

Synthesis through Unification [★]

Rajeev Alur¹, Pavol Černý², and Arjun Radhakrishna¹

¹ University of Pennsylvania

² University of Colorado Boulder

Abstract. Given a specification and a set of candidate programs (program space), the program synthesis problem is to find a candidate program that satisfies the specification. We present the synthesis through unification (STUN) approach, which is an extension of the counterexample guided inductive synthesis (CEGIS) approach. In CEGIS, the synthesizer maintains a subset S of inputs and a candidate program `Prog` that is correct for S . The synthesizer repeatedly checks if there exists a counterexample input c such that the execution of `Prog` is incorrect on c . If so, the synthesizer enlarges S to include c , and picks a program from the program space that is correct for the new set S .

The STUN approach extends CEGIS with the idea that given a program `Prog` that is correct for a subset of inputs, the synthesizer can try to find a program `Prog'` that is correct for the rest of the inputs. If `Prog` and `Prog'` can be *unified* into a program in the program space, then a solution has been found. We present a generic synthesis procedure based on the STUN approach and specialize it for three different domains by providing the appropriate unification operators. We implemented these specializations in prototype tools, and we show that our tools often performs significantly better on standard benchmarks than a tool based on a pure CEGIS approach.

1 Introduction

The task of program synthesis is to construct a program that satisfies a given declarative specification. The computer-augmented programming [17, 2] approach allows the programmers to express their intent in different ways, for instance by providing a partial program, or by defining the space of candidate programs, or by providing positive and negative examples and scenarios. This approach to synthesis is becoming steadily more popular and successful [5].

We propose a novel algorithmic approach for the following problem: given a specification, a set of candidate programs (a program space), and a set of all possible inputs (an input space), find a candidate program that satisfies the specification on all inputs from the input space. The basic idea of our approach is simple: if we have a candidate program that is correct only on a part of the input space, we can attempt to find a program that works on the rest of the

[★] This research was supported in part by the NSF under award CCF 1421752 and the Expeditions award CCF 1138996, by DARPA under agreement FA8750-14-2-0263, by the Simons Foundation, and by a gift from the Intel Corporation.

input space, and then unify the two programs. The unification operator must ensure that the resulting program is in the program space.

The program space is syntactically restricted to a set which can be specified using a typed grammar. If this grammar contains `if` statements, and its expression language is expressive enough, then a simple unification operator exists. A program `Prog` for inputs that satisfy an expression C , and a program `Prog'` that works on the rest of the inputs can be unified into `if (C) then Prog else Prog'`. Even when `if` statements are not available, different unification operators may exist. These unification operators may be preferable to unification through `if` statements due to efficiency reasons. However, such unification operators may not be complete — it might not be possible to unify two given programs. We present an approach that deals with such cases with appropriate backtracking.

Our approach, which we dub STUN, works as follows: its first step is to choose a program `Prog` that works for a subset \mathcal{I}_G of the input space. This step can be performed by any existing method, for instance by multiple rounds of the CEGIS loop [16]. The STUN procedure then makes a recursive call to itself to attempt to synthesize a program `Prog'` for inputs on which `Prog` is incorrect. An additional parameter is passed to the recursive call — unification constraints that ensure that the program `Prog'` obtained from the recursive call is unifiable with `Prog`. If the recursive call succeeds, programs `Prog` and `Prog'` can be unified, and the solution to the original problem was found. If the recursive call fails, then we need to backtrack, and choose another candidate for program `Prog`. In this case, we also use a form of conflict-driven learning.

Problem domains. We instantiate the STUN approach to three different problem domains: bit-vector expressions, separable specifications for conditional linear arithmetic expressions, and non-separable specifications for conditional linear arithmetic expressions. In each domain, we provide a suitable unification operator, and we resolve the nondeterministic choices in the STUN algorithm.

We first consider the domain of bit-vector expressions. Here, the challenge is the absence of `if`-conditionals, which makes the unification operator harder to define. We represent bit-vector programs as (expr, ρ) , where `expr` is a bit-vector expression over input variables and additional auxiliary variables, and ρ is a constraint over the auxiliary variables. Two such pairs (expr_1, ρ_1) and (expr_2, ρ_2) can be unified if there exists a way to substitute the auxiliary variables in `expr`₁ and `expr`₂ to make the expressions equal, and the substitution satisfies the conjunction of ρ_1 and ρ_2 . A solver based on such a unification operator has comparable performance on standard benchmarks [1] as existing solvers.

For the second and third domain we consider, the program space is the set of conditional linear-arithmetic expressions (CLEs) over rationals. The difference between the two domains is in the form of specifications. Separable specifications are those where the specification only relates an input and its corresponding output. In contrast, the non-separable specifications can place constraints over outputs that correspond to different inputs. For instance, $x > 0 \implies f(x+2) = f(x)+7$ is a non-separable specification, as it relates outputs for multiple inputs.

The second domain of separable specifications and CLEs over rationals is an ideal example for STUN, as the unification operator is easy to implement using conditions of CLEs. We obtain an efficient implementation where partial

solutions are obtained by generalization of input-output examples, and such partial solutions are then unified. Our implementation of this procedure is order-of-magnitude faster on standard benchmarks than the existing solvers.

The third domain of non-separable specifications for CLEs requires solving constraints for which finding a solution might need an unbounded number of unification steps before convergence. We therefore implement a widening version of the unification operator, further demonstrating the generality of the STUN approach. Our implementation of this procedure performs on par with existing solvers on standard benchmarks.

Comparing CEGIS and STUN. The key conceptual difference between existing synthesis methods (CEGIS) and our STUN approach is as follows: CEGIS gradually collects a set of input-output examples (by querying the specification), and then finds a solution that matches all the examples. The STUN approach also collects input-output examples by querying the specification, but it finds a (general) solution for each of them separately, and then unifies the solutions. The STUN method has an advantage if solutions for different parts of the input space are different. In other words, CEGIS first combines subproblems, and then solves, while STUN first solves, and then combines solutions. The reason is that such solutions can be in many cases easily unifiable (if for instance the program space has `if` conditionals), but finding the whole solution at once for examples from the different parts of input space (as CEGIS requires) is difficult.

Summary. The main contributions of this work are two-fold. First, we propose a new approach to program synthesis based on unification of programs, and we develop a generic synthesis procedure using this approach. Second, we instantiate the STUN synthesis procedure to the domains of bit-vector expressions, and conditional linear expressions with separable and non-separable specifications. We show that in all cases, our solver has comparable performance to existing solvers, and in some cases (conditional linear-arithmetic expressions with separable specifications), the performance on standard benchmarks is several orders of magnitude better. This demonstrates the potential of the STUN approach.

2 Overview

In this section, we first present a simplified view of synthesis by unification (the UNIF loop), which works under very strong assumptions. We then describe what extensions are needed, and motivate our STUN approach.

UNIF loop. Let us fix a specification *Spec*, a *program space* \mathcal{P} (a set of candidate programs), and an *input space* \mathcal{I} . The program synthesis problem is to find a program in \mathcal{P} that satisfies the specification for all inputs in \mathcal{I} .

A classical approach to synthesis is the counterexample-guided inductive synthesis (CEGIS) loop. We choose the following presentation for CEGIS in order to contrast it with UNIF. In CEGIS (depicted in Figure 1), the synthesizer maintains a subset $\mathcal{J} \subseteq \mathcal{I}$ of inputs and a candidate program $\text{Prog} \in \mathcal{P}$ that is correct for \mathcal{J} . If $\mathcal{J} = \mathcal{I}$, i.e., if Prog is correct for all inputs in \mathcal{I} , the CEGIS loop terminates and returns Prog . If there is an input on which Prog is incorrect, the first step is to find such an input c . The second step is to find a program that is correct for both c and all the inputs in \mathcal{J} . In Figure 1, this is done in the call to `syntFitAll`). This process is then repeated until \mathcal{J} is equal to \mathcal{I} .




Fig. 1: CEGIS loop for input space \mathcal{I} and specification $Spec$




Fig. 2: UNIF loop for input space \mathcal{I} and specification $Spec$

The unification approach to synthesis is based on a simple observation: if we have a program \mathbf{Prog} that is correct for a subset \mathcal{J} of inputs (as in CEGIS), the synthesizer can try to find a program \mathbf{Prog}' that is correct for some of the inputs in $\mathcal{I} \setminus \mathcal{J}$, and then attempt to unify \mathbf{Prog} and \mathbf{Prog}' into a program in the program space \mathcal{P} . We call the latter option the UNIF loop. It is depicted in Figure 2. In more detail, the UNIF loop works as follows. We first call `syntFitSome` in order to synthesize a program \mathbf{Prog}' that works for some inputs in \mathcal{I} but not in \mathcal{J} . Let \mathcal{J}' be the set of those inputs in $\mathcal{I} \setminus \mathcal{J}$ for which \mathbf{Prog}' satisfies $Spec$.

Next, we consider two programs $\mathcal{J} \cdot \mathbf{Prog}$ and $\mathcal{J}' \cdot \mathbf{Prog}'$, where the notation $\mathcal{J} \cdot \mathbf{Prog}$ denotes a program that on inputs in \mathcal{J} behaves as \mathbf{Prog} , and on other inputs its behavior is undefined. We need to unify the two programs to produce a program (in the program space \mathcal{P}) which is defined on $\mathcal{J} \cup \mathcal{J}'$. The unification operator denoted by \oplus , and the unified program is obtained as $\mathcal{J} \cdot \mathbf{Prog} \oplus \mathcal{J}' \cdot \mathbf{Prog}'$. If the program space is closed under **if** conditionals, and if \mathbf{Prog} and \mathbf{Prog}' are in \mathcal{P} , then the unification is easy. We obtain **if** \mathcal{J} **then** \mathbf{Prog} **else if** \mathcal{J}' **then** \mathbf{Prog}' **else** \perp . Note that we abuse notation here: the symbols \mathcal{J} and \mathcal{J}' , when used in programs, denote expressions that define the corresponding input spaces.

Example 1. Consider the following specification for the function `max`.

$$Spec = f(x, y) \geq x \wedge f(x, y) \geq y \wedge (f(x, y) = x \vee f(x, y) = y)$$

The input space \mathcal{I} is the set of all pairs of integers. The program space \mathcal{P} is the set of all programs in a simple **if**-language with linear-arithmetic expressions.

We demonstrate the UNIF loop (Figure 2) on this example. We start with an empty program \perp . The program works for no inputs (i.e., the input space is \emptyset), so we start with the pair (\perp, \emptyset) at the top of Figure 2. As $\emptyset \neq \mathcal{I}$, we go to the right-hand side of Figure 2, and call the procedure `syntFitSome`.

We now describe the procedure `syntFitSome`($\mathcal{K}, Spec$) for the linear arithmetic domain. It takes two parameters: a set of inputs \mathcal{K} , and a specification $Spec$, and returns a pair $(\mathcal{J}', \mathbf{Prog}')$ consisting of a set $\emptyset \neq \mathcal{J}' \subseteq \mathcal{K}$ and a program \mathbf{Prog}' which is correct on \mathcal{J}' . We pick an input-output example from the input space \mathcal{K} . This can be done by using a satisfiability solver to obtain a model of $Spec$. Let us assume that the specification is in CNF. An input-output example satisfies at least one atom in each clause. Let us pick those atoms. For instance, for the example $(2, 3) \rightarrow 3$, we get the following conjunction G of atoms: $G \equiv f(x, y) \geq x \wedge f(x, y) \geq y \wedge f(x, y) = y$. We now generate a solution for the input-output example and G . For linear arithmetic, we could “solve for $f(x, y)$ ”, i.e. replace $f(x, y)$ by t and solve for t . Let us assume that the solution \mathbf{Prog}_0 that we obtain is a function that on any input (x, y) returns y . We then plug the solution \mathbf{Prog}_0 to G , and simplify the resulting formula in order to obtain G_0 ,

Algorithm 1 The STUN (synthesis through unification) procedure

Input: Specification $Spec$, Program space \mathcal{P} , Input space \mathcal{I} , outer unification constraints (OUCs) ψ

Output: $Prog \in \mathcal{P}$ s.t. $\forall inp \in \mathcal{I} : Prog[inp] \models Spec$ and $Prog \models \psi$, or **None**

Global variables: learned unification constraints (LUCs) β , initialized to **true**

```
1:  $\varphi \leftarrow \text{true}$  // CEGIS constraints
2: if  $\mathcal{I} = \emptyset$  return  $\top, \text{true}$  // input space is empty, base case of recursion
3: while true do
4:  $(Prog, \text{timeOut}) \leftarrow \text{Generate}(\mathcal{P}, Spec, \mathcal{I}, \varphi, \psi, \beta)$  // generate next candidate
5: if  $Prog = \text{None}$  then
6:   if  $\neg \text{timeOut}$  then
7:      $\beta \leftarrow \beta \wedge \text{LearnFrom}(Spec, \psi, \beta)$  //learn unification constraints
8:   return None //no solution exists
9:    $inp \leftarrow \text{PickInput}(\mathcal{I}, Prog)$  //take a positive- or a counter-example
10:  if  $Prog[inp] \not\models Spec$  then
11:     $\varphi \leftarrow \varphi \wedge \text{project}(Spec, inp)$  //get a constraint from a counter-example
12:  else
13:     $\mathcal{I}_G, \mathcal{I}_B \leftarrow \text{splitInpSpace}(Spec, Prog, inp)$  //  $\mathcal{I}_G \subseteq \{inp' \mid Prog[inp'] \models Spec\}$ 
// and  $inp \in \mathcal{I}_G$ , so  $\mathcal{I}_B \subsetneq \mathcal{I}$  and we can make a recursive call
14:     $Prog' \leftarrow \text{STUN}(Spec, \psi \wedge \text{UnifConstr}(\mathcal{I}_G, Prog), \mathcal{P}, \mathcal{I}_B)$  //recursive call
15:    if  $Prog' \neq \text{None}$  return  $\mathcal{I}_G \cdot Prog \oplus \mathcal{I}_B \cdot Prog'$  //return the unified program
```

where G_0 is $y \geq x$. G_0 defines the set of inputs for which the solution is correct. We have thus effectively obtain the pair $(G_0, Prog_0)$ that the function returns (this effectively represents the program **if** $y \geq x$ **then** y **else** \perp).

In the second iteration, we now call the function $\text{syntFitSome}(\mathcal{K}, Spec)$ with the parameter $\mathcal{K} = \neg G_0$. We now ask for an input-output example where the input satisfies $\neg G_0$. Let us say we obtain $(5, 4)$, with output 5. By a similar process as above, we obtain a program $Prog_1$ that for all inputs (x, y) returns x , and works for all input that satisfy $G_1 \equiv x \geq y$.

The next step of the STUN loop asks us to perform the unification $(G_0 \cdot Prog_0) \oplus (G_1 \cdot Prog_1)$. Given that we have **if** conditionals in the language, this step is simple. We unify the two programs to obtain: **if** $y \geq x$ **then** y **else** x .

From the UNIF loop to STUN The main assumption that the UNIF loop makes is that the unification operator \oplus always succeeds. We already mentioned that this is the case when the program space is closed under **if** conditionals. If the program space is not closed under **if** conditionals, or if we do not wish to use this form of unification for other reasons, then the UNIF loop needs to be extended. An example of a program space that is not closed under **if** conditionals, and that is an interesting synthesis target, are bit-vector expressions.

The STUN algorithm extends UNIF with backtracking (as explained in the introduction, this is needed since the unification can fail), and at each level, a CEGIS loop can be used in syntFitSome . The CEGIS and UNIF loops are thus combined, and the combination can be fine-tuned for individual domains.

3 Synthesis through Unification algorithm

Overview. The STUN procedure is presented in Algorithm 1. The input to the algorithm consists of a specification $Spec$, a program space \mathcal{P} , input space \mathcal{I} ,

and *outer unification constraints (OUCs)* ψ . OUCs are constraints on the program space which are needed if the synthesized program will need to be unified with an already created program. The algorithm is implemented as a recursive (backtracking) procedure STUN. At each level, a decision is tried: a candidate program that satisfies OUCs is generated, and passed to the recursive call. If the recursive call is successful, the returned program is unified with the current candidate. If the recursive call is unsuccessful, it records *learned unification constraint (LUCs)* to the global variable β , ensuring progress.

Algorithm description. The algorithm first checks whether the input space is empty (this is the base case of our recursion). If so, we return a program \top (Line 2), a program which can be unified with any other program.

If the input space \mathcal{I} is not empty, we start the main loop (Line 3). In the loop, we need to generate a program **Prog** (Line 4) that works for a nonempty subset of \mathcal{I} . The generated program has to satisfy “CEGIS” constraints φ (that ensure that the program is correct on previously seen inputs at this level of recursion), OUCs ψ that ensure that the program is unifiable with programs already created in the upper levels of recursion, and LUCs β , which collects constraints learned from the lower levels of recursion. If the call to **Generate** fails (i.e., returns **None**), we exit this level of recursion, and learn constraints unification constraints that can be inferred from the failed exploration (Line 7). The only exception is when **Generate** fails due to a timeout, in which case we are not sure whether the task was unrealizable, and so no constraints are learned. Learning the constraints (computed by the function *LearnFrom*) is a form of conflict-driven learning.

Once a program **Prog** is generated, we need to check whether it works for all inputs in \mathcal{I} . If it does not, we need to decide whether to improve **Prog** (in a CEGIS-like way), or generate a program **Prog'** that works for inputs on which **Prog** does not work. The decision is made as follows. We pick an input **inp** and check whether the program **Prog** is correct on **inp** (Line 10). If **Prog** is not correct on **inp**, then we have found a counterexample, and we use it to strengthen our CEGIS constraints (Line 11). We refer to this branch as CEGIS-like branch.

If **Prog** is correct on **inp**, then we know that **Prog** is correct for at least one input, and we can make a recursive call to generate a program that is correct for the inputs for which **Prog** is not. We refer to this branch as the UNIF-like branch. The first step is to split the input space \mathcal{I} into the set \mathcal{I}_G (an underapproximation of the set of inputs on which **Prog** works containing at least **inp**), and \mathcal{I}_B , the rest of the inputs (Line 13). We can now make the recursive call on \mathcal{I}_B (Line 14). We pass the OUCs ψ to the recursive call, in addition to the information that the returned program will need to be unified with **Prog** (this is accomplished by adding *UnifConstr*($\mathcal{I}_G, \mathbf{Prog}$)). If the recursive call does not find a program (i.e., returns **Prog'** = **None**), then the loop continues, and another candidate is generated. If the recursive call successfully returns a program **Prog'**, this program is unified with **Prog** (Line 15). In more detail, we have a program **Prog** that works on inputs in \mathcal{I}_G , and a program **Prog'** that works on inputs in \mathcal{I}_B , and we unify them with the unification operator \oplus to produce $\mathcal{I}_G \cdot \mathbf{Prog} \oplus \mathcal{I}_B \cdot \mathbf{Prog}'$. We know that the unification operator will succeed, as the unification constraint *UnifConstr*($\mathcal{I}_G, \mathbf{Prog}$) was passed to the recursive call.

The input choice (line 9), here nondeterministic, can be tuned for individual domains to favor positive- or counter-examples, and hence, CEGIS or UNIF.

Example 2. Consider a specification that requires that the right-most bit set to 1 in the input bit-vector is reset to 0. This problem comes from the Hacker’s Delight collection [20]. A correct solution is, for instance, given by the expression $x \& (x - 1)$. We illustrate the STUN procedure on this example. The full STUN procedure for the bit-vector domain will be presented in Section 4.

Unification. The unification operator $\mathcal{I}_G \cdot \text{Prog} \oplus \mathcal{I}_B \cdot \text{Prog}'$ works as follows. $\mathcal{I}_G \cdot \text{Prog}$ and $\mathcal{I}_B \cdot \text{Prog}'$ can be unified if there exists a way to substitute the constants c_i and c'_i occurring in Prog and Prog' with sub-expressions expr_i and expr'_i such that after the substitution, Prog and Prog' are equal to the same program Prog^* , and for all input in \mathcal{I}_G , $\text{expr}_i[i] = c_i$ and for all inputs in \mathcal{I}_B , $\text{expr}'_i[i] = c'_i$. Note that this is a (very) simplified version of the unification operator introduced in the next section. It is used here to illustrate the algorithm.

Unification gone wrong. Let us assume that the **Generate** function at Line 4 generates the program $x | 0$ (this can happen if say the simpler programs already failed). Note that $|$ is the bitwise or operator. Now let us assume that at Line 9, we pick the input 0. The program matches *Spec* at this input. The set \mathcal{I}_G is $\{0\}$, and we go to the recursive call at Line 14 for the rest of the input space, with the constraint that the returned program must be unifiable with $x | 0$. In the recursive call, **Generate** is supposed to find a program that is unifiable with $x | 0$, i.e., of the form $x | c$ for some constant c . Further, for the recursive call to finally succeed (i.e., take the else branch at Line 12), we need this program to be correct on some input other than $x = 0$. However, as it can be seen, there is no such program and input. Hence, the procedure eventually backtracks while adding a constraint that enforces that the program $x | 0$ will no longer be attempted.

Unification gone right. After the backtracking, with the additional constraint, the program generation procedure is forbidden from generating the program $x | 0$. The **Generate** procedure instead generates say $x \& -1$. As before, for the recursive call to finally succeed, the program generation procedure is asked to find a program unifiable with $x \& -1$ (i.e., of the form $x \& c$) that works for an input other than 0. Let us assume that generated program in the next level of recursion is $x \& 4$; one input for which this is correct is $x = 5$. Attempting to unify these functions, the unification operator is asked to find an expression expr such that $\text{expr}[0/x] = -1$ and $\text{expr}[5/x] = 4$. One such candidate for expr is $x - 1$. This leads to a valid solution $x \& (x - 1)$ to the original synthesis problem.

Soundness. The procedure $\text{splitInpSpace}(\text{Spec}, \text{Prog}, \text{inp})$ is sound if for every invocation, it returns a pair $(\mathcal{I}_G, \mathcal{I}_B)$ such that $\{\text{inp}\} \subseteq \mathcal{I}_G \subseteq \{\text{inp}' \mid \text{Prog}[\text{inp}'] \models \text{Spec}\} \wedge \mathcal{I}_B = \mathcal{I} \setminus \mathcal{I}_G$. The unification operator \oplus is sound w.r.t. *Spec* and \mathcal{P} if for programs Prog_1 and Prog_2 satisfying *Spec* on inputs in \mathcal{I}_1 and \mathcal{I}_2 , respectively, the program $\mathcal{I}_1 \cdot \text{Prog}_1 \oplus \mathcal{I}_2 \cdot \text{Prog}_2$ is in \mathcal{P} and that it satisfies *Spec* on $\mathcal{I}_1 \cup \mathcal{I}_2$. The procedure STUN is *sound* if for all inputs $\mathcal{P}, \mathcal{I}, \text{Spec}, \psi$, it returns a program Prog such that $\text{Prog} \in \mathcal{P}$ and that $\forall \text{inp} \in \mathcal{I} : \text{Prog}[\text{inp}] \models \text{Spec}$.

Theorem 1. *Let us fix specification *Spec* and program space \mathcal{P} . If splitInpSpace and the unification operator \oplus are sound, then the STUN procedure is sound.*

Domains and Specifications. We instantiate STUN approach to three domains: bit-vector expressions, separable specifications for conditional linear-

arithmetic expressions, and non-separable specifications for conditional linear arithmetic expressions. Separable specifications are those where the specification relates an input and its corresponding output, but does not constrain outputs that correspond to different inputs. Formally, we define separable specifications syntactically — they are of the form $f(x) = o \wedge \Phi(o, x)$, where x is the tuple of all input variables, o is the output variable, f is the function being specified, and Φ is a formula. For example, the specification $Spec \equiv f(x, y) \geq x \wedge f(x, y) \geq y$ is separable as $Spec = (f(x, y) = o) \wedge (o \geq x \wedge o \geq y)$, and the specification $f(0) = 1 \vee f(1) = 1$ is a non-separable specification.

Notes about implementation. We have implemented the STUN procedure for each of the three domains described above as a suite of tools. In each case, we evaluate our tool on the benchmarks from the SyGuS competition 2014 [1], and compare the performance of our tool against the enumerative solver ESOLVER [2, 18]. The tool ESOLVER was the overall winner in the SyGuS competition 2014, and hence, is a good yardstick that represents the state of the art.

4 Domain: Bit-Vector Expressions

The first domain to which we apply the STUN approach is the domain of bit-vector expressions specified by separable specifications. Each bit-vector expression is either an input variable, a constant, or a standard bit-vector operator applied to two sub-expressions. This syntax does not have a top level `if-then-else` operator that allows unification of any two arbitrary programs.

Here, we instantiate the **Generate** procedure and the unification operator of Algorithm 1 to obtain a nondeterministic synthesis procedure (nondeterministic mainly in picking inputs that choose between the CEGIS-like and UNIF-like branches). Later, we present a practical deterministic version of the algorithm.

Representing candidate programs. In the following discussion, we represent programs using an alternative formalism that lets us lazily instantiate constants in the program. This representation is for convenience only—the procedure can be stated without using it. Formally, a *candidate bit-vector program* **Prog** over inputs v_1, \dots, v_n is a tuple $\langle \mathbf{expr}, \rho \rangle$ where: (a) **expr** is a bit-vector expression over $\{v_1, \dots, v_n\}$ and auxiliary variables $\{\mathbf{SubProg}_0, \dots, \mathbf{SubProg}_m\}$ such that each $\mathbf{SubProg}_i$ occurs exactly once in **expr**; and (b) ρ is a satisfiable constraint over $\mathbf{SubProg}_i$'s. Variables $\mathbf{SubProg}_i$ represent constants of **expr** whose exact values are yet to be synthesized, and ρ is a constraint on their values. Intuitively, in the intermediate steps of the algorithm, instead of generating programs with explicit constants, we generate programs with symbolic constants along with constraints on them. A concrete program can be obtained by replacing the symbolic constants with values from some satisfying assignment of ρ .

Unification. As mentioned briefly in Section 3, two candidate programs are unifiable if the constants occurring in the expressions can be substituted with sub-expressions to obtain a common expression. However, the presence of symbolic constants requires a more involved definition of the unification operator. Further, note that the symbolic constants in the two programs do not have to be the same. Formally, programs $\mathbf{Prog} = \langle \mathbf{expr}, \rho \rangle$ and $\mathbf{Prog}' = \langle \mathbf{expr}', \rho' \rangle$ over input spaces \mathcal{I} and \mathcal{I}' are unifiable if:

- There exists an expression \mathbf{expr}^* that can be obtained from \mathbf{expr} by replacing each variable SubProg_i in \mathbf{expr} by an expression \mathbf{expr}_i , over the formal inputs $\{v_1, \dots, v_n\}$ and new auxiliary variables $\{\text{SubProg}_1^*, \dots, \text{SubProg}_k^*\}$. Further, the same expression \mathbf{expr}^* should also be obtainable from \mathbf{expr}' by replacing each of its sub-programs $\text{SubProg}'_i$ by an expression \mathbf{expr}'_i .
- Constraint $\rho^* = \bigwedge_{\mathcal{V}} \rho[\forall i. \mathbf{expr}_i[\mathcal{V}]/\text{SubProg}_i] \wedge \bigwedge_{\mathcal{V}'} \rho'[\forall i. \mathbf{expr}'_i[\mathcal{V}']/\text{SubProg}'_i]$ is satisfiable. Here, \mathcal{V} and \mathcal{V}' range over inputs from \mathcal{I} and \mathcal{I}' , respectively.

If the above conditions hold, one possible unified program $\mathcal{I} \cdot \text{Prog} \oplus \mathcal{I}' \cdot \text{Prog}'$ is $\text{Prog}^* = (\mathbf{expr}^*, \rho^*)$. Intuitively, in the unified program, each SubProg_i is replaced with a sub-expression \mathbf{expr}_i , and further, ρ^* ensures that the constraints from the individual programs on the value of these sub-expressions are satisfied.

Example 3. The programs $\text{Prog} = (x \ \& \ \text{SubProg}_0, \text{SubProg}_0 = -1)$ and $\text{Prog}' = (x \ \& \ \text{SubProg}'_0, \text{SubProg}'_0 = 4)$ over the input spaces $\mathcal{I} = (x = 0)$ and $\mathcal{I}' = (x = 5)$ can be unified into $(x \ \& \ (x - \text{SubProg}_0^*), (0 - \text{SubProg}_0^* = -1) \wedge (5 - \text{SubProg}_0^* = 4))$. Here, both SubProg_0 and $\text{SubProg}'_0$ are replaced with $x - \text{SubProg}_0^*$ and the constraints have been instantiated with inputs from corresponding input spaces.

Unification constraints. In this domain, an outer unification constraint ψ is given by a candidate program Prog_T . Program $(\mathbf{expr}, \rho) \models \psi$ if $\text{Prog}_T = (\mathbf{expr}_T, \rho_T)$ and \mathbf{expr} can be obtained from \mathbf{expr}_T by replacing each SubProg_i^T with appropriate sub-expressions. A learned unification constraint β is given by $\bigwedge \text{Not}(\text{Prog}_F^i)$. Program $(\mathbf{expr}, \rho) \models \beta$ if for each $\text{Prog}_F^i = (\mathbf{expr}_F, \rho_F)$, there is no substitution of SubProg_i^F 's that transforms \mathbf{expr}_F to \mathbf{expr} . Intuitively, a Prog satisfies $\psi = \text{Prog}_T$ and $\beta = \bigwedge \text{Not}(\text{Prog}_F^i)$ if Prog can be unified with Prog_T and cannot be unified with any of Prog_F^i . Boolean combinations of unification constraints can be easily defined. In Algorithm 1, we define $\text{UnifConstr}(\mathcal{I}, \text{Prog}) = \text{Prog}$ and $\text{LearnFrom}(\text{Spec}, \psi, \beta) = \text{Not}(\psi)$. Note that using the alternate representation for programs having symbolic constants lets us have a very simple LearnFrom that just negates ψ – in general, a more complex LearnFrom might be needed.

Program generation. A simple *Generate* procedure enumerates programs, ordered by size, and checks if the expression satisfies all the constraints.

Theorem 2. *Algorithm 1 instantiated with the procedures detailed above is a sound synthesis procedure for bit-vector expressions.*

A practical algorithm. We instantiate the non-deterministic choices in the procedure from Theorem 2 to obtain a deterministic procedure. Intuitively, this procedure maintains a set of candidate programs and explores them in a fixed order based on size. Further, we optimize the program generation procedure to only examine programs that satisfy the unification constraints, instead of following a generate-and-test procedure. Additionally, we eliminate the recursive call in Algorithm 1, and instead store the variables \mathcal{I}_G locally with individual candidate programs. Essentially, we pass additional information to convert the recursive call into a tail call. Formally, we replace ρ in the candidate programs with $\{(\mathcal{V}_0, \rho_0), \dots, (\mathcal{V}_k, \rho_k)\}$ where \mathcal{V}_i 's are input valuations that represent \mathcal{I}_G from previous recursive calls. Initially, the list of candidate programs contains the program $(\text{SubProg}_0, \emptyset)$. In each step, we pick the first candidate (say

$(\text{expr}, \{(\mathcal{V}_0, \rho_0), \dots\})$) and concretize expr to expr^* by substituting SubProg_i 's with values from a model of $\bigwedge_i \rho_i$. If expr^* satisfies Spec , we return it.

Algorithm 2 A deterministic STUN algorithm for bit-vector expressions

```

1:  $Candidates \leftarrow \langle (\text{SubProg}_0, \emptyset) \rangle$ 
2: while true do
3:    $(\text{expr}, \{(\mathcal{V}_0, \rho_0), \dots, (\mathcal{V}_n, \rho_n)\}) \leftarrow Candidates[0]$ 
4:    $\text{expr}^* \leftarrow \text{substitute}(\text{expr}, \text{getModel}(\bigwedge_i \rho_i))$ 
5:   if  $\nexists \text{inp} : \text{expr}^*[\text{inp}] \models \text{Spec}$  return  $\text{expr}^*$ 
6:    $\rho_{\text{inp}} \leftarrow \text{concretize}(\text{expr}, \text{Spec}, \text{inp})$  where  $\text{expr}^*[\text{inp}] \not\models \text{Spec}$ 
7:   if  $\neg \text{Satisfiable}(\rho_{\text{inp}})$  then
8:      $Candidates \leftarrow \text{tail}(Candidates)$ 
9:   else
10:     $Candidates[0] \leftarrow (\text{expr}, \{(\mathcal{V}_0, \rho_0), \dots, (\mathcal{V}_n, \rho_n)\} \cup \{(\mathcal{V}_{\text{inp}}, \rho_{\text{inp}})\})$ 
11:    if  $\neg \text{Satisfiable}(\bigwedge \rho_i \wedge \rho_{\text{inp}})$  then
12:       $Candidates \leftarrow \text{tail}(Candidates)$ 
13:      for all  $\text{SubProg}_i \in \text{AuxVariables}(\text{expr})$ ,  $\text{expr}' \leftarrow \text{LevelOneExprs}()$  do
14:         $Candidates \leftarrow \text{append}(Candidates, \text{substitute}(\text{Prog}, (\text{SubProg}_i, \text{expr}')))$ 

```

Otherwise, there exists an input inp on which expr^* is incorrect. We obtain a new constraint ρ_{inp} on SubProg_i 's by substituting the input and the expression expr^* in the specification Spec . If ρ_{inp} is unsatisfiable, there are no expressions which can be substituted for SubProg_i 's to make expr correct on inp . Hence, the current candidate is eliminated—this is equivalent to a failing recursive call in the non-deterministic version.

Instead, if ρ_{inp} is satisfiable, it is added to the candidate program. Now, if $\bigwedge \rho_i \wedge \rho_{\text{inp}}$ is unsatisfiable, the symbolic constants SubProg_i 's cannot be instantiated with explicit constants to make expr correct on all the seen inputs \mathcal{V}_i . However, SubProg_i 's can possibly be instantiated with other sub-expressions. Hence, we replace the current candidate with programs where each SubProg_i is replaced with a small expression of the form $\text{operator}(e_1, e_2)$ where e_1 and e_2 are either input variables or fresh SubProg_i variables. Note that while substituting these expression for SubProg_i in ρ_j , the input variables are replaced with the corresponding values from \mathcal{V}_j .

Informally, each (expr, ρ_i) is a candidate program generated at one level of the recursion in the non-deterministic algorithm and each valuation \mathcal{V}_i is the corresponding input-space. An iteration where ρ_{inp} is unsatisfiable is a case where there is no program that is correct on inp is unifiable with the already generated program, and an iteration where $\bigwedge \rho_i \wedge \rho_{\text{inp}}$ is unsatisfiable when the unification procedure cannot replace the symbolic constants with explicit constants, but instead has to search through more complex expressions for the substitution.

Theorem 3. *Algorithm 2 is a sound and complete synthesis procedure for bit-vector expressions.*

Experiments. We implemented Algorithm 2 in a tool called AUK and evaluated it on benchmarks from the bit-vector track of SyGuS competition 2014 [1]. For the full summary of results, see the full version [3]. For easy benchmarks (where both tools take < 1 second), ESOLVER is faster than AUK. However, on larger benchmarks, the performance of AUK is better. We believe that these results are

due to ESOLVER being able to enumerate small solutions extremely fast, while AUK starts on the expensive theory reasoning. On larger benchmarks, AUK is able to eliminate larger sets of candidates due to the unification constraints while ESOLVER is slowed down by the sheer number of candidate programs.

5 Domain: CLEs with Separable Specifications

We now apply the STUN approach to the domain of conditional linear arithmetic expressions (CLEs). A program Prog in this domain is either a linear expression over the input variables or is $\text{if}(\text{cond}) \text{Prog} \text{ else } \text{Prog}'$, where cond is a boolean combination of linear inequalities. This is an ideal domain for the UNIF loop due to the natural unification operator that uses the if-then-else construct. Here, we present our algorithm for the case where the variables range over rationals. Later, we discuss briefly how to extend the technique to integer variables.

Unification. Given two CLEs Prog and Prog' , and input spaces \mathcal{I} and \mathcal{I}' , we define $\mathcal{I} \cdot \text{Prog} \oplus \mathcal{I}' \cdot \text{Prog}'$ to be the program if $(\mathcal{I}) \text{Prog}$ else if $(\mathcal{I}') \text{Prog}'$ else \perp . Note that we assume that \mathcal{I} and \mathcal{I}' are expressed as linear constraints. Here, since any two programs can be unified, unification constraints are not used.

Program Generation. Algorithm 3 is the program generation procedure *Generate* for CLEs for rational arithmetic specifications. Given a specification Spec and input space \mathcal{I} , it first generates a concrete input-output example such that the input is in \mathcal{I} and the example satisfies Spec . Then, it generalizes the input-output pair to a program as follows. From each clause of Spec , we pick one disjunct that evaluates to true for the current input-output pair. Each disjunct that constrains the output can be expressed as $o \text{ op } \phi$ where $\text{op} \in \{\leq, \geq, <, >\}$ and ϕ is a linear expression over the input variables. Recall from the definition of separable specifications that o is the output variable that represents the output of the function to be synthesized. Each such inequality gives us either an upper or a lower bound (in terms of input variables) on the output variable. These bounds are evaluated using the input-output example, and the strictest upper and lower bounds are chosen. The algorithm then returns an expression Prog that respects these strictest bounds. We define the *SplitInpSpace* procedure from Algorithm 1 as follows: input space \mathcal{I}_G is obtained by substituting the program Prog into the disjuncts, and \mathcal{I}_B is obtained as $\mathcal{I} \wedge \neg \mathcal{I}_G$.

Algorithm 3 Procedure *Generate*

Require: Specification Spec in CNF, Input space \mathcal{I}

Ensure: Candidate program Prog

- 1: **if** $\mathcal{I} = \emptyset$ **return** \top
 - 2: $pe_x \leftarrow \text{getModel}(\mathcal{I} \wedge \text{Spec})$
 - 3: $LB \leftarrow -\infty, UB \leftarrow \infty$
 - 4: **for all** *Clause* of Spec **do**
 - 5: Pick *Disjunct* in *Clause* such that $\text{Disjunct}[pe_x]$ holds
 - 6: **if** o occurs in *Disjunct* and $\text{Disjunct} \equiv (o \text{ op } \phi)$ **then**
 - 7: **case** $\text{op} \in \{\leq, <\} \wedge UB[pe_x] > \phi[pe_x]$: $UB \leftarrow \phi$
 - 8: **case** $\text{op} \in \{\geq, >\} \wedge LB[pe_x] < \phi[pe_x]$: $LB \leftarrow \phi$
 - 9: **return** $(LB + UB)/2$
-

Theorem 4. *Algorithm 1 instantiated with the procedures detailed above is a sound and complete synthesis procedure for conditional linear rational arithmetic expressions specified using separable specifications.*

Extension to integers. The above procedure cannot be directly applied when variables range over integers instead of rationals. Here, each disjunct can be put into the form $c \cdot o \text{ op } \phi$ where c is a positive integer and ϕ is a linear integer expression over inputs. For rationals, this constraint can be normalized to obtain $o \text{ op } \frac{1}{c}\phi$. In the domain of integers, $\frac{1}{c}\phi$ is not necessarily an integer.

There are two possible ways to solve this problem. A simple solution is to modify the syntax of the programs to allow floor $\lfloor \cdot \rfloor$ and ceiling $\lceil \cdot \rceil$ functions. Then, $c \cdot o \leq \phi$ and $c \cdot o \geq \phi$ can be normalized as $o \leq \lfloor \phi/c \rfloor$ and $o \geq \lceil \phi/c \rceil$. The generation procedure can then proceed using these normalized expressions. The alternative approach is to use a full-fledged decision procedure for solving the constraints of the form $o \text{ op } \frac{1}{c}\phi$. However, this introduces divisibility constraints into the generated program. For a detailed explanation on this approach and techniques for eliminating the divisibility constraints, see [14].




Fig. 3: Results on separable linear integer benchmarks

Experiments. We implemented the above procedure in a tool called PUFFIN and evaluated it on benchmarks from the linear integer arithmetic track with separable specifications from the SyGuS competition 2014. The results on three classes of benchmarks (`maxn`, `array_searchn`, and `array_sumn`) have been summarized in Figure 3. The `maxn` benchmarks specify a function that outputs the maximum of n input variables (the illustrative example from Section 2 is `max2`). Note that the SyGuS competition benchmarks only go up to `max5`. The `array_searchn` and `array_sumn` benchmarks respectively specify functions that search for a given input in an array, and check if the sum of two consecutive elements in an array is equal to a given value. In all these benchmarks, our tool significantly outperforms ESOLVER and other CEGIS-based solvers. This is because the CEGIS solvers try to generate the whole program at once, which is a complex expression, while our solver combines simple expressions generated for parts of the input spaces where the output expression is simple.

6 Domain: Non-Separable Specifications for CLEs

Here, we consider CLEs specified by non-separable specifications. While this domain allows for simple unification, non-separable specifications introduce complications. Further, unlike the previous domains, the problem itself is undecidable.

First, we define what it means for a program `Prog` to satisfy a non-separable specification on an input space \mathcal{I} . In further discussion, we assume that the

program to be synthesized is represented by the function f in all specifications and formulae. We say that **Prog** satisfies $Spec$ on \mathcal{I} if $Spec$ holds whenever the inputs to f in each occurrence in $Spec$ belong to \mathcal{I} . For example, program $\mathbf{Prog}(i)$ satisfies $Spec \equiv f(x) = 1 \wedge x' = x + 1 \implies f(x') = 1$ on the input space $0 \leq i \leq 2$ if $(0 \leq x \leq 2 \wedge 0 \leq x' \leq 2) \implies Spec[f \leftarrow \mathbf{Prog}]$ holds, i.e., we require $Spec$ to hold when both x and x' belong to the input space.

Unification and Unification constraints. The unification operator we use is the same as in Section 5. However, for non-separable specifications, the outputs produced by **Prog** on \mathcal{I} may constrain the outputs of \mathbf{Prog}' on \mathcal{I}' , and hence, we need non-trivial unification constraints. An outer unification constraint ψ is a sequence $\langle (\mathcal{I}_0, \mathbf{Prog}_0), (\mathcal{I}_1, \mathbf{Prog}_1), \dots \rangle$ where \mathcal{I}_i 's and \mathbf{Prog}_i 's are input spaces and programs, respectively. A learned unification constraint β is given by $\bigwedge \rho_i$ where each ρ_i is a formula over f having no other free variables. Intuitively, \mathcal{I}_i and \mathbf{Prog}_i fix parts of the synthesized function, and ρ_i 's enforce the required relationships between the outputs produced by different \mathbf{Prog}_i 's. Formally, $\mathbf{Prog} \models \psi$ if its outputs agree with each \mathbf{Prog}_i on \mathcal{I}_i and $\mathbf{Prog} \models \beta$ if $\bigwedge \rho_i[\mathbf{Prog}/f]$ holds.

Program Generation. The *Generate* procedure works using input-output examples as in the previous section. However, it is significantly more complex due to the presence of multiple function invocations in $Spec$. Intuitively, we replace all function invocations except one with the partial programs from the unification constraints and then solve the arising separable specification using techniques from the previous section. We explain the procedure in detail using an example.

Example 4. Consider the specification $Spec$ given by $x \neq y \implies f(x) + f(y) = 10$. Here, the only solution is the constant function 5. Now, assume that the synthesis procedure has guessed that \mathbf{Prog}_0 given by $\mathbf{Prog}_0(i) = 0$ is a program that satisfies $Spec$ for the input space $\mathcal{I}_0 \equiv i = 0$.

The unification constraint $\psi_0 = \langle (\mathbf{Prog}_0, \mathcal{I}_0) \rangle$ is passed to the recursive call to ensure that the synthesized function satisfies $f(0) = 0$. The program generation function in the recursive call works as follows: it replaces the invocation $f(x)$ in $Spec$ with the partial function from ψ to obtain the constraint $(x = 0 \wedge x \neq y \implies \mathbf{Prog}_0(0) + f(y) = 10)$. Solving to obtain the next program and input space, we get $\mathbf{Prog}_1(i) = 10$ for the input space $\mathcal{I}_1 \equiv i = 1$. Now, the unification constraint passed to the next recursive call is $\psi = \langle (\mathbf{Prog}_0, \mathcal{I}_0), (\mathbf{Prog}_1, \mathcal{I}_1) \rangle$.

Again, instantiating $f(x)$ with \mathbf{Prog}_0 and \mathbf{Prog}_1 in the respective input spaces, we obtain the constraint $(x = 0 \wedge x \neq y \implies \mathbf{Prog}_0(x) + f(y) = 10) \wedge (x = 1 \wedge x \neq y \implies \mathbf{Prog}_1(x) + f(y) = 10)$. Now, this constraint does not have a solution—for $y = 2$, there is no possible value for $f(y)$. Here, a reason $\beta = \rho_0$ (say $\rho_0 \equiv f(1) = f(0)$) is learnt for the unsatisfiability and added to the learned constraint. Note that this conflict-driven learning is captured in the function *LearnFrom* in Algorithm 1. Now, in the parent call, no program satisfies β as well as $\psi = \langle (\mathbf{Prog}_0, \mathcal{I}_0), (\mathbf{Prog}_1, \mathcal{I}_1) \rangle$. By a similar unsatisfiability analysis, we get $\rho_1 \equiv f(0) = 5$ as the additional learned constraint. Finally, at the top level, with $\beta \equiv f(0) = f(1) \wedge f(0) = 5$, we synthesize the right value for $f(0)$.

Example 5 (Acceleration). Let $Spec \equiv (0 \leq x, y \leq 2 \implies f(x, y) = 1) \wedge (x = 4 \wedge y = 0 \implies f(x, y) = 0) \wedge (f(x, y) = 1 \wedge (x', y') = (x + 2, y + 2) \implies f(x', y') = 1)$.

The synthesis procedure first obtains the candidate program $\text{Prog}_0(i, j) = 1$ on the input space $\mathcal{I}_0 \equiv 0 \leq i \leq 1 \wedge 0 \leq j \leq 1$. The recursive call is passed $(\text{Prog}_0, \mathcal{I}_0)$ as the unification constraint and generates the next program fragment $\text{Prog}_1(i, j) = 1$ on the input space $\mathcal{I}_1 \equiv 0 \leq i - 2 \leq 2 \wedge 0 \leq j - 2 \leq 2$. Similarly, each further recursive call generates $\text{Prog}_n(i, j) = 1$ on the input space \mathcal{I}_n given by $0 \leq i - 2 * n \leq 2 \wedge 0 \leq j - 2 * n \leq 2$. The sequence of recursive calls do not terminate. To overcome this problem, we use an accelerating widening operator. Intuitively, it generalizes the programs and input spaces in the unification constraints to cover more inputs. In this case, the acceleration operator we define below produces the input space $\mathcal{I}^* \equiv 0 \leq i \wedge 0 \leq j \wedge -2 \leq i - j \leq 2$. Proceeding with this widened constraint lets us terminate with the solution program.

Acceleration. The *accelerating widening operator* ∇ operates on unification constraints. In Algorithm 1, we apply ∇ to the unification constraints being passed to the recursive call on line 14, i.e., we replace the expression $\psi \wedge \text{UnifConstr}(\mathcal{I}_G, \text{Prog})$ with $\nabla(\psi \wedge \text{UnifConstr}(\mathcal{I}_G, \text{Prog}), \beta)$.

While sophisticated accelerating widening operators are available for partial functions (see, for example, [9, 11]), in our implementation, we use a simple one. Given an input unification constraint $\langle (\mathcal{I}_0, \text{Prog}_0), \dots, (\mathcal{I}_n, \text{Prog}_n) \rangle$, the accelerating widening operator works as follows: (a) If $\text{Prog}_n \neq \text{Prog}_j$ for all $j < n$, it returns the input. (b) Otherwise, $\text{Prog}_n = \text{Prog}_j$ for some $j < n$ and we widen the domain where Prog_n is applicable to \mathcal{I}^* where $\mathcal{I}_j \cup \mathcal{I}_n \subseteq \mathcal{I}^*$. Intuitively, we do this by letting $\mathcal{I}^* = \nabla(\mathcal{I}_i, \mathcal{I}_j)$ where ∇ is the widening join operation for convex polyhedra abstract domain [10]. However, we additionally want Prog_n on \mathcal{I}^* to not cause any violation of the learned constraints $\beta = \bigwedge \rho_i$. Therefore, we use a widening operator with bounds on the convex polyhedral abstract domain instead of the generic widening operator. The bounds are obtained from the concrete constraints. We do not describe this procedure explicitly, but present an example below. The final output returned is $\langle (\mathcal{I}_0, \text{Prog}_0), \dots, (\mathcal{I}^*, \text{Prog}_n) \rangle$.

Example 6. Consider the specification $\text{Spec} = f(0) = 1 \wedge (f(x) = 1 \wedge 0 \leq x \leq 10 \implies f(x + 1) = 1) \wedge (f(12) = 0)$. After two recursive calls, we get the unification constraint $\psi = \langle (i = 0, \text{Prog}_0(i) = 1), (i = 1, \text{Prog}_1(i) = 1) \rangle$. Widening, we generalize the input spaces $i = 0$ and $i = 1$ to $\mathcal{I}^* = (i \geq 0)$. However, further synthesis fails due to the clause $f(12) = 0$ from Spec , and we obtain a learned unification constraint $\beta \equiv f(12) = 0$ at the parent call.

We then obtain an additional bound for the unification as replacing f by Prog_1 violates $f(12) = 0$. With this new bound, the widening operator returns the input space $\mathcal{I}^* = (12 > i \geq 0)$, which allows us to complete the synthesis.

Theorem 5. *Algorithm 1 instantiated with the procedures described above is a sound synthesis procedure for conditional linear expressions given by non-separable specifications.*

Experiments. We implemented the above procedure in a tool called RAZORBILL and evaluated it linear integer benchmarks with non-separable specifications from SyGuS competition 2014. For the full summary of results, see the full version [3]. As for the bit-vector benchmarks, on small benchmarks (where both tools finish in less than 1 second), ESOLVER is faster. However, on larger

benchmarks, RAZORBILL can be much faster. As before, we hypothesize that this is due to eSOLVER quickly enumerating small solutions before the STUN based solver can perform any complex theory reasoning.

7 Concluding Remarks

Related work. Algorithmic program synthesis became popular a decade ago with the introduction of CEGIS [17]. Much more recently, syntax-guided synthesis [2] framework, where the input to synthesis is a program space and a specification, was introduced, along with several types of solvers. Our synthesis problem falls into this framework, and our solvers solve SyGuS problem instances. Kuncak et al. [14] present another alternative (non-CEGIS) solver for linear arithmetic constraints.

STUN is a general approach to synthesis. For instance, in the domain of synthesis of synchronization [4, 19, 6, 7, 13], the algorithm used can be presented as an instantiation of STUN. The approach is based on an analysis of a counterexample trace that infers a fix in the form of additional synchronization. The bug fix works for the counterexample and possibly for some related traces. Such bug fixes are then unified similarly as in the STUN approach.

A synthesis technique related to STUN is based on version-space algebras [12, 15]. There, the goal is to compose programs that works on a part of a single input (say a string) to a transformation that would work for the complete single input. In contrast, STUN unifies programs that work for different parts of the input space. The combination of the two approaches could thus be fruitful.

The widening operator has been introduced in [8], and has been widely used in program analysis, but not in synthesis. We proposed to use it to accelerate the process in which STUN finds solutions that cover parts of the input space. Use of other operators such as narrowing is worth investigating.

Limitations. We mentioned that the simple unification operator based on if statements might lead to inefficient code. In particular, if the specification is given only by input-output examples, the resulting program might be a sequence of conditionals with conditions corresponding to each example. That is why we proposed a different unification operator for the bit-vector domain, and we plan to investigate unification further. Furthermore, a limitation of STUN when compared to CEGIS is that designing unification operators requires domain knowledge (knowledge of the given program space).

Future work. We believe STUN opens several new directions for future research. First, we plan to investigate unification operators for domains where the programs have loops or recursion. This seems a natural fit for STUN, because if for several different input we find that the length of the synthesized sequence of instructions in the solution depends on the size of the input, then the unification operator might propose a loop in the unified solution. Second, systems that at runtime prevent deadlocks or other problems can be thought of as finding solutions for parts of the input space. A number of such fixes could then be unified into a more general solution. Last, we plan to optimize the prototype solvers we presented. This is a promising direction, as even our current prototypes have comparable or significantly better performance than the existing solvers.

References

1. SyGuS competition 2014. <http://www.sygus.org/SyGuS-COMP2014.html>
2. Alur, R., Bodík, R., Juniwal, G., Martin, M., Raghothaman, M., Seshia, S., Singh, R., Solar-Lezama, A., Torlak, E., Udupa, A.: Syntax-guided synthesis. In: FMCAD. pp. 1–17 (2013)
3. Alur, R., Černý, P., Radhakrishna, A.: Synthesis through unification. CoRR abs/1505.05868 (2015), <http://arxiv.org/abs/1104.4306>
4. Bloem, R., Hofferek, G., Könighofer, B., Könighofer, R., Außerlechner, S., Spörk, R.: Synthesis of synchronization using uninterpreted functions. In: FMCAD. pp. 35–42 (2014)
5. Bodík, R., Jobstmann, B.: Algorithmic program synthesis: introduction. STTT 15(5-6), 397–411 (2013)
6. Černý, P., Henzinger, T., Radhakrishna, A., Ryzhyk, L., Tarrach, T.: Efficient synthesis for concurrency by semantics-preserving transformations. In: CAV. pp. 951–967 (2013)
7. Černý, P., Henzinger, T., Radhakrishna, A., Ryzhyk, L., Tarrach, T.: Regression-free synthesis for concurrency. In: CAV, pp. 568–584 (2014)
8. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: POPL. pp. 238–252 (1977)
9. Cousot, P., Cousot, R.: An abstract interpretation framework for termination. In: POPL. pp. 245–258 (2012)
10. Cousot, P., Halbwegs, N.: Automatic discovery of linear restraints among variables of a program. In: POPL. pp. 84–96 (1978)
11. Cousot, P., Cousot, R.: Relational abstract interpretation of higher order functional programs (extended abstract). In: JTASPEFT/WSA. pp. 33–36 (1991)
12. Gulwani, S.: Automating string processing in spreadsheets using input-output examples. In: POPL. pp. 317–330 (2011)
13. Gupta, A., Henzinger, T., Radhakrishna, A., Samanta, R., Tarrach, T.: Succinct representation of concurrent trace sets. In: POPL15. pp. 433–444 (2015)
14. Kuncak, V., Mayer, M., Piskac, R., Suter, P.: Complete functional synthesis. In: PLDI. pp. 316–329 (2010)
15. Lau, T., Domingos, P., Weld, D.: Version space algebra and its application to programming by demonstration. In: ICML. pp. 527–534 (2000)
16. Solar-Lezama, A.: Program sketching. STTT 15(5-6), 475–495 (2013)
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. Udupa, A., Raghavan, A., Deshmukh, J.V., Mador-Haim, S., Martin, M.M.K., Alur, R.: TRANSIT: specifying protocols with concolic snippets. In: PLDI. pp. 287–296 (2013)
19. Vechev, M.T., Yahav, E., Yorsh, G.: Abstraction-guided synthesis of synchronization. In: POPL. pp. 327–338 (2010)
20. Warren, H.S.: Hacker’s Delight. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (2002)