Relatively Complete and Efficient Partial Quantifier Elimination

Estifanos Getachew(✉), Arie Gurfinkel, and Richard Trefler

University of Waterloo, Waterloo, Canada
{esgetach,agurfink,trefler}@uwaterloo.ca

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. Considerably generalizing the previous approach, we identify a subclass of *local theories* in which *partial quantifier elimination* can be performed efficiently. We present $\mathcal{T}$-QEL a parametrized polynomial time algorithm that is a sound extension of QEL and 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.

1 Introduction

Several automated reasoning tasks, including quantified SMT solving [1], exists/forall solving [3,4], program synthesis [17], model checking [15,9], and CHC solving [22] utilize quantifier elimination. Regardless, the intractability of complete quantifier elimination procedures, see e.g. [16,12,2], 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 [8], that resulted in significant performance improvements in Z3 [19] and the CHC solver Spacer [15].

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 work was supported, in part, by Individual Discovery Grants from the Natural Sciences and Engineering Research Council of Canada.

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 x_i 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* [10,7] 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 [10,11,18]. 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 [7] 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 [7], 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 [7]. 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 $s_0, s_1, \ldots$ 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 theories[1], 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.

[1] Strictly speaking, we are considering an extended notion of *locality* from [7].

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 $\mathcal{T}_{eq}$.

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.

Contributions. The contributions of this paper 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.
- 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* with respect to the locally ground theories.

2 Basic Notions and Notation

First Order Logic [5]. We assume the usual concepts from First order logic (FOL) with equality such as model, satisfaction, logical consequence, and a theory, as presented in e.g. [5,6]. We briefly summarize the notation used here.

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 $\alpha : \mathcal{V} \to A$. We denote by $t^{\mathcal{I}}$ the interpretation of the term t under $\mathcal{I}$. We use the notation $\mathrm{FV}(\varphi)$ 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 $f(t_1, \ldots, t_n)$. 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^F the set of all function symbols in Σ, and by Σ^T the set of all terms generated from Σ. Given a function $\lambda : \mathcal{P}(\Sigma^T) \to \mathcal{P}(\Sigma^T)$ 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 $f(t_1, \ldots, t_n), f(s_1, \ldots, s_n) \in \Psi$ we have $f(t_1, \ldots, t_n) \, R \, f(s_1, \ldots, s_n)$.

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, denoted by $\mathrm{Cn}\,\Phi = \{\varphi \mid \Phi \models \varphi\}$, is a theory. 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 $\mathcal{T}$.

Horn Theories and Direct products [14]. We use the term basic Horn clause for formulas of the form $\bigwedge \psi_i \to \psi$, where ψ_i, ψ are atoms, to distinguish them from (universal) Horn clauses $\forall \bar{x} \, (\bigwedge \psi_i \to \psi)$. Given a non-empty index set I and a family of Σ-structures $(\mathfrak{A}_i)_{i \in I}$, we use the notation $\Pi_{i \in I} \, \mathfrak{A}_i$ for the *direct product (product)* of the given family of structures. For background on direct products of models of Horn theories and products in general, see Ch. 9 of [14].

Egraphs [20,8]. An egraph is a well-known data structure to compactly represent a set of terms and an equivalence relation on those terms [20]. We assume that graphs have an ordered successor relation and use $n[i]$ to denote the ith successor (child) of a node n. We denote by $deg(n)$, the out degree of a node n, i.e., the number of edges leaving n.

Definition 1. *(Egraph [8]) For a given signature Σ, an egraph is a tuple $G = \langle N, E, L, root \rangle$, where*

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

Given an egraph G, we let $term_G : N \to \Sigma^T$ 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 $egraph(\varphi)$ the egraph of φ built by the standard procedure given in [20], see Section 2 of [8] for more.

3 Overview

Intuitively, $\mathcal{T}$-QEL, the algorithm presented in Section 5, 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, we want to do this in *polynomial time*, and we want to maintain *relative completeness* — that is, we want to eliminate every variable that has an entailed ground definition. We illustrate our approach in the following example.[2]

Example 1. 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

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

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 $S_\varphi/{\sim}$ of the subterms. Let $\sim$ be the congruence relation defined on the subterm set of φ, S_φ, as: $t_1 \sim t_2$ 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 $S_\varphi/{\sim}$, e.g. $[x_1] = \{ x_1, f(c) \}$.
Construct a ground term. Next, we pick for each equivalence class in $S_\varphi/{\sim}$ (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: $f(c) \approx f(c) \land f(x_2) \approx g(f(c)) \land g(g(f(c))) \leq$

[2] The intuitive explanation presented here slightly differs from the presentation of $\mathcal{T}$-QEL in Sec. 5. See the explanation for Algorithm 1($\mathcal{T}$-QEL) provided in Sec. 5.

$g(g(f(c))) \wedge g(g(f(c))) \leq g(g(f(c)))$. 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 $S_\varphi/\!\sim$ on an egraph. Egraphs give us a compact way to represent congruence relations. For local theories such as $\mathcal{T}_{po}$, representing the partition $S_\varphi/\!\sim$ 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 Section 5.

4 Locally Ground Theories

In this section, 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. We introduce the notion of being a *constructively ground term* with respect to a conjunction of literals φ and a theory $\mathcal{T}$. We give a definition for the *locally ground theories* based on this notion.

We adopt here the notion of a *restricted derivation* on which the definition of local theories (in FOL without equality) [10] 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).

A local theory in FOL with equality [7] 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.[3,4] 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.

Theorem 2. *(Ganzinger [7]). 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 [10] motivates the proof theoretic definition we

[3] Note that, only ground formulas are considered in [7] and [10]. 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.

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

adopt below. We note here that, the definition of locality we give below subsumes the original notion of *locality* given in [10,7], 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 $\lambda : \mathcal{P}(\Sigma^T) \to \mathcal{P}(\Sigma^T)$ 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 3. *(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 $\theta_1, \ldots, \theta_n$ s.t. θ_n is θ and we have for each θ_i* ***(i)*** *$S_{\theta_i} \subseteq S_\varphi^\lambda$ and,* ***(ii)*** *there exists a Horn clause $\forall \bar{x}\, (\psi_1 \wedge \cdots \wedge \psi_k \to \psi) \in \Phi \cup \Phi_{eq}$ 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 [10]. 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 $f(t_1, \ldots, t_n)$ 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 4. *(Locally Ground Theory). Call a theory $\mathcal{T}$ locally ground, if there exists a proper function λ such that* ***(i)*** *$\mathcal{T}$ is a local theory with λ, 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. We state the following variant of McKinsey's lemma for theories axiomatized by Horn clauses. This will help us ignore the negative literals when considering the logical consequences of a (satisfiable) conjunction of literals, as shown in Corollary 6. The proof involves noting that models of Horn formulas are closed under direct products. Other similar formulations of McKinsey's lemma can be found in [14,13].

Lemma 5. *(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 $\psi \in \Psi(\bar{x})$.*

Corollary 6. *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 $\bigwedge \varphi_i \models_{\mathcal{T}} \psi(\bar{x})$.*

4.1 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* [21] as shown in Section 4.2. 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 $\mathcal{P}(\Sigma^T_{po})$. Moreover, we simply write $\sim$ and $[-]$ in place of $\sim^{\varphi}_{\mathcal{T}_{po},\lambda_p}$ and $[t]^{\varphi}_{\mathcal{T}_{po},\lambda_p}$, when there is no ambiguity.

Lemma 7. *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 7 and proceed to show that $\mathcal{T}_{po}$ is a locally ground theory. After which we present a proof of the lemma.

Theorem 8. *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 [24], 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 7, $\mathcal{I}_p$ is a model of φ and $\mathcal{T}_{po}$, hence $t^{\mathcal{I}_p} = s^{\mathcal{I}_p}$. By (ii) in Lemma 7, $t^{\mathcal{I}_p} = [t] = s^{\mathcal{I}_p}$ and hence $s^{\mathcal{I}_p} \neq \zeta$. Thus, by (iii) in lemma 7, $s^{\mathcal{I}_p} = [r]$ for some c-ground term r, and hence $[t]$ is ground. □

Proof. (of Lemma 7) 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 $a_1, \ldots, a_n \in A_p$,
$f^{\mathfrak{A}_p}(a_1, \ldots, a_n) = [f(s_1, \ldots, s_n)]$ if $\exists\, \bar{s}, f(s_1, \ldots, s_n) \in S_\varphi$ and $\forall i, a_i = [s_i]$;
otherwise $f^{\mathfrak{A}_p}(a_1, \ldots, a_n) = \zeta$.
(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 $\Gamma_\varphi \cup \Phi_{po} \Vdash t_1 \leq t_2$.[5]

Finally, set $\mathcal{I}_p = (\mathfrak{A}_p, \alpha_p)$, where the variable assignment $\alpha_p : \mathcal{V} \to A_p$ 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 7. 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 $c^{\mathcal{I}} = [c]$ or $c \notin S_\varphi$ and $c^{\mathcal{I}} = \zeta$. Otherwise, let t be the term $f(t_1, \ldots, t_n)$ 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 $t^{\mathcal{I}_p} = f^{\mathfrak{A}_p}([u_1], \ldots, [u_n]) = [f(s_1, \ldots, s_n)]$ 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 $f(s_1, \ldots, s_n)$ is a constructively ground term and the conclusion follows.

Finally, we show that $\mathcal{I}_p$ is a model of both $\mathcal{T}_{po}$ and φ. Note that, for any t_1, t_2 in S_φ, we have by Corollary 6 and the fact that Φ_{po} is a local axiomatization of $\mathcal{T}_{po}$, $[t_1] \leq^{\mathfrak{A}_p} [t_2]$ iff $\varphi \models_{\mathcal{T}_{po}} t_1 \leq t_2$. Hence, $\leq^{\mathfrak{A}_p}$ is reflexive, transitive, and antisymmetric on A_p. Therefore, $\mathcal{I}_p \models \Phi_{po}$, and thus is a model of $\mathcal{T}_{po}$. 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 $\varphi \models_{\mathcal{T}_{po}} \varphi_i$, and that for each $t \in S_\varphi$, $t^{\mathcal{I}_p} = [t]$. Therefore, it follows $\mathcal{I}_p \models_{\mathcal{T}_{po}} \varphi$. □

The construction is motivated by Shostak's decision procedure [23] for the quantifier free fragment of the theory of equality. However, unlike [23] we don't

[5] Equivalently, $[t_1] \leq^{\mathfrak{A}_p} [t_2]$ iff $\varphi \models_{\mathcal{T}_{po}} t_1 \leq t_2$.

use the entire Herbrand universe, and simply interpret/map terms to their equivalence classes under $\sim$.

4.2 Recursively Defined Data Structures

We now show that the theory of Recursively Defined Data Structures, $\mathcal{T}_{rd}$, [21] without the acyclicity axioms is locally ground. We take as our working signature $\Sigma_{rd} = (\Sigma^F \cup \{ cr, sr_1, \ldots, sr_k \}, \varnothing)$ 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:

$$\forall x\; cr(sr_1(x), \ldots, sr_k(x)) \approx x \qquad \text{(Construction)}$$
$$\forall \bar{x}\; sr_i(cr(x_1, \ldots, x_k)) \approx x_i \quad \text{for } i = 1, \ldots, k \qquad \text{(Selection)}$$

The approach we take in showing that $\mathcal{T}_{rd}$ is a locally ground theory is similar to the approach we took in Section 4.1. Following Oppen's method in [21], we construct an infinite model of $\mathcal{T}_{rd}$. Let $\lambda_r : \mathcal{P}(\Sigma^T_{rd}) \to \mathcal{P}(\Sigma^T_{rd})$ be the function such that for each $S \in \mathcal{P}(\Sigma^T_{rd})$,

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

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

Lemma 9. *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 $S^{\lambda_r}_{\varphi}/\sim$ 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 $[t] \in Z \cap S^{\lambda_r}_{\varphi}/\sim$, $[t]$ is ground.*

We defer the proof of Lemma 9 and proceed to show that $\mathcal{T}_{rd}$ is a locally ground theory. After which we present a proof of the lemma.

Theorem 10. *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 $\Phi_{rd}[S^{\lambda_r}_{\psi}] \models \psi$. 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

9, 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 9, $t^{\mathcal{I}_r} = [t] = s^{\mathcal{I}_r}$, moreover, by (iii) in Lemma 9, $s^{\mathcal{I}_r} \in Z$. Hence, by definition of Z, $[t]$ is ground. □

Proof. (of Lemma 9). 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 $i \in \{1, \ldots, k\}$.

(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:
$cr_0(a_1, \ldots, a_k) = [cr(t_1, \ldots, t_k)]$ if $\exists\, \bar{t}, cr(t_1, \ldots, t_k) \in S_\varphi^{\lambda_r}$ s.t. $\forall\, i, a_i = [t_i]$;
otherwise $cr_0(a_1, \ldots, a_k)$ is undefined.
(iii) For $i = 1, \ldots, k$, define the partial functions $sr_{i,0}$ on A_0 as follows:
$sr_{i,0}(a) = [sr_i(t)]$ if $\exists\, t, sr_i(t) \in S_\varphi^{\lambda_r}$ and $a = [t]$; otherwise $sr_{i,0}(a)$ is undefined.

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

(i) Let $S_{n-1} = \{\zeta_1^a, \ldots, \zeta_k^a \mid a \in (A_{n-1} \backslash \mathrm{dom}\, sr_{i,n-1}) \text{ and } \zeta_i^a \notin (A_{n-1} \cup A_{n-1}^k)\}$
(ii) Let $A_n = A_{n-1} \cup (A_{n-1}^k \backslash \mathrm{dom}\, cr_{n-1}) \cup S_{n-1}$.
(iii) Define cr_n to be a partial function on A_n that is an extension of cr_{n-1}, such that for each $\bar{b} = (b_1, \ldots, b_k) \in A_n^k \backslash \mathrm{dom}\, cr_{n-1}$:
$cr_n(\bar{b}) = \bar{b}$ if $\bar{b} \in A_{n-1}^k \backslash \mathrm{dom}\, cr_{n-1}$; otherwise $cr_n(\bar{b}) = a$ if $\bar{b} = (\zeta_1^a, \ldots, \zeta_k^a)$;
and finally, $cr_n(\bar{b})$ is undefined for every other case.
(iv) Similarly, for $i = 1, \ldots, k$, define the partial function $sr_{i,n}$ on A_n, that is an extension of $sr_{i,n-1}$, such that, for all $a \in A_n \backslash \mathrm{dom}\, sr_{i,n-1}$:
$sr_{i,n}(a) = \zeta_i^a$ if $a \in A_{n-1} \backslash \mathrm{dom}\, sr_{i,n-1}$; otherwise $sr_{i,n}(a) = a_i$ if
$a = (a_1, \ldots, a_k) \in A_{n-1}^k \backslash \mathrm{dom}\, cr_{n-1}$; and finally $sr_{i,n}(a)$ 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 $\bar{a} = (a_1, \ldots, a_k)$ in $(A_r)^k$,
$cr^{\mathfrak{A}_r}(a_1, \ldots, a_k) = cr_n(a_1, \ldots, a_k)$, for the least index n, s.t. $\bar{a} \in \mathrm{dom}\, cr_n$.
(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 $a \in dom\, sr_n$.
(iv) For each k-ary function symbol $f \in \Sigma^F$ and $a_1, \ldots, a_k$ in A_r,
$f^{\mathfrak{A}_r}(a_1, \ldots, a_k) = [f(s_1, \ldots, s_k)]$ if $\exists\, \bar{s},\ f(s_1, \ldots, s_k) \in S_\varphi$ and $\forall i, a_i = [s_i]$;
and otherwise $f^{\mathfrak{A}_r}(a_1, \ldots, a_k) = \zeta$.

Finally, define the variable assignment $\alpha_r : \mathcal{V} \to A_r$, 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 [21] the following:

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

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 9. 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:

$$Z_0 = \{\, \zeta \,\} \cup \{\, [t] \in S_\varphi^{\lambda_r}/{\sim} \mid [t] \text{ is ground} \,\} \text{ and for } n > 0,$$
$$Z_n = Z_{n-1} \cup (Z_{n-1}^k \backslash \mathrm{dom}\; cr_{n-1}) \cup \bigcup_{\zeta_i^a \in S_{n-1}} \{\, \zeta_i^a \mid a \in Z_{n-1} \,\}$$

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 $s = f(t_1, \ldots, t_n)$ for some $f \in \Sigma^F$, the argument is the same as in Lemma 7. 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 $a \in \mathrm{dom}\; sr_{i,n}$, 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 $a \in A_{n-1} \backslash \mathrm{dom}\; sr_{i,n-1}$ in which case $s^{\mathcal{I}_r} = sr_{i,n}(a) = \zeta_i^a$ and hence $s^{\mathcal{I}_r} \in Z_n$, or $a = (a_1, \ldots, a_k) \in A_{n-1}^k \backslash \mathrm{dom}\, cr_{n-1}$ and $sr_{i,n}(a) = a_i \in Z_{n-1}$.

Similarly, for the case $s = cr(t_1, \ldots, t_k)$ we have by the induction hypothesis, $t_i^{\mathcal{I}_r} \in Z_{n_i}$. Let n be the least index such that $\bar{b} = (t_1^{\mathcal{I}_r}, \ldots, t_k^{\mathcal{I}_r}) \in \mathrm{dom}\, cr_n$, 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 $\bar{b} \in A_{n-1}^k \backslash \mathrm{dom}\, cr_{n-1}$, in which case $s^{\mathcal{I}_r} = cr_n(\bar{b}) = \bar{b}$, and hence $s^{\mathcal{I}_r} \in Z_n$, or $\bar{b} = (\zeta_1^a, \ldots, \zeta_k^a)$, and $s^{\mathcal{I}_r} = cr_n(\bar{b}) = a \in Z_{n-1}$. □

5 Partial Quantifier Elimination

In this section, 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)** $\mathrm{FV}(\psi) \subseteq \mathrm{FV}(\varphi)$ and **(iii)** for all $x \in \mathrm{FV}(\varphi)$, if x has an entailed ground definition (i.e., $\varphi \models_{\mathcal{T}} x \approx s$ with s an arbitrary ground term), then $x \notin \mathrm{FV}(\psi)$.

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 $\mathcal{T} = \mathrm{Cn}\,\Phi$. 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 $S_\varphi/{\sim}$ 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 $S_\varphi/{\sim}$ 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 $\psi_1 \wedge \cdots \wedge \psi_k \to \theta$, 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.

We refer to intermediate graphs arising from the construction that might not satisfy the congruence condition (c) in the definition of egraphs, as *partial egraphs*. For simplicity of presentation, we follow [8] and allow labelling nodes with predicate symbols and introduce the fresh constant symbols $\top$ and $\bot$.

Definition 11. *(Representability). Given a (partial) egraph, $G = (N, E, L, root)$, 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 $root(n_{t_1}) = root(n_{t_2})$, 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 $P(t_1, \ldots, t_k)$ or $\neg P(t_1, \ldots, t_k)$, and for some node w in N labelled with P, we have $w[i] = n_{t_i}$, and $root(w) = n_\top$ if γ is positive and $root(w) = n_\bot$ otherwise.*

Let $G_0 = egraph(\varphi)$, 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 $\forall\, \bar{x}\, (\psi_1 \wedge \cdots \wedge \psi_m \to P(t_1, \ldots, t_n))$ where each ψ_i is an atom and P is different from the equality symbol, then let $\Pi_k^\psi = \{\, \bar{u} \in N_{k-1}^n \mid term(u_j) = t_j\sigma$ and each $\psi_i\sigma$ is representable in G_{k-1} for some substitution $\sigma \,\}$ and
$\Pi_k = \bigcup \{\, (\bar{u}, P) \mid \bar{u} \in \Pi_k^\psi$ for some ψ of the above form$\,\}$. Then, form
$G_k = (N_k, E_k, L_k, root_k)$ as follows:

(a) $N_k = N_{k-1} \cup \{\, w_{(\bar{u},\, P)} \mid (\bar{u}, P) \in \Pi_k \,\}$
(b) $L_k = L_{k-1} \cup \{\, (w_{(\bar{u},\, P)}, P) \mid w_{(\bar{u},\, P)} \in N_k \backslash N_{k-1} \,\}$
(c) $E_k = E_{k-1} \cup \{\, (w_{(\bar{u},\, P)}, u_i) \mid w_{(\bar{u},\, P)} \in N_k \backslash N_{k-1}$ 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

Algorithm 1 $\mathcal{T}$-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: $G \leftarrow egraph(\varphi)$
2: $G \rightsquigarrow_\Phi G^*$
3: $\beta \leftarrow G^*.find\_defs()$
4: $core \leftarrow G^*.find\_core(\beta)$
5: **return** $G^*.to\_formula(\beta, N^* \backslash core)$

equivalence classes of any 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 $\forall\, \bar{x}(\psi_1 \wedge \cdots \wedge \psi_m \rightarrow t_1 \approx t_2)$, where each ψ_i is an atom. Then let
$\Delta_k^\psi = \{\, \bar{u} \in N_{k-1}^2 \mid term(u_j) = t_j\sigma$ and each $\psi_i\sigma$ is representable in G_{k-1} for some substitution $\sigma\,\}$

Let $\Delta_k = \bigcup \Delta_k^\psi$ and let $(u_1, v_1), \ldots, (u_r, v_r) \in \Delta_k$ 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:

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

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

For a (partial) egraph G, let $\mathcal{X}_G = \{\ \gamma \mid \gamma$ 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 12. *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.*

Our algorithm $\mathcal{T}$-QEL is presented in Algorithm 1. $\mathcal{T}$-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 [8]. 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 $S_\varphi/{\sim}$ of the subterms induced by $\sim$ on the egraph. We note here that once $S_\varphi/{\sim}$ 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 $S_\varphi/{\sim}$ is represented on the egraph G^*. First, we prove the following Lemma.

Lemma 13. *(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. ($\Longrightarrow$) 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 Corollary 6, 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, $\theta_1, \ldots, \theta_d$, 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 $\forall \bar{x} (\psi_1 \wedge \cdots \wedge \psi_n \rightarrow \theta)$ 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 $P(t_1, \ldots, t_n)$, 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, $P(t_1, \ldots, t_n)$ 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 jth 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 $\forall \bar{x}(\psi_1 \wedge \cdots \wedge \psi_n \rightarrow \psi)$ 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}.

($\Longleftarrow$) 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 $G_0 =$ *egraph*(φ) hence $\varphi \models_\mathcal{T} \gamma$. Now, for the inductive step, if γ is $P(t_1, \ldots, t_n)$ that is different from an equality atom, then there are nodes $\bar{u} \in \Pi_k^\psi$ and $w_{(\bar{u},\,P)} \in N_k \backslash N_{k-1}$ for some $\psi \in \Phi \cup \Phi_{eq}$ where ψ is of the form $\forall \bar{x}(\psi_1 \wedge \cdots \wedge \psi_n \rightarrow P(t'_1, \ldots, t'_n))$ and $term(u_i) = t'_i\sigma = t_i$ 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, $\varphi \models_\mathcal{T} P(t'_1, \ldots, t'_n)\sigma$, and thus $\varphi \models_\mathcal{T} P(t_1, \ldots, t_n)$ 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 $\forall \bar{x} (\psi_1 \wedge \cdots \wedge \psi_n \rightarrow t'_1 \approx t'_2)$ s.t. $term(u_i) = t'_i\sigma$, (u_1, u_2) is the jth element,(u_j, v_j), in the enumeration of Δ_k and $root_k^j(n_{t_1}) = root_k^j(u_j)$ and $root_k^j(n_{t_2}) = root_k^j(v_j)$. 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. □

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

Corollary 14. *Assume $\mathcal{T}$ is a locally ground theory and φ a $\mathcal{T}$-satisfiable conjunction of literals. Then, for any subterms $t, s \in S_\varphi$, $t \sim^{\varphi}_{\mathcal{T},\lambda} s$ iff $root^*(n_t) = root^*(n_s)$.*

Theorem 15. *(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 $\mathcal{T}$.

6 Conclusion and Future Work

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, 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.

Disclosure of Interest. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bjørner, N.S., Janota, M.: Playing with quantified satisfaction. LPAR (short papers) **35**, 15–27 (2015)
2. Davenport, J.H., Heintz, J.: Real quantifier elimination is doubly exponential. Journal of Symbolic Computation **5**(1-2), 29–35 (1988)
3. Dutertre, B.: Yices 2.2. In: Biere, A., Bloem, R. (eds.) 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. Lecture Notes in Computer Science, vol. 8559, pp. 737–744. Springer (2014), `https://doi.org/10.1007/978-3-319-08867-9_49`
4. Dutertre, B.: Solving exists/forall problems with yices. In: Workshop on satisfiability modulo theories (2015)
5. Enderton, H.B.: A mathematical introduction to logic. Harcourt/Academic Press, Burlington, MA, second edn. (2001)
6. Fitting, M.: First-Order Logic and Automated Theorem Proving, Second Edition. Graduate Texts in Computer Science, Springer (1996), `https://doi.org/10.1007/978-1-4612-2360-3`
7. Ganzinger, H.: 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. pp. 81–90. IEEE (2001)
8. Garcia-Contreras, I., Hari Govind, V.K., Shoham, S., Gurfinkel, A.: Fast approximations of quantifier elimination. In: Computer aided verification. Part II, Lecture Notes in Comput. Sci., vol. 13965, pp. 64–86. Springer (2023), `https://doi.org/10.1007/978-3-031-37703-7_4`
9. Ghilardi, S., Ranise, S.: MCMT: A model checker modulo theories. In: Giesl, J., Hähnle, R. (eds.) Automated Reasoning, 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6173, pp. 22–29. Springer (2010), `https://doi.org/10.1007/978-3-642-14203-1_3`
10. Givan, R., Mcallester, D.: Polynomial-time computation via local inference relations. ACM Transactions on Computational Logic (TOCL) **3**(4), 521–541 (2002)
11. Givan, R., McAllester, D.A.: New results on local inference relations. In: Nebel, B., Rich, C., Swartout, W.R. (eds.) Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR'92). Cambridge, MA, USA, October 25-29, 1992. pp. 403–412. Morgan Kaufmann (1992)
12. Haase, C., Krishna, S.N., Madnani, K., Mishra, O.S., Zetzsche, G.: An efficient quantifier elimination procedure for presburger arithmetic. In: Bringmann, K., Grohe, M., Puppis, G., Svensson, O. (eds.) 51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024, Tallinn, Estonia. LIPIcs, vol. 297, pp. 142:1–142:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2024), `https://doi.org/10.4230/LIPIcs.ICALP.2024.142`
13. Hodges, W.: Logical features of Horn clauses. In: Handbook of logic in artificial intelligence and logic programming, Vol. 1, pp. 449–503. Oxford Sci. Publ., Oxford Univ. Press, New York (1993)
14. Hodges, W.: Model Theory. Encyclopedia of Mathematics and its Applications 42, Cambridge University Press, 1 edn. (1993)
15. Komuravelli, A., Gurfinkel, A., Chaki, S.: Smt-based model checking for recursive programs. In: Biere, A., Bloem, R. (eds.) 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. Lecture Notes in Computer Science, vol. 8559, pp. 17–34. Springer (2014), https://doi.org/10.1007/978-3-319-08867-9_2
16. Kroening, D., Strichman, O.: Decision Procedures - An Algorithmic Point of View, Second Edition. Texts in Theoretical Computer Science. An EATCS Series, Springer (2016), https://doi.org/10.1007/978-3-662-50497-0
17. Kuncak, V., Mayer, M., Piskac, R., Suter, P.: Complete functional synthesis. In: Zorn, B.G., Aiken, A. (eds.) Proceedings of the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2010, Toronto, Ontario, Canada, June 5-10, 2010. pp. 316–329. ACM (2010), https://doi.org/10.1145/1806596.1806632
18. McAllester, D.A.: Automatic recognition of tractability in inference relations. J. ACM **40**(2), 284–303 (1993), https://doi.org/10.1145/151261.151265
19. de Moura, L.M., Bjørner, N.S.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) 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. Lecture Notes in Computer Science, vol. 4963, pp. 337–340. Springer (2008), https://doi.org/10.1007/978-3-540-78800-3_24
20. Nelson, G., Oppen, D.C.: Fast decision algorithms based on union and find. In: 18th Annual Symposium on Foundations of Computer Science (Providence, R.I., 1977), pp. 114–119. IEEE, Long Beach, CA (1977)
21. Oppen, D.C.: Reasoning about recursively defined data structures. J. ACM **27**(3), 403–411 (Jul 1980), https://doi.org/10.1145/322203.322204
22. Riley, D., Fedyukovich, G.: Multi-phase invariant synthesis. In: Roychoudhury, A., Cadar, C., Kim, M. (eds.) 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. pp. 607–619. ACM (2022), https://doi.org/10.1145/3540250.3549166
23. Shostak, R.E.: An algorithm for reasoning about equality. Communications of the ACM **21**(7), 583–585 (1978)
24. Sofronie-Stokkermans, V.: Hierarchic reasoning in local theory extensions. In: Nieuwenhuis, R. (ed.) Automated Deduction - CADE-20, 20th International Conference on Automated Deduction, Tallinn, Estonia, July 22-27, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3632, pp. 219–234. Springer (2005), https://doi.org/10.1007/11532231_16

A Missing Proofs

Lemma 5. *(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 let $\mathfrak{A}_\psi$ be a structure s.t. $\mathfrak{A}_\psi \models (\bigwedge \Phi \wedge \neg\psi)(\bar{a}_\psi)$ for some $\bar{a}_\psi \in A^n_\psi$. Now let $\mathfrak{B}$ be the direct product of the structures, i.e., $\mathfrak{B} = \Pi_{\psi\in\Psi} \mathfrak{A}_\psi$, and $\bar{b} = (b_1, \dots, b_n)$ be a tuple in B such that $b_i(\psi) = \bar{a}_\psi(i)$. Then by Theorem 9.1.5 in [14] we have models of Horn formulas are preserved under direct products, hence $\mathfrak{B} \models \bigwedge \Phi(\bar{b})$, thus from our assumption $\mathfrak{B} \models \bigvee \Psi(\bar{b})$, and we have $\mathfrak{B} \models \psi(\bar{b})$ for some $\psi \in \Psi$. Again by Theorem 9.1.5 $\mathfrak{A}_\psi \models \psi(\bar{a}_\psi)$. Contradiction. □

Lemma 12. *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. □

Given an egraph G, the class of a node $n \in N$, $class_G(n) := \rho_{root}(n)$, is the set of all nodes that are equivalent to n. We denote by $children(n)$ the set of nodes with an incoming edge from n. We now adopt two definitions from [8] below.

Definition 16. *(Constructively Ground Node [8]) 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) $deg(n) > 0$ 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, $class(n)$, *ground* if it contains a c-ground node.

Definition 17. *(Admissible Representative Function [8]) Given an egraph $G = \langle N, E, L, root \rangle$, a representative function $\beta : N \to N$ is admissible for G if*

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

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

As a consequence of Corollary 14 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 14 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 $G^* = \mathit{egraph}(\gamma)$. Note that as φ is satisfiable, so is γ. Below we show using Lemma 13 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 18. *Let γ and φ^* be as given above, then $\models \varphi^* \leftrightarrow \gamma$.*

Proof. We proceed in two steps, first we show $\models \varphi^* \to \gamma$ then we show $\models \gamma \to \varphi^*$. *i.* ($\models \varphi^* \to \gamma$) We show $\models \varphi^* \to \gamma_i$. 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 $\mathit{egraph}(\gamma) = G^*$, hence by Lemma 13 $\varphi \models_{\mathcal{T}} \gamma_i$. As $\mathcal{T}$ is a locally ground theory, $\Gamma_\varphi \cup \Phi \Vdash \gamma_i$, hence $\gamma_i \in D$ and $\models \varphi^* \to \gamma_i$. Otherwise, γ_i is a negative literal and similarly, γ_i is representable in $\mathit{egraph}(\gamma) = G^*$. Moreover, as φ is satisfiable, the negative literals represented in G^* are exactly the negative literals represented in $G_0 = \mathit{egraph}(\varphi)$. Which in turn are exactly the negative literals occurring in φ. Therefore, γ_i occurs in φ and thus $\models \varphi^* \to \gamma_i$.

ii. ($\models \gamma \to \varphi^*$) Similarly, we show $\models \gamma \to \theta_i$ 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 13 θ_i is representable in $G^* = \mathit{egraph}(\gamma)$, 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 $G^* = \mathit{egraph}(\gamma)$. Similarly, as γ is satisfiable θ_i occurs in γ and hence $\gamma \models \theta_i$. $\square$

We will use Lemma 18 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 [8]. 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 19. *(Lemma 13 [8]). Representative functions β computed by `find_defs` are admissible functions that are maximally ground.*

Theorem 20. *(Theorem 1 [8]). 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 21. *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 14 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 [8]. □

Theorem 15. *(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 18. First note that $\models_{\mathcal{T}} \varphi \leftrightarrow \varphi^*$, furthermore, we have by Lemma 18, $\models \gamma \leftrightarrow \varphi^*$ giving us $\models_{\mathcal{T}} \varphi \leftrightarrow \gamma$. By Lemma 19 we have β is an admissible representative function and hence by Theorem 20 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) $\mathrm{FV}(\psi) \subseteq \mathrm{FV}(\varphi)$. Trivial.
(iii) Assume $\varphi(\bar{x}) \models_{\mathcal{T}} x_i \approx t$ for some ground term t, then by Lemma 19 we have β is an admissible representative function that is maximally ground. Hence, by Lemma 21 we have the representative selected $\beta(n_{x_i})$ is c-ground and `ntt(`$\beta(n_x)$`)` is ground. Thus, x_i is successfully eliminated from ψ, the output of G^*`.to_formula`.

□