

Using Symmetries to Lift Satisfiability Checking

Pierre Carbonnelle¹, Gottfried Schenner², Maurice Bruynooghe¹, Bart Bogaerts³, Marc Denecker¹

¹KU Leuven, Belgium

²Siemens, Austria

³Vrije Universiteit Brussels, Belgium

{pierre.carbonnelle, maurice.bruynooghe, marc.denecker}@kuleuven.be,
gottfried.schenner@siemens.com, bart.bogaerts@vub.be

Abstract

We analyze how symmetries can be used to compress structures (also known as interpretations) onto a smaller domain without loss of information. This analysis suggests the possibility to solve satisfiability problems in the compressed domain for better performance. Thus, we propose a 2-step novel method: (i) the sentence to be satisfied is automatically translated into an equisatisfiable sentence over a “lifted” vocabulary that allows domain compression; (ii) satisfiability of the lifted sentence is checked by growing the (initially unknown) compressed domain until a satisfying structure is found. The key issue is to ensure that this satisfying structure can always be expanded into an uncompressed structure that satisfies the original sentence to be satisfied.

We present an adequate translation for sentences in typed first-order logic extended with aggregates. Our experimental evaluation shows large speedups for generative configuration problems. The method also has applications in the verification of software operating on complex data structures. Our results justify further research in automatic translation of sentences for symmetry reduction.

1 Introduction

In made-to-order manufacturing, the *configuration problem* is the problem of finding a configuration of components that satisfies the customer requirements and feasibility constraints (Felfernig et al. 2014). Such problems can be solved by choosing a formal vocabulary and by representing the customer requirements and the feasibility criteria as a logic sentence to be satisfied. A structure satisfying the sentence (a *model*) represents an acceptable configuration.

Methods to solve configuration problems do not scale well, and various heuristics have been used to improve performance (Schenner and Taupe 2017). Configurations often have components that are interchangeable. They are the source of many redundancies in the search space that negatively impact performance. The standard approach is to add symmetry breaking rules (Crawford et al. 1996). Here, we use another approach: we reformulate the problem to reduce symmetries (Gent, Petrie, and Puget 2006). This approach has been less studied, and is more an art than a science.

We observed that solutions to configuration problems can be compressed in what we call a “lifted model”, and that the solution of the original problem can be obtained by expanding the lifted model. This suggested that configuration problems, and, more generally, satisfiability problems, could be solved in the compressed domain: if the compressed domain is significantly smaller, it could lead to better performance.

As a toy example, consider the following pigeonhole problem: “given 10 pigeons and 5 pigeonholes, assign each pigeon to a pigeonhole such that each hole has (at most) 2 pigeons.” With the appropriate vocabulary, a solution is:

$$Pigeon = \{p_1, \dots, p_{10}\} \quad (1)$$

$$Hole = \{h_1, \dots, h_5\} \quad (2)$$

$$isIn = \{(p_1, h_1), (p_2, h_2), \dots, (p_5, h_5), \\ (p_6, h_1), (p_7, h_2), \dots, (p_{10}, h_5)\} \quad (3)$$

Note the symmetries: another solution is obtained by exchanging 2 pigeons or 2 holes in the interpretation of *isIn*.

This solution can be compressed to:

$$Pigeon = \{\mathbf{p}_1\} \quad (4)$$

$$Hole = \{\mathbf{h}_1\} \quad (5)$$

$$mul = \{\mathbf{p}_1 \mapsto 10, \mathbf{h}_1 \mapsto 5\} \quad (6)$$

$$isIn = \{(\mathbf{p}_1, \mathbf{h}_1)\} \quad (7)$$

where the *mul* function indicates how many concrete domain elements each “lifted” domain element represents, and thus allows the domain compression. The lifted pigeon (resp. hole) represents 10 concrete pigeons (resp. 5 concrete holes).

Even though the compressed structure only has 2 domain elements ($\mathbf{p}_1, \mathbf{h}_1$, excluding the naturals), it contains enough information to allow us to *expand* it into a model isomorphic to the original model. Following the theory we develop in Section 3, the expansion is as follows:

$$\mathbf{p}_1 \rightsquigarrow p_1, \dots, p_{10}$$

$$\mathbf{h}_1 \rightsquigarrow h_1, \dots, h_5$$

$$(\mathbf{p}_1, \mathbf{h}_1) \rightsquigarrow \{(p_1, h_1), \dots, (p_5, h_5), (p_6, h_1), \dots, (p_{10}, h_5)\}$$

Furthermore, in Section 4, we present a translation of sentences in typed first-order logic extended with aggregates into sentences that allow domain compression. We show that

a sentence in that language is equisatisfiable with its translation, given the number of concrete elements in each type: if either one is satisfiable, the other one is too. This analysis allows us to solve satisfiability problems in the compressed domain. The novel method operates in two steps: (i) the sentence to be satisfied is automatically translated into a “lifted” sentence over the “lifted” vocabulary; (ii) satisfiability of the lifted sentence is checked by step-wise growing the compressed domain¹ until a satisfying structure is found. Crucially, this satisfying structure can always be expanded into a model of the original sentence.

Notice that the lifted model for a variation of the pigeonhole problem above with 100 times more pigeons and holes has the same number of lifted domain elements as the base case: the multiplicity of each lifted domain element is simply multiplied by 100. As a result, and unlike traditional symmetry breaking methods, our method solves that pigeonhole problem in constant time with respect to the domain size (excluding the naturals).

We evaluate this method by comparing the time needed to find solutions for generative configuration problems discussed in the literature. Our method has significantly better performance than the traditional one for problems whose solution can be substantially compressed.

Our paper is structured as follows: after introducing our notation, we analyze how symmetries can be used to compress structures onto a smaller domain without loss of information (Section 3), and describe how to lift a concrete sentence so that it is equisatisfiable with the lifted sentence (Section 4); we describe the method, evaluate it on generative configuration problems (Section 5), and discuss applications in Boolean Algebra of sets with Presburger Arithmetic (BAPA) (Section 6) before concluding with a discussion.

2 Preliminaries

This section introduces the logic language supported by our method, and the concept of permutation of a structure.

Typed First Order Logic With Aggregates

We call FO(Type, Aggregate) the language we support. We assume familiarity with first order logic (Enderton 1972). A *vocabulary* Σ is a set of type, predicate, and function symbols. Predicates and functions have a type signature (e.g. $f : \bar{T} \rightarrow T$, where \bar{T} denotes a tuple of types). Some symbols are pre-defined: type \mathbb{B} (booleans), \mathbb{Q} (rationals), equality, arithmetic operators and arithmetic comparisons. *Terms* and *formulae* are constructed from symbols according to the usual FO syntactic rules. We also allow sum aggregates (written as $\sum_{\bar{x} \in \bar{T}: \phi} t$ or as $\sum_{\bar{x} \in \bar{T}} (t \text{ if } \phi)$). A cardinality aggregate $\#\{\bar{x} \in \bar{T} \mid \phi\}$ is a shorthand for a sum aggregate whose term t is 1. Quantification and aggregation can only be over finite types. Terms and formulas must be well-typed. A formula without free-variable is called a *sentence*.

A Σ -*structure* I , consists of a domain and an *interpretation* of each symbol of vocabulary Σ . The interpretations

¹An iterative method is required because the size of the compressed domain is not known in advance.

of types are disjoint and finite (except \mathbb{Q}^I which is infinite). The interpretation p^I of a predicate p is a set of tuples \bar{d} of domain elements of appropriate types. The interpretation f^I of a function f is a set of pairs (\bar{d}, d) , also denoted by $\bar{d} \mapsto d$.

An *extended structure* is a structure with a variable assignment, i.e., with a mapping from variable \bar{x} to values \bar{d} , denoted $[\bar{x} : \bar{d}]$. The value of a formula or term in an extended structure is computed according to the usual FO semantic rules, which require that the interpretation of function symbols be total functions over their domain. A Σ -structure is a *model* of a sentence if the value of the sentence is **true** in the structure, i.e., if it *satisfies* it. *Satisfiability checking* in typed FO logic is the problem of deciding whether a sentence has a model, given the interpretation of the types.

Permutations and Orbits

A *permutation* of a set is a bijection from that set to itself. We denote a permutation by π , and its inverse by π^{-1} . The *identity* permutation, π^0 , maps every element of the set to itself. The *order* of a permutation is the smallest positive number n such that π^n is the identity permutation. Hence, the n th permutation of an element is equal to its “ n th modulo the order” permutation.

Cycles are permutations that map elements in a cyclic fashion. A cycle is denoted by $(d^1 d^2 \dots d^n)$. A permutation over a finite set has a unique decomposition into *disjoint cycles*. Its order is the least common multiple (*lcm*) of the length of its cycles.

The permutation of a tuple of elements \bar{d} is denoted by $\pi(\bar{d})$ and is equal to $(\pi(d_1), \dots, \pi(d_n))$. The π -*orbit* of an element d (resp. of a tuple of elements \bar{d}) is the set of its repeated permutations $o_\pi(d) = \{\pi^i(d) \mid i \in \mathbb{N}\}$ (resp. $o_\pi(\bar{d}) = \{\pi^i(\bar{d}) \mid i \in \mathbb{N}\}$).

3 Lossless Compression of Structures

In this section, we discuss how, and when, a concrete structure with symmetries can be compressed into a lifted structure over a smaller domain without loss of information, and how a lifted structure can be expanded into a concrete one.

Symmetries in a structure I are described by a domain permutation π (Devriendt et al. 2016): it is a permutation of the domain of I that maps numbers to themselves, and other domain elements in \bar{T}^I to domain elements in T^I . Since numbers are mapped to themselves by the permutation, and since types besides \mathbb{Q} are finite, the permutation is composed of cycles.

A domain permutation *induces a structure transformation*. The interpretation in the transformed structure, denoted $\pi(I)$, is defined as follows:

- the type of domain element d in $\pi(I)$ is its type in I ;
- a tuple \bar{d}' is in the transformed interpretation of predicate symbol p if and only if $\pi^{-1}(\bar{d}')$ is in the original interpretation of p :

$$p^{\pi(I)} = \{\bar{d}' \mid \pi^{-1}(\bar{d}') \in p^I\} = \{\pi(\bar{d}) \mid \bar{d} \in p^I\};$$
- a tuple $(\bar{d}' \mapsto d')$ is in the transformed interpretation of function symbol f if and only if $\pi^{-1}(\bar{d}') \mapsto \pi^{-1}(d')$ is

in the original interpretation of f :

$$\begin{aligned} f^{\pi(I)} &= \{(\bar{d}' \mapsto d') \mid (\pi^{-1}(\bar{d}') \mapsto \pi^{-1}(d')) \in f^I\} \\ &= \{(\pi(\bar{d}) \mapsto \pi(d)) \mid (\bar{d} \mapsto d) \in f^I\} \end{aligned}$$

Definition 1 (Automorphism) A permutation π that transforms a structure into itself, i.e., such that $I = \pi(I)$, is an automorphism of the structure.

Every structure has at least one automorphism: the identity domain permutation. Note that an automorphism maps the interpretation of any constant to itself, i.e., $c^I() = \pi(c^I())$. Thus, the length of the cycle containing $c^I()$ is 1.

We introduce the new concept of *backbone*, which plays a critical role in the compression that we propose. Essentially, a *backbone* for an automorphism of I is a set of domain elements obtained by picking one element in each cycle, such that the interpretation of any symbol can be reconstructed from its interpretation restricted to the backbone, by applying the automorphism repeatedly. Formally:

Definition 2 (Backbone) A backbone for an automorphism π of I is a subset S of the domain of I such that:

- each cycle C of π has exactly one element in S ;
- for each predicate $p/n \in \Sigma$, p^I is the union of the π -orbits of the tuples in $p^I \cap S^n$, i.e.,

$$p^I = \bigcup_{\bar{d} \in p^I \cap S^n} o_\pi(\bar{d}) = \{\pi^i(\bar{d}) \mid \bar{d} \in p^I \cap S^n, i \in \mathbb{N}\} \quad (8)$$

- for each function $f/n \in \Sigma$, f^I is the union of the π -orbits of the tuples in $f^I \cap (S^n \times S)$, i.e.,

$$\begin{aligned} f^I &= \bigcup_{(\bar{d} \mapsto d) \in f^I \cap (S^n \times S)} o_\pi(\bar{d} \mapsto d) \\ &= \{\pi^i(\bar{d} \mapsto d) \mid (\bar{d} \mapsto d) \in f^I \cap (S^n \times S), i \in \mathbb{N}\} \end{aligned} \quad (9)$$

Example 1 For the pigeonhole example in the introduction of the paper, a backbone of automorphism $(p_1 \cdots p_{10})(h_1 \cdots h_5)$ is $S = \{p_1, h_1\}$. Another is $S = \{p_2, h_2\}$.

In all structures, the set of domain elements is a *trivial* backbone for the identity automorphism. However, not all automorphisms have a backbone.

Example 2 Let $\{a, b\}$ be a (concrete) domain with one type T . In structure I_1 (resp. I_2), function symbol $f : T \rightarrow T$ is interpreted as $\{a \mapsto a, b \mapsto b\}$ (resp. $\{a \mapsto b, b \mapsto a\}$). Permutation (ab) is an automorphism of both structures. The only two subsets S satisfying the first condition of backbone for (ab) are $\{a\}$ and $\{b\}$. Since f^{I_1} can be reconstructed from $f^{I_1} \cap (S \times S)$ for both subsets S , both subsets are backbones of I_1 . However, none is a backbone of I_2 (because $f^{I_2} \cap (S \times S) = \emptyset$ for both candidate sets S).

A backbone enables us to lift a structure into a structure with a smaller domain, as we now describe.

Definition 3 (Lifted vocabulary) For a vocabulary Σ , its lifted vocabulary Σ_l consists of the symbols of Σ and, for any type T of Σ , a symbol $mul_T : T \rightarrow \mathbb{N}$, called the multiplicity function for T .

We will drop the subscript T in the function symbol mul_T when this is unambiguous.

Definition 4 (Lifted structure) Let I be a Σ -structure with an automorphism π having backbone S . A lifted structure L derived from I is a Σ_l -structure such that:

- its domain is S , called the lifted domain;
- for each type predicate T , $T^L = T^I \cap S$;
- for each predicate symbol p/n , $p^L = p^I \cap S^n$;
- for each function symbol f/n , $f^L = f^I \cap (S^n \times S)$;
- for each $l \in S$, $mul^L(l) = |o_\pi(l)|$, the size of the π -orbit of l in I .

Example 3 Continuing the pigeonhole problem, the lifted structure is described by Equations 4-7.

Given the lifted structure L derived from I and the automorphism π on the concrete domain I , one can reconstruct I , i.e., one can *expand* the lifted interpretations of the type, predicate and function symbols, essentially by closing them under repeated application of π . Formally,

$$exp_\pi(T^L) = \bigcup_{l \in T^L} o_\pi(l) = \{\pi^i(l) \mid l \in T^L, i \in \mathbb{N}\} \quad (10)$$

$$exp_\pi(p^L) = \bigcup_{\bar{l} \in p^L} o_\pi(\bar{l}) = \{\pi^i(\bar{l}) \mid \bar{l} \in p^L, i \in \mathbb{N}\} \quad (11)$$

$$\begin{aligned} exp_\pi(f^L) &= \bigcup_{(\bar{l} \mapsto l) \in f^L} o_\pi(\bar{l} \mapsto l) \\ &= \{\pi^i(\bar{l} \mapsto l) \mid (\bar{l} \mapsto l) \in f^L, i \in \mathbb{N}\}. \end{aligned} \quad (12)$$

Example 4 Continuing the pigeonhole problem, the expansion of the lifted structure described by Equations 4-7 for automorphism $(p_1 \cdots p_{10})(h_1 \cdots h_5)$ is the structure described by Equations 1-3.

In our approach, we need to find lifted structures that can be expanded into concrete ones. We observe that there is a simple construction to expand any lifted Σ_l -structure L with lifted domain S that results sometimes (but not always) in a concrete structure I with an automorphism π having backbone S . This construction is as follows.

Definition 5 (Expansion of lifted domain) An expansion of the domain $S = \{l_i \mid i \in \mathbb{N}\}$ of the lifted Σ_l -structure L is a set $D = \{d_i^j \mid \exists i, j \in \mathbb{N} : l_i \in S \wedge 1 \leq j \leq Mul^L(l_i)\}$ such that the d_i^j are distinct and $\forall i : d_i^1 = l_i$.

We call d_i^1 a *base* element. Notice that lifted domain elements with multiplicity zero have no image in D .² We define a permutation on the set D having the cycles $(l_i, d_i^2 \dots d_i^{Mul^L(l_i)})$:

Definition 6 (Permutation of the expanded domain)

The permutation of an expansion D of a lifted domain is the function $\pi : D \mapsto D$ such that $\pi(d_i^j) = d_i^{j+1}$ for $1 \leq j < Mul^L(l_i) - 1$, and $\pi(d_i^{Mul^L(l_i)}) = d_i^1$.

²Null multiplicities will prove useful for iterative methods in Section 5

The expansion of a lifted tuple \bar{l} in a lifted structure is its π -orbit:

$$\text{exp}_\pi(\bar{l}) = o_\pi(\bar{l}) = \{\pi^i(\bar{l}) \mid i \in \mathbb{N}\} \quad (13)$$

if none of the expansion of its elements is empty, and is empty otherwise. The concrete interpretation of symbols is obtained by expanding the tuples in their lifted interpretations, as in Equations (10-12).

It follows from the correspondence between Definition 2 and Equations (10-12) that, if π and I are constructed from L in this way, and if I is actually a structure, then S is a backbone of π in I . When L only has strictly positive multiplicities, a lifted structure derived from the expansion of L will be isomorphic to L : in some sense, the compression is lossless. However, I may not be a structure because the interpretation of a function symbol in I might not be a total function on its domain, as the following example shows.

Example 5 Consider the lifted structure L with domain $S = \{a, b\}$, with types $A^L = \{a\}$, $B^L = \{b\}$, and for $f : A \times A \mapsto B$, $f^L = \{(a, a) \mapsto b\}$. Finally, let $\text{mul}(a) = \text{mul}(b) = 2$. The expanded domain is $\{a^1, a^2, b^1, b^2\}$ with permutation $(a^1 a^2)(b^1 b^2)$. The expansion of f^L is $\{(a^1, a^1) \mapsto b^1, (a^2, a^2) \mapsto b^2\}$. There is no entry for (a^1, a^2) , so the expansion is not a total function.

Consider now the lifted structure L with domain $S = \{a, b\}$, with type $A^L = \{a\}$, $B^L = \{b\}$, for $f : A \mapsto B$, $f^L = \{a \mapsto b\}$. Finally, $\text{mul}(a) = 1$, $\text{mul}(b) = 2$. The expanded domain is $\{a^1, b^1, b^2\}$ with permutation $(a^1)(b^1 b^2)$. The expansion of f^L is $\{a^1 \mapsto b^1, a^1 \mapsto b^2\}$. This expansion gives two different values for a^1 , so it is not a function.

To distinguish lifted structures that can be expanded into concrete ones from those that cannot, we introduce the notion of regularity.

Definition 7 (Regular function, regular lifted structure) The lifted interpretation of a function symbol is regular if its expansion defines a total function in the concrete domain, i.e., if it specifies exactly one image for every tuple in the concrete domain. A regular lifted structure is a lifted structure in which the interpretation of each function symbol is regular.

Now, we define the expansion of a regular lifted structure:

Definition 8 (Expansion of a regular lifted structure)

Let L be a regular lifted structure over a lifted vocabulary. Then, I , the expansion of L , is the structure over the concrete vocabulary defined as follows:

- its domain is the expansion of the lifted domain, having permutation π derived from L ;
- for each type symbol T , $T^I = \text{exp}_\pi(T^L)$;
- for each predicate symbol p , $p^I = \text{exp}_\pi(p^L)$;
- for each function symbol f , $f^I = \text{exp}_\pi(f^L)$.

The expansion of a regular lifted structure does not involve any search. The time needed for this expansion is generally negligible (e.g., less than 0.1 sec for 10,000 concrete domain elements).

To further characterize regular functions, we introduce the concept of regular tuples. We use $\text{Mul}(l_1, \dots, l_n)$ (resp.

$\text{Lcm}(l_1, \dots, l_n)$) as a shorthand for the product of the multiplicities of l_i (resp. their least common multiple lcm). First, we observe that the size of the expansion (Equation 13) of a tuple \bar{l} is finite: it is the order of the permutation defined by the cycles of its elements, i.e., $\text{Lcm}(l_1, \dots, l_n)$.³ Also, the expansion has at most $\text{Mul}(l_1, \dots, l_n)$ tuples; it is then the cross-product of the expansions of its elements. When those two numbers are identical, we say that the tuple is regular.

Definition 9 (Regular lifted tuple) A lifted tuple is regular if and only if its expansion is the cross-product of the expansion of its elements.

Example 6 Let (a, b) be a tuple of two lifted domain elements with $\text{mul}(a) = 2$ and $\text{mul}(b) = 4$. Its expansion is $\{(a^1, b^1), (a^2, b^2), (a^1, b^3), (a^2, b^4)\}$, of size $\text{Lcm}(2, 4) = 4$. Note that, e.g., the tuple (a^1, b^2) does not belong to the expansion: thus, (a, b) is not regular. It is regular when $\text{mul}(b) = 0$ (the expansion of b and of (a, b) are empty), or when, e.g., $\text{mul}(b) = 3$ (the expansion is $\{(a^1, b^1), (a^2, b^2), (a^1, b^3), (a^2, b^1), (a^1, b^2), (a^2, b^3)\}$, of size $\text{Lcm}(2, 3) = 6$).

Nullary and unary tuples are always regular. An n -ary tuple is regular when one of its elements has multiplicity zero, or every pair of its elements have multiplicities that are coprime.

Proposition 1 (Regular function) A function f^L is regular if, for all tuples \bar{l} in the domain of f^L , it holds that (i) \bar{l} is regular, and (ii) the multiplicity of \bar{l} is a multiple of the multiplicity of its image.

Proof 1 First, we show that the expansion of f^L gives at least one value for every tuple \bar{d} in the concrete domain of f . Each element d_i of \bar{d} is in the expansion of a lifted domain element $l_i \in T_i^L$. Tuple $\bar{l} = (l_1, \dots, l_n) \in \bar{T}^L$ is in the lifted domain of f^L as f^L is total; it is regular by (i), hence \bar{d} is in its expansion and $f^I(\bar{d})$ is in $\text{exp}_\pi(f^L(\bar{l}))$.

Next, we show that the expansion gives at most one value for every tuple \bar{d} in the concrete domain of f . Let \bar{d} be in the expansion of \bar{l} , with $\text{Lcm}(\bar{l}) = m = \text{Mul}(\bar{l})$ by (i). We thus have $0 < m$. The expansion of f^L contains the pairs $\pi^i(\bar{l} \mapsto f^L(\bar{l}))$ and $\pi^{i+n \times m}(\bar{l} \mapsto f^L(\bar{l}))$, for any \bar{l} in the domain of f^L , and for any i and n . The first element of these two pairs are identical by definition of m ; the second elements are identical by (ii).

4 Translation Into a Lifted Sentence

We now present a translation of a sentence in FO(Type, Aggregate) into a sentence that allows domain compression, such that the translation is equisatisfiable with the original.

The translation $\chi(\phi)$ of a sentence ϕ is the conjunction of (i) the transformed sentence $\chi(\phi)$, and (ii) sentences expressing regularity conditions.

The transformation $\chi(e)$ of an expression e is defined recursively in Table ???. The left column shows the possible syntactical forms in the concrete sentence; the middle column shows the transformation; the third column shows the

³Taking the convention that the lcm of a tuple of numbers containing 0 is 0.

regularity constraints added to the translation. The bottom part of the table shows the regularity constraints added for each function symbol in the vocabulary to ensure that the function interpretations are regular (Proposition 1).

Note that sum aggregates and quantified formulas are transformed by specialized rules (Rules 18, 26, 27) when possible (see table footnote), and by general rules (Rules 19, 28, 29) otherwise. Generally, it is beneficial to do equivalence-preserving transformations of sentences to obtain sentences of the form allowing application of specialized rules. The specialized rules do not require the translation of $p(\bar{t}, s)$ by Rule 21, thus avoiding the regularity constraint $RC(p(\chi(\bar{t}, s)))$. E.g., for the atom $t = s$, this regularity constraint would enforce $mul(t) = mul(s) = 1$ for each tuple (t, s) in the lifted equality relation, significantly reducing the possibility of compression.

The transformation of a sentence consists of adding $0 < Mul(\bar{x})$ filters (to cope with lifted domains elements with an empty expansion) and of multiplying the term t in an aggregate term with some decompression factor: in Rule 19, it is the number $Mul(\bar{x})$ of possible concrete variable assignments; in Rule 18, it is that number multiplied by the fraction of concrete assignments that make $p(\bar{t}, s)$ true. The regularity condition added for an atom (Rule 21) ensures the translated atom is equisatisfiable with the original atom.

Example 7 The sentence "at most 2 pigeons in each hole":

$$\begin{aligned} &\forall h \in Hole : \#\{p \in Pigeon \mid isIn(p, h)\} \leq 2. \\ \text{Its transformation by Rule 18 (with } \phi = \text{true) and 28 is:}^4 \\ &\forall h \in Hole : 0 < mul(h) \Rightarrow \\ &\quad \sum_{p \in Pigeon} \left(\frac{Lcm(p, h)}{mul(h)} \text{ if } 0 < mul(p) \wedge isIn(p, h) \right) \leq 2. \end{aligned}$$

Theorem 1 (Equisatisfiability) An $FO(\text{Type}, \text{Aggregate})$ sentence is equisatisfiable with its translation, given the number of concrete elements in each type.

If I is a model of the sentence, then L constructed by extending I by setting all multiplicities to one is a model of the translated sentence: indeed, the added constraints are trivially satisfied, and the translated sentence is equivalent to the original one. Proving the converse is long and complex. It is proved by structural induction of two invariants over the syntactic tree of the sentence to be satisfied: (i) a transformed formula $\chi(\phi)$ is true in L under some variable assignment $[\bar{x} : \bar{l}_x]$ if and only if ϕ is true in I under any variable assignments $[\bar{x} : \bar{d}_x]$ such that each d_{x_i} is in the expansion of l_{x_i} (ii) similarly, if a transformed term $\chi(t)$ has value l in the lifted structure L under some variable assignment $[\bar{x} : \bar{l}_x]$, then, the expansion of the value l contains the value of the term t in I (the expansion of L), for any variable assignment $[\bar{x} : \bar{d}_x]$ such that each d_{x_i} is in the expansion of l_{x_i} . This property holds only when the regularity constraints given in Table ?? hold in the lifted structure. This explains why these constraints are added to the transformed sentence. The proof is in the supplementary material (Carbannelle et al. 2023).

⁴Recall that a cardinality aggregate is a shorthand for a sum aggregate whose term is 1.

5 Evaluation of the Method

Implementation The goal of the evaluation is to show that there are satisfiability problems where substantial compression of the domain is possible, and that the lifted models can indeed be expanded into concrete ones. A problem is solved iteratively, starting with empty lifted domains. Given a domain, the lifted sentence is reduced to a propositional sentence and its satisfiability is determined with a standard satisfiability solver capable of arithmetic reasoning. If the sentence is unsatisfiable with this domain, the sentence is reduced to a minimal unsatisfiable formula (Lynce and Marques-Silva 2004), and the domains of the types used in that formula are extended with one element. This process is repeated until a model of the sentence is found (it does not terminate if the original sentence is unsatisfiable for any domain size, unless one imposes an upper limit on the size of lifted domains).

In many experiments, it was sufficient to support the special Rules 18, 26-27 for the case where the atom p is of the form $t_1 = s$ (e.g., $holeOf(p) = h$). Then, the transformation of a sum aggregate per Rule 18 simplifies to:

$$\frac{1}{mul(\chi(s))} \sum_{\bar{x} \in \bar{T}} (Mul(\bar{x}) \times \chi(t)) \quad (32)$$

if $0 < Mul(\bar{x}) \wedge \chi(t_1) = \chi(s) \wedge \chi(\phi)$

This formula does not use the lcm function, which is not supported natively by solvers. In other experiments, we used an interpretation table for lcm .

Problems are expressed in $FO(\cdot)^5$, a Knowledge Representation language with support for types, subtypes, and aggregates.⁶ A $FO(\cdot)$ (lifted) sentence is translated by IDP-Z3 (Carbannelle et al. 2022) for use with the Z3 SMT solver (de Moura and Bjørner 2008). The overhead of this translation is negligible.

The source code of our examples is available on GitLab⁷. Tests were run using Z3 v4.12.1 on an Intel Core i7-8850H CPU at 2.60GHz with 12 cores, running Ubuntu 22.04 with 16 GB RAM. We run a modified version of IDP-Z3 v0.10.8 on Python 3.10.

Pigeonhole problem To validate the approach, we first consider the satisfiability problem of assigning each pigeon to one pigeonhole, such that each pigeonhole holds (at most) 2 pigeons. Function $holeOf : Pigeon \mapsto Hole$ is used to represent this relation.

When there are twice as many pigeons as holes, the lifted solution has 1 lifted pigeon and 1 lifted hole, as shown in the introduction of the paper. As expected, the time needed to solve the lifted problem is almost constant when it is satisfiable. The correct multiplicities are found quickly by Z3 using a sub-solver specialized for arithmetic. For example, with 10,000 pigeons, the lifted sentence is solved in only

⁵<https://fo-dot.readthedocs.io/>

⁶Subtypes are subsets of types, and are declared as unary predicates in $FO(\cdot)$. Subtypes can be used where types are used: in the type signature of a symbol and in quantification.

⁷<https://gitlab.com/pierre.carbannelle/idp-z3-generative>

Expression e	Transformation $\chi(e)$	Regularity condition
Term		
x	x	(14)
$c()$	$c()$	(15)
$f(\bar{t})$	$f(\chi(\bar{t}))$	(16)
$t_1 \oplus t_2, \oplus \in \{+, -, \times, \div\}$	$\chi(t_1) \oplus \chi(t_2)$	(17)
$\sum_{\bar{x} \in \bar{T}: p(\bar{t}, s) \wedge \phi} t^{(*)}$	$\sum_{\bar{x} \in \bar{T}} \left(\frac{Mul(\bar{x}) Lcm(\chi(\bar{t}), \chi(s))}{Lcm(\chi(\bar{t}), mul(\chi(s)))} \times \chi(t) \text{ if } 0 < Mul(\bar{x}) \wedge p(\chi(\bar{t}), \chi(s)) \wedge \chi(\phi) \right)$	(18)
$\sum_{\bar{x} \in \bar{T}: \phi} t$	$\sum_{\bar{x} \in \bar{T}} (Mul(\bar{x}) \times \chi(t) \text{ if } 0 < Mul(\bar{x}) \wedge \chi(\phi))$	(19)
Formula		
$p()$	$p()$	(20)
$p(\bar{t})$	$p(\chi(\bar{t}))$	$RC(p(\chi(\bar{t})))^{(**)}$ (21)
$t_1 \sim t_2, \sim \in \{<, >, \leq, \geq\}$	$\chi(t_1) \sim \chi(t_2)$	(22)
$\phi_1 \otimes \phi_2, \otimes \in \{\wedge, \vee, \Rightarrow, \Leftrightarrow\}$	$\chi(\phi_1) \otimes \chi(\phi_2)$	(23)
$\neg \phi$	$\neg \chi(\phi)$	(24)
$T \triangleq \{t_1, \dots, t_n\}$	$\sum_{x \in T} mul(x) = n$	(25)
$\forall \bar{x} \in \bar{T} : p(\bar{t}, s) \Rightarrow \phi^{(*)}$	$\forall \bar{x} \in \bar{T} : 0 < Mul(\bar{x}) \wedge p(\chi(\bar{t}), \chi(s)) \Rightarrow \chi(\phi)$	(26)
$\exists \bar{x} \in \bar{T} : p(\bar{t}, s) \wedge \phi^{(*)}$	$\exists \bar{x} \in \bar{T} : 0 < Mul(\bar{x}) \wedge p(\chi(\bar{t}), \chi(s)) \wedge \chi(\phi)$	(27)
$\forall \bar{x} \in \bar{T} : \phi$	$\forall \bar{x} \in \bar{T} : 0 < Mul(\bar{x}) \Rightarrow \chi(\phi)$	(28)
$\exists \bar{x} \in \bar{T} : \phi$	$\exists \bar{x} \in \bar{T} : 0 < Mul(\bar{x}) \wedge \chi(\phi)$	(29)
for each function $f : \bar{T} \rightarrow T$		$\forall \bar{x} \in \bar{T} : Mul(\bar{x}) = Lcm(\bar{x})$ (30)
		$\forall \bar{x} \in \bar{T} : \exists n \in \mathbb{N} : Mul(\bar{x}) = n \times mul(f(\bar{x}))$ (31)

^(*) Rules (18, 26, 27) are applied only when $vars(\bar{t}) \subseteq \{\bar{x}\}$ and $vars(s) \cap \{\bar{x}\} = \emptyset$; Rules (19, 28, 29) are applied otherwise.

^(**) $RC(p(\chi(\bar{t})))$ is defined as $\forall \bar{x} \in \bar{T} : p(\chi(\bar{t})) \Rightarrow Mul(\bar{x}) = Lcm(\bar{x}) \vee Mul(\chi(\bar{t})) = Lcm(\chi(\bar{t}))$, where $\bar{x} = vars(p(\chi(\bar{t})))$.

Table 1: Terms, formulas, and their translation

0.05 sec and the expansion of the lifted model into a concrete one in 0.1 sec. With the same solver and the original sentence, the solution time increases exponentially (4 sec to solve the problem for 30 pigeons). We are aware that symmetry breaking can reduce the complexity, but to the best of our knowledge, solving time is at least linear in the number of pigeons for the best symmetry breaking solvers.

We also validated our method on pigeonhole problems where the relation between pigeon and holes is represented by a binary predicate. The translated sentence uses an interpretation table for lcm . Experiments confirm equisatisfiability, but the complexity is quadratic because of the lcm table. Finding a more efficient implementation has been left for future work.

Generative configuration problems Generative configuration problems (GCP) are configuration problems in which the number of some components has to be found: the number of elements in some types is not known in advance. An iterative method is thus always required to find them. When the compressed domain is smaller than the concrete domain, the number of iterations needed to solve the lifted sentence is smaller than the number of iterations for the original sentence, leading to better performance. Hence, GCP is a good application domain for our method.

We evaluate our solving method on three representative GCP discussed in the literature:

- the House Configuration and Reconfiguration problem (Friedrich et al. 2011),
- the Organized Monkey Village (Reger, Suda, and Voronkov 2016)
- the Rack problem (Feinerer, Salzer, and Sisel 2011; Comploi-Taupe, Francescutto, and Schenner 2022),

These problems are expressed using only an equality predicate and unary symbols, and most regularity constraints introduced by our method are trivially satisfied.

The top half of Table ?? shows results with Z3 (with automatic translation of the sentence), the bottom half with the OR-tools solver⁸ (with manual translation). A more detailed table is provided in the supplementary material (Carbonnelle et al. 2023).

Solutions to problems with higher suffixes have more components than similar problems with lower ones. The table shows near constant time performance on the Rack ABCD problems. The lifted methods solve each of the 20 occurrences of the ABCD Racks problem in (Comploi-

⁸<https://google.github.io/or-tools/>, used via CPMpy (Guns 2019).

Lifted?	Time		#	
	Lifted	Orig.	Lifted	Orig.
HCP 1	0.2	0.2	6	10
HCP 3	0.22	1.5	6	27
HRP 1	0.27	0.23	7	10
HRP 3	0.30	6.41	8	25
Monkey 1	0.28	184.66	6	20
Monkey 4	0.27	T	6	52
Rack 2011	0.24	8.82	5	23
Rack A5	0.46	0.86	5	19
Rack A10	0.41	37.52	5	38
Rack A1000	0.54	T	4	2250
Rack ABCD4	1.55	4.37	13	20
Rack ABCD8	2.69	T	13	40
Rack ABCD20	3.31	T	13	100
Rack A1	0.61	0.75	5	7
Rack A2	0.64	1.54	5	9
Rack A3	0.62	3.69	5	11
Rack A4	0.65	36.9	5	13
Rack A5	0.65	T	5	19

Table 2: Wall clock time in seconds to solve configuration problems, and number (#) of used domain elements in models. Yes: lifted sentence, No: original sentence, T: timeout after 200 sec.

Taupe, Francescutto, and Schenner 2022) in less than 5 seconds (instead of 6 minutes on average in that paper, using an ASP solver). These results show that our method has significant merits for solving problems with symmetries and a preponderance of unary symbols.

6 Boolean Algebra of Sets with Presburger Arithmetic

Lifted domain elements represent disjoint sets of concrete domain elements. A model search in the lifted domain can be seen as a model search involving sets. Hence, our work is highly related to Boolean Algebra of sets with Presburger Arithmetic (BAPA), a logic that can express constraints on the cardinality of sets, of their unions and of their intersections (Kuncak and Rinard 2007; Suter, Steiger, and Kuncak 2011; Bansal et al. 2016). Some problems from verification of properties of software operating on complex data structures contain fragments that belong to BAPA.

A sample BAPA statement is $|A| > 1 \wedge A \subseteq B \wedge |B \cap C| \leq 2$, where A, B, C are sets, and $|A|$ is the cardinality of A . The equivalent expression in FO(Type, Aggregate) is $(\#\{d : A(d)\} > 1) \wedge (\forall d : A(d) \Rightarrow B(d)) \wedge (\#\{d : B(d) \wedge C(d)\} \leq 2)$, where A, B, C are now unary predicates over a (unique) type whose interpretation is to be found. This expression can be lifted and solved using our approach (see the “theories/BAPA” folder in our repository⁷). In general, any BAPA sentence can be converted to a concrete FO(Type, Aggregate) sentence that only uses unary predicates, and that can be lifted without regularity constraints. Hence, our approach offers a simple way to solve BAPA problems using any solver capable of reasoning over the rationals. The performance advantage should be evaluated in future work.

On the other hand, the conversion of a concrete FO(Type, Aggregate) sentence to BAPA logic is challenging because FO(Type, Aggregate) is more expressive: it allows n-ary relations, functions, sum aggregates, and product of cardinalities. Thus, a BAPA solver could not be used to solve, e.g., generative configuration problems. Still, extensions of BAPA solvers to handle finite n-ary relations have been implemented in CVC4 (Meng et al. 2017).

A simple approach to represent structures in BAPA is to use disjoint subsets of the concrete domain, called *Venn regions*, so that the cardinalities of any set of interest is the sum of the cardinalities of its Venn regions. Unfortunately, the number of Venn regions grows exponentially with the number of sets of interest. Hence, various methods have been developed to reduce this growth, e.g., by creating new Venn regions lazily when required (Bansal et al. 2016). Venn regions are similar to our lifted domain elements, and the iterative method that creates new Venn region is similar, to some extent, to our iterative method that creates new lifted elements when required.

Our method cannot prove the unsatisfiability of a sentence. By contrast, efficient methods have been proposed to prove the unsatisfiability of BAPA sentences (Suter, Steiger, and Kuncak 2011).

7 Discussion

Unlike the traditional approach of adding symmetry breaking conditions to a formula to accelerate satisfiability checking, we *automatically translate* the formula to a form with fewer symmetries. Our results demonstrate the benefits of this approach for problems with symmetries and a preponderance of unary symbols, and justify further research in automatic translation for symmetry reduction.

Much work remains to be done in evaluating the method and determining problem areas where it is useful. Moreover, the relationship with BAPA logic is worth further exploration because, unlike our method, BAPA can identify unsatisfiable sentences and our method does not terminate for such problems. Efficient implementation of the *lcm* function is another area of research. Worth exploring is also the relevance for related areas such as model counting.

Ideally, any compression of a model of a sentence (Section 3) should be a model of the lifted sentence. It is unlikely that this is achievable, as it requires a translation that avoids all regularity constraints apart from those for functions. This requires being able to predict the fraction of concrete variable assignments (in the expansion of a lifted assignment) that make a concrete formula true, given that the translated formula is true with the lifted assignment. In particular, filter formulas (in an aggregate) having both free and quantified variables are problematic.

Still, many refinements of the translation in Table ?? are feasible. We have already worked out some refinements, but they cannot be presented within the current space constraints. They will be presented in the PhD thesis of one of the authors (Carbannelle 2024).

Acknowledgments

This research received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

References

- Bansal, K.; Reynolds, A.; Barrett, C.; and Tinelli, C. 2016. A new decision procedure for finite sets and cardinality constraints in SMT. In *International Joint Conference on Automated Reasoning*, 82–98. Springer.
- Carbannelle, P. 2024. *Standard, Interactive and Lifted Model Expansion using FO(·)*. Ph.D. thesis, KU Leuven, Belgium.
- Carbannelle, P.; Schenner, G.; Bruynooghe, M.; Bogaerts, B.; and Denecker, M. 2023. Using Symmetries to Lift Satisfiability Checking. arXiv:2311.03424.
- Carbannelle, P.; Vandeveld, S.; Vennekens, J.; and Denecker, M. 2022. IDP-Z3: a reasoning engine for FO(·). *CoRR*, abs/2202.00343.
- Comptoi-Taupe, R.; Francescutto, G.; and Schenner, G. 2022. Applying incremental answer set solving to product configuration. In et al., A. F., ed., *SPLC '22: 26th ACM Int'l Systems and Software Product Line Conference*, 150–155. ACM.
- Crawford, J. M.; Ginsberg, M. L.; Luks, E. M.; and Roy, A. 1996. Symmetry-Breaking Predicates for Search Problems. In Aiello, L. C.; Doyle, J.; and Shapiro, S. C., eds., *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR'96)*, Cambridge, Massachusetts, USA, November 5-8, 1996, 148–159. Morgan Kaufmann.
- de Moura, L.; and Bjørner, N. 2008. Z3: An Efficient SMT Solver. In et al., C. R. R., ed., *Tools and Algorithms for the Construction and Analysis of Systems, 14th Int'l Conference*, volume 4963 of *Lecture Notes in Computer Science*, 337–340. Springer.
- Devriendt, J.; Bogaerts, B.; Bruynooghe, M.; and Denecker, M. 2016. On local domain symmetry for model expansion. *Theory Pract. Log. Program.*, 16(5-6): 636–652.
- Enderton, H. B. 1972. *A mathematical introduction to logic*. Academic Press. ISBN 978-0-12-238450-9.
- Feinerer, I.; Salzer, G.; and Sisel, T. 2011. Reducing Multiplicities in Class Diagrams. In et al., J. W., ed., *MODELS 2011.*, volume 6981 of *Lecture Notes in Computer Science*, 379–393. Springer.
- Felfernig, A.; Hotz, L.; Bagley, C.; and Tiuhonen, J. 2014. *Knowledge-based configuration: From research to business cases*. Newnes.
- Friedrich, G.; Ryabokon, A.; Falkner, A. A.; Haselböck, A.; Schenner, G.; and Schreiner, H. 2011. (Re)configuration using Answer Set Programming. In et al., K. M. S., ed., *Proceedings of the IJCAI 2011 Workshop on Configuration*, volume 755 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Gent, I. P.; Petrie, K. E.; and Puget, J.-F. 2006. Symmetry in constraint programming. *Foundations of Artificial Intelligence*, 2: 329–376.
- Guns, T. 2019. Increasing modeling language convenience with a universal n-dimensional array, Cppy as python-embedded example. In *Proceedings of the 18th workshop on Constraint Modelling and Reformulation at CP (Modref 2019)*, volume 19.
- Kuncak, V.; and Rinard, M. C. 2007. Towards Efficient Satisfiability Checking for Boolean Algebra with Presburger Arithmetic. In Pfenning, F., ed., *CADE-21, 2007*, volume 4603 of *Lecture Notes in Computer Science*, 215–230. Springer.
- Lynce, I.; and Marques-Silva, J. 2004. On Computing Minimum Unsatisfiable Cores. In *SAT 2004 - The Seventh International Conference on Theory and Applications of Satisfiability Testing, 10-13 May 2004, Vancouver, BC, Canada, Online Proceedings*.
- Meng, B.; Reynolds, A.; Tinelli, C.; and Barrett, C. W. 2017. Relational Constraint Solving in SMT. In de Moura, L., ed., *CADE 26, 2017*, volume 10395 of *Lecture Notes in Computer Science*, 148–165. Springer.
- Reger, G.; Suda, M.; and Voronkov, A. 2016. Finding Finite Models in Multi-sorted First-Order Logic. In et al., N. C., ed., *Theory and Applications of Satisfiability Testing - SAT 2016*, volume 9710 of *Lecture Notes in Computer Science*, 323–341. Springer.
- Schenner, G.; and Taupe, R. 2017. Techniques for solving large-scale product configuration problems with ASP. In *Proceedings of the 19th International Configuration Workshop*, 12–19.
- Suter, P.; Steiger, R.; and Kuncak, V. 2011. Sets with Cardinality Constraints in Satisfiability Modulo Theories. In et al., R. J., ed., *VMCAI 2011*, volume 6538 of *Lecture Notes in Computer Science*, 403–418. Springer.