Verifying Catamorphism-Based Contracts using Constrained Horn Clauses[*]

EMANUELE DE ANGELIS and **MAURIZIO PROIETTI**

IASI-CNR, Rome, Italy

(*e-mails:* emanuele.deangelis@iasi.cnr.it, maurizio.proietti@iasi.cnr.it)

FABIO FIORAVANTI

DEc, University of Chieti-Pescara, Pescara, Italy

(*e-mail:* fabio.fioravanti@unich.it)

ALBERTO PETTOROSSI

DICII, University of Rome "Tor Vergata", Rome, Italy

(*e-mail:* adp@iasi.cnr.it)

Abstract

We address the problem of verifying that the functions of a program meet their contracts, specified by pre/postconditions. We follow an approach based on *constrained Horn clauses* (CHCs) by which the verification problem is reduced to the problem of checking satisfiability of a set of clauses derived from the given program and contracts. We consider programs that manipulate *algebraic data types* (ADTs) and a class of contracts specified by catamorphisms, that is, functions defined by simple recursion schemata on the given ADTs. We show by several examples that state-of-the-art CHC satisfiability tools are not effective at solving the satisfiability problems obtained by direct translation of the contracts into CHCs. To overcome this difficulty, we propose a transformation technique that removes the ADT terms from CHCs and derives new sets of clauses that work on basic sorts only, such as integers and booleans. Thus, when using the derived CHCs there is no need for induction rules on ADTs. We prove that the transformation is sound, that is, if the derived set of CHCs is satisfiable, then so is the original set. We also prove that the transformation always terminates for the class of contracts specified by catamorphisms. Finally, we present the experimental results obtained by an implementation of our technique when verifying many non-trivial contracts for ADT manipulating programs.

1 Introduction

Many program verification techniques are based on the classical axiomatic approach proposed by Hoare (1969), where the functional correctness of a program is specified

---

by a pair of assertions of first order logic: a *precondition*, which is assumed to hold on the program variables before execution, and a *postcondition*, which is expected to hold after execution. This pair of assertions is often referred to as a *contract* (Meyer 1992), and many programming languages provide built-in support for contracts associated with function definitions (see, for instance, Ada Booch and Bryan 1994, Ciao Hermenegildo *et al.* 2012, and Scala Odersky *et al.* 2011). In order to prove that all program functions meet their contracts, *program verifiers* generate verification conditions, that is, formulas of first order logic that have to be discharged by a theorem prover. Recent developments of Satisfiability Modulo Theory (SMT) solvers (de Moura and Bjørner 2008; Barrett *et al.* 2011; Komuravelli *et al.* 2014; Hojjat and Rümmer 2018) provide support for proving verification conditions in a wide range of logical theories that axiomatize data types, such as booleans, uninterpreted functions, linear integer or real arithmetic, bit vectors, arrays, strings, algebraic data types, and heaps. Among the program verifiers that use SMT solvers as a back-end, we mention Boogie (Barnett *et al.* 2006), Dafny (Leino 2013), Leon (Suter *et al.* 2011), Stainless (Hamza *et al.* 2019), and Why3 (Filliâtre and Paskevich 2013). There are, however, various issues that remain to be solved when following this approach to contract verification. For programs manipulating ADTs, like lists or trees, one such issue is that the verifier often has to generate suitable loop invariants whose verification may require the extension of SMT solvers with inductive proof rules (Reynolds and Kuncak 2015).

An alternative approach is based on translating the contract verification problem into an equivalent satisfiability problem for *constrained Horn clauses*[1] (CHCs), that is, Horn clauses extended with logical theories that axiomatize data types like the ones mentioned above (Jaffar and Maher 1994; Grebenshchikov *et al.* 2012; Bjørner *et al.* 2015; De Angelis *et al.* 2021). For clauses extended with theories on basic sorts, such as the theories of boolean values and linear integer arithmetic, various state-of-the-art CHC solvers are available. Among them, let us mention Eldarica (Hojjat and Rümmer 2018) and SPACER (Komuravelli *et al.* 2014) that are quite effective in checking clause satisfiability. For clauses defined on ADTs, some solvers that can handle them, have been recently proposed. They are based on the extension of the satisfiability algorithms by induction rules (Unno *et al.* 2017; Yang *et al.* 2019), tree automata (Kostyukov *et al.* 2021), and abstractions (Govind *et al.* 2022).

In this paper, we present a method for proving the satisfiability of CHCs defined on ADTs that avoids the need of extending the satisfiability algorithms and, instead, follows a transformational approach (De Angelis *et al.* 2018, 2022a). A set of CHCs is transformed, by applying the fold/unfold rules (Etalle and Gabbrielli 1996; Tamaki and Sato 1984), into a new set of CHCs such that: (i) the ADT terms are no longer present, and hence no induction rules are needed to reason on them, and (ii) the satisfiability of the derived set implies the satisfiability of the original set. The transformational approach has the advantage of separating the concern of dealing with ADTs (which we face at transformation time) from the concern of dealing with simpler, non-inductive constraint theories (which we face at solving time by applying CHC solvers that support basic sorts only). We show that the transformational approach is well suited for a significant class of verification problems where program contracts are specified by means of

---

[1] In recent verification literature, the term *constrained Horn clauses* is often used instead of *constraint logic programs*, as the focus is on their logical meaning rather than their execution as programs.

catamorphisms, that is, functions defined by a simple structural recursion schema over the ADTs manipulated by the program (Meijer *et al.* 1991; Suter *et al.* 2010).

The main contributions of this paper are the following. (i) We define a class of CHCs that represent ADT manipulating programs and their contracts. No restrictions are imposed on programs, while contracts can be specified by means of catamorphisms only (see Section 4). (ii) We define an algorithm that, by making use of the given contract specifications as lemmas, transforms a set of CHCs into a new set of CHCs without ADT terms such that, if the transformed clauses are satisfiable, so are the original ones, and hence the contracts specified by the original clauses are valid (see Section 5). (iii) Unlike previous work (De Angelis *et al.* 2018, 2022a), we prove that the transformation algorithm *terminates* for all sets of CHCs in the given class, and it introduces in a fully automatic way new predicates corresponding to loop invariants (see Section 5). (iv) Finally, by using a prototype implementation of our method, we prove many nontrivial contracts relative to programs that manipulate lists and trees (see Section 6).

2 Preliminaries on constrained Horn clauses

We consider CHCs defined in a many-sorted first order language with equality that includes the language of linear integer arithmetic (*LIA*) and boolean expressions (*Bool*). For notions not recalled here we refer to the literature (Jaffar and Maher 1994; Bjørner *et al.* 2015). A *constraint* is a quantifier-free formula c, where the linear integer constraints may occur as subexpressions of boolean constraints, according to the SMT approach (Barrett *et al.* 2009). The formula c is constructed as follows:

$c ::= d \,|\, B \,|\, true \,|\, false \,|\, {\sim} c \,|\, c_1 \,\&\, c_2 \,|\, c_1 \lor c_2 \,|\, c_1 \Rightarrow c_2 \,|\, c_1 = c_2 \,|\, ite(c, c_1, c_2) \,|\, t = ite(c, t_1, t_2)$
$d ::= t_1 = t_2 \,|\, t_1 \geq t_2 \,|\, t_1 > t_2 \,|\, t_1 \leq t_2 \,|\, t_1 < t_2$

where B is a boolean variable and t, possibly with subscripts, is a *LIA* term of the form $a_0 + a_1 X_1 + \cdots + a_n X_n$ with integer coefficients $a_0, \ldots, a_n$ and variables $X_1, ..., X_n$. The "$\sim$" symbol denotes negation. The ternary function *ite* denotes the if-then-else operator. The equality "$=$" symbol is used for both integers and booleans.

An *atom* is a formula of the form $p(t_1, \ldots, t_m)$, where p is a predicate symbol not occurring in $LIA \cup Bool$, and $t_1, \ldots, t_m$ are first order terms. A *constrained Horn clause* (or a CHC, or simply, a *clause*) is an implication of the form $H \leftarrow c, G$. The conclusion (or *head*) H is either an atom or *false*, the premise (or *body*) is the conjunction of a constraint c and a (possibly empty) conjunction G of atoms. A clause is called a *goal* if its head is *false*, and a *definite clause*, otherwise. Without loss of generality, we assume that every atom occurring in the body of a clause has distinct variables (of any sort) as arguments. By $vars(e)$ we denote the set of all variables occurring in an expression e. Given a formula φ, we denote by $\forall(\varphi)$ its *universal closure*. Let $\mathbb{D}$ be the usual interpretation for the symbols of theory $LIA \cup Bool$. By $M(P)$ we denote the *least* $\mathbb{D}$-model of a set P of definite clauses (Jaffar and Maher 1994). In the examples, we will use the Prolog syntax and the teletype font. Moreover, we will often prefer writing `B1` and `~B2`, instead of the equivalent constraints `B1=true` and `B2=false`, respectively.

3 A motivating example

The CHC translation of a contract verification problem for a functional or an imperative program (Grebenshchikov *et al.* 2012; De Angelis *et al.* 2021) produces three sets of

```
                    /* ------ Program Reverse ------ */
 1. rev([],[]).
 2. rev([H|T],R) :- rev(T,S), snoc(S,H,R).
 3. snoc([],X,[X]).
 4. snoc([X|Xs],Y,[X|Zs]) :- snoc(Xs,Y,Zs).
                    /* ------ Program properties ------ */
 5. is_asorted([],Res) :- Res.
 6. is_asorted([H|T],Res) :- Res = (IsDefHdT => (H=<HdT & ResT)),
                             hd(T,IsDefHdT,HdT), is_asorted(T,ResT).
 7. is_dsorted([],Res) :- Res.
 8. is_dsorted([H|T],Res) :- Res = (IsDefHdT => (H>=HdT & ResT)),
                             hd(T,IsDefHdT,HdT), is_dsorted(T,ResT).
 9. hd([],IsDefHd,Hd) :- ~IsDefHd & Hd=0.    /* hd computes the head of a list.        */
10. hd([H|T],IsDefHd,Hd) :- IsDefHd & Hd=H.  /* IsDefHd=true iff the list is not empty. */
11. leq_all(X,[],Res) :- Res.                /* leq_all(X,L,true) iff for all Y in L, X=<Y. */
12. leq_all(X,[H|T],Res) :- Res = (X=<H & R), leq_all(X,T,R).
                    /* ------ Contracts in goal form ------ */
13. false :- (BL & ~BR), rev(L,R), is_asorted(L,BL), is_dsorted(R,BR).
14. false :- (BA & BX & ~BC), snoc(A,X,C), is_dsorted(A,BA),
             leq_all(X,A,BX), is_dsorted(C,BC).
```

Fig. 1. The set *Reverse* of CHCs. For technical reasons (see Definition 1) all program
properties are defined by *total functions*. In particular, for the empty list, the `hd` function
returns the arbitrarily chosen value `0` (which is never used).

clauses, as shown in Figure 1, where we refer to a program that reverses a list of integers
(we omit the source functional program for lack of space). The first set (clauses `1`–`4`)
is the translation of the operational semantics of the program. The second set (clauses
`5`–`12`) is the translation of the properties needed for specifying the contracts. The third
set (goals `13`–`14`) is the translation of the contracts for the `rev` and `snoc` functions. In
particular, goal `13` is the translation of the contract for `rev`, which, written in relational
form, is the following universally quantified implication:

$\forall$ L,R. is_asorted(L,true) $\wedge$ rev(L,R) $\rightarrow$ is_dsorted(R,true)

The atoms is_asorted(L,true) and is_dsorted(R,true) are the precondition and the post-
condition for `rev`, respectively, stating that, if a list L of integers is sorted in ascending
order with respect to the "$\leq$" relation, then the list R computed by the `rev` function for
input L is sorted in descending order.

The problem of checking the validity of the contracts for the functions `rev` and `snoc`
is reduced to the problem of proving the satisfiability of the set *Reverse* of clauses shown
in Figure 1. The set *Reverse* is indeed satisfiable, but state-of-the-art CHC solvers, such
as Eldarica and SPACER, fail to prove its satisfiability. This is basically due to the fact
that those solvers lack any form of inductive reasoning on lists and, moreover, they do
not use the information about the validity of the contract for `snoc` during the proof of
satisfiability of the goal representing the contract for `rev`.

The algorithm we will present in Section 5 transforms the set *Reverse* of clauses into
a new set *TransfReverse* of clauses (see Figure 2) without occurrences of list terms, such
that if *TransfReverse* is satisfiable, so is *Reverse*. Since in the set *TransfReverse* there are
only integer and boolean terms, no induction rule is needed for proving its satisfiability.

The transformation works by introducing, for each predicate p representing a program
function (in our case, `rev` and `snoc`), a new predicate symbol *newp* (in our case, `new3` and
`new7`) defined in terms of p together with predicates defining program properties used in
the contracts (in our case, is_asorted, is_dsorted, hd, and leq_all). The arguments of

```
T1. new7(A,B,C,D,E,F,G,H,D,I,J) :- A & B=D & C=(K=>((D>=L) & M)) & E & ~F &
              G=0 & H & (J=((I=<D) & N)) & M & ~K & L=0 & N.
T2. new7(A,B,C,D,E,F,G,H,D,I,J) :- A & B=K & C=(L=>((K>=M) & N)) &
              E=((D=<K) & T) & F & G=K & H=(P=>((K>=Q) & R)) & J=((I=<K) & S) & (R & T)=>N,
              new7(L,M,N,D,T,P,Q,R,D,I,S).
T3. new3(A,B,C,D,E,F) :- A & C & ~D & E=0 & F.
T4. new3(A,B,C,D,E,F) :- D & E=G & F=(H=>((G=<I) & J)) & J=>K & (K & L)=>A,
              new3(K,G,L,H,I,J), new7(M,N,A,G,L,T,P,K,G,B,C).
T5. false :- (A & ~B), new3(B,C,D,E,F,A).
```

Fig. 2. The set *TransfReverse* of transformed CHCs. The clauses shown here are those derived
from clauses 1-12 of *Reverse* and goal 13. The clauses derived from goal 14 are listed in
Appendix A of the extended version of this paper (De Angelis *et al.* 2022b).

newp are the variables of basic sorts occurring in the body of its defining clause, and
hence *newp* specifies a relation among the values of the catamorphisms that are applied
to *p*. For example, the transformation algorithm introduces the following predicate new3:

```
new3(K,D,J,G,H,I) :- is_asorted(F,I), hd(F,G,H), rev(F,C),
                     is_dsorted(C,K), leq_all(D,C,J).
```

Then by applying the fold/unfold transformation rules, the algorithm derives a recursive
definition of *newp*. During the transformation, the algorithm makes use of the user-
provided contracts as lemmas, thus adding new constraints that ease the subsequent
satisfiability proof. By construction, the recursive definition of *newp* is a set of clauses
that do *not* manipulate ADTs. Note that while the contract specifications are provided
by the users, the introduction of the new predicate definitions, which is the key step
in our transformation algorithm is done in a fully automatic way, as we will show in
Section 5.

   If the predicates defining program properties are in the class of catamorphisms (for-
mally defined in Section 4), then our transformation is guaranteed to terminate. Thus,
we eventually get, as desired, a set of clauses that are defined on basic sorts only, whose
satisfiability can be checked by CHC solvers that handle the *LIA* ∪ *Bool* theory. In our
example, both Eldarica and SPACER are able to show the satisfiability of *TransfReverse*.

4 Specifying contracts using catamorphisms

The notion of a catamorphism has been popularized in the field of functional program-
ming (Meijer *et al.* 1991) and many generalizations of it have been proposed in the
literature (Hinze *et al.* 2013). Catamorphisms have also been considered in the context of
many-sorted first order logic with recursively defined functions (Suter *et al.* 2010; Pham
et al. 2016; Govind *et al.* 2022), as we do in this paper.

   Let f be a predicate symbol whose $m+n$ arguments (for $m,n \geq 0$) have sorts $\alpha_1,\ldots,\alpha_m$,
$\beta_1,\ldots,\beta_n$, respectively. We say that f is *functional* from $\alpha_1 \times \ldots \times \alpha_m$ to $\beta_1 \times \ldots \times \beta_n$, with
respect to a set P of definite clauses, if $M(P) \models \forall X, Y, Z. f(X,Y) \wedge f(X,Z) \rightarrow Y = Z$,
where X is an m-tuple of distinct variables, and Y and Z are n-tuples of distinct variables.
X and Y are said to be tuples of the *input* and *output* variables of f, respectively.
Predicate f is said to be *total* if $M(P) \models \forall X \exists Y. f(X,Y)$. In what follows, a "total,
functional predicate" f from α to β will be called a "total function" and denoted by
$f \in [\alpha \rightarrow \beta]$ (the set P of clauses that define f will be understood from the context).

(A)
```
h(X,[],Res) :- base1(X,Res).
h(X,[H|T],Res) :-
    h(X,T,R),
    combine1(X,H,R,Res).
```

(B)
```
t(X,leaf,Res) :- base2(X,Res).
t(X,node(L,N,R),Res) :-
    t(X,L,RL), t(X,R,RR),
    combine2(X,N,RL,RR,Res).
```

Fig. 3. (A) List catamorphism. (B) Tree catamorphism.

(C)
```
h(X,[],Res) :- base3(X,Res).
h(X,[H|T],Res) :-
    h(X,T,R),
    f(X,T,Rf),
    combine3(X,H,R,Rf,Res).
```

(D)
```
t(X,leaf,Res) :- base4(X,Res).
t(X,node(L,N,R),Res) :-
    t(X,L,RL), t(X,R,RR),
    g(X,L,RLg), g(X,R,RRg),
    combine4(X,N,RL,RR,RLg,RRg,Res).
```

Fig. 4. (C) Generalized list catamorphism. (D) Generalized tree catamorphism.

Definition 1 (*Catamorphisms*)

A *list catamorphism*, shown in Figure 3(A), is a total function $h \in [\sigma \times \mathtt{list}(\beta) \to \varrho]$, where: (i) σ, β, and ϱ are (products of) basic sorts, (ii) $\mathtt{list}(\beta)$ is the sort of any list of elements each of which is of sort β, (iii) base1 is a total function in $[\sigma \to \varrho]$, and (iv) combine1 is a total function in $[\sigma \times \beta \times \varrho \to \varrho]$. Similarly, a (*binary*) *tree catamorphism* (or a *tree catamorphism*, for short) is a total function $t \in [\sigma \times \mathtt{tree}(\beta) \to \varrho]$ defined as shown in Figure 3(B).

The parameter X and some atoms in the body of the clauses in Figure 3 may be absent (see, for instance, the predicate hd in Figure 1). The definition of catamorphisms we consider here slightly extends the usual first order definitions (Suter *et al.* 2010; Pham *et al.* 2016; Govind *et al.* 2022) by allowing the parameter X, which gives an extra flexibility for specifying contracts. Catamorphisms with parameters have also been considered in functional programming (Hinze *et al.* 2013). To see an example, the predicate leq_all (see Figure 1) is a catamorphism in $[\mathtt{int} \times \mathtt{list}(\mathtt{int}) \to \mathtt{bool}]$, where: (i) base1(X,Res) is the function in $[\mathtt{int} \to \mathtt{bool}]$ defined by the constraint Res = true (i.e. it binds the output boolean variable Res to true), and (ii) combine1(X,H,R,Res) is the function in $[\mathtt{int} \times \mathtt{int} \times \mathtt{bool} \to \mathtt{bool}]$ defined by the $LIA \cup Bool$ constraint Res = (X=<H & R).

The schemata presented in Figure 3 can be extended by adding to the bodies of the clauses extra atoms that have the list tail or the left and right subtrees as arguments. These extensions are shown in Figure 4, where base3, combine3, base4, and combine4 are total functions on basic sorts, and f and g are defined by instances of the same schemata (C) and (D), respectively. Strictly speaking, these schemata are a CHC translation of the *zygomorphism* recursion schemata (Hinze *et al.* 2013), extended with parameter X. However, schemata (C) and (D) can be transformed into schemata (A) and (B), respectively (see Appendix B of the extended version of this paper (De Angelis *et al.* 2022b)), and hence we prefer not to introduce a different terminology and we call them simply catamorphisms.

Examples of list catamorphisms that are instances of the schema of Figure 3(A) are the functions is_asorted and is_dsorted shown in Figure 1 of Section 3. For instance, is_asorted is a catamorphism in $[\mathtt{list}(\mathtt{int}) \to \mathtt{bool}]$, where: (i) base3 is the constant function defined by the constraint Res = true, (ii) the auxiliary function f is

```
count(X,[],N) :- N = 0.
count(X,[H|T],N) :- count(X,T,NT), N = ite(X=H,NT+1,NT).
bstree(leaf,B) :- B.
bstree(node(L,N,R),B) :- bstree(L,BL), bstree(R,BR), treemax(L,IsDefL,MaxL),
    treemin(R,IsDefR,MinR),  (GrtLeft = (IsDefL => N>MaxL)) &
    (LessRight = (IsDefR => N<MinR)) & (B = (BL & BR & GrtLeft & LessRight)).
```

Fig. 5. The catamorphisms `count` and `bstree`.

`hd` $\in$ [`list(int)` $\rightarrow$ `bool`$\times$`int`], and (iii) `combine3` is the function in [`bool`$\times$`int`$\times$`int`$\times$`bool` $\rightarrow$ `bool`] defined by the $LIA \cup Bool$ constraint `Res = (IsDefHdT => (H=<HdT & ResT))`. Two more examples are given in Figure 5, where `count` counts the occurrences of a given element in a list, `bstree` checks whether or not a tree is a binary search tree (duplicate keys are not allowed), `treemax` and `treemin` compute, respectively, the maximum and the minimum element in a binary tree.

In the sets of CHCs we consider, we identify two disjoint sets of predicates: (1) the *program predicates*, defined by any set of CHCs not containing occurrences of catamorphisms, and (2) the *catamorphisms*, defined by instances of the schemata in Figure 4. An atom is said to be a *program atom* (or a *catamorphism atom*) if its predicate symbol is a program predicate (or a catamorphism, respectively).

Definition 2
A *contract* is a formula of the form (where implication is right-associative):

$(K) \qquad pred(Z) \rightarrow c, cata_1(X_1, T_1, Y_1), \ldots, cata_n(X_n, T_n, Y_n) \rightarrow d$

where: (i) *pred* is a program predicate and Z is a tuple of distinct variables, (ii) c is a constraint such that $vars(c) \subseteq \{X_1, \ldots, X_n, Z\}$, (iii) $cata_1, \ldots, cata_n$ are catamorphisms, (iv) $X_1, \ldots, X_n, Y_1, \ldots, Y_n$ are pairwise disjoint tuples of distinct variables of basic sort, (v) $T_1, \ldots, T_n$ are ADT variables occurring in Z, and (vi) d is a constraint, called the *postcondition* of the contract, such that $vars(d) \subseteq \{X_1, \ldots, X_n, Y_1, \ldots, Y_n, Z\}$.

The following are the contracts for `rev` and `snoc`.[2]
```
:- spec rev(L,R) ==> is_asorted(L,BL), is_dsorted(R,BR) => (BL=>BR).
:- spec snoc(A,X,C) ==> is_dsorted(A,BA), leq_all(X,A,BX), is_dsorted(C,BC)
                        => ((BX & BA) => BC).
```

Definition 3
Let *Catas* denote the conjunction $cata_1(X_1, T_1, Y_1) \wedge \ldots \wedge cata_n(X_n, T_n, Y_n)$ of the catamorphisms in the contract K (see Definition 2), and let P be a set of definite CHCs. We say that contract K is *valid* (with respect to the set P of CHCs) if $M(P) \models \forall (pred(Z) \wedge c \wedge Catas \rightarrow d)$.

Theorem 1 (Correctness of the CHC translation)
For contract K, let $\gamma(K)$ denote the goal *false* $\leftarrow \neg d, c, pred(Z), Catas$. Contract K is valid with respect to a set P of CHCs if and only if $P \cup \{\gamma(K)\}$ is satisfiable.

The proof of this theorem is given in Appendix C of the extended version of this paper (De Angelis *et al.* 2022b). The use of the catamorphism schemata (C) and (D)

---

[2] In concrete contract specifications we use the keyword "`spec`" and the two distinct implication symbols "`==>`" and "`=>`."

Algorithm $\mathcal{T}_{cata}$.

Input: A set P of definite clauses and a set Cns of contracts, one for each program predicate.
Output: A set $TransfCls$ of clauses (including goals) on basic sorts such that, if $TransfCls$ is satisfiable, then every contract in Cns is valid with respect to P.

$InCls := \{\gamma(K) \mid K \in Cns\}; \quad Defs := \emptyset; \quad OutCls := InCls;$
while $InCls \neq \emptyset$ **do** $Define(InCls, Defs, NewDefs);$
                                     $Unfold(NewDefs, P, UnfCls);$
                                     $Apply\text{-}Contracts(UnfCls, Cns, RCls);$
                                     $InCls := not\text{-}foldable(RCls, Defs);$
                                     $OutCls := OutCls \cup foldable(RCls, Defs);$
$Fold(OutCls, Defs, TransfCls)$

Fig. 6. The Transformation Algorithm $\mathcal{T}_{cata}$.

guarantees the termination of the transformation algorithm (see Theorem 2) and, at same time, allows the specification of many non-trivial contracts (see our benchmark in Section 6). Among the properties that cannot be specified by our notion of cata-morphisms, we mention ADT equality (indeed ADT equality has *more than one* ADT argument). Thus, in particular, the property $\forall$ `L,RR. double-rev(L,RR)` $\rightarrow$ `L=RR`, where `double-rev(L,RR)` holds if the conjunction "`rev(L,R),rev(R,RR)`" holds, cannot be writ-ten as a contract in our framework. We leave it for future work to identify larger classes of contracts that can be handled by our transformation-based approach.

5 Catamorphism-based transformation algorithm

In this section we present Algorithm $\mathcal{T}_{cata}$ (see Figure 6) which, given a set P of definite clauses manipulating ADTs and a set Cns of contracts, derives a set $TransfCls$ of CHCs with new predicates manipulating terms of basic sorts only, such that if $TransfCls$ is satisfiable, then $P \cup \{\gamma(K) \mid K \in Cns\}$ is satisfiable. By Theorem 1, the satisfiability of $TransfCls$ implies the validity of the contracts in Cns with respect to P.

During the **while-do** loop, $\mathcal{T}_{cata}$ iterates the *Define*, *Unfold*, and *Apply-Contracts* procedures as we now explain. Their formal definition is given in Figures 7, 8, and 9.

– Procedure *Define* works by introducing suitable new predicates defined by clauses, called *definitions*, of the form: $newp(U) \leftarrow c, A, Catas_A$, where U is a tuple of variables of basic sort, A is a program atom, and $Catas_A$ is a conjunction of catamorphism atoms whose ADT variables occur in A. Thus, $newp(U)$ defines the projection onto $LIA \cup Bool$ of the relation between the variables of the program atom A and the catamorphisms acting on the ADT variables of A. In particular, for each program atom A occurring in the body "c, G" of a definite clause or a goal in $InCls$, the *Define* procedure may *either* (i) introduce a new definition whose body consists of A, together with the conjunction of all catamorphism atoms in G that share an ADT variable with A, and the constraints on the input variables of A of basic sort (case *Project*) *or* (ii) extend an already introduced definition for the program predicate of A by (ii.1) adding new catamorphism atoms to its body and/or (ii.2) generalizing its constraint (case *Extend*). The new predicate definitions introduced by a single application of the *Define* procedure are collected in *NewDefs*, while the set of all definitions introduced during the execution of $\mathcal{T}_{cata}$ are collected in *Defs*.

Procedure *Define*(*InCls*, *Defs*, *NewDefs*)
Input: A set *InCls* of clauses and a set *Defs* of definitions.
Output: A set *NewDefs* of new definitions.

$NewDefs := \emptyset$;
for each clause C: $H \leftarrow c, G$ in *InCls* and each program atom A in G **do**

let $Catas_A = \bigwedge\{F \mid F$ is a catamorphism atom in G and $adt\text{-}vars(A) \cap adt\text{-}vars(F) \neq \emptyset\}$;

(*Clause is foldable*) **if** in *Defs* there is a clause $newp(U) \leftarrow d, A, B$, with $U = bvars(\{d, A, B\})$, such that: (i) $Catas_A$ is a subconjunction of B, and (ii) $\mathbb{D} \models \forall(c \rightarrow d)$, **then** skip;

(*Extend*) **else if** the maximal definition for the predicate of A in $Defs \cup NewDefs$ is a clause $newp(U) \leftarrow d, A, B$, such that: (i) $Catas_A$ is *not* a subconjunction of B, or (ii) $\mathbb{D} \not\models \forall(c \rightarrow d)$, **then** introduce definition $ExtD$: $extp(V) \leftarrow \alpha(d, c), A, B'$, where: (i) $extp$ is a new predicate symbol, (ii) $V = bvars(\{\alpha(d, c), A, B'\})$, and $B' = \bigwedge\{F \mid F$ occurs either in B or in $Catas_A\}$; $NewDefs := NewDefs \cup \{ExtD\}$; $Defs := Defs \cup \{ExtD\}$;

(*Project*) **else if** there is no clause in *Defs* with A in its body
then introduce definition D: $newp(U) \leftarrow \pi(c, I), A, Catas_A$, where: (i) $newp$ is a new predicate symbol, (ii) I is the tuple of the input variables of basic sort in $\{A, Catas_A\}$, and (iii) $U = bvars(\{\pi(c, I), A, Catas_A\})$; $NewDefs := NewDefs \cup \{D\}$; $Defs := Defs \cup \{D\}$;

Fig. 7. The *Define* procedure.

Procedure *Unfold*(*NewDefs*, P, *UnfCls*)
Input: A set *NewDefs* of definitions and a set P of definite clauses.
Output: A set *UnfCls* of clauses.

$UnfCls := NewDefs$;
(1) (*Unfold program atom*) **for** all clauses D: $newp(U) \leftarrow c, A, Catas_A$ in *UnfCls*, where A has a program predicate **do** $UnfCls := (UnfCls - \{D\}) \cup Unf(D, A, P)$;
(2) (*Unfold catamorphisms*) **while** there exists a clause C: $H \leftarrow d, L, B, R$ in *UnfCls*, for some conjunctions L and R of atoms, such that B is a catamorphism atom whose input argument of ADT sort is not a variable **do** $UnfCls := (UnfCls - \{C\}) \cup Unf(C, B, P)$;
(3) (*Apply Functionality*) **while** there exists a clause C: $H \leftarrow d, L, h(X, Y), h(X, Z), R$ in *UnfCls*, for some catamorphism h **do** $UnfCls := (UnfCls - \{C\}) \cup \{H \leftarrow d, Y = Z, L, h(X, Y), R\}$;

Fig. 8. The *Unfold* procedure.

– Then, Procedure *Unfold* (1) unfolds the program atoms occurring in the body of the definitions belonging to *NewDefs*, and then (2) unfolds all catamorphism atoms whose ADT argument is not a variable. Finally, (3) by the functionality property (see Section 4), repeated occurrences of a function with the same input are removed.

– Next, Procedure *Apply-Contracts* applies the contracts in *Cns* for the program predicates occurring in the body of the clauses generated by the *Unfold* procedure. This is done in two steps. First, (1) for each program atom A in a clause C obtained by unfolding, the procedure adds the catamorphism atoms (with free output variables) that are present in the contract for A and not in the body of C. Note that, since catamorphisms are *total* functions, their addition preserves satisfiability of clauses. Then, (2) if the constraints on the input variables of the added catamorphisms are satisfiable, the procedure adds also the postconditions of the contracts. Thus, the effect of the *Apply-Contracts* procedure is similar to the one produced by the application of user-provided lemmas.

The *Unfold* and *Apply-Contracts* procedures may generate clauses that are not *foldable*, that is, clauses whose body "c, G" contains a conjunction of a program atom and

---

Procedure *Apply-Contracts*(*UnfCls*, *Cns*, *RCls*)

Input: A set *UnfCls* of clauses and a set *Cns* of contracts, one for each program predicate.

Output: A set *RCls* of clauses.

$RCls := \emptyset$;

for each clause C: $H \leftarrow e, G$ in *UnfCls* **do**

  (1) (*Catamorphism Addition*) **for** each program atom A in G **do**

  - let K: $A \rightarrow c, cata_1(X_1, T_1, Y_1), \ldots, cata_n(X_n, T_n, Y_n) \rightarrow d$ be the contract in *Cns* for the predicate of A, with $Y_1, \ldots, Y_n$ variables not occurring in C;

  - **for** each catamorphism atom $cata_i(X_i, T_i, Y_i)$ in K **do**

     if there is no atom $cata_i(V, T_i, W)$ in G **then** add $cata_i(X_i, T_i, Y_i)$ to G;

  (2) (*Constraint Addition*) let E: $H \leftarrow e, G'$ be the clause obtained at the end of Step (1);

  - let $c_1, \ldots, c_k$ be the constraints on the input variables of the catamorphisms in the contracts used for deriving E, and let $d_1, \ldots, d_k$ be the corresponding contract postconditions;

  - let Z be the tuple of variables occurring in $\{c_1, \ldots, c_k\}$ and not in E;

  if $\mathbb{D} \models \forall (e \rightarrow \exists Z. \, c_1 \wedge \ldots \wedge c_k)$ **then** $RCls := RCls \cup \{H \leftarrow e, c_1, \ldots, c_k, d_1, \ldots, d_k, G'\}$

Fig. 9. The *Apply-Contracts* procedure.

Procedure *Fold*(*OutCls*, *Defs*, *TransfCls*)

Input: A set *OutCls* of clauses and a set *Defs* of definitions.

Output: A set *TransfCls* of clauses.

$MDefs := \{D \in Defs \mid D$ is a maximal definition for some program predicate$\}$;

$TbF := \{C \in OutCls \mid$ the head of C is either *false* or its predicate occurs in a head of *MDefs*$\}$;

for each clause C: $H \leftarrow e, G$ in *TbF* and each program atom A in G **do**

  - let $B_A = \bigwedge \{F \mid F$ is a catamorphism atom in G and *adt-vars*$(A) \cap$ *adt-vars*$(F) \neq \emptyset\}$;

  - let D: $newp(U) \leftarrow c, A, Catas_A$ be the definition in *MDefs* for the predicate of A, such that $\mathbb{D} \models \forall (e \rightarrow c)$ and B_A is a subconjunction of $Catas_A$;

  - replace A by $newp(U)$ in the body of C;

Remove all catamorphism atoms from the derived clauses and add them to *TransfCls*;

Fig. 10. The *Fold* procedure.

catamorphism atoms which is not a variant of the body of any definition in *Defs*, or, if there is such a definition in *Defs*, the constraint c does not imply the constraint occurring in that definition. By using the function *not-foldable*, those clauses are added to the set *InCls* of clauses to be further processed by the while-do loop, while the others, by using the function *foldable*, are added to the set *OutCls* of clauses that are output by the loop.

The termination of the while-do loop is guaranteed by the following two facts: (i) there are finitely many catamorphism atoms that can be added to the body of a definition, and (ii) by implementing constraint generalization through a *widening* operator (Cousot and Halbwachs 1978), a most general constraint is eventually computed. When Algorithm $\mathcal{T}_{cata}$ exits the while-do loop, it returns a set *OutCls* of clauses (which are all foldable) and a set *Defs* of new predicate definitions. Then, the *Fold* procedure (Figure 10) uses the definitions in *Defs* for removing ADT variables from the clauses in *OutCls*. By construction, (i) the head of each clause C in *OutCls* is either *false* or an atom $newq(V)$, where V is a tuple of variables of basic sort, and (ii) for each conjunction of a program atom A and catamorphism atoms B in the body of C sharing an ADT variable with A, there is in *Defs* a definition $newp(U) \leftarrow c, A, Catas_A$ such that B is a

subconjunction of $Catas_A$ and c is implied by the constraint in C. The *Fold* procedure removes all ADT variables by replacing in C the atom A by $newp(U)$ and removing the subconjunction B.

Let us introduce some notions used in the procedures *Define*, *Unfold*, *Apply-Contracts*, and *Fold*. Given a conjunction G of atoms, by $bvars(G)$ and $adt\text{-}vars(G)$ we denote the set of variables in G that have a basic sort and an ADT sort, respectively.

Definition 4
The *projection* of a constraint c onto a tuple V of variables is a constraint $\pi(c, V)$ such that: (i) $vars(\pi(c, V)) \subseteq V$ and (ii) $\mathbb{D} \models \forall (c \rightarrow \pi(c, V))$.

A *generalization* of two constraints c_1 and c_2 is a constraint, denoted $\alpha(c_1, c_2)$, such that $\mathbb{D} \models \forall (c_1 \rightarrow \alpha(c_1, c_2))$ and $\mathbb{D} \models \forall (c_2 \rightarrow \alpha(c_1, c_2))$.

Let D: $newp(U) \leftarrow c, A, Catas_A$ be a clause in *Defs*, where: (i) A is a program atom with predicate p, (ii) $Catas_A$ is a conjunction of catamorphism atoms, and (iii) c is a constraint on input variables of A, and U is a tuple of variables of basic sort. We say that D is *maximal* for p if, for all definitions $newq(V) \leftarrow d, A, B$ in *Defs*, we have that: (i) B is a subconjunction of $Catas_A$, (ii) $\mathbb{D} \models \forall (d \rightarrow c)$, and (iii) V is a subtuple of U.

For a concrete definition of a generalization operator, based on widening, we refer to the existing literature (De Angelis *et al.* 2021).

Note that, by the *Extend* case of the *Define* procedure, for every program predicate p occurring in *Defs*, there is a unique maximal definition.

Example 1 (*Reverse, continued*)
InCls is initialized to the set $\{13, 14\}$ of goals (see Figure 1). The while-do loop starts by applying the *Define* procedure to goal 13. (For lack of space, we will not show here the transformation steps starting from goal 14.) No definitions for predicate `rev` are present in *Defs*, and hence the case *Project* applies. Thus, the *Define* procedure introduces the following clause defining the new predicate `new1`:

```
D1. new1(BL,BR) :- is_asorted(L,BL), rev(L,R), is_dsorted(R,BR).
```

where: (i) the body is made out of the program atom `rev(L,R)` and the catamorphisms on the lists `L` and `R` occurring in goal 13, (ii) `BL` and `BR` are the variables of basic sort in the body of `D1`, and (iii) the projection of the constraint of goal 13 onto the (empty) tuple of input variables of basic sort of the body of `D1` is `true` (and thus, omitted). □

Definition 5 (*Unfolding*)
Let C: $H \leftarrow c, G_L, A, G_R$ be a clause, where A is an atom, and let P be a set of definite clauses with $vars(C) \cap vars(P) = \emptyset$. Let *Cls*: $\{K_1 \leftarrow c_1, B_1, \ldots, K_m \leftarrow c_m, B_m\}$, with $m \geq 0$, be the set of clauses in P, such that: for $j = 1, \ldots, m$, (i) there exists a most general unifier ϑ_j of A and K_j, and (ii) the conjunction of constraints $(c, c_j)\vartheta_j$ is satisfiable. We define: $Unf(C, A, P) = \{(H \leftarrow c, c_j, G_L, B_j, G_R)\vartheta_j \mid j = 1, \ldots, m\}$.

Example 2 (*Reverse, continued*)
The *Unfold* procedure first (1) unfolds the program atom `rev(L,R)` in clause `D1`, and then (2) unfolds the catamorphism atoms with non-variable ADT arguments. We get:

```
C1.  new1(A,B) :- A & B.
```

```
C2.  new1(A,B) :- A=(G=>((D=<H) & I)),
        is_asorted(F,I), hd(F,G,H), rev(F,C), snoc(C,D,E), is_dsorted(E,B).
```

Finally, functionality is not applicable, and *Unfold* terminates.   □

Example 3 (*Reverse, continued*)

The *Apply-Contracts* procedure first (1) adds the catamorphism atoms is_dsorted(C,K), leq_all(D,C,J) to the body of clause C2, and then (2) adds the postconditions of the contracts for rev and snoc. We get:

```
C3. new1(A,B) :- A=(G=>(D=<H & I)) & I=>K & (K&J)=>B, is_asorted(F,I), hd(F,G,H),
        rev(F,C), snoc(C,D,E), is_dsorted(C,K), leq_all(D,C,J), is_dsorted(E,B).
```

Now, in the second iteration of the while-do loop of $\mathcal{T}_{cata}$, clause C3 is processed by the *Define* procedure. The following new definition D2 relative to the program atom rev(F,C), is introduced according to the *Extend* case because the catamorphism atoms in C3 that share ADT variables with rev(F,C) is not a subset of the catamorphism atoms in D1.

```
D2. new3(K,D,J,G,H,I) :- is_asorted(F,I), hd(F,G,H), rev(F,C),
        is_dsorted(C,K), leq_all(D,C,J).
```
□

Example 4 (*Reverse, continued*)

When Algorithm $\mathcal{T}_{cata}$ exits the while-do loop, we get a set *OutCls* of clauses including:

```
C4. new3(A,B,C,D,E,F) :- (D & E=G & F=(H=>G=<I & J) & J=>K & (K&L)=>A),
        is_asorted(M,J), hd(M,H,I), rev(M,N), is_dsorted(N,K), leq_all(G,N,L),
        snoc(N,G,V), is_dsorted(V,A), leq_all(B,V,C).
```

and a set *Defs* of definitions including clause D2 and the following clause relative to snoc:

```
D3. new7(A,B,C,D,E,F,G,H,D,I,J) :- hd(K,A,B), is_dsorted(K,C), leq_all(D,L,E),
        hd(L,F,G), is_dsorted(L,H), snoc(L,D,K), leq_all(I,K,J).
```

By the *Fold* procedure, from clause C4, using D2 and D3, we get clause T4 of Figure 2. □

Theorem 2 (*Termination of Algorithm $\mathcal{T}_{cata}$*)

Let P be a set of definite clauses and *Cns* a set of contracts specified by catamorphisms. Then, Algorithm $\mathcal{T}_{cata}$ terminates for P and *Cns*.

Theorem 3 (*Soundness of Algorithm $\mathcal{T}_{cata}$*)

Let P and *Cns* be the input of Algorithm $\mathcal{T}_{cata}$, and let *TransfCls* be the output set of clauses. Then, every clause in *TransfCls* has basic sort, and if *TransfCls* is satisfiable, then all contracts in *Cns* are valid with respect to P.

The proofs of Theorems 2 and 3 are given in Appendix C of the extended version of this paper (De Angelis *et al.* 2022b). The converse of Theorem 3 does not hold. Thus, the unsatisfiability of *TransfCls* means that our transformation technique is unable to prove the validity of the contract at hand, and not necessarily that the contract is not valid.

```
bstdel(X,leaf,leaf).
bstdel(X,node(A,leaf,R),R) :- X=A.
bstdel(X,node(A,L,leaf),L) :- X=A.
bstdel(X,node(A,L,R),node(A1,L,T)) :- X=A & D, delmin(R,T,A1,D).
bstdel(X,node(A,L,R),node(A,L1,R)) :- X<A, bstdel(X,L,L1).
bstdel(X,node(A,L,R),node(A,L,R1)) :- A<X, bstdel(X,R,R1).
delmin(leaf,leaf,M,D) :- M=0 & ~D.
delmin(node(A,leaf,R),R,M,D) :- M=A & D.
delmin(node(A,node(B,U,V),R),node(A,T,R),M,D) :- D, delmin(node(B,U,V),T,M,D).
:- spec bstdel(X,T1,T2) ==> bstree(T1,B1), bstree(T2,B2) => (B1=B2).
```

Fig. 11. Deleting an element from a binary search tree: `bstdel` and its contract.

6 Experimental evaluation

In this section we describe a tool, called VeriCaT, implementing our verification method based on the use of catamorphisms and we present some case studies to which VeriCaT has been successfully applied (see https://fmlab.unich.it/vericat/ for details).

VeriCaT implements the two steps of our method: (i) the *Transf* step, which realizes the CHC transformation algorithm presented in Section 5, and (ii) the *CheckSat* step, which checks CHC satisfiability. For the *Transf* step, VeriCaT uses VeriMAP (De Angelis *et al.* 2014), a tool for fold/unfold transformation of CHCs. It takes as input a set of CHCs representing a program manipulating ADTs, together with its contracts, and returns a new set of CHCs acting on variables of basic sorts only. For the *CheckSat* step, VeriCaT uses SPACER to check the satisfiability of the transformed CHCs.

We have applied our method to prove the validity of contracts for programs implementing various algorithms for concatenating, permuting, reversing, and sorting lists of integers, and also for inserting and deleting elements in binary search trees. The contracts specify properties defined by catamorphisms such as: list length, tree size, tree height, list (or tree) minimum and maximum element, list sortedness (in ascending or descending order), binary search tree property, element sum, and list (or tree) content (defined as sets or multisets of elements). For instance, for the sorting programs implementing Bubblesort, Insertionsort, Mergesort, Quicksort, Selectionsort, and Treesort, VeriCaT was able to prove the following two contracts:

```
:- spec sort(Xs,Ys) ==> is_asorted(Ys,B) => B.
```
stating that the output `Ys` of `sort` is a list in ascending order, and
```
:- spec sort(Xs,Ys) ==> count(X,Xs,N1), count(X,Ys,N2) => N1=N2.
```
stating that the input and output of the `sort` program have the same multiset of integers.

As an example of a verification problem for a tree manipulating program, in Figure 11 we present: (i) the clauses for `bstdel(X,T1,T2)`, which deletes the element `X` from the binary search tree `T1`, thereby deriving the tree `T2`, and (ii) a contract for `bstdel` (catamorphism `bstree` is shown in Figure 5). VeriCaT proved that contract by first transforming the clauses of Figure 11 and the following goal that represents the contract for `bstdel`:

```
false :- ~(B1=B2), bstree(T1,B1), bstdel(X,T1,T2), bstree(T2,B2).
```

In Table 1 we summarize the results of our experiments performed on an Intel Xeon CPU E5-2640 2.00GHz with 64GB RAM under CentOS. The columns report the name of the program, the number of contracts proved by VeriCaT for each program, and the total time, in milliseconds, needed for the *Transf* and *CheckSat* steps. Finally, as a

Table 1. *Contracts proved by the VeriCaT and AdtRem tools. Times are in ms*

Program	VeriCaT: # of proved contracts	VeriCaT: time needed for *Transf*	VeriCaT: time needed for *CheckSat*	AdtRem: # of proved contracts
List Membership	3	4410	160	3
List Permutation	12	17,840	1150	12
List Concatenation	9	14,430	2280	4
Reverse	20	27,790	2350	12
Double Reverse	6	9810	1930	0
Reverse w/Accumulator	6	9780	380	3
Bubblesort	18	27,300	1270	18
Insertionsort	12	17,670	1300	11
Mergesort	25	35,810	2060	17
Quicksort (version 1)	19	27,200	3770	12
Quicksort (version 2)	18	25,930	4700	11
Selectionsort	20	31,160	2680	18
Treesort	10	15,250	3530	1
Binary Search Tree	13	21,330	4820	3
Total	191	285,710	32,380	125

baseline, we report the number of contracts proved by AdtRem (De Angelis *et al.* 2022a), which does not take advantage of the user-provided contract specifications and, instead, tries to discover lemmas by using the differential replacement transformation rule. Our results show that VeriCaT is indeed able to exploit the extra information provided by the contracts, and performs better than previous transformational approaches.

In order to compare the effectiveness of our method with that of other tools, we have also run solvers such as AdtInd (Yang *et al.* 2019), CVC4 extended with induction (Reynolds and Kuncak 2015), Eldarica (2.0.6), and SPACER (with Z3 4.8.12) on the CHC specifications (translated to SMT-LIB format) *before* the application of the *Transf* step. Eldarica and SPACER proved the satisfiability of the CHCs for 12 and 1 contracts, respectively, while the AdtInd and CVC4 did not solve any problem within the time limit of 300 s. However, it might be the case that better results can be achieved by those tools by using some different encodings of the verification problems. Finally, we ran the Stainless verifier (0.9.1) (Hamza *et al.* 2019) on a few manually encoded specifications of programs for reversing a list, deleting an element from a binary search tree, and sorting a list using the Quicksort algorithm. Stainless can prove some, but not all contracts of each of these specifications. For a more exhaustive comparison we would need an automated translator, which is not available yet, between CHCs and Stainless specifications.

7 Related work and conclusions

Many program verifiers are based on Hoare's axiomatic notion of correctness (Hoare 1969) and have the objective of proving the validity of pre/postconditions. Some of those verifiers use SMT solvers as a back-end to check verification conditions in various logical

theories (Barnett *et al.* 2006; Suter *et al.* 2011; Leino 2013; Filliâtre and Paskevich 2013; Hamza *et al.* 2019). In order to deal with properties of programs that manipulate ADTs, program verifiers may also be enhanced with some form of induction (e.g. Dafny Leino 2013), or be based on the unfolding of recursive functions (e.g. Leon Suter *et al.* 2011 and Stainless Hamza *et al.* 2019), or rely on an SMT solver that uses induction, such as the extension of CVC4 proposed by Reynolds and Kuncak (2015).

Catamorphisms were used to define decision procedures for suitable classes of SMT formulas (Suter *et al.* 2010; Pham *et al.* 2016), and a special form of integer-valued catamorphisms, called *type-based norms*, were used for proving termination of logic programs (Bruynooghe *et al.* 2007) and for *resource analysis* (Albert *et al.* 2020). The main difference of our approach with respect to these works is that we transform a set of clauses with catamorphisms into a new set of CHCs that act on the codomains of the catamorphisms, and in those clauses neither ADTs nor catamorphisms are present.

The contracts considered in this paper are similar to `calls` and `success` user-defined assertions supported by the Ciao logic programming system (Hermenegildo *et al.* 2012). Those assertions may refer to operational properties of logic programs, taking into account the order of execution and the extra-logical features of the language, while here we consider only the logical meaning of CHCs and contracts. By using abstract interpretation techniques, the CiaoPP preprocessor (Hermenegildo *et al.* 2005) can statically verify a wide range of `calls`/`success` assertions related to types, modes, non-failure, determinism, and typed-norms. However, CiaoPP suffers from some limitations in checking the validity of properties expressed by constrained types, such as the catamorphisms considered in this paper, for example, the properties of being a sorted list or a binary search tree.

The use of CHCs for program verification has become very popular and many techniques and tools for translating program verification problems into satisfiability problems for CHCs have been proposed (see, for instance, the surveys by Bjørner et al. 2015 and by De Angelis et al. 2021). However, as also shown in this paper, in the case of clauses with ADT terms, state-of-the-art CHC solvers have some severe limitations due to the fact that they do not include any proof technique for inductive reasoning on the ADT structures. Some approaches to mitigate these limitations include: (i) a proof system that combines inductive theorem proving with CHC solving (Unno *et al.* 2017), (ii) lemma generation based on syntax-guided synthesis from user-specified templates (Yang *et al.* 2019), (iii) invariant discovery based on finite tree automata (Kostyukov *et al.* 2021), and (iv) use of suitable abstractions (Govind *et al.* 2022).

Transformation-based approaches to the verification of CHC satisfiability on ADTs have been proposed in recent work (Mordvinov and Fedyukovich 2017; De Angelis *et al.* 2018; Kobayashi *et al.* 2020), with the aim to avoid the complexity of integrating CHC solving with induction. The transformational approach compares well with induction-based solvers, but it also shares similar issues for full mechanization, such as the need for lemma discovery (Yang *et al.* 2019; De Angelis *et al.* 2022a). The transformation technique we have proposed in this paper avoids the problem of lemma discovery by relying on user-specified contracts and, unlike previous work (De Angelis *et al.* 2018, 2022a), it guarantees the termination of the transformation for a large class of CHCs.

Our experiments show that the novel transformation technique we propose can successfully exploit the information supplied by the user-provided contracts, and indeed,

(i) it can increase the effectiveness of state-of-the-art CHC solvers in verifying contracts encoded as CHCs, and (ii) performs better than previous transformational approaches based on lemma discovery.

For future work, we plan to extend the practical applicability of our verification method by developing automatic translators to CHCs of programs and contracts written in the languages used by verifiers such as Dafny, Stainless, and Why3.

References

ALBERT, E., GENAIM, S., GUTIÉRREZ, R. AND MARTIN-MARTIN, E. 2020. A transformational approach to resource analysis with typed-norms inference. *Theory and Practice of Logic Programming 20*, 3, 310–357.

BARNETT, M., CHANG, B.-Y. E., DE LINE, R., JACOBS, B. AND LEINO, K. R. M. 2006. Boogie: A modular reusable verifier for object-oriented programs. In *Formal Methods for Components and Objects*. LNCS, vol. 4111. Springer, 364–387.

BARRETT, C., CONWAY, C. L., DETERS, M., HADAREAN, L., JOVANOVIC, D., KING, T., REYNOLDS, A. AND TINELLI, C. 2011. CVC4. *23rd CAV*. LNCS, vol. 6806. Springer, 171–177.

BARRETT, C. W., SEBASTIANI, R., SESHIA, S. A. AND TINELLI, C. Satisfiability modulo theories. In *Handbook of Satisfiability*. Frontiers in Artificial Intelligence and Applications, vol. 185. IOS Press, 825–885.

BJØRNER, N., GURFINKEL, A., MCMILLAN, K. L. AND RYBALCHENKO, A. 2015. Horn clause solvers for program verification. In *Fields of Logic and Computation* (II). LNCS, vol. 9300. Springer, 24–51.

BOOCH, G. AND BRYAN, D. 1994. *Software Engineering with Ada*, 3rd ed. Series in Object-Oriented Software Engineering. Benjamin/Cummings.

BRUYNOOGHE, M., CODISH, M., GALLAGHER, J. P., GENAIM, S. AND VANHOOF, W. 2007. Termination analysis of logic programs through combination of type-based norms. *ACM Transactions on Programming Languages and Systems 29*, 2, 10–es.

COUSOT, P. AND HALBWACHS, N. 1978. Automatic discovery of linear restraints among variables of a program. In *5th POPL*. ACM, 84–96.

DE ANGELIS, E., FIORAVANTI, F., GALLAGHER, J. P., HERMENEGILDO, M. V., PETTOROSSI, A. AND PROIETTI, M. 2021. Analysis and transformation of constrained Horn clauses for program verification. *Theory and Practice of Logic Programming*, 1–69, doi: 10.1017/S1471068421000211.

DE ANGELIS, E., FIORAVANTI, F., PETTOROSSI, A. AND PROIETTI, M. 2014. VeriMAP: A tool for verifying programs through transformations. In *20th TACAS*. LNCS, vol. 8413. Springer, 568–574.

DE ANGELIS, E., FIORAVANTI, F., PETTOROSSI, A. AND PROIETTI, M. 2018. Solving Horn clauses on inductive data types without induction. *Theory and Practice of Logic Programming 18*, 3-4, 452–469.

DE ANGELIS, E., FIORAVANTI, F., PETTOROSSI, A. AND PROIETTI, M. 2022a. Satisfiability of constrained Horn clauses on algebraic data types: A transformation-based approach. *Journal of Logic and Computation 32*, 2, 402–442.

DE ANGELIS, E., FIORAVANTI, F., PETTOROSSI, A. AND PROIETTI, M. 2022b. Verifying catamorphism-based contracts using constrained Horn clauses (Extended version). CoRR. URL: https://doi.org/10.48550/arXiv.2205.06236.

DE MOURA, L. M. AND BJØRNER, N. 2008. Z3: An efficient SMT solver. In *14th TACAS*. LNCS, vol. 4963. Springer, 337–340.

Etalle, S. and Gabbrielli, M. 1996. Transformations of CLP modules. *Theoretical Computer Science 166*, 101–146.

Filliâtre, J.-C. and Paskevich, A. 2013. Why3 - Where programs meet provers. In *22nd ESOP*. LNCS, vol. 7792. Springer, 125–128.

Govind V. K., H., Shoham, S. and Gurfinkel, A. 2022. Solving constrained Horn clauses modulo algebraic data types and recursive functions. *Proceedings of the ACM on Programming Languages 6*, POPL, 1–29.

Grebenshchikov, S., Lopes, N. P., Popeea, C. and Rybalchenko, A. 2012. Synthesizing software verifiers from proof rules. In *33rd PLDI*, 405–416.

Hamza, J., Voirol, N. and Kuncak, V. 2019. System FR: Formalized foundations for the Stainless verifier. *ACM on Programming Languages 3*, OOPSLA, 166:1–166:30.

Hermenegildo, M., Bueno, F., Carro, M., López-García, P., Mera, E., Morales, J. F. and Puebla, G. 2012. An overview of Ciao and its design philosophy. *Theory and Practice of Logic Programming 12*, 1–2, 219–252.

Hermenegildo, M. V., Puebla, G., Bueno, F. and López-García, P. 2005. Integrated program debugging, verification, and optimization using abstract interpretation (and the Ciao system preprocessor). *Science of Computer Programming 58*, 1-2, 115–140.

Hinze, R., Wu, N. and Gibbons, J. 2013. Unifying structured recursion schemes. In *ICFP 2013*. ACM, 209–220.

Hoare, C. A. R. 1969. An axiomatic basis for computer programming. *CACM 12*, 10, 576–580, 583.

Hojjat, H. and Rümmer, P. 2018. The ELDARICA Horn solver. In *Formal Methods in Computer Aided Design, FMCAD*. IEEE, 1–7.

Jaffar, J. and Maher, M. 1994. Constraint logic programming: A survey. *Journal of Logic Programming 19/20*, 503–581.

Kobayashi, N., Fedyukovich, G. and Gupta, A., 2020. Fold/unfold transformations for fixpoint logic. In *26th TACAS*. LNCS, vol. 12079. Springer, 195–214.

Komuravelli, A., Gurfinkel, A. and Chaki, S. 2014. SMT-based model checking for recursive programs. In *26th CAV*. LNCS, vol. 8559. Springer, 17–34.

Kostyukov, Y., Mordvinov, D. and Fedyukovich, G. 2021. Beyond the elementary representations of program invariants over algebraic data types. In *42nd PLDI*. ACM, 451–465.

Leino, K. R. M. 2013. Developing verified programs with Dafny. In *International Conference on Software Engineering*. IEEE Press, 1488–1490.

Meijer, E., Fokkinga, M. M. and Paterson, R. 1991. Functional programming with bananas, lenses, envelopes and barbed wire. In *5th ACM Conference on Functional Programming Languages and Computer Architecture*. LNCS, vol. 523. Springer, 124–144.

Meyer, B. 1992. Applying "Design by Contract". *Computer 25*, 10, 40–51.

Mordvinov, D. and Fedyukovich, G. 2017. Synchronizing constrained Horn clauses. In *LPAR-21*. EPiC Series in Computing, vol. 46. EasyChair, 338–355.

Odersky, M., Spoon, L. and Venners, B. 2011. *Programming in Scala: A Comprehensive Step-by-Step Guide*, 2nd ed. Artima, Sunnyvale, CA, USA.

Pham, T., Gacek, A. and Whalen, M. W. 2016. Reasoning about algebraic data types with abstractions. *Journal of Automated Reasoning 57*, 4, 281–318.

Reynolds, A. and Kuncak, V. 2015. Induction for SMT solvers. In *16th VMCAI*. LNCS, vol. 8931. Springer, 80–98.

Suter, P., Dotta, M. and Kuncak, V. 2010 Decision procedures for algebraic data types with abstractions. In *37th POPL*. ACM, 199–210.

Suter, P., Köksal, A. S. and Kuncak, V. 2011. Satisfiability modulo recursive programs. In *18th SAS*. LNCS, vol. 6887. Springer, 298–315.

TAMAKI, H. AND SATO, T. 1984. Unfold/fold transformation of logic programs. In *2nd ICLP*. Uppsala University, Sweden, 127–138.

UNNO, H., TORII, S. AND SAKAMOTO, H. 2017. Automating induction for solving Horn clauses. In *29th CAV*. LNCS, vol. 10427. Springer, 571–591.

YANG, W., FEDYUKOVICH, G. AND GUPTA, A. 2019. Lemma synthesis for automating induction over algebraic data types. In *25th CP*. LNCS, vol. 11802. Springer, 600–617.