

Symbolic Analysis via Semantic Reinterpretation

Junghee Lim¹, Akash Lal¹, and Thomas Reps^{1,2}

¹ University of Wisconsin; Madison, WI; USA. {junghee, akash, reps}@cs.wisc.edu

² GrammaTech, Inc.; Ithaca, NY; USA.

Abstract. The paper presents a novel technique to create implementations of the basic primitives used in symbolic program analysis: *forward symbolic evaluation*, *weakest liberal precondition*, and *symbolic composition*. We used the technique to create a system in which, for the cost of writing just *one* specification—an interpreter for the programming language of interest—one obtains automatically-generated, mutually-consistent implementations of all *three* symbolic-analysis primitives. This can be carried out even for languages with pointers and address arithmetic. Our implementation has been used to generate symbolic-analysis primitives for x86 and PowerPC.

1 Introduction

The use of symbolic-reasoning primitives for *forward symbolic evaluation*, *weakest liberal precondition* (\mathcal{WLP}), and *symbolic composition* has experienced a resurgence in program-analysis tools because of the power that they provide when exploring a program’s state space. The semantics of such operations are easy to state; for instance, if $\tau(\sigma, \sigma')$ is a formula that represents the semantics of an instruction, then $\mathcal{WLP}(\tau, \phi)$ is $\forall \sigma'. (\tau(\sigma, \sigma') \Rightarrow \phi(\sigma'))$. However, this formula uses quantification over states—i.e., *second-order quantification*—whereas SMT solvers, such as Yices and Z3, support only *quantifier-free first-order* logic. Hence, such a formula cannot be used directly.

For a simple language that has only `int`-valued variables, it is easy to recast matters in first-order logic. For instance, the \mathcal{WLP} of postcondition ϕ with respect to an assignment statement $var = rhs$; can be obtained by substituting rhs for all (free) occurrences of var in ϕ : $\phi[var \leftarrow rhs]$. For real-world programming languages, however, the situation is more complicated. For instance, for languages with pointers, Morris’s rule of substitution [11] requires taking into account all possible aliasing combinations.

The standard approach to implementing each of the symbolic-analysis primitives for a programming language of interest (which we call the *subject* language) is to create hand-written *translation procedures*—one per symbolic-analysis primitive—that convert subject-language commands into appropriate formulas. With this approach, a system can contain subtle inconsistency bugs if the different translation procedures adopt different “views” of the semantics. Such bugs are easy to introduce because each translation procedure must encode the subject language’s semantics; however, the encodings for symbolic execution, \mathcal{WLP} , and symbolic composition have different flavors.

Our own interest is in analyzing machine code, such as x86 and PowerPC. Unfortunately, machine-code instruction sets have hundreds of instructions, as well

as other complicating factors, such as the use of separate instructions to set flags (based on the condition that is tested) and to branch according to the flag values; the ability to perform address arithmetic and dereference computed addresses; etc. To appreciate the need for tool support for creating symbolic-analysis primitives for real machine-code languages, consult Section 3.2 of the Intel manual (<http://download.intel.com/design/processor/manuals/253666.pdf>), and imagine writing three separate encodings of each instruction’s semantics to implement symbolic execution, \mathcal{WLP} , and symbolic composition. Some tools (e.g., [7, 3]) need an instruction-set emulator, in which case a fourth encoding of the semantics is also required.

To address these issues, this paper presents a way to automatically obtain mutually-consistent, correct-by-construction implementations of symbolic primitives, by *generating* them from a specification of the subject language’s concrete semantics. More precisely, we present a method to obtain quantifier-free, first-order-logic formulas for (a) symbolic execution of a single command, (b) \mathcal{WLP} with respect to a single command, and (c) symbolic composition for a class of formulas that express state transformations. The generated implementations are guaranteed to be mutually consistent, and also to be consistent with an instruction-set emulator (for concrete execution) that is generated from the same specification of the subject language’s concrete semantics.

Primitives (a) and (b) immediately extend to compound operations over a given program path for use in forward and backwards symbolic execution, respectively; see §6. (The design of client algorithms that use such primitives to perform state-space exploration is an orthogonal issue that is outside the scope of this paper.)

Semantic Reinterpretation. Our approach is based on factoring the concrete semantics of a language into two parts: (i) a *client* specification, and (ii) a semantic *core*. The interface to the core consists of certain base types, function types, and operators, and the client is expressed in terms of this interface. Such an organization permits the core to be *reinterpreted* to produce an alternative semantics for the subject language. The idea of exploiting such a factoring comes from the field of abstract interpretation [4], where semantic reinterpretation has been proposed as a convenient tool for formulating abstract interpretations [12, 10] (see §2).

Achievements and Contributions. We used the approach described in the paper to create a “Yacc-like” tool for generating mutually-consistent, correct-by-construction implementations of symbolic-analysis primitives for instruction sets (§7). The input is a specification of an instruction set’s concrete semantics; the output is a triple of C++ functions that implement the three symbolic-analysis primitives. The tool has been used to generate such primitives for x86 and PowerPC. To accomplish this, we leveraged an existing tool, TSL [9], as the implementation platform for defining the necessary reinterpretations. However, we wish to stress that the ideas presented in the paper are not TSL-specific; other ways of implementing the necessary reinterpretations are possible (see §2).

The contributions of this paper lie in the insights that went into defining the specific reinterpretations that we use to obtain mutually-consistent, correct-by-construction implementations of the symbolic-analysis primitives, and the discovery that \mathcal{WLP} could be obtained by using two different reinterpretations working in tandem. The paper’s other contributions can be summarized as follows:

- We present a new application for semantic reinterpretation, namely, to create implementations of the basic primitives for symbolic reasoning (§4 and §5). In particular, two key insights allowed us to obtain the primitives for \mathcal{WLP} and symbolic composition. The first insight was that we could apply semantic reinterpretation in a new context, namely, to the interpretation function of a *logic* (§4). The second insight was to define a particular form of state-transformation formula to be a first-class notion in the logic, which allows such formulas (i) to be reinterpreted, and (ii) to serve as a replacement domain in various reinterpretations (§4).
- We show how reinterpretation can automatically create a \mathcal{WLP} primitive that implements Morris’s rule of substitution for a language with pointers [11] (§4).
- We conducted an experiment on real x86 code using the generated primitives (§7).

For expository purposes, simplified languages are used throughout. Our discussion of machine code (§3.3 and §5) is based on a greatly simplified fragment of the x86 instruction set; however, our implementation (§7) works on code from real x86 programs compiled from C++ source code, including C++ STL, using Visual Studio.

Organization. §2 presents the basic principles of semantic reinterpretation by means of an example in which reinterpretation is used to create abstract transformers for abstract interpretation. §3 defines the logic that we use, as well a simple source-code language (PL) and an idealized machine-code language (MC). §4 discusses how to use reinterpretation to obtain the three symbolic-analysis primitives for PL. §5 addresses reinterpretation for MC. §6 explains how other language constructs beyond those found in PL and MC can be handled. §7 describes our implementation and the experiment carried out with it. §8 discusses related work. App. A presents correctness proofs.

2 Semantic Reinterpretation for Abstract Interpretation

This section presents the basic principles of semantic reinterpretation in the context of abstract interpretation. We use a simple language of assignments, and define the concrete semantics and an abstract sign-analysis semantics via semantic reinterpretation.

Example 1. [Adapted from [10].] Consider the following fragment of a denotational semantics, which defines the meaning of assignment statements over variables that hold signed 32-bit `int` values (where \oplus denotes exclusive-or):

$$\begin{array}{ll} I \in Id & E \in Expr ::= I \mid E_1 \oplus E_2 \mid \dots \\ S \in Stmt ::= I = E; & \sigma \in State = Id \rightarrow Int32 \end{array}$$

$$\begin{array}{ll} \mathcal{E} : Expr \rightarrow State \rightarrow Int32 & \mathcal{I} : Stmt \rightarrow State \rightarrow State \\ \mathcal{E}[[I]]\sigma = \sigma I & \mathcal{I}[[I = E;]]\sigma = \sigma[I \mapsto \mathcal{E}[[E]]\sigma] \\ \mathcal{E}[[E_1 \oplus E_2]]\sigma = \mathcal{E}[[E_1]]\sigma \oplus \mathcal{E}[[E_2]]\sigma & \end{array}$$

The specification given above can be factored into client and core specifications by introducing a domain *Val*, as well as operators *xor*, *lookup*, and *store*. The client specification is defined by

$$xor : Val \rightarrow Val \rightarrow Val \quad lookup : State \rightarrow Id \rightarrow Val \quad store : State \rightarrow Id \rightarrow Val \rightarrow State$$

$$\begin{array}{ll} \mathcal{E} : Expr \rightarrow State \rightarrow Val & \mathcal{I} : Stmt \rightarrow State \rightarrow State \\ \mathcal{E}[[I]]\sigma = lookup \ \sigma \ I & \mathcal{I}[[I = E;]]\sigma = store \ \sigma \ I \ \mathcal{E}[[E]]\sigma \\ \mathcal{E}[[E_1 \oplus E_2]]\sigma = \mathcal{E}[[E_1]]\sigma \ xor \ \mathcal{E}[[E_2]]\sigma & \end{array}$$

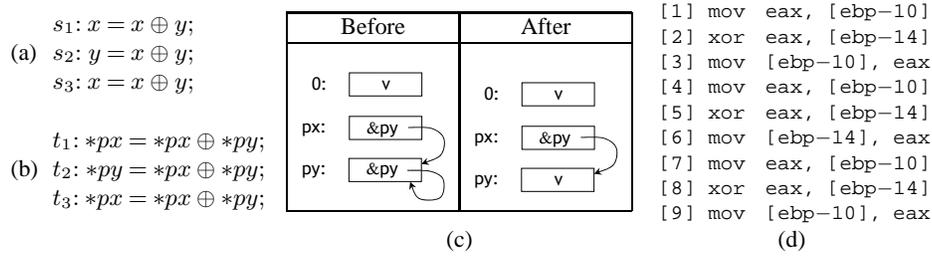


Fig. 1. (a) Code fragment that swaps two `ints`; (b) code fragment that swaps two `ints` using pointers; (c) possible before and after configurations for code fragment (b): the swap is unsuccessful due to aliasing; (d) x86 machine code corresponding to (a).

For the concrete (or “standard”) semantics, the semantic core is defined by

$$\begin{aligned}
 v \in Val_{std} &= Int32 & lookup_{std} &= \lambda\sigma.\lambda I.\sigma I & xor_{std} &= \lambda v_1.\lambda v_2.v_1 \oplus v_2 \\
 State_{std} &= Id \rightarrow Val & store_{std} &= \lambda\sigma.\lambda I.\lambda v.\sigma[I \mapsto v]
 \end{aligned}$$

Different abstract interpretations can be defined by using the same client semantics, but giving a different interpretation of the base types, function types, and operators of the core. For example, for sign analysis, assuming that `Int32` values are represented in two’s complement, the semantic core is reinterpreted as follows:³

$$\begin{aligned}
 v \in Val_{abs} &= \{neg, zero, pos\}^\top & xor_{abs} &= \lambda v_1.\lambda v_2.
 \end{aligned}$$

		v_2	
		neg	$zero$
v_1	neg	\top	neg
	$zero$	neg	$zero$
	pos	neg	pos
	\top	\top	\top

$$\begin{aligned}
 State_{abs} &= Id \rightarrow Val_{abs} \\
 lookup_{abs} &= \lambda\sigma.\lambda I.\sigma I \\
 store_{abs} &= \lambda\sigma.\lambda I.\lambda v.\sigma[I \mapsto v]
 \end{aligned}$$

For the code fragment shown in Fig. 1(a), which swaps two `ints`, sign-analysis reinterpretation creates abstract transformers that, given the initial abstract state $\sigma_0 = \{x \mapsto neg, y \mapsto pos\}$, produce the following abstract states:

$$\begin{aligned}
 \sigma_0 &:= \{x \mapsto neg, y \mapsto pos\} \\
 \sigma_1 &:= \mathcal{I}[s_1 : x = x \oplus y;] \sigma_0 = store_{abs} \sigma_0 x (neg \ xor_{abs} \ pos) = \{x \mapsto neg, y \mapsto pos\} \\
 \sigma_2 &:= \mathcal{I}[s_2 : y = x \oplus y;] \sigma_1 = store_{abs} \sigma_1 y (neg \ xor_{abs} \ pos) = \{x \mapsto neg, y \mapsto neg\} \\
 \sigma_3 &:= \mathcal{I}[s_3 : x = x \oplus y;] \sigma_2 = store_{abs} \sigma_2 x (neg \ xor_{abs} \ neg) = \{x \mapsto \top, y \mapsto neg\}.
 \end{aligned}$$

Semantic Reinterpretation Versus Standard Abstract Interpretation. Semantic reinterpretation [12, 10] is a form of abstract interpretation [4], but differs from the way abstract interpretation is normally applied: in standard abstract interpretation, one reinterprets the constructs of each *subject language*; in contrast, with semantic reinterpretation one reinterprets the constructs of the *meta-language*. Standard abstract interpretation helps in creating semantically sound *tools*; semantic reinterpretation helps in creating semantically sound *tool generators*. In particular, if you have N subject languages and M analyses, with semantic reinterpretation you obtain $N \times M$ analyzers by writing just $N + M$ specifications: concrete semantics for N subject languages and M

³ For the two’s-complement representation, $pos \ xor_{abs} \ neg = neg \ xor_{abs} \ pos = neg$ because, for all combinations of values represented by pos and neg , the high-order bit of the result is set, which means that the result is always negative. However, $pos \ xor_{abs} \ pos = neg \ xor_{abs} \ neg = \top$ because the concrete result could be either 0 or positive, and $zero \sqcup pos = \top$.

reinterpretations. With the standard approach, one must write $N \times M$ abstract semantics.

Semantic Reinterpretation Versus Translation to a Common Intermediate Representation. The mapping of a client specification to the operations of the semantic core that one defines in a semantic reinterpretation resembles a translation to a common intermediate representation (CIR) data structure. Thus, another approach to obtaining “systematic” reinterpretations that are similar to semantic reinterpretations—in that they apply to multiple subject languages—is to translate subject-language programs to a CIR, and then create various interpreters that implement different abstract interpretations of the node types of the CIR data structure. Each interpreter can be applied to (the translation of) programs in any subject language L for which one has defined an L -to-CIR translator. Compared with interpreting objects of a CIR data type, the advantages of semantic reinterpretation (i.e., reinterpreting the constructs of the *meta-language*) are

1. The presentation of our ideas is simpler because one does not have to introduce an additional language of trees for representing CIR objects.
2. With semantic reinterpretation, there is no explicit CIR data structure to be interpreted. In essence, semantic reinterpretation removes a level of interpretation, and hence generated analyzers should run faster.

To some extent, however, the decision to explain our ideas in terms of semantic reinterpretation is just a matter of presentational style. The goal of the paper is *not* to argue the merits of semantic reinterpretation *per se*; on the contrary, the goal is to present *particular interpretations that yield three desirable symbolic-analysis primitives* for use in program-analysis tools. Semantic reinterpretation is used because it allows us to present our ideas in a concise manner. The ideas introduced in §4 and §5 can be implemented using semantic reinterpretation—as we did (see §7); alternatively, they can be implemented by defining a suitable CIR datatype and creating appropriate interpretations of the CIR’s node types—again using ideas similar to those presented in §4 and §5.

3 A Logic and Two Programming Languages

3.1 L : A Quantifier-Free Bit-Vector Logic with Finite Functions

The logic L is quantifier-free first-order bit-vector logic over a vocabulary of constant symbols ($I \in Id$) and function symbols ($F \in FuncId$). Strictly speaking, we work with various instantiations of L , denoted by $L[PL]$ and $L[MC]$, in which the vocabularies of function symbols are chosen to describe aspects of the values used by, and computations performed by, the programming languages PL and MC, respectively.

We distinguish the syntactic symbols of L from their counterparts in PL (§2 and §3.2) by using boxes around L ’s symbols.

$$\begin{array}{ll}
 c \in C_{In32} = \{0, 1, \dots\} & op2_L \in BinOp_L = \{\boxed{+}, \boxed{-}, \boxed{\oplus}, \dots\} \\
 bop_L \in BoolOp_L = \{\boxed{\&\&}, \boxed{\parallel}, \dots\} & rop_L \in RelOp_L = \{\boxed{=}, \boxed{\neq}, \boxed{<}, \boxed{>}, \dots\}
 \end{array}$$

The rest of the syntax of $L[\cdot]$ is defined as follows:

$$\begin{aligned}
& \text{const} : C_{Int32} \rightarrow \text{Val} \\
& \text{cond}_L : \text{BVal} \rightarrow \text{Val} \rightarrow \text{Val} \rightarrow \text{Val} \\
& \text{lookupId} : \text{LogicalStruct} \rightarrow \text{Id} \rightarrow \text{Val} \\
& \text{binop}_L : \text{BinOp}_L \rightarrow (\text{Val} \times \text{Val} \rightarrow \text{Val}) \\
& \text{relop}_L : \text{RelOp}_L \rightarrow (\text{Val} \times \text{Val} \rightarrow \text{BVal}) \\
& \text{boolop}_L : \text{BoolOp}_L \rightarrow (\text{BVal} \times \text{BVal} \rightarrow \text{BVal}) \\
& \text{lookupFuncId} : \text{LogicalStruct} \rightarrow \text{FuncId} \rightarrow (\text{Val} \rightarrow \text{Val}) \\
& \text{access} : (\text{Val} \rightarrow \text{Val}) \times \text{Val} \rightarrow \text{Val} \\
& \text{update} : ((\text{Val} \rightarrow \text{Val}) \times \text{Val} \times \text{Val}) \rightarrow (\text{Val} \rightarrow \text{Val}) \\
\\
& \mathcal{T} : \text{Term} \rightarrow \text{LogicalStruct} \rightarrow \text{Val} \qquad \mathcal{F} : \text{Formula} \rightarrow \text{LogicalStruct} \rightarrow \text{BVal} \\
& \mathcal{T}[\![c]\!] \iota = \text{const}(c) \qquad \mathcal{F}[\![\mathbb{T}]\!] \iota = \mathbb{T} \\
& \mathcal{T}[\![I]\!] \iota = \text{lookupId } \iota I \qquad \mathcal{F}[\![\mathbb{F}]\!] \iota = \mathbb{F} \\
& \mathcal{T}[\![T_1 \text{ op}_L T_2]\!] \iota = \mathcal{T}[\![T_1]\!] \iota \text{ binop}_L(\text{op}_L) \mathcal{T}[\![T_2]\!] \iota \qquad \mathcal{F}[\![T_1 \text{ rop}_L T_2]\!] \iota = \mathcal{T}[\![T_1]\!] \iota \text{ relop}_L(\text{rop}_L) \mathcal{T}[\![T_2]\!] \iota \\
& \mathcal{T}[\![\text{ite}(\varphi, T_1, T_2)]\!] \iota = \text{cond}_L(\mathcal{F}[\![\varphi]\!] \iota, \mathcal{T}[\![T_1]\!] \iota, \mathcal{T}[\![T_2]\!] \iota) \qquad \mathcal{F}[\![\neg]\!] \varphi_1 \iota = \neg \mathcal{F}[\![\varphi_1]\!] \iota \\
& \mathcal{T}[\![FE(T)]\!] \iota = \text{access}(\mathcal{F}\mathcal{E}[\![FE]\!] \iota, \mathcal{T}[\![T]\!] \iota) \qquad \mathcal{F}[\![\varphi_1 \text{ bop}_L \varphi_2]\!] \iota = \mathcal{F}[\![\varphi_1]\!] \iota \text{ boolop}_L(\text{bop}_L) \mathcal{F}[\![\varphi_2]\!] \iota \\
\\
& \mathcal{F}\mathcal{E} : \text{FuncExpr} \rightarrow \text{LogicalStruct} \rightarrow (\text{Val} \rightarrow \text{Val}) \\
& \mathcal{F}\mathcal{E}[\![F]\!] \iota = \text{lookupFuncId } \iota F \\
& \mathcal{F}\mathcal{E}[\![FE_1[T_1 \mapsto T_2]]\!] \iota = \text{update}(\mathcal{F}\mathcal{E}[\![FE_1]\!] \iota, \mathcal{T}[\![T_1]\!] \iota, \mathcal{T}[\![T_2]\!] \iota) \\
\\
& \mathcal{U} : \text{StructUpdate} \rightarrow \text{LogicalStruct} \rightarrow \text{LogicalStruct} \\
& \mathcal{U}[\![\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\}]\!] \iota = ((\iota \uparrow 1)[I_i \mapsto \mathcal{T}[\![T_i]\!] \iota], (\iota \uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[\![FE_j]\!] \iota])
\end{aligned}$$

Fig. 2. The factored semantics of L .

$$I \in \text{Id}, T \in \text{Term}, \varphi \in \text{Formula}, F \in \text{FuncId}, FE \in \text{FuncExpr}, U \in \text{StructUpdate}$$

$$\begin{aligned}
T ::= & c \mid I \mid T_1 \text{ op}_L T_2 \mid \text{ite}(\varphi, T_1, T_2) \mid FE(T) & FE ::= & F \mid FE_1[T_1 \mapsto T_2] \\
\varphi ::= & \mathbb{T} \mid \mathbb{F} \mid T_1 \text{ rop}_L T_2 \mid \neg \varphi_1 \mid \varphi_1 \text{ bop}_L \varphi_2 & U ::= & (\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\})
\end{aligned}$$

A *FuncExpr* of the form $FE_1[T_1 \mapsto T_2]$ denotes a *function-update expression*. An expression of the form $(\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\})$ is called a *structure-update expression*. The subscripts i and j implicitly range over certain index sets, which will be omitted to reduce clutter. To emphasize that I_i and F_j refer to next-state quantities, we sometimes write structure-update expressions with primes: $(\{I'_i \leftrightarrow T_i\}, \{F'_j \leftrightarrow FE_j\})$. $\{I'_i \leftrightarrow T_i\}$ specifies the updates to the interpretations of the constant symbols and $\{F'_j \leftrightarrow FE_j\}$ specifies the updates to the interpretations of the function symbols (see below). Thus, a structure-update expression $(\{I'_i \leftrightarrow T_i\}, \{F'_j \leftrightarrow FE_j\})$ can be thought of as a kind of restricted 2-vocabulary (i.e., 2-state) formula $\bigwedge_i (I'_i = T_i) \wedge \bigwedge_j (F'_j = FE_j)$. We define U_{id} to be $(\{I' \leftrightarrow I \mid I \in \text{Id}\}, \{F' \leftrightarrow F \mid F \in \text{FuncId}\})$.

Semantics of L . The semantics of $L[\cdot]$ is defined in terms of a *logical structure*, which gives meaning to the *Id* and *FuncId* symbols of the logic's vocabulary.

$$\iota \in \text{LogicalStruct} = (\text{Id} \rightarrow \text{Val}) \times (\text{FuncId} \rightarrow (\text{Val} \rightarrow \text{Val}))$$

$(\iota \uparrow 1)$ assigns meanings to constant symbols, and $(\iota \uparrow 2)$ assigns meanings to function symbols. $((p \uparrow 1)$ and $(p \uparrow 2)$ denote the 1st and 2nd components, respectively, of a pair p .)

The factored semantics of L is presented in Fig. 2. Motivated by the needs of later sections, we retain the convention from §2 of working with the domain *Val* rather than *Int32*. Similarly, we also use *BVal* rather than *Bool*. The standard interpretations of binop_L , relop_L , and boolop_L are as one would expect, e.g., $v_1 \text{ binop}_L(\oplus) v_2 = v_1 \text{ xor } v_2$, etc. The standard interpretations for lookupId_{std} and $\text{lookupFuncId}_{std}$ select from the first and second components, respectively, of a *LogicalStruct*: $\text{lookupId}_{std} \iota I =$

$$\begin{array}{l}
v \in \text{Val} \\
l \in \text{Loc} = \text{Val} \\
\sigma \in \text{State} = \text{Store} \times \text{Env} \\
\\
\text{const} : C_{\text{Int32}} \rightarrow \text{Val} \\
\text{cond} : \text{BVal} \rightarrow \text{Val} \rightarrow \text{Val} \rightarrow \text{Val} \\
\text{lookupState} : \text{State} \rightarrow \text{Id} \rightarrow \text{Val} \\
\text{lookupEnv} : \text{State} \rightarrow \text{Id} \rightarrow \text{Loc} \\
\text{lookupStore} : \text{State} \rightarrow \text{Loc} \rightarrow \text{Val} \\
\text{updateStore} : \text{State} \rightarrow \text{Loc} \rightarrow \text{Val} \rightarrow \text{State} \\
\\
\mathcal{E} : \text{Expr} \rightarrow \text{State} \rightarrow \text{Val} \\
\mathcal{E}[\![c]\!] \sigma = \text{const}(c) \\
\mathcal{E}[\![I]\!] \sigma = \text{lookupState } \sigma \ I \\
\mathcal{E}[\![\&I]\!] \sigma = \text{lookupEnv } \sigma \ I \\
\mathcal{E}[\![*E]\!] \sigma = \text{lookupStore } \sigma \ (\mathcal{E}[\![E]\!] \sigma) \\
\mathcal{E}[\![E_1 \text{ op2 } E_2]\!] \sigma = \mathcal{E}[\![E_1]\!] \sigma \ \text{binop}(\text{op2}) \ \mathcal{E}[\![E_2]\!] \sigma \\
\mathcal{E}[\![\text{BE ? } E_1 : E_2]\!] \sigma = \text{cond}(\mathcal{B}[\![\text{BE}]\!] \sigma, \mathcal{E}[\![E_1]\!] \sigma, \mathcal{E}[\![E_2]\!] \sigma) \\
\\
\mathcal{B} : \text{BoolExpr} \rightarrow \text{State} \rightarrow \text{BVal} \\
\mathcal{B}[\![\top]\!] \sigma = \top \\
\mathcal{B}[\![\mathbb{F}]\!] \sigma = \mathbb{F} \\
\mathcal{B}[\![E_1 \text{ rop } E_2]\!] \sigma = \mathcal{E}[\![E_1]\!] \sigma \ \text{relop}(\text{rop}) \ \mathcal{E}[\![E_2]\!] \sigma \\
\mathcal{B}[\![\neg \text{BE}_1]\!] \sigma = \neg \mathcal{B}[\![\text{BE}_1]\!] \sigma \\
\mathcal{B}[\![\text{BE}_1 \text{ bop } \text{BE}_2]\!] \sigma = \mathcal{B}[\![\text{BE}_1]\!] \sigma \ \text{boolop}(\text{bop}) \ \mathcal{B}[\![\text{BE}_2]\!] \sigma \\
\\
\mathcal{I} : \text{Stmt} \rightarrow \text{State} \rightarrow \text{State} \\
\mathcal{I}[\![I = E;]\!] \sigma = \text{updateStore } \sigma \ (\text{lookupEnv } \sigma \ I) \ (\mathcal{E}[\![E]\!] \sigma) \\
\mathcal{I}[\![*I = E;]\!] \sigma = \text{updateStore } \sigma \ (\mathcal{E}[\![I]\!] \sigma) \ (\mathcal{E}[\![E]\!] \sigma) \\
\mathcal{I}[\![S_1 \ S_2;]\!] \sigma = \mathcal{I}[\![S_2]\!] (\mathcal{I}[\![S_1]\!] \sigma)
\end{array}$$

Fig. 3. The factored semantics of PL.

$(\iota \uparrow 1)(I)$ and $\text{lookupFuncId}_{\text{std}} \iota F = (\iota \uparrow 2)(F)$. The standard interpretations for *access* and *update* select from, and store to, a map, respectively.

Let $U = (\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\})$. Because $\mathcal{U}[\![U]\!] \iota$ retains from ι the value of each constant I and function F for which an update is not defined explicitly in U (i.e., $I \in (\text{Id} - \{I_i\})$ and $F \in (\text{FuncId} - \{F_j\})$), as a notational convenience we sometimes treat U as if it contains an identity update for each such symbol; that is, we say that $(U \uparrow 1)I = I$ for $I \in (\text{Id} - \{I_i\})$, and $(U \uparrow 2)F = F$ for $F \in (\text{FuncId} - \{F_j\})$.

3.2 PL : A Simple Source-Level Language

PL is the language from §2, extended with some additional kinds of `int`-valued expressions, an address-generation expression, a dereferencing expression, and an indirect-assignment statement. Note that arithmetic operations can also occur inside a dereference expression; i.e., PL allows arithmetic to be performed on addresses (including bitwise operations on addresses: see Ex. 2).

$$S \in \text{Stmt}, E \in \text{Expr}, \text{BE} \in \text{BoolExpr}, I \in \text{Id}, c \in C_{\text{Int32}}$$

$$\begin{array}{l}
c ::= 0 \mid 1 \mid \dots \\
S ::= I = E; \mid *I = E; \mid S_1 \ S_2 \quad \text{BE} ::= \top \mid \mathbb{F} \mid E_1 \ \text{rop} \ E_2 \mid \neg \text{BE}_1 \mid \text{BE}_1 \ \text{bop} \ \text{BE}_2 \\
E ::= c \mid I \mid \&I \mid *E \mid E_1 \ \text{op2} \ E_2 \mid \text{BE ? } E_1 : E_2
\end{array}$$

Semantics of PL. The factored semantics of PL is presented in Fig. 3. The semantic domain *Loc* stands for *locations* (or memory addresses). We identify *Loc* with the set *Val* of values. A state $\sigma \in \text{State}$ is a pair (η, ρ) , where, in the standard semantics, *environment* $\eta \in \text{Env} = \text{Id} \rightarrow \text{Loc}$ maps identifiers to their associated locations and *store* $\rho \in \text{Store} = \text{Loc} \rightarrow \text{Val}$ maps each location to the value that it holds.

The standard interpretations of the operators used in the PL semantics are

$$\begin{array}{l}
const : C_{Int32} \rightarrow Val \\
store_{reg} : State \rightarrow register \rightarrow Val \rightarrow State \\
lookup_{reg} : State \rightarrow register \rightarrow Val \\
store_{flag} : State \rightarrow flagName \rightarrow BVal \rightarrow State \\
lookup_{flag} : State \rightarrow flagName \rightarrow BVal \\
\\
\mathcal{R} : reg \rightarrow State \rightarrow Val \\
\mathcal{R}[\![r]\!] \sigma = lookup_{reg}(\sigma, r) \\
\\
\mathcal{K} : flagName \rightarrow State \rightarrow BVal \\
\mathcal{K}[\![ZF]\!] \sigma = lookup_{flag}(\sigma, ZF) \\
\\
cond : BVal \rightarrow Val \rightarrow Val \rightarrow Val \\
store_{mem} : State \rightarrow Val \rightarrow Val \rightarrow State \\
lookup_{mem} : State \rightarrow Val \rightarrow Val \\
store_{eip} : State \rightarrow State \\
store_{eip} = \lambda \sigma. store_{reg}(\sigma, EIP, \mathcal{R}[\![EIP]\!] \sigma binop(+) 4) \\
\\
\mathcal{O} : src_operand \rightarrow State \rightarrow Val \\
\mathcal{O}[\![Indirect(r, c)]\!] \sigma = lookup_{mem}(\sigma, \mathcal{R}[\![r]\!] \sigma binop(+) const(c)) \\
\mathcal{O}[\![DirectReg(r)]\!] \sigma = \mathcal{R}[\![r]\!] \sigma \\
\mathcal{O}[\![Immediate(c)]\!] \sigma = const(c) \\
\\
\mathcal{I} : instruction \rightarrow State \rightarrow State \\
\mathcal{I}[\![MOV(Indirect(r, c), so)]\!] \sigma = store_{eip}(store_{mem}(\sigma, \mathcal{R}[\![r]\!] \sigma binop(+) const(c), \mathcal{O}[\![so]\!] \sigma)) \\
\mathcal{I}[\![MOV(DirectReg(r), so)]\!] \sigma = store_{eip}(store_{reg}(\sigma, r, \mathcal{O}[\![so]\!] \sigma)) \\
\mathcal{I}[\![CMP(do, so)]\!] \sigma = store_{eip}(store_{flag}(\sigma, ZF, \mathcal{O}[\![do]\!] \sigma binop(-) \mathcal{O}[\![so]\!] \sigma relop(=) 0)) \\
\mathcal{I}[\![XOR(do:Indirect(r, c), so)]\!] \sigma = store_{eip}(store_{mem}(\sigma, \mathcal{R}[\![r]\!] \sigma binop(+) const(c), \mathcal{O}[\![do]\!] \sigma binop(\oplus) \mathcal{O}[\![so]\!] \sigma)) \\
\mathcal{I}[\![XOR(do:DirectReg(r), so)]\!] \sigma = store_{eip}(store_{reg}(\sigma, r, \mathcal{O}[\![do]\!] \sigma binop(\oplus) \mathcal{O}[\![so]\!] \sigma)) \\
\mathcal{I}[\![JZ(do)]\!] \sigma = store_{reg}(\sigma, EIP, cond(\mathcal{K}[\![ZF]\!] \sigma, \mathcal{R}[\![EIP]\!] \sigma binop(+) 4, \mathcal{O}[\![do]\!] \sigma))
\end{array}$$

Fig. 4. The factored semantics of MC.

$$\begin{array}{l}
BVal_{std} = BVal \\
Val_{std} = Int32 \\
Loc_{std} = Int32 \\
\eta \in Env_{std} = Id \rightarrow Loc_{std} \\
\rho \in Store_{std} = Loc_{std} \rightarrow Val_{std} \\
\\
cond_{std} = \lambda b. \lambda v_1. \lambda v_2. (b ? v_1 : v_2) \\
lookupState_{std} = \lambda(\eta, \rho). \lambda I. \rho(\eta(I)) \\
lookupEnv_{std} = \lambda(\eta, \rho). \lambda I. \eta(I) \\
lookupStore_{std} = \lambda(\eta, \rho). \lambda l. \rho(l) \\
updateStore_{std} = \lambda(\eta, \rho). \lambda l. \lambda v. (\eta, \rho[l \mapsto v])
\end{array}$$

3.3 MC: A Simple Machine-Code Language

MC is based on the x86 instruction set, but greatly simplified to have just four registers, one flag, and four instructions.

$$\begin{array}{l}
r \in register, do \in dst_operand, so \in src_operand, i \in instruction \\
\\
r ::= EAX \mid EBX \mid EBP \mid EIP \quad do ::= Indirect(r, Val) \mid DirectReg(r) \\
flagName ::= ZF \quad so ::= do \cup Immediate(Val) \\
instruction ::= MOV(do, so) \mid CMP(do, so) \mid XOR(do, so) \mid JZ(do)
\end{array}$$

Semantics of MC. The factored semantics of MC is presented in Fig. 4. It is similar to the semantics of PL, although MC exhibits two features not part of PL: there is an explicit program counter (*EIP*), and MC includes the typical feature of machine-code languages that a branch is split across two instructions (*CMP* . . . *JZ*). An MC state $\sigma \in State$ is a triple $(mem, reg, flag)$, where mem is a map $Val \rightarrow Val$, reg is a map $register \rightarrow Val$, and $flag$ is a map $flagName \rightarrow BVal$. We assume that each instruction is 4 bytes long; hence, the execution of a *MOV*, *CMP* or *XOR* increments the program-counter register *EIP* by 4. *CMP* sets the value of *ZF* according to the difference of the values of the two operands; *JZ* updates *EIP* depending on the value of flag *ZF*.

4 Symbolic Analysis for PL via Reinterpretation

A PL state (η, ρ) can be modeled in $L[PL]$ by using a function symbol F_ρ for store ρ , and a constant symbol $c_x \in Id$ for each PL identifier x . (To reduce clutter, we will use x for such constants instead of c_x .) Given $\iota \in LogicalStruct$, the constant symbols and

their interpretations in ι correspond to environment η , and the interpretation of F_ρ in ι corresponds to store ρ .

Symbolic Evaluation. A primitive for forward symbolic-evaluation must solve the following problem: *Given the semantic definition of a programming language, together with a specific statement s , create a logical formula that captures the semantics of s .* The following table illustrates how the semantics of PL statements can be expressed as $L[\text{PL}]$ structure-update expressions:

PL	$L[\text{PL}]$
$x = 17;$	$(\emptyset, \{F'_\rho \leftrightarrow F_\rho[x \mapsto 17]\})$
$x = y;$	$(\emptyset, \{F'_\rho \leftrightarrow F_\rho[x \mapsto F_\rho(y)]\})$
$x = *q;$	$(\emptyset, \{F'_\rho \leftrightarrow F_\rho[x \mapsto F_\rho(F_\rho(q))]\})$

To create such expressions automatically using semantic reinterpretation, we use formulas of logic $L[\text{PL}]$ as a reinterpretation domain for the semantic core of PL. The base types and the state type of the semantic core are reinterpreted as follows (our convention is to mark each reinterpreted base type, function type, and operator with an overbar): $\overline{Val} = \text{Term}$, $\overline{BVal} = \text{Formula}$, and $\overline{State} = \text{StructUpdate}$. The operators used in PL's meaning functions \mathcal{E} , \mathcal{B} , and \mathcal{T} are reinterpreted over these domains as follows:

- The arithmetic, bitwise, relational, and logical operators are interpreted as syntactic constructors of $L[\text{PL}]$ Terms and Formulas, e.g., $\overline{\text{binop}(\oplus)} = \lambda T_1. \lambda T_2. T_1 \boxed{\oplus} T_2$. Straightforward simplifications are also performed; e.g., $0 \boxed{\oplus} a$ simplifies to a , etc. Other simplifications that we perform are similar to ones used by others, such as the preprocessing steps used in decision procedures (e.g., the ite-lifting and read-over-write transformations for operations on functions [5]).
- $\overline{\text{cond}}$ residuates an $\overline{\text{ite}(\cdot, \cdot, \cdot)}$ Term when the result cannot be simplified to a single branch.

The other operations used in the PL semantics are reinterpreted as follows:

$$\begin{array}{ll}
\overline{\text{lookupState}} : \text{StructUpdate} \rightarrow \text{Id} \rightarrow \text{Term} & \overline{\text{lookupState}} = \lambda U. \lambda I. ((U \uparrow 2) F_\rho) ((U \uparrow 1) I) \\
\overline{\text{lookupEnv}} : \text{StructUpdate} \rightarrow \text{Id} \rightarrow \text{Term} & \overline{\text{lookupEnv}} = \lambda U. \lambda I. (U \uparrow 1) I \\
\overline{\text{lookupStore}} : \text{StructUpdate} \rightarrow \text{Term} \rightarrow \text{Term} & \overline{\text{lookupStore}} = \lambda U. \lambda T. ((U \uparrow 2) F_\rho)(T) \\
\overline{\text{updateStore}} : \text{StructUpdate} \rightarrow \text{Term} \rightarrow \text{Term} \rightarrow \text{StructUpdate} & \\
\overline{\text{updateStore}} = \lambda U. \lambda T_1. \lambda T_2. ((U \uparrow 1), (U \uparrow 2)[F_\rho \mapsto ((U \uparrow 2) F_\rho)[T_1 \mapsto T_2]]) &
\end{array}$$

By extension, this produces functions $\overline{\mathcal{E}}$, $\overline{\mathcal{B}}$, and $\overline{\mathcal{T}}$ with the following types:

Standard	Reinterpreted
$\mathcal{E}: \text{Expr} \rightarrow \text{State} \rightarrow \text{Val}$	$\overline{\mathcal{E}}: \text{Expr} \rightarrow \text{StructUpdate} \rightarrow \text{Term}$
$\mathcal{B}: \text{BoolExpr} \rightarrow \text{State} \rightarrow \text{BVal}$	$\overline{\mathcal{B}}: \text{BoolExpr} \rightarrow \text{StructUpdate} \rightarrow \text{Formula}$
$\mathcal{T}: \text{Stmt} \rightarrow \text{State} \rightarrow \text{State}$	$\overline{\mathcal{T}}: \text{Stmt} \rightarrow \text{StructUpdate} \rightarrow \text{StructUpdate}$

Function $\overline{\mathcal{T}}$ translates a statement s of PL to a phrase in logic $L[\text{PL}]$.

Example 2. The steps of symbolic execution of Fig. 1(b) via semantic reinterpretation, starting with an *StructUpdate* that corresponds to Fig. 1(c) are shown in Fig. 5. The *StructUpdate* U_4 can be considered to be the 2-vocabulary formula $F'_\rho = F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto v]$, which expresses a state change that does not usually perform a successful swap.

$$\begin{aligned}
U_1 &= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto \text{py}]) \\
\bar{\mathcal{I}}[*px = *px \boxplus *py;]U_1 &= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto (\bar{\mathcal{E}}[*px]U_1 \boxplus \bar{\mathcal{E}}[*py]U_1)]) \\
&= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto (\text{py} \boxplus \text{py})]) \\
&= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto 0]) = U_2 \\
\bar{\mathcal{I}}[*py = *px \boxplus *py;]U_2 &= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto (\bar{\mathcal{E}}[*px]U_2 \boxplus \bar{\mathcal{E}}[*py]U_2)][\text{px} \mapsto \text{py}][\text{py} \mapsto 0]) \\
&= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto (0 \boxplus v)][\text{px} \mapsto \text{py}][\text{py} \mapsto 0]) \\
&= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto 0]) = U_3 \\
\bar{\mathcal{I}}[*px = *px \boxplus *py;]U_3 &= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto (\bar{\mathcal{E}}[*px]U_3 \boxplus \bar{\mathcal{E}}[*py]U_3)]) \\
&= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto (0 \boxplus v)]) \\
&= (\emptyset, F_\rho \leftrightarrow F_\rho[0 \mapsto v][\text{px} \mapsto \text{py}][\text{py} \mapsto v]) = U_4
\end{aligned}$$

Fig. 5. Symbolic execution of Fig. 1(b) via semantic reinterpretation, starting with an *StructUpdate* that corresponds to the “Before” column of Fig. 1(c).

WLP. $\mathcal{WLP}(s, \varphi)$ characterizes the set of states σ such that the execution of s starting in σ either fails to terminate or results in a state σ' such that $\varphi(\sigma')$ holds. For a language that only has `int`-valued variables, the \mathcal{WLP} of a postcondition (specified by formula φ) with respect to an assignment statement $var = rhs$; can be expressed as the formula obtained by substituting rhs for all (free) occurrences of var in φ : $\varphi[var \leftarrow rhs]$.

For a language with pointer variables, such as PL, syntactic substitution is not adequate for finding \mathcal{WLP} formulas. For instance, suppose that we are interested in finding a formula for the \mathcal{WLP} of postcondition $x = 5$ with respect to $*p = e$. It is not correct merely to perform the substitution $(x = 5)[*p \leftarrow e]$. That substitution yields $x = 5$, whereas the \mathcal{WLP} depends on the execution context in which $*p = e$; is evaluated:

- If p points to x , then the \mathcal{WLP} formula should be $e = 5$.
- If p does not point to x , then the \mathcal{WLP} formula should be $x = 5$.

The desired formula can be expressed informally as $((p = \&x) ? e : x) = 5$.

For a program fragment that involves multiple pointer variables, the \mathcal{WLP} formula may have to take into account all possible aliasing combinations. This is the essence of Morris’s rule of substitution [11]. One of the most important features of our approach is its ability to create correct implementations of Morris’s rule of substitution automatically—and basically for free.

Example 3. In $L[\text{PL}]$, such a formula would be expressed as shown below on the right. (This formula will be created using semantic reinterpretation in Ex. 4.)

	Informal	$L[\text{PL}]$
Query	$\mathcal{WLP}(*p = e, x = 5)$	$\mathcal{WLP}(*p = e, F_\rho(x) \boxplus 5)$
Result	$((p = \&x) ? e : x) = 5$	$ite(F_\rho(p) \boxplus x, F_\rho(e), F_\rho(x)) \boxplus 5$

To create primitives for \mathcal{WLP} and symbolic composition via semantic reinterpretation, we again use $L[\cdot]$ as a reinterpretation domain; however, there is a trick: in contrast with what is done to generate symbolic-evaluation primitives, we use the *StructUpdate* type of $L[\cdot]$ to reinterpret the meaning functions \mathcal{U} , \mathcal{FE} , \mathcal{F} , and \mathcal{T} of $L[\cdot]$ itself! By this means, the “alternative meaning” of a *Term/Formula/FuncExpr/StructUpdate* is a (usually different) *Term/Formula/FuncExpr/StructUpdate* in which some substitution and/or simplification has taken place. The general scheme is outlined in the following table:

Meaning function(s)	Type reinterpreted	Replacement type	Function created
$\overline{\mathcal{I}}, \overline{\mathcal{E}}, \overline{\mathcal{B}}$	<i>State</i>	<i>StructUpdate</i>	Symbolic evaluation
$\overline{\mathcal{F}}, \overline{\mathcal{T}}$	<i>LogicalStruct</i>	<i>StructUpdate</i>	\mathcal{WLP}
$\overline{\mathcal{U}}, \overline{\mathcal{FE}}, \overline{\mathcal{F}}, \overline{\mathcal{T}}$	<i>LogicalStruct</i>	<i>StructUpdate</i>	Symbolic composition

In §3.1, we defined the semantics of $L[\cdot]$ in a form that would make it amenable to semantic reinterpretation. However, one small point needs adjustment: in §3.1, the type signatures of *LogicalStruct*, *lookupFuncId*, *access*, *update*, and \mathcal{FE} include occurrences of $Val \rightarrow Val$. This was done to make the types more intuitive; however, for reinterpretation to work, an additional level of factoring is necessary. In particular, the occurrences of $Val \rightarrow Val$ need to be replaced by *FVal*. The standard semantics of *FVal* is $Val \rightarrow Val$; however, for creating symbolic-analysis primitives, *FVal* is reinterpreted as *FuncExpr*.

The reinterpretation used for $\overline{\mathcal{U}}$, $\overline{\mathcal{FE}}$, $\overline{\mathcal{F}}$, and $\overline{\mathcal{T}}$ is similar to what was used for symbolic execution of PL programs:

- $\overline{Val} = Term$, $\overline{BVal} = Formula$, $\overline{FVal} = FuncExpr$, and $\overline{LogicalStruct} = StructUpdate$.
- The arithmetic, bitwise, relational, and logical operators are interpreted as syntactic *Term* and *Formula* constructors of L (e.g., $\overline{binop_L(\oplus)} = \lambda T_1. \lambda T_2. T_1 \oplus T_2$) although straightforward simplifications are also performed.
- $\overline{cond_L}$ residuates an *ite*(\cdot, \cdot, \cdot) *Term* when the result cannot be simplified to a single branch.
- $\overline{lookupId}$ and $\overline{lookupFuncId}$ are resolved immediately, rather than residuated:
 - $\overline{lookupId}(\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\}) I_k = T_k$
 - $\overline{lookupFuncId}(\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\}) F_k = FE_k$.
- \overline{access} and \overline{update} are discussed below.

By extension, this produces reinterpreted meaning functions $\overline{\mathcal{U}}$, $\overline{\mathcal{FE}}$, $\overline{\mathcal{F}}$, and $\overline{\mathcal{T}}$.

Somewhat surprisingly, we do not need to introduce an explicit operation of substitution for our logic because *a substitution operation is produced as a by-product of reinterpretation*. In particular, in the standard semantics for L , the return types of meaning function \mathcal{T} and helper function *lookupId* of the semantic core are both *Val*. However, in the reinterpreted semantics, a \overline{Val} is a *Term*—i.e., something *symbolic*—which is used in subsequent computations. Thus, when $\iota \in LogicalStruct$ is reinterpreted as $U \in StructUpdate$, the reinterpretation of formula φ via $\overline{\mathcal{F}}[\varphi]U$ substitutes *Terms* found in U into φ : $\overline{\mathcal{F}}[\varphi]U$ calls $\overline{\mathcal{T}}[T]U$, which may call $\overline{lookupId} U I$; the latter would return a *Term* fetched from U , which would be a subterm of the answer returned by $\overline{\mathcal{T}}[T]U$, which in turn would be a subterm of the answer returned by $\overline{\mathcal{F}}[\varphi]U$.

To create a formula for \mathcal{WLP} via semantic reinterpretation, we make use of both $\overline{\mathcal{F}}$, the reinterpreted logic semantics, and $\overline{\mathcal{T}}$, the reinterpreted programming-language semantics. The \mathcal{WLP} formula for φ with respect to statement s is obtained by performing the following computation:

$$\mathcal{WLP}(s, \varphi) = \overline{\mathcal{F}}[\varphi](\overline{\mathcal{T}}[s]U_{id}).$$

For a proof of correctness of this method, see App. A.2.

To understand how pointers are handled during the \mathcal{WLP} operation, the key reinterpretations to concentrate on are the ones for the operations of the semantic core of $L[PL]$ that manipulate *FVals* (i.e., arguments of type $Val \rightarrow Val$)—in particular, *access* and *update*. We want \overline{access} and \overline{update} to enjoy the following semantic properties:

$$\begin{aligned}\mathcal{T}[\overline{\text{access}}(FE_0, T_0)]\iota &= (\mathcal{F}\mathcal{E}[\![FE_0]\!]\iota)(\mathcal{T}[\![T_0]\!]\iota) \\ \mathcal{T}[\overline{\text{update}}(FE_0, T_0, T_1)]\iota &= (\mathcal{F}\mathcal{E}[\![FE_0]\!]\iota)[\mathcal{T}[\![T_0]\!]\iota \mapsto \mathcal{T}[\![T_1]\!]\iota]\end{aligned}$$

Note that these properties require evaluating the results of $\overline{\text{access}}$ and $\overline{\text{update}}$ with respect to an arbitrary $\iota \in \text{LogicalStruct}$. As mentioned earlier, it is desirable for reinterpreted base-type operations to perform simplifications whenever possible, when they construct *Terms*, *Formulas*, *FuncExprs*, and *StructUpdates*. However, because the value of ι is unknown, $\overline{\text{access}}$ and $\overline{\text{update}}$ operate in an uncertain environment.

To use semantic reinterpretation to create a \mathcal{WLP} primitive that implements Morris’s rule, simplifications are performed by $\overline{\text{access}}$ and $\overline{\text{update}}$ according to the definitions given below, where \equiv , \neq , and \doteq denote *equality-as-terms*, *definite-disequality*, and *possible-equality*, respectively.

$$\begin{aligned}\overline{\text{access}}(F, k_1) &= F(k_1) \\ \overline{\text{access}}(FE[k_2 \mapsto d_2], k_1) &= \begin{cases} d_2 & \text{if } (k_1 \equiv k_2) \\ \overline{\text{access}}(FE, k_1) & \text{if } (k_1 \neq k_2) \\ \text{ite}(k_1 \boxed{\equiv} k_2, d_2, \overline{\text{access}}(FE, k_1)) & \text{if } (k_1 \doteq k_2) \end{cases} \\ \overline{\text{update}}(F, k_1, d_1) &= F[k_1 \mapsto d_1] \\ \overline{\text{update}}(FE[k_2 \mapsto d_2], k_1, d_1) &= \begin{cases} FE[k_1 \mapsto d_1] & \text{if } (k_1 \equiv k_2) \\ \overline{\text{update}}(FE, k_1, d_1)[k_2 \mapsto d_2] & \text{if } (k_1 \neq k_2) \\ FE[k_2 \mapsto d_2][k_1 \mapsto d_1] & \text{if } (k_1 \doteq k_2) \end{cases}\end{aligned}$$

(The possible-equality tests, “ $k_1 \doteq k_2$ ”, are really “otherwise” cases of three-pronged comparisons.) The possible-equality case for $\overline{\text{access}}$ introduces *ite* terms. As illustrated in Ex. 4, it is these *ite* terms that cause the reinterpreted operations to account for possible aliasing combinations, and thus are the reason that the semantic-reinterpretation method automatically carries out the actions of Morris’s rule of substitution [11].

Example 4. We now demonstrate how semantic reinterpretation produces the $L[\text{PL}]$ formula for $\mathcal{WLP}(*p = e, x = 5)$ claimed in Ex. 3.

$$\begin{aligned}U &:= \overline{\mathcal{T}}[\![*p = e]\!]U_{Id} \\ &= \overline{\text{updateStore}}(U_{Id}, \overline{\mathcal{E}}[\![p]\!]U_{Id}, \overline{\mathcal{E}}[\![e]\!]U_{Id}) \\ &= \overline{\text{updateStore}}(U_{Id}, \text{lookupState}(U_{Id}, p), \overline{\text{lookupState}}(U_{Id}, e)) \\ &= \overline{\text{updateStore}}(U_{Id}, F_\rho(p), F_\rho(e)) \\ &= ((U_{Id} \uparrow 1), F_\rho \mapsto F_\rho[F_\rho(p) \mapsto F_\rho(e)]) \\ \mathcal{WLP}(*p = e, F_\rho(x) \boxed{\equiv} 5) &= \overline{\mathcal{F}}[\![F_\rho(x) \boxed{\equiv} 5]\!]U \\ &= \overline{\mathcal{T}}[\![F_\rho(x)]\!]U \boxed{\equiv} \overline{\mathcal{T}}[\![5]\!]U \\ &= \overline{\text{access}}(\overline{\mathcal{F}\mathcal{E}}[\![F_\rho]\!]U, \overline{\mathcal{T}}[\![x]\!]U) \boxed{\equiv} 5 \\ &= \overline{\text{access}}(\text{lookupFuncId}(U, F_\rho), \text{lookupId}(U, x)) \boxed{\equiv} 5 \\ &= \overline{\text{access}}(F_\rho[F_\rho(p) \mapsto F_\rho(e)], x) \boxed{\equiv} 5 \\ &= \text{ite}(F_\rho(p) \boxed{\equiv} x, F_\rho(e), \overline{\text{access}}(F_\rho, x)) \boxed{\equiv} 5 \\ &= \text{ite}(F_\rho(p) \boxed{\equiv} x, F_\rho(e), F_\rho(x)) \boxed{\equiv} 5\end{aligned}$$

Note how the case for $\overline{\text{access}}$ that involves a possible-equality comparison causes an *ite* term to arise that tests “ $F_\rho(p) \boxed{\equiv} x$ ”. The test determines whether the value of p is the address of x , which is the only aliasing condition that matters for this example.

Symbolic Composition. The goal of symbolic composition is to have a method that, given two symbolic representations of state changes, computes a symbolic representation of their composed state change. In our approach, each state change is represented in

logic $L[\cdot]$ by an *StructUpdate*, and the method computes a new *StructUpdate* that represents their composition. To accomplish this, $L[\cdot]$ is used as a reinterpretation domain, exactly as for \mathcal{WLP} . Moreover, \overline{U} turns out to be exactly the symbolic-composition function that we seek. In particular, \overline{U} works as follows:

$$\overline{U}[\{\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\}\}]U = ((U\uparrow 1)[I_i \mapsto \overline{T}[T_i]U], (U\uparrow 2)[F_j \mapsto \overline{\mathcal{F}\mathcal{E}}[FE_j]U])$$

Example 5. For the swap-code fragment from Fig. 1(a), we can demonstrate the ability of \overline{U} to perform symbolic composition by showing that

$$\overline{T}[\{s_1; s_2; s_3\}]U_{id} = \overline{U}[\overline{T}[\{s_3\}]U_{id}](\overline{T}[\{s_1; s_2\}]U_{id}).$$

First, consider the left-hand side. It is not hard to show that $\overline{T}[\{s_1; s_2; s_3\}]U_{id} = (\{x' \leftrightarrow y, y' \leftrightarrow x\}, \emptyset)$. Now consider the right-hand side. Let $U_{1,2}$ and U_3 be

$$\begin{aligned} U_{1,2} &= \overline{T}[\{s_1; s_2\}]U_{id} = (\{x' \leftrightarrow x \oplus y, y' \leftrightarrow x\}, \emptyset) \\ U_3 &= \overline{T}[\{s_3\}]U_{id} = (\{x' \leftrightarrow x \oplus y, y' \leftrightarrow y\}, \emptyset). \end{aligned}$$

We want to compute

$$\begin{aligned} \overline{U}[\{U_3\}]U_{1,2} &= \overline{U}[(\{x' \leftrightarrow x \oplus y, y' \leftrightarrow y\}, \emptyset)]U_{1,2} \\ &= ((U_{1,2}\uparrow 1)[x \mapsto \mathcal{T}[x \oplus y]U_{1,2}, y \mapsto \mathcal{T}[y]U_{1,2}], \emptyset) \\ &= ((U_{1,2}\uparrow 1)[x \mapsto (x \oplus y) \oplus x, y \mapsto x], \emptyset) \\ &= ((U_{1,2}\uparrow 1)[x \mapsto y, y \mapsto x], \emptyset) \\ &= (\{x' \leftrightarrow y, y' \leftrightarrow x\}, \emptyset) \end{aligned}$$

Therefore, $\overline{T}[\{s_1; s_2; s_3\}]U_{id} = \overline{U}[\{U_3\}]U_{1,2}$.

5 Symbolic Analysis for MC via Reinterpretation

To obtain the three symbolic-analysis primitives for MC, we use a reinterpretation of MC's semantics that is essentially identical to the reinterpretation for PL, modulo the fact that the semantics of PL is written in terms of the combinators *lookupEnv*, *lookupStore*, and *updateStore*, whereas the semantics of MC is written in terms of *lookup_{reg}*, *store_{reg}*, *lookup_{flag}*, *store_{flag}*, *lookup_{mem}*, and *store_{mem}*.

The base types are redefined as $\overline{BVal} = \text{Formula}$, $\overline{Val} = \text{Term}$, $\overline{State} = \text{StructUpdate}$, where the vocabulary for *LogicalStructs* is $(\{ZF, EAX, EBX, EBP, EIP\}, \{F_{mem}\})$. Lookup and store operations for MC, such as $\overline{lookup_{mem}}$ and $\overline{store_{mem}}$, are handled the same way that $\overline{lookupStore}$ and $\overline{updateStore}$ are handled for PL.

Example 6. Fig. 1(d) shows the MC code that corresponds to the swap code in Fig. 1(a): lines 1–3, lines 4–6, and lines 7–9 correspond to lines 1, 2, and 3 of Fig. 1(a), respectively. For the MC code in Fig. 1(d), $\overline{T}_{MC}[\text{swap}]U_{id}$ produces the *StructUpdate*

$$\left(\{EAX' \leftrightarrow F_{mem}(EBP \boxed{-} 14)\}, \left\{ F'_{mem} \leftrightarrow F_{mem}[EBP \boxed{-} 10] \mapsto F_{mem}(EBP \boxed{-} 14)[EBP \boxed{-} 14] \mapsto F_{mem}(EBP \boxed{-} 10) \right\} \right)$$

Fig. 1(d) illustrates why it is essential to be able to handle address arithmetic: an access on a source-level variable is compiled into machine code that dereferences an address in the stack frame computed from the frame pointer (*EBP*) and an offset. This example shows that \overline{T}_{MC} is able to handle address arithmetic correctly.

<pre> [1] void foo(int e, int x, int* p) { [2] ... [3] *p = e; [4] if(x == 5) [5] goto ERROR; [6] }</pre>	<pre> [1] mov eax, p; [2] mov ebx, e; [3] mov [eax], ebx; [4] cmp x, 5; [5] jz ERROR; [6] ... [7] ERROR: ...</pre>
(a)	(b)

Fig. 6. (a) A simple source-code fragment written in PL; (b) the MC code for (a).

To create a formula for the \mathcal{WLP} of φ with respect to instruction i via semantic reinterpretation, we use the reinterpreted MC semantics $\bar{\mathcal{T}}_{MC}$, together with the reinterpreted $L[MC]$ meaning function $\bar{\mathcal{F}}_{MC}$, where $\bar{\mathcal{F}}_{MC}$ is created via the same approach used in §4 to reinterpret $L[PL]$. $\mathcal{WLP}(i, \varphi)$ is obtained by performing $\bar{\mathcal{F}}_{MC}[\![\varphi]\!] (\bar{\mathcal{T}}_{MC}[\![i]\!] U_{id})$.

Example 7. Fig. 6(a) shows a source-code fragment; Fig. 6(b) shows the corresponding MC code. (To simplify the MC code, source-level variable names are used.) In Fig. 6(a), the largest set of states just before line [3] that cause the branch to ERROR to be taken at line [4] is described by $\mathcal{WLP}(*p = e, x = 5)$. In Fig. 6(b), an expression that characterizes whether the branch to ERROR is taken is $\mathcal{WLP}(s_{[1]-[5]}, (EIP \sqsubseteq c_{[7]}))$, where $s_{[1]-[5]}$ denotes instructions [1]–[5] of Fig. 6(b), and $c_{[7]}$ is the address of ERROR. Using semantic reinterpretation, $\bar{\mathcal{F}}_{MC}[\![(EIP \sqsubseteq c_{[7]})]\!] (\bar{\mathcal{T}}[\![s_{[1]-[5]}\!] U_{id}])$ produces the formula $(ite((F_{mem}(p) \sqsubseteq x), F_{mem}(e), F_{mem}(x)) \sqsubseteq 5) \sqsubseteq 0)$, which, transliterated to informal source-level notation, is $((p = \&x) ? e : x) - 5 = 0$.

Even though the branch is split across two instructions, \mathcal{WLP} can be used to recover the branch condition. $\mathcal{WLP}(\text{cmp } x, 5; \text{ jz ERROR}, (EIP \sqsubseteq c_{[7]}))$ returns the formula $ite(((F_{mem}(x) \sqsubseteq 5) \sqsubseteq 0), c_{[7]}, c_{[6]}) \sqsubseteq c_{[7]}$ as follows:

$$\begin{aligned}
\bar{\mathcal{T}}[\![\text{cmp } x, 5]\!] U_{id} &= (\{ZF' \leftrightarrow (F_{mem}(x) \sqsubseteq 5) \sqsubseteq 0\}, \emptyset) = U_1 \\
\bar{\mathcal{T}}[\![\text{jz ERROR}]\!] U_1 &= (\{EIP' \leftrightarrow ite(((F_{mem}(x) \sqsubseteq 5) \sqsubseteq 0), c_{[7]}, c_{[6]})\}, \emptyset) = U_2 \\
\bar{\mathcal{F}}_{MC}[\![EIP \sqsubseteq c_{[7]}\!] U_2 &= ite(((F_{mem}(x) \sqsubseteq 5) \sqsubseteq 0), c_{[7]}, c_{[6]}) \sqsubseteq c_{[7]}
\end{aligned}$$

Because $c_{[7]} \neq c_{[6]}$, this simplifies to $(F_{mem}(x) \sqsubseteq 5) \sqsubseteq 0$ —i.e., in source-level terms, $(x - 5) = 0$.

6 Other Language Constructs

Branching. Ex. 7 illustrated a \mathcal{WLP} computation across a branch. We now illustrate forward symbolic execution across a branch.

Suppose that an if-statement is represented by $IfStmt(BE, Int32, Int32)$, where BE is the condition and the two $Int32$ s are the addresses of the true-branch and false-branch, respectively. Its factored semantics would specify how the value of the program counter PC changes:

$$\mathcal{I}[\![IfStmt(BE, c_T, c_F)]\!] \sigma = \text{updateStore } \sigma \text{ PC cond}(\mathcal{B}[\![BE]\!] \sigma, \text{const}(c_T), \text{const}(c_F)).$$

In the reinterpretation for symbolic evaluation, the $StructUpdate$ U obtained by $\mathcal{I}[\![IfStmt(BE, c_T, c_F)]\!] U_{id}$ would be $(\{PC' \leftrightarrow ite(\varphi_{BE}, c_T, c_F)\}, \emptyset)$, where φ_{BE} is the *Formula* obtained for BE under the reinterpreted semantics. To obtain the branch condition for a specific branch, say the true-branch, we evaluate $\bar{\mathcal{F}}[\![PC \sqsubseteq c_T]\!] U$. The result

```

Formula ObtainPathConstraintFormula(Path  $\pi$ ) {
  Formula  $\varphi = \mathbb{T}$ ; // Initial path-constraint formula
  StructUpdate  $U = U_{id}$ ; // Initial symbolic state-transformer
  let [ $PC_1 : i_1, PC_2 : i_2, \dots, PC_n : i_n, PC_{n+1} : \text{skip}$ ] =  $\pi$  in
  for ( $k = 1; k \leq n; k++$ ) {
     $U = \mathcal{I}[[i_k]]U$ ; // Symbolically execute  $i_k$ 
    if ( $i_k$  is a branch instruction)
       $\varphi = \varphi \ \&\& \ \mathcal{F}[[PC = PC_{k+1}]]U$ ; // Conjoin the branch condition for  $i_k$ 
  }
  return  $\varphi$ ;
}

```

Fig. 7. An algorithm to obtain a path-constraint formula that characterizes which initial states must follow path π .

is ($ite(\varphi_{BE}, c_T, c_F) \stackrel{\square}{=} c_T$), which (assuming that $c_T \neq c_F$) simplifies to φ_{BE} . (A similar formula simplification was performed in Ex. 7 on the result of the \mathcal{WLP} formula.)

Loops. One kind of intended client of our approach to creating symbolic-analysis primitives is hybrid concrete/symbolic state-space exploration [6, 13, 7, 3]. Such tools use a combination of concrete and symbolic execution to generate inputs that increase coverage. In such tools, a program-level loop is executed concretely a specific number of times as some path π is followed. The symbolic-execution primitive for a single instruction is applied to each instruction of π to obtain symbolic states at each point of π . A *path-constraint formula* that characterizes which initial states must follow π can be obtained by collecting the branch formula φ_{BE} obtained at each branch condition by the technique described above; the algorithm is shown in Fig. 7.

X86 String Instructions. X86 string instructions can involve actions that perform an *a priori* unbounded amount of work (e.g., the amount performed is determined by the value held in register *ECX* at the start of the instruction). This can be reduced to the loop case discussed above by giving a semantics in which the instruction itself is one of its two successors. In essence, the “microcode loop” is converted into an explicit loop.

Procedures. A call statement’s semantics (i.e., how the state is changed by the call action) would be specified with some collection of operations. Again, the reinterpretation of the state transformer is induced by the reinterpretation of each operation:

- For a call statement in a high-level language, there would be an operation that creates a new activation record. The reinterpretation of this would generate a fresh logical constant to represent the location of the new activation record.
- For a call instruction in a machine-code language, register operations would change the stack pointer and frame pointer, and memory operations would initialize fields of the new activation record. These are reinterpreted in exactly the same way that register and memory operations are reinterpreted for other constructs.

Dynamic Allocation. Two approaches are possible:

- The allocation package is implemented as a library. One can apply our techniques to the machine code from the library.
- If a formula is desired that is based on a high-level semantics, a call statement that calls `malloc` or `new` can be reinterpreted using the kind of approach used in other systems (a fresh logical constant denoting a new location can be generated).

7 Implementation and Evaluation

Implementation. Our implementation uses the TSL system [9]. (TSL stands for “Transformer Specification Language”.) The TSL language is a strongly typed, first-order functional language with a datatype-definition mechanism for defining recursive datatypes, plus deconstruction by means of pattern matching. Writing a TSL specification for an instruction set is similar to writing an interpreter in first-order ML. For instance, the meaning function \mathcal{I} of §3.3 is written as a TSL function

```
state interpInstr(instruction I, state S) {...};
```

where `instruction` and `state` are user-defined datatypes that represent the syntactic objects (in this case, instructions) and the semantic states, respectively.

We used TSL to (1) define the syntax of $L[\cdot]$ as a user-defined datatype; (2) create a reinterpretation based on $L[\cdot]$ formulas; (3) define the semantics of $L[\cdot]$ by writing functions that correspond to \mathcal{T} , \mathcal{F} , etc.; and (4) apply reinterpretation (2) to the meaning functions of $L[\cdot]$ itself. (We already had TSL specifications of x86 and PowerPC.)

TSL’s meta-language provides a fixed set of base-types; a fixed set of arithmetic, bitwise, relational, and logical operators; and a facility for defining map-types. Each TSL reinterpretation is defined over the *meta-language constructs*, by reinterpreting the TSL base-types, base-type operators, map-types, and map-type operators (i.e., *access* and *update*). When semantic reinterpretation is performed in this way, it is *independent* of any given subject language. Consequently, now that we have carried out steps (1)–(4), all three symbolic-analysis primitives can be generated automatically for a new instruction set IS merely by writing a TSL specification of IS , and then applying the TSL compiler. In essence, TSL act as a “Yacc-like” tool for generating symbolic-analysis primitives from a semantic description of an instruction set.

To illustrate the leverage gained by using the approach presented in this paper, the following table lists the number of (non-blank) lines of C++ that are generated from the TSL specifications of the x86 and PowerPC instruction sets. The number of (non-blank) lines of TSL are indicated in bold.

	TSL		C++	
	$\mathcal{I}[\cdot]$	$\mathcal{F}[\cdot] \cup \mathcal{T}[\cdot] \cup \mathcal{FE}[\cdot] \cup \mathcal{U}[\cdot]$	$\overline{\mathcal{I}}[\cdot]$	$\overline{\mathcal{F}}[\cdot] \cup \overline{\mathcal{T}}[\cdot] \cup \overline{\mathcal{FE}}[\cdot] \cup \overline{\mathcal{U}}[\cdot]$
x86	3,524	1,510	23,109	15,632
PowerPC	1,546	(already written)	12,153	15,632

The C++ code is emitted as a template, which can be instantiated with different interpretations. For instance, instantiations that create C++ implementations of $\mathcal{I}_{x86}[\cdot]$ and $\mathcal{I}_{PowerPC}[\cdot]$ (i.e., emulators for x86 and PowerPC, respectively) can be obtained trivially. Thus, for a hybrid concrete/symbolic tool for x86, our tool essentially furnishes 23,109 lines of C++ for the concrete-execution component and 23,109 lines of C++ for the symbolic-execution component. Note that the 1,510 lines of TSL that defines $\mathcal{F}[\cdot]$, $\mathcal{T}[\cdot]$, $\mathcal{FE}[\cdot]$, and $\mathcal{U}[\cdot]$ needs to be written only once.

In addition to the components for concrete and symbolic execution, one also obtains an implementation of \mathcal{WLP} —via the method described in §4—by calling the C++ implementations of $\overline{\mathcal{F}}[\cdot]$ and $\overline{\mathcal{I}}[\cdot]$: $\mathcal{WLP}(s, \varphi) = \overline{\mathcal{F}}[\varphi](\overline{\mathcal{I}}[s]U_{id})$. By Thm. A2, \mathcal{WLP} is guaranteed to be consistent with the components for concrete and symbolic execution (modulo bugs in the implementation of TSL).

Evaluation. Some tools that use symbolic reasoning employ formula transformations that are not faithful to the actual semantics. For instance, SAGE [7] uses an approximate x86 symbolic execution in which concrete values are used when non-linear operators or symbolic pointer dereferences are encountered. As a result, its symbolic execution of a path can produce an “unfaithful” path-constraint formula φ ; that is, φ can be unsatisfiable when the path is executable, or satisfiable when the path is not executable. Both situations are called a *divergence* [7]. Because the intended use of SAGE is to generate inputs that increase coverage, it can be acceptable for the tool to have a substantial divergence rate (due to the use of unfaithful symbolic techniques) if the cost of performing symbolic operations is lowered in most circumstances.

However, if we eventually hope to model check x86 machine code, implementations of faithful symbolic techniques will be required. Using faithful symbolic techniques could raise the cost of performing symbolic operations because faithful path-constraint formulas could end up being a great deal more complex than unfaithful ones. Thus, our experiment was designed to answer the question “What is the cost of using exact symbolic-execution primitives instead of unfaithful ones?”

It would have been an error-prone task to implement a faithful symbolic-execution primitive for x86 machine code manually. Using TSL, however, we were able to generate a faithful symbolic-execution primitive from an existing, well-tested TSL specification of the semantics of x86 instructions. We also generated an unfaithful symbolic-execution primitive that adopts SAGE’s approximate approach. We used these to create two symbolic-execution tools that perform state-space exploration—one that uses the faithful primitive, and one that uses the unfaithful primitive.

Although the presentation in earlier sections was couched in terms of simplified core languages, the implemented tools work with real x86 programs. Our experiments used six C++ programs, each exercising a single algorithm from the C++ STL, compiled under Visual Studio 2005. We compared the two tools’ divergence rates and running times (see Tab. 1). On average, the approximate version had 5.2X fewer constraints in

Name (STL)	# Tests	Trace #instrs	# branch	Faithful					Approximate				
				CE	SE	SMT	$ \varphi $	Div.	C+SE	SMT	$ \varphi $	Div.	Dist.
search	18	770	28	0.26	8.68	0.26	10.5	0%	9.13	0.10	4.8	61%	55%
random_shuffle	48	1831	51	0.59	21.6	0.17	27.3	0%	21.9	0.03	1.0	95%	93%
copy	5	1987	57	0.69	55.0	0.15	5.4	0%	55.8	0.03	1.0	60%	57%
partition	13	2155	76	0.72	26.4	0.43	35.2	0%	27.4	0.02	1.0	92%	58%
max_element	101	2870	224	0.94	17.0	3.59	153.0	0%	18.0	2.90	78.4	83%	6%
transform	11	10880	476	4.22	720.8	1.12	220.6	0%	713.6	0.03	1.0	82%	89%

Table 1. Experimental results. We report the number of tests executed, the average length of the trace obtained from the tests, and the average number of branches in the traces. For the faithful version, we report the average time taken for concrete (CE) and symbolic executions (SE). In the approximate version, these were done in lock step and their total time is reported in (C+SE). (All times are in seconds.) For each version, we also report the average time taken by the SMT solver (Yices), the average number of constraints found ($|\varphi|$), and the divergence rate. For the approximate version, we also show the average distance (in % of the total length of the trace) before a diverging test diverged.

φ , had a 79% divergence rate, and was about 2X faster than the faithful version; the faithful version reported no divergences.

8 Related Work

Symbolic analysis is used in many recent systems for testing and verification:

- Hybrid concrete/symbolic tools [6, 13, 7, 3] use a combination of concrete and symbolic execution to generate inputs that increase coverage.
- \mathcal{WLP} can be used to create new predicates that split part of a program’s abstract state space [1, 2].
- Symbolic composition is useful when a tool has access to a formula that summarizes a called procedure’s behavior [14]; re-exploration of the procedure is avoided by symbolically composing a path formula with the procedure-summary formula.

However, compared with the way such symbolic-analysis primitives are implemented in existing program-analysis tools, our work has one key advantage: it creates the core concrete-execution and symbolic-analysis components in a way that ensures by construction that they are *mutually consistent*. We are not aware of existing tools in which the concrete-execution and symbolic-analysis primitives are implemented in a way that guarantees such a consistency property. For instance, in the source code for B2 [8] (the next-generation Blast), one finds symbolic execution (*post*) and \mathcal{WLP} implemented with different pieces of code, and hence mutual consistency is not guaranteed. \mathcal{WLP} is implemented via substitution, with special-case code for handling pointers.

References

1. T. Ball, R. Majumdar, T. Millstein, and S. Rajamani. Automatic predicate abstraction of C programs. In *PLDI*, 2001.
2. N. Beckman, A. Nori, S. Rajamani, and R. Simmons. Proofs from tests. In *ISSTA*, 2008.
3. D. Brumley, C. Hartwig, Z. Liang, J. Newsome, P. Poosankam, D. Song, and H. Yin. Automatically identifying trigger-based behavior in malware. In *Botnet Analysis and Defense*. Springer, 2008.
4. P. Cousot and R. Cousot. Abstract interpretation. In *POPL*, 1977.
5. V. Ganesh and D. Dill. A decision procedure for bit-vectors and arrays. In *CAV*, 2007.
6. P. Godefroid, N. Klarlund, and K. Sen. DART: Directed automated random testing. In *PLDI*, 2005.
7. P. Godefroid, M. Levin, and D. Molnar. Automated whitebox fuzz testing. In *NDSS*, 2008.
8. R. Jhala and R. Majumdar. B2: Software model checking for C, 2009. www.cs.ucla.edu/~rupak/b2/.
9. J. Lim and T. Reps. A system for generating static analyzers for machine instructions. In *CC*, 2008.
10. K. Malmkjær. *Abstract Interpretation of Partial-Evaluation Algorithms*. PhD thesis, Dept. of Comp. and Inf. Sci., Kansas State Univ., 1993.
11. J. Morris. A general axiom of assignment. In M. Broy and G. Schmidt, editors, *Theor. Found. of Program. Methodology*. Reidel, 1982.
12. A. Mycroft and N. Jones. A relational framework for abstract interpretation. In *PADO*, 1985.
13. K. Sen, D. Marinov, and G. Agha. CUTE: A concolic unit testing engine for C. In *FSE*, 2005.
14. Y. Xie and A. Aiken. Saturn: A scalable framework for error detection using Boolean satisfiability. *TOPLAS*, 29(3), 2007.

A Appendix

In this section, we give correctness proofs for our generated primitives for symbolic execution, \mathcal{WLP} , and symbolic composition. These apply to the language PL (§3.2) and reinterpretation given in §4; the proofs for MC differ only slightly.

As a notational convenience, we do not distinguish between a *State* and a *LogicalStruct*. A *LogicalStruct* ι corresponds to the *State*: $((\iota\uparrow 1), (\iota\uparrow 2)F_\rho)$. Because, for PL, logical structures only contain the single function F_ρ , there is a one-to-one correspondence with states. Hence, whenever necessary (e.g. in the applications of $\mathcal{E}[\cdot]$, $\mathcal{B}[\cdot]$, and $\mathcal{I}[\cdot]$), we assume that that a *LogicalStruct* ι is coerced to $((\iota\uparrow 1), (\iota\uparrow 2)F_\rho)$.

A.1 Correctness of the Symbolic-Execution Method $\overline{\mathcal{I}}$

Lemma 1 (Relationship of $\overline{\mathcal{E}}$ to \mathcal{E} and $\overline{\mathcal{B}}$ to \mathcal{B}).

$$\begin{aligned} (1) \quad & \mathcal{T}[\overline{\mathcal{E}}[E]U]\iota = \mathcal{E}[E](\mathcal{U}[U]\iota) \\ (2) \quad & \mathcal{F}[\overline{\mathcal{B}}[BE]U]\iota = \mathcal{B}[BE](\mathcal{U}[U]\iota) \end{aligned}$$

Proof. The two lemmas are simultaneously proved using structural induction on E and BE , as shown below. Let U be $(\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\})$.

Note that the standard interpretations of *binop*, *relop*, and *boolop* coincide with those of *binop_L*, *relop_L*, and *boolop_L*. Thus, reasoning steps of the form $\text{binop}_L(\text{op2}_L) \rightsquigarrow \text{binop}(\text{op2})$ are shorthands for reasoning about each case, such as $\text{binop}_L(\boxed{+}) \rightsquigarrow \text{binop}(+)$, etc.

$$(1) \quad (i) \quad \mathcal{T}[\overline{\mathcal{E}}[c]U]\iota = \mathcal{T}[\overline{\text{const}}(c)]\iota = \mathcal{T}[c]\iota = \text{const}(c) = \mathcal{E}[c](\mathcal{U}[U]\iota)$$

(ii)

$$\begin{aligned} \text{lhs} & : \mathcal{T}[\overline{\mathcal{E}}[I]U]\iota = \mathcal{T}[\overline{\text{lookupState}} \ U \ I]\iota = \mathcal{T}[\overline{((U\uparrow 2)F_\rho)}((U\uparrow 1)I)]\iota \\ \text{rhs} & : \mathcal{E}[I](\mathcal{U}[U]\iota) \\ & = \mathcal{E}[I](\iota\uparrow 1)[I_i \mapsto \mathcal{T}[(U\uparrow 1)I_i]\iota, (\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota] \\ & = \text{lookupState}(\iota\uparrow 1)[I_i \mapsto \mathcal{T}[(U\uparrow 1)I_i]\iota, (\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota]I \\ & = ((\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota] (\mathcal{T}[(U\uparrow 1)I]\iota) \\ & = (\mathcal{F}\mathcal{E}[(U\uparrow 2)F_\rho]\iota) (\mathcal{T}[(U\uparrow 1)I]\iota) \\ & = \text{access}(\mathcal{F}\mathcal{E}[(U\uparrow 2)F_\rho]\iota, \mathcal{T}[(U\uparrow 1)I]\iota) \\ & = \mathcal{T}[\overline{((U\uparrow 2)F_\rho)}((U\uparrow 1)I)]\iota \end{aligned}$$

(iii)

$$\begin{aligned} \text{lhs} & : \mathcal{T}[\overline{\mathcal{E}}[\&I]U]\iota = \mathcal{T}[\overline{\text{lookupEnv}} \ U \ I]\iota = \mathcal{T}[\overline{(U\uparrow 1)I}]\iota \\ \text{rhs} & : \mathcal{E}[\&I](\mathcal{U}[U]\iota) \\ & = \mathcal{E}[\&I](\iota\uparrow 1)[I_i \mapsto \mathcal{T}[(U\uparrow 1)I_i]\iota, (\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota] \\ & = \text{lookupEnv}(\iota\uparrow 1)[I_i \mapsto \mathcal{T}[(U\uparrow 1)I_i]\iota, (\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota]I \\ & = \mathcal{T}[\overline{(U\uparrow 1)I}]\iota \end{aligned}$$

$$\begin{aligned}
& (iv) \\
\text{lhs} & : \mathcal{T}[\overline{\mathcal{E}}[*E]U]\iota = \mathcal{T}[\overline{\text{lookupStore } U}(\overline{\mathcal{E}}[E]U)]\iota = \mathcal{T}[(U\uparrow 2)F_\rho](\overline{\mathcal{E}}[E]U)\iota \\
\text{rhs} & : \mathcal{E}[*E](\mathcal{U}[U]\iota) \\
& = \mathcal{E}[*E](\iota\uparrow 1)[I_i \mapsto \mathcal{T}[(U\uparrow 1)I_i]\iota], (\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota] \\
& = \text{lookupStore}((\iota\uparrow 1)[I_i \mapsto \mathcal{T}[(U\uparrow 1)I_i]\iota], (\iota\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[(U\uparrow 2)F_j]\iota])(\mathcal{E}[E](\mathcal{U}[U]\iota)) \\
& = (\mathcal{F}\mathcal{E}[(U\uparrow 2)F_\rho]\iota)(\mathcal{E}[E](\mathcal{U}[U]\iota)) \\
& = (\mathcal{F}\mathcal{E}[(U\uparrow 2)F_\rho]\iota)(\mathcal{T}[\overline{\mathcal{E}}[E]U]\iota) // \text{by ind. via (1)} \\
& = \text{access}(\mathcal{F}\mathcal{E}[(U\uparrow 2)F_\rho]\iota, \mathcal{T}[\overline{\mathcal{E}}[E]U]\iota) \\
& = \mathcal{T}[(U\uparrow 2)F_\rho](\overline{\mathcal{E}}[E]U)\iota \\
& (v) \mathcal{T}[\overline{\mathcal{E}}[E_1 \text{ op2 } E_2]U]\iota \\
& = \mathcal{T}[\overline{\mathcal{E}}[E_1]U \text{ op2}_L \overline{\mathcal{E}}[E_2]U]\iota \\
& = \mathcal{T}[\overline{\mathcal{E}}[E_1]U]\iota \text{ binop}_L(\text{op2}_L) \mathcal{T}[\overline{\mathcal{E}}[E_2]U]\iota \\
& = \mathcal{E}[E_1](\mathcal{U}[U]\iota) \text{ binop}(\text{op2}) \mathcal{E}[E_2](\mathcal{U}[U]\iota) // \text{by ind. via (1)} \\
& = \mathcal{E}[E_1 \text{ op2 } E_2](\mathcal{U}[U]\iota) \\
& (vi) \mathcal{T}[\overline{\mathcal{E}}[BE ? E_1 : E_2]U]\iota \\
& = \mathcal{T}[\text{ite}(\overline{\mathcal{B}}[BE]U, \overline{\mathcal{E}}[E_1]U, \overline{\mathcal{E}}[E_2]U)]\iota \\
& = \text{cond}_L(\mathcal{F}[\overline{\mathcal{B}}[BE]U]\iota, \mathcal{T}[\overline{\mathcal{E}}[E_1]U]\iota, \mathcal{T}[\overline{\mathcal{E}}[E_2]U]\iota) \\
& = \mathcal{F}[\overline{\mathcal{B}}[BE]U]\iota ? \mathcal{T}[\overline{\mathcal{E}}[E_1]U]\iota : \mathcal{T}[\overline{\mathcal{E}}[E_2]U]\iota \\
& = \mathcal{B}[BE](\mathcal{U}[U]\iota) ? \mathcal{E}[E_1](\mathcal{U}[U]\iota) : \mathcal{E}[E_2](\mathcal{U}[U]\iota) // \text{by ind. via (1) and (2)} \\
& = \mathcal{E}[BE ? E_1 : E_2](\mathcal{U}[U]\iota) \\
& (2) (i) \mathcal{F}[\overline{\mathcal{B}}[\mathbb{T}]U]\iota = \mathcal{F}[\mathbb{T}]\iota = \mathbb{T} = \mathcal{B}[\mathbb{T}](\mathcal{U}[U]\iota) \\
& (ii) \mathcal{F}[\overline{\mathcal{B}}[\mathbb{F}]U]\iota = \mathcal{F}[\mathbb{F}]\iota = \mathbb{F} = \mathcal{B}[\mathbb{F}](\mathcal{U}[U]\iota) \\
& (iii) \mathcal{F}[\overline{\mathcal{B}}[E_1 \text{ rop } E_2]U]\iota \\
& = \mathcal{F}[\overline{\mathcal{E}}[E_1]U \text{ rop}_L \overline{\mathcal{E}}[E_2]U]\iota \\
& = \mathcal{T}[\overline{\mathcal{E}}[E_1]U]\iota \text{ relop}_L(\text{rop}_L) \mathcal{T}[\overline{\mathcal{E}}[E_2]U]\iota \\
& = \mathcal{E}[E_1](\mathcal{U}[U]\iota) \text{ relop}(\text{rop}) \mathcal{E}[E_2](\mathcal{U}[U]\iota) // \text{by ind. via (1)} \\
& = \mathcal{B}[E_1 \text{ rop } E_2](\mathcal{U}[U]\iota) \\
& (iv) \mathcal{F}[\overline{\mathcal{B}}[\neg BE_1]U]\iota = \mathcal{F}[\overline{\neg} \overline{\mathcal{B}}[BE_1]U]\iota = \neg \mathcal{F}[\overline{\mathcal{B}}[BE_1]U]\iota \\
& = \neg \mathcal{B}[BE_1](\mathcal{U}[U]\iota) // \text{by ind. via (2)} \\
& = \mathcal{B}[\neg BE_1](\mathcal{U}[U]\iota) \\
& (v) \mathcal{F}[\overline{\mathcal{B}}[BE_1 \text{ bop } BE_2]U]\iota \\
& = \mathcal{F}[\overline{\mathcal{B}}[BE_1]U \text{ bop}_L \overline{\mathcal{B}}[BE_2]U]\iota \\
& = \mathcal{F}[\overline{\mathcal{B}}[BE_1]U]\iota \text{ boolop}_L(\text{bop}_L) \mathcal{F}[\overline{\mathcal{B}}[BE_2]U]\iota \\
& = \mathcal{B}[BE_1](\mathcal{U}[U]\iota) \text{ boolop}(\text{bop}) \mathcal{B}[BE_2](\mathcal{U}[U]\iota) // \text{by ind. via (2)} \\
& = \mathcal{B}[BE_1 \text{ bop } BE_2](\mathcal{U}[U]\iota)
\end{aligned}$$

Theorem A1 For all $\iota \in \text{LogicalStruct}$, evaluating $\mathcal{U}[\overline{\mathcal{I}}[s]U]\iota$ is equivalent to running \mathcal{I} on s with an input state obtained from $\mathcal{U}[U]\iota$; that is,

$$\mathcal{U}[\overline{\mathcal{I}}[s]U]\iota = \mathcal{I}[s](\mathcal{U}[U]\iota).$$

Proof.

$$\begin{aligned}
(i) \mathcal{U}[\overline{\mathcal{T}}[I = E;]U]\iota &= \mathcal{U}[\overline{\text{updateStore } U} (\overline{\text{lookupEnv } U} I) (\overline{\mathcal{E}}[E]U)]\iota \\
&= \mathcal{U}[\overline{\text{updateStore } U} ((U\uparrow 1)I) (\overline{\mathcal{E}}[E]U)]\iota \\
&= \mathcal{U}[\overline{((U\uparrow 1), (U\uparrow 2)[F_\rho \mapsto ((U\uparrow 2)F_\rho)[(U\uparrow 1)I \mapsto \overline{\mathcal{E}}[E]U])}]]\iota \\
&= ((\mathcal{U}[U]\iota\uparrow 1), (\mathcal{U}[U]\iota\uparrow 2)[\mathcal{T}[(\mathcal{U}[U]\iota)I] \mapsto \mathcal{T}[\overline{\mathcal{E}}[E](\mathcal{U}[U]\iota)]]]) \\
&= ((\mathcal{U}[U]\iota\uparrow 1), (\mathcal{U}[U]\iota\uparrow 2)[(\mathcal{U}[U]\iota\uparrow 1)I \mapsto \mathcal{E}[E](\mathcal{U}[U]\iota)]) // \text{by Lem. 1(1)} \\
&= \text{updateStore } (\mathcal{U}[U]\iota) (\text{lookupEnv } (\mathcal{U}[U]\iota) I) (\mathcal{E}[E](\mathcal{U}[U]\iota)) \\
&= \mathcal{I}[I = E;](\mathcal{U}[U]\iota)
\end{aligned}$$

$$\begin{aligned}
(ii) \mathcal{U}[\overline{\mathcal{T}}[*I = E;]U]\iota &= \mathcal{U}[\overline{\text{updateStore } U} (\overline{\mathcal{E}}[I]U) (\overline{\mathcal{E}}[E]U)]\iota \\
&= \mathcal{U}[\overline{((U\uparrow 1), (U\uparrow 2)[F_\rho \mapsto ((U\uparrow 2)F_\rho)[\overline{\mathcal{E}}[I]U \mapsto \overline{\mathcal{E}}[E]U])}]]\iota \\
&= ((\mathcal{U}[U]\iota\uparrow 1), (\mathcal{U}[U]\iota\uparrow 2)[\mathcal{T}[\overline{\mathcal{E}}[I](\mathcal{U}[U]\iota) \mapsto \mathcal{T}[\overline{\mathcal{E}}[E](\mathcal{U}[U]\iota)]]]) \\
&= ((\mathcal{U}[U]\iota\uparrow 1), (\mathcal{U}[U]\iota\uparrow 2)[\mathcal{E}[I](\mathcal{U}[U]\iota) \mapsto \mathcal{E}[E](\mathcal{U}[U]\iota)]) // \text{by Lem. 1(1)} \\
&= \text{updateStore } (\mathcal{U}[U]\iota) (\mathcal{E}[I](\mathcal{U}[U]\iota)) (\mathcal{E}[E](\mathcal{U}[U]\iota)) \\
&= \mathcal{I}[*I = E;](\mathcal{U}[U]\iota)
\end{aligned}$$

$$\begin{aligned}
(iii) (\mathcal{U}[\overline{\mathcal{T}}[S_1 S_2]U]\iota) &= (\mathcal{U}[\overline{\mathcal{T}}[S_2](\overline{\mathcal{T}}[S_1]U)]\iota) \\
&= \mathcal{I}[S_2](\mathcal{U}[\overline{\mathcal{T}}[S_1]U]\iota) // \text{by induction} \\
&= \mathcal{I}[S_2](\mathcal{I}[S_1](\mathcal{U}[U]\iota)) // \text{by induction} \\
&= \mathcal{I}[S_1 S_2](\mathcal{U}[U]\iota)
\end{aligned}$$

A.2 Correctness of \mathcal{WLP}

Although \mathcal{WLP} is sometimes confused with the formula-manipulation operations used to obtain a formula that expresses it, or with the formula ψ that results, \mathcal{WLP} is really a semantic notion—the set of states *described* by ψ . For example, for any statement $s: \text{var} = \text{rhs}$; in a language that only has `int`-valued variables, and postcondition formula φ , the formula $\varphi[\text{var} \leftarrow \text{rhs}]$ obtained by substitution is not the only formula that expresses $\mathcal{WLP}(s, \varphi)$. In fact, there are an infinity of acceptable formulas. We characterize what constitutes an acceptable formula as follows:

Definition 1 (Acceptable \mathcal{WLP} Formula.) ψ is an acceptable formula for $\mathcal{WLP}(s, \varphi)$ iff, for all $\iota \in \text{LogicalStruct}$,

$$\mathcal{F}[\psi]\iota = \mathcal{F}[\varphi](\mathcal{I}[s]\iota).$$

That is, ψ holds in the pre-state structure ι exactly when φ holds in the post-state structure $(\mathcal{I}[s]\iota)$.

Lemma 2 (Relationship of $\overline{\mathcal{T}}$ to \mathcal{T} , $\overline{\mathcal{F}}$ to \mathcal{F} , $\overline{\mathcal{F}\mathcal{E}}$ to $\mathcal{F}\mathcal{E}$).

$$\begin{aligned}
(1) \mathcal{T}[\overline{\mathcal{T}}[T]U]\iota &= \mathcal{T}[T](\mathcal{U}[U]\iota) \\
(2) \mathcal{F}[\overline{\mathcal{F}}[\varphi]U]\iota &= \mathcal{F}[\varphi](\mathcal{U}[U]\iota) \\
(3) \mathcal{F}\mathcal{E}[\overline{\mathcal{F}\mathcal{E}}[FE]U]\iota &= \mathcal{F}\mathcal{E}[FE](\mathcal{U}[U]\iota)
\end{aligned}$$

Proof. The three lemmas are simultaneously proved using structural induction on T , φ , and FE , as shown below. Let U be $(\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\})$. (Thus, $T_i = (U \uparrow 1)I_i$ and $FE_j = (U \uparrow 2)F_j$.) Let f be $(\iota \uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[\![FE_j]\!] \iota]$.

$$(1) (i) \mathcal{T}[\overline{\mathcal{T}}[\![c]\!]U] \iota = \mathcal{T}[\![c]\!] \iota = \text{const}(c) = \mathcal{T}[\![c]\!](\mathcal{U}[\![U]\!] \iota)$$

(ii)

$$\begin{aligned} \text{lhs} &= \mathcal{T}[\overline{\mathcal{T}}[\![I]\!]U] \iota = \mathcal{T}[\overline{\text{lookupId}} U I] \iota = \mathcal{T}[(U \uparrow 1)I] \iota \\ \text{rhs} &= \mathcal{T}[\![I]\!](\mathcal{U}[\![U]\!] \iota) = \mathcal{T}[\![I]\!](\iota \uparrow 1)[I_i \mapsto \mathcal{T}[\![T_i]\!] \iota, f) \\ &= \text{lookupId}((\iota \uparrow 1)[I_i \mapsto \mathcal{T}[\![T_i]\!] \iota, f) I \\ &= \mathcal{T}[(U \uparrow 1)I] \iota \end{aligned}$$

$$\begin{aligned} (iii) \mathcal{T}[\overline{\mathcal{T}}[\![T_1 \text{ op2}_L T_2]\!]U] \iota &= \mathcal{T}[\overline{\mathcal{T}}[\![T_1]\!]U \text{ op2}_L \overline{\mathcal{T}}[\![T_2]\!]U] \iota \\ &= \mathcal{T}[\overline{\mathcal{T}}[\![T_1]\!]U] \iota \text{ binop}_L(\text{op2}_L) \mathcal{T}[\overline{\mathcal{T}}[\![T_2]\!]U] \iota \\ &= \mathcal{T}[\![T_1]\!](\mathcal{U}[\![U]\!] \iota) \text{ binop}_L(\text{op2}_L) \mathcal{T}[\![T_2]\!](\mathcal{U}[\![U]\!] \iota) // \text{ by ind. via (1)} \\ &= \mathcal{T}[\![T_1 \text{ op2}_L T_2]\!](\mathcal{U}[\![U]\!] \iota) \end{aligned}$$

$$\begin{aligned} (iv) \mathcal{T}[\overline{\mathcal{T}}[\![ite(\varphi, T_1, T_2)]\!]U] \iota &= \mathcal{T}[\![ite(\overline{\mathcal{F}}[\![\varphi]\!]U, \overline{\mathcal{T}}[\![T_1]\!]U, \overline{\mathcal{T}}[\![T_2]\!]U)] \iota \\ &= \text{cond}_L(\mathcal{F}[\overline{\mathcal{F}}[\![\varphi]\!]U] \iota, \mathcal{T}[\overline{\mathcal{T}}[\![T_1]\!]U] \iota, \mathcal{T}[\overline{\mathcal{T}}[\![T_2]\!]U] \iota) \\ &= \mathcal{F}[\overline{\mathcal{F}}[\![\varphi]\!]U] \iota ? \mathcal{T}[\overline{\mathcal{T}}[\![T_1]\!]U] \iota : \mathcal{T}[\overline{\mathcal{T}}[\![T_2]\!]U] \iota \\ &= \mathcal{F}[\varphi](\mathcal{U}[\![U]\!] \iota) ? \mathcal{T}[\![T_1]\!](\mathcal{U}[\![U]\!] \iota) : \mathcal{T}[\![T_2]\!](\mathcal{U}[\![U]\!] \iota) // \text{ by ind. via (1) and (2)} \\ &= \mathcal{F}[\varphi ? T_1 : T_2](\mathcal{U}[\![U]\!] \iota) \end{aligned}$$

$$\begin{aligned} (v) \mathcal{T}[\overline{\mathcal{T}}[\![FE(T)]\!]U] \iota &= \mathcal{T}[\overline{\mathcal{F}\mathcal{E}}[\![FE]\!]U(\overline{\mathcal{T}}[\![T]\!]U)] \iota \\ &= (\mathcal{F}\mathcal{E}[\overline{\mathcal{F}\mathcal{E}}[\![FE]\!]U] \iota)(\mathcal{T}[\overline{\mathcal{T}}[\![T]\!]U] \iota) \\ &= (\mathcal{F}\mathcal{E}[\![FE]\!](\mathcal{U}[\![U]\!] \iota))(\mathcal{T}[\![T]\!](\mathcal{U}[\![U]\!] \iota)) // \text{ by ind. via (3)} \\ &= \mathcal{T}[\![FE(T)]\!](\mathcal{U}[\![U]\!] \iota) \end{aligned}$$

$$(2) (i) \mathcal{F}[\overline{\mathcal{F}}[\![\mathbb{T}]\!]U] \iota = \mathcal{F}[\![\mathbb{T}]\!] \iota = \mathbb{T} = \mathcal{F}[\![\mathbb{T}]\!](\mathcal{U}[\![U]\!] \iota)$$

$$(ii) \mathcal{F}[\overline{\mathcal{F}}[\![\mathbb{F}]\!]U] \iota = \mathcal{F}[\![\mathbb{F}]\!] \iota = \mathbb{F} = \mathcal{F}[\![\mathbb{F}]\!](\mathcal{U}[\![U]\!] \iota)$$

$$\begin{aligned} (iii) \mathcal{F}[\overline{\mathcal{F}}[\![T_1 \text{ rop}_L T_2]\!]U] \iota &= \mathcal{F}[\overline{\mathcal{T}}[\![T_1]\!]U \text{ relop}_L(\text{rop}_L) \overline{\mathcal{T}}[\![T_2]\!]U] \iota \\ &= \mathcal{T}[\overline{\mathcal{T}}[\![T_1]\!]U] \iota \text{ relop}_L(\text{rop}_L) \mathcal{T}[\overline{\mathcal{T}}[\![T_2]\!]U] \iota \\ &= \mathcal{T}[\![T_1]\!](\mathcal{U}[\![U]\!] \iota) \text{ relop}_L(\text{rop}_L) \mathcal{T}[\![T_2]\!](\mathcal{U}[\![U]\!] \iota) // \text{ by ind. via (1)} \\ &= \mathcal{F}[\![T_1 \text{ rop}_L T_2]\!](\mathcal{U}[\![U]\!] \iota) \end{aligned}$$

$$\begin{aligned} (iv) \mathcal{F}[\overline{\mathcal{F}}[\![\neg]\!] \varphi_1] U] \iota &= \mathcal{F}[\![\neg]\!] \mathcal{F}[\![\varphi_1]\!] U] \iota \\ &= \neg \mathcal{F}[\![\varphi_1]\!] U] \iota \\ &= \neg \mathcal{F}[\varphi_1](\mathcal{U}[\![U]\!] \iota) // \text{ by ind. via (2)} \\ &= \mathcal{F}[\![\neg]\!] \varphi_1](\mathcal{U}[\![U]\!] \iota) \end{aligned}$$

$$\begin{aligned} (v) \mathcal{F}[\overline{\mathcal{F}}[\![\varphi_1 \text{ bop}_L \varphi_2]\!]U] \iota &= \mathcal{F}[\overline{\mathcal{F}}[\![\varphi_1]\!]U \text{ boolop}_L(\text{bop}_L) \overline{\mathcal{F}}[\![\varphi_2]\!]U] \iota \\ &= \mathcal{F}[\overline{\mathcal{F}}[\![\varphi_1]\!]U] \iota \text{ boolop}_L(\text{bop}_L) \mathcal{F}[\overline{\mathcal{F}}[\![\varphi_2]\!]U] \iota \\ &= \mathcal{F}[\varphi_1](\mathcal{U}[\![U]\!] \iota) \text{ boolop}_L(\text{bop}_L) \mathcal{F}[\varphi_2](\mathcal{U}[\![U]\!] \iota) // \text{ by ind. via (2)} \\ &= \mathcal{F}[\varphi_1 \text{ bop}_L \varphi_2](\mathcal{U}[\![U]\!] \iota) \end{aligned}$$

$$\begin{aligned}
(3) (i) \\
\text{lhs} &= \mathcal{F}\mathcal{E}[\overline{\mathcal{F}\mathcal{E}}[F]U]\iota = \mathcal{F}\mathcal{E}[\overline{\text{lookupId}} U F]\iota = \mathcal{F}\mathcal{E}[(U\uparrow 2)F]\iota \\
\text{rhs} &= \mathcal{F}\mathcal{E}[F](\mathcal{U}[U]\iota) \\
&= \mathcal{F}\mathcal{E}[F](\iota\uparrow 1)[I_i \mapsto \mathcal{T}[[T_i]\iota], f) \\
&= \text{lookupFuncId}((\iota\uparrow 1)[I_i \mapsto \mathcal{T}[[T_i]\iota], f) F \\
&= \mathcal{F}\mathcal{E}[(U\uparrow 2)F]\iota \\
(ii) & \mathcal{F}\mathcal{E}[\overline{\mathcal{F}\mathcal{E}}[FE_0[T_1 \mapsto T_2]]U]\iota \\
&= \mathcal{F}\mathcal{E}[(\overline{\mathcal{F}\mathcal{E}}[FE_0]U)(\overline{\mathcal{T}}[[T_1]U] \mapsto \overline{\mathcal{T}}[[T_2]U])]\iota \\
&= \mathcal{F}\mathcal{E}[(\overline{\mathcal{F}\mathcal{E}}[FE_0]U)]\iota[\mathcal{T}[\overline{\mathcal{T}}[[T_1]U]\iota] \mapsto \mathcal{T}[\overline{\mathcal{T}}[[T_2]U]\iota]] \\
&= \mathcal{F}\mathcal{E}[FE_0](\mathcal{U}[U]\iota)[\mathcal{T}[[T_1](\mathcal{U}[U]\iota)] \mapsto \mathcal{T}[[T_2](\mathcal{U}[U]\iota)]] // \text{by ind. via (1)} \\
&= \mathcal{F}\mathcal{E}[FE_0[T_1 \mapsto T_2]](\mathcal{U}[U]\iota)
\end{aligned}$$

Theorem A2 For any Stmt s and Formula $\varphi, \psi := \overline{\mathcal{F}}[\varphi](\overline{\mathcal{T}}[s]U_{id})$ is an acceptable \mathcal{WLP} formula for φ with respect to s .

Proof. For all $\iota \in \text{LogicalStruct}$,

$$\begin{aligned}
\mathcal{F}[\psi]\iota &= \mathcal{F}[\overline{\mathcal{F}}[\varphi](\overline{\mathcal{T}}[s]U_{id})]\iota \\
&= \mathcal{F}[\varphi](\mathcal{U}[\overline{\mathcal{T}}[s]U_{id}]\iota) // \text{by Lem. 2} \\
&= \mathcal{F}[\varphi](\mathcal{I}[s](\mathcal{U}[U_{id}]\iota)) // \text{by Thm. A1} \\
&= \mathcal{F}[\varphi](\mathcal{I}[s]\iota)
\end{aligned}$$

and therefore, by Defn. 1, $\overline{\mathcal{F}}[\varphi](\overline{\mathcal{T}}[s]U_{id})$ is an acceptable \mathcal{WLP} formula for φ with respect to s .

A.3 Correctness of the Symbolic-Composition Method $\overline{\mathcal{U}}$

Theorem A3 For all $U_1, U_2 \in \text{StructUpdate}$ and $\iota \in \text{LogicalStruct}$, $\mathcal{U}[\overline{\mathcal{U}}[U_2]U_1]\iota = \mathcal{U}[U_2](\mathcal{U}[U_1]\iota)$.

Proof. Let $U_2 = (\{I_i \leftrightarrow T_i\}, \{F_j \leftrightarrow FE_j\})$. Let I_k and F_m range over Id and $FuncId$, respectively.

$$\begin{aligned}
& \mathcal{U}[\overline{\mathcal{U}}[U_2]U_1]\iota \\
&= \mathcal{U}[\overline{((U_1\uparrow 1)[I_i \mapsto \overline{\mathcal{T}}[[T_i]U_1], (U_1\uparrow 2)[F_j \mapsto \overline{\mathcal{F}\mathcal{E}}[FE_j]U_1])}]\iota \\
&= \mathcal{U} \left[\overline{\left(\begin{array}{l} \{I_k \mapsto ((U_1\uparrow 1)[I_i \mapsto \overline{\mathcal{T}}[[T_i]U_1])I_k\}, \\ \{F_m \mapsto ((U_1\uparrow 2)[F_j \mapsto \overline{\mathcal{F}\mathcal{E}}[FE_j]U_1])F_m\} \end{array} \right)} \right] \iota \\
&= ((\iota\uparrow 1)[I_k \mapsto \mathcal{T}[\overline{((U_1\uparrow 1)[I_i \mapsto \overline{\mathcal{T}}[[T_i]U_1])I_k}]\iota], \\
&\quad (\iota\uparrow 2)[F_m \mapsto \mathcal{F}\mathcal{E}[\overline{((U_1\uparrow 2)[F_j \mapsto \overline{\mathcal{F}\mathcal{E}}[FE_j]U_1])F_m}]\iota]) \\
&= ((\iota\uparrow 1)[I_{k(\neq i)} \mapsto \mathcal{T}[\overline{((U_1\uparrow 1)I_k}]\iota][I_i \mapsto \mathcal{T}[\overline{\mathcal{T}}[[T_i]U_1}]\iota], \\
&\quad (\iota\uparrow 2)[F_{m(\neq j)} \mapsto \mathcal{F}\mathcal{E}[\overline{(U_1\uparrow 2)F_m}]\iota][F_j \mapsto \mathcal{F}\mathcal{E}[\overline{\mathcal{F}\mathcal{E}}[FE_j]U_1]]\iota) \\
&= // \text{by Lem. 2} \\
&\quad ((\iota\uparrow 1)[I_{k(\neq i)} \mapsto \mathcal{T}[\overline{(U_1\uparrow 1)I_k}]\iota][I_i \mapsto \mathcal{T}[[T_i](\mathcal{U}[U_1]\iota)], \\
&\quad (\iota\uparrow 2)[F_{m(\neq j)} \mapsto \mathcal{F}\mathcal{E}[\overline{(U_1\uparrow 2)F_m}]\iota][F_j \mapsto \mathcal{F}\mathcal{E}[FE_j](\mathcal{U}[U_1]\iota)]) \\
&= ((\mathcal{U}[U_1]\iota)\uparrow 1)[I_i \mapsto \mathcal{T}[[T_i](\mathcal{U}[U_1]\iota)], \\
&\quad ((\mathcal{U}[U_1]\iota)\uparrow 2)[F_j \mapsto \mathcal{F}\mathcal{E}[FE_j](\mathcal{U}[U_1]\iota))] \\
&= \mathcal{U}[U_2](\mathcal{U}[U_1]\iota)
\end{aligned}$$