

Deciding Functional Lists with Sublist Sets

Thomas Wies¹, Marco Muñoz², and Viktor Kuncak³

¹ New York University, New York, NY, USA
wies@cs.nyu.edu

² University of Freiburg, Germany
muniz@informatik.uni-freiburg.de

³ EPFL, Switzerland
viktor.kuncak@epfl.ch

Abstract. Motivated by the problem of deciding verification conditions for the verification of functional programs, we present new decision procedures for automated reasoning about functional lists. We first show how to decide in NP the satisfiability problem for logical constraints containing equality, constructor, selectors, as well as the transitive sublist relation. We then extend this class of constraints with operators to compute the set of all sublists, and the set of objects stored in a list. Finally, we support constraints on sizes of sets, which gives us the ability to compute list length as well as the number of distinct list elements. We show that the extended theory is reducible to the theory of sets with linear cardinality constraints, and therefore still in NP. This reduction enables us to combine our theory with other decidable theories that impose constraints on sets of objects, which further increases the potential of our decidability result in verification of functional and imperative software.

1 Introduction

Specifications using high-level data types, such as sets and algebraic data types have proven effective for describing the behavior of functional and imperative programs [14, 27]. Functional lists are particularly convenient and widespread in both programs and specifications. Efficient decision procedures for reasoning about lists can therefore greatly help automate software verification tasks.

Theories that allow only constructing and decomposing lists correspond to term algebras and have efficient decision procedures for quantifier-free fragments [1, 17]. However, these theories do not support list concatenation or sublists. Adding list concatenation makes the logic difficult because it subsumes the existential problem for word equations [6, 13, 19], which has been well-studied and is known to be difficult.

This motivates us to use as a starting point the logic of lists with a sublist (suffix) relation, which can express some (even if not all) of the properties expressible using list concatenation. We give an axiomatization of this theory where quantifiers can be instantiated in a complete and efficient way, following the methodology of local theory extensions [20]. Although local theory extensions have been applied to term algebras with certain recursive functions [21], they have not been applied to term algebras in the presence of the sublist operation. The general subterm relation in term algebras

was shown to be in NP using different techniques [23], without discussion of practical implementation procedures and without support for set operators. Several expressive logics of linked imperative data structures have been proposed [3, 11, 12, 24]. In these logics, variables range over graph nodes, as opposed to lists viewed as terms. In other words, the theories that we consider have an additional extensionality axiom, which ensures that no two list objects in the universe have identical tail and head. This axiom has non-trivial consequences on the set of satisfiable formulas and requires a new decision procedure. Our logic admits reasoning about set algebra constraints and cardinalities of sublist and content sets. In particular, the cardinality of the set of sublists of a list xs can be used to express the length xs . Decidable logics that allow reasoning about length of lists have been considered before [28]. However, our set algebra constraints are strictly more expressive and capture forms of quantification that are useful for the specification of complex properties.

Contributions. We summarize the contributions of our paper as follows:

- We give a set of local axioms for lists with sublist relation that admits efficient implementation in the spirit of [11, 24] and can leverage general implementation methods for local theory extensions [4, 5].
- We show how to extend this theory with an operator to compute the longest common suffix of two lists. We also give local axioms that give the decision procedure for the extended logic.
- We show how to further extend the theory by defining sets of elements that correspond to all sublists of a list, and then stating set algebra and size operations on such sets. Using a characterization of the models of this theory, we establish that the theory admits a reduction to the logic BAPA of sets with cardinality constraints [9, 10]. We obtain a decidable logic that supports reasoning about the contents of lists as well as about the number of elements in the list.

Impact on verification tools. We have found common functions in libraries of functional programming languages and interactive theorem provers that can be verified to meet a detailed specification using our logic. We discuss several examples in the paper. Moreover, the reduction to BAPA makes it possible to combine this logic with a number of other BAPA-reducible logics [18, 22, 25, 26]. Therefore, we believe that our logic will be a useful component of verification tools in the near future.

2 Examples

We describe our contributions through several examples. In the first two examples we show how we use our decision procedure to verify functional correctness of a function written in a functional programming notation similar to the Scala programming language [16]. In our third example we demonstrate the usefulness of our logic to increase the degree of automation in interactive theorem proving. Throughout this section we use the term *sublist* for a suffix of a list.

Example 1: dropping elements from a list. Our first example, listed in Figure 1, is the function `drop` of the `List` class in the Scala standard library (such functions also occur in

```

def drop[T](n: Int, xs: List[T]): List[T] = {
  if (n ≤ 0) xs
  else xs match {
    case nil ⇒ nil
    case cons(x, ys) ⇒ drop(n-1, ys)
  }
}
ensuring (zs ⇒ (n < 0 → zs = xs) ∧
           (n ≥ 0 ∧ length(xs) < n → zs = nil) ∧
           (n ≥ 0 ∧ length(xs) ≥ n → zs ≤ xs ∧ length(zs) = length(xs) - n))

```

Fig. 1: Function drop that drops the first n elements of a list xs

$$\begin{aligned}
& n > 0 \wedge xs \neq \text{nil} \wedge \text{cons}(x, ys) = xs \wedge zs \preceq ys \wedge \\
& (n - 1 \geq 0 \wedge \text{length}(ys) \geq n - 1 \rightarrow \text{length}(zs) = \text{length}(ys) - (n - 1)) \rightarrow \\
& \underbrace{zs \preceq xs}_{G_1} \wedge \underbrace{(n \geq 0 \wedge \text{length}(xs) \geq n \rightarrow \text{length}(zs) = \text{length}(xs) - n)}_{G_2}
\end{aligned}$$

Fig. 2: One of the verification conditions for the function drop

standard libraries for other functional languages, such as Haskell). The function takes as input an integer number n and a parametrized functional list xs . The function returns a functional list zs which is the sublist obtained from xs after dropping the initial n elements.

The **ensuring** statement specifies the postcondition of the function (a precondition is not required). The postcondition is expressed in our logic FLS^2 of functional lists with sublist sets shown in Figure 8. We consider the third conjunct of the postcondition in detail: it states that if the input n is a positive number and smaller than the length of xs then (1) the returned list zs is a sublist of the input list xs , denoted by $zs \preceq xs$, and (2) the length of zs is equal to the length of xs discounting the n dropped elements.

Deciding verification conditions. To verify the correctness of the drop function, we generate verification conditions and use our decision procedure to decide their validity. Figure 2 shows one of the generated verification conditions. This verification condition corresponds to the case when n is greater than 0 and xs is not the empty list. It is expressed in our logic of Functional Lists with Sublists and Sets (FLS^2).

We further split the verification condition into two subgoals G_1 and G_2 (see Figure 2). Each subgoal corresponds to one of the conjuncts in the postcondition of function drop. For proving subgoal G_1 we only need to reason about lists and sublists, but not about their lengths. We can prove this subgoal using the decision procedure for the simpler theory of Functional Lists with Sublists (FLS) that we present in Section 6. In the following, we concentrate on the more interesting subgoal G_2 .

Subgoal G_2 can be proved using the FLS^2 decision procedure presented in Section 7. The theory FLS^2 is a combination of the theory FLS and the theory of sets with linear cardinality constraints (BAPA) [9]. Our decision procedure follows the methodology of [25] that enables the combination of such set-sharing theories via reduction to BAPA. Figure 3 illustrates how this decision procedure proves subgoal G_2 . We first negate the subgoal and then eliminate the length function. For every list xs we encode its length $\text{length}(xs)$ using sublist sets as follows. We introduce a set variable Xs and

FLS fragment:

$$Xs = \sigma(xs) \wedge Ys = \sigma(ys) \wedge Zs = \sigma(zs) \wedge$$

$$xs \neq \text{nil} \wedge \text{cons}(x, ys) = xs \wedge zs \preceq ys$$

Projection onto shared sets Xs, Ys, Zs :

$$Zs \subseteq Ys \wedge Ys \subseteq Xs \wedge \text{card}(Xs) > 1 \wedge \text{card}(Xs) = \text{card}(Ys) + 1$$

BAPA fragment:

$$xs_length = \text{card}(Xs) - 1 \wedge ys_length = \text{card}(Ys) - 1 \wedge zs_length = \text{card}(Zs) - 1 \wedge$$

$$n > 0 \wedge (n - 1 \geq 0 \wedge ys_length \geq n - 1 \rightarrow zs_length = ys_length - (n - 1)) \wedge$$

$$n \geq 0 \wedge xs_length \geq n \wedge zs_length \neq xs_length - n$$

Projection onto shared sets Xs, Ys, Zs : $\text{card}(Xs) \neq \text{card}(Ys) + 1$

Fig. 3: Separated conjuncts for the negated subgoal G_2 of the VC in Figure 2 with the projections onto shared sets

define it as the set of all sublists of xs : $\{l. l \preceq xs\}$, which we denote by $\sigma(xs)$. We then introduce an integer variable xs_length that denotes the length of xs by defining $xs_length = \text{card}(Xs) - 1$, where $\text{card}(Xs)$ denotes the cardinality of set Xs . Note that we have to subtract 1, since nil is also a sublist of xs . We then purify the resulting formula and separate it into two conjuncts for the individual fragments. These two conjuncts are depicted in Figure 3. The two separated conjuncts share the set variables Xs, Ys , and Zs . After the separation the underlying decision procedure of each fragment computes a projection of the corresponding conjunct onto the shared set variables. These projections are the strongest BAPA consequences that are expressible over the shared sets in the individual fragments. After the projections have been computed, we check satisfiability of their conjunction using the BAPA decision procedure. In our example the conjunction of the two projections is unsatisfiable, which proves that G_2 is valid. In Section 7 we describe how to construct these projections onto set variables for the FLS² theory.

Example 2: greatest common suffix. Figure 4 shows our second example, a Scala function `gcs`, which takes as input two functional lists xs, ys and their corresponding lengths lxs, lys . This precondition is specified by the **require** statement. The function returns a pair (zs, lzs) such that zs is the greatest common suffix of the two input lists and lzs its length. This is captured by the postcondition. Our logic provides the operator $xs \sqcap ys$ that denotes the greatest common suffix of two lists xs and ys . Thus, we can directly express the desired property. Figure 5 depicts two constellations of lists xs, ys , and their greatest common suffix zs that may arise during the computation of `gcs`.

Figure 6 shows one of the verification conditions that are generated for the function `gcs`. This verification condition captures the case when the lists xs, ys are not empty, their lengths are equal, their head elements x, y are equal, and lzs is equal to $\text{length}(xs) - 1$. This situation is depicted on the right hand side of Figure 5, i.e., in this case the lists xs and ys are identical. The verification condition can again be split into two subgoals. We focus on subgoal G_1 . Figure 7 shows the separated conjuncts for this subgoal and their projections onto the shared set variables Xs, Ys, Zs , and $Z1s$. Using the BAPA decision procedure, we can again prove that the conjunction of the two projections is unsatisfiable.

```

def gcs[T](xs: List[T], lxs: Int, ys: List[T], lys: Int): (List[T], Int)
require (length(xs)=lxs ∧ length(ys)=lys) =
  (xs,ys) match {
  case (nil, _) ⇒ (nil, 0)
  case (_, nil) ⇒ (nil, 0)
  case (cons(x, x1s), cons(y, y1s)) ⇒
    if (lxs > lys) gcs(x1s, lxs-1, ys, lys)
    else if (lxs < lys) gcs(xs, lxs, y1s, lys-1)
    else {
      val (z1s, lz1s) = gcs(x1s, lxs-1, y1s, lys-1)
      if (x = y ∧ lz1s = (lxs - 1)) (cons(x, z1s), lz1s+1) else (z1s, lz1s)
    }
  }
} ensuring ((zs, lzs) ⇒ length(zs) = lzs ∧ zs = xs □ ys)

```

Fig. 4: Function gcs that computes the greatest common suffix of two lists

Fig. 5: Lists x_s , y_s and their greatest common suffix z_s

Example 3: interactive theorem proving. Given a complete specification of functions such as `drop` and `gcd` in our logic, we can use our decision procedure to automatically prove more complex properties about such functions. For instance, the function `drop` is not just defined in the Scala standard library, but also in the theory `List` of the Isabelle/HOL interactive theorem prover [15]. Consider the following property of function `drop`:

$$m \geq n \rightarrow \tau(\text{drop}(n, xs)) \subseteq \tau(\text{drop}(m, xs))$$

where the expression $\tau(xs)$ denotes the *content set* of a list xs , i.e., $\tau(xs) = \{\text{head}(l).l \preceq xs\}$. This property corresponds to Lemma `set.drop.subset.set.drop`, which is proved in the Isabelle theory `List`. Using the postcondition of function `drop` to eliminate all occurrences of this function in Lemma `set.drop.subset.set.drop` yields the following formula in our logic:

$$\begin{aligned}
& (n < 0 \rightarrow zs_n = xs) \wedge \\
& (n \geq 0 \wedge \text{length}(xs) < n \rightarrow zs_n = \text{nil}) \wedge \\
& (n \geq 0 \wedge \text{length}(xs) \geq n \rightarrow zs_n \preceq xs \wedge \text{length}(zs_n) = \text{length}(xs) - n) \wedge \\
& (m < 0 \rightarrow zs_m = xs) \wedge \\
& (m \geq 0 \wedge \text{length}(xs) < m \rightarrow zs_m = \text{nil}) \wedge \\
& (m \geq 0 \wedge \text{length}(xs) \geq m \rightarrow zs_m \preceq xs \wedge \text{length}(zs_m) = \text{length}(xs) - m) \wedge \\
& m \geq n \rightarrow \tau(zs_n) \subseteq \tau(zs_m)
\end{aligned}$$

The proof of lemma `set.drop.subset.set.drop` that is given in the Isabelle theory is not fully automated and involves the statement of an intermediate lemma. Instead, using our decision procedure we can prove the lemma directly and fully automatically. Our logic is, thus, useful to increase the degree of automation in interactive theorem proving.

$$\begin{aligned}
& \text{length}(xs) = lxs \wedge \text{length}(ys) = lys \wedge xs \neq \text{nil} \wedge ys \neq \text{nil} \wedge lxs = lys \wedge x = y \wedge \\
& \text{cons}(x, x1s) = xs \wedge \text{cons}(y, y1s) = ys \wedge lz1s = lxs - 1 \wedge \\
& \text{length}(z1s) = lz1s \wedge z1s = xs1 \sqcap y1s \wedge zs = \text{cons}(x, z1s) \wedge lzs = lz1s + 1 \rightarrow \\
& \underbrace{\text{length}(zs) = lzs}_{G_1} \wedge \underbrace{zs = xs \sqcap ys}_{G_2}
\end{aligned}$$

Fig. 6: One of the verification conditions for the function gcs

FLS fragment:

$$\begin{aligned}
& Xs = \sigma(xs) \wedge Ys = \sigma(ys) \wedge Zs = \sigma(zs) \wedge Z1s = \sigma(z1s) \wedge \\
& \text{cons}(x, x1s) = xs \wedge \text{cons}(y, y1s) = ys \wedge xs \neq \text{nil} \wedge x = y \wedge ys \neq \text{nil} \wedge \\
& z1s = x1s \sqcap y1s \wedge zs = \text{cons}(x, z1s)
\end{aligned}$$

Projection onto shared sets $Xs, Ys, Zs, Z1s$:

$$\begin{aligned}
& \text{card}(Xs) > 1 \wedge \text{card}(Ys) > 1 \wedge \text{card}(Zs) > 1 \wedge \\
& Z1s \subseteq Zs \wedge \text{card}(Zs) = \text{card}(Z1s) + 1 \wedge \\
& ((\text{card}(Z1s) = \text{card}(Xs) - 1 \vee \text{card}(Z1s) = \text{card}(Ys) - 1) \rightarrow Zs = Xs = Ys)
\end{aligned}$$

BAPA fragment:

$$\begin{aligned}
& xs_length = \text{card}(Xs) - 1 \wedge ys_length = \text{card}(Ys) - 1 \wedge \\
& zs_length = \text{card}(Zs) - 1 \wedge z1s_length = \text{card}(Z1s) - 1 \wedge \\
& xs_length = lxs \wedge ys_length = lys \wedge z1s_length = lz1s \wedge \\
& lxs = lys \wedge lz1s = lxs - 1 \wedge lzs = lz1s + 1 \wedge zs_length \neq lzs
\end{aligned}$$

Projection onto shared sets $Xs, Ys, Zs, Z1s$:

$$\text{card}(Z1s) = \text{card}(Xs) - 1 \wedge \text{card}(Z1s) = \text{card}(Ys) - 1 \wedge \text{card}(Zs) \neq \text{card}(Z1s) + 1$$

Fig. 7: Separated conjuncts for the negated subgoal G_1 of the VC in Figure 6 with projections onto the shared sets

3 Logic FLS² of Functional Lists with Sublists Sets

The grammar of our logic of functional lists with sublist sets is shown in Figure 8. It supports reasoning about lists built from list constructors and selectors, sublists, the length of lists, and cardinality and set algebraic constraints over the sets of sublists of lists $\sigma(l)$ as well their content sets $\tau(l)$.

The remainder of the paper is structured as follows. In Section 5 we first present the fragment FLS of the logic FLS² that only subsumes formulas over lists and sublists, but not sets, cardinalities, or length constraints. We then formally define the semantics of the logic FLS and give a decision procedure for its satisfiability problem in Section 6. Finally, in Section 7 we show how to use this decision procedure for a BAPA reduction that decides the full logic FLS².

4 Preliminaries

In the following, we define the syntax and semantics of formulas. We further recall the notions of partial structures and local theory extensions from [20].

Sorted logic. We present our problem in sorted logic with equality. A *signature* Σ is a tuple (S, Ω) , where S is a countable set of sorts and Ω is a countable set of function

$$\begin{aligned}
F &::= A_L \mid A_S \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid \neg F \\
A_L &::= T_L \preceq T_L \mid T_L = T_L \mid T_H = T_H \\
T_L &::= v_L \mid \text{nil} \mid \text{cons}(T_H, T_L) \mid \text{tail}(T_L) \mid T_L \sqcap T_L \\
T_H &::= v_H \mid \text{head}(T_L) \\
A_S &::= B_L = B_L \mid B_L \subseteq B_L \mid T_I = T_I \mid T_I < T_I \\
B_L &::= s_L \mid \emptyset \mid \{T_L\} \mid \sigma(T_L) \mid B_L \cup B_L \mid B_L \setminus B_L \\
B_H &::= s_H \mid \emptyset \mid \{T_H\} \mid \tau(T_L) \mid \text{head}[B_L] \mid B_H \cup B_H \mid B_H \setminus B_H \\
T_I &::= v_I \mid K \mid T_I + T_I \mid K \cdot T_I \mid \text{card}(B_L) \mid \text{card}(B_H) \mid \text{length}(T_L) \\
K &::= \dots -2 \mid -1 \mid 0 \mid 1 \mid 2 \dots
\end{aligned}$$

Fig. 8: Logic FLS² of lists, sublist, sublist sets, list contents, and size constraints

symbols f with associated arity $n \geq 0$ and associated sort $s_1 \times \dots \times s_n \rightarrow s_0$ with $s_i \in S$ for all $i \leq n$. Function symbols of arity 0 are called *constant symbols*. We assume that all signatures contain the sort `bool` and for every other sort $s \in S$ a dedicated equality symbol $=_s \in \Omega$ of sort $s \times s \rightarrow \text{bool}$. Note that we generally treat predicate symbols of sort s_1, \dots, s_n as function symbols of sort $s_1 \times \dots \times s_n \rightarrow \text{bool}$. Terms are built as usual from the function symbols in Ω and (sorted) variables taken from a countably infinite set X that is disjoint from Ω . A term t is said to be *ground*, if no variable appears in t . We denote by $\text{Terms}(\Sigma)$ the set of all ground Σ -terms.

A Σ -atom A is a Σ -term of sort `bool`. We use infix notation for atoms built from the equality symbol. A Σ -formula F is defined via structural recursion as either one of A , $\neg F_1$, $F_1 \wedge F_2$, or $\forall x : s.F_1$, where A is a Σ -atom, F_1 and F_2 are Σ -formulas, and $x \in X$ is a variable of sort $s \in S$. We typically drop the sort annotation (both for quantified variables and the equality symbols) if this does not cause any ambiguity. We use syntactic sugar for Boolean constants (`true`, `false`), disjunctions ($F_1 \vee F_2$), implications ($F_1 \rightarrow F_2$), and existential quantification ($\exists x.F_1$). We define literals and clauses as usual. A clause C is called *flat* if no term that occurs in C below a predicate symbol or the symbol $=$ contains nested function symbols. A clause C is called *linear* if (i) whenever a variable occurs in two non-variable terms in C that do not start with a predicate or the equality symbol, the two terms are identical, and if (ii) no such term contains two occurrences of the same variable.

Total and partial structures. Given a signature $\Sigma = (S, \Omega)$, a *partial Σ -structure* α is a function that maps each sort $s \in S$ to a non-empty set $\alpha(s)$ and each function symbol $f \in \Omega$ of sort $s_1 \times \dots \times s_n \rightarrow s_0$ to a partial function $\alpha(f) : \alpha(s_1) \times \dots \times \alpha(s_n) \rightarrow \alpha(s_0)$. If α is understood, we write just t instead of $\alpha(t)$ whenever this is not ambiguous. We assume that all partial structures interpret the sort `bool` by the two-element set of Booleans $\{0, 1\}$. We further assume that all structures α interpret the symbol $=_s$ by the equality relation on $\alpha(s)$. A partial structure α is called *total structure* or simply *structure* if it interprets all function symbols by total functions. For a Σ -structure α where Σ extends a signature Σ_0 with additional sorts and function symbols, we write $\alpha|_{\Sigma_0}$ for the Σ_0 -structure obtained by restricting α to Σ_0 .

Given a total structure α and a *variable assignment* $\beta : X \rightarrow \alpha(S)$, the evaluation $\llbracket t \rrbracket_{\alpha, \beta}$ of a term t in α, β is defined as usual. For a ground term t we typically write just $\llbracket t \rrbracket_{\alpha}$. A quantified variable of sort s ranges over all elements of $\alpha(s)$. From the interpre-

tation of terms the notions of satisfiability, validity, and entailment of atoms, formulas, clauses, and sets of clauses in total structures are derived as usual. In particular, we use the standard interpretations for propositional connectives of classical logic. We write $\alpha, \beta \models F$ if α satisfies F under β where F is a formula, a clause, or a set of clauses. Similarly, we write $\alpha \models F$ if F is valid in α . In this case we also call α a *model* of F . The interpretation $\llbracket t \rrbracket_{\alpha, \beta}$ of a term t in a partial structure α is as for total structures, except that if $t = f(t_1, \dots, t_n)$ for $f \in \Omega$ then $\llbracket t \rrbracket_{\alpha, \beta}$ is undefined if either $\llbracket t_i \rrbracket_{\alpha, \beta}$ is undefined for some i , or $(\llbracket t_1 \rrbracket_{\alpha, \beta}, \dots, \llbracket t_n \rrbracket_{\alpha, \beta})$ is not in the domain of $\alpha(f)$. We say that a partial structure α *weakly satisfies* a literal L under β , written $\alpha, \beta \models_w L$, if (i) L is an atom A and either $\llbracket A \rrbracket_{\alpha, \beta} = 1$ or $\llbracket A \rrbracket_{\alpha, \beta}$ is undefined, or (ii) L is a negated atom $\neg A$ and either $\llbracket A \rrbracket_{\alpha, \beta} = 0$ or $\llbracket A \rrbracket_{\alpha, \beta}$ is undefined. The notion of weak satisfiability is extended to clauses and sets of clauses as for total structures. A clause C (respectively, a set of clauses) is *weakly valid* in a partial structure α if α weakly satisfies α for all variable assignments β . We then call α a *weak partial model* of C .

Theories and local theory extensions. A *theory* \mathcal{T} for a signature Σ is simply a set of Σ -formulas. We consider theories $\mathcal{T}(\mathcal{M})$ defined as a set of Σ -formulas that are valid in a given set of models \mathcal{M} , as well as theories $\mathcal{T}(\mathcal{K})$ defined as a set of Σ -formulas that are consequences of a given set of formulas \mathcal{K} . In the latter case, we call \mathcal{K} the *axioms* of the theory $\mathcal{T}(\mathcal{K})$ and we often identify \mathcal{K} and $\mathcal{T}(\mathcal{K})$. In particular, we call $\mathcal{K} = \emptyset$ the *empty theory*.

In what follows, we consider theories that are defined by a set of axioms. Let $\Sigma_0 = (S, \Omega_0)$ be a signature and assume that signature $\Sigma_1 = (S, \Omega_0 \cup \Omega_1)$ extends Σ_0 by new function symbols Ω_1 . We call the function symbols in Ω_1 *extension symbols* and terms starting with extension symbols *extension terms*. Now, a theory \mathcal{T}_1 over Σ_1 is an *extension* of a theory \mathcal{T}_0 over Σ_0 , if \mathcal{T}_1 is obtained from \mathcal{T}_0 by adding a set of (universally quantified) clauses \mathcal{K} . In the following, when we refer to a set of ground clauses G , we assume they are over the signature $\Sigma_1^c = (S, \Omega_0 \cup \Omega_1 \cup \Omega_c)$ where Ω_c is a set of new constant symbols. Let \mathcal{K} be a set of (universally quantified) clauses. We denote by $\text{st}(\mathcal{K}, G)$ the set of all ground subterms that appear in \mathcal{K} or G and by $\mathcal{K}[G]$ the set of all instantiations of clauses in \mathcal{K} where variables appearing below extension terms have been instantiated by the terms in $\text{st}(\mathcal{K}, G)$. Then an extension $\mathcal{T}_1 = \mathcal{T}_0 \cup \mathcal{K}$ is a *local extension* if it satisfies condition (Loc):

- (Loc) For every finite set of ground clauses G , $G \cup \mathcal{T}_1 \models \text{false}$ iff there is no partial Σ_1^c -structure α such that $\alpha|_{\Sigma_0}$ is a total model of \mathcal{T}_0 , all terms in $\text{st}(\mathcal{K}, G)$ are defined in α , and α weakly satisfies $\mathcal{K}[G] \cup G$.

5 Logic FLS of Functional Lists with Sublists

We now define the logic of functional lists with sublists (FLS) and its accompanying theory. The logic FLS is given by all quantifier-free formulas over the signature $\Sigma_{\text{FLS}} = (S_{\text{FLS}}, \Omega_{\text{FLS}})$. The signature Σ_{FLS} consists of sorts $S_{\text{FLS}} = \{\text{bool}, \text{list}, \text{data}\}$ and function symbols $\Omega_{\text{FLS}} = \{\text{nil}, \text{cons}, \text{head}, \text{tail}, \sqcap, \preceq\}$. The sorts of the function symbols in Ω_{FLS} are shown in Figure 9. We use infix notation for the symbols \sqcap and \preceq .

The theory of functional lists with sublist relationship \mathcal{T}_{FLS} is the set of all formulas in FLS that are true in the canonical model of lists. We denote this canonical

$\text{nil} : \text{list}$	$\text{tail} : \text{list} \rightarrow \text{list}$	$\preceq : \text{list} \times \text{list} \rightarrow \text{bool}$
$\text{cons} : \text{data} \times \text{list} \rightarrow \text{list}$	$\text{head} : \text{list} \rightarrow \text{data}$	$\sqcap : \text{list} \times \text{list} \rightarrow \text{list}$

Fig. 9: Sorts of function symbols in the signature Σ_{FLS}

$$\begin{aligned}
\alpha_{\text{FLS}}(\text{list}) &\stackrel{\text{def}}{=} \mathbf{L} \stackrel{\text{def}}{=} \{ t \in \text{Terms}(\Sigma_{\text{L}}) \mid t : \text{list} \} \\
\alpha_{\text{FLS}}(\text{data}) &\stackrel{\text{def}}{=} \mathbf{D} \stackrel{\text{def}}{=} \{ t \in \text{Terms}(\Sigma_{\text{L}}) \mid t : \text{data} \} \\
\alpha_{\text{FLS}}(\text{cons}) &\stackrel{\text{def}}{=} \text{cons}_{\text{L}} \stackrel{\text{def}}{=} \lambda(d, l). \text{cons}(d, l) \\
\alpha_{\text{FLS}}(\text{nil}) &\stackrel{\text{def}}{=} \text{nil} \\
\alpha_{\text{FLS}}(\text{tail}) &\stackrel{\text{def}}{=} \text{tail}_{\text{L}} \stackrel{\text{def}}{=} \lambda l. \text{if } l = \text{nil} \text{ then nil else } l' \text{ where } l = \text{cons}(d, l') \\
\alpha_{\text{FLS}}(\text{head}) &\stackrel{\text{def}}{=} \text{head}_{\text{L}} \stackrel{\text{def}}{=} \lambda l. \text{if } l = \text{nil} \text{ then } d_1 \text{ else } d \text{ where } l = \text{cons}(d, l') \\
\alpha_{\text{FLS}}(\preceq) &\stackrel{\text{def}}{=} \lambda(l_1, l_2). l_1 \preceq_{\text{L}} l_2 \\
\alpha_{\text{FLS}}(\sqcap) &\stackrel{\text{def}}{=} \lambda(l_1, l_2). l_1 \sqcap_{\text{L}} l_2
\end{aligned}$$

Fig. 10: The canonical model α_{FLS} of functional lists with sublists

model by α_{FLS} . The structure α_{FLS} is the term algebra generated by the signature $\Sigma_{\text{L}} = (S_{\text{FLS}}, \{\text{cons}, \text{nil}, d_1, d_2, \dots\})$, where d_1, d_2, \dots are infinitely many constant symbols of sort data. The complete definition of α_{FLS} is given in Figure 10. The canonical model interprets the sort list as the set of all Σ_{L} -terms of sort list. We denote this set by \mathbf{L} . Likewise, the sort data is interpreted as the set of all Σ_{L} -terms of sort data. We denote this set by \mathbf{D} . The function symbols cons and nil are interpreted as the corresponding term constructors. The function symbols head and tail are interpreted as the appropriate selectors head_{L} and tail_{L} . The predicate symbol \preceq is interpreted as the sublist relation $\preceq_{\text{L}} \subseteq \mathbf{L} \times \mathbf{L}$ on lists. The sublist relation is defined as the inverse of the reflexive transitive closure of the tail selector function:

$$l_1 \preceq_{\text{L}} l_2 \stackrel{\text{def}}{\iff} (l_2, l_1) \in \{ (l, \text{tail}_{\text{L}}(l)) \mid l \in \mathbf{L} \}^*$$

The relation \preceq_{L} is a partial order on lists. In fact, one can show more.

Proposition 1. *The relation \preceq_{L} induces a meet-semilattice on the set \mathbf{L} .*

We denote by \sqcap_{L} the meet operator of the semilattice induced by \preceq_{L} . Given two lists l_1 and l_2 , the list $l_1 \sqcap_{\text{L}} l_2$ denotes the greatest common suffix of l_1 and l_2 . The structure α_{FLS} interprets the function symbol \sqcap as the operator \sqcap_{L} .

We further define the theory of all finite substructures of α_{FLS} . Let Σ_{FLSf} be the signature Σ_{FLS} without the function symbol cons and let α_{FLSf} be the structure α_{FLS} restricted to the signature Σ_{FLSf} . We call a finite subset L of \mathbf{L} sublist closed if for all $l \in L, l' \in \mathbf{L}, l' \preceq_{\text{L}} l$ implies $l' \in L$. For a finite sublist closed subset L of \mathbf{L} , the structure α_L is the finite total substructure of α_{L} induced by the restricted support sets $\alpha_L(\text{list}) \stackrel{\text{def}}{=} L$ and $\alpha_L(\text{data}) \stackrel{\text{def}}{=} \{ \text{head}_{\text{L}}(l) \mid l \in L \}$. We denote by $\mathcal{M}_{\text{FLSf}}$ the set of all such finite total substructures α_L of α_{FLSf} . The theory $\mathcal{T}_{\text{FLSf}}$ is the set of all FLS formulas that are true in all structures $\mathcal{M}_{\text{FLSf}}$.

6 Decision Procedure for FLS

In the following, we show that the theory \mathcal{T}_{FLS} is decidable. For this purpose we reduce the decision problem for \mathcal{T}_{FLS} to the decision problem of the theory $\mathcal{T}_{\text{FLSf}}$. We then give a finite first-order axiomatization of the theory $\mathcal{T}_{\text{FLSf}}$ and show that it is a local theory extension of the empty theory. In total, this implies that deciding satisfiability of a ground formula F with respect to the theory \mathcal{T}_{FLS} can be reduced to deciding satisfiability of F conjoined with finitely many ground instances of the first-order axioms of $\mathcal{T}_{\text{FLSf}}$.

6.1 Reducing FLS to FLSf

We first note that satisfiability of an FLS formula F in the canonical model can be reduced to checking satisfiability in the finite substructures, if the function symbol `cons` does not occur in F .

Proposition 2. *Let F be a quantifier-free Σ_{FLSf} -formula. Then F is satisfiable in α_{FLS} if and only if F is satisfiable in some structure $\alpha \in \mathcal{M}_{\text{FLSf}}$.*

The proof of Proposition 2 is similar to the proof of [7, Theorem 2].

We can now exploit the fact that, in the term algebra α_{FLS} , the constructor `consL` is uniquely determined by the functions `headL` and `tailL`. Let F be an FLS formula. Then we can eliminate an occurrence $F(\text{cons}(t_d, t_l))$ of function symbol `cons` in a term of F by rewriting $F(\text{cons}(t_d, t_l))$ into:

$$x \neq \text{nil} \wedge \text{head}(x) = t_d \wedge \text{tail}(x) = t_l \wedge F(x)$$

where x is a fresh variable of sort `list` that does not appear elsewhere in F . Let $\text{elimcons}(F)$ be the formula that results from rewriting recursively all appearances of function symbol `cons` in F . Clearly, in the canonical model α_{FLS} , the formulas F and $\text{elimcons}(F)$ are equisatisfiable. Thus, with Proposition 2 we can conclude.

Lemma 3. *Let F be an FLS formula. Then F is satisfiable in α_{FLS} if and only if $\text{elimcons}(F)$ is satisfiable in some structure $\alpha \in \mathcal{M}_{\text{FLSf}}$.*

6.2 Axiomatizing FLSf

We next show that there exists a first-order axiomatization $\mathcal{K}_{\text{FLSf}}$ of the theory $\mathcal{T}_{\text{FLSf}}$. The axioms $\mathcal{K}_{\text{FLSf}}$ are given in Figure 11. The free variables appearing in the formulas are implicitly universally quantified. We now explain each of these axioms and argue their soundness, i.e., that each axiom is true in the canonical model α_{FLS} .

The axiom `Pure` is a logical consequence of the following formula, which is true in the canonical model α_{FLS} : $\forall x. \text{cons}(\text{head}(x), \text{tail}(x)) = x \vee \text{nil} = x$. Hence, `Pure` is true in all finite total substructures of α_{FLS} . The axiom `NoCycle1` expresses that all lists have `nil` as a sublist. This axiom is true because all lists are constructed from `nil`. The axiom `NoCycle2` excludes self-cycles of the form $\text{tail}(x) = x$, with the exception of `nil`. The axioms `Refl`, `Trans`, and `AntiSym` express that \preceq is a partial order. These axioms

Pure: $\text{head}(x) = \text{head}(y) \wedge \text{tail}(x) = \text{tail}(y) \rightarrow x = y \vee x = \text{nil} \vee y = \text{nil}$	
NoCycle1: $\text{nil} \preceq x$	UnfoldL: $\text{tail}(x) \preceq x$
NoCycle2: $\text{tail}(x) = x \rightarrow x = \text{nil}$	UnfoldR: $x \preceq y \rightarrow x = y \vee x \preceq \text{tail}(y)$
Refl: $x \preceq x$	GCS1: $x \sqcap y \preceq x$
Trans: $x \preceq y \wedge y \preceq z \rightarrow x \preceq z$	GCS2: $x \sqcap y \preceq y$
AntiSym: $x \preceq y \wedge y \preceq x \rightarrow x = y$	GCS3: $z \preceq x \wedge z \preceq y \rightarrow z \preceq x \sqcap y$
Total: $y \preceq x \wedge z \preceq x \rightarrow y \preceq z \vee z \preceq y$	

Fig. 11: First-order axiomatization $\mathcal{K}_{\text{FLSf}}$ of the theory $\mathcal{T}_{\text{FLSf}}$

follow from Proposition 1. The axiom Total expresses the fact that, for a fixed list x , all sublists of x are totally ordered by the sublist relation. This axiom follows from the definition of \preceq_L as the reflexive transitive closure of a functional relation. The axioms UnfoldL and UnfoldR express that, by applying tail to either side of the sublist relation, we stay in the sublist relation. Finally, the axioms GCS1, GCS2, and GCS3 express that \sqcap is the greatest lower bound operator of the partial order \preceq . These axioms follow from the definition of \sqcap_L and Proposition 1.

Lemma 4. *The axioms $\mathcal{K}_{\text{FLSf}}$ are sound, i.e., for all $\alpha \in \mathcal{M}_{\text{FLSf}}$, $\alpha \models \mathcal{K}_{\text{FLSf}}$.*

As a prerequisite for proving completeness of the axioms, we next show that the finite models of the axioms $\mathcal{K}_{\text{FLSf}}$ are structurally equivalent to the finite substructures of the canonical model of functional lists.

Proposition 5. *Every finite model of $\mathcal{K}_{\text{FLSf}}$ is isomorphic to some structure in $\mathcal{M}_{\text{FLSf}}$.*

The proof is in Appendix B.1.

6.3 Locality of FLSf

We will now prove that the theory $\mathcal{K}_{\text{FLSf}}$ can be understood as a local theory extension and, at the same time, prove that $\mathcal{K}_{\text{FLSf}}$ is a complete axiomatization of the theory $\mathcal{T}_{\text{FLSf}}$.

In what follows, the signature Σ_{FLSf} is the signature of the theory extension $\mathcal{K}_{\text{FLSf}}$. We also have to determine the signature Σ_0 of the base theory \mathcal{T}_0 by fixing the extension symbols. We treat the function symbols $\Omega_e \stackrel{\text{def}}{=} \{\text{head}, \text{tail}, \sqcap\}$ as extension symbols, but the sublist relation \preceq as a symbol in the signature of the base theory, i.e. $\Sigma_0 \stackrel{\text{def}}{=} (\mathcal{S}_{\text{FLS}}, \{\text{nil}, \preceq\})$. The base theory itself is given by the axioms that define the sublist relation, but that do not contain any of the extension symbols, i.e.

$$\mathcal{T}_0 \stackrel{\text{def}}{=} \{\text{NoCycle1}, \text{Refl}, \text{Trans}, \text{AntiSym}, \text{Total}\} .$$

We further denote by $\mathcal{K}_e \stackrel{\text{def}}{=} \mathcal{K}_{\text{FLSf}} \setminus \mathcal{T}_0$ the extension axioms.

We now show that $\mathcal{K}_{\text{FLSf}} = \mathcal{T}_0 \cup \mathcal{K}_e$ is a local theory extension. As in the definition of local theory extensions in Section 4, for a set of ground clauses G , we denote by $\mathcal{K}_e[G]$ all instances of axioms \mathcal{K}_e where the variables occurring below extension symbols Ω_e are instantiated by all ground terms $\text{st}(\mathcal{K}_e, G)$ that appear in \mathcal{K}_e and G . Furthermore, we denote by Σ_{FLSf}^c the signature Σ_{FLSf} extended with finitely many new constant symbols Ω_c .

Lemma 6. *For every finite set of Σ_{FLSf}^c ground clauses G , if α is a partial Σ_{FLSf}^c -structure such that $\alpha|_{\Sigma_0}$ is a total model of \mathcal{T}_0 , all terms in $\text{st}(\mathcal{K}_e, G)$ are defined in α , and α weakly satisfies $\mathcal{K}_e[G] \cup G$ then there exists a finite total Σ_{FLSf}^c -structure that satisfies $\mathcal{K}_{\text{FLSf}} \cup G$.*

We sketch the proof of Lemma 6. Let α be a partial Σ_{FLSf}^c -structure as required in the lemma. We can obtain a finite partial substructure α' from α by restricting the interpretations of sorts `data` and `list` to the elements that are used in the interpretations of the ground terms $\text{st}(\mathcal{K}_e, G)$. Then α' is still a total model of \mathcal{T}_0 and still weakly satisfies $\mathcal{K}_e[G] \cup G$, since all axioms in $\mathcal{K}_{\text{FLSf}}$ are universal. We can then complete α' to a finite total model of $\mathcal{K}_{\text{FLSf}} \cup G$ as follows. First, for every $u \in \alpha'(\text{list})$ where $\alpha'(\text{head})$ is not defined, we can extend $\alpha'(\text{data})$ by a fresh element d_u and define $\alpha'(\text{head})(u) = d_u$. Now, let $u \in \alpha'(\text{list})$ such that $\alpha'(\text{tail})$ is not defined on u . If $u = \alpha'(\text{nil})$, we define $\alpha'(\text{tail})(u) = u$. Otherwise, from the fact that α' satisfies axioms `NoCycle1`, `AntiSym`, and `Total` we can conclude that there exists a maximal element $v \in \alpha'(\text{list}) \setminus \{u\}$ such that $(v, u) \in \alpha'(\preceq)$. However, we cannot simply define $\alpha'(\text{tail})(u) = v$. The resulting structure would potentially violate axiom `Pure`. Instead, we extend $\alpha'(\text{list})$ with a fresh element w and $\alpha'(\text{data})$ with a fresh element d_w , and define: $\alpha'(\text{head})(w) = d_w$, $\alpha'(\text{tail})(w) = v$, and $\alpha'(\text{tail})(u) = w$. We further extend the definition of $\alpha'(\preceq)$ for the newly added element w , as expected. The completion of $\alpha'(\sqcap)$ to a total function is then straightforward.

From Lemma 6 we can now immediately conclude that the theory $\mathcal{K}_{\text{FLSf}}$ satisfies condition `(Loc)`.

Theorem 7. *$\mathcal{K}_{\text{FLSf}}$ is a local theory extension of the theory \mathcal{T}_0 .*

Similarly, from Proposition 5 and Lemma 6, we can conclude that the axioms $\mathcal{K}_{\text{FLSf}}$ are complete.

Theorem 8. *$\mathcal{K}_{\text{FLSf}}$ is an axiomatization of the theory $\mathcal{T}_{\text{FLSf}}$, i.e., $\mathcal{T}(\mathcal{K}_{\text{FLSf}}) = \mathcal{T}_{\text{FLSf}}$.*

6.4 Deciding FLS

We now describe the decision procedure for deciding satisfiability of FLS formulas. Given an FLS formula F as input, the decision procedure proceeds as follows: (1) compute $\hat{F} = \text{elimcons}(\neg F)$, replace all variables in \hat{F} with fresh constant symbols, and transform the resulting formula into a set of ground clauses G ; and (2) use Theorem 7 and the reduction scheme for reasoning in local theory extensions [20], to reduce satisfiability of the set of clauses $\mathcal{K}_{\text{FLSf}} \cup G$ to satisfiability of an equisatisfiable formula in the Bernays-Schönfinkel-Ramsey class, which is decidable. The reduction scheme computes the set of clauses $\mathcal{T}_0 \cup \mathcal{K}_e[G] \cup G$ and then eliminates all occurrences of extension functions Ω_e in literals of clauses in this set. The resulting set of clauses contains only universally quantified variables, constants, relation symbols, and equality, i.e., it belongs to the Bernays-Schönfinkel-Ramsey class. Soundness and completeness of the decision procedure follows from Lemma 3, Theorems 7 and 8, and [20, Lemma 4].

Complexity. For formulas in the Bernays-Schönfinkel-Ramsey class that have a bounded number of universal quantifiers, the satisfiability problem is known to be NP-complete [2, page 258]. The only quantified variables appearing in the set of clauses obtained after the reduction step of the decision procedure are those that come from the axioms in $\mathcal{K}_{\text{FLSf}}$, more precisely, the axioms in \mathcal{T}_0 and the (partial) instantiations of the axioms in \mathcal{K}_e . In fact, we can write the clauses for these axioms in such a way that they use exactly 3 quantified variables. Finally, from the parametric complexity considerations in [20] follows that the size of the set of clauses obtained in the final step of our decision procedure is polynomial in the size of the input formula. It follows that the satisfiability problem for FLS is decidable in NP. NP-hardness follows immediately from the fact that FLS can express arbitrary propositional formulas.

Theorem 9. *The decision problem for the theory \mathcal{T}_{FLS} is NP-complete.*

7 Extension with Sets of Sublists and Content Sets

We next show decidability of the logic that extends FLS with constraints on sets of sublists and the contents of lists. We do this by reducing the extended logic to constraints on sets. For this we need a normal form of formulas in our logic. To obtain this normal form, we start from Σ_{FLSf}^2 partial models, but refine them further to be able to reduce them to constraints on disjoint sets. We then give a BAPA reduction [25] for each of these refined models.

Predecessor-Refined Partial Structures. Our normal form of an FLS formula F is given by a disjunction of certain partial models α of $\mathcal{K}_{\text{FLSf}}$. We call these models *predecessor-refined partial models*.

Definition 10. α is a *Predecessor-Refined Partial (PRP) Structure* if it is a partial substructure of a structure in $\mathcal{M}_{\text{FLSf}}$ and the following conditions hold in α for all elements $x, y \in \alpha(\text{list})$:

1. $x \preceq y$ is totally defined on $\alpha(\text{list})$
2. $(x \sqcap y) \in \alpha(\text{list})$. Moreover, if $x, y, (x \sqcap y)$ are three distinct elements, then there exists $x_1 \in \alpha(\text{list})$ such that $x_1 \preceq x$ and $\text{tail}(x_1) = (x \sqcap y)$.
3. if $x \neq y$, if $\text{tail}(x)$ and $\text{tail}(y)$ are defined and equal, then both $\text{head}(x)$ and $\text{head}(y)$ are defined.

Definition 11. With each PRP structure α we associate the conjunction of literals that are (strongly) satisfied in α . We call this formula a PRP conjunction.

Theorem 12. *Each FLS formula is equivalent to an existentially quantified finite disjunction of PRP conjunctions.*

The proof of the theorem is in Appendix B.2. We can compute the PRP structures for an FLS formula F by using a simple modification of the decision procedure for FLS presented in Section 6.4: instead of instantiating the axioms \mathcal{K}_e of the theory extension only with the ground subterms $\text{st}(K_e, G)$ appearing in the extension axioms K_e and

the clauses G generated from F , we instantiate the axioms with the larger set of ground terms Ψ defined as follows:

$$\begin{aligned}\Psi_0 &= \text{st}(K_e, G) \cup \{t_1 \sqcap t_2 \mid t_1, t_2 \in \text{st}(K_e, G)\} \\ \Psi &= \Psi_0 \cup \{\text{head}(t) \mid t \in \Psi_0\} \cup \{\text{pre}(t_1, t_2), \text{tail}(\text{pre}(t_1, t_2)) \mid t_1, t_2 \in \Psi_0\}\end{aligned}$$

Here pre is a fresh binary function symbol, which we introduce as a Skolem function for the existential variable x_1 in property 3 of PRP structures, i.e., we constrain pre using the following axiom:

$$\text{Pre} : \forall xy. x \neq y \wedge x \neq x \sqcap y \wedge y \neq x \sqcap y \rightarrow \text{pre}(x, y) \preceq x \wedge \text{tail}(\text{pre}(x, y)) = x \sqcap y$$

The PRP structures for F are then given by the partial models of $\mathcal{T}_0 \cup (K_e \cup \{\text{Pre}\})[\Psi] \cup G$ in which all terms in Ψ and \preceq are totally defined. These partial models can be computed using a tool such as H-PILoT [4].

Constraints on Sets of Sublists. Define $\sigma(y) = \{x. x \preceq y\}$. Our goal is to show that extending FLS with the $\sigma(_)$ operator and the set algebra of such sets yields in a decidable logic. To this extent, we consider an FLS formula F with free variables x_1, \dots, x_n and show that the defined relation on sets $\rho = \{(\sigma(x_1), \dots, \sigma(x_n)). F(x_1, \dots, x_n)\}$ is definable as $\rho = \{(s_1, \dots, s_n). G(s_1, \dots, s_n)\}$ for some quantifier-free BAPA [10] formula G . By Theorem 12, it suffices to show this property when F is a PRP conjunction, given by some PRP structure α . Figure 12 shows the generation of set constraints from a PRP structure. By replacing each $\sigma(x)$ with a fresh set variable s_x in the resulting constraint we obtain a formula in set algebra. We can check the satisfiability of such formulas following the algorithms in [9, 10] and generate explicit models or perform synthesis as in [8]. Soundness and completeness of the reduction in Figure 12 are proved in Appendices B.3 and B.4. In Appendix A we demonstrate the reduction for the example shown in Figure 3.

Among the consequences of this reduction is NP-completeness of a logic containing atomic formulas of FLS, along with formulas $s = \sigma(x)$, set algebra expressions containing $\subseteq, \cap, \cup, \setminus, =$ on sets, and the cardinality operator $\text{card}(s)$ that computes the size of the set s along with integer linear arithmetic constraints on such sizes. Because the length of the list x is equal to $\text{card}(\sigma(x)) - 1$, this logic also naturally supports reasoning about list lengths. We note that such a logic can also support a large class of set comprehensions of the form $S = \{x. F(x, y_1, \dots, y_n)\}$ when the atomic formulas within F are of the form $u \preceq v$ and at least one atomic formula of the form $x \preceq y_i$ occurs positively in disjunctive normal form of F . Because $\forall x. F$ is equivalent to $\text{card}(\{x. \neg F\}) = 0$, sets give us a form of universal quantification on top of FLS.

Additional Constraints on List Content. We next extend the previous constraints to impose set constraints not only on the set of sublists $\sigma(x)$ but also on the images of such sets under the head function. We define the list content function by $\tau(x) = \text{head}[\sigma(x) \setminus \{\text{nil}\}]$ where we define $\text{head}[s] = \{\text{head}(x) \mid x \in s\}$. We then obtain our full logic FLS² shown in Figure 8 that introduces constraints of the form $\text{head}[s] = v$ on top of FLS and constraints on sets of sublists. To show decidability of this logic, we

Input: a PRP structure α . **Output:** a set constraint G_α .

Step 1: Define the relation \preceq_1 as irreflexive transitive reduct of \preceq without the tail relation. Formally, for all $x, y \in \alpha(\text{list})$, define $x \preceq_1 y$ iff all of the following conditions hold: (1) $x \preceq y$, (2) $x \neq y$, (3) $\text{tail}(y)$ is undefined, and (4) there is no z in $\alpha(\text{list})$ such that x, y, z are distinct, $x \preceq z$, and $z \preceq y$.

Step 2: Introduce sets $S_{x,y}$ with the meaning $S_{x,y} = (\sigma(y) \setminus \sigma(x)) \setminus \{y\}$ and define $\text{Segs} = \{S_{x,y} \mid x \preceq_1 y\}$.

Step 3: Generate the conjunction \hat{G}_α of the following constraints:

1. $\sigma(\text{nil}) = \{\text{nil}\}$
2. $\sigma(y) = \{y\} \cup \sigma(x)$, for each x, y such that α satisfies $\text{tail}(y) = x$
3. $\sigma(y) = \{y\} \cup S_{x,y} \cup \sigma(x)$, for each x, y such that α satisfies $x \preceq_1 y$
4. $\text{disjoint}((S)_{S \in \text{Segs}}, (\{x\})_{x \in \alpha(\text{list})})$

Step 4: Existentially quantify over all Segs variables in \hat{G}_α . If the goal is to obtain a formula without Segs variables, replace each variable $S_{x,y}$ with $(\sigma(y) \setminus \sigma(x)) \setminus \{y\}$.

Step 5: Return the resulting formula G_α .

Fig. 12: Generation of set constraints from a PRP structure

use techniques inspired by [26] to eliminate the image constraints. The elimination procedure is shown in Figure 13. We use the properties of PRP structures that the elements for which $\text{tail}(x_L) = \text{tail}(x_R)$ holds have defined values $\text{head}(x_L)$ and $\text{head}(x_R)$. This allows us to enforce sufficient conditions on sets of sublists and sets of their heads to ensure that the axiom Pure can be enforced. The elimination procedure assumes that we have $\text{head}(s)$ expressions only in the cases where s is a combination of sets of the form $\sigma(x)$ and $\{x\}$, which ensures that s is a disjoint combination of polynomially many partitions. This restriction is not necessary [26], but is natural in applications and ensures the membership in NP.

8 Conclusion

We presented a new decidable logic that can express interesting properties of functional lists and has a reasonably efficient decision procedure. We showed that this decision procedure can be a useful to increase the degree of automation in verification tools and interactive theorem provers.

References

1. C. Barrett, I. Shikanian, and C. Tinelli. An abstract decision procedure for satisfiability in the theory of recursive data types. *Electronic Notes in Theoretical Computer Science*, 174(8):23–37, 2007.
2. E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.
3. A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. A logic-based framework for reasoning about composite data structures. In *CONCUR*, 2009.
4. C. Ihlemann and V. Sofronie-Stokkermans. System description: H-PILoT. In *CADE*, pages 131–139, 2009.
5. S. Jacobs. Incremental instance generation in local reasoning. In *CAV*, pages 368–382, 2009.

Input: a PRP structure α and an image constraint C .

Output: a set constraint C_α without $\text{head}[s]$ expressions

Step 1: Replace each $\tau(x)$ in C with $\text{head}[\sigma(x) \setminus \{\text{nil}\}]$.

Step 2: Let P_i be all sets of the form $\{x_i\}$ or S_{x_i, x_j} from Figure 12. If s is a boolean combination of expressions of the form $\sigma(x)$, $\{x\}$, let $J(s)$ be such that $s = \bigcup_{i \in J(s)} P_i$ is the decomposition of s into disjoint sets, derived from set equalities in Figure 12. Then replace each expression $\text{head}[s]$ with $\bigcup_{i \in J(s)} \text{head}[P_i]$.

Step 3: Replace each $\text{head}[P_i]$ with a fresh set variable Q_i and conjoin the result with the following constraints on the image sets Q_i :

1. $\text{card}(Q_i) \leq \text{card}(P_i)$
2. $Q_i = \emptyset \rightarrow P_i = \emptyset$
3. $Q_i \cap Q_j = \emptyset$, for each $x, y \in \alpha(\text{list})$ such that $P_i = \{x\}$, $P_j = \{y\}$, $x \neq y$ and $\text{tail}(x) = \text{tail}(y)$.

Step 4: Existentially quantify over all Q_i and return the resulting formula C_α .

Fig. 13: Eliminating $\text{head}[s]$ from image constraints by introducing additional constraints on top of Figure 12.

6. J. Jaffar. Minimal and complete word unification. *J. ACM*, 37(1):47–85, 1990.
7. V. Kuncak and D. Jackson. Relational analysis of algebraic datatypes. In *Joint 10th ESEC and 13th ACM SIGSOFT FSE*, 2005.
8. V. Kuncak, M. Mayer, R. Piskac, and P. Suter. Complete functional synthesis. In *PLDI*, 2010.
9. V. Kuncak, H. H. Nguyen, and M. Rinard. Deciding Boolean Algebra with Presburger Arithmetic. *J. of Automated Reasoning*, 2006.
10. V. Kuncak and M. Rinard. Towards efficient satisfiability checking for Boolean Algebra with Presburger Arithmetic. In *CADE-21*, 2007.
11. S. Lahiri and S. Qadeer. Back to the future: revisiting precise program verification using SMT solvers. In *POPL*, 2008.
12. T. Lev-Ami, N. Immerman, T. Reps, M. Sagiv, S. Srivastava, and G. Yorsh. Simulating reachability using first-order logic with applications to verification of linked data structures. In *CADE-20*, 2005.
13. G. Makanin. The problem of solvability of equations in a free semigroup. *Math. USSR Sbornik*, pages 129–198, 1977. (In AMS, (1979)).
14. H. H. Nguyen, C. David, S. Qin, and W.-N. Chin. Automated verification of shape, size and bag properties via separation logic. In *VMCAI*, 2007.
15. T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
16. M. Odersky, L. Spoon, and B. Venners. *Programming in Scala: a comprehensive step-by-step guide*. Artima Press, 2008.
17. D. C. Oppen. Reasoning about recursively defined data structures. In *POPL*, pages 151–157, 1978.
18. R. Piskac, P. Suter, and V. Kuncak. On decision procedures for ordered collections. Technical Report LARA-REPORT-2010-001, EPFL, 2010.
19. W. Plandowski. Satisfiability of word equations with constants is in PSPACE. *J. ACM*, 51(3), 2004.
20. V. Sofronie-Stokkermans. Hierarchic reasoning in local theory extensions. In *CADE*, pages 219–234, 2005.

21. V. Sofronie-Stokkermans. Locality results for certain extensions of theories with bridging functions. In *CADE*, 2009.
22. P. Suter, M. Dotta, and V. Kuncak. Decision procedures for algebraic data types with abstractions. In *POPL*, 2010.
23. 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.
24. T. Wies, M. Muñiz, and V. Kuncak. An efficient decision procedure for imperative tree data structures. In *CADE*, 2011.
25. T. Wies, R. Piskac, and V. Kuncak. Combining theories with shared set operations. In *FroCoS*, 2009.
26. K. Yessenov, V. Kuncak, and R. Piskac. Collections, cardinalities, and relations. In *VMCAI*, 2010.
27. K. Zee, V. Kuncak, and M. Rinard. Full functional verification of linked data structures. In *PLDI*, 2008.
28. T. Zhang, H. B. Sipma, and Z. Manna. Decision procedures for term algebras with integer constraints. *Inf. Comput.*, 204(10):1526–1574, 2006.

A BAPA Reduction for Drop Example

We use the procedure given in Section 7 to reduce the FLS conjunct shown in Figure 3 to a BAPA formula over the set variables Xs , Ys , Zs . Figure 14 depicts the PRP structures that are equivalent to the FLS conjunct in Figure 3. Every node denotes an element of sort list. The square nodes denote the nil element.




Fig. 14: PRP structures for the FLS conjunct in Figure 3

Figure 15 shows the set constraints generated from the four PRP structures. There is one disjunct for each PRP in Figure 14. In each disjunct we can now replace the sets $\sigma(xs)$, $\sigma(ys)$, and $\sigma(zs)$ by the shared set variables Xs , Ys , and Zs . Existentially quantifying over all remaining variables and applying quantifier elimination to the resulting disjunction then yields the BAPA formula

$$Zs \subseteq Ys \wedge Ys \subseteq Xs \wedge \text{card}(Xs) > 1 \wedge \text{card}(Xs) = \text{card}(Ys) + 1 .$$

This formula is the projection onto the shared set variables shown in Figure 3.

$$\begin{aligned}
& \exists S_{\text{nil},\text{ys}}. \sigma(\text{nil}) = \{\text{nil}\} \wedge \sigma(\text{ys}) = \{\text{ys}\} \cup S_{\text{nil},\text{ys}} \cup \sigma(\text{nil}) \wedge \sigma(\text{zs}) = \sigma(\text{ys}) \wedge \\
& \quad \sigma(\text{xs}) = \{\text{xs}\} \cup \sigma(\text{ys}) \wedge S_{\text{nil},\text{ys}} = (\sigma(\text{ys}) \setminus \sigma(\text{nil})) \setminus \{\text{ys}\} \wedge \\
& \quad \text{disjoint}(S_{\text{nil},\text{ys}}, \{\text{xs}\}, \{\text{ys}\}, \{\text{nil}\}) \\
\vee & \exists S_{\text{nil},\text{ys}}. \sigma(\text{nil}) = \{\text{nil}\} \wedge \sigma(\text{zs}) = \sigma(\text{nil}) \wedge \sigma(\text{ys}) = \{\text{ys}\} \cup S_{\text{nil},\text{ys}} \cup \sigma(\text{nil}) \wedge \\
& \quad \sigma(\text{xs}) = \{\text{xs}\} \cup \sigma(\text{ys}) \wedge S_{\text{nil},\text{ys}} = (\sigma(\text{ys}) \setminus \sigma(\text{nil})) \setminus \{\text{ys}\} \wedge \\
& \quad \text{disjoint}(S_{\text{nil},\text{ys}}, \{\text{xs}\}, \{\text{ys}\}, \{\text{nil}\}) \\
\vee & \exists S_{\text{zs},\text{ys}}, S_{\text{nil},\text{zs}}. \sigma(\text{nil}) = \{\text{nil}\} \wedge \sigma(\text{zs}) = \{\text{zs}\} \cup S_{\text{nil},\text{zs}} \wedge \sigma(\text{ys}) = \{\text{ys}\} \cup S_{\text{zs},\text{ys}} \cup \sigma(\text{zs}) \wedge \\
& \quad \sigma(\text{xs}) = \{\text{xs}\} \cup \sigma(\text{ys}) \wedge S_{\text{nil},\text{zs}} = (\sigma(\text{zs}) \setminus \sigma(\text{nil})) \setminus \{\text{zs}\} \wedge \\
& \quad S_{\text{zs},\text{ys}} = (\sigma(\text{ys}) \setminus \sigma(\text{zs})) \setminus \{\text{ys}\} \wedge \text{disjoint}(S_{\text{nil},\text{zs}}, S_{\text{zs},\text{ys}}, \{\text{xs}\}, \{\text{ys}\}, \{\text{zs}\}, \{\text{nil}\}) \\
\vee & \sigma(\text{nil}) = \{\text{nil}\} \wedge \sigma(\text{ys}) = \sigma(\text{nil}) \wedge \sigma(\text{zs}) = \sigma(\text{nil}) \wedge \sigma(\text{xs}) = \{\text{xs}\} \cup \sigma(\text{nil}) \wedge \\
& \quad \text{disjoint}(\{\text{xs}\}, \{\text{nil}\})
\end{aligned}$$

Fig. 15: Set constraints generated from PRP structures in Figure 14.

B Additional Proofs

B.1 Proof of Proposition 5

Proposition 5. *Every finite model of $\mathcal{K}_{\text{FLSf}}$ is isomorphic to some structure in $\mathcal{M}_{\text{FLSf}}$.*

Proof. Let α be a finite model of $\mathcal{K}_{\text{FLSf}}$. In the following, for a function symbol f , we use f_α as a short-hand for $\alpha(f)$. Similarly, we denote by tail_α^* the reflexive transitive closure of the function $\alpha(\text{tail})$.

First, note that the axioms Refl, Trans, AntiSym, and GCS1, GCS2, GCS3 ensure that $(\alpha(\text{list}), \preceq_\alpha, \sqcap_\alpha)$ is a meet-semilattice. We next prove that \preceq_α is the inverse of tail_α^* .

For proving $\text{tail}_\alpha^* \subseteq \preceq_\alpha^{-1}$, let $u, v \in \alpha(\text{list})$ such that $(u, v) \in \text{tail}_\alpha^*$. Then there exist u_1, \dots, u_n such that $u = u_1, v = u_n$, and for all $1 \leq i < n$, $\text{tail}_\alpha(u_i) = u_{i+1}$. If $n = 1$ then $u = v$ and by axiom Refl we immediately have $v \preceq_\alpha u$. If on the other hand $u \neq v$ then by UnfoldL we have for all $1 \leq i < n$, $\text{tail}_\alpha(u_i) \preceq_\alpha u_i$ and thus $u_{i+1} \preceq_\alpha u_i$. Using axiom Trans we then conclude by induction on i that for all $1 < i \leq n$, $u_i \preceq_\alpha u_1$. Hence, $v \preceq_\alpha u$.

For proving the other direction, let $u \in \alpha(\text{list})$ and let $S_u = \{v \mid v \preceq_\alpha u\}$. We show that for all $v \in S_u$, $(u, v) \in \text{tail}_\alpha^*$. Since α is finite, S_u is finite, as well. Thus, using axioms Total and AntiSym we can construct an enumeration u_1, \dots, u_n of the elements of S_u such that for all $1 \leq i < j \leq n$, $u_j \preceq_\alpha u_i$ but not $u_i \preceq_\alpha u_j$. In particular $u_1 = u$. We prove by induction on i that for all $1 \leq i \leq n$, $(u_1, u_i) \in \text{tail}_\alpha^*$. By reflexivity of tail_α^* we immediately have $(u_1, u_1) \in \text{tail}_\alpha^*$. Now assume that $(u_1, u_i) \in \text{tail}_\alpha^*$. Since $u_i \in S_u$, we know by UnfoldL and Trans that $\text{tail}_\alpha(u_i) \in S_u$. Hence, $\text{tail}_\alpha(u_i) = u_j$ for some $j \geq i$. By UnfoldR we know that for all $j \geq i$, $u_j = u_i$ or $u_j \preceq_\alpha \text{tail}_\alpha(u_i)$. Hence, by construction of the enumeration it follows that

either $u_i = \text{tail}_\alpha(u_i)$ or $u_{i+1} = \text{tail}_\alpha(u_i)$. Assume $u_i = \text{tail}_\alpha(u_i)$. Then from axiom NoCycle2 follows that $u_i = \alpha(\text{nil})$. However, we do not have $u_{i+1} \preceq_\alpha u_i$, which contradicts the fact that α satisfies axiom NoCycle1. Hence, we must have $u_{i+1} = \text{tail}_\alpha(u_i)$. Together with the induction hypothesis we then conclude $(u_1, u_{i+1}) \in \text{tail}_\alpha^*$.

We can now define a structure isomorphism ϕ between α and some structure $\alpha' \in \mathcal{M}_{\text{FLSf}}$, i.e., for each sort s , ϕ is a structure-preserving bijection from $\alpha(s)$ to $\alpha'(s)$. Since all structures agree on the interpretation of sort `bool`, we first define ϕ as the identity mapping on $\alpha(\text{bool})$. For the sort `data`, we let ϕ be some injective mapping from $\alpha(\text{data})$ to the set of data terms \mathbb{D} , such that $\phi(\alpha(\text{nil})) = d_1$. In order to define ϕ on $\alpha(\text{list})$, note that α satisfies axioms NoCycle1 and AntiSym. Thus, we know that the inverse of the relation tail_α^* is well-founded and $\alpha(\text{nil})$ is the smallest element of $\alpha(\text{list})$ with respect to this relation. We can hence recursively define ϕ as a mapping from $\alpha(\text{list})$ to \mathbb{L} as follows:

$$\phi(u) = \begin{cases} \text{nil} & \text{if } u = \alpha(\text{nil}) \\ \text{cons}(\phi(v), \phi(w)) & \text{if } \alpha(\text{tail})(u) = v \text{ and } \alpha(\text{head})(u) = w \end{cases}$$

From axiom Pure follows that ϕ is an injective mapping from $\alpha(\text{list})$ to \mathbb{L} . Now define the structure α' as follows: for all sorts $s \in \Sigma_{\text{FLSf}}$, let $\alpha'(s) = \phi(\alpha(s))$, and for all function symbols $f \in \Sigma_{\text{FLSf}}$ of sort $s_1 \times \dots \times s_n \rightarrow s$, let $\alpha'(f)(u_1, \dots, u_n) = \phi(\alpha(f)(\phi^{-1}(u_1), \dots, \phi^{-1}(u_n)))$. By construction $\alpha' \in \mathcal{M}_{\text{FLSf}}$ and α' is isomorphic to α . \square

B.2 Proof of Theorem 12

Consider a FLS formula F . It suffices to show that each partial substructure can be represented by finitely many PRP structures.

To ensure that a relation (such as $x \preceq y$) is defined, we perform case analysis on whether $x \preceq y$ holds or not. To ensure the remaining properties, proceed as follows.

Define an *ultimately converging triple* to be a triple of distinct elements (x_L, y, x_R) such that $y \preceq x_L$ and $y \preceq x_R$. Then define a *converging triple* to be an ultimately converging triple such that there is no distinct ultimately converging triple (x'_L, y', x'_R) with the property $x'_L \preceq x_L, z' \preceq z, x'_R \preceq x_R$.

An *unresolved converging triple* is a converging triple (x_L, y, x_R) such that $x_L \sqcap x_R$ is not defined in the structure. Given an unresolved converging triple (x_L, y, x_R) , we consider the case 1) $x_L \sqcap x_R = y$ and the case 2) $x_L \sqcap x_R = z$ for a fresh elements of the structure such that $x_L \sqcap x_R = z$ and $y \preceq z$. In each step of this process the number of unresolved converging triples reduces by one, so we can construct finitely many structures where all converging triples are resolved.

In the next step, we ensure that for all converging triples (x_L, y, x_R) we have that $\text{tail}(x_L)$ and $\text{tail}(x_R)$ are defined and equal to y . We call a triple for which this does not hold *non-refined*. Given a non-refined triple (x_L, y, x_R) then it cannot be the case that both $\text{tail}(x_L)$ or $\text{tail}(x_R)$ are undefined, otherwise there would be a converging triple $(\text{tail}(x_L), y, x_R)$ or $(x_L, y, \text{tail}(x_R))$. Suppose without loss of generality that $\text{tail}(x_L)$ is undefined. We then consider either $\text{tail}(x_L) = y$, or we introduce z such that $\text{tail}(z) = y$

and $z \preceq x_L$. By repeating this process at most twice for each non-refined triple, we ensure that all triples are both resolved and refined.

Suppose every converging triple (x_L, y, x_R) is refined. Then define both $\text{head}(x_L)$ and $\text{head}(y_L)$ to be either existing or fresh elements. The Pure axiom ensures $\text{head}(x_L) \neq \text{head}(y_L)$. As a result we obtain a finite set of PRP structures. By existentially quantifying over the freshly introduced elements we obtain a disjunction of PRP conjunctions equivalent to F . \square

B.3 Soundness of the Reduction in Figure 12

Consider a $\mathcal{M}_{\text{FLSf}}$ element whose substructure is α . We show that the generated formula G_α is a consequence of the PRP conjunction corresponding to α . The condition $\sigma(\text{nil}) = \{\text{nil}\}$ follows from the definition, as well as $\sigma(y) = \{y\} \cup \sigma(x)$ for $\text{tail}(y) = x$. To ensure $\sigma(y) = \{y\} \cup S_{x,y} \cup \sigma(x)$ we define $S_{x,y} = (\sigma(y) \setminus \sigma(x)) \setminus \{y\}$. It remains to show that for this definition the sets $S_{x,y}$ and the singleton sets $\{x\}$ are all disjoint. By construction $\{x\} \cap \{y\} = \emptyset$ for distinct x, y . Consider two distinct set variables $S_{x,y}$ and $S_{u,v}$. Let $z = x \sqcap u$. Then also $z = y \sqcap v$. By definition of $\sigma(y)$ and $\sigma(v)$, we have $\sigma(y) \cap \sigma(v) = \sigma(z)$. We also have $\sigma(z) \subseteq \sigma(x)$ and $\sigma(z) \subseteq \sigma(u)$. This implies that the sets $\sigma(y) \setminus \sigma(x)$ and $\sigma(v) \setminus \sigma(u)$ are disjoint, so $S_{x,y}$ and $S_{u,v}$ are also disjoint. Showing that $\{x\}$ and $S_{u,v}$ are disjoint is similarly straightforward.

B.4 Completeness of the Reduction in Figure 12

Given a partial structure α and the values of sets that satisfy the generated formula G_α in Figure 12, we extend α to a finite total structure in $\mathcal{M}_{\text{FLSf}}$. To do this, consider each pair x, y in the domain of α for which $x \preceq_1 y$. Let k be the size of the set $S_{x,y}$. If $k = 0$ then let $\text{tail}(y) = x$. Otherwise, introduce k fresh and distinct list elements z_1, \dots, z_k and extend tail such that

$$\text{tail}(y) = z_k \wedge \left(\bigwedge_{i=1}^{k-1} \text{tail}(z_{i+1}) = z_i \right) \wedge \text{tail}(z_1) = x$$

To ensure that the elements are distinct, define $\text{head}(z_i) = h_i$ where h_i are fresh head elements. Extend \preceq and \sqcap according to tail to ensure that the structure belongs to $\mathcal{M}_{\text{FLSf}}$. This completes the construction showing the completeness of a reduction from FLS to BAPA.