

Finding Finite Models in Multi-sorted First-Order Logic

Giles Rege¹(✉), Martin Suda¹, and Andrei Voronkov^{1,2,3}

¹ University of Manchester, Manchester, UK
`giles.reger@manchester.ac.uk`

² Chalmers University of Technology, Gothenburg, Sweden

³ EasyChair, Manchester, UK

Abstract. This work extends the existing MACE-style finite model finding approach to multi-sorted first-order logic. This existing approach iteratively assumes increasing domain sizes and encodes the related ground problem as a SAT problem. When moving to the multi-sorted setting each sort may have a different domain size, leading to an explosion in the search space. This paper focusses on methods to tame that search space. The key approach adds additional information to the SAT encoding to suggest which domains should be grown. Evaluation of an implementation of techniques in the Vampire theorem prover shows that they dramatically reduce the search space and that this is an effective approach to find finite models in multi-sorted first-order logic.

1 Introduction

There have been a number of approaches looking at finding finite models for First-Order Logic (FOL), however there has not been much work on finding such models for Multi-Sorted FOL where symbols are given *sorts*. We consider a model finding method, pioneered by MACE [12], that encodes the search as a SAT problem. We show how this method can be modified to deal directly with multi-sorted input, rather than translating the problem to the unsorted setting, which is the most common current method.

There are two main motivations for this work. Firstly, many problems are more naturally expressed in multi-sorted FOL than in unsorted FOL (although their theoretical expressive power is equivalent). Therefore, it is useful to be able to reason in this setting and translations from multi-sorted FOL to unsorted FOL often make this reasoning harder. Secondly, MACE-style model finders can use sort information to make the SAT encoding smaller. However, as we discuss below, finding finite models of multi-sorted formulas also presents significant challenges.

This work was supported by EPSRC grant EP/K032674/1. Martin Suda and Andrei Voronkov were partially supported by ERC Starting Grant 2014 SYMCAR 639270. Andrei Voronkov was also partially supported by the Wallenberg Academy Fellowship 2014 - TheProSE.

The MACE-style approach, later extended in the Paradox [5] work, involves selecting a domain size for the finite model, grounding the first-order problem with this domain and translating the resulting formulas into a SAT problem, which, if satisfied, gives a finite model of the selected size. Search for a finite model then involves considering iteratively larger domain sizes. In the multi-sorted setting it is necessary to consider the size of each sort separately. This can be demonstrated by the following example, which is an extension of the much used *Monkey Village* example [2, 4].

Example 1 (Organised Monkey Village). Imagine a village of monkeys where each monkey owns at least two bananas. As the monkeys are well-organised, each tree contains exactly three monkeys. Monkeys are also very friendly, so they pair up to make sure they will always have a partner. We can represent this problem as follows:

$$\begin{aligned}
 & (\forall M : \textit{monkey})(\textit{owns}(M, \mathbf{b}_1(M)) \wedge \textit{owns}(M, \mathbf{b}_2(M)) \wedge \mathbf{b}_1(M) \neq \mathbf{b}_2(M)) \\
 & (\forall M_1, M_2 : \textit{monkey})(\forall B : \textit{banana})(\textit{owns}(M_1, B) \wedge \textit{owns}(M_2, B) \rightarrow M_1 = M_2) \\
 & (\forall T : \textit{tree})(\exists M_1, M_2, M_3 : \textit{monkey})((\bigwedge_{i=1}^3 \textit{sits}(M_i) = T) \wedge \textit{distinct}(M_1, M_2, M_3)) \\
 & (\forall M_1, M_2, M_3, M_4 : \textit{monkey})(\forall T : \textit{tree})((\bigwedge_{i=1}^4 \textit{sits}(M_i) = T) \Rightarrow \neg \textit{distinct}(M_1, M_2, M_3, M_4)) \\
 & (\forall M : \textit{monkey})(\textit{partner}(M) \neq M \wedge \textit{partner}(\textit{partner}(M)) = M)
 \end{aligned}$$

where the predicates *owns* associates monkeys with bananas, the functions \mathbf{b}_1 and \mathbf{b}_2 witness the existence of each monkey's minimum two bananas, the function *sits* maps monkeys to the tree that they sit in, the function *partner* associates a monkey with its partner, and the (meta-)predicate *distinct* is true if all of its arguments are distinct.

This problem requires the domain of *monkey* to be exactly three times larger than the domain of *tree*, the domain of *banana* to be at least twice as large as the domain of *monkey*, and the domain of *monkey* to be even. The main model finding effort then becomes searching for an assignment of domain sizes that satisfies this problem. Here, the smallest such assignment is $|tree| = 2$, $|monkey| = 6$, and $|banana| = 12$. In the general case it is necessary to try all combinations of domain sizes. In the worst case this will mean trying a number of assignments exponential in the number of sorts.

The techniques introduced in this paper tackle this issue by introducing a number of ways to constrain this search space. The main contributions can be summarised as:

1. We show how the MACE-style approach (described in Sect. 3) can be extended to the multi-sorted setting via a novel extension of the SAT encoding (Sect. 5). This encoding uses information from the SAT solver to guide the search through the space of domain size assignments.
2. We use *monotonic sorts* introduced in [4] in a new way for the multi-sorted case to reduce the search space further (Sect. 5.5).
3. We utilise the saturation-based approach for first-order logic to detect further constraints on the search space introduced by injective and surjective functions (Sect. 6).

4. We present an alternative to (1), a complementary search strategy, utilising a different SAT encoding, that only needs to expand (and never shrink) the sizes of sort domains (Sect. 8).

These ideas have been realised within the VAMPIRE theorem prover [11] and evaluated on problems taken from the TPTP and SMT-LIB benchmark suites. Our experimental evaluation (Sect. 9) shows that these techniques can be used to (i) improve finite model finding in the unsorted setting, (ii) effectively and efficiently find finite models in the multi-sorted setting, and (iii) detect cases where no models exist.

2 Preliminaries

Multi-sorted First-Order Logic. We consider a multi-sorted first-order logic with equality. A term is either a variable, a constant, or a function symbol applied to terms. A literal is either a propositional symbol, a predicate applied to terms, an equality of two terms, or a negation of either. Function and predicate symbols are *sorted* i.e. their arguments (and the return value in the case of functions) have a unique *sort* drawn from a finite set of sorts S . We only consider well-sorted literals. There is an equality symbol per sort and equalities can only be between terms of the same sort. Formulas may use the standard notions of quantification and boolean connectives, but in this work we assume all formulas are *clausified* using standard techniques. A *clause* is a disjunction of literals where all variables are universally quantified (existentially quantified variables can be replaced by skolem functions during clausification).

SAT Solvers. The technique we present later will make use of a black-box SAT solver and we assume the reader is familiar with their general properties. We assume that a SAT solver supports *solving under assumptions* [7,8]. This means the SAT solver can be asked to search for a model of a set of clauses N additionally satisfying a conjunction of assumption literals A and is able, in case the answer is UNSAT, to provide a subset $A_0 \subseteq A$ of those assumptions which were sufficient for the unsatisfiability proof.

3 MACE-Style Finite Model Finding in an Unsorted Setting

We describe the finite model finding procedure in a single sorted setting. This is a variation of the approach taken by Paradox [5]. The general idea is to create, for each integer $n \geq 1$, a SAT problem that is satisfiable if the problem has a finite model of size n . To find a finite model we therefore iterate the approach for domain sizes $n = 1, 2, 3, \dots$

3.1 DC-Models

Let S be a set of clauses. Let us fix an integer $n \geq 1$. Let $DC = \{c_1, \dots, c_n\}$ be a set of distinct constants not occurring in S , we will call the elements of DC *domain constants*. We extend the language by adding the domain constants and say that an interpretation is a *DC-interpretation*, if (i) the domain of this interpretation is DC and (ii) every domain constant c_i is interpreted in it by itself. Every model of S that is also a *DC-interpretation* will be called a *DC-model* of S . It is not hard to argue that, if S has a model of size n , then it also has a *DC-model*. We say that S is *n-satisfiable* if it has a model of size n .

Let C be a clause. A *DC-instance* of C is a ground clause obtained by replacing every variable in C by a constant in DC . For example, if $p(x) \vee x = y$ is a clause and $n \geq 2$, then $p(c_1) \vee c_1 = c_2$ and $p(c_1) \vee c_1 = c_1$ are *DC-instances*, while $p(c_1) \vee c_2 = c_3$ is not a *DC-instance*. A clause with k different variables has exactly n^k *DC-instances*.

Theorem 1. *Let I be a DC-interpretation and C a clause. Then C is true in I if and only if all DC-instances of C are true in I .*

Let us denote by S^* the set of all *DC-instances* of the clauses in S . Consider an example. Let S consist of three clauses

$$p(b), \quad f(a) \neq b, \quad f(f(x)) = x.$$

The smallest model of S has a domain of size two. Take $n = 2$, then $DC = \{c_1, c_2\}$. By the above theorem, S has a model of size two if and only if S^* has a *DC-model*. The set S^* consists of four ground clauses:

$$p(b), \quad f(a) \neq b, \quad f(f(c_1)) = c_1, \quad f(f(c_2)) = c_2.$$

Note that *DC-models* are somehow similar to Herbrand models used in logic programming and resolution theorem proving, except that they are built using (domain) constants instead of all ground terms and *DC-instances* instead of ground instances.

Theorem 1 is not directly applicable to encode the existence of models of size n as a SAT problem, because *DC-instances* can contain complex terms. We will now introduce a special kind of ground atom which contains no complex subexpressions. We call a *principal term* any term of the form $f(d_1, \dots, d_m)$, where $m \geq 0$, f is a function symbol, which is not a domain constant, and d_1, \dots, d_m are domain constants. In our example there are four principal terms: $a, b, f(c_1), f(c_2)$. A ground atom is called *principal* if it either has the form $p(d_1, \dots, d_m)$ where $m \geq 0$, p is a predicate symbol different from equality and d_1, \dots, d_m are domain constants or has the form $t = d$, where t is a principal term and d a domain constant. We call a *principal literal* a principal atom or its negation.

Theorem 2. *Let I_1, I_2 be DC-interpretations. If they satisfy the same principal atoms, then I_1 coincides with I_2 .*

Theorem 1 reduces n -satisfiability of S to the existence of a DC -interpretation of the set S^* of ground clauses. Theorem 2 shows that DC -interpretations can be identified by the set of principal atoms true in them. What we will do next is to introduce a propositional variable for every principal atom and reduce the existence of a DC -model of S^* to satisfiability of a set of clauses using only principal literals.

3.2 The SAT Encoding

The main step in the reduction is to transform every non-ground clause C into an equivalent clause C' such that DC -instances of C' consist (almost) only of principal literals. We will explain what “almost” means below. This transformation is known as *flattening*.

Flattening. A literal is called *flat* if it has one of the following forms:

1. $p(x_1, \dots, x_m)$ or $\neg p(x_1, \dots, x_m)$, where $m \geq 0$ and p is a predicate symbol;
2. $f(x_1, \dots, x_m) = y$ or $f(x_1, \dots, x_m) \neq y$, where $m \geq 0$ and f is a function symbol, which is not a domain constant.
3. an equality between variables $x = y$.

Every DC -instance of a flat literal is either a principal literal (for the first two cases), or an equality $c_i = c_j$ between domain constants.

To flatten clauses in S , we first get rid of all inequalities between variables, replacing every clause of the form $x \neq y \vee C[x]$ by the equivalent clause $C[y]$. Then we repeatedly replace every clause $C[t]$, where t is not a variable and t occurs as an argument to a predicate or a function symbol, by the equivalent clause $t \neq x \vee C[x]$, where x is a fresh variable.

Our example clauses can be flattened as follows:

$$p(y) \vee b \neq y, \quad f(y_1) = y_2 \vee a \neq y_1 \vee b \neq y_2, \quad f(y) = x \vee f(x) \neq y$$

DC -Instances. We can now produce the DC -instances of each flattened clause $C[x_1, \dots, x_k]$. For our running example (with $n = 2$) this produces the following ten DC -instances:

$$\begin{array}{lll} p(c_1) \vee b \neq c_1 & f(c_1) = c_1 \vee a \neq c_1 \vee b \neq c_1 & f(c_1) = c_1 \vee f(c_1) \neq c_1 \\ p(c_2) \vee b \neq c_2 & f(c_1) = c_2 \vee a \neq c_1 \vee b \neq c_2 & f(c_1) = c_2 \vee f(c_2) \neq c_1 \\ & f(c_2) = c_2 \vee a \neq c_2 \vee b \neq c_2 & f(c_2) = c_2 \vee f(c_2) \neq c_2 \\ & f(c_2) = c_1 \vee a \neq c_2 \vee b \neq c_1 & f(c_2) = c_1 \vee f(c_1) \neq c_2 \end{array}$$

Note that all literals are principal. If we treat principal atoms as propositional variables, the two leftmost clauses can be satisfied by making $b \neq c_1$ and $b \neq c_2$ both true, but this violates the assumption that b should equal one of the domain constants. Additionally, the two rightmost topmost clauses can be satisfied by making $f(c_1) = c_1$ and $f(c_1) = c_2$ true but this violates the assumption that f is a function. We would like to prevent both situations. To do this we introduce additional definitions.

Functionality Definitions. For each principal term p and distinct domain constants d_1, d_2 we produce the following clause

$$p \neq d_1 \vee p \neq d_2,$$

These clauses are satisfied by every *DC*-interpretation and guarantee that all function symbols are interpreted as (partial) functions.

For our running example we introduce four new definitions:

$$a \neq c_1 \vee a \neq c_2, \quad b \neq c_1 \vee b \neq c_2, \quad f(c_1) \neq c_1 \vee f(c_1) \neq c_2, \quad f(c_2) \neq c_1 \vee f(c_2) \neq c_2$$

Totality Definitions. For each principal term p we produce the following clause

$$p = c_1 \vee \dots \vee p = c_n$$

These clauses are satisfied by every *DC*-interpretation of size n and guarantee, together with functionality axioms, that all function symbols are interpreted as total functions.

For our running example we introduce four new definitions:

$$a = c_1 \vee a = c_2, \quad b = c_1 \vee b = c_2, \quad f(c_1) = c_1 \vee f(c_1) = c_2, \quad f(c_2) = c_1 \vee f(c_2) = c_2$$

The resulting SAT clauses have a model, meaning that the original clauses have a finite model with a domain of size 2, which can be extracted from the SAT encoding.

Equalities Between Variables. Flattening can result in equalities between variables, that is, clauses of the form $C \vee x = y$. *DC*-instances of such clauses can have, in addition to principal literals, equalities between domain constants $d_1 = d_2$, which are not principal literals. Since we only want to deal with principal literals, we will get rid of such equalities in an obvious way: delete clauses containing tautologies $d = d$ and delete from clauses literals $d_1 = d_2$, where d_1 are distinct d_2 domain constants.

The following theorem underpins the SAT-based finite model building method:

Theorem 3. *Let S be a set of flat clauses and S' be the set of clauses obtained from S^* by removing equalities between domain constants as described above and adding all functionality and totality definitions. Then (i) all literals in S' are principal and (ii) S is n -satisfiable if and only if S' is propositionally satisfiable.*

Incrementality. In [5] the authors describe a method for incremental finite model finding which advocates keeping (parts of) the contents of the SAT solver when increasing n . However, in previous experiments we discovered that the technique of *variable and clause elimination* [6] is useful at reducing the size of the SAT problem. As this is not compatible with incremental solving, our general approach is non-incremental.

3.3 Reducing the Number of Variables

The number of instances produced is exponential in the number of variables in a flattened clause. We describe two approaches that aim to reduce this number.

Definition Introduction. This reduces the size of clauses produced by flattening. Complex ground subterms are removed from clauses by introducing definitions. For example, a clause $p(f(a, b), g(f(a, b)))$ becomes $p(e_1, e_2)$ and we introduce the definition clauses $e_1 = f(a, b)$ and $e_2 = g(e_1)$, where e_1, e_2 are new constants. One can also introduce definitions for non-ground subterms.

Clause Splitting. Clauses with k variables are split into subclauses having less than k variables each. New predicate symbols applied to the shared variables are then added to join the subclauses. For example, the clause $p(x, y) \vee q(y, z)$ with three variables is replaced by the two clauses $p(x, y) \vee s(y)$ and $\neg s(y) \vee q(y, z)$ where s is a new predicate symbol. These new clauses have two variables each. For large domain sizes splitting can drastically reduce the size of the resulting propositional problem. This was first used for finite model finding by Gandalf [17] and later in Eground [14] for EPR problems.

3.4 Symmetry Breaking

The SAT problem produced above can contain many symmetries. For example, every permutation of DC applied to a DC -model will give a DC -model, and there are $n!$ such permutations. We can (partially) break these symmetries as follows. Firstly, if the input contains constants a_1, \dots, a_l we can add the clauses

$$a_i \neq c_m \vee a_1 = c_{m-1} \vee \dots \vee a_{i-1} = c_{m-1}$$

for $1 < i \leq l$ and $1 < m \leq n$, where we have arbitrarily ordered the constants and captured the constraint that if the i -th constant is equal to a domain element then some earlier constant must be equal to the next smallest domain element. Secondly, we can tell the SAT solver about this order on constants by adding the clauses

$$a_i = c_1 \vee \dots \vee a_i = c_i$$

for $i \leq \min(m, n)$, which captures the constraint that the i -th constant must be equal to one of the first i -th domain elements. If $1 < m < n$ then we can also use principal terms other than constants in the second case, but not in the first.

3.5 Determining Unsatisfiability

If it is possible to detect the *maximum* domain size then it is possible to show there is no model for a formula if all domain sizes up to, and including, this maximum size have been explored. There are two straightforward ways to detect maximum domain sizes. Firstly, we can look for axioms such as $(\forall x)(x = a \vee x = b)$ and $(\forall x)(\forall y)(\forall z)(x = y \vee x = z \vee z = y)$. Both indicate that the problem has a maximum domain size of 2. Secondly, we can look for so-called EPR problems that only use constant function symbols, in this case, the domain size is bounded by the number of constants.

4 Previous Work in the Multi-sorted Setting

We review previous work related to finite model finding for multi-sorted FOL.

Translating Sorts Away. One approach to dealing with multi-sorted FOL is to translate the sorts away. We discuss two well-known translations, see [2] for further discussions of such translations.

Sort Predicates. One can *guard* the use of sorted variables by a *sort predicate* that indicates whether a variable is of that sort. This predicate can be set to false in a model for all constants not of the appropriate sort. For example, the last formula in the Organised Monkey Village problem can be rewritten using the sort predicate `isMonkey`.

$$(\forall M)(\text{isMonkey}(M) \rightarrow \text{partner}(M) \neq M \wedge \text{partner}(\text{partner}(M)) = M)$$

One also needs to add additional axioms that say that sorts are non-empty and that functions return the expected sort. For the *monkey* sort we need to add

$$(\exists M)(\text{isMonkey}(M)) \quad (\forall M)(\text{isMonkey}(\text{partner}(M))).$$

Sort Functions or Tags. One can *tag* all values of a sort using a *sort function* for that sort. The idea is that in a model the function can map all constants (of any sort) to a constant of the given sort. For example, the last formula from the Organised Monkey Village problem can be rewritten using `fm` as a sort function for *monkey*:

$$(\forall M)(f_m(\text{partner}(f_m(M))) \neq f_m(M) \wedge f_m(\text{partner}(f_m(\text{partner}(f_m(M))))) = f_m(M))$$

The authors of [2] suggest conditions that allow certain sort predicates and functions to be omitted. However, their arguments relate to resolution proofs and do not apply here.

Sorting it Out with Monotonicity. In [4] Claessen et al. introduce a monotonicity analysis and show how it can help translate multi-sorted formulas to unsorted ones by only applying the above translations to non-monotonic sorts. A sort τ is monotonic for a multi-sorted FOL formula ϕ if for any model of ϕ one can add an element to the domain of τ to produce another model of ϕ . For example, in the Organised Monkey Village example the *banana* sort is monotonic as we can add more bananas once we have enough. However, *monkey* and *tree* are not monotonic as increasing either requires more trees, monkeys and bananas.

In [4] they observe that if there is no positive equality between elements of a sort then a new domain constant can be added and made to behave like an existing domain constant and there is no way to detect this i.e. positive equalities are required to bound a sort. They refine this notion further by noting that a positive equality can be guarded by a predicate, if that predicate can be forced to be true for all new domain elements. They introduce a calculus and associated SAT encoding capturing these ideas that can be used to detect monotonic sorts, which we use in our work.

Using a Theory of Sort Cardinalities. In the single-sorted setting there is a family of techniques called SEM-style after the SEM model finder [18] based on constraint satisfaction methods. There exists a technique in this direction for the multi-sorted setting implemented in the CVC4 SMT solver [13]. The idea behind this approach is to introduce a theory of sort cardinality constraints and to incorporate this theory into the standard SMT solver structure. Briefly, this approach introduces cardinality constraints (upper bounds) for sorts and searches for a set of constraints that is consistent with the axioms. To check a cardinality constraint k for sort s , a congruence relation is built for s -terms and an attempt made to merge congruence classes so that there are at most k . Cardinality constraints are then increased if found to be inconsistent. Quantified formulas are then instantiated with representative constants from the equivalence classes.

5 A Framework for the Multi-sorted Setting

In this section we introduce our framework able to build models of multi-sorted formulas directly, in contrast to translating the sorts away. The key challenge is dealing with a large and growing search space of domain sizes.

5.1 Using Sorts in the SAT Encoding

The SAT encoding in Sect. 3 can be updated to become sort-aware. First, instead of the domain size n we use finite domain sizes n_s for every sort s . Second, instead of considering $DC = \{c_1, \dots, c_n\}$, we consider domains for each sort $DC_s = \{c_1, \dots, c_{n_s}\}$. We can now define n as a function (called *domain size assignment*) mapping each sort s to n_s and likewise, define DC as the function mapping each sort s to DC_s . After that we can speak about DC -models and n -satisfiability in the multi-sorted case.

All the definitions for the one-sorted case are modified to respect sorts. This means, in particular, that in a DC -instance of a clause a variable of a sort s can only be replaced by a domain constant in DC_s . For example, for the Organised Monkey Village problem (see page 324) we could consider the domain size assignment n such that $n_{tree} = 1$, $n_{monkey} = 2$ and $n_{banana} = 2$. The first formula in this description can be split into three clauses, the first of which would be flattened as $\text{owns}(M, x) \vee \mathbf{b}_1(M) \neq x$, which would have the two DC -instances $\text{owns}(c_1, c_1) \vee \mathbf{b}_1(c_1) \neq c_1$ and $\text{owns}(c_1, c_2) \vee \mathbf{b}_1(c_1) \neq c_2$. We can use c_1 for both monkeys and bananas here as monkeys and bananas are never compared. For this reason we can also break symmetries on a per-sort basis.

Once we have updated the SAT encoding, finite model finding can then proceed as before where we construct the SAT problem for the current domain size assignment, check for satisfiability, and then either return a model or repeat the process with an updated domain size assignment. The problem then becomes how to generate the next domain size assignment to try.

5.2 A Search Strategy

We will view the search space of domain size assignments as an infinite directed graph whose nodes are domain size assignments and the children of an assignment are all the nodes that have exactly one domain size that is one larger. Thus, the number of children of every node is the number of sorts. A child of a node n having a larger domain size than n for a sort s is called the s -child of n . The s -descendant relation is the transitive closure of the s -child relation.

A *search strategy* will explore this graph node by node in such a way that a node is always visited before its children. For each node n that we visit, we can either check n -satisfiability or *ignore* this node. To decide whether a node can be ignored, we will maintain a set of *constraints*. Abstractly, a constraint is a predicate on domain size assignments and nodes that do not satisfy the current set of constraints will be ignored. Concretely, we will use a language of (boolean combinations of) arithmetical comparison literals such as $|s| < b$, $|s| \leq b$, \dots to represent the constraints. Here b stands for a concrete integer and $|s|$ is a symbolic placeholder variable for the “intended size” of the domain of sort s . The semantics of this representation is the obvious one.

We will work with a queue Q of nodes and a set \mathcal{C} of *constraints*. Initially, Q consists of a single node assigning 1 to all sorts and \mathcal{C} is empty. We then repeat the following steps:

1. If Q is empty, return “unsatisfiable”.
2. Remove the node q from the front of Q . Do nothing if q was visited before at this step. Otherwise, continue with the following steps.
3. If q satisfies all constraints in \mathcal{C} , perform finite model finding for q , terminating if a model is found. In variations of this algorithm considered later, we can add some constraints to \mathcal{C} at this step: these constraints will be obtained by analyzing the proof of q -unsatisfiability.
4. Add to Q all children of q .

We will now introduce an important notion helping us to prevent exploring large parts of the search space. A constraint is said to have the s -beam property at a node n , if all s -descendants of n violate this constraint. For example, the constraint $|s| < 3$ has the s -beam property at any node q having $q_s = 2$. We can generalize this notion to more than one sort.

With this notion we can improve step 4 of the algorithm as follows:

- 4.' If there is a constraint in \mathcal{C} having an s -beam property at the s -child n of q , if n violates this constraint, add to Q all children of q apart from n .

For example, if we have the constraint $|s| < 3$ and $q_s = 2$, this constraint will prevent us from considering the s -child n of q having $n_s = 3$.

As a small refinement, we introduce a heuristic for deciding which node in the queue to consider next, rather than processing them in the first-in-first-out order. The idea is to estimate how difficult a domain size assignment is to check and to prioritise exploration of the easier parts of the search space. Under this variation,

Q is a priority queue ordered by some size measure of the corresponding SAT encoding (in the experiment, we measured size in the number of clauses). This setup is complete, as long as this size grows strictly from a parent to its child (which is trivially satisfied for number of clauses).

5.3 Encoding the Search Problem

We now show how an extension of the SAT encoding can be used to produce constraints and therefore indicate areas of the search space that should be avoided. This is done by marking certain clauses of the encoding with certain special variables and using the mechanism for solving under assumptions to detect which of these clauses were actually used in the unsatisfiability proof. We will design the names of these special marking variables in such a way that the detected set of used assumptions will immediately correspond to a (disjunctive) constraint.

Let us assume we are encoding for the domain size assignment n . For each sort s we introduce two new propositional variables “ $|s| > n_s$ ” and “ $|s| < n_s$ ”, which can be understood as stating that the intended size of the domain of s should be larger, respectively smaller, than current n_s . The marking of clauses is now done as follows.

The totality definition for each principal term p becomes

$$p = c_1 \vee \dots \vee p = c_{n_s} \vee “|s| > n_s”$$

i.e. either the principal term equals one of the domain constants or the domain is currently too small. *DC*-instances can be similarly updated. Let C be a *DC*-instance and let $\text{sorts}(C)$ be the set of sorts of variables occurring in C . We replace C by

$$C \vee \bigvee_{s \in \text{sorts}(C)} “|s| < n_s”$$

i.e. either the *DC*-instance holds or the domain is too large.

We then attempt to solve the updated SAT problem under the assumptions

$$A = \bigwedge_{s \in S} (\neg “|s| > n_s”) \wedge (\neg “|s| < n_s”)$$




Fig. 1. Finding a finite model for the Organised Monkey Village problem.

i.e. we assume that we are using the correct domain sizes. These added assumptions ensure that the logical meaning of the updated encoding is exactly the same as before. However, in the unsatisfiable case the solver now returns a subset $A_0 \subseteq A$ of the assumptions that were sufficient to establish unsatisfiability. Equivalently, $\neg A_0$ is a conflict clause over the marking variables implied by the encoded problem. This clause can now be understood as the newly derived constraint. We just need to interpret the marking variables in their “unquoted” form, i.e., as arithmetic comparison literals.

The argument why this interpretation is correct is best done with the set A_0 , which contains the marking variables negated. It consists of two main observations:

1. If A_0 contains $\neg“|s| > n_s”$, the unsatisfiability relies on a totality clause for the sort s . Because a totality clause gets logically stronger when the domain size is decreased, essentially the same unsatisfiability proof could be repeated for the domain size $|s|$ smaller or equal to the current n_s (given the other conditions from A_0).
2. If A_0 contains $\neg“|s| < n_s”$, the unsatisfiability relies on a *DC*-instance with a variable of sort s . Because we only add more instances of a clause if a domain size is increased, the same unsatisfiability proof would also work for the domain size $|s|$ greater or equal to the current n_s .

Thus at least one of the (atomic) constraints represented by the literals in A_0 must be violated by a domain size assignment, if we want to have a chance of finding a model.

5.4 An Example

Let us consider the Organised Monkey Village example (page 2). Running the initial search strategy on this problem requires checking 2,661 different domain sort assignments. Using the encoding described above means that only 18 assignments are tried. The search carried out by our approach is illustrated in Fig. 1. We give two projections of the 3-dimensional search space and the arrows show the parts of the search space ruled out by *s*-beam constraints. On each step the constraints rule out all but one neighbour, meaning that we take a direct path to the solution through the search space.

5.5 Using Monotonicity

In our framework we can use the notion of monotonicity (see Sect. 4) in two ways.

- *Collapsing Monotonic Sorts*. All monotonic sorts can be collapsed into a single sort as this sort can grow to the size of the largest monotonic sort. It is never safe to collapse a monotonic sort into a non-monotonic one as the monotonic sort may depend on the non-monotonic one. For example, whilst *banana* is a monotonic sort it must always be twice as large as the non-monotonic sort *monkey*.

- *Refining the Search Encoding.* If a sort s is monotonic then if no model exists for domain size n_s then no model can exist where n_s is smaller. We reflect this in our encoding by not marking DC -instances with marking variables for monotonic sorts. This leads to a derivation of potentially stronger constraints.

6 Detecting Constraints Between Sorts

In this section we discuss how properties of functions between sorts can be used to further constrain the domain size search space. Consider the following set of formulas

$$\begin{aligned} &\text{distinct}(a_1, a_2, a_3, a_4, a_5) && (\forall x : s_1)(f(x) \neq b) \\ &(\forall x, y : s_1)(f(x) = f(y) \rightarrow x = y) && (\forall x, y : s_2)(g(x) = g(y) \rightarrow x = y) \end{aligned}$$

where a_i are constants of sort s_1 , b is a constant of sort s_2 , $f : s_1 \rightarrow s_2$, and $g : s_2 \rightarrow s_3$. The previous approach would try increasing the size of each sort by 1, discovering that one sort must grow at each step. However, we can see from the bottom two formulas that f and g are injective and therefore that $|s_1| \leq |s_2| \leq |s_3|$, furthermore, the second formula tells us that f is non-surjective and therefore that $|s_1| < |s_2|$. Using these constraints we can immediately discount 9 of the 15 domain size assignments considered without them.

To find constraints between the sizes of sorts s_1 and s_2 we look for four cases:

1. If a function $f : s_1 \rightarrow s_2$ is injective, then $|s_1| \leq |s_2|$.
2. If a function $f : s_1 \rightarrow s_2$ is injective and non-surjective, then $|s_1| < |s_2|$.
3. If a function $f : s_1 \rightarrow s_2$ is surjective, then $|s_1| \geq |s_2|$.
4. If a function $f : s_1 \rightarrow s_2$ is surjective and non-injective, then $|s_1| > |s_2|$.

These constraints can be added to the constraints used in the search described in Sect. 5.

Our method for detecting bounds was inspired by Infinox [3], a method for showing no finite model can exist for unsorted FOL formulas if there is a strict bound within a sort. We detect bounds by attempting to prove properties of functions *between* sorts. For a unary function $f : s_1 \rightarrow s_2$ occurring in the problem we can simply make a claim such as

$$(\forall x : s_1)(\forall y : s_2)(f(x) = f(y) \rightarrow x = y) \wedge (\exists y : s_2)(\forall x : s_1)(f(x) \neq y)$$

for each case (this is case (2) above), and then ask whether this claim follows from the axioms of the input problem. For non-unary functions it is necessary to existentially quantify over one of the arguments, details of how to do this can be found in [3].

To check each claim C we could use standard techniques to check $A \models C$ where A are the input axioms. Any black box solver could be used for this. However, doing this on a per-claim basis is inefficient and we implement an optimisation of Vampire’s saturation loop to establish multiple claims in a single proof attempt. Recall that the saturation loop will search for consequences of

its input. Therefore, we saturate $A \cup \{C_i \rightarrow l_i\}$ where l_i is a fresh propositional symbol labelling claim C_i . If the unit l_i is derived then we can conclude that the claim C_i is a consequence of A . This approach was inspired by the consequence elimination mode of Vampire [9] (see this work for technical details).

7 Getting More Sorts

Previously we have seen how sort information can be used to reduce the size of the SAT encoding by only growing the domain sizes of sorts that need to be grown. In this section we recall a technique first described in [4] for inferring new sorts and explain how these new sorts can be useful.

Inferring Subsorts. Consider the Organised Monkey Village example. The *monkey* sort can be split into three separate subsorts as there are three parts (assigning bananas to monkeys, assigning monkeys to trees and assigning monkeys to their partners) where the signatures do not overlap. Abstractly, we can use different monkeys in these different places as they do not interact – later we will see why this is useful. To infer such subsorts we can use the standard union-find method on positions in the signature.

Using Inferred Subsorts. Claessen et al. [4] describe two uses for inferred subsorts:

–*Removing Instances.* If a subsort τ is monotonic and all function symbols with the return sort τ are constants, then we can bound the subsort by the number of constants. It is easy to argue that any ground clauses (instances, totality or functionality) for a domain constant larger than the bound of the subsort can be omitted as they will necessarily be equivalent to an existing clause. This helps reduce the size of the SAT encoding.

–*Symmetry breaking.* For the same reasons that symmetry breaking can occur per sort, symmetry breaking can now occur per inferred subsort. This is safe due to the above observation that values for different subsorts will never be compared.

Making Subsorts Proper Sorts. Proper sorts and inferred subsorts are treated differently as we only grow the sizes of proper sorts. If an inferred subsort is not bounded as described above then it is forced to grow to the same size as its parent sort. To understand why this can be problematic consider the FOL formula

$$\text{distinct}(a_1, \dots, a_{50}) \wedge (\forall x)(f(f(f(f(f(f(f(f(f(x)))))))))) \neq x)$$

which has an overall finite model size of 50. Establishing this finite model requires a SAT problem consisting of 1,187,577 clauses. However, there are two subsorts: that of the constants a_1 to a_{50} and that of f . The second subsort is monotonic and does not need to grow beyond size 3. If this had been declared as a separate sort then the required SAT encoding would only consist of 125,236 clauses.

It is only safe to treat an inferred subsort as a proper sort if we can translate any resulting model into one where elements of the inferred subsort belong to the original sort. This is possible when (i) the inferred subsort is monotonic, and

(ii) the size of the inferred subsort is not larger than the size of its parent sort. To ensure (ii) we add constraints to the search strategy in the same way as for sort bounds detected previously.

8 An Alternative Growing Search

The previous search strategy considers each domain size assignment separately (we therefore refer to it as a *pointwise* encoding). However, we can modify the encoding so that it captures the current assignment *and all smaller ones* at the same time. Thus we no longer talk of a domain size but rather of a domain size upper *bound*, as the parameter of the encoding. These bounds never need to shrink and thus grow monotonically for each sort. We call this encoding a *contour* encoding as we can think of it drawing a contour around the explored part and growing this outwards.

This alternative encoding works as follows. For each sort s with domain size bound n_s we introduce n_s propositional variables $bound_s(1)$ to $bound_s(n_s)$. Then instead of single totality constraint for each principal term p we introduce all totality constraints for domain sizes up to n_s guarded by the appropriate bound i.e.

$$p = c_1 \vee bound_s(1), \quad \dots, \quad p = c_1 \vee \dots \vee p = c_{n_s} \vee bound_s(n_s)$$

We guard *DC*-instances of clauses with negations of these guards in the following way. For each sort s let s_{max} be the index of the largest domain constant in this instance used to replace a variable of sort s . Then if s_{max} is defined, i.e. there is at least one such variable, and $s_{max} > 1$ we guard the instance with a literal $\neg bound_s(s_{max} - 1)$. For example, given a function symbol $f : s_1 \rightarrow s_2$, a constant $b : s_2$, and a flattened clause $f(x) \neq y \vee b \neq y$, its *DC*-instance $f(c_3) \neq c_1 \vee b \neq c_1$ would be guarded as $f(c_3) \neq c_1 \vee b \neq c_1 \vee \neg bound_{s_1}(2)$.

In this encoding the SAT solver can satisfy the clauses for a domain size smaller than n_s i.e. if it can satisfy a stricter totality constraint then it can effectively ignore some of the instances. As a further variation, if a sort is monotonic then we do not need to consider the possibility that a sort is smaller than its current bound. Therefore, we only need the largest totality constraint and do not need constraints on instances.

In a similar way as before, we solve the problem under the assumptions that the sort sizes are big enough i.e.

$$A = \bigwedge_{s \in S} \neg bound_s(n_s).$$

If this is shown unsatisfiable the subset of assumptions A_0 will suggest the sorts that could be grown; growing a sort not mentioned in A_0 would allow the same proof of unsatisfiability to be produced. If A_0 is empty, the SAT-solver has shown that the given first-order formula is unsatisfiable. Otherwise, we can either arbitrarily select a sort to grow out of the ones mentioned in A_0 . This approach is significantly different from the previous approach as now we only consider one

Table 1. Experimental results for unsorted problems.

	VAMPIRE					
	CVC4	Paradox	iProver	Ignore	Use	Expand
FOF+CNF: sat	1181	1444	1348	1421	1463	1503
FOF+CNF: unsat	-	-	1337	1400	1604	1628

next domain size assignment. However, the SAT problems may be considerably harder to solve as the SAT solver is now considering a much larger set of models. In essence, each new SAT problem contains all the previous ones as sub-problems.

Finally, if the SAT problem is satisfiable then the actual size of a sort s is given by its smallest totality constraint that is “enabled”; more precisely, by the smallest i such that $bound_s(i)$ is false in the computed model.

9 Experimental Evaluation

In this section we evaluate the different techniques for finite model finding in multi-sorted FOL described in this paper and compare our approach to other tools.

Experimental Setup. We considered two sets of problems. From the TPTP [16] library (version 6.3.0) we took unsorted problems in the FOF or CNF format. From the SMT-LIB library [1] we took problems from the UF (Uninterpreted Functions) logic. Experiments were run on the StarExec cluster [15], whose nodes are equipped with Intel Xeon 2.4 GHz processors and 128 GB of memory. For each experiment we will report the number of problems solved with the time limit of 60 s.

On satisfiable problems we compare our implementation with version 3.0 of Paradox [5] and version 1.4 of CVC4 [13]; Paradox does not establish unsatisfiability and CVC4 runs more than a finite model finding approach, making a comparison on unsatisfiable problems difficult. On the TPTP problems, we also compare to version 2.0 of iProver [10]. The techniques described in this paper were implemented in VAMPIRE.

Adding Sorts to Unsorted Problems. Our first experiment considers the effect of sort inference on unsorted problems. We consider three settings: (i) inferred subsorts are ignored, (ii) inferred subsorts are used to reduce the problem size and break symmetries only, and (iii) inferred subsorts are expanded to proper sorts where possible. Table 1 presents the results. This shows that sort information can be used to solve more problems. For satisfiable problems the best VAMPIRE strategy solves more problems than CVC4, Paradox or iProver. For both satisfiable and unsatisfiable problems, expanding subsorts into proper sorts and treating the problems as multi-sorted problems helps solve the most problems. We note that 4 problems found unsatisfiable using this approach could not be solved by any other technique in VAMPIRE, this is significant as VAMPIRE is one of the best theorem provers available for such problems.

Table 2. Experimental results for translations from multi-sorted to unsorted.

	Sort Predicates				Sort Functions			
	Plain	Monotonicity	Subsorts	Both	Plain	Monotonicity	Subsorts	Both
UF: sat	813	810	872	874	710	771	834	873
UF: unsat	101	112	221	232	67	67	171	171

Table 3. Experimental results for multi-sorted problems.

	CVC4	Pointwise				Contour			Without Constraints
		Default	Expand	Collapse	Bounds	Default	Collapse	Bounds	
UF: sat	764 (8)	795	789	901 (12)	810	886 (3)	899 (1)	886 (1)	154
UF: unsat	-	212	215	241	218	270	261	267	66

Removing Sorts from Sorted Problems. Next we consider the translation techniques described in Sect. 4 applied to multi-sorted problems. Table 2 shows the results of running variations of these translations on the multi-sorted UF problems described above. Plain applies the translation to the whole problem, ignoring subsorts in the result. With Monotonicity only non-monotonic sorts are translated and with Subsorts the resulting problem is solved using inferred subsorts. Both adds both variations.

These results show that, for these problems, sort predicates are more useful and that the techniques of monotonicity detection and subsort inference are useful in improving the translation and reasoning with it.

Finding Models of Multi-sorted Problems. Finally, we consider our framework for reasoning with multi-sorted problems directly. Table 3 gives the results for CVC4 and eight variations of the techniques presented in this paper (we use only CVC4 since Paradox does not work on sorted problems). At the top level these are split into the Pointwise and Contour encodings and a version where no constraints were added. Then Expand refers to subsort expansion (Sect. 7), Collapse refers to collapsing monotonic sorts together (Sect. 5.5), and Bounds refers to sort bound extraction (Sect. 6). These results can also be compared to Table 2 as the problems are the same.

The three main conclusions from this information are (i) overall the approach taken in this paper is able to solve more problems than the approach taken by CVC4, (ii) collapsing monotonic sorts is very useful, and (iii) including the search problem as part of the SAT encoding is vital. Bracketed numbers show unique problems solved by an approach. This shows that although CVC4 solves fewer problems it does solve some uniquely. The contour encoding was generally more successful, however in UF there are 15 and 19 problems that are only solvable using the pointwise and contour encodings respectively. As a further point, we note that the heuristic introduced on page 9 is useful, without it the default pointwise approach solved 61 fewer problems. Finally, comparing with the results in Table 2, we see that finding models for multi-sorted problems directly performs better than translating the problem to an unsorted one.

10 Conclusions and Further Work

We have introduced a new framework for MACE-style finite model finding for multi-sorted first-order logic. This involved two complementary SAT encodings that capture the search for a satisfying domain size assignment and techniques aimed at decreasing the size of this search space. We have demonstrated experimentally that these techniques are effective at improving finite model finding in the unsorted setting and finding finite models for multi-sorted first-order formulas. Further work will consider possible extensions to uninterpreted sorts and infinite, but finitely representable, models.

References

1. Barrett, C., Stump, A., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB) (2010). <http://www.SMT-LIB.org>
2. Blanchette, J.C., Böhme, S., Popescu, A., Smallbone, N.: Encoding monomorphic and polymorphic types. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013 (ETAPS 2013). LNCS, vol. 7795, pp. 493–507. Springer, Heidelberg (2013)
3. Claessen, K., Lillieström, A.: Automated inference of finite unsatisfiability. *J. Autom. Reasoning* **47**(2), 111–132 (2011)
4. Claessen, K., Lillieström, A., Smallbone, N.: Sort it out with monotonicity. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE 2011. LNCS, vol. 6803, pp. 207–221. Springer, Heidelberg (2011)
5. Claessen, K., Sörensson, N.: New techniques that improve MACE-style model finding. In: CADE-19 Workshop: Model Computation - Principles, Algorithms and Applications (2003)
6. Eén, N., Biere, A.: Effective preprocessing in SAT through variable and clause elimination. In: Bacchus, F., Walsh, T. (eds.) SAT 2005. LNCS, vol. 3569, pp. 61–75. Springer, Heidelberg (2005)
7. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) SAT 2003. LNCS, vol. 2919, pp. 502–518. Springer, Heidelberg (2004)
8. Eén, N., Sörensson, N.: Temporal induction by incremental SAT solving. *Electr. Notes Theor. Comput. Sci.* **89**(4), 543–560 (2003)
9. Hoder, K., Kovács, L., Voronkov, A.: Case studies on invariant generation using a saturation theorem prover. In: Batyrshin, I., Sidorov, G. (eds.) MICAI 2011, Part I. LNCS, vol. 7094, pp. 1–15. Springer, Heidelberg (2011)
10. Korovin, K.: iProver – an instantiation-based theorem prover for first-order logic (System Description). In: Armando, A., Baumgartner, P., Dowek, G. (eds.) IJCAR 2008. LNCS (LNAI), vol. 5195, pp. 292–298. Springer, Heidelberg (2008)
11. Kovács, L., Voronkov, A.: First-order theorem proving and VAMPIRE. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 1–35. Springer, Heidelberg (2013)
12. Mccune, W.: A Davis-Putnam Program and its Application to Finite First-Order Model Search: Quasigroup Existence Problems. Technical report, Argonne National Laboratory (1994)
13. Reynolds, A., Tinelli, C., Goel, A., Krstić, S.: Finite model finding in SMT. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 640–655. Springer, Heidelberg (2013)

14. Schulz, S.: A comparison of different techniques for grounding near-propositional CNF formulae. In: Proceedings of the Fifteenth International Florida Artificial Intelligence Research Society Conference, May 14–16, 2002, Pensacola Beach, Florida, USA, pp. 72–76 (2002)
15. Stump, A., Sutcliffe, G., Tinelli, C.: StarExec, a cross community logic solving service (2012). <https://www.starexec.org>
16. Sutcliffe, G.: The TPTP problem library and associated infrastructure. *J. Autom. Reasoning* **43**(4), 337–362 (2009)
17. Tammet, T.: Reasoning. *Gandalf. J. Autom* **18**(2), 199–204 (1997)
18. Zhang, J., Zhang, H.: SEM: a system for enumerating models. In: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, IJCAI 95, Montréal Québec, Canada, August 20–25 1995, vol. 2s, pp. 298–303 (1995)