

On Reductions for Synthesis Procedures

EPFL-REPORT-181031

Swen Jacobs¹, Viktor Kuncak², and Philippe Suter²

¹ Graz University of Technology, Austria

`swen.jacobs@iaik.tugraz.at`

² École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

`firstname.lastname@epfl.ch`

Abstract. A synthesis procedure acts as a compiler for declarative specifications. It accepts a formula describing a relation between inputs and outputs, and generates a function implementing this relation. This paper presents synthesis procedures for data structures. Our procedures are reductions that lift a synthesis procedure for the elements into synthesis procedures for containers storing these elements. We introduce a framework to describe synthesis procedures as systematic applications of inference rules. We show that, by interpreting both synthesis problems and programs as relations, we can derive and modularly prove transformation rules that are widely applicable, thus simplifying both the presentation and the correctness argument.

1 Introduction

Software synthesis is an active area of research [6, 17, 19, 22]. It has received increased attention recently, though it has been studied continuously since the early days of Computer Science [3, 11, 12, 16]. Our paper pursues synthesis of functions from inputs to outputs that are guaranteed to satisfy a given input/output relation expressed in a decidable logic. We call this approach complete functional synthesis [8, 9]. The appeal of this direction is that it can synthesize functions over unbounded domains, and that the produced code is guaranteed to satisfy the specification for the entire unbounded range of inputs.

Previous work described synthesis procedures for linear arithmetic and sets [8, 9] as well as extensions to unbounded bitvector constraints [4, 18]. In this paper we make a step towards systematic derivation of synthesis procedures by showing how inference rules that describe decision procedure steps, possibly for a combination of theories, can be generalized to synthesis procedures. Moreover, we derive new synthesis procedures for two relevant decidable theories of data structures: term algebras (algebraic data types), and the theory of arrays. The two synthesis procedures that we present are interesting in their own right. Synthesis for algebraic data types can be viewed as a generalization of the compilation of pattern matching, and is therefore a useful way to increase the expressive power of functional programs. Synthesis for arrays is useful also for

imperative programs. On the technical side, the synthesis for arrays is challenging because it requires iteration over the space of indices. It therefore illustrates the importance of synthesizing not only individual values that meet a constraint, but also functions that enumerate all values.

Our synthesis procedures are expressed as a set of modular transformation rules whose correctness can be checked in a straightforward way, and more easily implemented (even in foundational proof assistants). The transformations gradually evolve a constraint into a program. Sound rules for such transformations can be formulated for each decidable theory separately, but they can be interleaved for more efficient synthesis and more efficient synthesized programs. Our framework therefore contributes to the methodology for synthesis. We start from proof rules for a decision procedure, and extend them into transformation rule that can be viewed as a result of partially evaluating the execution of inference rules.

Contributions In summary, this paper makes the following contributions:

1. synthesis procedure for the theory of algebraic data types;
2. synthesis procedure for the theory of arrays;
3. a formalization of the above procedures in a framework supporting:
 - (a) proving correctness of synthesis steps, and
 - (b) combining synthesis procedures in a sound way.

We start by introducing our framework and illustrating it with a simple synthesis procedure for Presburger arithmetic. We then present the synthesis procedures for algebraic data types and for arrays.

2 Synthesis using Relation Transformations

A *synthesis problem* is a triple

$$\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket$$

where \bar{a} is a set of *input* variables, \bar{x} is a set of *output* variables and ϕ is a formula whose free variables are a subset of $\bar{a} \cup \bar{x}$. A synthesis problem denotes a binary relation $\{(\bar{a}, \bar{x}) \mid \phi\}$ between inputs and outputs. The goal of synthesis is to transform such relations until they become executable programs. Programs correspond to formulas of the form $P \wedge (\bar{x} = \bar{T})$ where $\text{vars}(P) \cup \text{vars}(\bar{T}) \subseteq \bar{a}$. We denote programs

$$\langle P \mid \bar{T} \rangle$$

We call the formula P a *precondition* and call the term \bar{T} a *program term*.

We use \vdash to denote the transformation on synthesis problems, so

$$\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \llbracket \bar{a} \langle \phi' \rangle \bar{x} \rrbracket \tag{1}$$

means that the problem $\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket$ can be transformed into the problem $\llbracket \bar{a} \langle \phi' \rangle \bar{x} \rrbracket$. The variables on the right-hand side are always the same as on the left-hand side.

Our goal is to compute, given \bar{a} , one value of \bar{x} that satisfies ϕ . We therefore define the soundness of (1) as a process that refines the binary relation given by ϕ into a smaller relation given by ϕ' , without reducing its domain. Expressed in terms of formulas, the conditions become the following:

$$\begin{array}{ll} \phi' \models \phi & \text{refinement} \\ \exists \bar{x}. \phi \models \exists \bar{x}. \phi' & \text{domain preservation} \end{array}$$

In other words, \vdash denotes domain-preserving refinement of relations. Note that the dual entailment $\exists \bar{x}. \phi' \models \exists \bar{x}. \phi$ also holds, but it follows from *refinement*. Note as well that \vdash is transitive.

Equivalences in the theory of interest immediately yield useful transformation rules: if ϕ and ϕ' are equivalent, (1) is sound. We can express fact as the following *inference rule*:

$$\frac{\models \phi_1 \leftrightarrow \phi_2}{\llbracket \bar{a} \langle \phi_1 \rangle \bar{x} \rrbracket \vdash \llbracket \bar{a} \langle \phi_2 \rangle \bar{x} \rrbracket} \quad (2)$$

In most cases we will consider transformations whose result is a program:

$$\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle$$

The correctness of such transformations reduces to

$$\begin{array}{ll} P \models \phi[\bar{x} \mapsto \bar{T}] & \text{refinement} \\ \exists \bar{x}. \phi \models P & \text{domain preservation} \end{array}$$

A *synthesis procedure* for a theory \mathcal{T} is given by a set of inference rules and a strategy for applying them such that every formula in the theory is transformed into a program.

2.1 Theory-Independent Inference Rules

We next introduce inference rules for a logic with equality. These rules are generally useful and are not restricted to a particular theory.

Equivalence. From the transitivity of \vdash and (2), we can derive a rule for synthesizing programs from equivalent predicates.

$$\frac{\llbracket \bar{a} \langle \phi_1 \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad \models \phi_1 \leftrightarrow \phi_2}{\llbracket \bar{a} \langle \phi_2 \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle}$$

Ground. In the case where no input variables are present, a synthesis problem is simply a satisfiability problem.

$$\frac{\mathcal{M} \models \phi}{\llbracket \emptyset \langle \phi \rangle \bar{x} \rrbracket \vdash \langle \top \mid \mathcal{M} \rangle} \quad \frac{\neg \exists \mathcal{M}. \mathcal{M} \models \phi}{\llbracket \emptyset \langle \phi \rangle \bar{x} \rrbracket \vdash \langle \perp \mid \perp \rangle}$$

(In these rules \mathcal{M} is a *model* for ϕ and should be thought of as a tuple of ground terms.) Note that the second rule can be generalized: even in the presence of input variables, if the synthesis predicate ϕ is unsatisfiable, then the generated program must be $\langle \perp \mid \perp \rangle$.

Assertions. Parts of a formula that only refer to input variables are essentially assertions and can be moved to the precondition.

$$\frac{\llbracket \bar{a} \langle \phi_1 \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad \text{vars}(\phi_2) \subseteq \bar{a}}{\llbracket \bar{a} \langle \phi_1 \wedge \phi_2 \rangle \bar{x} \rrbracket \vdash \langle \phi_2 \wedge P \mid \bar{T} \rangle}$$

Case Split. A top-level disjunction in the formula can be handled by deriving programs for both disjuncts and combining them with an if-then-else structure.

$$\frac{\llbracket \bar{a} \langle \phi_1 \rangle \bar{x} \rrbracket \vdash \langle P_1 \mid \bar{T}_1 \rangle \quad \llbracket \bar{a} \langle \phi_2 \rangle \bar{x} \rrbracket \vdash \langle P_2 \mid \bar{T}_2 \rangle}{\llbracket \bar{a} \langle \phi_1 \vee \phi_2 \rangle \bar{x} \rrbracket \vdash \langle P_1 \vee P_2 \mid \text{if}(P_1) \{ \bar{T}_1 \} \text{ else } \{ \bar{T}_2 \} \rangle}$$

Unused Input. Input variables that are not used in the constraint can be discarded.

$$\frac{\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad a_0 \notin \text{vars}(\phi)}{\llbracket a_0; \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle}$$

Unconstrained Output. Output variables that are not constrained by ϕ can be assigned any value.

$$\frac{\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad x_0 \notin \text{vars}(\phi)}{\llbracket \bar{a} \langle \phi \rangle x_0; \bar{x} \rrbracket \vdash \langle P \mid \text{any}; \bar{T}(\bar{a}) \rangle}$$

In the program, **any** designates a nullary function that returns any value of the appropriate type.

One-point. Whenever the value of an output variable is uniquely determined by an equality atom, it can be eliminated by a simple substitution.

$$\frac{\llbracket \bar{a} \langle \phi[x_0 \mapsto t] \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad x_0 \notin \text{vars}(t)}{\llbracket \bar{a} \langle x_0 = t \wedge \phi \rangle x_0; \bar{x} \rrbracket \vdash \langle P \mid \text{let } \bar{x} := \bar{T} \text{ in } t; \bar{x} \rangle}$$

Definition. The definition rule is in a sense dual to **One-point**, and is very convenient to give a name to a subterm appearing in a formula. Typical applications include purification, and unnesting of terms.

$$\frac{\llbracket \bar{a} \langle x_0 = t \wedge \phi[t \mapsto x_0] \rangle x_0; \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad x_0 \notin \text{vars}(t)}{\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \text{let } (x_0; \bar{x}) := \bar{T} \text{ in } \bar{x} \rangle}$$

Sequencing. The sequencing rule allows us to synthesize values for two groups of variables one after another. It fixes the values of some of the output variables, treating them temporarily as inputs, and then continues with the synthesis of the remaining ones.

$$\frac{\llbracket \bar{a}; \bar{x} \langle \phi \rangle \bar{y} \rrbracket \vdash \langle P_1 \mid \bar{T}_1 \rangle \quad \llbracket \bar{a} \langle P_1 \rangle \bar{x} \rrbracket \vdash \langle P_2 \mid \bar{T}_2 \rangle}{\llbracket \bar{a} \langle \phi \rangle \bar{x}; \bar{y} \rrbracket \vdash \langle P_2 \mid \text{let } \bar{x} := \bar{T}_2 \text{ in } \bar{x}; \bar{T}_1 \rangle}$$

Static Computation. A basic rule is to perform computational steps when possible.

$$\frac{\llbracket a_0; \bar{a} \langle F[t \mapsto a_0] \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad \text{vars}(t) \subseteq \bar{a}}{\llbracket \bar{a} \langle F \rangle \bar{x} \rrbracket \vdash \langle P \mid \text{let } a_0 := t \text{ in } \bar{T} \rangle}$$

An \bar{a} -input term is a term that is not a variable but can be directly evaluated given \bar{a} : its variables are among \bar{a} , and all of its function symbols are known computable functions.

Variable Transformation. The \vdash transformation preserves the variables. To show how we can change the set of variables soundly, we next present in our framework *variable transformation* by a computable function ρ [8], as an inference rule on two \vdash transformations.

$$\frac{\llbracket \bar{a} \langle \phi[\bar{x} \mapsto \rho(\bar{x}')] \rangle \bar{x}' \rrbracket \vdash \langle P \mid \bar{T} \rangle}{\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \rho(\bar{T}) \rangle}$$

Slightly more generally, we have the following.

$$\frac{\llbracket \bar{a} \langle \phi' \rangle \bar{x}' \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad \exists \bar{x}. \phi \models \exists \bar{x}'. \phi' \quad \phi' \models \phi[\bar{x} \mapsto \rho(\bar{x}')]}{\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \rho(\bar{T}) \rangle}$$

Existential Projection. A special case of variable transformation, where ρ simply projects out some of the variables.

$$\frac{\llbracket \bar{a} \langle \phi \rangle \bar{x}; \bar{x}' \rrbracket \vdash \langle P \mid \bar{T} \rangle}{\llbracket \bar{a} \langle \exists \bar{x}'. \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \text{let } (\bar{x}; \bar{x}') := \bar{T} \text{ in } \bar{x} \rangle}$$

3 Synthesis for Presburger Arithmetic

This section summarizes a simple version of a synthesis procedure for Presburger arithmetic using our current synthesis rules. Our goal is to give a complete procedure that is easy to prove correct, as opposed to one that generates efficient code. The reader will observe that our description reads like a description of quantifier elimination. Note, however, that the inference rules that we refer to are from the previous section and therefore also specify how to compute the corresponding program.

Unlike the procedure in [9] the procedure below does not perform efficient solving of equations, but could be refined to do so by adopting the description in [9] to our inference rules.

As in the preprocessing steps for simple quantifier elimination for Presburger arithmetic, the equivalences we use as rules include replacing $t_1 \neq t_2$ with $t_1 < t_2 \vee t_2 < t_1$. In principle, we can rewrite $t_1 = t_2$ into $t_1 \leq t_2 \wedge t_2 \leq t_1$ (see [9] for more efficient approaches). We rewrite $t_1 \leq t_2$ into $t_1 < t_2 + 1$. When needed we assume that we perform disjunctive splitting rule to obtain only a conjunction of literals. We also assume that we apply the sequencing rule to fix the remaining variables and only consider one output variable x . Thanks to the assertion rule, we assume that all literals contain x .

A rule that takes into account divisibility is the following:

$$\frac{\llbracket \bar{a} \langle \phi[kx \mapsto y] \wedge y \equiv_k 0 \rangle y \rrbracket \vdash \langle P \mid T \rangle \quad k \neq 0 \quad x \text{ in } \phi \text{ only as } kx}{\llbracket \bar{a} \langle \phi \rangle x \rrbracket \vdash \langle P \mid T/k \rangle}$$

The rule is a case of **Variable Transformation** with $\rho(y) = y/k$.

To enable the previous rule, we can ensure that all occurrences of a variable have the same coefficient by multiplying constraints by a positive constant (e.g., the least common multiple of all coefficients). These transformations are based on using (in a context) equivalences between $t_1 \bowtie t_2$ and $kt_1 \bowtie kt_2$, for $k > 0$ and $\bowtie \in \{=, <, >, \equiv_p\}$.

Using the rules so far, we can ensure that the coefficients of an output variable have a coefficient one. If such a variable occurs in an equality, we can eliminate it using the one-point inference rule. If the variable occurs only in inequalities, we perform the main step of the procedure.

Elimination of Inequalities Based on the discussion above, we can assume that the formula ϕ in the synthesis problem is of the form

$$\bigwedge_{i=1}^L a_i < x \wedge \bigwedge_{j=1}^U x < b_j \wedge \bigwedge_{i=1}^D x + t_i \equiv_{K_i} 0$$

We aim to replace ϕ with ϕ' such that

$$\llbracket \bar{a} \langle \phi \rangle x \rrbracket \vdash \llbracket \bar{a} \langle \phi' \rangle x \rrbracket \quad (3)$$

We define ϕ' as

$$\bigvee_{j=1}^L \bigvee_{i=1}^N (\phi \wedge x = a_j + i)$$

where N is the least common multiple of K_1, \dots, K_D . Clearly ϕ' is stronger than ϕ because each disjunct is stronger than ϕ , so it remains to argue about domain preservation. Suppose there exists values for x so that ϕ holds. Let a_j be the largest value among the values of lower bounds a_j and let I be such that $a_j + I \equiv_N x$ holds. Then letting x to be $a_j + I$ makes ϕ' true as well.

After performing disjunctive splitting, x can be eliminated using one-point rule. The correctness follows by (3) and the correctness of the one-point rule. The cases where some of the lower bounds do not exist can be treated similarly.

This completes the overview of synthesis of functions given by Presburger arithmetic relations.

Enumerating Solutions In addition to finding one solution \bar{x} such that ϕ holds, it is useful to be able to find all solutions, when this set is finite. When solving constraints at run time, a simple way to find all solutions is to maintain a list of previously found solutions $\bar{v}_1, \dots, \bar{v}_n$ for \bar{x} and add to ϕ an additional conjunct $\bigwedge_{i=1}^n \bar{x} \neq \bar{v}_i$, see [7].

One possible approach to compile this process is to enrich Presburger arithmetic with finite uninterpreted relations as parameters. This enables a synthesis procedure to, for example, take the set of previous solutions as the input. If R is such a finite-relation symbol or arity n and \bar{x} are n variables, we introduce an additional literal $\bar{x} \notin R$ into the logic, with the intention that R stores the previously found solutions. The elimination of inequalities in the presence of then produces terms that avoid the elements of R by considering not only the value $a_j + i$ for x but enumerating a larger number of solutions, $a_j + i + kN$, for multiple values of $k \geq 0$. Because R is known only at run time, the generated code contains a loop that increases $k \geq 0$ to allow x to leave the range of the corresponding coordinate of R . The value of k is bounded at run time by, for example, $\lceil (\max(R_x) - \min(R_x)) / N \rceil + 1$ where R_x is the projection of R onto the coordinate at which x appears in the literal $\bar{x} \notin R$. Code containing generated loops is not unusual; it also arises when synthesis for Presburger arithmetic is generalized to allow parametric coefficients K_i , see [9]. The generated loop is guaranteed to terminate.

4 Synthesis for Term Algebras

We present in this section a synthesis procedure for formulas over term algebras. We start by assuming a pure term algebra, and later extend the system to algebras with elements from foreign theories. In both cases, we present a series of normal forms and of inference rules, and argue that together they can be used together to form a synthesis procedure.

4.1 Pure Term Algebras

The grammar of atoms over our term algebra is given by the following two production rules, where c and F denote a constant and a function symbol from the algebraic signature, respectively:

$$\begin{aligned} A &::= T = T \mid T \neq T \mid \text{is}_c(T) \mid \text{is}_F(T) \\ T &::= x \mid c \mid F(\bar{T}) \mid F_i(T) \end{aligned}$$

In the following we assume that the algebra defines at least one constant and one non-nullary constructor function. Formulas are built from atoms with the usual propositional connectives. We use an extension of the standard theory of term algebras. The extension defines additional unary *tester* functions $\text{is}_c(\cdot)$ and $\text{is}_F(\cdot)$ for constant and functions in the algebraic signature respectively, and unary *selector* functions $F_i(\cdot)$, with $1 \leq i \leq n$ where n is the arity of F . These extra symbols form a definitional extension [5] given by the axioms:

$$\forall x. \text{is}_c(x) \leftrightarrow x = c \quad (4)$$

$$\forall x. \text{is}_F(x) \leftrightarrow \exists \bar{y}. x = F(\bar{y}) \quad (5)$$

$$\begin{aligned} \forall x, y. F_i(x) = y \leftrightarrow (\exists \bar{y}. y = \bar{y}[i] \wedge F(\bar{y}) = x) \\ \vee \neg(\exists \bar{y}. F(\bar{y}) = x) \wedge x = y \end{aligned} \quad (6)$$

Note that the case analysis in (6) is required only to make the selector functions total. In practice, we are only interested in cases where the selectors are applied to arguments of the proper type. We will therefore assume in the following that each selector application $F_i(x)$ is accompanied with a side condition $\text{is}_F(x)$.

Rewriting of tester and selector functions. By applying the axioms (4) and (5), we can rewrite all applications of a tester function into an existentially quantified equality over terms. We can similarly eliminating applications of testers by existentially quantifying over the arguments of the corresponding constructors. Using the **Existential Projection** rule, we can in turn consider the obtained synthesis problem as a quantifier-free one.

Elimination through unification. We can at any point apply *unification* to a set of equalities over terms. Unification rewrites a conjunction of term equations into either \perp , if the equations contain a cycle or an equality involving incompatible constructors, or into an equivalent conjunction of atoms $\bigwedge_i v_i = t_i$, where v_i is a variable and t_i is a term. This set of equations has the additional property that

$$\left(\bigcup_i \{v_i\} \right) \cap \left(\bigcup_i \text{vars}(t_i) \right) = \emptyset$$

In other words, it defines a set of variables as a set of terms built from a distinct set of variables [2]. This form is particularly suitable for applications of the **One-point** rule: indeed, whenever v_i is an output variable, we can apply it, knowing that v_i does not appear in t_i (or in any other equation).

Dual view. Unification allows us to eliminate output variables that are to the left of an equality. When instead an input variable appears in such position, we can resort to a dual form to eliminate output variables appearing in the right-hand side. We obtain the dual form by applying as much as possible the following two rules to term equalities:

$$\frac{t = c}{\text{is}_c(t)} \qquad \frac{t = F(t_1, \dots, t_n)}{F_1(t) = t_1 \wedge \dots \wedge F_n(t) = t_n \wedge \text{is}_F(t)}$$

Note that these are rewrite rules for formulas. Because they preserve the set of variables and equisatisfiability, they can be lifted to inference rules for programs using the **Equivalence** rule. Observe that at saturation, the generated atoms are of two kinds: 1) applications of tester predicates and 2) equalities between two terms, each containing at most one variable. In particular, all equalities between an output variable and a term are amenable to applications of the **One-point** rule.

Disequalities. Finally, we introduce a dedicated rule for the treatment of disequalities between terms. The rule is defined for disequalities over variables and constants in *conjunctive normal form* (CNF). From a conjunction of disequalities over terms, we can obtain this normal form by applying the following rewrite rules until saturation:

$$\frac{F(\bar{T}_1) \neq F(\bar{T}_2) \quad F \neq G}{\top} \qquad \frac{F(\bar{t}_1) \neq F(\bar{t}_2)}{t_1^1 \neq t_n^1 \vee \dots \vee t_1^2 \neq t_n^2}$$

Intuitively, the first rule captures the fact that terms built with distinct constructors are trivially distinct (note that this also captures distinct constants, which are nullary constructors). The second rule breaks down a disequality into a disjunction of disequalities over subterms.

To obtain witness terms from the normal form, it suffices to satisfy one disequality in each conjunct. We achieve this by eliminating one variable after another, applying for each a diagonalization principle. In the following rule ϕ_{CNF} denotes a CNF formula over atomic disequalities.

$$\frac{\llbracket \bar{a} \langle \phi_{\text{CNF}} \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T} \rangle \quad x_0 \notin \phi_{\text{CNF}}}{\llbracket \bar{a} \left\langle \begin{array}{c} (x_0 \neq t_1 \vee \dots) \\ \wedge \\ \dots \\ \wedge (x_0 \neq t_n \vee \dots) \\ \wedge \\ \phi_{\text{CNF}} \end{array} \right\rangle x_0; \bar{x} \rrbracket \vdash \langle P \mid \text{distinct}(t_1, \dots, t_n); \bar{T}(\bar{a}) \rangle}$$

In the generated program, **distinct** denotes an n-ary computable function that returns, at run time, a term distinct from all its arguments. Such a value is guaranteed to exist, since the term algebra is assumed to have at least one constructor (and it can be computed in time polynomial in the number of arguments).

Synthesis Procedure We now argue that the reductions to normal forms and the rules presented above are sufficient to form a complete synthesis procedure for a given pure term algebra. The procedure is given by the following steps:

1. Reduce an arbitrary propositional structure to a conjunction through applications of the **Case Split** rule.
2. Remove selectors and testers through rewrites and applications of **Existential Projection**.

3. Apply unification to all equalities, then apply **One-point** as often as possible. As a result, the only equalities remaining have the form $a = t$, where a is an input variable and $a \notin \text{vars}(t)$.
4. Rewrite into dual form, then apply **One-point** as much as possible. After applying **Assertions**, the problem is reduced to a conjunction of disequalities, each involving at least one output variable.
5. Transform the conjunction into CNF and eliminate all remaining variables by successive applications of the diagonalization rule.

We note that (for conjunctions of literals), the generated program always run in time polynomial in the size of the input terms, as they consist of a sequence of assignments, one for each output variable, and each term has polynomial size.

4.2 Reduction to an Interpreted Theory

We now consider the case of a term algebra defined over an interpreted theory \mathcal{T} . A canonical example is the algebra of integer lists, where \mathcal{T} is the theory of integers, and defined by the constant $\text{Nil} : \text{List}$ and the constructor $\text{Cons} : \mathbb{Z} \times \text{List} \rightarrow \text{List}$. In this theory, the selector function $\text{Cons}_1(\cdot)$, for instance, is of type $\text{List} \rightarrow \mathbb{Z}$. We show how to reduce a synthesis problem in the combination of theories to a synthesis problem in \mathcal{T} . We focus on the important differences with the previous case.

Purification. We can assume without loss of generality that constructor terms contain no subterms from \mathcal{T} other than variables. Indeed one can always apply the **Definition** rule to separate such terms.

Unification. Applying unification can result in derived equalities between variables of \mathcal{T} . These should simply be preserved in the reduced problem.

Dual view. Applying the rewriting into the dual view can result in derived equalities of the form $x = t$, where x is a variable from \mathcal{T} and t is an application of selectors to an input variable. Because \mathcal{T} cannot handle these selectors, we need to rewrite t into a simple variable. By using the **Definition** and **Sequencing** rules, we make this variable an input of the problem in \mathcal{T} .

Disequalities. Contrary to the pure case, we cannot always eliminate all conjuncts in the CNF by applying a diagonalization; we can eliminate variables that belong to the term algebra, but not variables of \mathcal{T} . Instead, for each equality $v = w$ over \mathcal{T} in the CNF, we introduce at the top-level a disjunction $v = w \vee v \neq w$, and apply the **Case Split** rule to encode a guess. This in essence compiles the guessing of the partitioning of shared variables that is traditionally introduced in a Nelson-Oppen setting [15]. Because **Case Split** preserves the relation entirely, this is a sound and complete reduction step.

Once all the disequalities have been handled, either through diagonalization if they are over algebraic terms, or by case-splitting if they are over \mathcal{T} variables, we have entirely reduced the synthesis problem into a synthesis problem for \mathcal{T} .

5 Synthesis for Arrays

In this section, we introduce a synthesis procedure for arrays. In contrast to other theories for which synthesis procedures have been introduced thus far, the standard array theory with unbounded arrays does not allow for quantifier elimination. Therefore, for a complete compile-time approach we need to restrict our inputs to arrays that are bounded in size at compile time. With a known finite bound on array size, the procedure reduces the synthesis problem to synthesis problems over indices and elements, in the same way as a satisfiability procedure for arrays reduces the satisfiability problem to these component theories.

If we do not know size bounds at compile time, we need to employ a mixed approach, which postpones some of the reasoning to run time. The reduction is the same as before, but now the component synthesis procedures not only return one solution of the synthesis problem, but instead an iterator over all possible solutions (given the limited knowledge about the inputs). Then, at runtime we need to go through these solutions and, if possible, find one that matches the array inputs. We will explain the latter approach in the following, and add some remarks on statically bounded arrays after that.

5.1 Preliminaries

We consider synthesis for a basic array theory, with the usual `read` and `write` operations. The theory is parameterized by an index theory \mathcal{T}_I and an element theory \mathcal{T}_E . We assume that our input formula ϕ is a conjunction of literals, and that we have synthesis procedures for these theories. For \mathcal{T}_I , we will also use the `find_all` procedure, which returns an iterator over all possible solutions of a synthesis problem. Additionally, we assume that we have a predicate \approx_I between arrays, where I can be any set of variables or constants, and $a \approx_I b$ evaluates to true iff a and b are equal up to (the elements stored at) indices $i \in I$. In particular, this also subsumes extensionality of arrays (with $I = \emptyset$).

Array Bounds In the following, we assume that our specification ϕ contains for every array variable a two special variables a_l, a_u , standing for the lower and upper bound of the array. Additionally, we assume that ϕ contains for every index variable i that is used to read or write into a the constraints $a_l \leq i \wedge i \leq a_u$. These constraints ensure that synthesized programs do not contain out-of-bounds array accesses. If a is an array parameter, then a_l, a_u are additional parameters.³

5.2 A Reduction-based Synthesis Procedure for Arrays

We introduce a synthesis procedure for arrays, based on a reduction to synthesis in the index theory \mathcal{T}_I and the element theory \mathcal{T}_E . Before we get to the

³ If we have $b = \text{write}(a, i, e)$ or $a \approx_I b$, then we assume that $a_l = b_l$ and $a_u = b_u$, i.e. we do not have to introduce multiple lower and upper bounds for “connected” arrays.

main synthesis procedure, we will first introduce a subprocedure `array_reduct`, that reduces a specification in the array theory to a specification in $\mathcal{T}_I \cup \mathcal{T}_E$, and a subprocedure `Impl_Eq`, that will be used to deduce implied equalities and disequalities, and is assumed to exist for \mathcal{T}_I and \mathcal{T}_E .

Array Reduction Procedure `array_reduct` takes as input a formula ϕ with input variables \bar{a} and output variables \bar{x} in the theory of arrays. It reduces ϕ to a formula $\phi_E \wedge \phi_I \wedge \text{Impl}$ in the combined theory $\mathcal{T}_I \cup \mathcal{T}_E$. Furthermore, it returns a set of definitions D for auxiliary variables introduced during the reduction, and classifies these auxiliary variables as in- or output variables:

1. substitute all array writes `write(a, i, e)` by fresh array variable b , add $b[i] = e \wedge a \approx_{\{i\}} b$ to ϕ , and the definition $b = \text{write}(a, i, e)$ to D ;
2. let Eq_A be the set of all literals $a \approx_I b$ in the transitive closure of \approx wrt. ϕ (if $a \approx_I b$ and $b \approx_J c$, then $a \approx_{I \cup J} c$); remove all literals $a \approx_I b$ from ϕ
3. for every array read $a[i]$ in ϕ and predicate $a \approx_J b$ in Eq_A , add a formula $(\bigwedge_{j \in J} i \neq j) \rightarrow a[i] = b[i]$ to `Impl`
4. for all terms $a[i]$ in ϕ , substitute all occurrences of $a[i]$ by a fresh element variable a_i , add for every pair of array reads $a[i], a[j]$ a formula $i = j \rightarrow a_i = a_j$ to `Impl`, and the definition $a_i = a[i]$ to D
5. finally, partition ϕ into a pure \mathcal{T}_E -part ϕ_E and a pure \mathcal{T}_I -part ϕ_I

Obtain \bar{a}' from \bar{a} by adding all auxiliary variables that stand for reads into array variables in \bar{a} , and removing all array variables, and obtain \bar{x}' from \bar{x} by adding all other auxiliary variables and removing all array variables. Then

$$\text{array_reduct}(\bar{a}, \bar{x}, \phi) = (\bar{a}', \bar{x}', \phi_E, \phi_I, \text{Impl}, Eq_A, D).$$

In terms of our inference rules, this amounts to several `Equivalence` steps (all added implications are consequences of ϕ), adding variables by `Definition` and considering them as parameters via `Sequencing`. This justifies applying `array_reduct` as a transformation of synthesis problems in our framework.

Synthesis Procedure for Arrays

1. Let $(\bar{a}', \bar{x}', \phi_E, \phi_I, \text{Impl}, Eq_A, D) := \text{array_reduct}(\bar{a}, \bar{x}, \phi)$.
2. We use the following `Enumerator` rule to remove D from the formula we need to consider. \bar{T}^* should be an iterator over all possible solutions, or a specialised decision procedure allowing us to adjust the solution once the parameter value is known. We apply the rule to remove all array parameters, along with definitions D that mention them. After synthesis in element and index theory, we will generate code that tests against D for possible solutions from index and element synthesis:

$$\frac{\llbracket \bar{a} \langle \phi \rangle \bar{x} \rrbracket \vdash \langle P \mid \bar{T}^* \rangle \quad v \notin \text{vars}(\psi)}{\llbracket v; \bar{a} \langle \phi \wedge \psi \rangle \bar{x} \rrbracket \vdash \left\langle \begin{array}{l} P \wedge \exists i. \psi[\bar{x} \mapsto \bar{T}^*.next^{(i)}] \\ \bar{x} := \bar{T}^*(\bar{a}); \\ \text{while}(\neg\psi(v; \bar{a}, \bar{x})) \\ \text{if}(\bar{T}^*.hasNext) \{ \\ \bar{x} := \bar{T}^*.next(\bar{a}) \\ \} \\ \text{else} \{ \text{return UNSAT} \} \end{array} \right\rangle}$$

3. Let E_{q_I} be the set of all equalities $i = j$ s.t. either $i = j$ or $i \neq j$ appears in an implication in **Impl**. After using **Equivalence** to introduce a disjunction over all possible valuations of equalities in E_{q_I} and **Sequencing** to consider index variables as inputs for now, we use **Case Split** to branch synthesis of element variables for all these cases:

- (a) let V_E be the set of truth valuations of elements of E_{q_I} , each described by a conjunction of literals $v \in V_E$ (containing for every $l \in E_{q_I}$ either l or $\neg l$). For every $v \in V_E$, obtain a new formula ϕ_v by adding to ϕ_E the succedent of all implications in **Impl** for which the antecedent is in v .
- (b) **\mathcal{T}_E -synthesis**: for each $v \in V_E$, let $\llbracket \bar{a}' \langle \phi_v \rangle \bar{x}_E \rrbracket \vdash \langle P_v \mid \bar{T}_v \rangle$, where \bar{x}_E are element variables, \bar{x}_I index variables from \bar{x}' .
- (c) Then $P_E := \phi_I \wedge \bigvee_{v \in V_E} (v \wedge P_v)$, and

$$\begin{aligned} T_E &:= \text{if } (v_1 \wedge P_{v_1}) T_{v_1}(\bar{a}) \\ &\quad \text{else if } (v_2 \wedge P_{v_2}) T_{v_2}(\bar{a}) \\ &\quad \dots \\ &\quad \text{else if } (v_n \wedge P_{v_n}) T_{v_n}(\bar{a}). \end{aligned}$$

4. Having closed all branches, we turn to **\mathcal{T}_I -synthesis**:
Let $\llbracket \bar{a}' \langle \phi_v \rangle \bar{x}_I \rrbracket \vdash \langle P_I \mid \bar{T}_I^* \rangle$. That is, instead of synthesizing one solution for the synthesis problem, we obtain an enumerator for all possible solutions, as explained in Sect. 3.⁴
5. Coming to the other branch of **Enumerator**, we generate

$$\begin{aligned} T_I &= \text{let } \bar{x}_I = T_I^* \text{ in} \\ &\quad \text{while}(\neg(D(\bar{a}, T_E(\bar{a}, \bar{x}_I); \bar{x}_I))) \\ &\quad \quad \text{if } T_I^*.hasNext \text{ then } \bar{x}_I = T_I^*.next \text{ else return UNSAT;} \\ &\quad \text{return } T_I^* \end{aligned}$$

6. Given T_I and T_E , synthesis of array variables is simple: let $a \in \bar{x}$ be an array variable, and I the set of all index variables i for which $a_i = a[i]$ is in D . For a given v , let J be the maximal subset of I s.t. $\forall i, j \in J. v \models i \neq j$. By construction, there must be an array b s.t. $a \approx_I b$ is in E_{q_A} . If b is not a parameter array, all positions not explicitly defined can be defined arbitrarily. Then the witness term T^a for variable a is defined by:

$$T^a(\bar{a}) := \text{write}(\dots (\text{write}(b, T_I^{j_1}(\bar{a}), T_E^{a_{j_1}}(\bar{a})) \dots), T_I^{j_n}(\bar{a}), T_E^{a_{j_n}}(\bar{a}))$$

⁴ If arrays (or all array indices in ϕ) are bounded, then all solutions can be computed statically. Otherwise, there may be infinitely many solutions and we need to be able to compute them during runtime.

where the T_I^j are witness terms for index variables $j \in J$, and the $T_E^{a_j}$ witness terms for the corresponding element variables. Altogether, we obtain (for array variables a^1, \dots, a^n in \bar{x})

$$T_A := T^{a^1}; \dots; T^{a^n}$$

This is the an array-specific step, the correctness of which is obvious from the correctness of array decision procedures using the same reduction.

7. Finally, let P_E^- be $\bigvee_{v \in V_E} (v[\bar{x}_I \mapsto T_I] \wedge P_v)$. Then

$$\begin{aligned} P &:= P_I \wedge P_E^- \wedge \exists i. D[\bar{x}' \mapsto (T_E(\bar{a}, T_I^*.\text{next}^{(i)}); T_I^*.\text{next}^{(i)})] \\ T &:= (T_E; T_I; T_A) \end{aligned}$$

Remarks For efficiency, it may be useful to deduce, both in \mathcal{T}_I and \mathcal{T}_E , equalities that are implied by ϕ at any time, and replacing clauses in `Impl` by their succedents if the antecedent is implied by ϕ . This will avoid unnecessary branching, speeding up synthesis and removing dead branches from the resulting code.

Theory Combination $\mathcal{T}_E \cup \mathcal{T}_I$ The reduction above assumes that theories \mathcal{T}_E and \mathcal{T}_I are *strongly disjoint*, i.e. they share not even the equality symbol. Alternatively, we can make the restriction that variables that are used for array reads may never be compared to variables that are used as elements. In this case, implications from congruence of array reads is the only connection between the theories, and \mathcal{T}_I -synthesis can run completely independent of \mathcal{T}_E -synthesis, provided the latter accounts for all possible cases of \mathcal{T}_I -equalities. If the theories are not strongly disjoint, we really need a synthesis procedure for the combined theory. In this case, we directly use the combined decision procedure to produce an iterator over all possible solutions in both element and index theory.

Statically Bounded Arrays If all arrays in ϕ are statically bounded, i.e. values of upper and lower bounds are known or can be computed at compile time, then we can statically compute all solutions for constraints in \mathcal{T}_I that are within array bounds. Thus, we will not have to integrate an iterator that computes additional solutions into our synthesized program, and can give a bound on the maximum number of traversals of the loop at compile time.

5.3 Complexity of Synthesis and Synthesized Code

Complexity of the array synthesis procedure is dominated by the branching on equalities of index variables: we may need exponentially (in the number of index variables) many calls to the synthesis procedure for \mathcal{T}_E . The array reduction itself is polynomial.

Correspondingly, size of the synthesized code is also exponential in the number of index variables, as it will contain branches for all these possibilities. Although only one of these branches will be explored at runtime, the worst-case

running time of the synthesized code will still be exponential in the number of index variables: for a size bound n on a given array, there may be n^i many solutions to the constraints in \mathcal{T}_I . In the worst case, the condition of the while-loop needs to be checked for all of these solutions.

5.4 Example of Array Synthesis

Suppose we want to synthesize a most general precondition P and program code s.t. for any input array a and bounds a_l, a_u that satisfy P , the synthesized code computes values for an array b and integer variables i, j, k such that the following is satisfied:⁵

$$\begin{aligned} \phi \equiv & \quad a_l = 0 \wedge i > a_l \wedge i = j + j \wedge i < a_u \wedge k \geq a_l \wedge k < i \\ & \quad \wedge a[i] > 0 \wedge a[k] \leq 0 \wedge b[i] > a[i-2] \wedge b[k] = a[i] \\ & \quad \wedge a' = \text{write}(a, i, e_1) \wedge b = \text{write}(a', k, e_2). \end{aligned}$$

Array reduction:

1. Let $D := \{a' = \text{write}(a, i, e_1), b = \text{write}(a', k, e_2)\}$, replace these writes in ϕ by

$$a'[i] = e_1 \wedge a' \approx_i a \wedge b[k] = e_2 \wedge b \approx_k a'.$$

2. From $a' \approx_i a \wedge b \approx_k a'$, conclude $a \approx_{\{i,k\}} b$. Thus,

$$Eq_A := \{a' \approx_i a, b \approx_k a', a \approx_{\{i,k\}} b\}.$$

3. From Eq_A we obtain (leaving out implications with false antecedents)

$$\text{Impl} = \left\{ \begin{array}{l} k \neq i \rightarrow a[k] = a'[k], \quad i-2 \neq i \rightarrow a[i-2] = a'[i-2], \\ i \neq k \rightarrow a'[i] = b[i], \quad i-2 \neq k \rightarrow a'[i-2] = b[i-2], \\ i-2 \neq i \wedge i-2 \neq k \rightarrow a[i-2] = b[i-2] \end{array} \right\}.$$

4. Replacing array reads by fresh variables, we add to Impl the set

$$\left\{ \begin{array}{l} i = k \rightarrow a_i = a_k, \quad i-2 = k \rightarrow a_{i-2} = a_k, \\ i = k \rightarrow a'_i = a'_k, \quad i-2 = k \rightarrow a'_{i-2} = a'_k, \\ i = k \rightarrow b_i = b_k, \quad i-2 = k \rightarrow b_{i-2} = b_k \end{array} \right\},$$

and to D the set

$$\left\{ \begin{array}{l} a_i = a[i], a_{i-2} = a[i-2], a_k = a[k], \\ a'_i = a'[i], a'_{i-2} = a'[i-2], a'_k = a'[k], \\ b_i = b[i], b_{i-2} = b[i-2], b_k = b[k] \end{array} \right\},$$

⁵ Note that the last two literals imply $a \approx_{\{i,k\}} b$, which in turn implies that there exist valuations for a', e_1, e_2 satisfying these literals. Thus, we can allow statements of the form $a \approx_I b$ in specifications, and replace them with a number of `write` definitions according to the size of I , with fresh element and array variables in every `write`.

and after replacement obtain a formula in $\mathcal{T}_I \cup \mathcal{T}_E$:

$$\begin{aligned} & i > a_l \wedge i = j + j \wedge i < a_u \wedge k \geq a_l \wedge k < i \\ & \wedge a_i > 0 \wedge a_k \leq 0 \wedge b_i > a_{i-2} \wedge b_k = a_i \\ & \wedge a'_i = e_1 \wedge b_k = e_2 \end{aligned}$$

5. Separating index and element constraints, we obtain

$$\begin{aligned} \phi_I & := i > a_l \wedge i = j + j \wedge i < a_u \wedge k \geq a_l \wedge k < i \\ \phi_E & := a_i > 0 \wedge a_k \leq 0 \wedge b_i > a_{i-2} \wedge b_k = a_i \wedge a'_i = e_1 \wedge b_k = e_2 \\ \bar{a}' & := (n, a_i, a_k, a_{i-2}) \\ \bar{x}' & := (i, j, k, a'_i, a'_k, a'_{i-2}, b_i, b_k, b_{i-2}, e_1, e_2) \end{aligned}$$

Implied equalities and disequalities: From ϕ_I we can conclude that $i \neq j, k \neq i$ and $i - 2 \neq i$. From ϕ_E we conclude that $a_i \neq a_k, b_i \neq a_{i-2}$.

Propagating equalities: $k \neq i$ implies $a_k = a'_k$ and $a'_i = b_i$. $i - 2 \neq i$ implies $a_{i-2} = a'_{i-2}$. In the opposite direction, $a_i \neq a_k$ implies $i \neq k$ (which we already knew). We get $Eq_E = \{a_i \neq a_k, b_i \neq a_{i-2}, a'_i = b_i, a_{i-2} = a'_{i-2}\}$ and $Eq_I = \{i \neq k, i - 2 \neq i\}$.

Branching: From the three equations that appear in antecedents of Impl , valuations for two are in Eq_I . Thus, we only branch on the valuation of $i - 2 = k$. Let $v_1 \equiv i - 2 = k$, which implies $a_{i-2} = a_k, a'_{i-2} = a'_k$ and $b_{i-2} = b_k$. Let $v_2 \equiv i - 2 \neq k$, which implies $a'_{i-2} = b_{i-2}$ and (together with $i - 2 \neq i$) $a_{i-2} = b_{i-2}$.

\mathcal{T}_E -synthesis: Assuming v_1 , we obtain the following valuations for variables \bar{x}_E :

$$e_2 := b_k := b_{i-2} := a_i, e_1 := b_i := a'_i := a_{i-2} + 1, a'_k := a'_{i-2} := a_k.$$

Assuming v_2 , valuations are the same except for $b_{i-2} := a'_{i-2} := a_{i-2}$. The synthesized precondition is in both cases $P_E \equiv a_i > 0 \wedge a_k \leq 0$.

\mathcal{T}_i -synthesis: We obtain $j := \lfloor \frac{i}{2} \rfloor$ and an iterator \bar{T}^* of solutions for (i, k) :

```

T* := (0,2)
T*.next = let (k,i) = T* in
  if(k+1<i) (k+1,i)
  else if(i+2<a_u) (0,i+2)
  else return UNSAT

```

along with a precondition $P_I \equiv a_u > 2$.

Array synthesis: Lifting the witness terms for elements to array b , we obtain

$$\begin{aligned} b & = \text{if}(i - 2 = k) \\ & \quad \text{write}(\text{write}(a, i, a[T_I^i - 2] + 1), T_I^k, a[T_I^i]) \\ & \quad \text{else write}(\text{write}(\text{write}(a, i, a[T_I^i - 2] + 1), T_I^i - 2, a[T_I^i]), T_I^k, a[T_I^i]) \end{aligned}$$

Result: Finally, we obtain the precondition

$$a_u > 2 \wedge \exists n. (i, k) = T_I^*.next^{(n)} \rightarrow a[T_I^i] > 0 \wedge a[T_I^k] \leq 0$$

and the program⁶ in Fig. 1 for computing i, j, k and b from a and a_u .

```

if( $a_u > 2$ ) {
  ( $i, k$ ) :=  $T_I^*$  in
  while ( $\text{not}(a[i] > 0 \ \&\& \ a[k] \leq 0)$ )
    if ( $T_I^*.hasNext$ )
      ( $i, k$ ) :=  $T_I^*.next$ 
    else throw new Exception("Unsatisfiable constraint.")
  let  $j = i / 2$  in
  if ( $i - 2 = k$ ) {
     $bi := a[i - 2] + 1$ 
     $bk := a[i]$ 
  } else {
     $bi := a[i - 2] + 1$ 
     $bk := a[i]$ 
     $bi2 := a[i - 2]$ 
  }
  if ( $i - 2 = k$ ) {  $b := \text{write}(\text{write}(\text{write}(a, i, bi), k, bk))$ 
  } else {  $b := \text{write}(\text{write}(\text{write}(a, i, bi), k, bk), i - 2, bi2)$ 
  }
  ( $b, i, j, k$ )
} else {
throw new Exception("Unsatisfiable constraint.")
}

```

Fig. 1. Example of code generated by array synthesis procedure

5.5 Example: Inverting Program Fragments

The synthesis procedure for arrays can also be used to invert given code fragments, e.g. for automatically obtaining a program that reverts (some of) the changes a given piece of code did to some data. Consider the following code fragment:

```

if( $a[i] == 0$ )
   $a[i] := a[i + 1]$ 
else if ( $a[i] == 1$ )
   $a[i] := a[0]$ 
else if ( $a[i] > 1$ )

```

⁶ The code can be significantly simplified by merging parts that are not affected by case distinctions.

```

    a[i]:=a[i]-1
  else a[i]:= a[i]+2

```

translates into

$$\begin{aligned}
 & (a_0[i] = 0 \wedge a_1 = \text{write}(a_0, i, a_0[i + 1])) \\
 & \vee (a_0[i] = 1 \wedge a_1 = \text{write}(a_0, i, a_0[0])) \\
 & \vee (a_0[i] > 1 \wedge a_1 = \text{write}(a_0, i, a_0[i - 1])) \\
 & \vee (a_0[i] < 0 \wedge a_1 = \text{write}(a_0, i, a_0[i + 2])),
 \end{aligned}$$

where a_0 refers to the pre-, and a_1 to the post-state value of array a . For synthesizing the inverted code, we assume that a_1 is the input, and a_0 the output. The synthesis procedure will return a piece of code

```

  if(a[i]==a[i+1])
    a[i]:=0
  else if(a[i]==a[0])
    a[i]:=1
  else if(a[i]>0)
    a[i]:=a[i]+1
  else a[i]:=a[i]-2

```

Since the relation given by the input code does not model a bijection, applying the inverted code after the input code will not result in exactly the same state. However, for a deterministic code, the resulting state will be equivalent with respect to the original piece of code: if we run the original program for the second time from such state, we will get the same final result as when running the program once.

6 Related Work

Term algebras admit quantifier elimination [5,21] and thus are natural candidates for synthesis. Our synthesis procedure is similar to quantifier elimination when it comes to eliminating variables that are constrained by an equality, with the additional requirement that the witness term be stored to serve in the program. However the treatment of disequalities is simplified: elimination procedures typically rewrite a disequality between a variable and a term into a disjunction of equalities between the same variable and terms constructed with different constructors [5, p.63sqq]. This has the advantage that the language of formulas needs not be extended, allowing for nested quantifiers to be eliminated one after the other. In our synthesis setting, this is not necessary, as we can rely on additional computable functions, as we have illustrated with the use of `distinct`, thus greatly simplifying the resulting program. A related area of research is compilation of unification in Prolog [1]. Note, however, that this process typically does not require handling of disequalities, so it deals with a simpler language.

Pattern-matching compilation is a task for which specialized procedures for term algebras have been developed [13,23]. When viewed through the prism of synthesis procedures, these algorithms can be thought of as procedures that are specialized for disjunctions of term equalities, and where the emphasis is put

on code reuse. We expect that using a combination of our synthesis procedures and common subexpression elimination techniques, one should be able to derive pattern-matching compilation schemes that would support, e.g., disjunctive patterns, non-linear patterns, and could take into account guards referring to integer predicates.

Our synthesis procedure for arrays is based on a reduction of constraints over arrays to constraints in the combined theory of indices and elements. In particular, our reduction is very close to the decision procedure for extensional arrays introduced by Stump et al. [20]. Combination of strongly disjoint theories is also used in the array decision procedure of de Moura and Bjørner [14], but the main focus of their work was to make array decision procedures more efficient by restricted application of fine-grained reduction rules. In the presence of unknown inputs, these techniques are not applicable in general.

Specialization of decision procedure for the purpose of predicate abstraction was considered in [10]. In addition to covering a different set of theories, our results are broader because our process generates not only a satisfiability check but also the values of variables.

7 Conclusions

We have presented synthesis procedures for two important theories: algebraic data types and arrays, and shown how they can be formulated in unified framework. Our contribution fills a previously unexplored area between two approaches for executing constraints. The first one is based on running SMT solvers at run time [7]. The second one uses quantifier elimination to enable compilation of specifications. In this paper we have shown that for two important theories supported by SMT solvers, compilation of constraints is also feasible. Moreover, much like SMT can be built and proved correct modularly, synthesis procedures can be combined as well using our framework.

As remarked in [9], compiled synthesis procedures could be viewed as a result of partial evaluation of the execution of a constraint solver at run time. This is a useful observation from a methodological point of view. However, it likely has similar limitations as an attempt to automatically transform an interpreter into a compiler. We therefore expect that the insights of researchers will continue to play a key role in designing synthesis procedures. These insights both take the form of understanding decidable logics, but also understanding how to solve certain classes of problems efficiently. Examples of manually deriving compiled code that can be more efficient than run-time search appear in both synthesis for term algebras and the synthesis of arrays. We can assume that the values in these theories are finitely generated by terms. Because these terms become known only at run time, it appears, at first, necessary to continue running decision procedure at run time. However, because the nature of processing steps is known at compile time, it was possible to generate loops instead of an invocation of a general-purpose constraint solver at run time.

References

1. Ait-Kaci, H.: Warren's Abstract Machine: A Tutorial Reconstruction. MIT Press (1991), available online as PDF
2. Baader, F., Snyder, W.: Unification theory. In: Handbook of Automated Reasoning. Elsevier (2001)
3. Flener, P.: Logic Program Synthesis from Incomplete Information. Kluwer Academic Publishers (1995)
4. Hamza, J., Jobstmann, B., Kuncak, V.: Synthesis for regular specifications over unbounded domains. In: FMCAD. pp. 101–109 (2010)
5. Hodges, W.: A Shorter Model Theory. Cambridge University Press (1997)
6. Jobstmann, B., Bloem, R.: Optimizations for ltl synthesis. In: FMCAD. pp. 117–124 (2006)
7. Köksal, A.S., Kuncak, V., Suter, P.: Constraints as control. In: POPL. pp. 151–164 (2012)
8. Kuncak, V., Mayer, M., Piskac, R., Suter, P.: Software synthesis procedures. CACM 55(2), 103–111 (2012)
9. Kuncak, V., Mayer, M., Piskac, R., Suter, P.: Complete functional synthesis. In: PLDI. pp. 316–329 (2010)
10. Lahiri, S.K., Ball, T., Cook, B.: Predicate abstraction via symbolic decision procedures. In: CAV. pp. 24–38 (2005)
11. Manna, Z., Waldinger, R.J.: Toward automatic program synthesis. CACM 14(3), 151–165 (1971)
12. Manna, Z., Waldinger, R.J.: A deductive approach to program synthesis. TOPLAS 2(1), 90–121 (1980)
13. Maranget, L.: Compiling pattern matching to good decision trees. In: ML. pp. 35–46 (2008)
14. de Moura, L.M., Bjørner, N.: Generalized, efficient array decision procedures. In: FMCAD. pp. 45–52 (2009)
15. Nelson, G., Oppen, D.C.: Simplification by cooperating decision procedures. TOPLAS 1(2), 245–257 (1979)
16. Smith, D.R.: KIDS: A semiautomatic program development system. TSE 16(9), 1024–1043 (1990)
17. Solar-Lezama, A., Tancau, L., Bodík, R., Seshia, S.A., Saraswat, V.A.: Combinatorial sketching for finite programs. In: ASPLOS. pp. 404–415 (2006)
18. Spielmann, A., Kuncak, V.: Synthesis for unbounded bit-vector arithmetic. In: IJCAR. pp. 499–513 (2012)
19. Srivastava, S., Gulwani, S., Foster, J.S.: From program verification to program synthesis. In: POPL. pp. 313–326 (2010)
20. Stump, A., Barrett, C.W., Dill, D.L., Levitt, J.R.: A decision procedure for an extensional theory of arrays. In: LICS. pp. 29–37 (2001)
21. Sturm, T., Weispfenning, V.: Quantifier elimination in term algebras: The case of finite languages. In: Ganzha, V.G., Mayr, E.W., Vorozhtsov, E.V. (eds.) Computer Algebra in Scientific Computing (CASC). TUM München (2002)
22. Vechev, M.T., Yahav, E., Yorsh, G.: Abstraction-guided synthesis of synchronization. In: POPL. pp. 327–338 (2010)
23. Wadler, P.: The Implementation of Functional Programming Languages: Efficient Compilation of Pattern-matching, chap. 5, pp. 78–103 (1987)