

Synthesizing Formal Semantics from Executable Interpreters

ANONYMOUS AUTHOR(S)*

Program verification and synthesis frameworks that allow one to customize the language in which one is interested typically require the user to provide a formally defined semantics for the language. Because writing a formal semantics can be a daunting and error-prone task, this requirement stands in the way of such frameworks being adopted by non-expert users. We present an algorithm that can automatically synthesize inductively defined syntax-directed semantics when given (i) a grammar describing the syntax of a language and (ii) an executable (closed-box) interpreter for computing the semantics of programs in the language of the grammar. Our algorithm synthesizes the semantics in the form of Constrained-Horn Clauses (CHCs), a natural, extensible, and formal logical framework for specifying inductively defined relations that has recently received widespread adoption in program verification and synthesis. The key innovation of our synthesis algorithm is a Counterexample-Guided Synthesis (CEGIS) approach that breaks the *hard* problem of synthesizing a set of constrained Horn clauses into small, tractable expression-synthesis problems that can be dispatched to existing SyGuS synthesizers. Our tool SYNANTIC synthesized inductively-defined formal semantics from 14 interpreters for languages used in program-synthesis applications. When synthesizing formal semantics for one of our benchmarks, SYNANTIC unveiled an inconsistency in the semantics computed by the interpreter for a language of regular expressions; fixing the inconsistency resulted in a more efficient semantics and, for some cases, in a 1.2x speedup for a synthesizer solving synthesis problems over such a language.

1 INTRODUCTION

Recent work on frameworks for program verification and program synthesis has created tools that are parametric in the language that is supported [5, 11, 13]. A user of such a framework must define the language of interest by giving both a syntactic specification and a formal semantic specification of the language. The semantic specification assigns a meaning to each program in the language. However, for most programming languages, and even for simple ones used in program-synthesis applications, it is usually a demanding task to create a *formal* semantics that defines the behaviors of the programs in the language. Obstacles include: (i) the language's semantics might only be documented in natural language, and thus may be ambiguous (or worse, inconsistent), and (ii) the sheer level of detail that is involved in writing such a semantics.

Synthesizing Formal Semantics from Interpreters. In this paper, we propose an alternative approach—based on synthesis—that is applicable to any programming language for which a compiler or interpreter exists. Such infrastructure serves as an operational semantics for the language, albeit one for which anything other than closed-box access would be difficult. Thus, we take closed-box access as a given, and ask the following question:

Is it possible to use an existing compiler or interpreter for a language L to create a formal semantics for L automatically?

In this paper, we assume that the given compiler or interpreter is capable of executing any program or subprogram in language L .

This question is natural, but answering it formally requires one to address two key challenges.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

2475-1421/2018/1-ART1 \$15.00

<https://doi.org/>

50 First, in what formalism should the formal semantics be expressed? The right formalism should
 51 be expressive enough to capture common semantics, yet structured enough to allow synthesis to
 52 be possible. Furthermore, the formalism should not be tied to any specific programming language—
 53 i.e., it should be language-agnostic.

54 Second, how can the synthesis problem be broken down into simple enough small problems
 55 for which one can design a practical approach? The representation of the semantics of most pro-
 56 gramming languages is usually very large, and a monolithic synthesis approach that does not take
 57 advantage of the compositionality of semantics definitions is bound to fail.

58 *Our Approach.* In this paper, we address both of these challenges and present an algorithm that
 59 can automatically synthesize an inductively defined syntax-directed semantics when given (i) a
 60 grammar describing the syntax of the language, and (ii) an executable (closed-box) interpreter for
 61 computing the semantics of programs in the language on given inputs.

62 To address the first of the aforementioned challenges, we choose to synthesize the formal se-
 63 mantics in the form of Constrained Horn Clauses (CHCs), a well-studied fragment of first-order
 64 logic that already provides the foundation of SemGuS [6, 11], a domain- and solver-agnostic frame-
 65 work for defining arbitrary synthesis problems. CHCs can naturally express a big-step operational
 66 semantics, structured as an inductive definition over a language’s abstract syntax, which makes
 67 them appropriate for compositional reasoning.

68 For example, the operational semantics for an assignment to a variable x in an imperative pro-
 69 gramming language can be written as the following CHC:

$$\frac{\llbracket e \rrbracket(s_I) = r_1 \quad s_0 = s_1 \wedge r_0 = s_0[x \mapsto r_1]}{\llbracket x := e \rrbracket(s_0) = r_0}$$

70 The CHC is defined inductively in terms of the semantics of the child term e .

71 To address the second aforementioned challenge, we take advantage of the inductive structure
 72 of CHCs and design a synthesis algorithm that inductively synthesizes the semantics of programs
 73 in the grammar, starting from simple base constructs and moving up to more complex inductively-
 74 defined constructs. For each construct in the language, our algorithm uses a counter-example-
 75 guided inductive synthesis (CEGIS) loop to synthesize the semantic rule—i.e., the CHC—for that
 76 construct. For each construct, we use input-output valuations obtained by calling the closed-box
 77 interpreter to approximate the behavior of its child terms. Such an approximation allows us to
 78 synthesize the semantics construct-by-construct, rather than all at once, which converts the prob-
 79 lem of synthesizing semantics into many smaller problems that only have to synthesize part of the
 80 overall semantics.

81 To evaluate our approach, we implemented it in a tool called SYNANTIC. Our evaluation of
 82 SYNANTIC involved synthesizing the semantics for languages with a wide variety of features, in-
 83 cluding assignments, conditionals, while loops, bit-vector operations, and regular expressions. The
 84 evaluation revealed that our approach not only can help synthesize semantics of non-trivial lan-
 85 guages but can also help debug existing semantics.

86 *Contributions.* Our work makes the following contributions:

- 87 • We introduce a new kind of synthesis problem: the *semantics-synthesis problem* (Section 3).
- 88 • We devise an algorithm for solving semantics-synthesis problems (Section 4). In this algo-
 89 rithm, we harness an example-based program synthesizer (specifically a SyGuS solver) to
 90 synthesize the constraint in each CHC.

- We implement our algorithm in a tool, called SYNANTIC, which also supports an optimization for multi-output productions, i.e., productions whose semantic constraints include multiple output variables (Section 5).
- We evaluate SYNANTIC on a range of different language benchmarks from the program-synthesis literature. For one benchmark, the SYNANTIC-generated semantics revealed an inconsistency in the way the original semantics had been formalized. Fixing the inconsistency in the semantics resulted in a more efficient semantics and a speedup (in some case 1.2x) for a synthesizer solving synthesis problems over such a language (Section 6)

Section 2 illustrates how our algorithm synthesizes the semantics of an imperative while-loop language. Section 7 discusses related work. Section 8 concludes.

2 ILLUSTRATIVE EXAMPLE

Suppose we have an imperative language IMP (cf. Example 2.1), and we want to synthesize an IMP program to automate some programming tasks. However, we find that there is no synthesizer for IMP. Next, we remember that the SemGuS framework can synthesize programs for user-defined languages [11]. However, when we start trying to write our synthesis problem using SemGuS, we quickly realize IMP's semantics is defined using an interpreter and we do not have a formal-logic-based semantics (i.e., a set of Constrained Horn Clauses), which SemGuS requires. At this point we are stuck, and cannot synthesize any IMP program unless we manually write the IMP semantics as a set of CHCs, a tedious and error-prone task.

In this paper, we consider the problem of synthesizing a formal (logical) semantics for a language from an executable interpreter. We use the IMP language described in Example 2.1 as a running example in this section to illustrate our algorithm (and return to it in Section 4).

Example 2.1 (Syntactic Definition of IMP). Consider the grammar G_{IMP_n} that defines the syntax of IMP for programs with n variables x_1, \dots, x_n :

$$\begin{aligned}
 S &::= x_1 := E \mid \dots \mid x_n := E \mid S ; S \mid \text{ite } B S S \mid \text{while } B \text{ do } S \\
 &\quad \mid \text{do } S \text{ while } B \mid \text{repeat } S \text{ until } B \\
 B &::= \text{false} \mid \text{true} \mid \neg B \mid B \wedge B \mid B \vee B \mid E < E \\
 E &::= 0 \mid 1 \mid x_1 \mid \dots \mid x_n \mid E + E \mid E - E
 \end{aligned}$$

The IMP language consists of arithmetic and Boolean expressions, statements for assignment to the variables x_1 through x_n , sequential composition, if-then-else, and various looping constructs. IMP also comes equipped with an executable interpreter \mathcal{I}_{IMP} that assigns to each term $t \in \mathcal{L}(G)$ its standard (denotational) semantics (e.g., arithmetic and Boolean expressions are evaluated as in linear integer arithmetic, $x_1 := e$ takes as input a state, and outputs the input state with x_i 's value updated by the result of evaluating e , etc.).

Suppose that we did not know the semantics of IMP *a priori*; that is, suppose that we only have access to the interpreter \mathcal{I}_{IMP} . How can we synthesize a formal semantics for each program in G_{IMP} using the interpreter? A naïve approach would randomly generate a large set of terms and inputs, and try to learn a function mapping inputs to outputs for each term. However, this approach would only provide a semantics for the enumerated terms, and fails to generalize to the entire language. A less naïve approach might attempt to form a monolithic synthesis problem to synthesize a semantic function for each production of the grammar that satisfies a set of generated example terms and input-output pairs. However, it is known that synthesizers scale exceptionally poorly in the size of the desired output [3], even for IMP_1 , which has only 17 productions, this approach would be practically impossible.

148 *Nullary productions.* One of the key innovations of our approach is that we synthesize the se-
 149 mantics on a per-production basis, i.e., working one production at a time. We start by synthesizing
 150 a semantics for nullary (leaf) productions. For IMP_1 , this means we synthesize a semantics for the
 151 productions `0`, `1`, `x1`, `false`, and `true` before we synthesize the semantics of any other productions.
 152 For a nullary production p , we synthesize a semantics of the form:

$$153 \quad \frac{x_0^{\text{out}} = f(x_0^{\text{in}})}{\text{Sem}(p, x_0^{\text{in}}, x_0^{\text{out}})}$$

154
 155
 156
 157 which states that, because the term p has no sub-terms, the output is only a function of the input
 158 x_0^{in} . In our approach, we use a Counter-Example-Guided Synthesis (CEGIS) approach to synthe-
 159 size a function f that captures the behavior of \bar{I}_{IMP} on production p . Within the CEGIS loop, we
 160 synthesize a candidate function f , then verify if it is consistent with \bar{I}_{IMP} (e.g., on a larger number
 161 of inputs x_0^{in}). If f is consistent, then we have successfully learned the semantics of p ; otherwise,
 162 the verifier generates a counter-example and a new candidate semantic function f .

163
 164 *Inductively synthesizing semantics.* Next, our approach synthesizes the semantics for other arith-
 165 metic and Boolean expressions. In this step, we inductively synthesize the semantics of productions
 166 by reusing the semantics of previously learned productions to learn the semantics of new produc-
 167 tions. At this point, we may assume that we know the semantics of all nullary productions. For
 168 instance, suppose that we wish to next learn the semantics of `+`. At first, our algorithm generates
 169 examples favoring terms like `1 + 1`, `x + 1`, etc. that contains sub-terms whose semantics have al-
 170 ready been learned. For $t_1 + t_2$, our algorithm generates a semantics that can rely on the semantics
 171 of its sub-terms t_1 and t_2 . Specifically, the semantics of $t_1 + t_2$ takes the following form:

$$172 \quad \frac{\begin{array}{ccc} \text{sem}(t_1, x_1^{\text{in}}, x_1^{\text{out}}) & \text{sem}(t_2, x_2^{\text{in}}, x_2^{\text{out}}) & \\ x_1^{\text{in}} = f_1(x_0^{\text{in}}) & x_2^{\text{in}} = f_2(x_0^{\text{in}}, x_1^{\text{out}}) & x_0^{\text{out}} = f_0(x_0^{\text{in}}, x_1^{\text{out}}, x_2^{\text{out}}) \end{array}}{\text{sem}(t_1 + t_2, x_0^{\text{in}}, x_1^{\text{out}})}$$

173
 174
 175
 176
 177 which states that the semantics of $t_1 + t_2$ is inductively defined in terms of the semantics of t_1
 178 and the semantics of t_2 . The semantics enforces a left-to-right evaluation order:¹ the rule expresses
 179 that the input to t_1 , x_1^{in} , is a function of $t_1 + t_2$'s input, x_0^{in} , and similarly that t_2 's input, x_2^{in} , is
 180 a function of $t_1 + t_2$'s input, x_0^{in} , and t_1 's output, x_1^{out} . Finally, it also expresses that the $t_1 + t_2$'s
 181 output, x_0^{out} , is a function of its input, x_0^{in} , and the outputs of t_1 (x_1^{out}) and t_2 (x_2^{out}).

182 When the semantics of a sub-term t_i is known (e.g., for nullary productions), we substitute its
 183 learned semantics for $\text{sem}(t_i, x_i^{\text{in}}, x_i^{\text{out}})$; otherwise, we approximate its semantics using examples.
 184 Again, we use a CEGIS loop to generate examples for the entire term $t_1 + t_2$, as well as any sub-
 185 terms whose exact semantics have not yet been synthesized (e.g., for a sub-term that uses `+` or `-`).
 186 The process proceeds analogously for most other productions in IMP .

187
 188 *Semantically recursive productions.* The final interesting case is for `while` loops, for
 189 which the semantics is recursive on the term itself. For semantically recursive produc-
 190 tions, we assume that the semantics can make a recursive call (i.e., effectively acting
 191 as if the term itself is a sub-term). We additionally synthesize a predicate determin-
 192 ing if the recursive call should be made or not. For `while b do s`, we synthesize two

193
 194
 195 ¹We show how to overcome this restriction in Section 5.1.

semantic rules, one in which the recursive call is made, and one in which it is not.

$$\frac{\text{sem}(\mathbf{b}, x_1^{\text{in}}, x_1^{\text{out}}) \quad \text{sem}(\mathbf{s}, x_2^{\text{in}}, x_2^{\text{out}}) \quad \neg \text{Pred}_{\text{rec}}(x_0^{\text{in}}, x_1^{\text{out}}, x_2^{\text{out}})}{x_1^{\text{in}} = f_1(x_0^{\text{in}}) \quad x_2^{\text{in}} = f_2(x_0^{\text{in}}, x_1^{\text{out}}) \quad x_0^{\text{out}} = f_0(x_0^{\text{in}}, x_1^{\text{out}}, x_2^{\text{out}})} \\ \text{sem}(\text{while } \mathbf{b} \text{ do } \mathbf{s}, x_0^{\text{in}}, x_1^{\text{out}})$$

$$\frac{\text{sem}(\mathbf{b}, x_1^{\text{in}}, x_1^{\text{out}}) \quad \text{sem}(\mathbf{s}, x_2^{\text{in}}, x_2^{\text{out}}) \quad \text{sem}(\text{while } \mathbf{b} \text{ do } \mathbf{s}, x_3^{\text{in}}, x_3^{\text{out}}) \quad \text{Pred}_{\text{rec}}(x_0^{\text{in}}, x_1^{\text{out}}, x_2^{\text{out}})}{x_1^{\text{in}} = f_1(x_0^{\text{in}}) \quad x_2^{\text{in}} = f_2(x_0^{\text{in}}, x_1^{\text{out}}) \quad x_3^{\text{in}} = f_2(x_0^{\text{in}}, x_1^{\text{out}}, x_2^{\text{out}}) \quad x_0^{\text{out}} = f_0(x_0^{\text{in}}, x_1^{\text{out}}, x_2^{\text{out}}, x_3^{\text{out}})} \\ \text{sem}(\text{while } \mathbf{b} \text{ do } \mathbf{s}, x_0^{\text{in}}, x_1^{\text{out}})$$

As with the previous productions, our algorithm uses a CEGIS loop to synthesize a candidate semantics of the above form, verify its correctness, and generate a counter-example if the candidate semantics is incorrect. While we may employ learned semantics for sub-terms, recursive calls to a sub-term must be approximated using examples because we are still in the process of learning its semantics. We formally define the semantics-synthesis problem that we solve in Section 3 and explain how our synthesis algorithm works in Section 4.

Multi-output productions. In the above **while**-loop example, we saw that the function f_0 had four inputs that must be considered when synthesizing a term to instantiate f_0 . As the number of input variables and the size of the desired result grows, synthesis scales poorly. In the above examples, the notation is not showing the full picture. For IMP_n all input and (most) output variables are an n -tuple of variables representing a state of an IMP_n program. Even for just IMP_2 , f_0 has twice as many inputs.

To address this problem, we allow synthesizing the semantics of each output of a production independently. For example, consider the production $\mathbf{x}_0 := \mathbf{t}$ (for IMP_2). We generate a semantics using two constraints F and G , independently. The constraint F (resp. G) represents the pair of functions f_0 and f_1 (resp. g_0 and g_1).

$$\frac{\text{sem}(\mathbf{t}, x_1^{\text{in}}, x_1^{\text{in}}) \quad x_1^{\text{in}} = f_1(x_0^{\text{in}}) \quad x_0^{\text{out}} = f_0(x_0^{\text{in}}, x_1^{\text{out}})}{\text{sem}(\mathbf{x}_0 := \mathbf{t}, x_0^{\text{in}}, x_0^{\text{out}})} F$$

$$\frac{\text{sem}(\mathbf{t}, x_1^{\text{in}}, x_1^{\text{in}}) \quad x_1^{\text{in}} = g_1(x_0^{\text{in}}) \quad x_0^{\text{out}} = g_0(x_0^{\text{in}}, x_1^{\text{out}})}{\text{sem}(\mathbf{x}_0 := \mathbf{t}, x_0^{\text{in}}, x_0^{\text{out}})} G$$

By independently synthesizing F and G , we reduce the burden on the underlying synthesizer; however, now the synthesizer is allowed to return an F and G for which $f_1 \neq g_1$. Thus, F and G have inconsistent inputs being provided to the child-term t . We use an SMT solver to determine if f_1 and g_1 are consistent for each of the example inputs to the term $\mathbf{x}_0 := \mathbf{t}$. If so, we will return either f_0, g_0, f_1 (or f_0, g_0, g_1) because f_1 and g_1 are consistent on all examples—i.e., when evaluated on the same example they return equal outputs—otherwise, we discover that f_1 and g_1 are inconsistent on some input and add a new constraint to ensure that the same pair of functions f_1 and g_1 cannot be synthesized again. This optimization is further discussed in Section 5.3.

3 PROBLEM DEFINITION

In this paper, we consider the problem of synthesizing a formal logical semantics for a deterministic language from an executable interpreter. While there are many possible ways to logically define a semantics, we are interested in an approach that is language-agnostic and inductive. The SemGuS synthesis framework has proposed using Constrained Horn Clauses as a way of defining program semantics that meets both of our desiderata. Concretely, SemGuS already supports synthesis for a large number of languages (which we consider in our experimental evaluation) by allowing a user to provide a user-defined semantics. As mentioned above, in SemGuS, semantics are defined

inductively on the structure of the grammar (i.e., per production/language construct) using logical relations represented as Constrained Horn Clauses (CHCs) [11]. In this paper, we follow suit and address the problem of learning a semantics of this form from an executable interpreter for the given language. This section formalizes the semantics-synthesis problem that we consider. We begin by detailing our representation of syntax (Section 3.1), interpreters (Section 3.2), semantics (Section 3.3), and semantic-equivalence oracles (Section 3.4). Finally, we formalize the semantics-synthesis problem in Section 3.4.

3.1 Syntax

We consider languages represented as regular tree grammars (RTGs). A **ranked alphabet** is a tuple $\langle \Sigma, rk_\Sigma \rangle$ that consists of a finite set of symbols Σ and a function $rk_\Sigma : \Sigma \rightarrow \mathbb{N}$ that associates every symbol with a rank (or arity). For any $n \geq 0$, $\Sigma^n \subseteq \Sigma$ denotes the set of symbols of rank n . The set of all (ranked) Trees over Σ is denoted by T_Σ . Specifically, T_Σ is the least set such that $\Sigma^0 \subseteq T_\Sigma$ and if $\sigma^k \in \Sigma^k$ and $t_1, \dots, t_k \in T_\Sigma$, then $\sigma^k(t_1, \dots, t_k) \in T_\Sigma$. In the remainder of the paper, we assume a fixed ranked alphabet $\langle \Sigma, rk_\Sigma \rangle$.

A **typed regular tree grammar** (RTG) is a tuple $G = \langle N, \Sigma, \delta, T, \theta, \tau \rangle$, where N is a finite set of non-terminal symbols of rank 0, Σ is a ranked alphabet, δ is a set of productions over a set of types T , and for each non-terminal $A \in N$, and θ_A (resp. τ_A) assigns A an input-type (resp. output-type) from T . Each production in δ takes the form:

$$A_0 \rightarrow \sigma(A_1, A_2, \dots, A_{rk_\Sigma(\sigma)})$$

where $A_i \in N$ and $\sigma \in \Sigma$. We use $\mathcal{L}(A)$ to denote the language of non-terminal A and $\delta(A)$ the set of all productions associated with A (i.e., all productions where A_0 is A). In the remainder, we assume a fixed grammar $G = \langle N, \Sigma, \delta, T, \theta, \tau \rangle$.

Example 3.1 (G_{IMP} as a Regular Tree Grammar). Consider the IMP language detailed in Section 2, G_{IMP} is a regular tree grammar that has been stylized to ease readability. For example, the non-terminals consist of the rank-0 symbols E , B , and S . The productions include $S \rightarrow x_1 := (E)$, $S \rightarrow ;(S, S)$, and $S \rightarrow \text{while}(B, S)$. For IMP_2 (IMP with two variables x_1 and x_2), θ_E is the type $\mathbb{Z} \times \mathbb{Z}$, representing the state of the two variables, and τ_E is \mathbb{Z} , representing the return type of arithmetic expressions.

3.2 Interpreters

We consider a class of deterministic executable interpreters—i.e., a program evaluator for which we may only observe input-output behavior.

Definition 3.2 (Interpreter). Formally, an *interpreter* for G maps each non-terminal $A \in N$ to a partial function $\mathcal{I}_A : (\mathcal{L}(A) \times \theta_A) \rightarrow \tau_A$ —with the interpretation that the interpreter maps a program $t \in \mathcal{L}(A)$ and input value $in \in \theta_A$ to some output $out \in \tau_A$ if and only if t starting with the input value in terminates with the output value out .

Example 3.3 (Interpreters for IMP_1). Recall the IMP language defined in Section 2. The interpreter \mathcal{I} for IMP consists of three base interpreters \mathcal{I}_E , \mathcal{I}_B , and \mathcal{I}_S , which are used to evaluate arithmetic expressions, Boolean expressions, and statements, respectively. Throughout this paper, we assume the interpreters for IMP_1 (and all IMP variants) evaluate according to the standard denotational semantics (e.g., 0 is the expression that always returns 0 regardless of input state; $+$ is mathematical $+$; **while b s** evaluates b , executes the loop body, and recurses if b evaluates to *true* and otherwise immediately terminates; etc.).

3.3 Semantics

We represent the big-step semantics of a language (defined by some grammar G) using a set of Constrained Horn Clauses (CHCs) within some background theory \mathcal{T} per production. While CHCs (at first glance) seem limiting, this formulation of semantics has been employed by the SemGuS framework to represent user-defined semantics for many languages [6, 11], including many variations of IMP, regular expressions, SyGuS expressions within the theory of bit vectors, algebraic data types, linear integer arithmetic.

Definition 3.4 (Constrained Horn Clause). A CHC (in theory \mathcal{T}) is a first-order formula of the form:

$$\forall \bar{x}_1, \dots, \bar{x}_n, \bar{x}. \phi \wedge R_1(\bar{x}_1) \wedge \dots \wedge R_n(\bar{x}_n) \Rightarrow H(\bar{x})$$

where R_1, \dots, R_n and H are uninterpreted relations, $\bar{x}_1, \dots, \bar{x}_n$ and \bar{x} are variables, and ϕ is a quantifier-free \mathcal{T} -constraint over the variables.

To specify the big-step semantics of a non-terminal $A \in N$ (for which the interpreter has type $\mathcal{I}_A : (\mathcal{L}(A) \times \theta_A) \rightarrow \tau_A$), we introduce the semantic relation $Sem_A(t_A, x_A^{in}, x_A^{out})$, where t_A is a variable representing elements of $\mathcal{L}(A)$, x_A^{in} is a variable of type θ_A , and x_A^{out} is a variable of type τ_A . Throughout this paper, we may also use $\llbracket t_A \rrbracket_{Sem}(x_A^{in}) = x_A^{out}$ to denote that $Sem_A(t_A, x_A^{in}, x_A^{out})$ holds.

Example 3.5 (Semantic relations). Consider the IMP₁ language introduced in Section 2; a semantics for IMP₁ uses the semantic relations:

$$Sem_E : \mathcal{L}(E) \times \mathbb{Z} \times \mathbb{Z} \rightarrow \text{bool} \quad Sem_B : \mathcal{L}(B) \times \mathbb{Z} \times \text{bool} \rightarrow \text{bool} \quad Sem_S : \mathcal{L}(S) \times \mathbb{Z} \times \mathbb{Z} \rightarrow \text{bool}$$

While CHCs are quite general and capable of defining both deterministic and non-deterministic semantics, we limit our scope to CHCs that represent deterministic semantics. Furthermore, for a grammar G , we assume that each production $A_0 \rightarrow \sigma(A_1, \dots, A_n) \in G$ evaluates sub-terms in a fixed order from left to right (i.e., for a term $p(t_1, \dots, t_n)$ sub-term t_1 is evaluated before t_2 , etc.). While this imposed order may seem too restrictive, we later show how this restriction can be lifted by considering all permutations of sub-terms.

Definition 3.6 (Semantic Rule, Semantic Constraint). Given a production $A_0 \rightarrow p(A_1, \dots, A_n)$ a **semantic rule** for p is a CHC of the form:

$$\frac{Sem_{A_1}(t_1, x_1^{in}, x_1^{out}) \quad \dots \quad Sem_{A_n}(t_n, x_n^{in}, x_n^{out}) \quad F(x_0^{in}, \dots, x_n^{in}, x_0^{out}, \dots, x_n^{out})}{Sem_{A_0}(p(t_1, \dots, t_n), x_0^{in}, x_0^{out})} \quad (1)$$

where F is constraint over theory \mathcal{T} , which we call a **semantic constraint**, that takes the form:

$$x_1^{in} = f_1(x_0^{in}) \wedge \dots \wedge x_n^{in} = f_n(x_1^{out}, \dots, x_{n-1}^{out}, x_0^{in}) \wedge x_0^{out} = f_0(x_1^{out}, \dots, x_n^{out}, x_0^{in}) \wedge P(x_0^{in}, x_0^{out}, \dots, x_n^{out}) \quad (2)$$

where each f_i is a function that returns a term of type θ_{A_i} for $i > 0$ and τ_{A_0} for $i = 0$. The semantic constraint also includes predicate $P(x_{A_0}^{in}, x_{A_1}^{out}, \dots, x_{A_n}^{out})$ that determines when the semantic rule is valid (e.g., for conditionals and loops).

344 *Example 3.7 (Semantics of do_while).* We give the semantics of the `do_while` IMP statement
 345 below:

$$\begin{array}{c}
 \frac{\llbracket s \rrbracket(x_1) = x'_1 \quad \llbracket b \rrbracket(x_2) = r_b}{\llbracket \text{do } s \text{ while } b \rrbracket(x_3) = x'_3} \quad \frac{r_b \quad x_1 = x_0 \quad x_2 = x'_1 \quad x_3 = x'_1 \quad x'_0 = x'_3}{\llbracket \text{do } s \text{ while } b \rrbracket(x_0) = x'_0} \\
 \frac{\llbracket s \rrbracket(x_1) = x'_1 \quad \llbracket b \rrbracket(x_2) = r_b \quad \neg r_b \quad x_1 = x_0 \quad x_2 = x'_1 \quad x'_0 = x'_1}{\llbracket \text{do } s \text{ while } b \rrbracket(x_0) = x'_0}
 \end{array}$$

350 The first rule executes the statement s and then, if the guard b is true recursively executes the
 351 whole loop and returns the resulting value. The second rule executes the statement s and then, if
 352 the guard b is false returns the output produced when executing the statement s .

349 3.4 Equivalence Oracle and Semantics Synthesis Problem

358 For a grammar G , a semantics Sem for G , and an interpreter \mathcal{I} for G , we define when Sem is
 359 equivalent to the semantics defined by interpreter \mathcal{I} via an *equivalence oracle*.

360 *Definition 3.8 (Equivalent, Equivalence Oracle).* Given an interpreter \mathcal{I} for a language G , a sub-
 361 grammar $G' \subseteq G$, and a semantics Sem for G' , we say that \mathcal{I} and Sem are **equivalent** on G' if and
 362 only if for every term $t \in \mathcal{L}(G')$, input $in \in \theta_A$, and output $out \in \tau$, we have:

$$364 \quad \mathcal{I}(t, in) = out \Leftrightarrow \llbracket t \rrbracket_{Sem}(in) = out$$

365 An **equivalence oracle** \mathcal{E} for \mathcal{I} is a function that takes as input a semantics Sem for G' and
 366 determines if Sem is equivalent to \mathcal{I} on G' . If Sem is not equivalent to \mathcal{I} , then \mathcal{E} returns an ex-
 367 ample $\langle in, t, out \rangle$ for which \mathcal{I} and Sem disagree—i.e., there is some term t and input in such that
 368 $\llbracket t \rrbracket_{Sem}(in) \neq \llbracket t \rrbracket_{\mathcal{I}}(in)$ —and otherwise returns *None* when Sem and \mathcal{I} are equivalent.

369 Given a language (a grammar and accompanying interpreter), the semantics synthesis problem
 370 is to find some semantics of the language that is equivalent to the interpreter. We formalize the
 371 semantics synthesis problem as follows:

372 *Definition 3.9 (Semantics-Synthesis Problem, Solution).* A **semantics-synthesis problem** is a
 373 tuple $\mathcal{P} \triangleq \langle G, \mathcal{I}, \mathcal{E} \rangle$, where G is a grammar, \mathcal{I} is an interpreter for G , and \mathcal{E} is an equivalence
 374 oracle for \mathcal{I} . A **solution** to the semantics-synthesis problem \mathcal{P} is a semantics Sem for G that is
 375 equivalent to \mathcal{I} as determined by \mathcal{E} .

376 4 SEMANTICS SYNTHESIS

377 This section presents an algorithm SEMSYNTH (Algorithm 1) to synthesize a semantics for a lan-
 378 guage from an executable interpreter. The input to SEMSYNTH is a semantics-synthesis problem
 379 consisting of (i) a grammar G , (ii) an executable interpreter \mathcal{I} for G , and (iii) an equivalence or-
 380 cle \mathcal{E} for \mathcal{I} . Upon termination, SEMSYNTH returns a semantics Sem for G that is equivalent to the
 381 executable interpreter \mathcal{I} as determined by the equivalence oracle \mathcal{E} .

382 Synthesizing a semantics for arbitrary languages comes with several challenges. In general, se-
 383 mantics are defined as complex recursively defined functions that provide an interpretation to
 384 every program within the language. Trying to directly synthesize such a semantics is already im-
 385 practical for relatively small languages, such as the IMP language defined in Example 2.1.

386 As described in Section 3.3, we consider semantics represented using logical relations defined by
 387 a set of Constrained Horn Clauses per production of G (cf. Definition 3.6). By formulating the de-
 388 sired semantics as CHCs per production, SEMSYNTH can inductively synthesize the semantics of G

Algorithm 1: Semantics-Synthesis Algorithm

```

393 Algorithm 1: Semantics-Synthesis Algorithm
394
395 1 Procedure SEMSYNTH ( $G, \mathcal{I}, \mathcal{E}$ )
396 2    $Order \leftarrow$  Some total ordering of productions of  $G$ ;
397 3    $N \leftarrow |G|$ ; //  $N$  the number of productions of  $G$ .
398   // Assume, each production is indexed 1 to  $N$  according to  $Order$ 
399 4    $Sem \leftarrow \lambda p. \perp$ ; // Maps each production to a semantics.
400 5    $E \leftarrow \lambda p. \emptyset$ ; // Maps each production to a set of examples.
401 6    $i \leftarrow 1$ ; // Start from production indexed by 1.
402 7   while  $i \leq N$  do // Synthesize semantics one production at a time according to  $Order$ .
403 8      $Sem[p_i] \leftarrow$  SYNTHSEMANTICCONSTRAINT( $Sem, p_i, E$ );
404 9      $cex, p_j \leftarrow$  VERIFY( $Sem, p_i, \mathcal{I}, \mathcal{E}$ );
405 10    if  $cex \neq None$  then // Counter-example found for production  $p_j$ .
406 11       $E[p_j] \leftarrow E[p_j] \cup \{cex\}$ ; // Update examples for production  $p_j$ .
407 12       $i \leftarrow j$ ; // Backtrack and resynthesize production  $p_j$ 's semantics.
408 13    else
409 14       $i \leftarrow i + 1$ ; // Proceed to synthesize next production's semantics.
410 15    return  $Sem$ ;

```

one production at a time. Furthermore, SEMSYNTH uses this flexibility to synthesize the semantics of simpler productions and fragments of the G before synthesizing the semantics of more complex productions/operators. Finally, by fixing the shape of the semantics (i.e., as a set of CHCs per production), SEMSYNTH reduces the monolithic synthesis problem to a series of first-order synthesis problems—specifically, by using a SyGuS or sketch-based synthesizer to synthesize the constraint of each semantic rule (CHC) defining the semantics of a production.

The remainder of this section is structured as follows: Section 4.1 provides a high-level overview of how SEMSYNTH solves semantic-synthesis problems, Sections 4.2 and 4.4 provide specifications for SYNTHSEMANTICCONSTRAINT and VERIFY, which synthesize semantic constraints from examples and verify candidate semantic constraints against the interpreter, respectively. Section 4.3 details how SEMSYNTH synthesizes the semantics of productions from examples. Finally, Section 4.5 explains how SEMSYNTH handles semantically recursive productions.

4.1 Overview of SEMSYNTH

SEMSYNTH (Algorithm 1) uses the counter-example-guided synthesis (CEGIS) paradigm to synthesize a semantics for G that is equivalent to \mathcal{I} according to the equivalence oracle \mathcal{E} . Throughout this section, we will use the IMP language from Example 2.1 to illustrate how SEMSYNTH operates.

Choosing an Order. To begin, SEMSYNTH determines an order to iterate over the productions of G . While any ordering is sound, we assume that SEMSYNTH picks an order $Order$ such that the following property holds. Suppose the grammar G is treated as a graph whose nodes are the productions and non-terminals of G and each production $A_0 \rightarrow p(A_1, \dots, A_n)$ induces an edge from A_0 to p and from p to each A_1 through A_n . If there is a path from p_i to p_j , but no path from p_j to p_i then p_i is ordered before p_j by $Order$.

For example, for the IMP language, $Order$ will order productions for arithmetic expressions, like `0`, `1`, and `x`, before the production `+`, and all productions for arithmetic expressions before the assignment-statement production `x :=`. Mutually recursive productions (e.g., sequencing `;`, `ite`, and `while`), may be explored in any order. For simplicity, we index each production by a natural

number from 1 to N (the number of productions in G). Next, SEMSYNTH initializes the synthesized semantics Sem (marking every production as initially undefined), and E to the example set (cf. definition 4.1) that maps each production to a set of examples.

Synthesizing a Candidate Semantics. After initialization, SEMSYNTH iteratively synthesizes the semantics of productions in the order defined by *Order*. SEMSYNTH employs a CEGIS loop to synthesize the semantics of each production. During each iteration, SEMSYNTH first synthesizes a candidate semantic constraint (cf. Definition 3.6) for production p_i using SYNTHSEMANTICCONSTRAINT. The procedure SYNTHSEMANTICCONSTRAINT returns some semantic constraint for p_i that satisfies the set of examples E . Section 4.2 provides a formal specification of SYNTHSEMANTICCONSTRAINT's operation.

SEMSYNTH then uses the procedure VERIFY to determine if the semantics synthesized thus far is consistent with the interpreter \mathcal{I} as determined by the equivalence oracle \mathcal{E} . A formal specification of VERIFY is provided in Section 4.4. If VERIFY determines that Sem is correct, then it returns *None*, and SEMSYNTH advances to attempt to synthesize the semantics of production p_{i+1} . If VERIFY determines that Sem does not match the semantics of \mathcal{I} , then it returns some example cex for which Sem and \mathcal{I} disagree. It additionally places blame on some production p_j 's synthesized semantics. SEMSYNTH then updates the set of examples E and backtracks to resynthesize the semantics of p_j .

4.2 Specification of SYNTHSEMANTICCONSTRAINT

Before formally specifying SYNTHSEMANTICCONSTRAINT (Section 4.2.3), we first define example sets (Section 4.2.1) and when a semantic constraint is consistent with an example set (Section 4.2.2).

4.2.1 Example Sets. For an interpreter \mathcal{I} , an example set E is a set of examples consistent with \mathcal{I} .

Definition 4.1 (Example set for interpreter \mathcal{I}). Given an interpreter \mathcal{I} for grammar G , an example set E for interpreter \mathcal{I} maps each production $A_0 \rightarrow p(A_1, \dots, A_n) \in G$ to a finite set of examples of the form $\langle in, p(t_1, \dots, t_n), out \rangle$, where $t_i \in L(A_i)$ and $\mathcal{I}(p(t_1, \dots, t_n), in) = out$.

Example 4.2 (Example set for IMP_1). Recall the interpreter \mathcal{I}_{IMP_1} described in Example 3.3 for language IMP_1 . An example set E for \mathcal{I}_{IMP_1} might include the examples $\langle 0, 0, 0 \rangle$, $\langle 1, 0, 0 \rangle$, $\langle 1, x := 0; x := x + 4, 4 \rangle$, and $\langle 10, \text{while } 0 < x \text{ do } x := x - 1, 0 \rangle$; however, an example set for \mathcal{I}_{IMP} could not include any example of the form $\langle n, \text{while } 0 < x \text{ do } x := x + 1, n' \rangle$ where n (the initial value of x) is some non-negative number. Since, $\text{while } 0 < x \text{ do } x := x + 1$ would not terminate on the input n . The example $\langle n, \text{while } 0 < x \text{ do } x := x + 1, n' \rangle$ would violate the assumption that E only contains examples consistent with the interpreter \mathcal{I}_{IMP} .

4.2.2 Example Consistency. In SEMSYNTH, we use the example set E to ensure that the semantic constraint returned by SYNTHSEMANTICCONSTRAINT is consistent with \mathcal{I} for at least the examples appearing in E .

Definition 4.3 (Consistency with Example Set). Given a production $A_0 \rightarrow p(A_1, \dots, A_n)$, a semantic rule R with semantic constraint F of the form defined in Definition 3.6, and example set E , we say R is **consistent** with E if and only if the semantic constraint F is consistent with E . Furthermore, the semantic constraint F is **consistent** with the example set E if for every example $\langle in_{A_0}, p(t_1, \dots, t_n), out_{A_0} \rangle \in E$ the following condition holds:

$$\forall x_0^{in}, \dots, x_n^{in}, x_0^{out}, \dots, x_n^{out} \cdot \left(\begin{array}{l} x_0^{in} = in_0 \\ \wedge \text{Summary}(t_1) \\ \dots \\ \wedge \text{Summary}(t_n) \\ \wedge F \end{array} \right) \Rightarrow x_0^{out} = out_0 \quad (3)$$

where $\text{Summary}(t_i) = \bigvee \{x_i^{\text{in}} = \text{in}_i \wedge x_i^{\text{out}} = \text{out}_i : \langle \text{in}_i, t_i, \text{out}_i \rangle \in E\}$ summarizes the semantics of t_i according to the examples found in E .

Example 4.4 (Example Consistency). Consider the production for the operator $+$, and the (correct) semantic constraint $F \triangleq x_1^{\text{in}} = x_0^{\text{in}} \wedge x_2^{\text{in}} = x_0^{\text{in}} \wedge x_0^{\text{out}} = x_1^{\text{out}} + x_2^{\text{out}}$; F is consistent with the examples $\langle 0, \mathbf{x}_0 + \mathbf{1}, 1 \rangle$, $\langle 0, \mathbf{x}_0, 0 \rangle$, and $\langle 0, \mathbf{1}, 1 \rangle$. Specifically, the following formula is valid:

$$\forall x_0^{\text{in}}, x_1^{\text{in}}, x_2^{\text{in}}, x_0^{\text{out}}, x_1^{\text{out}}, x_2^{\text{out}}. (x_0^{\text{in}} = 0 \wedge (x_1^{\text{in}} = 0 \wedge x_1^{\text{out}} = 0) \wedge (x_1^{\text{in}} = 0 \wedge x_1^{\text{out}} = 1) \wedge F) \Rightarrow x_0^{\text{out}} = 1.$$

4.2.3 Formal Specification of SYNTHSEMANTICCONSTRAINT. The procedure SYNTHSEMANTICCONSTRAINT takes as input the current semantics Sem , the production p_i whose semantics is to be synthesized, and the current example set E ; it returns a constraint F —of the form defined in Definition 3.6—defining a semantics for production p_i that is consistent with the example set E .

Example 4.5 (Synthesizing semantics of $\mathbf{x} :=$ consistent with examples). Recall that for the language IMP, the semantics of the production $\mathbf{x} :=$ is represented as (a set of) CHC rule(s) of the form:

$$\frac{\text{Sem}_E(e, x_1^{\text{in}}, x_0^{\text{out}}) \wedge x_1^{\text{in}} = f(x_0^{\text{in}}) \wedge x_0^{\text{out}} = g(x_0^{\text{in}}, x_1^{\text{out}})}{\text{Sem}_S(\mathbf{x} := e, x_0^{\text{in}}, x_0^{\text{out}})}$$

for some functions f and g (in the theory of linear integer arithmetic). The procedure call SYNTHSEMANTICCONSTRAINT($\text{Sem}, \mathbf{x} :=, E$) synthesizes the formulas $f(x_0^{\text{in}}) = t_f$ and $g(x_0^{\text{in}}, x_1^{\text{out}}) = t_g$, and returns the constraint $F \triangleq x_1^{\text{in}} = t_f \wedge x_0^{\text{out}} = t_g$ so that F is consistent with E .

We note that for functions expressible in a decidable first-order theory, this problem can be exactly encoded as a Syntax-Guided Synthesis (SyGuS) problem [2] and solved by a SyGuS solver (e.g., cvc5 [4]).

4.3 Synthesizing from Examples

To synthesize the semantics of a production, SEMSYNTH uses the counter-example guided inductive synthesis (CEGIS) paradigm. At a high level, SEMSYNTH synthesizes a candidate semantics—for productions p_1 through p_i , one production at a time—from a set of examples. If the candidate semantics is incorrect, a counter-example is produced (by VERIFY), which is added to the set of examples, and a new candidate semantics is synthesized. We illustrate how SEMSYNTH uses and generates examples to synthesize the semantics of a nullary production in Example 4.6.

4.4 Specification of VERIFY

The procedure VERIFY takes as input the currently synthesized semantics Sem , the production p_i whose semantics was just synthesized, the interpreter \mathcal{I} , and equivalence oracle \mathcal{E} ; it determines if Sem is equivalent to the interpreter \mathcal{I} for all terms in the sub-grammar G' that consists of only the productions p_1 through p_i . If VERIFY determines that Sem is not equivalent to \mathcal{I} , it returns a counter-example (i, t, o) and production p_j (with $1 \leq j \leq i$) such that t 's root production is p_j , $\llbracket t \rrbracket_{\mathcal{I}}(i) = o$, and $\llbracket t \rrbracket_{\text{Sem}}(i) \neq o$. Otherwise, VERIFY returns *None* to signify that Sem is equivalent to \mathcal{I} for all terms within the sub-grammar G' .

Example 4.6 (Synthesizing Semantics of 0 for G_{IMP}). Recall the IMP language in Example 2.1.

SEMSYNTH first synthesizes the semantics of the leaves of G_{IMP} . Assume that *Order* assigns the production $\mathbf{0}$ index 1 (i.e., it is the first production explored by SEMSYNTH). During the first iteration of SEMSYNTH, the example set is empty and SYNTHSEMANTICCONSTRAINT may return any constraint F of the form $x_0^{\text{out}} = f(x_0^{\text{in}})$. Assume that SYNTHSEMANTICCONSTRAINT returns the constraint $x_0^{\text{out}} = 1$. VERIFY returns the counter-example $\langle 0, \mathbf{0}, 0 \rangle$, and the example set E is updated.

In the next iteration, `SYNTHEMANTICCONSTRAINT` must return a constraint satisfying the updated example set. For example, suppose that `SYNTHEMANTICCONSTRAINT` returns the constraint $x_0^{out} = x_1^{in}$. Again, `VERIFY` determines that $x_0^{out} = x_1^{in}$ is incorrect and returns the new counter-example $\langle 1, 0, 0 \rangle$. The example set E is updated with the returned counter-example.

A new iteration of the loop is run. On this loop, `SYNTHEMANTICCONSTRAINT` must return a constraint that satisfies both of the previously returned examples. This time `SYNTHEMANTICCONSTRAINT` returns the constraint $x_0^{out} = 0$. This time, `VERIFY` determines that $x_0^{out} = 0$ is correct, and `SEMSYNTH` proceeds to synthesize the semantics of the next production (e.g., 1).

In Example 4.6, we see how `SEMSYNTH` handles nullary (leaf) productions. `SEMSYNTH` works nearly identically for most production rules (excluding semantically recursive productions like `while` loops). We demonstrate in Example 4.7 how `SEMSYNTH` synthesizes a semantics for non-nullary productions.

Example 4.7 (Synthesizing Semantics of Sequencing for IMP). Continuing from Example 4.6, `SEMSYNTH` proceeds and comes to the sequencing operator (i.e., for production $S \rightarrow ;(S, S)$). After several attempts at synthesizing the semantics of sequencing, $E(\cdot)$ contains the examples $\langle 0, x := 1; x := 0, 0 \rangle$, $\langle 0, x := 0; x := x + 1, 1 \rangle$, and $\langle 1, x := 0; (x := 1; x := x + 1), 2 \rangle$.

As in the nullary case, we summarise the semantics of each example sub-term using the example set E . `SEMSYNTH` then generates the formula specifying that the desired semantic constraint satisfies the example set E using the generated summaries, and produces a new semantic constraint using `SYNTHEMANTICCONSTRAINT`. On this iteration, `SYNTHEMANTICCONSTRAINT` returns the correct semantic constraint, `VERIFY` determines whether it is correct, and `SEMSYNTH` proceeds to synthesize a semantics for the next production.

4.5 Synthesizing Semantics for Semantically Recursive Productions

So far, we have seen how `SEMSYNTH` handles nullary productions and structurally recursive productions (e.g., `ite` and sequencing). However, we have not yet seen how to handle productions that are *semantically* recursive (e.g., `while` loops). To handle semantically recursive productions, we augment the form of the desired constraint to be synthesized: `SYNTHEMANTICCONSTRAINT` must synthesize a predicate P_{rec} and two base constraints F_{nonrec} and F_{rec} such that for every example $\langle in, p(t_1, \dots, t_n), out \rangle$, the following conditions hold:

$$\frac{\dots \quad \frac{Sem_{A_n}(t_n, x_{A_n}^{in}, x_{A_n}^{out}) \quad \neg P_{rec}(x_{A_0}^{in}, x_{A_1}^{out}, \dots, x_{A_n}^{out}) \quad F_{non-rec}(x_{A_0}^{in}, x_{A_1}^{out}, \dots, x_{A_n}^{out}) \quad x_{A_0}^{in} = in}{x_{A_0}^{out} = out} \quad non-rec}{Sem_{A_0}(p(t_1, \dots, t_n), x_{A_0}^{in}, x_{A_0}^{out}) \quad \frac{Sem_{A_1}(t_1, x_{A_1}^{in}, x_{A_1}^{out}) \quad \dots \quad Sem_{A_n}(t_n, x_{A_n}^{in}, x_{A_n}^{out}) \quad P_{rec}(x_{A_0}^{in}, x_{A_1}^{out}, \dots, x_{A_n}^{out}) \quad F_{rec}(x_{A_0}^{in}, x_{A_1}^{out}, \dots, x_{A_n}^{out}) \quad x_{A_0}^{in} = in}{x_{A_0}^{out} = out} \quad rec} \quad rec$$

where P_{rec} determines if the non-rec or rec condition should hold. The non-recursive case is similar to the conditions for non-semantically recursive statements (with the addition of asserting that P_{rec} is false). The recursive case, however additionally allows the semantics to make use of a recursive call to the program term. Other than the change in the shape of the desired semantics, `SEMSYNTH` remains unchanged.

Example 4.8 (Synthesizing semantics of while loops for IMP). Continuing from Example 4.7, `SEMSYNTH` eventually reaches the `while` production. We assume that the grammar G additionally annotates whether each production is semantically recursive.

589 After several more iterations, the set $E(\text{while})$ contains the examples $\langle 0, t, 0 \rangle$, $\langle 1, t, 0 \rangle$, and
 590 $\langle 1, t, 0 \rangle$, where t is the term `while 0 < x do x := x - 1`. In this iteration, `SYNTHEMANTIC-`
 591 `CONSTRAINT` gets called with a recursive summary of t containing the three examples, and the
 592 exact synthesized semantics for `x := x - 1` and `0 < x`.

593 In this iteration, `SYNTHEMANTICCONSTRAINT` finds the correct P_{rec} , $F_{non-rec}$ and F_{rec} . `VERIFY`
 594 determines that the result is indeed correct and the main loop of `SEMSYNTH` terminates (because
 595 `while` is the last production in the grammar). Finally, `SEMSYNTH` returns the synthesized semantics
 596 for each production.

597 Now that we have defined how `SEMSYNTH` handles semantically recursive productions, `SEMS-`
 598 `SYNTH` is fully specified. Theorem 4.9 states that `SEMSYNTH` is sound.

600 **THEOREM 4.9 (SEMSYNTH IS SOUND).** *For any semantics-synthesis problem $\mathcal{P} = \langle G, \mathcal{I}, \mathcal{E} \rangle$, if*
 601 *`SEMSYNTH`($G, \mathcal{I}, \mathcal{E}$) returns a semantics Sem , then Sem is a solution to \mathcal{P} .*

602 **PROOF.** In `SEMSYNTH`, we maintain the loop invariant that the synthesized semantics Sem is
 603 correct with respect to the oracle \mathcal{E} for productions p_1 through p_{i-1} . This condition trivially holds
 604 on the first iteration. To proceed to iteration $i + 1$, `VERIFY` must return *None*, which implies that
 605 Sem is correct for productions p_1 through p_i . Thus the invariant is maintained. Upon back-tracking,
 606 the invariant is trivially true (because it held for some greater iteration). Thus, upon termination
 607 with $i = N + 1$, Sem must be correct for all productions of G —i.e., Sem satisfies the given semantics-
 608 synthesis problem \mathcal{P} . \square

610 While Theorem 4.9 states the soundness of `SEMSYNTH`, it fails to show that `SEMSYNTH` will
 611 eventually synthesize a correct semantics. Theorem 4.10 states that `SEMSYNTH` makes progress.
 612 Intuitively, it ensures that once a semantics is explored during an iteration of `SEMSYNTH`, it is
 613 never explored in any future iterations of `SEMSYNTH`. However, the formal statement relaxes this
 614 condition because back-tracking may cause (a finite number of) future iterations to explore the
 615 same semantics.

616 **THEOREM 4.10 (SEMSYNTH MAKES PROGRESS).** *For any semantics-synthesis problem $\mathcal{P} = \langle G, \mathcal{I}, \mathcal{E} \rangle$,*
 617 *if `SEMSYNTH`($G, \mathcal{I}, \mathcal{E}$) is on iteration k of the main loop with current synthesized semantics Sem_k , then*
 618 *for some iteration $k_0 > k$, for all iterations $k' \geq k_0$, Sem will never take the value $Sem_{k'}$ again (i.e.,*
 619 *$Sem_{k'} \neq Sem_k$).*

621 **PROOF.** Assume the negation holds, i.e., “ $\forall k_0 > k, \exists k' \geq k_0, Sem_{k'} = Sem_k$ ”. Suppose that at
 622 iteration k , the example set is E_k , and at iteration k_0 , the example set is E_0 . Given that $Sem_{k'} = Sem_k$,
 623 those two semantics are defined for the same set of productions $\{p_1, \dots, p_c\}$. Thus, line 10–12 of
 624 Algorithm 1 must be executed at least once between the iterations k and k_0 (both inclusive). If
 625 we consider the first execution of line 11 after iteration k (and let k_1 denote the current iteration
 626 at that moment), we have that for production p_j , where $(1 \leq j \leq c)$, $E_k[p_j] \cup \{cex\} \subset E_0[p_j]$,
 627 where $cex = (in, t, out)$ satisfies $\llbracket t \rrbracket_{Sem_k}(in) = \llbracket t \rrbracket_{Sem_{k'}}(in) \neq out$. This situation will never happen,
 628 because after cex is added to E in iteration k_1 , on line 8, the constraints in Section 4.2 do not allow
 629 $\llbracket t \rrbracket_{Sem_{k'}}(in) \neq out$. Thus we reach a contradiction. \square

631 5 IMPLEMENTATION

632 This section gives details of `SYNANTIC`, which implements our approach to synthesizing semantics
 633 via the algorithm `SEMSYNTH`. `SYNANTIC` is developed in Scala (version 2.13), and uses `CVC5` (version
 634 1.0.3) to solve SyGuS problems—which are used within our implementation of `SYNTHEMANTIC-`
 635 `CONSTRAINT` to generate candidate semantic constraints. The remainder of this section is struc-
 636 tured as follows: Section 5.1 details how we implement `SYNTHEMANTICCONSTRAINT`. Section 5.2
 637

summarizes the implementation of `VERIFY`, and explains how we approximate an equivalence oracle for an interpreter. Section 5.3 presents an optimization of `SYNTHEMANTICCONSTRAINT` for productions with multiple outputs (i.e., where the output type of a production is a tuple).

5.1 Implementation of `SYNTHEMANTICCONSTRAINT`

In Section 4, `SEMSYNTH` is parameterized on the procedure `SYNTHEMANTICCONSTRAINT`. On line 8 of Algorithm 1, we assume that `SYNTHEMANTICCONSTRAINT` produces a semantic constraint F for production p_i that satisfies the example set E . To accomplish this task, we construct a SyGuS problem consisting of a grammar of allowable semantic constraints and a set of conditions to enforce that the semantic constraint is consistent with the example set. To handle productions whose semantics does not evaluate its child terms from left to right, we run in parallel a version of `SYNTHEMANTICCONSTRAINT` for each permutation of the child terms and immediately return upon any permutation's success. In practice, for all of our benchmarks, all the productions evaluate their children from left to right.

We defer discussion of the SyGuS grammars we use to Section 6.1 when we discuss each benchmark. The specification of the semantic constraint is exactly the condition specified in Equation (3).

5.2 Implementation of `VERIFY`

In Section 4, Algorithm 1 is also parameterized on the procedure `VERIFY` (line 9), which uses the equivalence oracle \mathcal{E} to determine if the learned semantics Sem is consistent with the interpreter for all terms that only uses the productions p_1 through p_i . In `SYNTANTIC`, we approximate an equivalence oracle using fuzzing. Specifically, we randomly generate terms and inputs and use the interpreter \mathcal{I} to generate an output. We then use the learned constraint for p_i to generate inputs to each sub-term (from left to right), and compute outputs for each using interpreter \mathcal{I} . In effect, we are computing a new example set E' , and testing the semantic constraints learned so far. If any example disagrees with the learned semantics of production p_j (for $1 \leq j \leq i$), the example and production p_j are returned as a counter-example.

When `VERIFY` fuzzes the semantics, it uses the interpreter to generate examples (i.e., terms with corresponding input-output examples). During example generation, we set a recursion limit of 1,000 recursive calls. We discard an input—i.e., we assume the program does not terminate—if its run exceeds the recursion depth.

5.3 Optimized `SYNTHEMANTICCONSTRAINT` for Multi-Output Productions

In Section 5.1, we described how `SYNTHEMANTICCONSTRAINT` produces and uses `CVC5` to solve a SyGuS problem to synthesize a semantic constraint that is consistent with the current example set. However, it is well known that SyGuS solvers scale poorly as a function of the size of the desired grammar/result. This issue is especially problematic when learning a semantic constraint for a language in which productions have multiple outputs (e.g., statements for `IMP` with more than one variable) and thus the grammar and resulting constraint grow with the number of outputs.

To address this issue, we modified `SYNTHEMANTICCONSTRAINT` to synthesize a constraint for each output independently. However, this process may lead to constraints that do not agree on the internal data flow of the constraints (i.e., the functions determining the input to each child term). To remedy this issue, our implementation of `SYNTHEMANTICCONSTRAINT` uses an additional `CEGIS` loop that resynthesizes the constraint for each output until all agree on the inputs to each child term. For simplicity, we explain how this optimization works for a production that has two outputs. Consider the case for $A_0 \rightarrow p(A_1, \dots, A_n)$ where $\tau_{A_0} \triangleq \tau_1 \times \tau_2$. In this scenario, our goal is to

687 synthesize two constraints F and G ,

$$688 \quad F \triangleq x_1 = f_1(x_0) \wedge \cdots \wedge x_n = f_n(x_0, x'_1, \dots, x'_{n-1}) \wedge x'_0 = f_0(x_0, x'_1, \dots, x'_n)$$

$$689 \quad G \triangleq x_1 = g_1(x_0) \wedge \cdots \wedge x_n = g_n(x_0, x'_1, \dots, x'_{n-1}) \wedge x'_0 = g_0(x_0, x'_1, \dots, x'_n)$$

691 To determine if F and G agree on each child term's input for example set E' , we generate the
692 formula ϕ shown below, for each example $\langle in, p(t_1, \dots, t_n), out \rangle \in E'(p)$:

$$693 \quad x_0^F = x_0^G = in \wedge Summary(t_1)(x_1^F, x_1'^F) \wedge \cdots \wedge Summary(t_1)(x_n^F, x_n'^F) \quad (4)$$

$$694 \quad \wedge Summary(t_1)(x_1^G, x_1'^G) \wedge \cdots \wedge Summary(t_1)(x_n^G, x_n'^G) \quad (5)$$

$$695 \quad \wedge F \wedge G \wedge (x_1^F \neq x_1^G \vee \cdots \vee x_n^F \neq x_n^G) \wedge \langle x_0'^F, x_0'^G \rangle = out \quad (6)$$

696 which asks if F and G agree on the input to each child term for the given example. To make this
697 concept concrete, consider the following example.

701 *Example 5.1 (Synthesizing Semantic Constraint for Multi-Output Production.)*. Consider the
702 task of synthesizing a semantics for $x_0 :=$ in the language IMP_2 , using the examples:
703 $\langle \langle 0, 1 \rangle, x_0 := x_1, \langle 1, 1 \rangle \rangle$, $\langle \langle 0, 1 \rangle, x_1, 1 \rangle$, $\langle \langle 1, 1 \rangle, x_1, 1 \rangle$.

704 For the above examples, `SYNTHEMANTICCONSTRAINT` might generate $F \triangleq x_{1,0}^{in} = x_{0,0}^{in} \wedge x_{1,1}^{in} =$
705 $x_{0,1}^{in} \wedge x_{0,0}^{out} = x_{1,1}^{out}$ and $G \triangleq x_{1,0}^{in} = x_{0,1}^{in} \wedge x_{1,1}^{in} = x_{0,1}^{in} \wedge x_{0,0}^{out} = x_{1,1}^{out}$, where $x_{i,j}^{in}$ is the j^{th} projection
706 of x_i^{in} . While both F and G are consistent with the examples, the data-flow of F is not consistent
707 with the data-flow of G (i.e., in F , $x_{1,0}^{in}$ is assigned $x_{0,0}^{in}$, while in G , $x_{1,0}^{in}$ is assigned $x_{0,1}^{in}$). We can
708 construct the formula in Equation (4) for F and G , and find out that in F , the variable $x_{1,0}^{in F}$ takes
709 the value 0, and in G , the variable $x_{1,0}^{in G}$ takes value 1. Thus, F and G are not consistent on data-
710 flows to children for the provided example. We generate a new condition for the next iteration of
711 `SYNTHEMANTICCONSTRAINT` that asserts $x_{0,0}^{in F} \neq 0 \vee x_{0,0}^{in G} \neq 1$.

712 In practice, we create a copy of each variable indexed by F and G , respectively, to avoid clashing
713 variable names when encoding the constraints F and G within a single formula. To check the
714 consistency of F and G 's data flows, we use `cvc5` to check the satisfiability of the formula ϕ in
715 Equation (4). If ϕ is unsatisfiable, then F and G must agree on the inputs of all child terms for
716 the given examples. If so, then we may return either $F \wedge x_{0,2} = g_0(\dots)$ or $G \wedge x_{0,1} = f_0(\dots)$ (i.e.,
717 because F and G agree on all child term inputs, we may use either to constrain the data-flow to
718 child terms).

719 If ϕ is satisfiable, then F and G do not agree on the input to all child terms. In this case, we find
720 a model that satisfies ϕ . If there is some subterm t_i such that there is no example $\langle in, t_i, out \rangle \in E$ such
721 that $in = M(x_i^F)$ or $in = M(x_i^G)$, then we add the example $\langle in, t, \mathcal{I}(t_i, in) \rangle$ to the set of examples,
722 and resynthesize the constraints F and G . Otherwise, we know that the sub-term summaries are
723 sufficient to fully specify both F and G for all examples in E . Thus, we must add a new constraint
724 that ensures the pair of constraints F and G are never synthesized again. To do this, we add a new
725 constraint $x_0^F \neq M(x_0^F) \vee x_0^G \neq M(x_0^G) \vee \cdots \vee x_n^F \neq M(x_n^F) \vee x_n^G \neq M(x_n^G)$, which ensures that the
726 input of at least one of the child terms for either F or G must change. A new candidate F and G
727 are then synthesized. The `CEGIS` loop continues until it finds a valid pair of F and G for the set of
728 examples.

731 6 EVALUATION

732 The goal of our evaluation is to answer the following questions:

733 **RQ1** Can `SYNTHEMANTIC` synthesize the semantics of non-trivial languages?

736 **RQ2** Where is time spent during synthesis?

737 **RQ3** Is the multi-output optimization from Section 5.3 effective?

738 **RQ4** How do synthesized semantics compare to manually written ones?

739 All experiments were run on a machine with an Intel(R) i9-13900K CPU and 32 GB of memory,
 740 running NixOS 23.10 and Scala 2.13.13. All experiments were allotted 2 hours, 4 cores of CPU, and
 741 24 GB of memory. Cvc5 version 1.0.3 is used for SMT solving and SyGuS function synthesis. For
 742 the total running time of each experiment, we report the median of 7 runs using different random
 743 seeds. For every language, we record whether SYNANTIC terminates within the given time limit of
 744 2 hours, and when it does, we also record the set of synthesized semantic rules. A language that
 745 does not terminate within the time limit on more than half of the seeds is reported as a timeout.
 746

747 6.1 Benchmarks

748 We collected 14 benchmarks from the two sources discussed below. For every language discussed
 749 in this section, we manually translated the semantics to a simple equivalent interpreter written in
 750 Scala; our goal was then to synthesize an appropriate CHC-based semantics from the interpreter.
 751 The one non-standard feature of our setup is that the interpreter must be capable of interpreting
 752 the programs derived from *any* nonterminal in the grammar.
 753

754 *SemGuS benchmarks.* Our first source of benchmarks is the SemGuS benchmark repository [11].
 755 This dataset contains SemGuS synthesis problems where each problem consists of a grammar of
 756 terms, a set of CHCs inductively defining the semantics of terms in the grammar, and a specification
 757 that the synthesized program should meet. For our purposes, we ignored the specification and
 758 collected the grammar plus semantics for 10 distinct languages that appear in the repository. We
 759 do not consider languages that contain abstract data types (e.g., stacks) or require a large range of
 760 inputs (e.g., ASCII characters) due to their poor support by the SyGuS solver. These 10 languages
 761 gave us 10 benchmarks.

762 Some of the languages used in the SemGuS benchmark set are parametric (denoted by a param-
 763 eter k), meaning that the semantics is slightly different based on a given parameter (e.g., number
 764 of program variables for IMP and length of the input string for regular expressions). For these
 765 benchmarks, we ran SYNANTIC on an increasing sequence of parameter values and reported the
 766 largest parameter value for which SYNANTIC succeeds.

767 $\text{REGEX}(k)$ is a language for matching regular expressions on strings of length k ; Given a regu-
 768 lar expression r and string s of length k (index starts from 0), the semantic functions produce a
 769 Boolean matrix $M \in \text{Bool}^{(k+1) \times (k+1)}$ such that $M_{i,j} = \text{true}$ iff the substring $s_{i..j-1}$ matches regular
 770 expression r —here $s_{i..i}$ denotes the empty string, and by definition, $M_{i,j} = \text{false}$ for $i \geq j$.

771 $\text{CNF}(k)$, $\text{DNF}(k)$ and $\text{CUBE}(k)$ are languages of Boolean formulas (of the syntactic kind indicated
 772 by their names, i.e., conjunctive normal form, disjunctive normal form, and cubes) involving up to
 773 k variables.

774 $\text{IMP}(k)$ is an imperative language that contains common control flow structures, such as condi-
 775 tionals and while loops, for programs with k integer variables. Note that IMP includes operators
 776 such as `while` and `do_while` for which the semantics involves semantically recursive productions
 777 (Section 4.5). The complete semantics of $\text{IMP}(k)$ can be found in the supplementary material.

778 INTARITH is a benchmark about basic integer calculations, like addition, multiplication, and
 779 conditional selection. It also includes three constants whose value can be specified in the input to
 780 the semantic relations.

781 $\text{BVSIMPLE}(k)$ describes bit-vector operations involving k bit-vector constants.
 782 $\text{BVSIMPLEIMP}(m, n)$ is essentially a variant of $\text{BVSIMPLE}(k)$ that augments the language with
 783 let-expressions. Parameters m and n mean that the language can use up to m bit-vector constants
 784

785 and n bit-vector variables. $\text{BVSATURATE}(k)$ and $\text{BVSATURATEIMP}(k)$ use the same syntaxes as
 786 $\text{BVSIMPLE}(k)$ and $\text{BVSIMPLEIMP}(k)$, respectively, but operations use a saturating semantics that
 787 never overflows or underflows.

788 *Attribute-grammar synthesis* [10]. Our second source of benchmarks is from the Panini tool for
 789 synthesizing attribute grammars [10]. An attribute grammar (AG) associates each nonterminal
 790 of an underlying context-free grammar with some number of *attributes*. Each production has a
 791 set of attribute-definition rules (sometimes called *semantic actions*) that specify how the value of
 792 one attribute of the production is set as a function of the values of other attributes of the produc-
 793 tion. In a given derivation tree of the AG, each node has an associated set of *attribute instances*.
 794 The attribute-definition rules are used to obtain a consistent assignment of values to the tree’s
 795 attribute instances: each attribute instance has a value equal to its defining function applied to the
 796 appropriate (neighboring) attribute instances of the tree. Effectively, AGs assign a semantics to
 797 programs via attributes, and the underlying attribute-definition rules can be captured via CHCs.
 798 While there are AG extensions to handle circular AGs [9, 14]—i.e., AGs in which some derivation
 799 trees have attribute instances that are defined in terms of themselves—the work of Kalita et al.
 800 concerns non-circular AGs.

801 Kalita et al. [10] present 12 benchmarks. We ignored 4 benchmarks that are either (i) not publicly
 802 accessible, or (ii) use semantic functions that cannot be expressed in SMT-LIB and are thus beyond
 803 what can be synthesized using a SyGuS solver—e.g., complex data structures, or (iii) identical to
 804 existing benchmarks from other sources. We did not run their tool on our benchmarks because our
 805 problem is more general than theirs, supporting a wider range of language semantics: the scope of
 806 our work includes recursive semantics, which can be handled only indirectly in a system such as
 807 theirs (which supports only non-circular AGs)—i.e., by introducing powerful hard-to-synthesize
 808 recursive functions that effectively capture an entire construct’s semantics. The running time is
 809 also not directly comparable, because Kalita et al.’s approach uses user-provided sketches (i.e.,
 810 partial solutions to each semantic action), which simplifies the synthesis problem. In contrast, in
 811 our work we do not assume that a sketch is provided for the semantic constraints and instead
 812 consider general SyGuS grammars.

813 The remaining 8 benchmarks of Kalita et al. are consolidated as 4 languages (i.e., giving us
 814 four benchmarks). ITEEXPR is a language of basic integer operations, comparison expressions, and
 815 ternary if-then-else expressions (not statements). Our ITEEXPR benchmark subsumes benchmarks
 816 B3, B4, and B5 of Kalita et al. because their only differences stem from whether the expression
 817 is written in prefix, postfix, or infix notation. For SYNANTIC , such surface-syntax differences are
 818 unimportant because SYNANTIC uses regular tree grammars to express a language’s abstract syntax,
 819 and the underlying abstract syntax of prefix, postfix, and infix expressions is the same. BINOP is
 820 a language of binary strings (combined from benchmarks B1 and B2 of Kalita et al.), along with
 821 built-in functions for popcount (counting the number of ones) and binary-to-decimal conversion.
 822 CURRENCY is a language for currency exchange and calculation. DIFF is a language for computing
 823 finite differences. Because the original benchmark from Kalita et al. involves differentiation and
 824 real numbers (which are not supported by existing SyGuS solvers), we modified the benchmark
 825 to perform the related operation of finite differencing over integer-valued functions. Specifically,
 826 for a function f , its finite difference is defined as $\Delta f = f(x + 1) - f(x)$. Starting from here, finite
 827 differences for sums and products can be obtained compositionally, e.g., $\Delta(u \cdot v) = u(x)\Delta v(x) +$
 828 $v(x + 1)\Delta u(x)$.

829
 830 *SyGuS grammars*. For each semantic function, we also provided a grammar for the SyGuS solver,
 831 which contains the operators of the underlying logical theory and any specific functions that must
 832 appear in the target semantics.

834 For instance, for all benchmarks using the logic fragment NIA, we allow the use of basic in-
 835 teger operations and integer constants, along with language-specific operations like conditional
 836 operators (if-then-else).

837 For the languages DIFF and CURRENCY we did not include conditional operators, because they
 838 do not appear in the semantics.

839 For BVSATURATED and BVIMPSATURATED we provided operators for detecting overflow and
 840 underflow.

841 Lastly, for languages known to be free of side effects, we modified the SyGuS grammars to forbid
 842 data flow between siblings, and only allow parent-to-child and child-to-parent assignments.

843

844

6.2 RQ1: Can SYNANTIC Synthesize the Semantics of Non-trivial Languages?

845 Table 1 presents a highlight of the results of running SYNANTIC on each benchmark (column 1)
 846 for each production rule (column 2). For the parametric languages, we ran each benchmark up
 847 to the largest parameter k for which the solver timed out and reported the running time and
 848 other metrics for the largest such k (more details below). The third column provides the median
 849 number of CEGIS iterations taken to synthesize each production, and the fourth column provides
 850 the median number of $\langle in, term, out \rangle$ counterexamples found for one production rule. We take the
 851 median of total execution time on one production rule and list it in column 7. Columns 5–6 are
 852 breakdowns of the total time into time for SyGuS solving and time for SMT solving. To summarize,
 853 SYNANTIC could synthesize complete semantics for $11/14 \approx 79\%$ of benchmark languages.

854 For REGEX(k) ($k = 2, \dots, 8$), SYNANTIC could synthesize a semantics for up to $k = 2$. For CNF(k)
 855 ($k = 4, \dots, 8$), DNF(k) ($k = 4, \dots, 8$), and CUBE(k) ($k = 4, \dots, 11$), SYNANTIC could synthesize
 856 semantics for all parameters included in the SemGuS benchmarks. For IMP(k) ($k = 1, 2$), SYNAN-
 857 TIC could synthesize a semantics up to $k = 2$. For the bit vector benchmarks, SYNANTIC could
 858 synthesize a semantics for BVSIMPLE(k) up to $k = 3$, and a semantics for BVIMPSIMPLE(m, n)
 859 ($(m, n) \in \{(1, 2), (3, 3)\}$) up to $m = 1$ and $n = 2$.

860 For all these parametric cases that timeout, the number of input and output variables in semantic
 861 functions is large: 10 inputs and 10 outputs for REGEX(3).

862 Additionally, SYNANTIC timed out for the benchmarks DIFF, BVSATURATED, and BVIMPSATU-
 863 RATED.² For DIFF, 4 of the 7 runs resulted in a timeout, so DIFF is reported as a timeout (even
 864 though at least one run could synthesize the semantics of all the productions). For the 4 runs that
 865 timed out, SYNANTIC can solve the semantics of 5 of the 6 productions in the grammar. SYNANTIC
 866 could synthesize the semantics of 9/18 productions for BVIMPSATURATED, and 10/17 productions
 867 for BVSATURATED in at least one run.

868 In benchmarks that timed out, the time-out happened during a call to the SyGuS solver—i.e., the
 869 functions to be synthesized were too complex (more details in Section 6.3).

870

871 *Finding:* To answer RQ1, SYNANTIC can synthesize semantics for many non-trivial languages as
 872 long as the semantics does not involve very large functions (more than 20 terms).

873

874

6.3 RQ2: Where is Time Spent during Synthesis?

875 *SyGuS vs SMT Time.* Table 2 also presents the breakdown of how much time the solver spends
 876 solving SyGuS problems (to find candidate functions) and calling SMT solvers (to compute com-
 877 plete summaries). Among all the benchmarks, a median of 15.09% of the total solving time is spent
 878 on SyGuS problems, and a median of 20.67% of the time is spent solving SMT queries. However,
 879 for the slowest 10% production rules (>31.92 s), the median of SyGuS solving time grows to 65.99%,
 880 which indicates that SyGuS contributes to most of the execution time on slow-running cases.

881

882


²Data for some languages are only listed in the supplementary material.

Table 1. Detailed results for selected benchmarks. See supplementary material for the full list of results.


Lang.	Rule	# Iter.	# Ex	SyGuS (s)	SMT (s)	Total (s)
	$E \rightarrow 0$	1	1	0.01	0.01	0.05
	$E \rightarrow 1$	1	1	0.01	0.01	0.04
	$S \rightarrow x --$	2	2	0.06	0.02	0.11
	$S \rightarrow y --$	2	2	0.11	0.03	0.17
	$B \rightarrow f$	1	1	0.01	0.01	0.06
	$S \rightarrow x ++$	2	2	0.04	0.03	0.11
	$S \rightarrow y ++$	2	2	0.12	0.02	0.16
	$B \rightarrow t$	1	1	0.01	0.02	0.13
	$E \rightarrow x$	2	2	0.01	0.01	0.04
	$E \rightarrow y$	1	1	0.01	0.01	0.04
	$S \rightarrow x := E$	2	2	0.10	3.23	6.17
	$S \rightarrow y := E$	2	2	0.04	3.22	6.19
	$B \rightarrow \neg B$	3	3	0.02	2.49	5.26
	$E \rightarrow E + E$	4	3	0.05	8.52	14.83
	$E \rightarrow E - E$	5	2	0.13	8.03	13.83
	$B \rightarrow E < E$	8	5	0.08	7.50	13.66
	$B \rightarrow B \wedge B$	4	4	0.03	5.33	11.71
	$B \rightarrow B \vee B$	4	4	0.05	4.61	8.99
	$S \rightarrow S ; S$	5	3	4.55	15.00	72.53
	$S \rightarrow \text{do_while } S B$	27	35	858.50	257.33	1374.13
	$S \rightarrow \text{while } B S$	9	7	16.88	122.41	266.80
	$S \rightarrow \text{ite } B S S$	11	5	525.28	33.88	628.71
	$B \rightarrow 0$	1	1	0.01	0.01	0.07
	$B \rightarrow 1$	1	1	0.01	0.01	0.22
	$B \rightarrow x$	2	2	0.01	0.01	0.08
	$N \rightarrow \text{atom } B$	2	2	0.09	0.04	0.30
	$M \rightarrow \text{atom}' B$	3	3	0.07	0.05	0.26
	$S \rightarrow \text{bin2dec } M$	2	2	0.02	0.09	0.30
	$S \rightarrow \text{count } N$	2	2	0.04	0.05	0.24
	$N \rightarrow \text{concat } N B$	5	5	8.61	0.22	10.31
	$M \rightarrow \text{concat}' M B$	5	5	288.81	0.23	308.50
	$\text{Start} \rightarrow \text{eval } R$	3	3	0.02	4.43	13.40
	$R \rightarrow ?$	3	3	3.84	0.07	4.07
	$R \rightarrow a$	4	4	11.10	0.07	11.53
	$R \rightarrow b$	5	5	11.63	0.06	12.01
	$R \rightarrow \epsilon$	1	1	0.07	0.07	2.38
	$R \rightarrow \emptyset$	1	1	0.19	0.07	0.46
	$R \rightarrow !R$	5	5	2.85	15.77	77.36
	$R \rightarrow R^*$	6	6	0.99	13.06	31.91
	$R \rightarrow R \cdot R$	24	24	333.71	72.58	495.45
	$R \rightarrow R R$	10	10	10.96	59.54	140.82

90% of the per-production semantics are solved within 31.92 s. The 10 rules that take longer than 31.92 s to be synthesized are all non-leaf rules and their partial semantic constraint fall into the following three categories: (i) 3 of them contain large integers or complex SMT primitives (e.g., 32-bit integer division); (ii) 3 of them involve large logical formulas with sizes ranging between 8 and 24 subterms, e.g., formulas representing 3×3 matrix multiplication or other matrix operations; (iii) 4 of them contain two or more input and output variables, e.g., `while` and `do_while`. In particular, SYNANTIC takes 1374.13 s to synthesize the CHC for `do_while` because there can be many possible ways to modify the data flow between its child terms, and this aspect will incur in many CEGIS iterations. In all of the above cases, as expected from known limitations of CVC5, the SyGuS solver accounts for most of the execution time—74.51% of the total running time is spent calling the SyGuS solver and the last call to SyGuS solver takes on average 27.45% of the total running time.

Relation to CEGIS Iterations and Size of Solutions. Table 2 hints that the cost of synthesizing a semantics may be proportional to the number of CEGIS iterations, which in general is a good indicator of the complexity of a formula (and of how expressive the underlying SyGuS grammar is). Additionally, the cost should also be proportional to the size of synthesized parts in the SyGuS



(a) Time vs. semantic constraint size



(b) Time vs. partial semantic constraint size

Fig. 1. Plots relating the time to synthesize the semantics of one production rule vs final semantic constraint solution size (a) and partial semantic constraint solution size (b). We only included selected slowest benchmarks due to graph size limit.



Fig. 2. Execution time per iteration for `do_while`

Fig. 3. Speedup provided by optimization

problems, which directly indicates formula complexity. We plotted Figure 1 to better understand those relations by using the data from some slowest benchmarks.

Figure 1a shows the relationship between the time for synthesizing a per-production rule semantics and the size of the final semantics. For the same language, the time grows exponentially with the increase in the size of the final solution. Figure 1b shows that the time also grows exponentially with the increase in solution size for per-output partial semantic constraint.

Since the performance varies heavily across different benchmarks, to better understand the impact of CEGIS iterations, we focus our attention on one hard benchmark, `IMP(k)` where $k = 2$. Specifically, we analyze the time taken to synthesize the semantic rule for `do_while`, which was one of the hardest productions in our benchmark set (2,500s). Figure 2 provides a stack plot detailing the running time for all 16 CEGIS iterations needed to synthesize `do_while`. As expected, as more examples are accumulated by CEGIS iterations, the SyGuS solver requires more time. The execution time for different parts is plotted by the areas of different colors. We can conclude that for the rule of `do_while`, SyGuS solver takes 64.3% of the execution time.

$$\begin{array}{c}
981 \\
982 \\
983 \\
984 \\
985 \\
986 \\
987 \\
988 \\
989 \\
990 \\
991 \\
992 \\
993 \\
994 \\
995 \\
996 \\
997 \\
998 \\
999 \\
1000 \\
1001 \\
1002 \\
1003 \\
1004 \\
1005 \\
1006 \\
1007 \\
1008 \\
1009 \\
1010 \\
1011 \\
1012 \\
1013 \\
1014 \\
1015 \\
1016 \\
1017 \\
1018 \\
1019 \\
1020 \\
1021 \\
1022 \\
1023 \\
1024 \\
1025 \\
1026 \\
1027 \\
1028 \\
1029
\end{array}$$

$$\begin{array}{c}
\frac{\llbracket e_1 \rrbracket(s) = \begin{pmatrix} M_{0,0} & M_{0,1} \\ & M_{1,1} \end{pmatrix} \quad \llbracket e_2 \rrbracket(s) = \begin{pmatrix} M'_{0,0} & M'_{0,1} \\ & M'_{1,1} \end{pmatrix}}{\llbracket e_1 \cdot e_2 \rrbracket(s) = \begin{pmatrix} (M_{0,0} \wedge M'_{0,0}) & (M_{0,0} \wedge M'_{0,1}) \vee (M_{0,1} \wedge M'_{1,1}) & (M_{0,0} \wedge M'_{0,2}) \vee (M_{0,1} \wedge M'_{1,2}) \vee (M_{0,2} \wedge M'_{2,2}) \\ & (M_{1,1} \wedge M'_{1,1}) & (M_{1,1} \wedge M'_{1,2}) \vee (M_{1,2} \wedge M'_{2,2}) \\ & & (M_{2,2} \wedge M'_{2,2}) \end{pmatrix}} \text{CONCATM} \\
\text{(a) Manually written semantics} \\
\frac{\llbracket e_1 \rrbracket(s) = \begin{pmatrix} M_{0,0} & M_{0,1} \\ & M_{1,1} \end{pmatrix} \quad \llbracket e_2 \rrbracket(s) = \begin{pmatrix} M'_{0,0} & M'_{0,1} \\ & M'_{1,1} \end{pmatrix}}{\llbracket e_1 \cdot e_2 \rrbracket(s) = \begin{pmatrix} M_{2,2} \wedge M'_{1,1} & (M_{0,0} \wedge M'_{0,1}) \vee (M'_{0,0} \wedge M_{0,1}) & (M_{2,2} \wedge M'_{0,2}) \vee (M_{0,2} \wedge M'_{0,0}) \vee (M_{0,1} \wedge M'_{1,2}) \\ & (M_{0,0} \vee M_{2,2}) \wedge M'_{2,2} & (M'_{1,1} \wedge M_{1,2}) \vee (M_{1,1} \wedge M'_{1,2}) \\ & & (M_{1,1} \wedge M'_{0,0}) \end{pmatrix}} \text{CONCATS} \\
\text{(b) Synthesized Semantics}
\end{array}$$

Fig. 4. Manually-written and synthesized semantics for CONCAT in REGEX(2)

Finding: To answer RQ2, SYNANTIC spends most of the time (71.78%) solving SyGuS problems, and the time is affected by the size of the candidate semantic function.

6.4 RQ3: Is the Multi-output Optimization from Section 5.3 Effective?

Figure 3 compares the running time of SYNANTIC with and without the multi-output optimization (Section 5.3) on all the runs of our tools for the 7 different random seeds.

With the optimization turned off, SYNANTIC timed out on 10 more runs (specifically all the 7 runs for REGEX and 3 more runs for DIFF). All the benchmarks for which disabling the optimization caused a timeout have 3 or more output variables. Comparing Figure 1a and Figure 1b shows how the semantic functions used in the REGEX benchmarks are very large (up to size 50), but thanks to the optimization, our algorithm only has to solve SyGuS problems on formulas of size at most 15.

On the runs that terminated both with and without the optimization, the non-optimized algorithm is on average 8% faster—i.e., the two versions of the algorithm have comparable performance. However, for 15/98 runs the optimization results in a 20% or more slowdown. When inspecting these instances, we observed that the multi-output optimization spent many iterations synchronizing the many possible data flows for productions where the final term was actually small but many variables were involved—e.g., sequential composition in IMP(2).

Finding: The multi-output optimization from Section 5.3 is effective for languages with 3 or more output variables in their semantics.

6.5 RQ4: How do Synthesized Semantics Compare to Manually Written Ones?

The synthesized semantics for almost all of our benchmarks are either identical to the original manually constructed one, or each CHC in the synthesized semantics is logically equivalent to the CHC of the original semantics.

The one exception is the semantics synthesized for the language of REGEX(2), for which the individual CHCs for OR, CONCAT, NEG, and STAR are not logically equivalent to the manually-written ones. For instance, consider the Concat rule for the semantics of concatenation. For this construct, the manually written CHC is shown in Figure 4a, whereas SYNANTIC synthesizes the CHC shown in Figure 4b. The two CHCs are not logically equivalent. For example, if the children evaluate to the matrices $M = \begin{pmatrix} true & false & false \\ & false & false \\ & & true \end{pmatrix}$ and $M' = \begin{pmatrix} true & false & false \\ & false & false \\ & & true \end{pmatrix}$, the outputs computed by the manually

$$\frac{\begin{array}{l} \llbracket e_1 \rrbracket(s) = \left(M_\epsilon, \begin{pmatrix} M_{0,0} & M_{0,1} \\ & M_{1,1} \end{pmatrix} \right) \quad \llbracket e_2 \rrbracket(s) = \left(M'_\epsilon, \begin{pmatrix} M'_{0,0} & M'_{0,1} \\ & M'_{1,1} \end{pmatrix} \right) \\ \llbracket e_1 \cdot e_2 \rrbracket(s) = \left(M_\epsilon \wedge M'_\epsilon, \begin{pmatrix} (M_\epsilon \wedge M'_{0,0}) \vee (M_{0,0} \wedge M'_\epsilon) & (M_\epsilon \wedge M'_{0,1}) \vee (M_{0,0} \wedge M'_{1,1}) \vee (M_{0,1} \wedge M'_\epsilon) \\ (M_\epsilon \wedge M'_{1,1}) \vee (M_{1,1} \wedge M'_\epsilon) \end{pmatrix} \right) \end{array}}{\text{CONCAT}}$$

Fig. 5. New semantics for CONCAT in REGEXIMP.

written CHC and the synthesized CHC are $M_{man} = \begin{pmatrix} true & false & false \\ & false & false \\ & & true \end{pmatrix}$, and $M_{syn} = \begin{pmatrix} false & false & false \\ & true & false \\ & & false \end{pmatrix}$, respectively, which have different values on the diagonal.

When inspecting the two rules, we realized that the example matrices M and M' shown above cannot actually be produced by the semantic rules for regular expressions. In particular, the examples require different Boolean values to appear on the diagonal of one 3×3 matrix. However, all the elements on the diagonal represent the semantics of the regular expression on the empty string, so they must all have the same value! We note that this inconsistency in the semantics can also be observed without a reference semantics to compare against because different runs of the algorithm could return logically inequivalent CHCs—in fact, such inequivalence was how we initially discovered the inconsistency.

SYNANTIC helped us discover an inefficiency in the semantics that was being used in the standard regular expressions benchmarks in the SemGuS repository. We thus modified the interpreter so that for the example above it only produces a 2×2 matrix $M = \begin{pmatrix} M_{0,1} & M_{0,2} \\ & M_{1,2} \end{pmatrix}$ (corresponding to the non-empty substrings of the input string) and a *single* variable M_ϵ to denote whether the regular expression should accept the empty string (instead of the previous multiple copies of logically equivalent variables). This semantics reduces the total number of variables in the semantic domain from 6 to 4 in this example.

We call this new semantics REGEXIMP (see Figure 5 for an example). After modifying the interpreter to produce this new semantics, SYNANTIC synthesized the corresponding CHCs in a median of 1968.00 s.

To check whether the semantics REGEXIMP is indeed more efficient than the original semantics REGEX, we modified all the 28 regular-expression synthesis benchmarks appearing in the SemGuS benchmark set. Each of these benchmarks requires one to find a regular expression that accepts some examples and rejects others.

We then used the Ks2 enumeration-based synthesizer to try to solve all the benchmarks with either of the two semantics. Because Ks2 enumerates programs of increasing size and uses the semantics to execute them and discard invalid program candidates, we conjectured that executing programs faster allows Ks2 to explore the search space faster.

Ks2 was faster at solving synthesis problems with the REGEXIMP semantics than with the REGEX ones (although both solved the same set of benchmarks). Although the speedup over all benchmarks is only 1.1x, the new semantics REGEXIMP was *particularly beneficial* for the harder synthesis problems. When considering the 13 benchmarks for which synthesis using the REGEX took longer than one second, the speedup increased to 1.18x.

Finding: SYNANTIC synthesized semantics that were identical to the manually written ones for 13/14 benchmarks. When SYNANTIC found a logically inequivalent semantics, it unveiled a performance bug.

7 RELATED WORK

Synthesis of Recursive Programs At a high level, the semantics-synthesis problem we consider is similar to a number of works on synthesizing recursively defined programs [7, 8, 12, 15]. In effect, a semantics for a recursively defined grammar is a recursive program assigning meaning to programs within the language. Both Farzan et al. [7], Farzan and Nicolet [8] use recursion skeletons to reduce their task from synthesizing a recursive program to synthesizing a non-recursive program. Our use of semantic constraints plays a similar role. While both of their techniques assume programs are only structurally recursive (i.e., no recursion on the program term itself), and our framework explicitly allows for program terms that are self-recursive (e.g., `while` loops in IMP).

Similar to the approach used by Miltner et al. [15] to synthesize simple recursive programs, SEMSYNTH employs a bottom-up approach to synthesis (i.e., we first synthesize semantics for nullary productions before moving on to other productions). However, unlike Miltner et al., SEMSYNTH is well-defined for any ordering of production rules and targets a more complex setting—i.e., synthesizing program semantics. Finally, Lee and Cho [12] synthesize recursive procedures from examples by first synthesizing blocks of straight-line code. This approach is similar in fashion to how we synthesize semantics by synthesizing semantic constraints. Unlike SEMSYNTH, Lee and Cho do not use CEGIS to perform synthesis. Instead, they use a finite number of input examples to discriminate between recursive programs within the desired search space.

Datalog Synthesis Albarghouthi et al. [1] synthesize Datalog programs (i.e., Horn clauses) with SMT solvers, whereas Si et al. [16] use a syntax-guided approach. In our work, we use constrained Horn clauses, which are strictly more expressive than Datalog programs, to denote semantics. Aside from the fact that the Datalog synthesis problem considers different inputs (i.e., the data), CHC also contains a function in a theory \mathcal{T} (such as LIA or BV), which we have to synthesize.

Synthesizing attribute grammars Kalita et al. [10] proposed a sketch-based method for synthesizing attribute grammars. When provided with a context-free grammar, their tool can automatically create appropriate *semantic actions* from sketches of attribute grammars. Instead of semantic actions, in our work we use CHCs to express program semantics. Our approach can model recursive semantics whereas the technique by Kalita et al. is limited to non-circular attribute grammars. Additionally, while their method requires providing a distinct program sketch (i.e., a partial program) for each production, our approach only requires providing a (fairly general) SyGuS grammar for each nonterminal in the language.

8 CONCLUSION

Writing logical semantics for a language can be a difficult task and our work supplies a method to automatically synthesize a language’s semantics from an executable interpreter that is treated as a closed-box. By generating example terms and input-output pairs from the interpreter, we use a SyGuS solver to synthesize semantic rules. Our evaluation shows that the approach applies to a wide range of language features, e.g., recursive semantic functions with multiple outputs.

As discussed in Section 2, one motivation for this work is to be able to generate automatically the kind of semantics that is needed to create a program synthesizer using the SemGuS framework. In our algorithm, we harness a SyGuS solver to synthesize the constraint in each CHC—i.e., we harness SyGuS in service to SemGuS—which limits us to synthesizing constraints that are written in theories that SyGuS supports. Going forward, we would like to make use of “higher-level” theories, supporting such abstractions as stores or algebraic data types. As SemGuS-based synthesizers and verifiers improve, we might be able to satisfy this wish by using SemGuS in service to SemGuS! That is, we could extend SYNANTIC to use SemGuS solvers to synthesize semantic constraints.

REFERENCES

- [1] Aws Albarghouthi, Paraschos Koutris, Mayur Naik, and Calvin Smith. 2017. Constraint-based synthesis of datalog programs. In *Principles and Practice of Constraint Programming: 23rd International Conference, CP 2017, Melbourne, VIC, Australia, August 28–September 1, 2017, Proceedings 23*. Springer, 689–706.
- [2] Rajeev Alur, Rastislav Bodík, Garvit Juniwal, Milo M. K. Martin, Mukund Raghothaman, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. 2013. Syntax-guided synthesis. In *Formal Methods in Computer-Aided Design, FMCAD 2013, Portland, OR, USA, October 20–23, 2013*. IEEE, 1–8. <https://ieeexplore.ieee.org/document/6679385/>
- [3] Rajeev Alur, Arjun Radhakrishna, and Abhishek Udupa. 2017. Scaling enumerative program synthesis via divide and conquer. In *International conference on tools and algorithms for the construction and analysis of systems*. Springer, 319–336.
- [4] Haniel Barbosa, Clark Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Andrew Reynolds, Ying Sheng, Cesare Tinelli, and Yoni Zohar. 2022. cvc5: A Versatile and Industrial-Strength SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, Dana Fisman and Grigore Rosu (Eds.). Springer International Publishing, Cham, 415–442.
- [5] Xiaohong Chen and Grigore Rosu. 2019. A Semantic Framework for Programming Languages and Formal Analysis. In *Engineering Trustworthy Software Systems - 5th International School, SETSS 2019, Chongqing, China, April 21–27, 2019, Tutorial Lectures (Lecture Notes in Computer Science, Vol. 12154)*, Jonathan P. Bowen, Zhiming Liu, and Zili Zhang (Eds.). Springer, 122–158. https://doi.org/10.1007/978-3-030-55089-9_4
- [6] Loris D’Antoni, Qinheping Hu, Jinwoo Kim, and Thomas Reps. 2021. Programmable program synthesis. In *Computer Aided Verification: 33rd International Conference, CAV 2021, Virtual Event, July 20–23, 2021, Proceedings, Part I 33*. Springer, 84–109.
- [7] Azadeh Farzan, Danya Lette, and Victor Nicolet. 2022. Recursion synthesis with unrealizability witnesses. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 244–259.
- [8] Azadeh Farzan and Victor Nicolet. 2021. Counterexample-Guided Partial Bounding for Recursive Function Synthesis. In *Computer Aided Verification: 33rd International Conference, CAV 2021, Virtual Event, July 20–23, 2021, Proceedings, Part I 33*. Springer, 832–855.
- [9] Larry G. Jones. 1990. Efficient Evaluation of Circular Attribute Grammars. *ACM Trans. Program. Lang. Syst.* 12, 3 (1990), 429–462. <https://doi.org/10.1145/78969.78971>
- [10] Pankaj Kumar Kalita, Miriyala Jeevan Kumar, and Subhajit Roy. 2022. Synthesis of semantic actions in attribute grammars. In *2022 Formal Methods in Computer-Aided Design (FMCAD)*. IEEE, 304–314.
- [11] Jinwoo Kim, Qinheping Hu, Loris D’Antoni, and Thomas Reps. 2021. Semantics-guided synthesis. *Proceedings of the ACM on Programming Languages* 5, POPL (2021), 1–32.
- [12] Woosuk Lee and Hangeol Cho. 2023. Inductive synthesis of structurally recursive functional programs from non-recursive expressions. *Proceedings of the ACM on Programming Languages* 7, POPL (2023), 2048–2078.
- [13] Junghee Lim and Thomas W. Reps. 2013. TSL: A System for Generating Abstract Interpreters and its Application to Machine-Code Analysis. *ACM Trans. Program. Lang. Syst.* 35, 1 (2013), 4:1–4:59. <https://doi.org/10.1145/2450136.2450139>
- [14] Eva Magnusson and Görel Hedin. 2003. Circular Reference Attributed Grammars - Their Evaluation and Applications. In *Workshop on Language Descriptions, Tools and Applications, LDTA@ETAPS 2003, Warsaw, Poland, April 12–13, 2003 (Electronic Notes in Theoretical Computer Science, Vol. 82)*, Barrett R. Bryant and João Saraiva (Eds.). Elsevier, 532–554. [https://doi.org/10.1016/S1571-0661\(05\)82627-1](https://doi.org/10.1016/S1571-0661(05)82627-1)
- [15] Anders Miltner, Adrian Trejo Nuñez, Ana Brendel, Swarat Chaudhuri, and Isil Dillig. 2022. Bottom-up synthesis of recursive functional programs using angelic execution. *Proceedings of the ACM on Programming Languages* 6, POPL (2022), 1–29.
- [16] Xujie Si, Woosuk Lee, Richard Zhang, Aws Albarghouthi, Paraschos Koutris, and Mayur Naik. 2018. Syntax-guided synthesis of datalog programs. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 515–527.

A SEMANTICS FOR LANGUAGES USED IN BENCHMARK

In this section, we present the semantics synthesized by our tool SYNANTIC for any languages referenced within the main text. Appendix A.1 provides the synthesized semantics of SemGuS benchmarks and Appendix A.2 presents the synthesized semantics of the attribute grammar benchmarks.

A.1 SemGuS Benchmarks

The SemGuS suite of benchmarks consists of a total of 11 languages. We present the synthesized semantics of each as follows:

- (1) $CNF(k)$ is depicted in Figure 6.
- (2) $DNF(k)$ is depicted in Figure 7.
- (3) $CUBE(k)$ is depicted in Figure 8.
- (4) $INTARITH$ is depicted in Figure 9.
- (5) $REGEX(2)$ is depicted in Figure 10.
- (6) $REGEXSIMP()$ is depicted in Figure 11.
- (7) IMP is depicted in Figures 12 and 13.
- (8) $BVSIMPLE(k)$ is depicted in Figure 14.
- (9) $BVSATURATED(k)$ is depicted in Figure 15.
- (10) $BVIMPSIMPLE(m, n)$ is depicted in Figure 16.
- (11) $BVIMPSATURATED(m, n)$ is depicted in Figure 17.

A.2 Attribute-Grammar Synthesis

The suite of attribute-grammar benchmarks from [Kalita et al. 2022] consists of four languages which we present as follows:

- (1) $BINOP$ is presented in Figure 18.
- (2) $CURRENCY$ is presented in Figure 19.
- (3) $DIFF$ is presented in Figure 20.
- (4) $ITEEXPR$ is presented in Figure 21.

B BENCHMARK DATA

We present the full detailed evaluation results for all languages and productions rules in Table 2. For each production we present the number of CEGIS iterations, number of generated examples, and execution time (*i*) to solve SyGuS problems, (*ii*) to solve SMT queries, and (*iii*) overall. For each column we report the number for the median run based on the total run time. See Section 6 for full description of experimental setup.

$$\begin{array}{c}
1226 \\
1227 \\
1228 \\
1229 \\
1230 \\
1231 \\
1232 \\
1233 \\
1234 \\
1235 \\
1236 \\
1237 \\
1238 \\
1239
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, (k-1)}{\llbracket v_i \rrbracket(x_0, \dots, x_{k-1}) = x_i} \text{VARATOM} \qquad \frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{var } v \rrbracket(x_0, \dots, x_{k-1}) = r} \text{VAR} \\
\frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{nvar } v \rrbracket(x_0, \dots, x_{k-1}) = \neg r} \text{NOTVAR} \qquad \frac{\llbracket c \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{clause } c \rrbracket(x_0, \dots, x_{k-1}) = r} \text{CLAUSE} \\
\frac{\llbracket c \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket b \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket c \wedge b \rrbracket(x_0, \dots, x_{k-1}) = r_1 \wedge r_2} \text{AND} \\
\frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket c \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket v \vee c \rrbracket(x_0, \dots, x_{k-1}) = r_1 \vee r_2} \text{OR}
\end{array}$$

Fig. 6. Semantics of CNF(k)

$$\begin{array}{c}
1241 \\
1242 \\
1243 \\
1244 \\
1245 \\
1246 \\
1247 \\
1248 \\
1249 \\
1250 \\
1251 \\
1252 \\
1253 \\
1254 \\
1255 \\
1256
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, (k-1)}{\llbracket v_i \rrbracket(x_0, \dots, x_{k-1}) = x_i} \text{VARATOM} \qquad \frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{var } v \rrbracket(x_0, \dots, x_{k-1}) = r} \text{VAR} \\
\frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{nvar } v \rrbracket(x_0, \dots, x_{k-1}) = \neg r} \text{NOTVAR} \qquad \frac{\llbracket c \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{conj } c \rrbracket(x_0, \dots, x_{k-1}) = r} \text{CONJUNCTION} \\
\frac{\llbracket c \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket b \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket c \vee b \rrbracket(x_0, \dots, x_{k-1}) = r_1 \vee r_2} \text{OR} \\
\frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket c \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket v \wedge c \rrbracket(x_0, \dots, x_{k-1}) = r_1 \wedge r_2} \text{AND}
\end{array}$$

Fig. 7. Semantics of DNF(k)

$$\begin{array}{c}
1257 \\
1258 \\
1259 \\
1260 \\
1261 \\
1262 \\
1263 \\
1264 \\
1265 \\
1266 \\
1267 \\
1268 \\
1269 \\
1270 \\
1271 \\
1272 \\
1273 \\
1274
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, (k-1)}{\llbracket v_i \rrbracket(x_0, \dots, x_{k-1}) = x_i} \text{VARATOM} \qquad \frac{\llbracket v \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \text{var } v \rrbracket(x_0, \dots, x_{k-1}) = r} \text{VAR} \\
\frac{\llbracket b \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket b \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket b \wedge b \rrbracket(x_0, \dots, x_{k-1}) = r_1 \wedge r_2} \text{AND}
\end{array}$$

Fig. 8. Semantics of CUBE(k)

$$\begin{array}{l}
1275 \\
1276 \\
1277 \\
1278 \\
1279 \\
1280 \\
1281 \\
1282 \\
1283 \\
1284 \\
1285 \\
1286 \\
1287 \\
1288 \\
1289 \\
1290 \\
1291 \\
1292 \\
1293 \\
1294 \\
1295 \\
1296 \\
1297 \\
1298 \\
1299 \\
1300 \\
1301 \\
1302 \\
1303 \\
1304 \\
1305 \\
1306 \\
1307 \\
1308 \\
1309 \\
1310 \\
1311 \\
1312 \\
1313 \\
1314 \\
1315 \\
1316 \\
1317 \\
1318 \\
1319 \\
1320 \\
1321 \\
1322 \\
1323
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, 3}{\llbracket i \rrbracket(v_0) = i} \text{INTLITERAL} \quad \frac{}{\llbracket t \rrbracket(v_0) = true} \text{TRUE} \quad \frac{}{\llbracket f \rrbracket(v_0) = false} \text{FALSE} \\
\\
\frac{}{\llbracket x \rrbracket(v_0) = v_0} \text{VARX} \quad \frac{}{\llbracket y \rrbracket(v_0) = v_0} \text{VARY} \quad \frac{}{\llbracket z \rrbracket(v_0) = v_0} \text{VARZ} \\
\\
\frac{\llbracket b \rrbracket(v_0) = r_0 \quad \llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2 \quad r_0}{\llbracket \text{ite } b \ e_1 \ e_2 \rrbracket(v_0) = r_1} \text{ITE1} \\
\\
\frac{\llbracket b \rrbracket(v_0) = r_0 \quad \llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2 \quad \neg r_0}{\llbracket \text{ite } b \ e_1 \ e_2 \rrbracket(v_0) = r_2} \text{ITE2} \\
\\
\frac{\llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2}{\llbracket e_1 + e_2 \rrbracket(v_0) = r_1 + r_2} \text{PLUS} \quad \frac{\llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2}{\llbracket e_1 \times e_2 \rrbracket(v_0) = r_1 \cdot r_2} \text{MULTIPLY} \\
\\
\frac{\llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2}{\llbracket e_1 < e_2 \rrbracket(v_0) = (r_1 < r_2)} \text{LESSTHAN} \quad \frac{\llbracket b_1 \rrbracket(v_0) = r_1 \quad \llbracket b_2 \rrbracket(v_0) = r_2}{\llbracket \text{and } b_1 \ b_2 \rrbracket(v_0) = (r_1 \wedge r_2)} \text{AND} \\
\\
\frac{\llbracket b_1 \rrbracket(v_0) = r_1 \quad \llbracket b_2 \rrbracket(v_0) = r_2}{\llbracket \text{or } b_1 \ b_2 \rrbracket(v_0) = (r_1 \vee r_2)} \text{OR} \quad \frac{\llbracket b \rrbracket(v_0) = r}{\llbracket \text{not } b \rrbracket(v_0) = (\neg r)} \text{NOT}
\end{array}$$

Fig. 9. Semantics of INTARITH

$$\begin{array}{c}
1324 \\
1325 \\
1326 \\
1327 \\
1328 \\
1329 \\
1330 \\
1331 \\
1332 \\
1333 \\
1334 \\
1335 \\
1336 \\
1337 \\
1338 \\
1339 \\
1340 \\
1341 \\
1342 \\
1343 \\
1344 \\
1345 \\
1346 \\
1347 \\
1348 \\
1349 \\
1350 \\
1351 \\
1352 \\
1353 \\
1354 \\
1355 \\
1356 \\
1357 \\
1358 \\
1359 \\
1360 \\
1361 \\
1362 \\
1363 \\
1364 \\
1365 \\
1366 \\
1367 \\
1368 \\
1369 \\
1370 \\
1371 \\
1372
\end{array}$$

$$\begin{array}{c}
M := \begin{pmatrix} M_{0,0} & M_{0,1} & M_{0,2} \\ & M_{1,1} & M_{1,2} \\ & & M_{2,2} \end{pmatrix} \quad M' := \begin{pmatrix} M'_{0,0} & M'_{0,1} & M'_{0,2} \\ & M'_{1,1} & M'_{1,2} \\ & & M'_{2,2} \end{pmatrix} \quad \frac{\llbracket e \rrbracket(s) = M}{\llbracket \text{eval } e \rrbracket(s) = M_{0,2}} \text{ EVAL} \\
\\
\frac{}{\llbracket \epsilon \rrbracket(s) = \begin{pmatrix} \text{true} & \text{false} & \text{false} \\ & \text{true} & \text{false} \\ & & \text{true} \end{pmatrix}} \text{ EPS} \quad \frac{}{\llbracket \phi \rrbracket(s) = \begin{pmatrix} \text{false} & \text{false} & \text{false} \\ & \text{false} & \text{false} \\ & & \text{false} \end{pmatrix}} \text{ PHI} \\
\\
\frac{}{\llbracket \mathbf{a} \rrbracket(s) = \begin{pmatrix} \text{false} & (s_0=\mathbf{a}) & \text{false} \\ & \text{false} & (s_1=\mathbf{a}) \\ & & \text{false} \end{pmatrix}} \text{ CHARA} \quad \frac{}{\llbracket \mathbf{b} \rrbracket(s) = \begin{pmatrix} \text{false} & (s_0=\mathbf{b}) & \text{false} \\ & \text{false} & (s_1=\mathbf{b}) \\ & & \text{false} \end{pmatrix}} \text{ CHARB} \\
\\
\frac{}{\llbracket ? \rrbracket(s) = \begin{pmatrix} \text{false} & (s_0 \neq \epsilon) & \text{false} \\ & \text{false} & (s_1 \neq \epsilon) \\ & & \text{false} \end{pmatrix}} \text{ ANY} \quad \frac{\llbracket e_1 \rrbracket(s) = M \quad \llbracket e_2 \rrbracket(s) = M'}{\llbracket e_1 + e_2 \rrbracket(s) = \begin{pmatrix} (M_{0,0} \vee M'_{0,0}) & (M_{0,1} \vee M'_{0,1}) & (M_{0,2} \vee M'_{0,2}) \\ & (M_{1,1} \vee M'_{1,1}) & (M_{1,2} \vee M'_{1,2}) \\ & & (M_{2,2} \vee M'_{2,2}) \end{pmatrix}} \text{ OR} \\
\\
\frac{\llbracket e_1 \rrbracket(s) = M \quad \llbracket e_2 \rrbracket(s) = M'}{\llbracket e_1 \cdot e_2 \rrbracket(s) = \begin{pmatrix} (M_{0,0} \wedge M'_{0,0}) & \bigvee_{i=0 \dots 1} (M_{0,i} \wedge M'_{i,1}) & \bigvee_{i=0 \dots 2} (M_{0,i} \wedge M'_{i,2}) \\ & (M_{1,1} \wedge M'_{1,1}) & \bigvee_{i=1 \dots 2} (M_{1,i} \wedge M'_{i,2}) \\ & & (M_{2,2} \wedge M'_{2,2}) \end{pmatrix}} \text{ CONCAT} \\
\\
\frac{\llbracket e \rrbracket(s) = M}{\llbracket e^* \rrbracket(s) = \begin{pmatrix} \text{true} & M_{0,1} & M_{0,2} \vee (M_{0,1} \wedge M_{1,2}) \\ & \text{true} & M_{1,2} \\ & & \text{true} \end{pmatrix}} \text{ STAR} \quad \frac{\llbracket e \rrbracket(s) = M}{\llbracket e^* \rrbracket(s) = \begin{pmatrix} \neg M_{0,0} & \neg M_{0,1} & \neg M_{0,2} \\ & \neg M_{1,1} & \neg M_{1,2} \\ & & \neg M_{2,2} \end{pmatrix}} \text{ NEG}
\end{array}$$

Fig. 10. Semantics of REGEX

$$\begin{array}{c}
1373 \\
1374 \\
1375 \\
1376 \\
1377 \\
1378 \\
1379 \\
1380 \\
1381 \\
1382 \\
1383 \\
1384 \\
1385 \\
1386 \\
1387 \\
1388 \\
1389 \\
1390 \\
1391 \\
1392 \\
1393 \\
1394 \\
1395 \\
1396 \\
1397 \\
1398 \\
1399 \\
1400 \\
1401 \\
1402 \\
1403 \\
1404 \\
1405 \\
1406 \\
1407 \\
1408 \\
1409 \\
1410 \\
1411 \\
1412 \\
1413 \\
1414 \\
1415 \\
1416 \\
1417 \\
1418 \\
1419 \\
1420 \\
1421
\end{array}$$

$$\begin{array}{l}
M := \begin{pmatrix} M_{0,0} & M_{0,1} & M_{0,2} \\ & M_{1,1} & M_{1,2} \\ & & M_{2,2} \end{pmatrix} \qquad M' := \begin{pmatrix} M'_{0,0} & M'_{0,1} & M'_{0,2} \\ & M'_{1,1} & M'_{1,2} \\ & & M'_{2,2} \end{pmatrix} \qquad \frac{\llbracket e \rrbracket(s) = (M_\epsilon, M)}{\llbracket \text{eval } e \rrbracket(s) = M_{0,1}} \text{ EVAL} \\
\\
\frac{}{\llbracket \epsilon \rrbracket(s) = (\text{true}, \begin{pmatrix} \text{false} & \text{false} \\ & \text{false} \end{pmatrix})} \text{ EPS} \qquad \frac{}{\llbracket \phi \rrbracket(s) = (\text{false}, \begin{pmatrix} \text{false} & \text{false} \\ & \text{false} \end{pmatrix})} \text{ PHI} \\
\\
\frac{}{\llbracket \mathbf{a} \rrbracket(s) = (\text{false}, \begin{pmatrix} s_0=\mathbf{a} & \text{false} \\ & s_0=\mathbf{a} \end{pmatrix})} \text{ CHARA} \qquad \frac{}{\llbracket \mathbf{b} \rrbracket(s) = (\text{false}, \begin{pmatrix} s_0=\mathbf{b} & \text{false} \\ & s_0=\mathbf{b} \end{pmatrix})} \text{ CHARB} \\
\\
\frac{}{\llbracket ? \rrbracket(s) = (\text{false}, \begin{pmatrix} s_0 \neq \epsilon & \text{false} \\ & s_0 \neq \epsilon \end{pmatrix})} \text{ ANY} \qquad \frac{\llbracket e_1 \rrbracket(s) = (M_\epsilon, M) \quad \llbracket e_2 \rrbracket(s) = (M'_\epsilon, M')}{\llbracket e_1 + e_2 \rrbracket(s) = (M_\epsilon \vee M'_\epsilon, \begin{pmatrix} M_{0,0} \vee M'_{0,0} & M_{0,1} \vee M'_{0,1} \\ M_{1,1} \vee M'_{1,1} \end{pmatrix})} \text{ OR} \\
\\
\frac{\llbracket e_1 \rrbracket(s) = (M_\epsilon, M) \quad \llbracket e_2 \rrbracket(s) = (M'_\epsilon, M')}{\llbracket e_1 \cdot e_2 \rrbracket(s) = (M_\epsilon \wedge M'_\epsilon, \begin{bmatrix} (M_\epsilon \wedge M'_{0,0}) \vee (M_{0,0} \wedge M'_\epsilon) & (M_\epsilon \wedge M'_{0,1}) \vee (M_{0,0} \wedge M'_{1,1}) \vee (M_{0,1} \wedge M'_\epsilon) \\ (M_\epsilon \wedge M'_{1,1}) \vee (M_{1,1} \wedge M'_\epsilon) \end{bmatrix})} \text{ CONCAT} \\
\\
\frac{\llbracket e \rrbracket(s) = (M_\epsilon, M)}{\llbracket e^* \rrbracket(s) = (\text{true}, \begin{pmatrix} M_{0,0} & M_{0,1} \vee (M_{0,0} \wedge M_{1,1}) \\ & M_{1,1} \end{pmatrix})} \text{ STAR} \qquad \frac{\llbracket e \rrbracket(s) = (M_\epsilon, M)}{\llbracket e^* \rrbracket(s) = (\neg M_\epsilon, \begin{bmatrix} \neg M_{0,0} & \neg M_{0,1} \\ & \neg M_{1,1} \end{bmatrix})} \text{ MEG}
\end{array}$$

Fig. 11. Semantics of REGEXIMP

$$\begin{array}{c}
1422 \quad \frac{}{\llbracket 0 \rrbracket(v_0, v_1) = 0} \text{CONST0} \quad \frac{}{\llbracket 1 \rrbracket(v_0, v_1) = 1} \text{CONST1} \quad \frac{}{\llbracket \text{t} \rrbracket(v_0, v_1) = \text{true}} \text{CONSTT} \\
1423 \\
1424 \\
1425 \quad \frac{}{\llbracket \text{f} \rrbracket(v_0, v_1) = \text{false}} \text{CONSTF} \quad \frac{}{\llbracket \text{x} \rrbracket(v_0, v_1) = v_0} \text{VARX} \quad \frac{}{\llbracket \text{y} \rrbracket(v_0, v_1) = v_1} \text{VARY} \\
1426 \\
1427 \\
1428 \quad \frac{\llbracket e_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket e_2 \rrbracket(v_0, v_1) = v_2}{\llbracket e_1 + e_2 \rrbracket(v_0, v_1) = v_1 + v_2} \text{PLUS} \quad \frac{\llbracket e_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket e_2 \rrbracket(v_0, v_1) = v_2}{\llbracket e_1 - e_2 \rrbracket(v_0, v_1) = v_1 - v_2} \text{MINUS} \\
1429 \\
1430 \\
1431 \quad \frac{\llbracket e_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket e_2 \rrbracket(v_0, v_1) = v_2 \quad v_1 < v_2}{\llbracket e_1 < e_2 \rrbracket(v_0, v_1) = \text{true}} \text{LESSTHANTTRUE} \\
1432 \\
1433 \\
1434 \quad \frac{\llbracket e_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket e_2 \rrbracket(v_0, v_1) = v_2 \quad v_1 \geq v_2}{\llbracket e_1 < e_2 \rrbracket(v_0, v_1) = \text{false}} \text{LESSTHANFALSE} \\
1435 \\
1436 \\
1437 \\
1438 \quad \frac{\llbracket b_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket b_2 \rrbracket(v_0, v_1) = v_2}{\llbracket b_1 \text{ and } b_2 \rrbracket(v_0, v_1) = v_1 \wedge v_2} \text{BOOLAND} \\
1439 \\
1440 \\
1441 \quad \frac{\llbracket b_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket b_2 \rrbracket(v_0, v_1) = v_2}{\llbracket b_1 \text{ or } b_2 \rrbracket(v_0, v_1) = v_1 \vee v_2} \text{BOOLOR} \quad \frac{\llbracket b \rrbracket(v_0, v_1) = v}{\llbracket \text{not } b \rrbracket(v_0, v_1) = \neg v} \text{BOOLNOT} \\
1442 \\
1443 \\
1444 \\
1445 \quad \frac{\llbracket e \rrbracket(v_0, v_1) = v}{\llbracket \text{x} := e \rrbracket(v_0, v_1) = v} \text{ASSIGNX} \quad \frac{\llbracket e \rrbracket(v_0, v_1) = v}{\llbracket \text{y} := e \rrbracket(v_0, v_1) = v} \text{ASSIGNY} \\
1446 \\
1447 \\
1448 \\
1449 \quad \frac{}{\llbracket \text{x} ++ \rrbracket(v_0, v_1) = v_0 + 1} \text{INCX} \quad \frac{}{\llbracket \text{y} ++ e \rrbracket(v_0, v_1) = v_1 + 1} \text{INCY} \\
1450 \\
1451 \\
1452 \quad \frac{}{\llbracket \text{x} -- \rrbracket(v_0, v_1) = v_0 - 1} \text{DECX} \quad \frac{}{\llbracket \text{y} -- \rrbracket(v_0, v_1) = v_1 - 1} \text{DECY} \\
1453 \\
1454 \\
1455 \\
1456 \\
1457 \\
1458 \\
1459 \\
1460 \\
1461 \\
1462 \\
1463 \\
1464 \\
1465 \\
1466 \\
1467 \\
1468 \\
1469 \\
1470
\end{array}$$

Fig. 12. Semantics of IMP(2), part 1

$$\begin{array}{c}
1471 \\
1472 \\
1473 \\
1474 \\
1475 \\
1476 \\
1477 \\
1478 \\
1479 \\
1480 \\
1481 \\
1482 \\
1483 \\
1484 \\
1485 \\
1486 \\
1487 \\
1488 \\
1489 \\
1490 \\
1491 \\
1492 \\
1493 \\
1494 \\
1495 \\
1496 \\
1497 \\
1498 \\
1499 \\
1500 \\
1501 \\
1502 \\
1503 \\
1504 \\
1505 \\
1506 \\
1507 \\
1508 \\
1509 \\
1510 \\
1511 \\
1512 \\
1513 \\
1514 \\
1515 \\
1516 \\
1517 \\
1518 \\
1519
\end{array}$$

$$\frac{\llbracket s_1 \rrbracket(v_0, v_1) = v'_0 \quad \llbracket s_2 \rrbracket(v'_0) = v''_0}{\llbracket s_1; s_2 \rrbracket(v_0, v_1) = v''_0} \text{SEQ}$$

$$\frac{\llbracket b \rrbracket(v_0, v_1) = v \quad \llbracket s_1 \rrbracket(v_0, v_1) = v_1 \quad \llbracket s_2 \rrbracket(v_0, v_1) = v_2}{\llbracket \text{if } b \text{ then } s_1 \text{ else } s_2 \rrbracket(v_0, v_1) = v ? v_1 : v_2} \text{ITE}$$

$$\frac{\llbracket b \rrbracket(v_0, v_1) = \text{true} \quad \llbracket s \rrbracket(v_0, v_1) = v_1 \quad \llbracket \text{while } b \text{ do } s \rrbracket(v_1) = v_2}{\llbracket \text{while } b \text{ do } s \rrbracket(v_0, v_1) = v_2} \text{WHILELOOP}$$

$$\frac{\llbracket b \rrbracket(v_0, v_1) = \text{false}}{\llbracket \text{while } b \text{ do } s \rrbracket(v_0, v_1) = v_0} \text{WHILEEND}$$

$$\frac{\llbracket s \rrbracket(v_0, v_1) = v_1 \quad \llbracket b \rrbracket(v_1) = \text{true} \quad \llbracket \text{do } s \text{ while } b \rrbracket(v_1) = v_2}{\llbracket \text{do } s \text{ while } b \rrbracket(v_0, v_1) = v_2} \text{DOWHILELOOP}$$

$$\frac{\llbracket s \rrbracket(v_0, v_1) = v_1 \quad \llbracket b \rrbracket(v_1) = \text{false}}{\llbracket \text{do } s \text{ while } b \rrbracket(v_0, v_1) = v_1} \text{DOWHILEEND}$$

Fig. 13. Semantics of IMP(2), part 2

$$\begin{array}{l}
1520 \quad \frac{i = 0, 1, \dots, (k-1)}{\llbracket v_i \rrbracket(x_0, \dots, x_{k-1}) = x_i} \text{VARATOM} \quad \frac{}{\llbracket 0 \rrbracket(x_0, \dots, x_{k-1}) = 0x00000000} \text{BvZERO} \\
1521 \\
1522 \\
1523 \\
1524 \quad \frac{}{\llbracket 1 \rrbracket(x_0, \dots, x_{k-1}) = 0x00000001} \text{BvONE} \\
1525 \\
1526 \\
1527 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket e_1 < e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_1 <_{\text{unsigned}} r_2} \text{BvULT} \\
1528 \\
1529 \\
1530 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket e_1 \geq e_2 \rrbracket(x_0, \dots, x_{k-1}) = \neg(r_1 <_{\text{unsigned}} r_2)} \text{BvUGE} \\
1531 \\
1532 \\
1533 \\
1534 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket e_1 \leq e_2 \rrbracket(x_0, \dots, x_{k-1}) = \neg(r_2 <_{\text{unsigned}} r_1)} \text{BvULE} \\
1535 \\
1536 \\
1537 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2 \quad \odot \in \{\&, |, \oplus, \gg, \ll\}}{\llbracket e_1 \odot e_2 \rrbracket(x_0, \dots, x_{k-1}) = (r_1 \odot r_2)} \text{BvBITWISE} \\
1538 \\
1539 \\
1540 \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r \quad r \neq 0x00000000}{\llbracket \text{any_bit } e \rrbracket(x_0, \dots, x_{k-1}) = 0x00000001} \text{ANYBIT1} \\
1541 \\
1542 \\
1543 \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r \quad r = 0x00000000}{\llbracket \text{any_bit } e \rrbracket(x_0, \dots, x_{k-1}) = 0x00000000} \text{ANYBIT0} \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \sim e \rrbracket(x_0, \dots, x_{k-1}) = \sim e} \text{BvNOT} \\
1544 \\
1545 \\
1546 \\
1547 \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \neg e \rrbracket(x_0, \dots, x_{k-1}) = \neg e} \text{BvNEG} \\
1548 \\
1549 \\
1550 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2 \quad \odot \in \{+, -, \times, \div\}}{\llbracket e_1 \odot e_2 \rrbracket(x_0, \dots, x_{k-1}) = (r_1 \odot r_2)} \text{BvARITH} \\
1551 \\
1552 \\
1553 \\
1554 \\
1555 \\
1556 \\
1557 \\
1558 \\
1559 \\
1560 \\
1561 \\
1562 \\
1563 \\
1564 \\
1565 \\
1566 \\
1567 \\
1568
\end{array}$$

Fig. 14. Semantics of BvSIMPLE(k)

$$\begin{array}{l}
1569 \quad \frac{i = 0, 1, \dots, (k-1)}{\llbracket v_i \rrbracket(x_0, \dots, x_{k-1}) = x_i} \text{VARATOM} \quad \frac{}{\llbracket 0 \rrbracket(x_0, \dots, x_{k-1}) = 0x00000000} \text{BVZERO} \\
1570 \\
1571 \\
1572 \\
1573 \quad \frac{}{\llbracket 1 \rrbracket(x_0, \dots, x_{k-1}) = 0x00000001} \text{BVONE} \\
1574 \\
1575 \\
1576 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket e_1 < e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_1 <_{\text{unsigned}} r_2} \text{BVULT} \\
1577 \\
1578 \\
1579 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket e_1 \geq e_2 \rrbracket(x_0, \dots, x_{k-1}) = \neg(r_1 <_{\text{unsigned}} r_2)} \text{BVUGE} \\
1580 \\
1581 \\
1582 \\
1583 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2}{\llbracket e_1 \leq e_2 \rrbracket(x_0, \dots, x_{k-1}) = \neg(r_2 <_{\text{unsigned}} r_1)} \text{BVULE} \\
1584 \\
1585 \\
1586 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2 \quad \odot \in \{\&, |, \oplus, \gg, \ll\}}{\llbracket e_1 \odot e_2 \rrbracket(x_0, \dots, x_{k-1}) = (r_1 \odot r_2)} \text{BVBITWISE} \\
1587 \\
1588 \\
1589 \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r \quad r \neq 0x00000000}{\llbracket \text{any_bit } e \rrbracket(x_0, \dots, x_{k-1}) = 0x00000001} \text{ANYBIT1} \\
1590 \\
1591 \\
1592 \\
1593 \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r \quad r = 0x00000000}{\llbracket \text{any_bit } e \rrbracket(x_0, \dots, x_{k-1}) = 0x00000000} \text{ANYBIT0} \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \sim e \rrbracket(x_0, \dots, x_{k-1}) = \sim e} \text{BVNOT} \\
1594 \\
1595 \\
1596 \\
1597 \quad \frac{\llbracket e \rrbracket(x_0, \dots, x_{k-1}) = r}{\llbracket \neg e \rrbracket(x_0, \dots, x_{k-1}) = \neg e} \text{BVNEG} \\
1598 \\
1599 \\
1600 \quad \frac{\llbracket e_1 \rrbracket(x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket(x_0, \dots, x_{k-1}) = r_2 \quad \odot \in \{+, -, \times, \div\}}{\llbracket e_1 \odot e_2 \rrbracket(x_0, \dots, x_{k-1}) = (r_1 \odot_{\text{sat}} r_2)} \text{BVARITH} \\
1601 \\
1602 \quad a \odot_{\text{sat}} b = a \odot b \quad \text{(when no overflow or underflow)} \\
1603 \\
1604 \quad a \odot_{\text{sat}} b = 0xffffffff \quad \text{(when overflow happens)} \\
1605 \\
1606 \quad a \odot_{\text{sat}} b = 0x00000000 \quad \text{(when underflow happens)} \\
1607 \\
1608 \\
1609 \\
1610 \\
1611 \\
1612 \\
1613 \\
1614 \\
1615 \\
1616 \\
1617
\end{array}$$

Fig. 15. Semantics of BVSATURATED(k)

$$\begin{array}{c}
1618 \\
1619 \\
1620 \\
1621 \\
1622 \\
1623 \\
1624 \\
1625 \\
1626 \\
1627 \\
1628 \\
1629 \\
1630 \\
1631 \\
1632 \\
1633 \\
1634 \\
1635 \\
1636 \\
1637 \\
1638 \\
1639 \\
1640 \\
1641 \\
1642 \\
1643 \\
1644 \\
1645 \\
1646 \\
1647 \\
1648 \\
1649 \\
1650 \\
1651 \\
1652 \\
1653 \\
1654 \\
1655 \\
1656 \\
1657 \\
1658 \\
1659 \\
1660 \\
1661 \\
1662 \\
1663 \\
1664 \\
1665 \\
1666
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, (m-1)}{\llbracket \mathbf{v}_i \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = u_i} \text{CONSTATOM} \\
\frac{i = 0, 1, \dots, (n-1)}{\llbracket \mathbf{o}_i \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = x_i} \text{VARATOM} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r}{\llbracket \mathbf{o}_i := e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = (u_0, \dots, u_{m-1}, x_0, \dots, x_{i-1}, r, x_{i+1}, \dots, x_{n-1})} \text{ASSIGN} \\
\frac{\llbracket s_1 \rrbracket (\vec{u}, \vec{x}) = (\vec{u}', \vec{x}') \quad \llbracket s_1 \rrbracket (\vec{u}'', \vec{x}'') = (\vec{u}'', \vec{x}'')}{\llbracket s_1 ; s_2 \rrbracket (\vec{u}, \vec{x}) = (\vec{u}'', \vec{x}'')} \text{SEQ} \quad \vec{u} = (u_0, \dots, u_{m-1}), \vec{x} = (x_0, \dots, x_{n-1}) \\
\frac{}{\llbracket \mathbf{0} \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000000} \text{BVZERO} \\
\frac{}{\llbracket \mathbf{1} \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000001} \text{BVONE} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2}{\llbracket e_1 < e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 <_{\text{unsigned}} r_2} \text{BVULT} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2}{\llbracket e_1 \geq e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \neg(r_1 <_{\text{unsigned}} r_2)} \text{BVUGE} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2}{\llbracket e_1 \leq e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \neg(r_2 <_{\text{unsigned}} r_1)} \text{BVULE} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2 \quad \odot \in \{\&, |, \oplus, \gg, \ll\}}{\llbracket e_1 \odot e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = (r_1 \odot r_2)} \text{BVBITWISE} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r \quad r \neq 0\text{x}00000000}{\llbracket \text{any_bit } e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000001} \text{ANYBIT1} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r \quad r = 0\text{x}00000000}{\llbracket \text{any_bit } e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000000} \text{ANYBIT0} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r}{\llbracket \sim e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \sim e} \text{BVNOT} \quad \frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r}{\llbracket \neg e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \neg e} \text{BVNEG} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2 \quad \odot \in \{+, -, \times, \div\}}{\llbracket e_1 \odot e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = (r_1 \odot r_2)} \text{BVARITH}
\end{array}$$

Fig. 16. Semantics of BVIMP SIMPLE(k)

$$\begin{array}{c}
1667 \\
1668 \\
1669 \\
1670 \\
1671 \\
1672 \\
1673 \\
1674 \\
1675 \\
1676 \\
1677 \\
1678 \\
1679 \\
1680 \\
1681 \\
1682 \\
1683 \\
1684 \\
1685 \\
1686 \\
1687 \\
1688 \\
1689 \\
1690 \\
1691 \\
1692 \\
1693 \\
1694 \\
1695 \\
1696 \\
1697 \\
1698 \\
1699 \\
1700 \\
1701 \\
1702 \\
1703 \\
1704 \\
1705 \\
1706 \\
1707 \\
1708 \\
1709 \\
1710 \\
1711 \\
1712 \\
1713 \\
1714 \\
1715
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, (m-1)}{\llbracket \mathbf{v}_i \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = u_i} \text{CONSTATOM} \\
\frac{i = 0, 1, \dots, (n-1)}{\llbracket \mathbf{o}_i \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = x_i} \text{VARATOM} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r}{\llbracket \mathbf{o}_i := e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = (u_0, \dots, u_{m-1}, x_0, \dots, x_{i-1}, r, x_{i+1}, \dots, x_{n-1})} \text{ASSIGN} \\
\frac{\llbracket s_1 \rrbracket (\vec{u}, \vec{x}) = (\vec{u}', \vec{x}') \quad \llbracket s_1 \rrbracket (\vec{u}'', \vec{x}'') = (\vec{u}'', \vec{x}'')}{\llbracket s_1 ; s_2 \rrbracket (\vec{u}, \vec{x}) = (\vec{u}'', \vec{x}'')} \text{SEQ} \quad \vec{u} = (u_0, \dots, u_{m-1}), \vec{x} = (x_0, \dots, x_{n-1}) \\
\frac{}{\llbracket \mathbf{0} \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000000} \text{BVZERO} \\
\frac{}{\llbracket \mathbf{1} \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000001} \text{BVONE} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2}{\llbracket e_1 < e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 <_{\text{unsigned}} r_2} \text{BVULT} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2}{\llbracket e_1 \geq e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \neg(r_1 <_{\text{unsigned}} r_2)} \text{BVUGE} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2}{\llbracket e_1 \leq e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \neg(r_2 <_{\text{unsigned}} r_1)} \text{BVULE} \\
\frac{\llbracket e_1 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_1 \quad \llbracket e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r_2 \quad \odot \in \{\&, |, \oplus, \gg, \ll\}}{\llbracket e_1 \odot e_2 \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = (r_1 \odot r_2)} \text{BVBITWISE} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r \quad r \neq 0\text{x}00000000}{\llbracket \text{any_bit } e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000001} \text{ANYBIT1} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r \quad r = 0\text{x}00000000}{\llbracket \text{any_bit } e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = 0\text{x}00000000} \text{ANYBIT0} \\
\frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r}{\llbracket \sim e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \sim e} \text{BVNOT} \quad \frac{\llbracket e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = r}{\llbracket \neg e \rrbracket (u_0, \dots, u_{m-1}, x_0, \dots, x_{n-1}) = \neg e} \text{BVNEG} \\
\frac{\llbracket e_1 \rrbracket (x_0, \dots, x_{k-1}) = r_1 \quad \llbracket e_2 \rrbracket (x_0, \dots, x_{k-1}) = r_2 \quad \odot \in \{+, -, \times, \div\}}{\llbracket e_1 \odot e_2 \rrbracket (x_0, \dots, x_{k-1}) = (r_1 \odot_{\text{sat}} r_2)} \text{BVARITH}
\end{array}$$

Fig. 17. Semantics of $\text{BvIMP SATURATED}(k)$
Proc. ACM Program. Lang., Vol. 1, No. CONF, Article 1. Publication date: January 2018.

$$\begin{array}{l}
1716 \\
1717 \quad \overline{\llbracket 1 \rrbracket(v_0) = 1} \text{ ONE} \quad \overline{\llbracket 0 \rrbracket(v_0) = 0} \text{ ZERO} \quad \overline{\llbracket x \rrbracket(v_0) = v_0} \text{ VARX} \quad \overline{\llbracket n \rrbracket(v_0) = r} \text{ COUNTBIT} \\
1718 \quad \overline{\llbracket \text{count } n \rrbracket(v_0) = r} \\
1719 \\
1720 \quad \overline{\llbracket m \rrbracket(v_0) = r} \text{ BINTODEC} \quad \overline{\llbracket n \rrbracket(v_0) = r_1 \quad \llbracket b \rrbracket(v_0) = r_2} \text{ COUNTBITCONCAT} \\
1721 \quad \overline{\llbracket \text{bin2dec } m \rrbracket(v_0) = r} \quad \overline{\llbracket \text{concat } n b \rrbracket(v_0) = r_1 + \text{int}(r_2)} \\
1722 \\
1723 \quad \overline{\llbracket m \rrbracket(v_0) = r_1 \quad \llbracket b \rrbracket(v_0) = r_2} \text{ BINTODECCONCAT} \quad \overline{\llbracket b \rrbracket(v_0) = r} \text{ COUNTBITATOM} \\
1724 \quad \overline{\llbracket \text{concat}' m b \rrbracket(v_0) = 2 \cdot r_1 + \text{int}(r_2)} \quad \overline{\llbracket \text{atom } b \rrbracket(v_0) = r} \\
1725 \\
1726 \\
1727 \quad \overline{\llbracket b \rrbracket(v_0) = r} \text{ BINTODECATOM} \\
1728 \quad \overline{\llbracket \text{atom}' b \rrbracket(v_0) = r} \\
1729
\end{array}$$

Fig. 18. Semantics of BINOP

$$\begin{array}{l}
1730 \\
1731 \\
1732 \\
1733 \quad \overline{\llbracket 0 \rrbracket(v_0) = 0} \text{ ZERO} \quad \overline{\llbracket 1 \rrbracket(v_0) = 1} \text{ ONE} \quad \overline{\llbracket 2 \rrbracket(v_0) = 2} \text{ TWO} \quad \overline{\llbracket 4 \rrbracket(v_0) = 4} \text{ FOUR} \\
1734 \\
1735 \\
1736 \quad \overline{\llbracket 8 \rrbracket(v_0) = 8} \text{ EIGHT} \quad \overline{\llbracket x \rrbracket(v_0) = v_0} \text{ VARX} \quad \overline{\llbracket k_1 \rrbracket(v_0) = r_1 \llbracket k_2 \rrbracket(v_0) = r_2} \text{ SCALARPLUS} \\
1737 \quad \overline{\llbracket k_1 +_k k_2 \rrbracket(v_0) = r_1 + r_2} \\
1738 \\
1739 \quad \overline{\llbracket s_1 \rrbracket(v_0) = r_1 \llbracket s_2 \rrbracket(v_0) = r_2} \text{ CURRENCYPLUS} \quad \overline{\llbracket s_1 \rrbracket(v_0) = r_1 \llbracket s_2 \rrbracket(v_0) = r_2} \text{ CURRENCYSUBTRACT} \\
1740 \quad \overline{\llbracket s_1 + s_2 \rrbracket(v_0) = r_1 + r_2} \quad \overline{\llbracket s_1 - s_2 \rrbracket(v_0) = r_1 - r_2} \\
1741 \\
1742 \quad \overline{\llbracket s \rrbracket(v_0) = r_1 \llbracket k \rrbracket(v_0) = r_2} \text{ CURRENCYTIMESSCALAR} \quad \overline{\llbracket k \rrbracket(v_0) = r} \text{ CURRENCYJPY} \\
1743 \quad \overline{\llbracket s \times k \rrbracket(v_0) = r_1 \cdot r_2} \quad \overline{\llbracket \text{jpy } k \rrbracket(v_0) = r} \\
1744 \\
1745 \\
1746 \quad \overline{\llbracket k \rrbracket(v_0) = r} \text{ CURRENCYCNY} \quad \overline{\llbracket k \rrbracket(v_0) = r} \text{ CURRENCYUSD} \\
1747 \quad \overline{\llbracket \text{cny } k \rrbracket(v_0) = 21 \cdot r} \quad \overline{\llbracket \text{usd } k \rrbracket(v_0) = 152 \cdot r} \\
1748 \\
1749 \\
1750
\end{array}$$

Fig. 19. Semantics of CURRENCY

$$\begin{array}{l}
1751 \\
1752 \quad \overline{\llbracket 0 \rrbracket(v_0) = (0, 0, 0)} \text{ ZERO} \quad \overline{\llbracket 1 \rrbracket(v_0) = (1, 1, 0)} \text{ ONE} \quad \overline{\llbracket 2 \rrbracket(v_0) = (2, 2, 0)} \text{ TWO} \\
1753 \\
1754 \\
1755 \quad \overline{\llbracket x \rrbracket(v_0) = (v_0, v_0 + 1, 0)} \text{ VARX} \quad \overline{\llbracket e_1 \rrbracket(v_0) = (r_1, s_1, t_1) \llbracket e_2 \rrbracket(v_0) = (r_2, s_2, t_2)} \text{ PLUS} \\
1756 \quad \overline{\llbracket e_1 + e_2 \rrbracket(v_0) = (r_1 + r_2, s_1 + s_2, t_1 + t_2)} \\
1757 \\
1758 \quad \overline{\llbracket e_1 \rrbracket(v_0) = (r_1, s_1, t_1) \llbracket e_2 \rrbracket(v_0) = (r_2, s_2, t_2)} \text{ MULTIPLY} \\
1759 \quad \overline{\llbracket e_1 \times e_2 \rrbracket(v_0) = (r_1 \cdot r_2, s_1 \cdot s_2, r_1 \cdot t_2 + r_2 \cdot s_1)} \\
1760 \\
1761 \\
1762 \\
1763 \\
1764
\end{array}$$

Fig. 20. Semantics of DIFF

$$\begin{array}{c}
1765 \\
1766 \\
1767 \\
1768 \\
1769 \\
1770 \\
1771 \\
1772 \\
1773 \\
1774 \\
1775 \\
1776 \\
1777 \\
1778 \\
1779 \\
1780 \\
1781 \\
1782 \\
1783 \\
1784 \\
1785 \\
1786 \\
1787 \\
1788 \\
1789 \\
1790 \\
1791 \\
1792 \\
1793 \\
1794 \\
1795 \\
1796 \\
1797 \\
1798 \\
1799 \\
1800 \\
1801 \\
1802 \\
1803 \\
1804 \\
1805 \\
1806 \\
1807 \\
1808 \\
1809 \\
1810 \\
1811 \\
1812 \\
1813
\end{array}$$

$$\begin{array}{c}
\frac{i = 0, 1, \dots, 8}{\llbracket i \rrbracket(v_0) = i} \text{INTLITERAL} \quad \frac{}{\llbracket x \rrbracket(v_0) = v_0} \text{VARX} \quad \frac{\llbracket e \rrbracket(v_0) = r}{\llbracket \text{expr } e \rrbracket(v_0) = r} \text{EXPR} \\
\\
\frac{\llbracket b \rrbracket(v_0) = r_0 \quad \llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2 \quad r_0}{\llbracket \text{ite } b \ e_1 \ e_2 \rrbracket(v_0) = r_1} \text{ITE1} \\
\\
\frac{\llbracket b \rrbracket(v_0) = r_0 \quad \llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2 \quad \neg r_0}{\llbracket \text{ite } b \ e_1 \ e_2 \rrbracket(v_0) = r_2} \text{ITE2} \\
\\
\frac{\llbracket e \rrbracket(v_0) = r_1 \quad \llbracket f \rrbracket(v_0) = r_2}{\llbracket e + f \rrbracket(v_0) = r_1 + r_2} \text{PLUS} \quad \frac{\llbracket e \rrbracket(v_0) = r_1 \quad \llbracket f \rrbracket(v_0) = r_2}{\llbracket e - f \rrbracket(v_0) = r_1 - r_2} \text{MINUS} \\
\\
\frac{\llbracket f \rrbracket(v_0) = r}{\llbracket \text{atom } f \rrbracket(v_0) = r} \text{ATOM} \quad \frac{\llbracket f \rrbracket(v_0) = r_1 \quad \llbracket g \rrbracket(v_0) = r_2}{\llbracket f * g \rrbracket(v_0) = r_1 \cdot r_2} \text{MULTIPLY} \\
\\
\frac{\llbracket f \rrbracket(v_0) = r_1 \quad \llbracket g \rrbracket(v_0) = r_2}{\llbracket f \div g \rrbracket(v_0) = r_1 \div r_2} \text{DIVIDE} \quad \frac{\llbracket g \rrbracket(v_0) = r}{\llbracket \text{num } g \rrbracket(v_0) = r} \text{ATOM} \\
\\
\frac{\llbracket e_1 \rrbracket(v_0) = r_1 \quad \llbracket e_2 \rrbracket(v_0) = r_2 \quad \Theta \in \{<, \leq, >, \geq, =, \neq\}}{\llbracket e_1 \Theta e_2 \rrbracket(v_0) = (r_1 \Theta r_2)} \text{CMP}
\end{array}$$

Fig. 21. Semantics of ITEXPR

Table 2. Evaluation results with optimization turned on.³

Lang.	Rule	# Iter.	# Ex	SyGuS (s)	SMT (s)	Total (s)
1814						
1815						
1816						
1817	$E \rightarrow 0$	1	1	0.01	0.02	0.29
1818	$E \rightarrow 1$	1	1	0.01	0.01	0.14
1819	$E \rightarrow v_0$	1	1	0.01	0.01	0.13
1819	$E \rightarrow v_1$	1	1	0.01	0.02	0.11
1819	$E \rightarrow v_2$	1	1	0.02	0.01	0.09
1820	$E \rightarrow \neg E$	2	2	0.09	1.09	1.92
1821	$E \rightarrow \sim E$	2	2	0.06	0.88	1.65
1821	$E \rightarrow \text{any bit } E$	4	4	1.93	0.93	3.64
1822	$E \rightarrow E + E$	9	2	1.98	0.89	4.12
1822	$E \rightarrow E \& E$	6	2	2.67	1.47	7.31
1823	$E \rightarrow E \div E$	6	2	1.39	37.08	62.33
1824	$E \rightarrow E = E$	19	6	69.51	4.16	77.67
1824	$E \rightarrow E \gg E$	9	3	1.72	2.05	5.54
1825	$E \rightarrow E \times E$	9	2	2.31	0.69	3.98
1826	$E \rightarrow E E$	10	3	2.55	1.19	4.98
1826	$E \rightarrow E \ll E$	9	3	1.60	3.13	6.71
1827	$E \rightarrow E - E$	9	2	1.28	1.61	4.61
1827	$S \rightarrow E \geq E$	14	6	4.73	2.34	10.39
1828	$S \rightarrow E \leq E$	13	5	2.55	1.84	6.75
1828	$S \rightarrow E < E$	6	5	0.46	1.54	3.62
1829	$E \rightarrow E \oplus E$	9	3	2.00	1.56	5.10
1830						
1830	$E \rightarrow 0$	1	1	0.01	0.01	0.15
1831	$E \rightarrow 1$	1	1	0.01	0.01	0.06
1831	$E \rightarrow v_0$	1	1	0.01	0.01	0.05
1832	$E \rightarrow v_1$	1	1	0.01	0.01	0.05
1833	$E \rightarrow E \& E$	5	3	0.46	0.94	3.41
1833	$E \rightarrow E \gg E$	6.0	3.0	1.07	2.67	5.65
1834	$E \rightarrow E \times E$	10.0	6.0	404.76	3.09	414.49
1834	$E \rightarrow E E$	5.5	3.0	1.62	1.30	4.31
1835	$E \rightarrow E \ll E$	6.0	3.0	0.70	2.03	4.41
1835	$E \rightarrow E - E$	6.0	5.0	5.46	3.07	12.22
1836	$E \rightarrow E \oplus E$	5.5	2.5	0.64	1.03	2.94
1837						
1837	$E \rightarrow 0$	1	1	0.01	0.02	0.36
1838	$E \rightarrow 1$	1	1	0.01	0.03	0.17
1838	$E \rightarrow o0$	1	1	0.01	0.02	0.10
1839	$E \rightarrow o1$	1	1	0.01	0.01	0.07
1840	$E \rightarrow v0$	1	1	0.01	0.01	0.14
1840	$S \rightarrow o0 := E$	1	1	0.35	0.49	1.27
1840	$S \rightarrow o1 := E$	2	2	0.32	0.38	1.47
1841	$E \rightarrow \neg E$	2	2	0.07	0.79	1.64
1841	$E \rightarrow \sim E$	2	2	0.08	0.61	1.47
1842	$E \rightarrow \text{any bit } E$	4	4	1.30	0.79	2.72
1843	$E \rightarrow E + E$	5	2	1.19	0.92	3.48
1843	$E \rightarrow E \& E$	7	3	3.65	1.46	8.17
1844	$E \rightarrow E \div E$	7	3	2.58	35.00	60.76
1845	$E \rightarrow E = E$	20	6	83.47	19.99	108.10
1846	$E \rightarrow E \gg E$	9	3	2.52	2.19	6.40
1846	$E \rightarrow E \times E$	9	3	2.48	0.86	4.39
1847	$E \rightarrow E E$	9	3	2.00	1.10	4.30
1847	$E \rightarrow E \ll E$	10	3	1.83	2.59	6.72
1848	$E \rightarrow E - E$	7	2	1.71	1.52	4.99
1848	$B \rightarrow S \geq S$	26	8	18.24	2.87	25.36
1849	$B \rightarrow S \leq S$	23	6	13.59	1.06	16.75
1849	$B \rightarrow S < S$	6	5	0.33	1.02	2.67
1850	$E \rightarrow E \oplus E$	6	2	1.34	1.10	3.50
1850	$S \rightarrow S ; S$	13	3	535.19	15.61	590.21
1851						
1851	$V \rightarrow v0$	2	2	0.01	0.01	0.12
1852	$V \rightarrow v1$	2	2	0.01	0.01	0.06
1853	$V \rightarrow v10$	3	3	0.02	0.01	0.05
1853	$V \rightarrow v2$	2	2	0.01	0.01	0.05
1854	$V \rightarrow v3$	2	2	0.01	0.01	0.05
1854	$V \rightarrow v4$	3	3	0.01	0.01	0.03
1855	$V \rightarrow v5$	3	3	0.01	0.01	0.04
1856	$V \rightarrow v6$	3	3	0.01	0.01	0.06
1856	$V \rightarrow v7$	3	4	0.02	0.01	0.06
1857	$V \rightarrow v8$	3	3	0.02	0.01	0.05
1857	$V \rightarrow v9$	3	4	0.01	0.01	0.05
1858	$B \rightarrow \text{var } V$	4	4	0.06	0.51	1.07
1858	$B \rightarrow B \wedge B$	116	8	1810.40	4.43	1831.92
1859						
1860						
1861						
1862						

³Note: The label (Ti) in the language name means the language timeouts under i runs.

1863							
1864	Dirf(T4)	$E \rightarrow 0$	1	1	0.01	0.01	0.11
1865		$E \rightarrow 1$	1	1	0.01	0.01	0.06
1866		$E \rightarrow 2$	1	1	0.02	0.01	0.06
1867		$E \rightarrow x$	2	2	0.35	0.01	0.41
1868		$E \rightarrow E \times E$	5	5	92.68	1.00	95.15
1869		$E \rightarrow E + E$	3	3	9.12	2.09	13.47
1870	BVIMPSAr.(1,2)(T7)	$E \rightarrow 0$	1	1	0.01	0.02	0.37
1871		$E \rightarrow 1$	1	1	0.01	0.01	0.14
1872		$E \rightarrow o0$	1	1	0.01	0.01	0.12
1873		$E \rightarrow o1$	1	1	0.01	0.01	0.09
1874		$E \rightarrow v0$	1	1	0.01	0.02	0.14
1875		$E \rightarrow E + E$	16	3	31.60	1.11	34.15
1876	$E \rightarrow E \& E$	6	3	1.75	1.89	7.24	
1877	CNF(8)	$V \rightarrow v0$	2	2	0.01	0.01	0.12
1878		$V \rightarrow v1$	2	2	0.01	0.01	0.04
1879		$V \rightarrow v2$	2	2	0.01	0.01	0.05
1880		$V \rightarrow v3$	2	3	0.01	0.01	0.04
1881		$V \rightarrow v4$	2	2	0.01	0.01	0.03
1882		$V \rightarrow v5$	3	3	0.01	0.01	0.04
1883		$V \rightarrow v6$	3	3	0.01	0.01	0.04
1884		$V \rightarrow v7$	3	4	0.01	0.01	0.04
1885		$B \rightarrow \text{clause } C$	4	4	0.03	0.27	0.48
1886		$C \rightarrow \text{nvar } V$	5	5	0.05	0.32	0.70
1887		$C \rightarrow \text{var } V$	4	4	0.05	0.31	0.74
1888	$B \rightarrow C \wedge B$	39	6	30.52	0.56	31.83	
1889	$C \rightarrow V \vee C$	41	8	37.03	0.69	38.62	
1890	DNF(8)	$V \rightarrow v0$	2	2	0.01	0.01	0.13
1891		$V \rightarrow v1$	2	2	0.01	0.01	0.04
1892		$V \rightarrow v2$	2	2	0.01	0.01	0.04
1893		$V \rightarrow v3$	2	3	0.01	0.01	0.04
1894		$V \rightarrow v4$	2	2	0.01	0.01	0.03
1895		$V \rightarrow v5$	3	3	0.01	0.01	0.04
1896		$V \rightarrow v6$	3	3	0.01	0.01	0.03
1897		$V \rightarrow v7$	3	4	0.01	0.01	0.05
1898		$B \rightarrow \text{conj } C$	4	4	0.05	0.30	0.57
1899		$C \rightarrow \text{nvar } V$	5	5	0.05	0.33	0.71
1900		$C \rightarrow \text{var } V$	4	4	0.05	0.32	0.76
1901	$C \rightarrow V \wedge C$	33	7	28.84	0.36	29.75	
1902	$B \rightarrow C \vee B$	72	6	93.47	0.79	95.62	
1903	Imp(2)	$E \rightarrow 0$	1	1	0.01	0.01	0.05
1904		$E \rightarrow 1$	1	1	0.01	0.01	0.04
1905		$S \rightarrow x --$	2	2	0.06	0.02	0.11
1906		$S \rightarrow y --$	2	2	0.11	0.03	0.17
1907		$B \rightarrow f$	1	1	0.01	0.01	0.06
1908		$S \rightarrow x ++$	2	2	0.04	0.03	0.11
1909		$S \rightarrow y ++$	2	2	0.12	0.02	0.16
1910		$B \rightarrow t$	1	1	0.01	0.02	0.13
1911		$E \rightarrow x$	2	2	0.01	0.01	0.04
1912		$E \rightarrow y$	1	1	0.01	0.01	0.04
1913		$S \rightarrow x := E$	2	2	0.10	3.23	6.17
1914		$S \rightarrow y := E$	2	2	0.04	3.22	6.19
1915		$B \rightarrow \neg B$	3	3	0.02	2.49	5.26
1916		$E \rightarrow E + E$	4	3	0.05	8.52	14.83
1917		$E \rightarrow E - E$	5	2	0.13	8.03	13.83
1918		$B \rightarrow E < E$	8	5	0.08	7.50	13.66
1919		$B \rightarrow B \wedge B$	4	4	0.03	5.33	11.71
1920		$B \rightarrow B \vee B$	4	4	0.05	4.61	8.99
1921	$S \rightarrow S ; S$	5	3	4.55	15.00	72.53	
1922	$S \rightarrow \text{do_while } S B$	27	35	858.50	257.33	1374.13	
1923	$S \rightarrow \text{while } B S$	9	7	16.88	122.41	266.80	
1924	$S \rightarrow \text{ite } B S S$	11	5	525.28	33.88	628.71	

1912							
1913		$E \rightarrow 0$	1	1	0.01	0.01	0.05
1914		$E \rightarrow 1$	1	1	0.01	0.01	0.04
1915		$E \rightarrow 2$	1	1	0.01	0.01	0.05
1916		$E \rightarrow 3$	1	1	0.03	0.04	0.10
1917		$B \rightarrow f$	1	1	0.01	0.02	0.07
1918		$B \rightarrow t$	1	1	0.01	0.03	0.16
1919		$E \rightarrow x$	2	2	0.01	0.03	0.09
1920		$E \rightarrow y$	2	2	0.01	0.02	0.07
1921		$E \rightarrow z$	2	2	0.02	0.03	0.09
1922		$B \rightarrow \neg B$	4	4	0.06	5.09	15.22
1923		$E \rightarrow E \times E$	3	3	1.38	12.51	22.58
1924		$E \rightarrow E + E$	3	3	1.60	11.42	21.82
1925		$B \rightarrow E < E$	6	6	0.73	11.20	26.87
1926		$B \rightarrow B \wedge B$	5	5	0.08	8.08	14.95
1927		$B \rightarrow B \vee B$	4	4	0.05	7.70	14.19
1928		$E \rightarrow \text{ite } B E E$	4	4	0.78	13.54	31.00
1929							
1930		$G \rightarrow 0$	1	1	0.01	0.01	0.27
1931		$G \rightarrow 1$	1	1	0.01	0.01	0.11
1932		$G \rightarrow 2$	1	1	0.01	0.01	0.09
1933		$G \rightarrow 3$	1	1	0.05	0.01	0.13
1934		$G \rightarrow 4$	1	1	0.01	0.01	0.10
1935		$G \rightarrow 5$	1	1	0.05	0.01	0.13
1936		$G \rightarrow 6$	1	1	0.09	0.01	0.18
1937		$G \rightarrow 7$	1	1	0.18	0.01	0.23
1938		$G \rightarrow 8$	1	1	0.01	0.01	0.05
1939		$G \rightarrow x$	1	1	0.02	0.01	0.07
1940		$E \rightarrow \text{atom } F$	2	2	0.03	0.25	0.55
1941		$S \rightarrow \text{expr } E$	1	1	0.02	0.10	0.20
1942		$F \rightarrow \text{num } G$	1	1	0.03	0.30	0.69
1943		$F \rightarrow F \times G$	2	2	1.19	0.77	3.33
1944		$E \rightarrow E + F$	2	2	1.25	0.25	1.79
1945		$E \rightarrow E - F$	2	2	1.12	0.27	1.68
1946		$F \rightarrow F \div G$	4	3	1.92	0.94	3.88
1947		$B \rightarrow E = E$	5	4	0.09	0.24	0.71
1948		$B \rightarrow E \geq E$	5	5	1.79	0.33	2.60
1949		$B \rightarrow E > E$	5	5	0.18	0.26	0.79
1950		$B \rightarrow E \leq E$	6	6	0.24	0.56	1.48
1951		$B \rightarrow E < E$	5	5	0.12	0.30	0.85
1952		$B \rightarrow E \neq E$	6	6	5.35	0.26	6.11
1953		$S \rightarrow \text{ite } B E E$	3	3	0.29	0.29	0.92
1954							
1955		$R \rightarrow ?$	3	3	3.84	0.07	4.07
1956		$R \rightarrow a$	4	4	11.10	0.07	11.53
1957		$R \rightarrow b$	5	5	11.63	0.06	12.01
1958		$R \rightarrow \epsilon$	1	1	0.07	0.07	2.38
1959		$R \rightarrow \emptyset$	1	1	0.19	0.07	0.46
1960		$\text{Start} \rightarrow \text{eval } R$	3	3	0.02	4.43	13.40
1961		$R \rightarrow !R$	5	5	2.85	15.77	77.36
1962		$R \rightarrow R^*$	6	6	0.99	13.06	31.91
1963		$R \rightarrow R \cdot R$	24	24	333.71	72.58	495.45
1964		$R \rightarrow R R$	10	10	10.96	59.54	140.82
1965							
1966		$B \rightarrow 0$	1	1	0.01	0.01	0.07
1967		$B \rightarrow 1$	1	1	0.01	0.01	0.22
1968		$B \rightarrow x$	2	2	0.01	0.01	0.08
1969		$N \rightarrow \text{atom } B$	2	2	0.09	0.04	0.30
1970		$M \rightarrow \text{atom}' B$	3	3	0.07	0.05	0.26
1971		$S \rightarrow \text{bin2dec } M$	2	2	0.02	0.09	0.30
1972		$S \rightarrow \text{count } N$	2	2	0.04	0.05	0.24
1973		$N \rightarrow \text{concat } N B$	5	5	8.61	0.22	10.31
1974		$M \rightarrow \text{concat}' M B$	5	5	288.81	0.23	308.50
1975							
1976		$K \rightarrow 0$	1	1	0.01	0.01	0.03
1977		$K \rightarrow 1$	1	1	0.01	0.01	0.02
1978		$K \rightarrow 2$	1	1	0.01	0.01	0.02
1979		$K \rightarrow 4$	1	1	0.01	0.01	0.02
1980		$K \rightarrow 8$	1	1	0.01	0.01	0.02
1981		$K \rightarrow x$	2	2	0.01	0.01	0.10
1982		$S \rightarrow \text{cny } K$	3	3	21.55	0.28	22.48
1983		$S \rightarrow \text{jpy } K$	1	1	0.01	0.10	0.22
1984		$S \rightarrow \text{usd } K$	2	2	19.65	0.21	20.27
1985		$S \rightarrow S \times K$	2	2	0.03	0.11	0.24
1986		$S \rightarrow S + S$	2	2	0.03	0.14	0.33
1987		$S \rightarrow S - S$	2	2	0.04	0.16	0.38
1988		$K \rightarrow K +_k K$	3	3	0.28	0.35	1.19
1989							
1990							
1991							
1992							
1993							
1994							
1995							
1996							
1997							
1998							
1999							
2000							