

Finite Quantification in Hierarchic Theorem Proving

Peter Baumgartner¹ Joshua Bax¹ Uwe Waldmann²

¹ NICTA* and Australian National University, Canberra, Australia
Peter.Baumgartner@nicta.com.au

² MPI für Informatik, Saarbrücken, Germany
uwe@mpi-inf.mpg.de

Abstract. Many applications of automated deduction require reasoning in first-order logic modulo background theories, in particular some form of integer arithmetic. A major unsolved research challenge is to design theorem provers that are “reasonably complete” even in the presence of free function symbols ranging into a background theory sort. In this paper we consider the case when all variables occurring below such function symbols are quantified over a finite subset of their domains. We present a non-naive decision procedure for background theories extended this way on top of black-box decision procedures for the EA-fragment of the background theory. In its core, it employs a *model-guided* instantiation strategy for obtaining pure background formulas that are equi-satisfiable with the original formula. Unlike traditional finite model finders, it avoids exhaustive instantiation and, hence, is expected to scale better with the size of the domains. Our main results in this paper are a correctness proof and first experimental results.

1 Introduction

Many applications of automated deduction require reasoning in first-order logic modulo background theories, in particular some form of integer arithmetic. A major unsolved research challenge is to design theorem provers that are “reasonably complete” for quantified formulas, in particular in presence of free function symbols ranging into a background theory sort (“free BG-sorted operators”, for short). Such formulas arise frequently when reasoning on data structures with specific properties, e.g., *symmetric* arrays over integers and *sorted* lists over integers. Modelling such data structures is easy when full quantification and free integer-sorted function symbols are available to axiomatize the array access function and the list head function respectively.

Unfortunately, (refutationally) complete theorem proving in the presence of free BG-sorted operators is intractable in general. For instance, just adding one free predicate symbol to linear integer arithmetic results in a Π_1^1 -hard validity problem [13]. Theorem proving approaches hence have to circumvent this problem in one way or the other. On the one hand, SMT-solvers [19] generally use instantiation heuristics [11,17] for reducing the input problem to a quantifier-free one, and these are complete only in rather restricted cases [12]. On the other hand, approaches rooted in first-order theorem proving either are incomplete; do not accept free BG-sorted operators at all [14,22,10,5,6] or, are complete only for certain fragments or under certain conditions [3,1,15,7,8].

* NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

In practice, lack of completeness is a major concern in, e.g., software verification applications, which frequently require disproving non-valid proof obligations. In such cases, incomplete theorem provers run out of resources or report “unknown” instead of detecting counter-satisfiability. We address this problem by working with quantification over *finite* segments of the background sorts, e.g., the integers. Our underlying methodology assumes that from a user’s point of view, data structures over the integers can often be supplanted by data structures over reasonably large finite segments of the integers, say, from $-\text{Maxint}$ to $+\text{Maxint}$, as good-enough approximations. As no other restrictions apply our method should be widely applicable in practice. Moreover, our method is also refutationally sound wrt. the standard semantics. That is, whenever our algorithm determines unsatisfiability wrt. finite domains, the given clause set is also unsatisfiable wrt. unbounded domains. Because of that, our approach can be seen as an extension of current quantifier instantiation heuristics by being able to determine satisfiability wrt. finite domains.

If all quantifiers range over finite domains, decidability can be recovered in a trivial way by exhaustive instantiation and calling a suitable SMT-solver afterwards. Of course, this naive approach does not scale with the domain size and cannot be expected to work well in practice. This problem has often been observed in the context of finite-model finding [23,24,16,9,4,21,20]. While our method is also based on instantiation, it is (often) far less prolific than the naive method.

More precisely, our method accepts as input a set of *finitely quantified clauses*. A clause is finitely quantified if every variable occurring below a free BG-sorted operator is quantified over a finite segment of its domain. The core idea is to give the free BG-sorted operators a *default* interpretation that is then stepwise refined. This default interpretation maps every free BG-sorted operator to a constant function, and refinements are done by finding exceptions to that in a conflict-driven way. After each refinement the given clause set is transformed into a certain form whose satisfiability can be decided by existing reasoners in a black-box fashion. Suitable reasoners are, e.g., theorem provers implementing hierarchic superposition [3,8] and, with one more simple transformation step, SMT-solvers for the EA-fragment of the background theory. The procedure stops after finitely many (hopefully few) refinement steps, either with a representation of a model or a set of ground instances obtained from exceptions which demonstrates the unsatisfiability of the given clause set.

In the following we preview our method with an example. Let N be the following set of finitely quantified clauses:

- $$\begin{array}{ll}
 (1) \quad \text{read}(\text{write}(a, i, x), i) \approx x & (4) \quad 1 \leq m \wedge m < 1000 \\
 (2) \quad \text{read}(\text{write}(a, i, x), j) \approx \text{read}(a, j) \vee i \approx j & (5) \quad \text{read}(a, m) < \text{read}(a, m + 1) \\
 (3) \quad \text{read}(a, i) \leq \text{read}(a, j) \vee \neg(i < j) \vee i \notin [1..1000^i] \vee j \notin [1..1000^j]
 \end{array}$$

where $t \in [l..h]$ abbreviates the formula $l \leq t \wedge t \leq h$ for any integer-sorted terms t, l and h . Variables are typeset in italics, e.g, x , and operators in sans-serif, e.g., read , a and m . The axioms (1) and (2) are the standard axioms for integer-sorted arrays with integer indices. The axiom (3) states that the array a is sorted within the domain $[1..1000]$ for i and j . Annotating the upper bounds as 1000^i and 1000^j facilitates replacing them with different values for a given variable, see below. The clauses (4) introduce an integer constant m within range as stated. The task is to check whether N is satisfiable wrt. the expected semantics, which it is.

In order to check satisfiability with hierarchic superposition the input clause set has to be *sufficiently complete* (cf. Section 2). In the example, sufficient completeness means that every ground `read`-term must be provably equal in pure first-order logic to some background term. With respect to axioms (1) and (2) there is nothing to do and we exclude them from the transformation. The clauses (3) and (5) constrain the interpretation of terms of the form `read(a, t)` but do not enforce sufficient completeness. Achieving sufficient completeness for *ground* clauses like (5) is easy, one just needs to add “definitions” like (5b) `read(a, m) ≈ n0` and (5c) `read(a, m + 1) ≈ n1` where n_0 and n_1 are fresh integer-sorted parameters (symbolic constants) and replace the clause (5) by (5a) $n_0 < n_1$. Indeed, our transformation does all that. (Our earlier calculus in [8] also includes such a transformation.)

The more difficult part concerns the non-ground clause (3). Our procedure generalizes the above mechanism of introducing definitions and applying them to the non-ground case. For that, it uses a candidate model which initially is the *default interpretation* that maps all `read`-terms of a particular shape to the *same* arbitrary symbolic constant. This results in the following transformation of clause (3):

$$\begin{aligned} (3a) \quad & n_3 \leq n_4 \vee \neg(i < j) \vee i \notin [1..1000^i] \vee j \notin [1..1000^j] \\ (3b) \quad & \text{read}(a, i) \approx n_3 \vee i \notin [1..1000^i] \quad (3c) \quad \text{read}(a, j) \approx n_4 \vee j \notin [1..1000^j] \end{aligned}$$

Clauses (3b) and (3c) are the definitions for the default interpretation, one per occurrence of a `read`-term in (3), and clause (3a) is clause (3) after applying these definitions.

The new clause set $N_1 = \{(1), (2), (3a)–(3c), (4), (5a)–(5c)\}$ now needs to be checked for satisfiability. Because the clause set N_1 is sufficiently complete and hierarchic superposition decides the underlying fragment, we get a definite result.

The clause set N_1 is in fact unsatisfiable. Because this only means that N is not satisfied using the current model candidate, the search for a model needs to continue. This is done by refining the default interpretation at a critical point that is responsible for unsatisfiability. In terms of the example the algorithm searches for sub-intervals of the domains for i and j that result in a satisfiable set, and it finds a point outside some such sub-interval that when added back gives unsatisfiability again. Intuitively, the sub-intervals represent points that do *not* need fixing to get a model, while the point outside is offending and needs to be fixed. Concretely, it finds the sub-intervals obtained by replacing 1000^j by 0^j and 1000^i by 999^i everywhere in N_1 . Let N_2 be the resulting (satisfiable) set. The point identified by our algorithm will be the number 1000 for the variable i . That is, replacing 999^i by 1000^i in N_2 gives an unsatisfiable set again.

These tests suggest excluding the point 1000 for i from the default interpretation (which we write as $[1..1000^i] \setminus \{1000\}$) and that we should provide a separate definition at the point 1000. The corresponding transformation of clause (3) hence looks as follows:

$$\begin{aligned} (3a1) \quad & n_{31} \leq n_4 \vee \neg(i < j) \vee i \notin [1..1000^i] \setminus \{1000\} \vee j \notin [1..1000^j] \\ (3a2) \quad & n_{32} \leq n_4 \vee \neg(1000 < j) \vee j \notin [1..1000^j] \\ (3b1) \quad & \text{read}(a, i) \approx n_{31} \vee i \notin [1..1000^i] \setminus \{1000\} \quad (3c) \quad \text{read}(a, j) \approx n_4 \vee j \notin [1..1000^j] \\ (3b2) \quad & \text{read}(a, 1000) \approx n_{32} \end{aligned}$$

Clauses (3b1) and (3b2) provide the modified definitions and clauses (3a1) and (3a2) are the correspondingly rewritten versions of (3). Let $N_3 = \{(1), (2), (3a1)–(3c), (4), (5a)–(5c)\}$ be the result of the current transformation step.

The clause set N_3 is still unsatisfiable. In the next round, the new upper bounds required for the clauses in N_3 to have satisfiability are 999^j and 1000^i . Transforming clause (3) wrt. the points 1000 for j and 1000 for i from the previous step gives:

$$\begin{aligned}
(3a1) \quad n_{31} &\leq n_{41} \vee \neg(i < j) \vee i \notin [1..1000^i] \setminus \{1000\} \vee j \notin [1..1000^j] \setminus \{1000\} \\
(3a2) \quad n_{32} &\leq n_{41} \vee \neg(1000 < j) \vee j \notin [1..1000^j] \setminus \{1000\} \\
(3a3) \quad n_{31} &\leq n_{42} \vee \neg(i < 1000) \vee i \notin [1..1000^i] \setminus \{1000\} \\
(3a4) \quad n_{32} &\leq n_{42} \vee \neg(1000 < 1000) \\
(3b1) \quad \text{read}(a, i) &\approx n_{31} \vee i \notin [1..1000^i] \setminus \{1000\} & (3b2) \quad \text{read}(a, 1000) &\approx n_{32} \\
(3c1) \quad \text{read}(a, j) &\approx n_{41} \vee j \notin [1..1000^j] \setminus \{1000\} & (3c2) \quad \text{read}(a, 1000) &\approx n_{42}
\end{aligned}$$

Let $N_4 = \{(1), (2), (3a1) - (3c2), (4), (5a) - (5c)\}$ be the result of the current transformation step. This time, N_4 is satisfiable, and so is N , with the same models. If I is any such model we have $I(m) = 999$, $I(\text{read}(a, i)) = k$, for some integer k and all $i = 1..999$, and $I(\text{read}(a, 1000)) = l$ for some integer $l > k$.

The whole example is solved after two iterations of transformation steps. In general, each transformation step needs $O(m \cdot \log(n))$ theorem prover calls to determine the sub-intervals and the next point as explained above, where m is the number of variables in the given clause set after making clauses variable-disjoint and n is the size of the largest domain. In total, with $m = 2$ and $n = 1000$ this accounts for $2 \cdot (m \cdot \log(n)) \leq 40$ theorem prover calls, however each one rather simple and quickly executed. By contrast, the full ground instantiation of the clauses (3)-(5) has a size of $n^m = 10^6$ which, in general, grows too quickly for current theorem provers or SMT solvers.

Related Work. Related work comes from several directions. Procedures for computing models of first-order logic formulas *without background theories* have a long tradition in automated reasoning. MACE-style model finding [9] utilizes translation into propositional SAT or into EPR [4] for deciding satisfiability wrt. a given candidate domain size k ; SEM-style model finding [23,24,16] utilizes constraint solving techniques, again wrt. k . The main problem is scalability wrt. both the domain size k and the number of variables in the input clause set, which severely limits the applicability of both styles in practice. Recently, Reynolds, Tinelli, Goel, Krstić, Deters and Barrett proposed a finite model finding procedure in the SMT framework that addresses this problem by on-demand instantiation techniques [21,20]. This way, their work is conceptually somewhat related to ours, but, unlike ours, they allow quantification only over variables ranging into the free sort. An extension for quantifying variables over background domains such as the integers does not seem straightforward and is left as future work in [21].

Heuristic instantiation is the state of the art technique for handling quantified formulas in SMT-solvers [11,17]. These heuristics perform impressively well in practice, but in general are incomplete even for pure first-order logic. Ge and deMoura [12] propose a technique where the ground terms used for instantiation come from solving certain set constraints. They obtain completeness results for the fragment where every variable occurs only as an argument of a free function or predicate symbol. Interestingly, they also use the notion of a default interpretation in a similar way as we do. However, even with certain extensions their approach remains incomparable to ours. For example, terms like $f(x + y)$ are disallowed, but are acceptable in our approach when x and y are finitely quantified.

Regarding first-order theorem proving, Weidenbach and Kruglow [15] consider the case when all background-sorted terms are ground, similarly to our calculus in [8]. In [7] we have identified a certain syntactic fragment that enables complete reasoning.

2 Hierarchic Theorem Proving

Hierarchic superposition [3,8] is a calculus for automated reasoning in a hierarchic combination of first-order logic and some background theory, for instance some form of arithmetic. We consider the following scenario:³

We assume that we have a *background (“BG”)* prover that accepts as input a set of clauses over a *BG signature* $\Sigma_B = (\mathcal{E}_B, \mathcal{Q}_B)$, where \mathcal{E}_B is a set of *BG sorts* and \mathcal{Q}_B is a set of *BG operators*. Terms/clauses over Σ_B and BG-sorted variables are called *BG terms/clauses*. The BG prover decides the satisfiability of Σ_B -clause sets w. r. t. a *BG specification*, that is, a class of term-generated Σ_B -interpretations (called *BG models*) that is closed under isomorphisms. We assume that \mathcal{Q}_B contains a set of distinguished constant symbols $\mathcal{Q}_B^D \subseteq \mathcal{Q}_B$ that has the property that any two distinct $d_1, d_2 \in \mathcal{Q}_B^D$ are interpreted by different elements in every BG model. We refer to these constant symbols as *(BG) domain elements*. We also assume that Σ_B contains an infinite number of *parameters*, that is, additional constant symbols that may be interpreted freely by arbitrary elements of the appropriate domain. In examples we use $\{0, 1, 2, \dots\}$ to denote BG domain elements, $\{+, -, <, \leq\}$ to denote (non-parameter) BG operators, and the possibly subscripted letters $\{x, y\}$ and $\{\alpha, \beta\}$ to denote variables and parameters, respectively. We assume that the BG specification is the class of all models of linear integer arithmetic (LIA).

For technical reasons, we assume that equality is the only predicate symbol in our language and that any non-equational atom $p(t_1, \dots, t_n)$ is encoded as an equation $p(t_1, \dots, t_n) \approx true_p$. When we write, say, $x \leq y$, this should always be taken as a shorthand for an equation as above.

The *foreground (“FG”)* theorem prover accepts as input clauses over a signature $\Sigma = (\mathcal{E}, \mathcal{Q})$, where $\mathcal{E}_B \subseteq \mathcal{E}$ and $\mathcal{Q}_B \subseteq \mathcal{Q}$. The sorts in $\mathcal{E}_F = \mathcal{E} \setminus \mathcal{E}_B$ and the operator symbols in $\mathcal{Q}_F = \mathcal{Q} \setminus \mathcal{Q}_B$ are called *FG sorts* and *FG operators*. We use $\{a, b, c, f, g\}$ to denote FG operators. We call a Σ -term a *FG term*, if it is not a BG term, that is, if it contains at least one FG operator or FG variable (and analogously for equations, literals, or clauses). We emphasize that for a FG operator $f : \xi_1 \dots \xi_n \rightarrow \xi_0$ in \mathcal{Q}_F any of the ξ_i may be a BG sort, and that consequently FG terms may have BG sorts. Every FG operator f with a BG range sort $\xi_0 \in \mathcal{E}_B$ is called a *free BG-sorted (FG) operator*.

After abstracting out certain BG terms that occur as subterms of FG terms,⁴ the FG prover saturates the set of Σ -clauses using the inference rules of hierarchic superposition, such as, e. g.,

$$\text{Negative superposition} \quad \frac{l \approx r \vee C \quad s[u] \not\approx t \vee D}{\text{abstr}((s[r] \not\approx t \vee C \vee D)\sigma)}$$

³ Due to a lack of space, we can only give a brief overview of the calculus and of the semantics of hierarchic specifications. We refer to [8] for the details.

⁴ *Abstracting out* a term t that occurs in a clause $C[t]$ means replacing $C[t]$ by $x \not\approx t \vee C[x]$ for a new variable x .

if (i) neither l nor u is a BG term, (ii) u is not a variable, (iii) σ is an mgu of l and u , (iv) σ maps all BG variables to BG terms, (v) $r\sigma \not\approx l\sigma$, (vi) $(l \approx r)\sigma$ is strictly maximal in $(l \approx r \vee C)\sigma$, (vii) the first premise does not have selected literals, (viii) $t\sigma \not\approx s\sigma$, and (ix) if the second premise has selected literals, then $s \not\approx t$ is selected in the second premise, otherwise $(s \not\approx t)\sigma$ is maximal in $(s \not\approx t \vee D)\sigma$.

These differ from the standard superposition inference rules [2] mainly in that only the FG parts of clauses are overlapped and that any BG clauses derived during the saturation are instead passed to the BG prover. The BG prover implements an inference rule

$$\text{Close} \frac{C_1 \quad \cdots \quad C_n}{\square} \quad \text{if } C_1, \dots, C_n \text{ are BG clauses and } \{C_1, \dots, C_n\} \text{ is unsatisfiable w. r. t. the BG specification.}$$

As soon as one of the two provers detects a contradiction, the input clause set has been shown to be unsatisfiable w. r. t. *conservative extensions of the BG specification*, i. e., Σ -interpretations whose restriction to Σ_B is a model of the BG specification. Below we refer to satisfiability in this sense as *B-satisfiability*.

There are two requirements for the refutational completeness of hierarchic superposition. The first one is a variant of *sufficient completeness*: We must be able to prove that every ground BG-sorted FG term is equal to some BG term – or at least every ground BG-sorted FG term that occurs in some non-redundant ground instance of an input clause. Sufficient completeness of a set of Σ -clauses is a property that is not even recursively enumerable. For certain classes of Σ -clause sets, however, it is possible to establish sufficient completeness automatically [15,8]: If all BG-sorted FG terms are ground, it suffices to add a *definition* $\alpha_t \approx t$ for every BG-sorted FG term t occurring in a clause $C[t]$, where α_t is a new parameter (BG constant); afterwards $C[t]$ can be replaced by $C[\alpha_t]$.

Since we can only pass *finite* clause sets to a BG prover, there is a second requirement for refutational completeness, namely the compactness of the BG specification. A specification is called *compact*, if every set of formulas that is unsatisfiable w. r. t. the specification has a finite unsatisfiable subset.

3 Finite Domain Transformation

We are interested in refutationally complete hierarchic theorem proving in the presence of free BG-sorted FG operators. Unfortunately, just adding one free predicate symbol to linear integer arithmetic results in a Π_1^1 -hard validity problem. To circumvent this problem, we work with a modified semantics and introduce a concept of finite quantification of BG variables. This allows us to remove all free BG-sorted FG operators by a *finite domain transformation*, introduced next, and use existing reasoning methods as decision procedures on the result.

Let $\xi \in \Xi_B$ be a BG sort. By a *finite ξ -domain* Δ we mean any possibly empty finite set $\{d_1, \dots, d_n\} \subseteq \Omega_B^D$ of ξ -sorted domain elements d_i . Set membership in Δ can be expressed by a BG formula $\mathcal{F}_\Delta[x]$ in one free ξ -sorted variable x whose extension is exactly the set Δ , in every \mathcal{B} -interpretation. One can always take $\mathcal{F}_\Delta[x] = x \approx d_1 \vee \dots \vee$

$x \approx d_n$, but if supported by the BG logic, as in the case of integer arithmetic, it may be advantageous to use “compact” representations like $\mathcal{F}_\Delta[x] = 1 \leq x \wedge x \leq 20$ instead.

We use set-theoretic expressions for finite ξ -domains, in particular of the form $\Delta \setminus \Gamma$, where Γ is a finite set of domain elements of the proper sort. In the previous example, e.g., $\mathcal{F}_{\Delta \setminus \{3,5\}}[x] = 1 \leq x \wedge x \leq 20 \wedge x \neq 3 \wedge x \neq 5$. Instead of $\mathcal{F}_\Delta[x]$ and $\mathcal{F}_{\Delta \setminus \Gamma}[x]$ we generally write $x \in \Delta$ and $x \in \Delta \setminus \Gamma$, respectively, and $x \notin \Delta$ and $x \notin \Delta \setminus \Gamma$ for their negations. We call these expressions *domain predicates* and treat them as literals in clauses instead of expanding them.

Definition 3.1. A *finitely quantified clause* is a Σ -clause of the form $D \vee x_1 \notin \Delta_{x_1} \vee \dots \vee x_n \notin \Delta_{x_n}$ such that D does not contain domain predicates, $n \geq 0$, $x_i \neq x_j$ for $1 \leq i < j \leq n$, and every variable occurring below a free BG-sorted operator in D is among x_1, \dots, x_n .

For example, $f(x+1) > \alpha + y \vee y > 0 \vee x \notin [1..1000]$ is finitely quantified.

Example 3.2. Let N consist of the following two finitely quantified clauses:

- (C₁) $f(x_1) > x_1 \vee x_1 \notin [1..1000]$
- (C₂) $f(x_2 + 3) < 10 \vee \neg(x_2 > 2) \vee x_2 \notin [1..1000]$

We have formally $\Delta_{x_1} = \Delta_{x_2} = [1..1000]$, and in C_1 the pseudo-literal $x_1 \notin [1..1000]$ is short for $\neg(1 \leq x_1 \leq 1000)$. □

Where $\mathbf{x} = (x_1, \dots, x_n)$ let $\Delta_{\mathbf{x}}$ denote the \mathbf{x} -indexed list $(\Delta_{x_1}, \dots, \Delta_{x_n})$ of sets of domain elements. We extend usual set operations pointwise to \mathbf{x} -indexed lists $\Pi_{\mathbf{x}}$ and $\Delta_{\mathbf{x}}$ of sets of domain elements. For instance $\Pi_{\mathbf{x}} \subseteq \Delta_{\mathbf{x}}$ iff $\Pi_x \subseteq \Delta_x$, for each $x \in \mathbf{x}$.

We are going to define the earlier mentioned finite domain transformation for evaluating finitely quantified clauses under a given interpretation. It takes as input a finitely quantified clause $C[\Delta_{\mathbf{x}}]$ and sets of points $\Pi_{\mathbf{x}}$ that provide possible exceptions to interpreting the free BG-sorted operators as the constant function on the domains $\Delta_{\mathbf{x}}$ as specified by the default interpretation.

Definition 3.3 (Finite Domain Transformation). Let $C[\Delta_{\mathbf{x}}] = D \vee x_1 \notin \Delta_{x_1} \vee \dots \vee x_n \notin \Delta_{x_n}$ be a finitely quantified clause and $\Pi_{\mathbf{x}} \subseteq \Delta_{\mathbf{x}}$ a list of sets of domain elements.

Let $\text{Cls}_C := \emptyset$ and $\text{Def}_C := \emptyset$ be initially empty sets of Σ -clauses. For every partition $\{y_1, \dots, y_k\} \uplus \{z_1, \dots, z_l\}$ of $\{x_1, \dots, x_n\}$ do the following:

For all substitutions $\gamma = [z_1 \mapsto d_1, \dots, z_l \mapsto d_l]$ such that $d_m \in \Pi_{z_m}$:

1. Let $E := D\gamma$
2. While E has the form $E[t]$ where t is a minimal term with a free BG-sorted operator at the top-level do the following:
 - (a) Let α be a fresh parameter
 - (b) Add to Def_C the clause $t \approx \alpha \vee y_1 \notin \Delta_{y_1} \setminus \Pi_{y_1} \vee \dots \vee y_k \notin \Delta_{y_k} \setminus \Pi_{y_k}$
 - (c) Set $E := E[\alpha]$
3. Add to Cls_C the clause $E \vee y_1 \notin \Delta_{y_1} \setminus \Pi_{y_1} \vee \dots \vee y_k \notin \Delta_{y_k} \setminus \Pi_{y_k}$

The result is the pair $FD(C, \Pi_{\mathbf{x}}) = (\text{Cls}_C, \text{Def}_C)$, the *finite domain transformation of C* .

By the minimality of t in (2) we mean that no proper subterm of t is built with a free BG-sorted operator. The finite domain transformation removes from the given finitely quantified clause C every occurrence of a term t built with some free BG-sorted symbol. Recall from Definition 3.1 that all variables in t are among $\mathbf{x} = (x_1, \dots, x_n)$. The removal of t distinguishes whether x_i is interpreted as an element of $\Delta_i \setminus \Pi_i$ or as an element $d_i \in \Pi_i$. This is done in all possible ways by exhaustive partitioning of the variables \mathbf{x} and exhausting the substitution γ for all possible assignments for x_i . The set $\Delta_i \setminus \Pi_i$ specifies those domain elements for which the interpretation of t is undistinguished, and the set Π_i specifies those domain elements for which the interpretation of t is distinguished, by taking different parameters α per substitution γ . In step (b) corresponding definitions for t are put into Def_C . Step (c) applies these definitions to the current clause E .

Example 3.2 (continued). Let $\Pi_{(x_1)} = (\{9\})$. Then $\text{FD}(C_1, \Pi_{(x_1)})$ consists of the clauses

$$\begin{array}{ll} (C_{11}) & \alpha_1 > x_1 \vee x_1 \notin [1..1000] \setminus \{9\} & (C_{13}) & \alpha_2 > 9 \\ (C_{12}) & f(x_1) \approx \alpha_1 \vee x_1 \notin [1..1000] \setminus \{9\} & (C_{14}) & f(9) \approx \alpha_2 \end{array}$$

where $\text{Cls}_{C_1} = \{C_{11}, C_{13}\}$ and $\text{Def}_{C_1} = \{C_{12}, C_{14}\}$. The left clauses stem from partitioning $\{x_1\}$ as $\{y_1\} \uplus \emptyset$, and the right clauses from $\emptyset \uplus \{z_1\}$. There are two occurrences of $\Delta_{x_1} = [1..1000]$. \square

There are no restrictions on nesting free BG-sorted operators, although none of our examples shows that. For example, a literal like $f(x + g(y, \beta)) > f(y) + y$ is perfectly acceptable. The possible nesting of free BG-sorted operators necessitates the while-loop in step (2) in Definition 3.3; removing all of them in a single step is not possible.

The sets of domain elements Δ_x occurring in clauses in $\text{FD}(C, \Pi_x)$ are all within pseudo-literals of the form $x \notin \Delta_x \setminus \Pi_x$. Hence, both Cls_C and Def_C are of the form $\text{Cls}_C[\Delta_x]$ and $\text{Def}_C[\Delta_x]$. Moreover, in $\text{FD}(C, \Pi_x)$, every free BG-sorted operator f occurs only in a clause of the form $f(t_1, \dots, t_n) \approx \alpha \vee D$ in Def_C where no t_i and no literal in D contains any free BG-sorted operator.

The finite domain transformation is generalized to clause sets by taking the union of the finite domain transformations applied to its members. More precisely, let $N = \{C_1[\Delta_{x_1}], \dots, C_m[\Delta_{x_m}]\}$ be a finite set of finitely quantified clauses. Let us assume the clauses in N have been renamed apart, so that the lists of variables \mathbf{x}_i are pairwise disjoint, for all $i = 1..m$. By definition, each \mathbf{x}_i consists of pairwise different variables, too. This allows us to take \mathbf{x} as the concatenation of all \mathbf{x}_i 's and to write Δ_x for the concatenation of all Δ_{x_m} 's. The clause set N hence is of the form $N[\Delta_x]$. Now let $(\text{Cls}_{C_i}, \text{Def}_{C_i}) = \text{FD}(C_i, \Pi_x)$ and define $\text{FD}(N, \Pi_x) = (\text{Cls}_N, \text{Def}_N)$ where $\text{Cls}_N = \bigcup_{i=1..m} \text{Cls}_{C_i}$ and $\text{Def}_N = \bigcup_{i=1..m} \text{Def}_{C_i}$. Below, we usually denote $\text{FD}(N, \Pi_x)$ as a single clause set $M[\Delta_x] = \text{Cls}_N \cup \text{Def}_N$. The following result follows immediately:

Proposition 3.5. *Let $N[\Delta_x]$ be a set of finitely quantified clauses and $\Pi_x \subseteq \Delta_x$. Then $\text{FD}(N, \Pi_x)$ is sufficiently complete.*

Proposition 3.5 is one of the ingredients that allows us to argue for hierarchic superposition [8] as a decision procedure for \mathcal{B} -satisfiability of the clause sets $\text{FD}(C, \Pi_x)$. We also need a termination argument for derivations (compactness, cf. Section 2, is unproblematic then). This is easy, for instance, in the absence of non-ground FG-sorted operators only finitely many superposition steps exist and all of these are between the clauses in

Def_C , and then only at the top-level- that is, between the literals $f(t_1, \dots, t_n) \approx \alpha$. Alternatively one can use SMT-solvers after removing all free BG-operators by exhaustive application of a superposition-like inference rule that from premises $f(t_1, \dots, t_n) \approx \alpha \vee D$ and $f(s_1, \dots, s_n) \approx \beta \vee E$ derives the clause $s_1 \neq t_1 \vee \dots \vee s_n \neq t_n \vee \alpha \approx \beta \vee D \vee E$. In general, hierarchic superposition can be used if it is guaranteed to terminate on $\text{FD}(C, \Pi_{\mathbf{x}})$. This applies, e.g., to the example in the introduction.

The notation $M[\Delta_{\mathbf{x}}]$ makes it easy to modify the sets Δ_x in pseudo-literals in clauses in M . More precisely, if $\Delta_{\mathbf{x}} = (\dots, \Delta_x, \dots)$ for some $x \in \mathbf{x}$ and Γ is a set of domain elements with the same sort as x , we denote by $\Delta_{\mathbf{x}}[x \mapsto \Gamma]$ the update of $\Delta_{\mathbf{x}}$ at index x by Γ , i.e., the list (\dots, Γ_x, \dots) . Correspondingly, $C[\Delta_{\mathbf{x}}[x \mapsto \Gamma]]$ is the clause that is obtained from $C[\Delta_{\mathbf{x}}]$ by replacing Δ_x by Γ_x everywhere. For clause sets $N[\Delta_{\mathbf{x}}]$ we define $N[\Delta_{\mathbf{x}}[x \mapsto \Gamma]]$ analogously.

Example 3.2 (continued). The clause set N is of the form $N[\Delta_{\mathbf{x}}]$ where $\mathbf{x} = (x_1, x_2)$ and $\Delta_{x_1} = \Delta_{x_2} = [1..1000]$. Now let $\Pi_{\mathbf{x}} = (\{9\}, \{6\})$. Then $M[\Pi_{\mathbf{x}}] = \text{FD}(N, \Pi_{\mathbf{x}}) = (\text{Cls}_{C_1} \cup \text{Cls}_{C_2}) \cup (\text{Def}_{C_1} \cup \text{Def}_{C_2})$ where $\text{Cls}_{C_2} = \{C_{21}, C_{23}\}$, $\text{Def}_{C_2} = \{C_{22}, C_{24}\}$ and

$$\begin{array}{ll} (C_{21}) & \alpha_3 < 10 \vee \neg(x_2 > 2) \vee x_2 \notin [1..1000] \setminus \{6\} & (C_{23}) & \alpha_4 < 10 \vee \neg(6 > 2) \\ (C_{22}) & f(x_2 + 3) \approx \alpha_3 \vee x_2 \notin [1..1000] \setminus \{6\} & (C_{24}) & f(6 + 3) \approx \alpha_4 \end{array}$$

The clause set $M[\Delta_{\mathbf{x}}[x_2 \mapsto \emptyset]] = M[(\{9\}, \emptyset)]$ is obtained by replacing the two occurrences of $\Delta_{x_2} = [1..1000]$ in C_{21} and C_{22} by the empty interval $[\]$. \square

We conclude this section with some lemmas that will be needed in the proof of the main correctness result, Theorem 4.2 below. In each of them, $N[\Delta_{\mathbf{x}}]$ is a set of finitely quantified clauses, $\Pi_{\mathbf{x}} \subseteq \Delta_{\mathbf{x}}$, $(\text{Cls}_N, \text{Def}_N) = \text{FD}(N, \Pi_{\mathbf{x}})$, and $M = \text{Cls}_N \cup \text{Def}_N$.

Lemma 3.7. $\text{Cls}_N \cup \text{Def}_N$ is \mathcal{B} -satisfiable iff $N \cup \text{Def}_N$ is \mathcal{B} -satisfiable.

Proof. For the if-direction assume that $N \cup \text{Def}_N$ is \mathcal{B} -satisfiable. It suffices to show that $N \cup \text{Cls}_N \cup \text{Def}_N$ is \mathcal{B} -satisfiable. Observe that all clauses in Cls_N can be seen to be obtained by paramodulation inferences from clauses in $N \cup \text{Def}_N$, which are all logical consequences of $N \cup \text{Def}_N$.

For the only-if direction assume that $\text{Cls}_N \cup \text{Def}_N$ is \mathcal{B} -satisfiable. The definitions in Def_N are exhaustive in the sense that any instance C of a finitely quantified clause in N obtained by ground instantiation with domain elements is congruent with some clause in Cls_N obtained by paramodulation with clauses in Def_N . This entails that $N \cup \text{Cls}_N \cup \text{Def}_N$ is \mathcal{B} -satisfiable, and hence so is $N \cup \text{Def}_N$. \square

Lemma 3.8. If $M[\emptyset_{\mathbf{x}}]$ is \mathcal{B} -unsatisfiable then N and N' are \mathcal{B} -unsatisfiable, where N' is obtained from N by removing from all clauses all domain predicates.

Proof. Assume that $M[\emptyset_{\mathbf{x}}]$ is \mathcal{B} -unsatisfiable. Every clause in $M[\Delta_{\mathbf{x}}]$ that contains a pseudo-literal of the form $x \notin \Delta_x \setminus \Pi_x$, for some $x \in \mathbf{x}$, becomes a tautology in $M[\emptyset_{\mathbf{x}}]$ after replacing $x \notin \Delta_x \setminus \Pi_x$ by $x \notin \emptyset \setminus \Pi_x$. Deleting all these tautologies leaves us with a (\mathcal{B} -unsatisfiable) set $M' \subseteq M[\emptyset_{\mathbf{x}}]$. All clauses in M' are either ground definitions in Def_N of the form $t \approx \alpha$ (cf. Definition 3.3), or clauses in Cls_N that are obtained by (repeated) paramodulation of the sub-clause D of a clause $C \in N$ (cf. again Definition 3.3) such that all instantiated domain predicates in the instance $C\gamma$ are satisfied. Clearly, adding such definitions to N preserves \mathcal{B} -satisfiability. The \mathcal{B} -unsatisfiability of both N and N' then follows from the soundness of paramodulation. \square

Lemma 3.9. *Let Γ_x be a vector of sets of domain elements of the proper sorts. For every $x \in \mathbf{x}$ and $d \in \Pi_x$, if $M[\Gamma_x]$ is \mathcal{B} -satisfiable then $M[\Gamma_x[x \mapsto \Gamma_x \cup \{d}]]$ is \mathcal{B} -satisfiable.*

Proof. All occurrences of Γ_x in clauses in $M[\Gamma_x]$ are within pseudo-literals of the form $x \notin \Gamma_x \setminus \Pi_x$. We are given $d \in \Pi_x$. It follows trivially that $\Gamma_x \setminus \Pi_x$ and $(\Gamma_x \cup \{d\}) \setminus \Pi_x$ are the same sets, which immediately entails the claim. \square

Example 3.2 (continued). Let $M[\Delta_{(x_1)}] = \text{FD}(C_1, \Pi_{(x_1)})$ from above. Let $\Gamma_{(x_1)} = ([5..500])$ and $d = 9$. Then $M[\Gamma_{(x_1)}[x_1 \mapsto \Gamma_{x_1} \cup \{d}]]$ consists of the clauses

$$\begin{array}{ll} (C'_{11}) & \alpha_1 > x_1 \vee x_1 \notin ([5..500] \cup \{9\}) \setminus \{9\} & (C_{13}) & \alpha_2 > 9 \\ (C'_{12}) & f(x_1) \approx \alpha_1 \vee x_1 \notin ([5..500] \cup \{9\}) \setminus \{9\} & (C_{14}) & f(9) \approx \alpha_2 \end{array}$$

Lemma 3.9 requires $d \in \Pi_x$. Adding d to Γ_x does not change anything, as d is again removed from $\Gamma_x \cup \{d\}$: the sets $([5..500] \cup \{9\}) \setminus \{9\}$ and $[5..500] \setminus \{9\}$ are the same. \square

4 Checking Satisfiability

Next we define a procedure `checkSAT` for checking the \mathcal{B} -satisfiability of sets of finitely quantified clauses. It repeatedly applies the finite domain transformation wrt. growing sets of exception points. It stops if a transformed set has been found that is either \mathcal{B} -satisfiable or serves to demonstrate \mathcal{B} -unsatisfiability.

```

1 algorithm checkSAT( $N[\Delta_x]$ )
2 // returns " $\mathcal{B}$ -satisfiable" or " $\mathcal{B}$ -unsatisfiable"
3 var  $\Pi_x := \emptyset_x$  // The current set of exceptions
4 while true {
5   let  $M = \text{FD}(N, \Pi_x)$ 
6   if  $M$  is  $\mathcal{B}$ -satisfiable
7     return " $\mathcal{B}$ -satisfiable" // justified by Lemma 3.7
8   else if  $M[\emptyset_x]$  is  $\mathcal{B}$ -unsatisfiable
9     return " $\mathcal{B}$ -unsatisfiable" // justified by Lemma 3.8
10  else {
11    let  $(x, d) = \text{find}(M)$ 
12     $\Pi_x := \Pi_x[x \mapsto \Pi_x \cup \{d\}]$ 
13  }
14 }

```

```

1 algorithm find( $M[\Delta_x]$ )
2 // assume  $\mathbf{x}$  is of the form  $(x_1, \dots, x_n)$ 
3 // returns a pair  $(x, d)$  such that  $x \in \mathbf{x}$  and  $d \in \Delta_x \setminus \Pi_x$ 
4 for  $i = 1$  to  $n$  {
5   if  $M[\emptyset_{(x_1, \dots, x_i)} \cdot \Delta_{(x_{i+1}, \dots, x_n)}]$  is  $\mathcal{B}$ -satisfiable {
6     By binary search on  $\Delta_{x_i}$  determine a maximal subset  $\Gamma \subseteq \Delta_{x_i}$ 
7     such that  $M[\emptyset_{(x_1, \dots, x_{i-1})} \cdot \Gamma_{x_i} \cdot \Delta_{(x_{i+1}, \dots, x_n)}]$  is  $\mathcal{B}$ -satisfiable
8     Let  $d \in \Delta_{x_i}$  be any point adjacent to  $\Gamma$ .
9     Return  $(x_i, d)$  // From Lemma 3.9 it follows  $d \in \Delta_x \setminus \Pi_x$  as claimed.
10  }
11 }

```

We tacitly assume that the \mathcal{B} -satisfiability tests in `checkSAT` and `find` are effective. This is always the case, e.g., if there are no FG operators other than free BG-sorted operators and the EA-fragment of the background theory is decidable.

Let us go through the run of `checkSAT(N)`, where $N = \{C_1, C_2\}$ from Example 3.2. Let $\Pi_{\mathbf{x}}^1 = (\emptyset, \emptyset)$ be the initially empty set of exceptions set in line 3. For $M^1 = \text{FD}(N, \Pi_{\mathbf{x}}^1)$ in line 5 none of the termination cases applies. The call of `find(M^1)` in line 11 returns the pair $(x_1, 9)$. The underlying set Γ is $[1..8]$, which is determined by binary search for a maximal sub-range of $\Delta_{x_1} = [1..1000]$. Indeed, Γ is maximal as $M^1([1..8], \Delta_{x_2})$ is \mathcal{B} -satisfiable and $M^1([1..8] \cup \{9\}, \Delta_{x_2})$ is \mathcal{B} -unsatisfiable. The adjacent point hence is $d = 9$.⁵ The updated set $\Pi_{\mathbf{x}}^2$ in `checkSAT` now is $(\{9\}, \emptyset)$ and we get $M^2[\Delta_{\mathbf{x}}] = \text{FD}(N, \Pi_{\mathbf{x}}^2)$ in the next iteration. Again, the termination tests do not apply and `find(M^2)` is called. This time $M^2((\emptyset, \Delta_{x_2}))$ is \mathcal{B} -unsatisfiable and the result of `find(M^2)` is $(x_2, 6)$.

The updated set $\Pi_{\mathbf{x}}^3$ hence is $(\{9\}, \{6\})$ and $M^3[\Delta_{\mathbf{x}}] = \text{FD}(N, \Pi_{\mathbf{x}}^3)$ consists of the clauses C_{11} – C_{14} and C_{21} – C_{24} already shown above. In the next iteration, the set $M^3[\emptyset_{\mathbf{x}}]$ is built, which is obtained by replacing the sets $\Delta_{x_1} = \Delta_{x_2} = [1..1000]$ everywhere by the empty interval []:

$$\begin{array}{ll}
(C'_{11}) & \alpha_1 > x_1 \vee x_1 \notin [] \setminus \{9\} & (C_{13}) & \alpha_2 > 9 \\
(C'_{12}) & f(x_1) \approx \alpha_1 \vee x_1 \notin [] \setminus \{9\} & (C_{14}) & f(9) \approx \alpha_2 \\
(C'_{21}) & \alpha_3 < 10 \vee \neg(x_2 > 2) \vee x_2 \notin [] \setminus \{6\} & (C_{23}) & \alpha_4 < 10 \vee \neg(6 > 2) \\
(C'_{22}) & f(x_2 + 3) \approx \alpha_3 \vee x_2 \notin [] \setminus \{6\} & (C_{24}) & f(6 + 3) \approx \alpha_4
\end{array}$$

By construction, all clauses affected by the replacement are tautological. Yet, the set $M^3[\emptyset_{\mathbf{x}}]$ is \mathcal{B} -unsatisfiable, which can be seen easily from the clauses in the right column. The algorithm stops and returns “unsatisfiable”. This is indeed correct, as, by construction, the remaining non-tautological clauses contain and use definitions for *ground* instances of the f-terms only. Because of that, our method is sound wrt. \mathcal{B} -unsatisfiability even for non-finitely quantified clause sets as expressed in Lemma 3.8 above.

The following lemma guarantees that `find` behaves as claimed in its comment.

Lemma 4.1. *Algorithm `find` is well-defined. More precisely, whenever `find` is called from `checkSAT` on line 11 then the if-clause in the for-loop in `find` is executed for some i , and the result (x_i, d) such that $x_i \in \mathbf{x}$ and $d \in \Delta_{x_i} \setminus \Pi_{x_i}$ exists.*

Proof. Assume `find(M[\Delta_{\mathbf{x}}])` is executed and that \mathbf{x} is of the form (x_1, \dots, x_n) . Because the test in line 8 in `checkSAT` has not applied it follows that the condition in line 5 in `find` is satisfied for some i in $1, \dots, n$. Among all these values, the if-clause is executed for the least one. That is, $M[\emptyset_{(x_1, \dots, x_i)} \cdot \Delta_{(x_{i+1}, \dots, x_n)}]$ is \mathcal{B} -satisfiable and $M[\emptyset_{(x_1, \dots, x_j)} \cdot \Delta_{(x_{j+1}, \dots, x_n)}]$ is \mathcal{B} -unsatisfiable, for all j with $1 \leq j < i \leq n$. Because $i \geq 1$ we can rewrite the former and obtain that $M[\emptyset_{(x_1, \dots, x_{i-1})} \cdot \emptyset_{x_i} \cdot \Delta_{(x_{i+1}, \dots, x_n)}]$ is \mathcal{B} -satisfiable. Furthermore, $M[\emptyset_{(x_1, \dots, x_{i-1})} \cdot \Delta_{x_i} \cdot \Delta_{(x_{i+1}, \dots, x_n)}]$ is \mathcal{B} -unsatisfiable: if $i = 1$ this follows from the fact that the test in line 6 in `checkSAT` has not applied, and if $i > 1$ this follows from the minimality of i . It follows that a (maximal) set $\Gamma \subseteq \Delta_{x_i}$ and hence an element $d \in \Delta_{x_i}$ exist as claimed. \square

⁵ Notice that `find` searches for the set Γ wrt. the whole set $M = \text{FD}(N, \Pi_{\mathbf{x}}) = \text{Cls}_N \cup \text{Def}_N$. It would be tempting to fix Cls_N and search only wrt. Def_N (or vice versa) but this would be unsound.

For termination of `checkSAT`, instead of determining the pair (x, d) in line 11 by the call to `find`, one could choose any (x, d) such that the current set Π_x grows. An advantage of using `find`, however, is that the relevant ground instances of the clauses $C_1[x_1]$ and $C_2[x_2]$, which are $C_1[9]$ and $C_2[6]$, have been found through semantic guidance by refining the default interpretation in only two steps.

In general terms, `checkSAT/find` realizes a *heuristic* that tries to search for a model by deviating from the current interpretation only when a conflict arises. The conflict is identified by the point d for the variable x_i in Line 8 of `find`. The next round of `checkSAT` continues with the correspondingly updated current interpretation by adding d to Π_x , which may stop now with “satisfiable”, “unsatisfiable” or continue the search.

We summarize the essential properties of `checkSAT` in our main result as follows.

Theorem 4.2 (Correctness of `checkSAT`). *For any set N of finitely quantified clauses, `checkSAT(N)` terminates with the correct result “ \mathcal{B} -satisfiable” or “ \mathcal{B} -unsatisfiable” for N . Moreover, in case of “ \mathcal{B} -unsatisfiable” the non-domain restricted version of N is \mathcal{B} -unsatisfiable, which is obtained from N by removing from all clauses all domain predicates.*

Proof. Termination follows from the fact that `find` always returns some pair (x, d) such that $x \in \mathbf{x}$ and $d \in \Delta_x \setminus \Pi_x$, as shown in Lemma 4.1. Hence, the set Π_x grows monotonically in line 12 in `checkSAT` and there are only finitely many elements in Δ_x available for that. Correctness follows from the lemmas in Section 3 as referenced in the comments in `checkSAT`. \square

5 Experimental Results

We have implemented the `checkSAT/find` algorithm on top of the hierarchic superposition prover Beagle [8].⁶ The implementation is prototypical and currently serves only to try out the ideas in the paper. Table 1 summarizes the experiments we carried out. We have tried six problems, some of them with varying domain sizes. The problems (1) and (6) are \mathcal{B} -unsatisfiable, the others \mathcal{B} -satisfiable. The “Problem” column contains the individual clause sets. The column “ $|\Delta|$ ” gives the size of the finite domains uniformly used in the problem clauses, e.g., $|\Delta| = 50$ means the range $[1..50]$. The column “#Iter” is the number of while-loop iterations in `checkSAT` needed to solve the problem for the given Δ . The column “#TP” is the number of theorem prover calls (Beagle calls) stemming from the various \mathcal{B} -satisfiability checks in `checkSAT/find`. Finally, “Time” is the total CPU time needed to solve the problem. All experiments were carried out on a Linux desktop with a quad-core Intel i7 cpu running at 2.8 GHz. For comparison, we have also run Microsoft’s SMT-solver Z3 [18] on our examples, using the obvious formula representation of the domains Δ .

Some comments on the individual problems. Problem (1) is trivially solved, for any Δ . In fact, the default interpretation is sufficient for that. Notice that the variable y is not finitely quantified (and does not need to be). Z3 reports “unknown” on problem (1), but, surprisingly it solves the essentially same problem $f(x) > y \vee y < 0$ quickly. Problem (2) is meant to showcase our algorithm in conjunction with Beagle’s theorem proving

⁶ <http://users.cecs.anu.edu.au/~baumgart/systems/beagle/>

#	Problem	$ D $	#Iter	#TP	Time
1	$f(x) > 1 + y \vee y < 0 \vee x \notin D$	any	1	1	<1
2	$g(x) \approx x \vee g(x) \approx x + 1 \vee \neg(x \geq 0)$ $g(x) \approx -x \vee \neg(x < 0)$ $f(x) < g(x) \vee x \notin D$	10	9	79	33
3	$f(x_1, x_2, x_3, x_4) > x_1 + x_2 + x_3 + x_4 \vee$ $x_1 \notin D \vee x_2 \notin D \vee x_3 \notin D \vee x_4 \notin D$	any	1	1	<1
4	$f(x) \approx x \vee x \notin D$	10	2	7	<1
	$f(5) \approx 8$	20	2	9	<1
	$f(8) \approx 5$	50	2	13	1.3
		100	2	15	1.8
		200	2	17	3.0
		500	2	20	2.8
		1000	2	22	3.8
5	see Section 1	10	3	15	3.7
		20	3	17	3.9
		50	3	19	2.6
		100	3	21	3.6
		200	3	23	3.7
		500	3	25	4.0
		1000	3	27	4.0
6	see Example 3.2	10	4	32	1.6
		20	4	73	2.9
		50	6	389	7.4
		100	6	795	9.5
		200	6	1973	22

Table 1. Experimental results.

capabilities. The function symbol g is “sufficiently complete” defined by the first two clauses, and only the third clause containing the function symbol f needs finite quantification. Z3 could not solve this problem within three minutes. We devised problem (3) to get some insight into Z3’s capabilities on the problems we are interested in. While it is trivial for our approach, Z3 seems to instantiate the clause in problem (3). Clearly, there is a scalability issue here, as for about $|D| > 60$ the problem becomes unsolvable in reasonable time.

As a side note, we found Z3’s performance impressive, and it could solve problems (4)–(6) in very short time. Indeed, we plan to integrate Z3 in our approach and expect much better performance on many problems (Beagle’s theory reasoning component is a rather slow implementation of Cooper’s quantifier elimination algorithm.)

Problem (4) is a simple test of the default interpretation/exception mechanism. Problem (5) is the one in the Introduction, and problem (6) is our running example.

Both problem (4) and (5) scale very well, as expected, because both are proven satisfiable using the default interpretation and a fixed number of exception points. In problem (4) these are easily discovered from the problem and in (5) the exceptions are quickly discovered by the search.

Problem (6) scales poorly since there is always a small subset of instances which are satisfiable (namely when $x_2 \notin [3..6]$ in (C_2)), but many combinations of subsets of A_{x_2} are checked before it is determined that this is a maximal satisfiable set- and the number of combinations scales with the domain size. This behaviour could possibly be circumvented with better heuristics in the search procedure.

6 Conclusions

We have presented a method for deciding hierarchic satisfiability, or satisfiability modulo theories, of first-order clause sets where all variables are quantified over finite subsets of background domains. The method tries to construct a model by stepwise amending a default interpretation in a conflict-driven way by utilizing a decision procedure for the EA-fragment of the background theory. It may also terminate with a set of ground instances witnessing that no model exists. For space reasons and for clarity we have focused in this paper on the basic principles and leave extensions for future work. Here are some ideas.

Richer input language: One important extension concerns foreground-sorted variables and operators, like the array-sorted variable a and the write-operator in clauses (1) and (2) in the introduction. In the example we got away without further modifications because the axioms (1) and (2) do not pose problems for sufficient completeness and for termination of hierarchic superposition. The question is under which conditions this is possible in general. One could also try enumerating finite segments of the foreground domains in a Herbrand fashion, similarly as with background domains.

Our method can also be applied to certain richer syntactic fragments that require a full-fledged theorem prover for hierarchic specifications instead of a decision procedure for the background theory. However, this would “reverse” the common architecture by invoking that foreground reasoner from within an outer loop. This is problematic, however, because the foreground reasoner might not terminate or be incomplete. To fix that, it should be possible under certain conditions to instead integrate the `checkSAT` as an inference rule into, say, hierarchic superposition and apply it only to finitely quantified clauses as defined above. (This would directly generalize the Define-rule in [8].)

Alternative default interpretation: Taking the constant function as the default interpretation for free BG-sorted operators is not always a good choice. For example, for the clause $f(x) \approx x \vee x \notin [1..1000]$ our method needs to amend the default interpretation at every point. Fortunately, any interpretation can be used as a default, and the identity function as the default interpretation for f leads immediately to a model. (On the other hand, in this example f is already sufficiently defined and could possibly be excepted from the transformation in the first place.)

Bernays-Schönfinkel fragment: The hierarchic superposition calculus can immediately be instantiated with, say, an instance-based method for deciding background theories that are given as a set of EPR-clauses. Our method, or the extensions above, could possibly be used to integrate arithmetic reasoners, instance-based methods and superposition in a beneficial way.

References

1. E. Althaus, E. Kruglov, and C. Weidenbach. Superposition modulo linear arithmetic sup(la). In S. Ghilardi and R. Sebastiani, eds., *FroCos*, 2009, *LNCS 5749*, pp. 84–99. Springer.
2. L. Bachmair and H. Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *Journal of Logic and Computation*, 4(3):217–247, 1994.
3. L. Bachmair, H. Ganzinger, and U. Waldmann. Refutational theorem proving for hierarchic first-order theories. *Appl. Algebra Eng. Commun. Comput.*, 5:193–212, 1994.
4. P. Baumgartner, A. Fuchs, H. de Nivelle, and C. Tinelli. Computing finite models by reduction to function-free clause logic. *Journal of Applied Logic*, 7(1):58–74, 2009.
5. P. Baumgartner, A. Fuchs, and C. Tinelli. ME(LIA) – Model Evolution With Linear Integer Arithmetic Constraints. In I. Cervesato, H. Veith, and A. Voronkov, eds., *Proceedings of the 15th International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR’08)*, 2008, *LNAI 5330*, pp. 258–273. Springer.
6. P. Baumgartner and C. Tinelli. Model evolution with equality modulo built-in theories. In N. Bjoerner and V. Sofronie-Stokkermans, eds., *CADE-23 – The 23rd International Conference on Automated Deduction*, 2011, *LNAI 6803*, pp. 85–100. Springer.
7. P. Baumgartner and U. Waldmann. Hierarchic superposition: Completeness without compactness. In M. Kosta and T. Sturm, eds., *MACIS 2013 –Fifth International Conference on Mathematical Aspects of Computer and Information Sciences*, 2013, pp. 8–12.
8. P. Baumgartner and U. Waldmann. Hierarchic superposition with weak abstraction. In M. P. Bonacina, ed., *CADE-24 – The 24th International Conference on Automated Deduction*, 2013, *LNAI 7898*, pp. 39–57. Springer.
9. K. Claessen and N. Sörensson. New techniques that improve mace-style finite model building. In P. Baumgartner and C. G. Fermüller, eds., *CADE-19 Workshop: Model Computation – Principles, Algorithms, Applications*, 2003.
10. H. Ganzinger and K. Korovin. Theory Instantiation. In *Proceedings of the 13th Conference on Logic for Programming Artificial Intelligence Reasoning (LPAR’06)*, 2006, *LNCS 4246*, pp. 497–511. Springer.
11. Y. Ge, C. Barrett, and C. Tinelli. Solving quantified verification conditions using satisfiability modulo theories. In F. Pfenning, ed., *Proceedings of the 21st International Conference on Automated Deduction (CADE-21)*, Bremen, Germany, 2007, Lecture Notes in Computer Science. Springer.
12. Y. Ge and L. M. de Moura. Complete instantiation for quantified formulas in satisfiability modulo theories. In A. Bouajjani and O. Maler, eds., *CAV*, 2009, *LNCS 5643*, pp. 306–320. Springer.
13. J. Halpern. Presburger Arithmetic With Unary Predicates is Π_1^1 -Complete. *Journal of Symbolic Logic*, 56(2):637–642, 1991.
14. K. Korovin and A. Voronkov. Integrating linear arithmetic into superposition calculus. In *Computer Science Logic (CSL’07)*, 2007, *LNCS 4646*, pp. 223–237. Springer.
15. E. Kruglov and C. Weidenbach. Superposition decides the first-order logic fragment over ground theories. *Mathematics in Computer Science*, pp. 1–30, 2012.
16. W. McCune. Mace4 reference manual and guide. Technical Report ANL/MCS-TM-264, Argonne National Laboratory, 2003.
17. L. M. de Moura and N. Bjørner. Efficient e-matching for smt solvers. In F. Pfenning, ed., *CADE*, 2007, *LNCS 4603*, pp. 183–198. Springer.
18. L. M. de Moura and N. Bjørner. Z3: An efficient smt solver. In C. R. Ramakrishnan and J. Rehof, eds., *TACAS*, 2008, *LNCS 4963*, pp. 337–340. Springer.
19. R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: from an Abstract Davis-Putnam-Logemann-Loveland Procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, 2006.

20. A. Reynolds, C. Tinelli, A. Goel, and S. Krstić. Finite model finding in SMT. In N. Sharygina and H. Veith, eds., *Proceedings of the 25th International Conference on Computer Aided Verification (St Petersburg, Russia)*, 2013, LNCS 8044, pp. 640–655. Springer.
21. A. Reynolds, C. Tinelli, A. Goel, S. Krstić, M. Deters, and C. Barrett. Quantifier instantiation techniques for finite model finding in SMT. In M. P. Bonacina, ed., *Proceedings of the 24th International Conference on Automated Deduction (Lake Placid, NY, USA)*, 2013, LNCS 7898, pp. 377–391. Springer.
22. P. Rümmer. A constraint sequent calculus for first-order logic with linear integer arithmetic. In I. Cervesato, H. Veith, and A. Voronkov, eds., *Proceedings of the 15th International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR'08)*, 2008, LNAI 5330, pp. 274–289. Springer.
23. J. Slaney. Finder (finite domain enumerator): Notes and guide. Technical Report TR-ARP-1/92, Australian National University, Automated Reasoning Project, Canberra, 1992.
24. J. Zhang and H. Zhang. Sem: a system for enumerating models. In C. Mellish, ed., *IJCAI-95 — Proceedings of the 14th International Joint Conference on Artificial Intelligence, Montreal*, 1995, pp. 298–303. Morgan Kaufmann.