On Relational Analysis of Algebraic Datatypes

Viktor Kuncak and Daniel Jackson
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
{vkuncak,dnj}@mit.edu

MIT CSAIL Technocal Report No 985

Abstract

We present a technique that enables the use of finite model finding to check the satisfiability of certain formulas whose intended models are infinite. Such formulas arise when using the language of sets and relations to reason about structured values such as algebraic datatypes. The key idea of our technique is to identify a natural syntactic class of formulas in relational logic for which reasoning about infinite structures can be reduced to reasoning about finite structures. As a result, when a formula belongs to this class, we can use existing finite model finding tools to check whether the formula holds in the desired infinite model.

1 Introduction

A new kind of analysis has become popular in the last decade in which a system is examined by considering all small cases within some bound. The rationale is that flaws are revealed more readily by this method than by conventional testing: exhausting a huge space of small cases works better than considering a much smaller suite of cases, even if it includes larger ones.

Model checking is the preeminent example of this approach, and bounds the set of reachable states and sometimes also the length of execution traces. The success of model checking in hardware verification has generated great interest in applying it to software. Most model checkers, though, offer only rudimentary support for data structures, so most applications of model checking to software until now have focused on control properties, and data has either been ignored or abstracted away.

To handle data structures effectively within this context, a reduction to small cases is needed. With such a reduction, no special abstractions for data would be needed, and the same bounding mechanism used for trace length, for example, could be applied to the size of data structures.

How should data structures be represented in such an analysis? A relational representation is very attractive, because it fits both the analyses that are widely used at the low level, and the object-oriented view of a program at the high level. Symbolic model checkers such as SMV [18] already represent the state as a bit vector; the adjacency matrix representation of a relation is therefore easily integrated. In the object-oriented view of program state, the heap is a graph, with objects as nodes and fields of objects as edges between objects—in other words, a collection of relations, one per field. This view now predominates, because it's simple and easily accounts for sharing (a shared object simply being in the range of two relations).

An important question to ask, therefore, is whether this relational viewpoint can accommodate a general theory of data structures. Can arbitrary structural properties be naturally expressed and analyzed by the small case approach? This question is not only of theoretical interest. It has arisen repeatedly amongst advanced users of one tool, the Alloy language and its associated analyzer [1], as they have discovered scenarios in which Alloy's relational encoding does not seem to capture their intuitions about data structures.

This paper's aim is to resolve this issue, not only for Alloy, but more generally for any tool that relies on small case analysis of a relational encoding. This includes not only model checkers (such as SMV [18] and NuSMV [6]), but also specification analysis tools based on constraints (such as ProB [23] and the Bremen USE tool [10]), and indeed potentially to any tool that encodes data relationally.

To frame the problem rigorously, a characterization of data structures independent of the relational viewpoint is needed. For this purpose, we use the theory of algebraic datatypes, which corresponds to the way most programmers think about structured values in datastructures, and is the basis for their implementation in many current programming languages.

We start by explaining the standard encoding of algebraic datatypes using relations in first order logic. This encoding is faithful, but it suffers from a major drawback: it requires all models of a formula to be infinite. Consequently, an analysis based on finite cases cannot be applied. To remedy this, we remove the logical axiom that is responsible for making models infinite. Surprisingly, most analyses performed in the absence of this axiom are still sound. There are, however, analyses that will produce spurious counterexamples. The principal contribution of this paper is a simple syntactic criterion that guarantees that a formula being analyzed will not suffer from this problem. The criterion is easy to understand, and could be applied automatically by a tool, warning the user when an analysis on the relational encoding might produce results that do not correspond to the full theory of algebraic datatypes.

The contributions of the paper are:

- A recognition of the problem of handling data structures in relational encodings, with positive and nega-

Version compiled April 11, 2005, 10:45am.

tive illustrations;

- A rigorous formulation of the problem, in terms of encoding algebraic datatype axioms in first order logic, and invariance of formula semantics under the exclusion of the 'generator' axiom that generates infinite models;
- A simple and effective syntactic criterion characterizing the class of formulas for which an analysis involving only finite tests is sound and complete.

2 An Example

In this section, we motivate the problem with an example of a relational encoding of a simple algebraic datatype. We show how the omission of a 'generator' axiom can cause spurious counterexamples, but its inclusion results in inconsistency, making the encoding useless. The challenge is to determine under what conditions the axiom can be omitted while remaining faithful to the theory of algebraic datatypes.

The example will be given in Alloy [17], a modelling language based on a simple first-order logic with relational operators. Although our work was motivated by Alloy, our results apply more broadly, and the rest of the paper presents our theory in a standard logic that has no Alloy-specific features.

A datatype for lists would be declared in a language such as ML [31] like this:

```
datatype List = Nil | Cons of Element * List
```

where `List` is the datatype being declared, `Element` is the type corresponding to the elements, and `Nil` and `Cons` are *constructors*, with no arguments and two arguments respectively.

In Alloy, `List` and `Element` are represented as top-level sets (called 'signatures' in Alloy). `Nil` is a singleton set – the set containing the empty list.

```
sig Element {}
sig List {}
one sig Nil extends List {}
```

`Cons` is represented by a set `Cons` and two *selectors*, `elt` and `rest`:

```
sig Cons extends List {
    elt: Element,
    rest: List
}
```

The `extends` syntax makes `Cons` a subset of `List`, disjoint from `Nil`. The selectors are semantically just relations from the set `Cons` to the sets `Element` and `List` respectively.

The function `cons` can be written as an Alloy predicate

```
pred cons (e: Element, l: List, c: Cons) {
   c.elt = e and c.rest = l
}
```

which associates with an element `e` and a list `l` any object `c` in `Cons` such that `e` and `l` are the element and rest components of `c`.

Now let's consider checking some putative theorems. This assertion says that the element of a list just created with an application of cons is the element used in the application:

```
assert A {
  all e: Element, l: List, c: Cons |
    cons (e, l, c) => e = c.elt
}
```

This holds trivially—the consequent being contained in the hypothesis—and so, when when checked by the Alloy Analyzer, it yields no counterexamples. In contrast, suppose we check the assertion

```
assert B {
  all e: Element, l: List, c, c': Cons |
    (cons (e, l, c) and cons (e, l, c')) => c = c'
}
```

which claims that `cons` is deterministic. The Alloy Analyzer will give a counterexample such as this:

```
List = {L0, L1, L2}
Cons = {L1, L2}
Nil = {L0}
Element = {E0}
elt = {(L1, E0), (L2, E0)}
rest = {(L1, L0), (L2, L0)}
e = {E0}
l = {L0}
c = {L1}
c' = {L2}
```

in which `cons` produces the two lists `L1` and `L2` which are structurally identical but nevertheless distinct list objects.

This might be acceptable for some applications, but if we wanted to model the kind of list used in languages such as ML, in which equality is structural (and identity of cells therefore cannot be observed), we could add an axiom

```
fact {
  all l, l': List |
    l.elt = l'.elt and l.rest = l'.rest => l = l'
 }
```

requiring that structurally identical lists have the same identity, ensuring that the assertion `B` is now valid.

Continuing with model exploration, we notice that in some cases Alloy generates cyclic lists as counterexamples of our properties. Because cyclic lists are not algebraic datatypes, we introduce the fact

```
fact {
  no l : List | l in l.^rest
 }
```

ensuring that `elt` is an acyclic relation.

So far so good. Consider, however, an assertion claiming that `cons` is total:

```
assert C {
  all e: Element, l: List | some c: Cons |
    cons (e, l, c)
}
```

Given our intuition about algebraic datatypes, we expect this assertion to be valid. But in the relational setting, given the facts stated so far, the assertion `C` is actually invalid. The Alloy Analyzer will generate a counterexample such as:

```
List = {L0}
Cons = {}
Nil = {L0}
Element = {E0}
elt = {}
rest = {}
e = {E0}
l = {L0}
```

The problem, roughly speaking, is that the Analyzer is free to construct a counterexample in which there aren't enough lists.

Suppose, following our previous strategy, we attempt to add an axiom to rule out counterexamples of this form:

```
fact {
  all l: List, e: Element | some c: Cons |
      c.elt = e and c.rest = l
}
```

This 'generator' axiom ensures that the selector relations are complete: for any combination of list and element, it requires the existence of list for which they are components. This will indeed rule out the counterexample above. Unfortunately, however, it rules out *all* (nonempty) counterexamples, even to a manifestly false assertion (for example that 0 equals 1). The problem is that this axiom has no finite models, and is therefore *inconsistent* in a setting in which only finite models are considered.

The key idea of this paper is that a finite checker cannot incorporate this axiom, but must nevertheless be able to handle algebraic datatypes. The question then is what class of formulas have the same models whether or not this axiom is included. The contribution of this paper is a characterization of a class of such formulas that is both expressive and easily checked syntactically.

Assertion `C`, it turns, will not be in this class. The culprit is the innermost quantification. The nesting of quantifiers is not in itself problematic; rather, the problem is that the quantification is not bounded. If instead, it were bounded by an expression in terms of variables bound in an outer quantifier, no spurious counterexample would be generated, even in the absence of the generator axiom. For example, the assertion

```
assert D {
  all c: Cons | some l: c.*rest, e: Element |
    cons (e, l, c)
}
```

saying that each list is the result of an application of cons to some sublist, is valid, as expected (the expression `c.*rest` is the application of the reflexive transitive closure of `rest` to `c`, giving the set of `c`'s sublists).

3 Logic and Algebraic Datatypes

This section introduces our two formalisms: term algebras, a theory of algebraic datatypes, and a first-order logic with transitive closure. We show how a term algebra can be straightforwardly translated into first-order logic, using four axioms. One of these is a generator axiom that causes all models to be infinite. In the following section, we establish our main result about when this axiom can be omitted.

Throughout the paper, we use binary trees as the canonical example of algebraic datatypes. Being simple and familiar, trees serve well as a pedagogical example. In addition, because trees can represent all other algebraic datatypes, there is no loss of generality.

$$
\begin{array}{rcl}
S & ::= & S_V \mid S_1 \,\mathsf{setOp}\, S_2 \mid S.R \mid \{x_1, \ldots, x_n\} \\
R & ::= & R_V \mid R_1 \,\mathsf{setOp}\, R_2 \mid R_1.R_2 \mid \Delta \mid {}^\frown R \mid *R \\
\mathsf{setOp} & ::= & \cup \mid \cap \mid \setminus \\
A & ::= & (x_1, x_2) \in R \mid x \in S \mid \\
 & & S_1 = S_2 \mid R_1 = R_2 \mid x_1 = x_2 \\
B & ::= & A \mid B_1 \wedge B_2 \mid \neg B_1 \\
F & ::= & B \mid \forall x :: \mathsf{sort}.F \mid \exists x :: \mathsf{sort}.F \mid F_1 \wedge F_2 \mid \neg F_1 \\
\mathsf{sort} & ::= & \mathsf{Tree} \mid \mathsf{Object}
\end{array}
$$

Figure 1: A Logic with Transitive Closure

Algebraic datatypes and term algebras. We consider structures that contain two kinds of values, or 'sorts': 1) a Tree sort, corresponding to algebraic datatypes, and 2) an Object sort, corresponding to all remaining values. In a programming language such as ML [31], we would define this algebraic datatype using a declaration such as:

```
datatype Tree = Nil | Node of Tree * Object * Tree
datatype Object = Obj1 | Obj2 | ... | ObjN
```

Note that the datatype Tree has an infinite set of values, because there is no bound on the size of a tree. On the other hand, we assume that we have already finitized the set Object corresponding to the remaining values. If our structure had only values of sort Object, we could use existing techniques to search for models of formulas (such as [16]); our goal, however, is to reason about the structures that also contain values of the sort Tree. One view of the results of this paper is that we show how to effectively finitize algebraic datatypes, without making the conclusions derived in the finitization meaningless with respect to the intended world of arbitrarily large datatypes.

Algebraic datatypes have proven to be useful not only in programming languages, but also in model theory, where they correspond to term algebras [26, Chapter 23], [13, Section 1.3]. Term algebras are algebras in which values are interpreted syntactically: given a term t without free variables (a *ground term*), the interpretation of t is t itself. The term algebras corresponding to the Tree datatype are generated by:

1) a constant Nil of sort Tree, and

2) a binary constructor Node of type

$$\mathsf{Tree} \times \mathsf{Object} \times \mathsf{Tree} \to \mathsf{Tree}$$

We next present the logic we use to express the properties of structures containing terms and objects; we use this logic to write the formulas that we wish to analyze.

A logic with transitive closure. We consider a fragment of first-order logic with transitive closure. The syntax of our logic is in Figure 1: the nonterminal S denotes set-valued expressions, R denotes relation-valued expressions,

$$M = (T, O, \iota), \quad \alpha : \mathsf{Vars} \to T \cup O$$

$$\begin{aligned}
[\![S_1 \ \mathsf{setOp}\ S_2]\!]^{M,\alpha} &= [\![S_1]\!]^{M,\alpha}\ \mathsf{setOp}\ [\![S_2]\!]^{M,\alpha} \\
[\![S.R]\!]^{M,\alpha} &= \{y \mid \exists x \in [\![S]\!]^{M,\alpha}.\ (x,y) \in [\![R]\!]^{M,\alpha}\} \\
[\![\{x_1,\ldots,x_n\}]\!]^{M,\alpha} &= \{\alpha(x_1),\ldots,\alpha(x_n)\} \\
[\![R_1 \ \mathsf{setOp}\ R_2]\!]^{M,\alpha} &= [\![R_1]\!]^{M,\alpha}\ \mathsf{setOp}\ [\![R_2]\!]^{M,\alpha} \\
[\![R_1.R_2]\!]^{M,\alpha} &= [\![R_1]\!]^{M,\alpha} \circ [\![R_2]\!]^{M,\alpha} \\
&= \{(x,z) \mid \exists y.\ (x,y) \in [\![R_1]\!]^{M,\alpha} \wedge \\
&\qquad (y,z) \in [\![R_2]\!]^{M,\alpha}\} \\
[\![\Delta]\!]^{M,\alpha} &= \{(x,x) \mid x \in T\} \\
[\![S_1 = S_2]\!]^{M,\alpha} &= ([\![S_1]\!]^{M,\alpha} = [\![S_2]\!]^{M,\alpha}) \\
[\![R_1 = R_2]\!]^{M,\alpha} &= ([\![R_1]\!]^{M,\alpha} = [\![R_2]\!]^{M,\alpha}) \\
[\![x_1 = x_2]\!]^{M,\alpha} &= (\alpha(x_1) = \alpha(x_2)) \\
[\![(x_1,x_2) \in R]\!]^{M,\alpha} &= (\alpha(x_1),\alpha(x_2)) \in [\![R]\!]^{M,\alpha} \\
[\![x \in R]\!]^{M,\alpha} &= \alpha(x) \in [\![S]\!]^{M,\alpha} \\
[\![\widehat{\ }R]\!]^{M,\alpha} &= \{\ (x_0,x_n) \mid \exists n \geq 1. \exists x_1,\ldots,x_{n-1} \in T. \\
&\qquad \textstyle\bigwedge_{i=1}^{n} (x_{i-1},x_i) \in [\![R]\!]^{M,\alpha}\} \\
[\![*R]\!]^{M,\alpha} &= [\![\Delta]\!]^{M,\alpha} \cup [\![\widehat{\ }R]\!]^{M,\alpha} \\
[\![A_1 \wedge A_2]\!]^{M,\alpha} &= [\![A_1]\!]^{M,\alpha} \wedge [\![A_2]\!]^{M,\alpha} \\
[\![\neg A]\!]^{M,\alpha} &= \neg [\![A]\!]^{M,\alpha} \\
[\![\forall x :: \mathsf{Tree}.F]\!]^{M,\alpha} &= \forall t \in T.\ [\![F]\!]^{M,\alpha'}, \quad \alpha' = \alpha[x := t] \\
[\![\forall x :: \mathsf{Object}.F]\!]^{M,\alpha} &= \forall o \in O.\ [\![F]\!]^{M,\alpha'}, \quad \alpha' = \alpha[x := o] \\
[\![\exists x :: \mathsf{Tree}.F]\!]^{M,\alpha} &= \exists t \in T.\ [\![F]\!]^{M,\alpha'}, \quad \alpha' = \alpha[x := t] \\
[\![\exists x :: \mathsf{Object}.F]\!]^{M,\alpha} &= \exists o \in O.\ [\![F]\!]^{M,\alpha'}, \quad \alpha' = \alpha[x := o]
\end{aligned}$$

Figure 2: Semantics for Logic of Figure 1

A denotes atomic formulas, B denotes quantifier-free formulas, and F denotes general (potentially quantified) formulas. The non-terminal S_V denotes sets (corresponding to one-place predicates), whereas the non-terminal S_R denotes binary relations (corresponding to two-place predicates). Notation $\widehat{\ }r$ denotes the irreflexive transitive closure of the binary relation r, whereas $*r$ denotes the reflexive transitive closure of r. (Among the expressions that we intentionally omit are the universal set, and relation inverse. These constructs make it difficult to ensure certain locality properties of the expressions that are useful for the formulation of our result.) We use the shorthand $\exists^1 x.F$ for the formula

$$\exists x.F(x) \wedge (\forall x, y.\ F(x) \wedge F(y) \Rightarrow x = y).$$

We interpret formulas in our logic over two-sorted structures $M = (T, O, \iota)$ where T is the domain of the sort Tree, O is the domain of the sort Object, and ι with domain $S_V \cup R_V$ interprets the built-in sets and relations of the structure M, so that $\iota(s) \subseteq T$ or $\iota(s) \subseteq O$ if s is a built-in set, and $\iota(r) \subseteq T^2$, $\iota(r) \subseteq O^2$, $\iota(r) \subseteq T \times O$, or $\iota(r) \subseteq O \times T$ if r is a built-in relation. The standard model-theoretic semantics of our logic is in Figure 2. The function $\alpha : \mathsf{Vars} \to T \cup O$ is a valuation that maps each variable to its value. If φ is a sentence (a formula with no free variables) then $[\![\varphi]\!]^{M,\alpha}$ does not depend on α, so when φ is a sentence and $[\![\varphi]\!]^{M,\alpha} = \mathsf{true}$, we say that M is a *model* of the sentence φ. A structure M is a model of a set of sentences iff M is a model of each of the sentences in the set.

Note that, although we use set-theoretic notation such as $x \in S$ and $(x, y) \in R$, our logic does not allow quantification over sets, and is no stronger than first-order logic with transitive closure. Recall, however, that first-order logic is very expressive. Indeed, first-order logic has been used as a foundation for set theory and all of mathematics [30]. Note also that the axioms and definitions used to represent many mathematical problems in first-order logic are often made under the assumption that the logic is interpreted over infinite structures. We view the results of this paper as a contribution towards reasoning about infinite structures using techniques that have proven to be effective for finite structures.

A language for term algebras. The language in Figure 1 presents a general logic over arbitrary sets and relations. We next turn to the question of choosing the sets and relations that are appropriate for describing term algebras.

Term algebras can be described using *constructor* relations, such as Node, or using *selector* relations, which are the inverse of the constructors [13, Section 2.6]. For our purpose, it is more convenient to use selectors. Because we consider a binary tree, we use the selectors left, $\mathsf{content}$, and right, where left and right denote the children of a node in the tree, and $\mathsf{content}$ denotes the Object value associated with a tree node. We represent selectors as binary relations that are partial functions defined on non-Nil terms. We define the relation node as the following shorthand:

$$\begin{aligned}(t, t_1, o, t_2) \in \mathsf{node} \iff\ & (t, t_1) \in \mathsf{left} \wedge (t, o) \in \mathsf{content} \wedge \\ & (t, t_2) \in \mathsf{right} \wedge\ t \neq \mathsf{Nil}\end{aligned}$$

We also use the $\mathsf{subterm}$ relation, defined using transitive closure:

$$\mathsf{subterm} \stackrel{\text{def}}{=} \widehat{\ }(\mathsf{left} \cup \mathsf{right})$$

The term model. We are interested in checking the satisfiability of formulas over the term algebra structure $M_T = (T_T, O, \iota_T)$ given as follows:

- T_T is the set of ground terms generated by constants Nil and Node; in other words, T is the least set such that
  1. $\mathsf{Nil} \in T_T$, and
  2. if $t_1, t_2 \in T_T$ and $o \in O$, then $\mathsf{Node}(t_1, o, t_2) \in T_T$.
- O is a finite set;
- ι_T is defined as follows:

$$\iota_T(\mathsf{Nil}) = \mathsf{Nil}$$
$$\iota_T(\mathsf{left}) = \{(\mathsf{Node}(t_1, o, t_2), t_1) \mid t_1, t_2 \in T_T, o \in O\}$$
$$\iota_T(\mathsf{content}) = \{(\mathsf{Node}(t_1, o, t_2), o) \mid t_1, t_2 \in T_T, o \in O\}$$
$$\iota_T(\mathsf{right}) = \{(\mathsf{Node}(t_1, o, t_2), t_2) \mid t_1, t_2 \in T_T, o \in O\}$$

Figure 3 sketches one part of the structure M_T.

Axioms for term algebras. We adopt the following axioms to describe the properties of of term algebras:

- *Selectors*: The binary relations left, $\mathsf{content}$, and right are total functions on the non-Nil elements of the sort Tree, and are undefined on Nil:

$$\begin{aligned} 1.\quad & \forall t :: \mathsf{Tree}.\ t \neq \mathsf{Nil} \Rightarrow (\exists^1 t_1 :: \mathsf{Tree}.\ (t, t_1) \in \mathsf{left}) \wedge \\ & \qquad (\exists^1 o :: \mathsf{Object}.\ (t, o) \in \mathsf{content}) \wedge \\ & \qquad (\exists^1 t_1 :: \mathsf{Tree}.\ (t, t_1) \in \mathsf{right}) \\ 2.\quad & \forall t_1 :: \mathsf{Tree}. \forall o :: \mathsf{Object}.\ (\mathsf{Nil}, t_1) \notin \mathsf{left} \wedge \\ & \qquad (\mathsf{Nil}, o) \notin \mathsf{content} \wedge \\ & \qquad (\mathsf{Nil}, t_1) \notin \mathsf{right} \end{aligned}$$

  We assume that a simple type system of our two-sorted language rules out the application of relations to elements of inappropriate sort; for example, if $t :: \mathsf{Tree}$ and $o :: \mathsf{Object}$, then $(t, o) \in \mathsf{left}$ is not a well-formed formula.

- *Uniqueness*: The defined relation node has the properties of a partial function:

$$\begin{aligned} & \forall t, t', t_1, t_2 :: \mathsf{Tree}. \forall o :: \mathsf{Object}. \\ & \quad (t, t_1, o, t_2) \in \mathsf{node} \wedge (t', t_1, o, t_2) \in \mathsf{node} \Rightarrow t = t' \end{aligned}$$

- *Generator*: The defined relation node has the properties of a total function:

$$\forall t_1, t_2 :: \mathsf{Tree}. \forall o :: \mathsf{Object}.\ \exists t :: \mathsf{Tree}.\ (t, t_1, o, t_2) \in \mathsf{node}$$

  (This axiom holds in M_T, but we will consider the consequences of omitting it from the axiomatization.)

- *Acyclicity*: A term is never a proper subterm of itself; that is, the subterm relation is acyclic:

$$\forall t :: \mathsf{Tree}.\ (t, t) \notin \mathsf{subterm}$$

We denote by SUGA the conjunction of the axioms above (taking the first letter of the name of each axiom).

Note that the SUGA axioms have no finite models. Namely, although not all models of SUGA are isomorphic, they all contain an infinite chain of elements $t_0, t_1, t_2, \ldots$ where $t_0 = \mathsf{Nil}$ and $(t_{i+1}, t_i, o, \mathsf{Nil}) \in \mathsf{node}$. These elements exist by the *Generator* axiom; the *Acyclicity* axiom guarantees that they are all distinct because they are ordered by the $\mathsf{subterm}$ relation.

In first-order logic without transitive closure, term algebras have a complete axiomatization [26, Chapter 23], [25]: there is a set of first-order sentences whose consequences are precisely the sentences that are true in the structure M_T (this set of sentences is infinite and requires some axiom schemas). However, even a complete axiomatization does not characterize the models up to isomorphism. For example, our SUGA axioms allow countable models with countably infinite paths of left and right that never terminate at Nil. The completeness of the axiomatization of term algebras is not of direct interest to us in any case, because a complete axiomatization forces the model to be infinite. Instead, we look for subclasses of formulas that can be checked on finite structures, and we show the soundness of our technique using a model-theoretic approach: we look at the truth value of the formulas in the desired term model M_T (as opposed to checking whether the formulas are a consequence of an axiomatization of M_T as in a proof-theoretic approach).

4 Finite Satisfiability Result

This section presents the main results of our paper, which enable the checking of properties of algebraic datatypes using finite models. The basic idea of our approach is the following: to prevent all models from being infinite, we drop the *Generator* axiom. Denote by SUA the conjunction of the remaining axioms (*Selectors*, *Uniqueness*, *Acyclicity*). It turns out that, among the finite structures, SUA characterizes precisely the substructures of the term model that are subterm-closed (if t is in the structure, then so is each subterm of t). Having proved this characterization, we identify a class of sentences whose validity in a finite model implies the validity in the full infinite model M_T.

We next define the notion of a sub-term closed finite substructure of M_T, illustrated in Figure 3. Intuitively, a finite substructure $M_0 = (T_0, O, \iota)$ of M_T is a structure obtained from M_T by selecting a finite set T_0 of terms and preserving all the relations between the terms in T_0. A structure is subterm-closed if a subtree of each tree in T_0 is also in T_0. More precisely, we have the following.

Consider a term model $M_T = (T_T, O, \iota_T)$. A *substructure of* M_T is a structure $M_0 = (T_0, O, \iota_0)$ where $T_0 \subseteq T_T$ and the relations given by ι_0 are restrictions of the corresponding relations given by ι_T, that is, $\iota_0(\mathsf{left}) = \iota_T(\mathsf{left}) \cap T_0^2$, $\iota_0(\mathsf{right}) = \iota_T(\mathsf{right}) \cap T_0^2$, and $\iota_0(\mathsf{content}) = \iota_T(\mathsf{content}) \cap T_0 \times O$. (We have for simplicity assumed that substructures have the same domain O of values of the sort Object.) A *subterm-closed finite substructure of* M_T is a finite substructure M_0 of M_T whose domain of terms T_0 satisfies the property $t \in T_0 \ \wedge\ (t, t_1) \in [\![\mathsf{subterm}]\!]^{M_T, \alpha} \Rightarrow t_1 \in T_0$ for all $t, t_1 \in T$.

We then have the following completeness theorem that explains why the SUA axioms are adequate. This theorem allows us to ensure that any model of SUA axioms has precisely the properties of a subterm-closed finite substructure

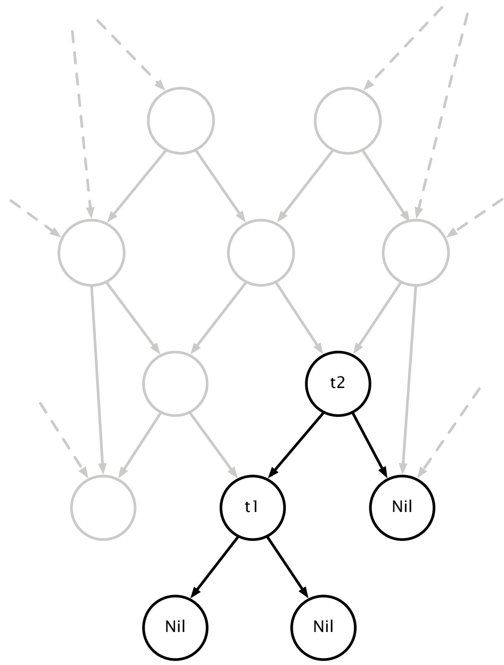

Figure 3: Sketch of the infinite structure M_T and an example of one of its finite subterm-closed substructures $M_0 = (T_0, \{o\}, \iota_0)$ where $T_0 = \{\mathsf{Nil}, t_1, t_2\}$ for $t_1 = \mathsf{Node}(\mathsf{Nil}, o, \mathsf{Nil})$ and $t_2 = \mathsf{Node}(t_1, \mathsf{Nil})$. The edges corresponding to the content relation are omitted for clarity.

of the term model M_T.

Theorem 1 (Axiomatization of Finite Substructures) *A two-sorted structure M is a model of* SUA *iff M is isomorphic to some subterm-closed finite substructure M_0 of M_T.*

The proof of Theorem 1 is in Section 8.

Having identified that the SUA axioms enforce that a finite structure "looks like" a term algebra M_T, we turn to the question of when is the case that checking a property in a finite model is sufficient to ensure that the property holds in the full model M_T. We first look at sentences that contain only existential quantifiers.

An *existential sentence* φ is a formula of the form

$$\exists t_1 :: \mathsf{Tree}.\ \ldots\ \exists t_n :: \mathsf{Tree}.\ \psi$$

where ψ is quantifier-free and the free variables of ψ are $t_1, \ldots, t_n$. The following is a fundamental property of structures (dual to the property that if a universal sentence holds in a structure then it also holds in its substructure).

Fact 1 *Let M_1 be a substructure of M and φ a purely existential sentence. If φ holds in M_1, then φ also holds in M.*

Informally, the property holds because the same existential witnesses from the smaller structure can be used in the larger structure.

An important consequence of Theorem 1 and Fact 1 is the following: if φ is an existential sentence and the conjunction $\mathsf{SUA} \land \varphi$ holds in a finite two-sorted model M_1, then φ holds in some subterm-closed finite substructure M_0 of M_T, so φ holds in the full term model M_T. We thus obtain a method to check whether a formula φ holds in M_T. In the rest of this section, we generalize this result by allowing arbitrary quantification in φ, as long as it is bounded by previously introduced values. The intuitive reason why this generalization is possible is that we are checking formulas on subterm-closed structures, which means that the bounded quantification has the same semantics in the substructure M_0 and in the full infinite structure M_T.

Let S be a set-valued term, denoted S in Figure 1, and suppose that S does not contain variable t. A *bounded universal term quantifier*

$$\forall_S\, t :: \mathsf{Tree}.\, F$$

is a shorthand for the formula

$$\forall t :: \mathsf{Tree}.\ t \in S \ \Rightarrow\ F$$

Dually, a *bounded existential term quantifier*

$$\exists_S\, t :: \mathsf{Tree}.\, F$$

is a shorthand for the formula

$$\exists t :: \mathsf{Tree}.\ t \in S \ \land\ F$$

Hence, bounded quantifiers are expressible in terms of the ordinary quantifiers, but are more restrictive. An *existential—bounded-universal sentence* requires each universal quantifier to be bounded by some set expression S. More precisely, we have the following definition:

Definition 1 *An* existential—bounded-universal sentence *(an* EBU *sentence) is a formula of the form*

$$Q_1 v_1 :: s_1.\ \ldots\ Q_n v_n :: s_n.\ \psi$$

where ψ is a quantifier-free formula (denoted B in Figure 1) and each $Q_i v_i :: s_i$ is a quantifier or a bounded quantifier of one of the following forms:

- *An existential term quantifier $\exists v_k ::$ Tree;*
- *A bounded universal term quantifier $\forall_S\, v_k ::$ Tree where the free variables of the set-valued term S are among the previously quantified variables $v_1, \ldots, v_{k-1}$;*
- *A bounded existential term quantifier $\exists_S\, v_k ::$ Tree where the free variables of the set-valued term S are among the previously quantified variables $v_1, \ldots, v_{k-1}$ (this quantifier is a special case of the existential term quantifier);*
- *A universal object quantifier $\forall v_k ::$ Object;*
- *An existential object quantifier $\exists v_k ::$ Object.*

We write $\varphi \in$ EBU to denote that φ is an EBU *sentence.*

Note the ways in which EBU sentences generalize purely existential sentences: not only is it possible to have arbitrary bounded quantifiers, it is also possible to introduce new unrestricted existential quantifiers, even after bounded universal quantifiers.

We are now ready to state our main theorem:

procedure TermSat
input: φ: an EBU sentence
output: if φ is true in some term model M_T:
a finite substructure M_0 of M_T where φ holds
if φ is false in all term models M_T:
no result (infinite loop)
begin
k := 1;
while (true) **do**
for each model $M_1 = (T, O, \iota)$ with $|T| + |O| = k$ **do**
if ($\mathsf{SUA} \wedge \varphi$) holds in M_1 **then return** M_1; **fi**
end
k := k + 1;
end
end TermSat.

Figure 4: A semidecision procedure for checking satisfiability of EBU sentences in the term model M_T.

Theorem 2 (Finite Satisfiability Theorem) *Let φ be an EBU sentence and M_T a term model. Then φ holds in M_T iff it holds in some subterm-closed finite substructure M_0 of M_T.*

The proof of Theorem 2 is in Section 8.

The identification of EBU sentences and the proof that they can be verified on finite models is the main contribution of this paper. In the sequel, we explore some of the consequences of this result.

5 Consequences

Given the results of the previous section, we can now answer the question we posed in Section 2. An analysis of an Alloy model in which algebraic datatypes are encoded relationally will yield sound counterexamples so long as the formula being checked is in EBU. A syntactic check for membership in EBU based on Definition 1 is easy to implement. The result extends to bounded model checkers (such as NuSMV [6]) whose analysis consists of finding a model of a formula. The language of formulas for such model checkers could be soundly extended to EBU sentences over algebraic datatypes.

Analysis procedure. The analysis procedure suggested by our result is shown in Figure 4: to check whether an EBU sentence φ holds in M_T, search for increasingly large finite models of $\mathsf{SUA}\wedge\varphi$. The procedure captures the spirit of analyses such as that performed by the Alloy Analyzer [16, 17]; in practice, the search for models would employ pruning and heuristics and would not require an exhaustive enumeration. The correctness and completeness of the procedure follows from the results of this section and is the main result of this paper:

- we can check the condition that a structure is a subterm-closed submodel of M_T by simply conjoining axioms SUA to φ, thanks to Theorem 1;
- we know that the existence of the returned finite model M_1 implies that φ holds in M_T, thanks to the ($\Longleftarrow$) (soundness) direction of Theorem 2;
- we know that if φ holds in model M_T, then the algorithm will find a finite model M_1 which proves this fact, thanks to the ($\Longrightarrow$) (completeness) direction of Theorem 2.

Closure under boolean operations. Having identified EBU sentences as a useful class of formulas for which the algorithm in Figure 4 is applicable, we next examine the following question. If $\varphi_1, \varphi_2 \in \mathsf{EBU}$, is there an effectively constructible sentence $\varphi \in \mathsf{EBU}$ such that the following equivalences hold in M_T:

- $\varphi \iff (\varphi_1 \wedge \varphi_2)$
- $\varphi \iff (\varphi_1 \vee \varphi_2)$
- $\varphi \iff (\neg\varphi_1)$
- $\varphi \iff (\varphi_1 \Rightarrow \varphi_2)$

It turns out that the answer to first two questions is "yes", whereas the answer to the last two questions is "no". In other words, EBU sentences are closed only under positive boolean combinations, but are not closed under negation or implication. We make this claim precise using the following two propositions.

Proposition 1 *Let $\varphi_1 \equiv \mathrm{BQ}_1.F_1$ and $\varphi_2 \equiv \mathrm{BQ}_2.F_2$ be EBU sentences where BQ_1 and BQ_2 denote sequences of quantifiers and bounded quantifiers and where F_1, F_2 are quantifier-free formulas. Let $\mathrm{BQ}_2'.F_2'$ be the result of renaming the variables in φ_2 so that they are all distinct from the variables in φ_1. Then*

$$\begin{aligned} \varphi_1 \wedge \varphi_2 &\iff \mathrm{BQ}_1.\mathrm{BQ}_2'.\, F_1 \wedge F_2' \\ \varphi_1 \vee \varphi_2 &\iff \mathrm{BQ}_1.\mathrm{BQ}_2'.\, F_1 \vee F_2' \end{aligned} \tag{1}$$

Moreover, $\mathrm{BQ}_1.\mathrm{BQ}_2'.F_1 \wedge F_2'$ and $\mathrm{BQ}_1.\mathrm{BQ}_2'.F_1 \vee F_2'$ are EBU sentences.

The condition (1) follows from the basic monotonicity property of quantifiers and operations $\wedge, \vee$. The fact that the concatenation of disjoint EBU sequences of quantifiers is again an EBU sequence of quantifiers follows from the definition of EBU sentences.

We next turn to the absence of the closure under negation and implication. We first note that the entire class of EBU sentences is undecidable.

Fact 2 *The problem of determining, given an EBU sentence φ, whether φ holds in M_T, is undecidable.*

Fact 2 follows from the fact that the subterm relation is definable in our logic, and from the undecidability result in [39, Section 4], which shows a reduction from the Post correspondence problem to the satisfiability of sentences with existential quantifiers and bounded universal quantifiers.

Fact 2 has two main consequences for this paper. The first consequence is that, from the viewpoint of computability, the semidecision procedure in Figure 4 is as good as we can hope for. The second consequence is the absence of closure under negation and implication, as given by the following proposition.

Proposition 2 *There is no effective algorithm that, given a sentence $\varphi \in \mathsf{EBU}$, constructs a sentence in EBU equivalent to $\neg\varphi$. Consequently, there is no such algorithm that constructs a sentence equivalent to $\varphi \Rightarrow \mathsf{false}$.*

The following is an indirect argument for Proposition 2: suppose that there is such an algorithm. Consider any EBU sentence φ. Let $\bar{\varphi}$ be the sentence computed by the algorithm, so that $\bar{\varphi}$ is equivalent to φ. Then either φ or $\bar{\varphi}$ holds in M_T, so if we run two copies of the procedure in Figure 4 in parallel, one with the input φ and the other one with the input $\bar{\varphi}$, then one of the algorithms will eventually terminate and we will conclude that φ is either true or false. This implies that the class of EBU formulas is decidable, contradicting Fact 2.

Arbitrary algebraic datatypes. We have presented our result for binary trees, but it applies to all algebraic datatypes, and, more generally, to any structured data. Indeed, it is easy in relational logic to reason about records and tuples that have an *a priory* bounded number of components: just introduce a new variable for each component. What the results of this paper imply is that we can also reason about structures, such as binary trees, that do not have an *a priory* bound on their size. It is not difficult to generalize the proofs of Theorem 1 and Theorem 2 to the case of any finite number of mutually recursive algebraic datatypes. Alternatively, we can encode any number of datatypes using binary trees. (Indeed, the experience with programming languages such as LISP [28] is convincing evidence that data structures can be represented using LISP-like lists, which are binary trees.) The idea of representing algebraic datatypes with trees is to replace each constructor application $C_k(t_1, \ldots, t_n, o_1, \ldots, o_m)$ with an expression

$$\begin{array}{l}\mathsf{Node}(f_k, o_0,\\ \quad \mathsf{Node}(t_1, o_1, \mathsf{Node}(t_2, o_2, \ldots \mathsf{Node}(t_P, o_P, \mathsf{Nil}) \ldots)))\end{array}$$

where f_k is a finite tree (of size $O(\log k)$) that encodes the name of the constructor C_k, where $P = \max(n, m)$, $t_i = \mathsf{Nil}$ if $i > n$, and $o_i = o_0$ if $i > m$. Here $o_0 \in O$ is some arbitrary fixed object from the set of uninterpreted objects. The corresponding selector relations are similarly definable using quantifier-free formulas in terms of selectors left and right, and so is the subterm relation. Note that, when reasoning about arbitrary algebraic datatypes, we are interested not in all possible trees, but only in the substructure of M_T which is the image of the embedding. In other words, we would like to ensure that the binary trees that represent the values of variables in formulas are consistent with the type system of the original algebraic datatypes. Luckily, this condition is expressible using our logic with transitive closure. Therefore, it suffices to restrict all quantified variables to the terms that satisfy this condition, and the resulting formula can be checked using the algorithm in Figure 4.

The scope of our result. After realizing that our technique applies to algebraic datatypes, a natural question to ask is: does the technique fundamentally depend on the properties of the structure of algebraic datatypes, such as the uniqueness of left and right relations (as given by the *Selectors* axiom), the uniqueness of the parent relation node (as given by the *Uniqueness* axiom), or even the acyclicity (given by the *Acyclicity* axiom)? When examining this question it is worthwhile to consider two separate questions:

- How do we generalize the notion of subterm-bounded substructures of M_T to the case of substructures of some other infinite structure M_∞ of interest? (The generalization of Theorem 2.)

  Suppose that we are interested in checking constraints over an infinite structure M_∞ with relation symbols $r_1, \ldots, r_n$. It turns out that the only essential requirement on the structure M_∞ is that, for some term variable t, the set $[\![\{t\}.{*}(r_1 \cup \ldots \cup r_n)]\!]^{M_\infty, \alpha}$ is finite for each valuation α. In other words, as long as the set of elements "below" each element of M_∞ is finite, we can use bounded quantification to reduce the truth value of EBU sentences in M_∞ to the satisfiability in finite substructures closed under the "below" relation. In particular, the technique applies to structures that contain shared elements and cycles.

- How do we axiomatize a class of finite structures of interest? (The generalization of Theorem 1.)

  From an algorithmic point of view, this question admits a wider spectrum of solutions than just the use of axioms in first-order logic with transitive closure (although the use of axioms may have an advantage in the context of constraint-solving tools). Indeed, given a family of finite structures of interest (in particular, given a family of finite subterm-closed substructures of M_∞) we can use any language of computable functions to define an executable test predicate that determines whether a finite structure is isomorphic to one of the finite structures of interest. In other words, we can use an algorithm specialized for a given problem to filter the finite structures of interest. This idea of using "executable predicates" appears in the form of run-time assertions in many programming languages and has found applications in software testing [5, 27].

Because of these generalizations, we expect our result to be applicable to a range of infinite structures.

6 Related Work

Constraint-checking tools. Because of its full automation, model checking approaches based on finitization of the problem space are very attractive. These approaches have had great success for control-intensive problems [6, 18] such as those arising in hardware verification. The complexity of software systems often comes from the data structures that they manipulate, and notations such as UML [35] have been used to describe such constraints. The Alloy notation [16,17] can also be used to describe such constraints; the Alloy Analyzer tool [1] can then search for the structures that satisfy these constraints. Our experience in using the Alloy notation and the analyzer to reason about structured values was the immediate inspiration for this paper. Because it establishes a general correspondence between satisfiability in finite and infinite models, our result is potentially applicable not only to Alloy, but also to tools such as MACE [29], Paradox [7], USE [10], ProB [23], RACER [40], and FaCT [14].

Algebraic datatypes. Our paper uses algebraic datatypes as a well-studied example of unbounded structured values. Algebraic datatypes are the basis of the algebraic approach to formal specification and verification [2, 4, 11, 12]. The use of the list algebraic datatype was pioneered by LISP [28]. User-defined algebraic datatypes go back to ML [31] and are used in variants such as Haskell [3] and Objective Caml [22].

Term algebras without transitive closure. The first-order theory of term algebras is decidable [25, 26, 37]. Because the interpretation of Object is a finite set, omitting the transitive closure from our logic makes formulas decidable even with arbitrary (not only bounded) quantifiers. The complexity of the resulting decision problem is non-elementary [8, 9] with the height of the tower of exponentials linear in the number of quantifier alternations in the formula [42]. More tractable classes of term algebras include the class of quantifier-free formulas [33]. Several decidable extensions of term algebras have been proved decidable [20, 21, 36, 38, 41], mostly using quantifier elimination techniques.

Term algebras with a subterm relation. Adding a subterm relation to the first-order language of term algebras makes the problem substantially more difficult. Indeed, [39] shows that even the satisfiability of formulas with bounded universal quantifiers is undecidable (although the satisfiability of the purely existential fragment with a subterm relation is still decidable). As we noted in Section 5, the undecidability result for algebras with subterms applies to our logic as well, because the subterm relation is expressible using transitive closure. A search for counterexamples is useful even for an undecidable logic (and is, in fact, at least as important as the search for counterexamples in decidable logics), and the results of this paper show how to perform such search for a useful class of formulas.

First-order logic with transitive closure. First-order logic with transitive closure is useful for reasoning about program data structures and has been used not only in Alloy [17], but also in shape analysis tools such as TVLA [24] and PALE [32]. Among the decidable fragments with transitive closure are monadic second-order logic [19,34] and some subclasses of the existential monadic second-order logic of graphs [15].

Complexity and bounded quantification. In their study of lower and upper bounds on the complexity of logical theories, Ferrante and Rackoff [9, Page 30] describe the notion of H-bounded structures for some function H, which enables a reduction of general quantifiers to bounded quantifiers. The existence of appropriate such function implies the decidability of the structure, so such H does not exists for term algebras with subterm relation. Our use of bounded quantification is different: we have syntactically imposed boundedness of universal quantifiers and showed that it implies the ability to use finite structures to reason about certain classes of formulas in infinite structures.

7 Conclusions

The language of sets and relations has proven to be a very powerful notation for modelling a range of structures arising in software design and analysis. Model finding tools have made this approach accessible and practical. So far, model finding tools have been restricted to arbitrarily large, but finite models. However, some useful structures are inherently infinite, in particular the algebraic datatypes such as lists and trees. Such structures are widely used in implementations and models of software, but when we try to apply existing tools to these structures, we are faced, in general, with either ruling out all models (which is sound, but entirely useless), or allowing the possibility that the tool returns unsound, meaningless models that do not apply to the desired infinite structures.

In this paper we have presented a useful and natural class of formulas for which the existence of a finite model reveals the satisfiability of the formula in the infinite structure. For this class of properties, we have proved that it is possible to partially axiomatize the desired structure in such a way that finite models are simply substructures of the desired infinite structure. In this way, concrete feedback from model finding tools can be brought to a range of ubiquitous data structures that would otherwise remain out of their scope.

References

[1] Alloy Analyzer web page. `http://alloy.mit.edu/`.

[2] Egidio Astesiano, Michel Bidoit, Hélène Kirchner, Bernd Krieg-Brückner, Peter D. Mosses, Donald Sannella, and Andrzej Tarlecki. CASL: The Common Algebraic Specification Language. *TCS*, 286(2):153–196, 2002.

[3] Richard Bird. *Introduction to Functional Programming using Haskell*. Prentice-Hall, Inc., 2nd edition, 1998.

[4] Adel Bouhoula, Jean-Pierre Jouannaud, and Jose Meseguer. Specification and proof in membership equational logic. *Theoretical Computer Science*, 236(1-2):35–132, 2000.

[5] Chandrasekhar Boyapati, Sarfraz Khurshid, and Darko Marinov. Korat: Automated testing based on Java predicates. In *Proc. International Symposium on Software Testing and Analysis*, pages 123–133, July 2002.

[6] A. Cimatti, E.M. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: a new Symbolic Model Verifier. In N. Halbwachs and D. Peled, editors, *Proceedings Eleventh Conference on Computer-Aided Verification (CAV'99)*, number 1633 in Lecture Notes in Computer Science, pages 495–499, Trento, Italy, July 1999. Springer.

[7] Koen Claessen and Niklas Sörensson. New techniques that improve MACE-style model finding. In *Model Computation*, 2003.

[8] Kevin J. Compton and C. Ward Henson. A uniform method for proving lower bounds on the computational complexity of logical theories. *Annals of Pure and Applied Logic*, 48(1):1–79, July 1990.

[9] Jeanne Ferrante and Charles W. Rackoff. *The Computational Complexity of Logical Theories*, volume 718 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979.

[10] Martin Gogolla, Jörn Bohling, and Mark Richters. Validating UML and OCL models in USE by automatic snapshot generation. *Journal on Software and System Modeling*, 2005. to appear.

[11] John V. Guttag. Abstract data types, then and now. In *Software Pioneers: Contributions to Software Engineering*, pages 442–452. Springer, 2002.

[12] John V. Guttag, Ellis Horowitz, and David R. Musser. Abstract data types and software validation. *Communications of the ACM*, 21(12):1048–1064, 1978.

[13] Wilfrid Hodges. *Model Theory*, volume 42 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1993.

[14] I. Horrocks. Using an expressive description logic: FaCT or fiction? In *KR'98*, pages 636–647, 1998.

[15] Neil Immerman, Alexander Moshe Rabinovich, Thomas W. Reps, Shmuel Sagiv, and Greta Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *CSL*, pages 160–174, 2004.

[16] Daniel Jackson. Automating first-order relational logic. In *Proc. ACM SIGSOFT Conf. Foundations of Software Engineering*, 2000.

[17] Daniel Jackson. Alloy: A lightweight object modelling notation. *ACM TOSEM*, 11(2):256–290, 2002.

[18] J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang. Symbolic Model Checking: 10^{20} States and Beyond. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science*, pages 1–33, Washington, D.C., 1990. IEEE Computer Society Press.

[19] Nils Klarlund, Anders Møller, and Michael I. Schwartzbach. MONA implementation secrets. In *Proc. 5th International Conference on Implementation and Application of Automata*. LNCS, 2000.

[20] Viktor Kuncak and Martin Rinard. On the theory of structural subtyping. Technical Report 879, Laboratory for Computer Science, Massachusetts Institute of Technology, 2003.

[21] Viktor Kuncak and Martin Rinard. Structural subtyping of non-recursive types is decidable. In *Eighteenth Annual IEEE Symposium on Logic in Computer Science*, 2003.

[22] Xavier Leroy. *The Objective Caml system, release 3.08*, July 2004.

[23] Michael Leuschel and Michael J. Butler. ProB: A model checker for B. In *FME*, pages 855–874, 2003.

[24] Tal Lev-Ami and Mooly Sagiv. Tvla: A system for implementing static analyses. In *Proc. 7th International Static Analysis Symposium*, 2000.

[25] Michael J. Maher. Complete axiomatizations of the algebras of the finite, rational, and infinite trees. *Proc. 3rd IEEE LICS*, 1988.

[26] Anatolii Ivanovic Mal'cev. *The Metamathematics of Algebraic Systems*, volume 66 of *Studies in Logic and The Foundations of Mathematics*. North Holland, 1971.

[27] Darko Marinov. *Automatic Testing of Software with Structurally Complex Inputs*. PhD thesis, MIT, 2005.

[28] J. McCarthy. Recursive functions of symbolic expressions and their computation by machine, part 1. *Comm. A.C.M.*, 3:184–195, 1960.

[29] W. McCune. MACE 2.0 Reference Manual and Guide. *ArXiv Computer Science e-prints*, June 2001.

[30] Elliott Mendelson. *Introduction to Mathematical Logic*. Chapman & Hall, London, 4th edition, 1997.

[31] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. The MIT Press, Cambridge, Mass., 1997.

[32] Anders Møller and Michael I. Schwartzbach. The Pointer Assertion Logic Engine. In *Proc. ACM PLDI*, 2001.

[33] Derek C. Oppen. Reasoning about recursively defined data structures. *Journal of the ACM*, 27(3), 1980.

[34] M. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.*, 141:1–35, 1969.

[35] James Rumbaugh, Ivar Jacobson, and Grady Booch. *The Unified Modelling Language Reference Manual*. Addison-Wesley, Reading, Mass., 1999.

[36] Tatiana Rybina and Andrei Voronkov. A decision procedure for term algebras with queues. *ACM Transactions on Computational Logic (TOCL)*, 2(2):155–181, 2001.

[37] Thomas Sturm and Volker Weispfenning. Quantifier elimination in term algebras: The case of finite languages. In V. G. Ganzha, E. W. Mayr, and E. V. Vorozhtsov, editors, *Computer Algebra in Scientific Computing (CASC)*, TUM Muenchen, 2002.

[38] Ralf Treinen. Feature trees over arbitrary structures. In Patrick Blackburn and Maarten de Rijke, editors, *Specifying Syntactic Structures*, chapter 7, pages 185–211. CSLI Publications and FoLLI, 1997.

[39] K. N. Venkataraman. Decidability of the purely existential fragment of the theory of term algebras. *Journal of the ACM (JACM)*, 34(2):492–510, 1987.

[40] Ralf Möller Volker Haarslev. RACER system description. In *IJCAR'01*, 2001.

[41] Ting Zhang, Henny B. Sipma, and Zohar Manna. Decision procedures for recursive data structures with integer constraints. In *2nd IJCAR*, volume 3097 of *LNCS*, pages 157–167, 2004.

[42] Ting Zhang, Henny B. Sipma, and Zohar Manna. Term algebras with length function and bounded quantifier alternation. In *TPHOLs'04*, volume LNCS, pages 321–336, 2004.

8 Proofs of Theorems

Theorem 1. *A two-sorted structure M is a model of* SUA *iff M is isomorphic to some subterm-closed finite substructure M_0 of M_T.*

Proof. We prove both directions of the equivalence.

$\Longleftarrow$): Suppose that a structure M is isomorphic to a subterm-closed model M_0 of M_T. Then M satisfies same formulas as M_0. Therefore, it suffices to verify that M_0 satisfies the SUA axioms *Selectors*, *Uniqueness*, *Acyclicity*. Axioms *Uniqueness* and *Acyclicity* hold in M_T, so they hold in M_0 as well: indeed, a relation that has two values in a substructure M_0 also has two values in the larger structure M_T, and a cycle in M_0 is also a cycle in M_T. Axiom *Selectors* holds because M_0 is subterm-closed: the components of every non-Nil term t in T_0 are also in T_0.

$\Longrightarrow$): Suppose that a finite two-sorted structure $M = (T, O, \iota)$ satisfies SUA axioms. We identify a subterm-closed finite structure $M_0 = (T_0, O, \iota_0)$ isomorphic to M by establishing a relation $f \subseteq T \times T_T$ and showing that g is an isomorphism where $g = f \cup \Delta_0$ and Δ_O is the identity relation on O. We define f using the following least fixpoint construction. Let $f_0 = \{(\iota(\mathsf{Nil}), \mathsf{Nil})\}$ and let

$$\begin{aligned} f_{i+1} = f_i \cup \{\ & (t', \mathsf{Node}(t_1, o, t_2)) \in f \mid \\ & (t', t'_1) \in \iota(\mathsf{left}), (t', t'_2) \in \iota(\mathsf{right}), \\ & (t', o) \in \iota(\mathsf{content}), \\ & (t'_1, t_1), (t'_2, t_2) \in f_i\} \end{aligned}$$

Then define $f = \cup_{i \geq 0} f_i$. In other words, we map $\iota(\mathsf{Nil})$ to Nil and we extend the relation by following parent relation in both M and M_0.

We next define a measure on the elements of structures M and M_0. Consider first an element $t \in T$ of structure M and consider any sequence of elements $t_0, t_1, \ldots$ such that $t_0 = t$ and $(t_i, t_{i+1}) \in \iota(\mathsf{left}) \cup \iota(\mathsf{right})$. Because M satisfies *Acyclicity* and T is finite, the sequence is finite. Moreover, because of the axiom *Selectors*, the sequence terminates at the element $\iota(\mathsf{Nil})$. For each element t, let $d(t)$ be the maximum of the lengths of all such sequences. We correspondingly define $d(t)$ for $t \in T_T$ of the structure M_0.

We then prove by induction on i the conjunction of the following properties:

P1) $\mathsf{dom}(f_i) = \{t' \in T \mid d(t') \leq i\}$

P2) each relation $g_i = f_i \cup \Delta_O$ is a partial isomorphism, that is, that g_i is an isomorphism between structures induced by the domain of g_i (denoted $\mathsf{dom}(g_i)$) and the range of g_i (denoted $\mathsf{ran}(g_i)$);

Base case. g_0 is trivially a partial isomorphism because $\mathsf{dom}(r) = \iota(\mathsf{Nil})$ and $\mathsf{ran}(r) = \mathsf{Nil} = \iota_0(\mathsf{Nil})$, so P1 holds. Moreover, from *Selectors* it follows that if $d(t) = 0$ then $\iota(t) = \mathsf{Nil}$, so P1 and P2 also hold.

Inductive step. Suppose that g_i satisfies P1 and P2; we show that g_{i+1} satisfies these properties as well.

- P1. Let $t' \in \mathsf{dom}(f_{i+1})$. By definition of f_{i+1}, there exist $t'_1, t'_2 \in \mathsf{dom}(f_i)$ such that $(t', t'_1) \in \iota(\mathsf{left})$ and $(t', t'_2) \in \iota(\mathsf{right})$. By inductive hypothesis, $d(t'_1) \leq i$ and $d(t'_2) \leq i$. By axiom *Selectors*, there are no elements x of T other than t'_1, t'_2 such that $(t, x) \in \iota(\mathsf{left}) \cup \iota(\mathsf{right})$. Therefore, $d(t) \leq 1 + \max(d(t'_1), d(t'_1)) \leq i+1$. Conversely, let $t' \in T$ be such that $d(t') \leq i+1$. If $d(t') \leq i$ then $t' \in \mathsf{dom}(f_i) \subseteq \mathsf{dom}(f_{i+1})$, so let $d(t') = i+1$. Then $t' \neq \mathsf{Nil}$ so by *Selectors* there exist unique elements $t'_1, t'_2 \in T$ and $o \in O$ such that $(t', t'_1) \in \iota(\mathsf{left})$, $(t', t'_2) \in \iota(\mathsf{right})$ and $(t', o) \in \iota(\mathsf{content})$. Because $1 + \max(d(t'_1), d(t'_1)) = d(t) \leq i+1$, we have $d(t'_1) \leq i$ and $d(t'_2) \leq i$. By induction hypothesis, $t'_1, t'_2 \in \mathsf{dom}(f_i)$. Again by induction hypothesis, f_i is a partial isomorphism, so there exist terms $t_1, t_2 \in T_T$ such that $(t'_1, t_1) \in f_i$ and $(t'_2, t_2) \in f_i$. By definition of f_{i+1} we have $(t', \mathsf{Node}(t_1, o, t_2)) \in f_{i+1}$, so $t' \in \mathsf{dom}(f_{i+1})$.
- P2, functionality. For $t' \in \mathsf{dom}(f_i)$ the property follows by inductive hypothesis, because f_{i+1} does not add new values to elements that are already in f_i. So let $(t', t), (t', \bar{t}) \in f_{i+1} \setminus f_i$. Then for both (t', t) and (t', t) there are elements $t'_1, t'_2 \in T$ such that $(t', t'_1) \in \iota(\mathsf{left})$ and $(t', t'_2) \in \iota(\mathsf{right})$, and by *Selectors* these elements are unique. Moreover, by induction hypothesis, f_i is functional, so t'_1, t'_2 are related via f_i to unique elements $t_1, t_2 \in T_T$. Therefore, $\mathsf{Node}(t_1, o, t_2)$ is the unique element t such that $(t', t) \in f_{i+1}$.
- P2, injectivity. By definition of f_{i+1}, there is exactly one element t with the property $(t, \mathsf{Nil}) \in f_{i+1}$, namely $\iota(\mathsf{Nil})$. Hence, injectivity can be violated only on non-Nil terms. Consider $t = \mathsf{Node}(t_1, o, t_2) \in \mathsf{ran}(f_{i+1})$. A tuple (t', t) is in f_{i+1} only if there are some $t'_1, t'_2 \in T$ such that $(t'_1, t_1), (t'_2, t_2) \in f_i$, $(t', t'_1) \in \iota(\mathsf{left})$, and $(t', t'_2) \in \iota(\mathsf{right})$. By induction hypothesis, such t'_1, t'_2 are unique because f_i is a partial isomorphism. Finally, by *Uniqueness*, there is at most one t' such that $(t', t'_1) \in \iota(\mathsf{left})$ and $(t', t'_2) \in \iota(\mathsf{right})$, so t' is unique.
- P2, Nil preservation. Clearly $(\iota(\mathsf{Nil}), \mathsf{Nil}) \in f_0 \subseteq f_{i+1}$, so interpretation of Nil is mapped to the interpretation of Nil.
- P2, left preservation. Let $(t', t), (t'_1, t_1) \in f_{i+1}$. We show that $(t', t'_1) \in \iota(\mathsf{left})$ iff $(t, t_1) \in \iota_0(\mathsf{left})$. If $t', t'_1 \in \mathsf{dom}(f_i)$, the property holds by induction hypothesis, so suppose that $t' \notin \mathsf{dom}(f_i)$ or $t'_1 \notin \mathsf{dom}(f_i)$.

  Suppose first $(t', t'_1) \in \iota(\mathsf{left})$. Then $t' \in \mathsf{dom}(f_{i+1} \setminus f_i)$ so by definition of f_{i+1} there exists $o \in O$, $\bar{t}'_1 \in T$ and $\bar{t}_1, \bar{t}_2 \in T_T$ such that $(t', \bar{t}'_1) \in \iota(\mathsf{left})$, $(\bar{t}'_1, \bar{t}_1) \in f_i$, and $(t', \mathsf{Node}(\bar{t}_1, o, \bar{t}_2) \in f_{i+1}$. By axiom *Selectors*, $\bar{t}'_1 = t'_1$, so $(t'_1, \bar{t}_1) \in f_i$. We have shown above that f_{i+1} is functional, so $\bar{t}_1 = t_1$. Furthermore, $t = \mathsf{Node}(t_1, o, \bar{t}_2)$. By definition of ι_T, $(t, t_1) \in \iota_T(\mathsf{left})$, as desired.

  Conversely, suppose that $(t, t_1) \in \iota_T(\mathsf{left})$. By definition of ι_T, this means there are $o \in O$ and $t_2 \in T_T$ such that $t = \mathsf{Node}(t_1, o, t_2)$. By definition of f_{i+1}, there are $\bar{t}', \bar{t}'_1 \in T$ such that $(\bar{t}', t) \in f_{i+1}$, $(\bar{t}'_1, t_1) \in f_{i+1}$ and $(\bar{t}', \bar{t}'_1) \in \iota(\mathsf{left})$. We have shown that f_{i+1} is injective, so $\bar{t}' = t'$ and $\bar{t}'_1 = t'_1$. Hence, we have $(t', t'_1) \in \iota(\mathsf{left})$, as desired.
- P2, right preservation. Analogous to the previous case.

This completes the inductive step. Given that properties hold for all i, let $n = \max\{d(t') \mid t' \in T\}$. Then $\mathsf{dom}(f_n) = T$ and $g_{n+1} = g_n$, so $g = g_n$. Let $M_0 = (T_0, O, \iota_0)$ where $T_0 = \mathsf{ran}(f_n)$ and $\iota_0(\mathsf{content}) = \{(t, o) \mid \exists t' \in T.\, (t', o) \in \iota(\mathsf{content})\}$. Then g is a bijection $T \cup O \to T_0 \cup O$, it preserves left and right because f_n does, and it preserves $\mathsf{content}$ by construction. Therefore, M_0 is the desired model and g is the desired isomorphism. ■

Theorem 2. *Let φ be an* EBU *sentence and M_T a term model. Then φ holds in M_T iff it holds in some subterm-closed finite substructure M_0 of M_T.*

Proof. We prove both directions of the equivalence.

$\Longleftarrow$) : Suppose that φ holds in a subterm-closed finite substructure $M_0 = (T_0, O, \iota_0)$ of $M_T = (T_T, O, \iota_T)$. When evaluating φ in M_T, for any witness for an existential quantifier we can pick the same witness in M_T as in M_0, because $T_0 \subseteq T_T$. Moreover, regardless whether they are interpreted in M_0 or M_T, the universal quantifiers range only over elements of T_0, so they still hold in M_0. We next make this argument more precise.

Observe the following properties of set-valued and relation-valued terms in our language, for every structure M and every valuation α:

- if R is a relation-valued expression, then

$$[\![R]\!]^{M,\alpha} \subseteq [\![*(\mathsf{left} \cup \mathsf{right})]\!]^{M,\alpha} \tag{2}$$

- if S is a set-valued term with free variables $x_1, \ldots, x_n$ on which α is defined, then

$$[\![S]\!]^{M,\alpha} \subseteq [\![\{x_1, \ldots, x_n\}.*(\mathsf{left} \cup \mathsf{right})]\!]^{M,\alpha} \tag{3}$$

These properties follow by induction on the size of the expressions R and S.

Note also that M_0 is a substructure of M_T, so by induction on size of R and S we have

$$\begin{array}{l} [\![R]\!]^{M_0} = [\![R]\!]^{M_T} \cap T_0^2 \\ [\![S]\!]^{M_0} = [\![S]\!]^{M_T} \cap T_0^2 \end{array} \tag{4}$$

In this inductive proof, the interesting case is showing (after applying the induction hypothesis)

$$([\![R_1]\!]^{M_T} \cap T_0^2) \circ ([\![R_2]\!]^{M_T} \cap T_0^2) = ([\![R_1]\!]^{M_T} \circ [\![R_2]\!]^{M_T}) \cap T_0^2$$

The $\subseteq$ inclusion holds by definition of the relation composition $\circ$, whereas the $\supseteq$ inclusion follows from (2) and the fact that M_0 is subterm closed.

We next show that the truth-value of a quantifier-free formula F is the same in M_0 and M_T when the free variables of F are interpreted in T_0. We show by induction on the structure of formula F the following claim:

For all $\alpha : \mathsf{Vars} \to T_0 \cup O$,

$$[\![F]\!]^{M_0,\alpha} = [\![F]\!]^{M_T,\alpha}. \tag{5}$$

Indeed, (5) holds for atomic formulas by condition (4) and the assumption that $\alpha(x) \in T_0 \cup O$. Moreover, this property is preserved by propositional combinations, so it holds for all boolean combinations.

Finally, given an EBU sentence φ, we prove the following relationship for all quantified subformulas F of φ: for all $\alpha : \mathsf{Vars} \to T_0 \cup O$, $[\![F]\!]^{M_0,\alpha}$ implies $[\![F]\!]^{M_T,\alpha} = \mathsf{true}$.

The base case corresponds to the previously proved case of quantifier-free formulas. We show that the condition is preserved under existential quantifiers, bounded universal quantifiers, and quantifiers over the finite set O. So suppose that $[\![F]\!]^{M_0,\alpha}$ implies $[\![F]\!]^{M_T,\alpha}$ for all $\alpha : \mathsf{Vars} \to T_0 \cup O$ and suppose that $\alpha : \mathsf{Vars} \to T_0 \cup O$ and $[\![F_1]\!]^{M_0,\alpha}$.

- Let $F_1 \equiv \exists v_t :: \mathsf{Tree}.\, F$. Then there exists $t \in T_0$ such that $[\![F]\!]^{M_0,\alpha'}$ where $\alpha = \alpha'[v_t := t]$. By induction hypothesis $[\![F]\!]^{M_T,\alpha'} = \mathsf{true}$, so $[\![F_1]\!]^{M_T,\alpha} = \mathsf{true}$.
- Let $F_1 \equiv \forall_S\, v_t :: \mathsf{Tree}.\, F$ for some set expression S. Then $[\![F]\!]^{M_0,\alpha[v_t:=t]} = \mathsf{true}$ for each $t \in [\![S]\!]^{M_0,\alpha}$. From (3), (4), $\alpha : \mathsf{Vars} \to T_0 \cup O$, and the fact that M_0 is subterm-closed, we conclude $[\![S]\!]^{M_T,\alpha} = [\![S]\!]^{M_0,\alpha} \subseteq T_0$. Consider arbitrary $t \in [\![S]\!]^{M_T,\alpha}$. Then $t \in [\![S]\!]^{M_0,\alpha}$, so $[\![F]\!]^{M_0,\alpha[v_t:=t]} = \mathsf{true}$. Because $t \in T_0$, by induction hypothesis $[\![F]\!]^{M_T,\alpha[v_t:=t]} = \mathsf{true}$. This proves $[\![F_1]\!]^{M_T,\alpha} = \mathsf{true}$.
- The cases $F_1 \equiv \exists v_o :: \mathsf{Object}.\, F$ and $F_1 \equiv \forall v_o :: \mathsf{Object}.\, F$ are straightforward because the quantifiers are monotonic and the structures M_T and M_0 have the same domain of uninterpreted objects O.

This completes the proof of one direction of our statement. Note that we have not relied on the fact that M_T is full term model. In fact, this direction still holds for M_0 and M_1 where M_0 is a substructure of M_1 and M_1 is a substructure of M_T: if the EBU sentence holds in M_0, then it also holds in the larger substructure M_1. We will use this generalization in the proof of the converse direction.

$\Longrightarrow$) : Let φ be EBU sentence. We prove by induction that for all subformulas F of φ the following holds: for each $\alpha : \mathsf{Vars} \to T_T$, if $[\![F]\!]^{M_T,\alpha} = \mathsf{true}$, then there exists a finite subterm-closed model M_0 and a valuation α_0 such that $\alpha_0(x_i) = \alpha(x_i)$ for each variable x_i free in F, such that $[\![F]\!]^{M_0,\alpha_0} = \mathsf{true}$. The proof of this claim is by induction on the number of quantifiers in F.

For the base case, assume that F is quantifier-free, and let $x_1, \ldots, x_n$ be the variables of F. Then let $T_0 = [\![\{x_1, \ldots, x_n\}.*(\mathsf{left} \cup \mathsf{right})]\!]^{M_T,\alpha}$ and let M_0 be the substructure of M_T induced by T_0. Let $\alpha_0(x_i) = \alpha(x_i)$ for $1 \le i \le n$ and let $\alpha_0(v) = \alpha(x_1)$ for $v \notin \{x_1, \ldots, x_n\}$. Then $\alpha_0 : \mathsf{Vars} \to T_0 \cup O$, so by (5) we have $[\![F]\!]^{M_0,\alpha_0}$; we have thus identified the desired M_0 and α_0.

For the inductive step, assume that claim holds for formula F, we prove that it holds for F_1 which is the result of quantifying F. Suppose that $[\![F_1]\!]^{M_T,\alpha}$ holds. We consider several cases.

- $F_1 \equiv \exists v_t :: \mathsf{Tree}.\, F$. Then there exists $t \in T_T$ such that $[\![F]\!]^{M_T,\alpha'}$ where $\alpha' = \alpha[v_t := t]$. By induction hypothesis, there exists α_0 that agrees with α' on free variables of F and a finite subterm-closed model M_0 such that $[\![F]\!]^{M_0,\alpha_0}$. This means that $[\![F_1]\!]^{M_0,\alpha_0}$, and α_0 certainly agrees with α on the variables free in F_1.
- $F_1 \equiv \forall_S\, v_t :: \mathsf{Tree}.\, F$. Let $\bar{S} = [\![S]\!]^{M_T,\alpha}$. Assume first $\bar{S} \neq \emptyset$. Then for each $t \in \bar{S}$, if $\alpha(t) = \alpha[v_t := t]$, then $[\![F]\!]^{M_T,\alpha(t)}$, so by induction hypothesis there exists a model $M_0(t) = (T_0(t), O, \iota_0(t))$ and a valuation $\alpha_0(t)$ such that $[\![F]\!]^{M_0(t),\alpha_0(t)}$, and $\alpha_0(t)$ agrees with $\alpha(t)$ on the free variables of F. Then let

  $$T_0' = \bar{S} \cup \bigcup_{t \in \bar{S}} T_0(t)$$

  Let T_0 be the subterm closure of T_0', given by $T_0 = T_0 \cup \{t \mid \exists t' \in T_0'.\ (t', t) \in [\![\mathsf{subterm}]\!]^{M_T,\alpha}\}$. The union T_0' is finite because $\bar{S}$ is finite, and each $T_0(t)$ is finite. Therefore, the subterm closure T_0 is finite, so there exists a finite subterm-closed structure $M_0 = (T_0, O, \iota_0)$. By the generalized version of the ($\Longrightarrow$) direction, because $M_0(t)$ is a substructure of M_0, we have that $[\![F]\!]^{M_0,\alpha_0(t)}$ for each $t \in \bar{S}$. Because we have $[\![S]\!]^{M_0,\alpha} = \bar{S}$ we conclude $[\![F_1]\!]^{M_0,\alpha_0(t_1)}$ where $t_1 \in \bar{S}$ is arbitrary.

  Next, consider the special case $\bar{S} = \emptyset$. Let

  $$T_0 = [\![\{x_1, \ldots, x_n\}.*(\mathsf{left} \cup \mathsf{right})]\!]^{M_T,\alpha},$$

  where $x_1, \ldots, x_n$ are the free variables of S, and consider the corresponding model $M_0 = (T_0, O, \iota_0)$. Then $[\![S]\!]^{M_0,\alpha} = \emptyset$, so $[\![F_1]\!]^{M_0,\alpha}$.

- $F_1 \equiv \exists v_o :: \mathsf{Object}. F$. This case is analogous to the case $F_1 \equiv \exists v_t :: \mathsf{Tree}.F$.

- $F_1 \equiv \forall v_o :: \mathsf{Object}. F$. This case is similar to the case $F_1 \equiv \forall_S v_t :: \mathsf{Tree}.F$, but slightly simpler. For each $o \in O$, if $\alpha(o) = \alpha[v_o := o]$, then $[\![F]\!]^{M_T,\alpha(o)}$, so by induction hypothesis there exists a model $M_0(o) = (T_0(o), O, \iota_0(o))$ and a valuation $\alpha_0(o)$ such that $[\![F]\!]^{M_0(o),\alpha_0(o)}$, and $\alpha_0(o)$ agrees with $\alpha(o)$ on free variables of F. Then let

$$T_0 = \bigcup_{o \in O} T_0(o)$$

  The union T_0 is finite because $\bar{S}$ is finite, and each $T_0(t)$ is finite. T_0 is also subterm closed because each $T_0(o)$ is subterm-closed. Therefore, there exists a finite subterm-closed structure $M_0 = (T_0, O, \iota_0)$. By the generalized version of the $(\Longrightarrow)$ direction, because $M_0(o)$ is a substructure of M_0, we have that $[\![F]\!]^{M_0,\alpha_0(o)}$ for each $o \in O$. We conclude $[\![F_1]\!]^{M_0,\alpha_0(o_1)}$ where $o_1 \in O$ is arbitrary.

■