conclusion follows. If γ_i is an atom, then we have by the construction of egraphs, γ_i is representable in the $\text{egraph}(\gamma) = G^*$, hence by Lemma 7.1.3 $\varphi \models_{\mathcal{T}} \gamma_i$. As \mathcal{T} is a locally ground theory, $\Gamma_\varphi \cup \Phi \Vdash \gamma_i$, hence $\gamma_i \in D$ and $\models \varphi^* \rightarrow \gamma_i$. Otherwise, γ_i is a negative literal and similarly, γ_i is representable in $\text{egraph}(\gamma) = G^*$. Moreover, as φ is satisfiable, the negative literals represented in G^* are exactly the negative literals represented in $G_0 = \text{egraph}(\varphi)$. Which in turn are exactly the negative literals occurring in φ . Therefore, γ_i occurs in φ and thus $\models \varphi^* \rightarrow \gamma_i$.

ii. ($\models \gamma \rightarrow \varphi^*$) Similarly, we show $\models \gamma \rightarrow \theta_i$ where θ_i is a literal occurring in φ^* . Assume θ_i is an atom, then as $\varphi \models_{\mathcal{T}} \theta_i$ we have by Lemma 7.1.3 θ_i is representable in $G^* = \text{egraph}(\gamma)$, and hence by construction and completeness of egraphs $\gamma \models \theta_i$. Otherwise, θ_i is a negative literal occurring in φ , and hence is represented in $G^* = \text{egraph}(\gamma)$. Similarly, as γ is satisfiable θ_i occurs in γ and hence $\gamma \models \theta_i$. \square

We will use Lemma 7.2.3 to later establish that the result of \mathcal{T} -QEL($\exists \bar{x} \varphi(\bar{x}), \Phi$) is a quantifier reduction of $\exists \bar{x} \varphi(\bar{x})$. We restate here two results from [20]. The second one gives necessary and sufficient conditions for termination during formula extraction from egraphs. Additionally, it asserts the result of `to_formula` is a quantifier reduction in \mathcal{T}_{eq} (i.e., the theory axiomatized by the empty set in our current context).

Lemma 7.2.4. (Lemma 13 [20]). *Representative functions β computed by `find_defs` are admissible functions that are maximally ground.*

Theorem 7.2.5. (Theorem 1 [20]). Let G be the egraph of some conjunction of literals φ , core computed by QEL, and β an arbitrary representative function. Then, the function `G.to_formula`($\beta, G.Nodes() \setminus core$) terminates with result ψ such that $\models \exists \bar{x} \varphi \leftrightarrow \exists \bar{x} \psi$ iff β is admissible for G .

Below, we show that for nodes whose associated terms have an entailed ground definition, their representative selected by a maximally ground representative function is c-ground.

Lemma 7.2.6. *Let G^* be the completion of $G = egraph(\varphi)$, β an admissible representative function that is maximally ground for G^* . Then, for any $t \in S_\varphi$ and its associated node $n_t \in N^*$, if $\varphi \models_{\mathcal{T}} t \approx s$, for some ground term s , then $\beta(n_t)$ is c-ground and `ntt`($\beta(n_t)$) is ground.*

Proof. Assume that $\varphi \models_{\mathcal{T}} t \approx s$ for some ground term s , then as \mathcal{T} is a locally ground theory we have for some c-ground term $r \in S_\varphi$, $r \sim t$. By Corollary 7.1.4 we have $root^*(n_t) = root^*(n_r)$, hence $n_r \in class(n_t)$. Now as r is a c-ground term we have n_r is a c-ground node, thus by definition of maximally ground $\beta(n_t)$ is c-ground. The rest of the proof for `ntt`($\beta(n_t)$) being ground is the same as Theorem 2 in [20]. \square

Theorem 7.2.7. (Soundness and Relative Completeness). Let \mathcal{T} be a locally ground theory axiomatized by Φ and $\varphi(\bar{x})$ a \mathcal{T} -satisfiable conjunction of literals. Then,

- (a) The result of \mathcal{T} -QEL($\exists \bar{x} \varphi(\bar{x}), \Phi$) is a quantifier reduction of $\exists \bar{x} \varphi(\bar{x})$.
- (b) The algorithm \mathcal{T} -QEL is relatively complete for \mathcal{T} .

Proof. Let $\psi(\bar{y})$ be the result of \mathcal{T} -QEL($\exists \bar{x} \varphi(\bar{x}), \Phi$). We show that conditions (i), (ii) and (iii) given in the initial definition of [partial quantifier elimination](#) are satisfied for $\psi(\bar{y})$. In what follows, we let β be the representative function computed by QEL.

- (i) $\models_{\mathcal{T}} \exists \bar{x} \varphi(\bar{x}) \leftrightarrow \exists \bar{y} \psi(\bar{y})$. Let φ^* and γ be as given in Lemma 7.2.3. First note that $\models_{\mathcal{T}} \varphi \leftrightarrow \varphi^*$, furthermore, we have by Lemma 7.2.3, $\models \gamma \leftrightarrow \varphi^*$ giving us $\models_{\mathcal{T}} \varphi \leftrightarrow \gamma$. By Lemma 7.2.4 we have β is an admissible representative function and hence by Theorem 7.2.5 $G^*.to_formula$ terminates with result ψ s.t., $\models \exists \bar{y} \psi \leftrightarrow \exists \bar{x} \gamma$. The conclusion then follows, $\models_{\mathcal{T}} \exists \bar{x} \varphi(\bar{x}) \leftrightarrow \exists \bar{y} \psi(\bar{y})$.
- (ii) $FV(\psi) \subseteq FV(\varphi)$. Trivial.
- (iii) Assume $\varphi(\bar{x}) \models_{\mathcal{T}} x_i \approx t$ for some ground term t , then by Lemma 7.2.4 we have β is an admissible representative function that is maximally ground. Hence, by Lemma 7.2.6 we have the representative selected $\beta(n_{x_i})$ is c-ground and $ntt(\beta(n_{x_i}))$ is ground. Thus, x_i is successfully eliminated from ψ , the output of $G^*.to_formula$.

□

Chapter 8

Conclusion and Related Work

In this thesis, we studied *local theories* focusing on their proof theoretic and semantic characterization. In Chapter 4, we explored the proof theoretic characterization to show polynomial time decidability of the uniform word problem. We further showed that local theories provide a characterization of the complexity class P, i.e., any PTIME problem can be encoded as a uniform word problem for local theories. In Chapter 5, we gave a semantic characterization of local theories, focusing on the polynomial time decidability of the uniform word problem.

In Chapter 6, we identified a subclass of local theories, called *locally ground theories*, in which we can perform partial quantifier elimination efficiently while maintaining relative completeness. By lifting the proof theoretic characterization of locality to FOL with equality, in Chapter 7, we gave the polynomial time algorithm, \mathcal{T} -QEL. We showed that \mathcal{T} -QEL is sound in general and relatively complete for the locally ground theories. We showed several theories, which were previously shown to exhibit locality properties, were also locally ground.

For future work, it would be interesting to see how relative completeness is preserved for other theories and under combination of theories. Moreover, we leave for future work investigating the potential applications as an efficient preprocessing step for arithmetical theories.

8.1 Related Work

As was shown in Section 4.2 the class of local theories is undecidable. That is, we cannot uniformly decide whether a given theory is indeed a local theory. One strand of work found in McAllester’s [37], and Basin and Ganzinger’s [1, 2] is that of identifying subclasses of local theories in which membership in the class is (semi-)decidable. By generalizing the notion of locality, Basin and Ganzinger introduce the notion of order locality in [1, 2]. Where we consider different well-founded orderings other than the subterm ordering inherent in local theories. They show that saturation up to redundancy under ordered resolution can detect locality in many cases.

Another closely related concept is the notion of *local theory extensions* [45]. Unlike the local theories that were considered in this thesis, in local theory extensions, one extends the signature Σ_0 , of a given *base theory* \mathcal{T}_0 , to the signature Σ_1 by adding new function symbols. The properties of these new function symbols are axiomatized using a set of universal Horn clauses Φ , after which one considers the *theory extension* $\mathcal{T}_1 = \text{Cn}(\mathcal{T}_0 \cup \Phi)$. The extension \mathcal{T}_1 is said to be *local* if entailment of a ground clause in \mathcal{T}_1 can be reduced to entailment in \mathcal{T}_0 by considering only the instances of Φ in which the subterms are restricted to the subterms of the clause. This can be seen as a generalization of the notion of locality.

Local theory extensions have recently received significant attention. Recent work in *local theory extensions* has focused on interpolation and symbol elimination [50], decision procedures and combination of local theory extensions [3, 48, 27], and applications in mathematics and verification [49, 47, 46].

References

- [1] David A. Basin and Harald Ganzinger. Complexity analysis based on ordered resolution. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, pages 456–465. IEEE Computer Society, 1996.
- [2] David A. Basin and Harald Ganzinger. Automated complexity analysis based on ordered resolution. *J. ACM*, 48(1):70–109, 2001.
- [3] Markus Bender and Viorica Sofronie-Stokkermans. Decision procedures for theories of sets with measures. In Leonardo de Moura, editor, *Automated Deduction - CADE 26 - 26th International Conference on Automated Deduction, Gothenburg, Sweden, August 6-11, 2017, Proceedings*, volume 10395 of *Lecture Notes in Computer Science*, pages 166–184. Springer, 2017.
- [4] Joel Berman and Willem J. Blok. Equational dependencies. *J. Inf. Process. Cybern.*, 28(4):213–223, 1992.
- [5] Nikolaj S Bjørner and Mikolás Janota. Playing with quantified satisfaction. *LPAR (short papers)*, 35:15–27, 2015.
- [6] Stanley Burris. Polynomial time uniform word problems. *Math. Log. Q.*, 41:173–182, 1995.
- [7] Stanley Burris and John Lawrence. The equivalence problem for finite rings. *J. Symb. Comput.*, 15(1):67–71, 1993.
- [8] Chen Chung Chang and H. Jerome Keisler. *Model Theory*. North Holland, Amsterdam, Netherlands, 1973.
- [9] David C Cooper. Programs for mechanical program verification. *Machine Intelligence*, 6(1), 1971.

- [10] David C Cooper. Theorem proving in arithmetic without multiplication. *Machine intelligence*, 7(91-99):300, 1972.
- [11] James H Davenport and Joos Heintz. Real quantifier elimination is doubly exponential. *Journal of Symbolic Computation*, 5(1-2):29–35, 1988.
- [12] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Z3: an efficient SMT solver. In C. R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
- [13] Bruno Dutertre. Yices 2.2. In Armin Biere and Roderick Bloem, editors, *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, volume 8559 of *Lecture Notes in Computer Science*, pages 737–744. Springer, 2014.
- [14] Bruno Dutertre. Solving exists/forall problems with yices. In *Workshop on satisfiability modulo theories*, 2015.
- [15] Herbert B. Enderton. *A mathematical introduction to logic*. Harcourt/Academic Press, Burlington, MA, second edition, 2001.
- [16] Trevor Evans. The word problem for abstract algebras. *Journal of The London Mathematical Society-second Series*, pages 64–71, 1951.
- [17] Ronald Fagin. Generalized first-order spectra and polynomial-time recognizable sets. *Complexity of computation*, 7:43–73, 1974.
- [18] Melvin Fitting. *First-Order Logic and Automated Theorem Proving, Second Edition*. Graduate Texts in Computer Science. Springer, 1996.
- [19] Harald Ganzinger. Relating semantic and proof-theoretic concepts for polynomial time decidability of uniform word problems. In *Proceedings 16th Annual IEEE Symposium on Logic in Computer Science*, pages 81–90. IEEE, 2001.
- [20] Isabel Garcia-Contreras, V. K. Hari Govind, Sharon Shoham, and Arie Gurfinkel. Fast approximations of quantifier elimination. In *Computer aided verification. Part II*, volume 13965 of *Lecture Notes in Comput. Sci.*, pages 64–86. Springer, 2023.

- [21] Silvio Ghilardi and Silvio Ranise. MCMT: A model checker modulo theories. In Jürgen Giesl and Reiner Hähnle, editors, *Automated Reasoning, 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010. Proceedings*, volume 6173 of *Lecture Notes in Computer Science*, pages 22–29. Springer, 2010.
- [22] Robert Givan and David Mcallester. Polynomial-time computation via local inference relations. *ACM Transactions on Computational Logic (TOCL)*, 3(4):521–541, 2002.
- [23] Robert Givan and David A. McAllester. New results on local inference relations. In Bernhard Nebel, Charles Rich, and William R. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR'92). Cambridge, MA, USA, October 25-29, 1992*, pages 403–412. Morgan Kaufmann, 1992.
- [24] Christoph Haase, Shankara Narayanan Krishna, Khushraj Madnani, Om Swostik Mishra, and Georg Zetsche. An efficient quantifier elimination procedure for presburger arithmetic. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024, Tallinn, Estonia*, volume 297 of *LIPICs*, pages 142:1–142:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
- [25] Wilfrid Hodges. Logical features of Horn clauses. In *Handbook of logic in artificial intelligence and logic programming, Vol. 1*, Oxford Sci. Publ., pages 449–503. Oxford Univ. Press, New York, 1993.
- [26] Wilfrid Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications 42. Cambridge University Press, 1 edition, 1993.
- [27] Carsten Ihlemann and Viorica Sofronie-Stokkermans. On hierarchical reasoning in combinations of theories. In Jürgen Giesl and Reiner Hähnle, editors, *Automated Reasoning, 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010. Proceedings*, volume 6173 of *Lecture Notes in Computer Science*, pages 30–45. Springer, 2010.
- [28] Harry B. Hunt III, Daniel J. Rosenkrantz, and Peter A. Bloniarz. On the computational complexity of algebra on lattices. *SIAM J. Comput.*, 16(1):129–148, 1987.
- [29] Neil Immerman. Relational queries computable in polynomial time. *Inf. Control.*, 68(1-3):86–104, 1986.

- [30] Neil Immerman. Descriptive complexity: A logician’s approach to computation. *Notices of the American Mathematical Society*, 42(10):1127–1133, 1995.
- [31] Neil Immerman. *Descriptive complexity*. Graduate texts in computer science. Springer, 1999.
- [32] Anvesh Komuravelli, Arie Gurfinkel, and Sagar Chaki. Smt-based model checking for recursive programs. In Armin Biere and Roderick Bloem, editors, *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, volume 8559 of *Lecture Notes in Computer Science*, pages 17–34. Springer, 2014.
- [33] Dexter Kozen. Complexity of finitely presented algebras. In John E. Hopcroft, Emily P. Friedman, and Michael A. Harrison, editors, *Proceedings of the 9th Annual ACM Symposium on Theory of Computing, May 4-6, 1977, Boulder, Colorado, USA*, pages 164–177. ACM, 1977.
- [34] Daniel Kroening and Ofer Strichman. *Decision Procedures - An Algorithmic Point of View, Second Edition*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2016.
- [35] Viktor Kuncak, Mikael Mayer, Ruzica Piskac, and Philippe Suter. Complete functional synthesis. In Benjamin G. Zorn and Alex Aiken, editors, *Proceedings of the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2010, Toronto, Ontario, Canada, June 5-10, 2010*, pages 316–329. ACM, 2010.
- [36] John W. Lloyd. *Foundations of Logic Programming, 2nd Edition*. Springer, 1987.
- [37] David A. McAllester. Automatic recognition of tractability in inference relations. *J. ACM*, 40(2):284–303, 1993.
- [38] Jasper Nalbach, Valentin Promies, Erika Ábrahám, and Paul Kobialka. Fmplex: A novel method for solving linear real arithmetic problems. *arXiv preprint arXiv:2310.00995*, 2023.
- [39] Sara Negri and Jan von Plato. *Order and lattice theory. In Proof Analysis: A Contribution to Hilbert’s Last Problem*, page 50–67. Cambridge University Press, 2011.
- [40] Greg Nelson and Derek C. Oppen. Fast decision algorithms based on union and find. In *18th Annual Symposium on Foundations of Computer Science (Providence, R.I., 1977)*, pages 114–119. IEEE, Long Beach, CA, 1977.

- [41] Derek C Oppen. A 222pn upper bound on the complexity of presburger arithmetic. *Journal of Computer and System Sciences*, 16(3):323–332, 1978.
- [42] Derek C. Oppen. Reasoning about recursively defined data structures. *J. ACM*, 27(3):403–411, July 1980.
- [43] Daniel Riley and Grigory Fedyukovich. Multi-phase invariant synthesis. In Abhik Roychoudhury, Cristian Cadar, and Miryung Kim, editors, *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2022, Singapore, Singapore, November 14-18, 2022*, pages 607–619. ACM, 2022.
- [44] Robert E Shostak. An algorithm for reasoning about equality. *Communications of the ACM*, 21(7):583–585, 1978.
- [45] Viorica Sofronie-Stokkermans. Hierarchic reasoning in local theory extensions. In Robert Nieuwenhuis, editor, *Automated Deduction - CADE-20, 20th International Conference on Automated Deduction, Tallinn, Estonia, July 22-27, 2005, Proceedings*, volume 3632 of *Lecture Notes in Computer Science*, pages 219–234. Springer, 2005.
- [46] Viorica Sofronie-Stokkermans. Hierarchical reasoning for the verification of parametric systems. In Jürgen Giesl and Reiner Hähnle, editors, *Automated Reasoning, 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010. Proceedings*, volume 6173 of *Lecture Notes in Computer Science*, pages 171–187. Springer, 2010.
- [47] Viorica Sofronie-Stokkermans. Hierarchical reasoning and model generation for the verification of parametric hybrid systems. In Maria Paola Bonacina, editor, *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, volume 7898 of *Lecture Notes in Computer Science*, pages 360–376. Springer, 2013.
- [48] Viorica Sofronie-Stokkermans. On combinations of local theory extensions. In Andrei Voronkov and Christoph Weidenbach, editors, *Programming Logics - Essays in Memory of Harald Ganzinger*, volume 7797 of *Lecture Notes in Computer Science*, pages 392–413. Springer, 2013.
- [49] Viorica Sofronie-Stokkermans. Hierarchical reasoning in local theory extensions and applications. In Franz Winkler, Viorel Negru, Tetsuo Ida, Tudor Jebelean, Dana Petcu, Stephen M. Watt, and Daniela Zaharie, editors, *16th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC 2014*,

Timisoara, Romania, September 22-25, 2014, pages 34–41. IEEE Computer Society, 2014.

- [50] Viorica Sofronie-Stokkermans. On interpolation and symbol elimination in theory extensions. *Log. Methods Comput. Sci.*, 14(3), 2018.
- [51] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math*, 5:285–309, 1955.
- [52] Rüdiger Thiele. Hilbert’s twenty-fourth problem. *Am. Math. Mon.*, 110(1):1–24, 2003.
- [53] Ruediger Thiele and Larry Wos. Hilbert’s twenty-fourth problem. *J. Autom. Reason.*, 29(1):67–89, 2002.
- [54] Moshe Y. Vardi. The complexity of relational query languages (extended abstract). In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 137–146. ACM, 1982.
- [55] Volker Weispfenning. The complexity of linear problems in fields. *Journal of Symbolic Computation*, 5(1):3–27, 1988.
- [56] Volker Weispfenning. Complexity and uniformity of elimination in presburger arithmetic. In *Proceedings of the 1997 international symposium on Symbolic and algebraic computation*, pages 48–53, 1997.