

Gradual Program Verification

Johannes Bader¹, Jonathan Aldrich², and Éric Tanter³

¹ Microsoft Corporation, Redmond, USA,
jobader@microsoft.com

² Institute for Software Research, Carnegie Mellon University, Pittsburgh, USA
jonathan.aldrich@cs.cmu.edu

³ PLEIAD Lab, Computer Science Dept (DCC), University of Chile, Santiago, Chile
etanter@dcc.uchile.cl

Abstract. Both static and dynamic program verification approaches have significant disadvantages when considered in isolation. Inspired by research on gradual typing, we propose *gradual verification* to seamlessly and flexibly combine static and dynamic verification. Drawing on general principles from abstract interpretation, and in particular on the recent Abstracting Gradual Typing methodology of Garcia *et al.*, we systematically derive a gradual verification system from a static one. This approach yields, by construction, a gradual verification system that is compatible with the original static system, but overcomes its rigidity by resorting to dynamic verification when desired. As with gradual typing, the programmer can control the trade-off between static and dynamic checking by tuning the (im)precision of pre- and postconditions. The formal semantics of the gradual verification system and the proofs of its properties, including the gradual guarantees of Siek *et al.*, have been fully mechanized in the Coq proof assistant.

1 Introduction

Program verification techniques have the potential to improve the correctness of programs, by exploiting pre- and postconditions specified in formulas drawn from a given logic, such as Hoare logic [8]. Unfortunately, traditional approaches to verification have a number of shortcomings, as illustrated next.

Example 1.

```
int withdraw(int balance, int amount)
  requires (balance ≥ amount) ensures (balance ≥ 0) {
    return balance - amount; // returns the new balance
  }

int balance := 100;
balance := withdraw(balance, 30);
balance := withdraw(balance, 40);
```

In this case, we reason about a variable `balance` representing some bank account. The contract (pre- and postconditions) of `withdraw` specifies that it may only be called if the balance is high enough to withdraw the given amount, ensuring that no negative balance is reached. There are a number of ways to verify Example 1. We briefly discuss static and dynamic verification, including hybrid approaches. We then introduce *gradual verification* as an approach that has the potential to overcome a number of their shortcomings.

Static verification. Formal methods like Hoare logic are used to establish *statically* that a program is *valid*, *i.e.* satisfies its specification. In Example 1, the static verifier proves both that `withdraw` itself complies with its contract and that the three statements below are valid, *e.g.* that the precondition of `withdraw` is satisfied prior to both calls.

A lack of detailed contracts may prevent the verifier from establishing that a program is valid. In Example 1, verification of the second call to `withdraw` in fact fails: after the first call, the verifier knows from the postcondition that (`balance` ≥ 0), which is insufficient to justify that (`balance` ≥ 40) as required for the second call. Deriving such knowledge would require a stronger postcondition such as `balance = old(balance) - amount`. However, this is not the postcondition that was provided by the programmer, perhaps intentionally (*e.g.* if the intent was to focus on some weaker correctness properties) or perhaps due to limited expressiveness of the underlying logic (notation such as `old(x)` may not exist). In general, a verification tool might also fail to prove program properties due to undecidability of the underlying logic or practical limitations of the specific tool implementation.

Hoare logic has been extended to more powerful logics like separation logic [15] and implicit dynamic frames [20]. Yet, the requirement of rigorous annotation of contracts remains an issue in these settings. Due to space limitations and to capture the core ideas of gradual verification, this paper focuses on a simple Hoare logic. We have formalized an extension to implicit dynamic frames and implemented a prototype, which can both be found at <http://olydis.github.io/GradVer/impl/HTML5wp/>

Dynamic verification. An alternative approach is to use *dynamic* verification to ensure that a program adheres to its specification at runtime, by turning the contract into *runtime checks*. A contract violation causes a runtime exception to be thrown, effectively preventing the program from entering a state that contradicts its specification. In Example 1, a dynamic verification approach would not raise any error because the `balance` is in fact sufficient for both calls to succeed. Note that because contracts are checked at runtime, one can even use arbitrary programs as contracts, and not just formulas drawn from a logic [6].

Meyer’s Design by Contract methodology [12] integrated writing contracts in this way as an integral part of the design process, with the Eiffel language automatically performing dynamic contract verification [11]. Dynamic verification has also notably been used to check JML specifications [3], and has been extended to the case of separation logic by Nguyen *et al.* [14]. Note that unlike the static approach, the dynamic approach only requires programmers to encode the

properties they care about as pre- and postconditions, and does not require extra work for the sake of avoiding false negatives. However, the additional checks impose runtime overhead that may not always be acceptable. Furthermore, violations of the specification are no longer detected ahead of time.

Hybrid approaches. Recognizing that static and dynamic checking have complementary advantages, some approaches to combine them have emerged. In particular, with the Java Modeling Language (JML) [2] and Code Contracts [5], it is possible to use the same specifications for either static or dynamic verification. Additionally, Nguyen *et al.* explored a hybrid approach to reduce the overhead of their approach to runtime checking for separation logic, by exploiting static information [14].

Although useful, these techniques do not support a smooth continuum between static and dynamic verification. With the JML approach, engineers enable static *or* dynamic verification; the two checking regimes do not interact. Nguyen *et al.* use the static checker to optimize runtime checks, but do not try to report static verification failures because it is difficult to distinguish failures due to contradictions in the specification (which the developer should be warned about) from failures due to leaving out parts of the specification (which could have been intentional underspecification, and thus should not produce a warning). Their runtime checking approach also requires the specification of heap footprints to match in pre- and post-conditions, which like many static checking approaches forces programmers to do extra specification work to avoid false negatives.

Gradual verification. Because this tension between static and dynamic verification is reminiscent of the tension between static and dynamic type checking, we propose to draw on research on *gradual typing* [18,17,7] to develop a flexible approach to program verification, called *gradual verification*. Gradual typing supports both static and dynamic checking and the entire spectrum in between, driven by the precision of programmer annotations [19]. Similarly, gradual verification introduces a notion of *imprecise contracts*, supporting a continuum between static and dynamic verification. A static checker can analyze a gradually-specified program and warn the programmer of inconsistencies between specifications and code, including contracts that are intended to be fully precise but are not strong enough, as well as contracts that contradict one another despite possible imprecision in each. On the other hand, the static checker will not produce warnings that arise from a contract that is intentionally imprecise; in these cases, runtime checking is used instead. Programmers can rely on a *gradual guarantee* stating that reducing the precision of specifications never breaks the verifiability (and reduceability) of a program. This guarantee, originally formulated by Siek *et al.* in the context of gradual types [19], ensures that programmers can choose their desired level of precision without artificial constraints imposed by the verification technology.

It is worth noting that the similarly named work “The Gradual Verifier” [1] focuses on *measuring the progress of static verification*. Their verification technique “GraVy” is neither sound nor complete and does not comply with the gradual guarantee.

Gradual verification is not only useful in cases of missing information (*e.g.* when reusing a library that is not annotated) but also to overcome limitations of the static verification system as motivated by Example 1. Furthermore, programmers can gradually evolve and refine static annotations. As they do so, they are rewarded by progressively *increased static correctness guarantees* and progressively *decreased runtime checking*, supporting a pay-as-you-go cost model.

Specifically, we support imprecision by introducing an unknown formula $?$ that acts like a wildcard during static verification. Semantically, the static checker will optimistically accept a formula containing $?$ as long as there exists some interpretation of $?$ that makes the formula valid. As we learn more information about the program state at runtime, the dynamic checker ensures that some valid instantiation still exists. Crucially, the unknown formula can be combined with static formulas, forming *imprecise* formulas. For instance, going back to Example 1, we can write the imprecise postcondition $(\text{balance} \geq 0) \wedge ?$ in order to enable gradual reasoning, resulting in an optimistic interpretation of $?$ as $(\text{balance} \geq 40)$ when statically proving the precondition of the second call. At runtime, this interpretation is checked, to ensure soundness.

Note that the postcondition we suggest is only *partially unknown*, preserving the *static knowledge* $(\text{balance} \geq 0)$. This not only allows us to prove certain goals (*e.g.* $(\text{balance} \neq -10)$) without requiring any dynamic checks, but also to statically reject programs that provably contradict this knowledge (*e.g.* if a subsequent call had $\text{balance} = -10$ as precondition).

Contributions. This paper is the first development of the ideas of gradual typing in the context of program logics for verification. More precisely, we first introduce a simple statically-verified language called SVL, along with its associated program logic. We then adapt the *Abstracting Gradual Typing* methodology (AGT) [7] to the verification setting and show in section 3 how the static semantics of a gradually-verified language GVL can be derived from SVL using principles of abstract interpretation. Section 4 develops GVL’s dynamic semantics. Here, we deviate from the AGT approach and instead propose injecting a minimal amount of runtime assertion checks, yielding a pay-as-you-go cost model. Finally, Section 5 briefly discusses GVL_{IDF}, an extension of our approach to heap-allocated objects and an extended logic with implicit dynamic frames [20].

Limitations. Our approach for dynamic semantics requires assertions to be evaluable at runtime, naturally limiting the logic usable for annotations. The AGT methodology (based on combining the proof-trees at runtime) is not restricted that way, so it may be the ideal starting point for gradual verification in presence of higher-order logic assertions.

The formal semantics of GVL and the proofs of its properties have been fully mechanized in the Coq proof assistant and can be found at <http://olydis.github.io/GradVer/impl/HTML5wp/>. The site also includes a report with the formal treatment of the extended logic, as well as an interactive online prototype of GVL_{IDF}. Due to limited space, some figures contain only selected parts of definitions. Complete definitions can be found online as well.

$program ::= \overline{procedure} s$	$s \in \text{STMT} ::= \text{skip} \mid s_1; s_2 \mid T x \mid x := e$
$procedure ::= T m(\overline{T x}) \text{contract} \{ s \}$	$\mid x := m(x) \mid \text{assert } \phi$
$contract ::= \text{requires } \phi \text{ ensures } \phi$	$e \in \text{EXPR} ::= v \mid x \mid (e \oplus e)$
$T ::= \text{int}$	$x \in \text{VAR} ::= \text{result} \mid \text{ident} \mid \text{old}(\text{ident})$
$\oplus ::= + \mid - \mid \dots$	$v \in \text{VAL} ::= n$
$\odot ::= = \mid \neq \mid < \mid \dots$	$\phi \in \text{FORMULA} ::= \text{true} \mid (e \odot e) \mid \phi \wedge \phi$
and syntactic sugar	$\text{return } e \stackrel{\text{def}}{=} \text{result} := e \quad \text{and} \quad T x := e \stackrel{\text{def}}{=} T x; x := e$

Fig. 1. SVL: Syntax

2 SVL: Statically Verified Language

In the following sections, we describe a simple statically verified language called SVL. We formalize its syntax, semantics and soundness.

2.1 Syntax

Figure 1 shows the syntax of SVL. Programs consist of a collection of procedures and a designated statement resembling the entry point (“main procedure”). We include the empty statement, statement sequences, variable declarations, variable assignments, procedure calls, and assertions. All statements are in A-normal form, which is not essential semantically but does simplify the formalism. Procedures have contracts consisting of a pre- and postcondition, which are formulas. Formulas can be the constant **true**, binary relations between expressions, and a conjunction \wedge . Expressions can occur within a formula or variable assignment, and consist of variables, constants and arithmetic operations.⁴

For the remainder of this work we only consider well-formed programs: variables are declared before use, procedure names are unique and contracts only contain variables that are in scope. More specifically, a precondition may only contain the procedure’s parameters x_i , a postcondition may only contain the special variable **result** and the dummy variables **old**(x_i).

To simplify the presentation of semantics, we will give rules for procedures that have exactly one parameter.

2.2 Dynamic Semantics

We now describe the dynamic semantics of SVL. SVL has a small-step semantics $\cdot \longrightarrow \cdot : \text{STATE} \rightarrow \text{STATE}$ (see Fig. 2) that describes discrete transitions between program states. Program states that are not in the domain of this partial function are said to be *stuck*, which happens if an assertion does not hold or a contract is

⁴ Our approach is directly applicable to, say, further control structures, a richer type system or formulas that are arbitrary boolean expressions.

violated before/after a call. In Section 2.3, we define a static verification system whose soundness result implies that valid SVL programs do not get stuck.

Program states. Program states consist of a stack, *i.e.* $\text{STATE} = \text{STACK}$ where:

$$S \in \text{STACK} ::= E \cdot S \mid \text{nil} \quad \text{where} \quad E \in \text{STACKFRAME} = \text{ENV} \times \text{STMT}$$

A stack frame consists of a local variable environment $\rho \in \text{ENV} = \text{VAR} \rightarrow \text{VAL}$ and a continuation statement.

Evaluation. An expression e is evaluated according to a big-step evaluation relation $\rho \vdash e \Downarrow v$, yielding value v using local variable environment $\rho \in \text{ENV}$ of the top-most stack-frame. The definition is standard: variables are looked up in ρ , and the resulting values are combined according to standard arithmetic rules. Example: $[x \mapsto 3] \vdash x + 5 \Downarrow 8$

The evaluation of a formula in a given environment is specified by the predicate $\cdot \models \cdot \subseteq \text{ENV} \times \text{FORMULA}$. We assume standard evaluation semantics for standard concepts like equality. We also say that a formula *describes* a certain (infinite) set of environments (exactly the environments under which it holds), yielding natural definitions for formula satisfiability and implication.

Definition 1 (Denotational Formula Semantics).

Let $\llbracket \cdot \rrbracket : \text{FORMULA} \rightarrow \mathcal{P}(\text{ENV})$ be defined as $\llbracket \phi \rrbracket \stackrel{\text{def}}{=} \{ \rho \in \text{ENV} \mid \rho \models \phi \}$

Definition 2 (Formula Satisfiability). A formula ϕ is satisfiable if and only if $\llbracket \phi \rrbracket \neq \emptyset$. Let $\text{SATFORMULA} \subset \text{FORMULA}$ be the set of satisfiable formulas.

Definition 3 (Formula Implication). $\phi_1 \Rightarrow \phi_2$ if and only if $\llbracket \phi_1 \rrbracket \subseteq \llbracket \phi_2 \rrbracket$

Reduction rules. We define a standard small-step reduction semantics for statements (Fig. 2). **SSASSERT** ensures that assertions are stuck if the asserted formula is not satisfied. **SSCALL** sets up a new stack frame and makes sure that the procedure’s precondition is satisfied by the newly set up context. Similarly, **SSCALLEXIT** ensures that the postcondition is satisfied before returning control to the call site. Note our use of auxiliary functions **procedure** and **mpost** in order to retrieve a procedure’s definition or postcondition using that procedure’s name. Formally, we assume all rules and definitions are implicitly parameterized with the “executing” program $p \in \text{PROGRAM}$ from which to extract this information. When required for disambiguation, we explicitly annotate reduction arrows with the executing program p , as in \rightarrow_p .

Note that **SSCALL** also initializes **old**(x'), which allows assertions and most importantly the postcondition to reference the parameter’s initial value. In reality, no additional memory is required to maintain this value since it is readily available as $\rho(x)$, *i.e.* it lives in the stack frame of the call site. For a program to be well-formed, it may not write to **old**(x') in order to enable this reasoning.

$$\begin{array}{c}
 \frac{\rho \models \phi}{\langle \rho, \mathbf{assert} \ \phi \rangle \cdot S \longrightarrow \langle \rho, \mathbf{skip} \rangle \cdot S} \text{SSASSERT} \\
 \\
 \frac{\rho \vdash e \Downarrow v \quad \rho' = \rho[x \mapsto v]}{\langle \rho, x := e \rangle \cdot S \longrightarrow \langle \rho', \mathbf{skip} \rangle \cdot S} \text{SSASSIGN} \\
 \\
 \frac{\text{procedure}(m) = T_r \ m(T \ x') \ \mathbf{requires} \ \phi_p \ \mathbf{ensures} \ \phi_q \ \{ r \} \quad \rho \vdash x \Downarrow v \quad \rho' = [x' \mapsto v, \mathbf{old}(x') \mapsto v] \quad \rho' \models \phi_p}{\langle \rho, y := m(x); s \rangle \cdot S \longrightarrow \langle \rho', r \rangle \cdot \langle \rho, y := m(x); s \rangle \cdot S} \text{SSCALL} \\
 \\
 \frac{\text{post}(m) = \phi_q \quad \rho' \models \phi_q}{\langle \rho', \mathbf{skip} \rangle \cdot \langle \rho, y := m(x); s \rangle \cdot S \longrightarrow \langle \rho[y \mapsto \rho'(\mathbf{result})], s \rangle \cdot S} \text{SSCALEXIT}
 \end{array}$$

Fig. 2. SVL: Small-step semantics (selected rules)

$$\begin{array}{ll}
 \text{WLP}(\mathbf{skip}, \phi) & = \phi & \text{WLP}(s_1; s_2, \phi) & = \text{WLP}(s_1, \text{WLP}(s_2, \phi)) \\
 \text{WLP}(x := e, \phi) & = \phi[e/x] & \text{WLP}(\mathbf{assert} \ \phi_a, \phi) & = \phi_a \wedge \phi \\
 \text{WLP}(y := m(x), \phi) & = \max_{\Rightarrow} \{ \phi' \mid y \notin \text{FV}(\phi') \wedge \\
 & \quad \phi' \Rightarrow \text{mpre}(m)[x/\text{mparam}(m)] \wedge \\
 & \quad (\phi' \wedge \text{mpost}(m)[x, y/\text{old}(\text{mparam}(m)), \mathbf{result}]) \Rightarrow \phi \}
 \end{array}$$

Fig. 3. SVL: Weakest precondition (selected rules)

2.3 Static Verification

We define the static verification of SVL contracts through a weakest liberal precondition calculus [4]. This syntax-directed approach (compared to, say, Hoare logic, which has an existential in its sequence rule) will be useful for the dynamic semantics of our gradual language (will be pointed out again later).

Definition 4 (Valid Procedure).

A procedure with contract **requires** ϕ_p **ensures** ϕ_q , parameter x and body s is considered valid if $\phi_p \Rightarrow \text{WLP}(s, \phi_q)[x/\mathbf{old}(x)]$ holds.

We define $\text{WLP} : \text{STMT} \times \text{FORMULA} \rightarrow \text{FORMULA}$ as shown in Figure 3. WLP is standard for the most part. The rule for calls computes a maximal formula ϕ' (i.e. minimum information content) that is sufficient to imply both the procedure's precondition and the overall postcondition ϕ with the help of the procedure's postcondition.

Definition 5 (Valid Program). A program with entry point statement s is considered valid if $\mathbf{true} \Rightarrow \text{WLP}(s, \mathbf{true})$ holds and all procedures are valid. ⁵

⁵ Note that one can demand more than **true** to hold at the final state by simply ending the program with an assertion statement.

$$\begin{aligned}
\text{sWLP}(s \cdot \text{nil}, \phi) &= \text{WLP}(s, \phi) \cdot \text{nil} \\
\text{sWLP}(s \cdot (y := m(x); s') \cdot \bar{s}, \phi) &= \text{WLP}(s, \text{mpost}(m)) \cdot \text{sWLP}((y := m(x); s') \cdot \bar{s}, \phi)
\end{aligned}$$

Fig. 4. Weakest precondition across call boundaries

Example 2 (Static Checker of SVL). We demonstrate the resulting behavior of SVL’s static checker using example 1, but with varying contracts:

```

requires (balance ≥ amount)
ensures (result = old(balance) - old(amount))
  withdraw is valid since the WLP of the body, given the postcondition, is
  (balance - amount = old(balance) - old(amount)). Substitution gives
  (balance - amount = balance - amount) which is trivially implied by
  the precondition. The overall program is also valid since the main proce-
  dure’s WLP is (100 ≥ 70) which is implied by true.
requires (balance ≥ amount) ensures (result ≥ 0)    (as in example 1)
  withdraw is valid since the body’s WLP is (balance - amount ≥ 0) which
  matches the precondition. However, the program is not valid: The WLP of the
  second call is (balance ≥ 40) which is not implied by the postcondition of
  the first call. As a result, the WLP of the entire main procedure is undefined.
requires (balance ≥ 0) ensures (result ≥ 0)
  Validating withdraw fails since the body’s WLP (same as above) is not im-
  plied by the precondition.

```

2.4 Soundness

Verified programs should respect contracts and assertions. We have formulated the runtime semantics of SVL such that they get stuck if contracts or assertions are violated. As a result, *soundness* means that valid programs do not get stuck. In particular, we can use a syntactic progress/preservation argument [22].

If the WLP of a program is satisfied by the current state, then execution will not get stuck (progress) and after each step of execution, the WLP of the remaining program is again satisfied by the new state (preservation). We use a progress and preservation formulation of soundness not just because it allows us to reason about non-terminating programs, but mainly because this will allow us to naturally deal with the runtime checking needs of gradual verification.

To simplify reasoning about states with multiple stack frames, we extend the definition of WLP to accept a stack of statements and return a stack of preconditions, as shown in Figure 4. Note that WLP as defined previously can only reason about procedure calls atomically since an element of STMT cannot encode intermediate states of an ongoing procedure call. In contrast sWLP works across call boundaries by accepting a stack of statements and recursively picking up the postconditions of procedures which are currently being executed.

While before we defined what makes procedures as a whole valid, we can now validate arbitrary intermediate program states, e.g. we can say that

$$\text{sWLP} \left(\begin{array}{c} \text{return balance - amount} \\ \text{b2} := \text{withdraw}(\text{b1}, \text{a}) \quad , (\text{b2} \neq -1) \wedge (\text{a} = 4) \\ \text{nil} \end{array} \right) = \begin{array}{c} (\text{balance} - \text{amount} \geq 0) \\ (\text{b1} \geq \text{a}) \wedge (\text{a} = 4) \\ \text{nil} \end{array}$$

where `withdraw` is defined as in example 1. If \bar{s} are the continuations of some arbitrary program state $\pi \in \text{STATE}$, then $\text{sWLP}(\bar{s}, \text{true})$ is the precondition for \bar{s} . If $\text{sWLP}(\bar{s}, \text{true})$ holds in the variable environments $\bar{\rho}$ of π , respectively, then soundness guarantees that the program does not get stuck. In the following, we extend the notion of validity to arbitrary intermediate program states in order to formalize progress and preservation. Validity of states is an invariant that relates the static and dynamic semantics of valid SVL programs.

Definition 6 (Valid state). *We call the state $\langle \rho_n, s_n \rangle \cdot \dots \cdot \langle \rho_1, s_1 \rangle \cdot \text{nil} \in \text{STATE}$ valid if $\rho_i \models \text{sWLP}_i(s_n \cdot \dots \cdot s_1 \cdot \text{nil}, \text{true})$ for all $1 \leq i \leq n$. ($\text{sWLP}_i(\bar{s}, \phi)$ is the i -th component of $\text{sWLP}(\bar{s}, \phi)$)*

Validity of the initial program state follows from validity of the program (Def. 5).

Proposition 1 (SVL: Progress). *If $\pi \in \text{STATE}$ is a valid state and $\pi \notin \{\langle \rho, \text{skip} \rangle \cdot \text{nil} \mid \rho \in \text{ENV}\}$ then $\pi \longrightarrow \pi'$ for some $\pi' \in \text{STATE}$.*

Proposition 2 (SVL: Preservation). *If π is a valid state and $\pi \longrightarrow \pi'$ for some $\pi' \in \text{STATE}$ then π' is a valid state.*

3 GVL: Static Semantics

Having defined SVL, we can now derive its gradual counterpart GVL, which supports gradual program verification thanks to *imprecise* contracts. We follow the abstract interpretation perspective on gradual typing [7], AGT for short. In this sense, we introduce *gradual formulas* as formulas that can include the *unknown formula*, denoted $?$:

$\tilde{\phi} ::= \phi \mid \phi \wedge ?$ and standalone formula $?$ as syntactic sugar for $\text{true} \wedge ?$

We define $\tilde{\text{FORMULA}}$ as the set of all gradual formulas. The syntax of GVL is unchanged save for the use of gradual formulas in contracts: $\text{contract} ::= \text{requires } \phi \text{ ensures } \tilde{\phi}$. In Sections 3.2 to 3.4 we *lift* the predicates and functions SVL uses for verification from the static domain to the gradual domain, yielding a gradual verification logic for GVL.

3.1 Interpretation of Gradual Formulas

We call ϕ in $\phi \wedge ?$ *static part* of the imprecise formula and define a helper function $\text{static} : \tilde{\text{FORMULA}} \rightarrow \text{FORMULA}$ that extracts the static part of a gradual formula, i.e. $\text{static}(\phi) = \phi$ and $\text{static}(\phi \wedge ?) = \phi$. Following the AGT approach, a gradual formula is given meaning by concretization to *the set of static formulas* that it represents.

Definition 7 (Concretization of gradual formulas).

$\gamma : \tilde{\text{FORMULA}} \rightarrow \mathcal{P}^{\text{FORMULA}}$ is defined as:

$$\begin{aligned} \gamma(\phi) &= \{ \phi \} \\ \gamma(\phi \wedge ?) &= \{ \phi' \in \text{SATFORMULA} \mid \phi' \Rightarrow \phi \} \quad \text{if } \phi \in \text{SATFORMULA} \\ \gamma(\phi \wedge ?) &\text{ undefined otherwise} \end{aligned}$$

A fully-precise formula concretizes to the singleton set. Importantly, we only concretize imprecise formulas to precise formulas that are *satisfiable*. Note that the concretization of any gradual formula always implies the static part of that formula. The notion of *precision* between formulas, reminiscent of the notion of precision between gradual types [19], is naturally induced by concretization [7]:

Definition 8 (Precision). $\tilde{\phi}_1$ is more precise than $\tilde{\phi}_2$, written $\tilde{\phi}_1 \sqsubseteq \tilde{\phi}_2$ if and only if $\gamma(\tilde{\phi}_1) \subseteq \gamma(\tilde{\phi}_2)$.

3.2 Lifting Predicates

The semantics of SVL makes use of predicates that operate on formulas, namely formula implication and formula evaluation. As GVL must operate on gradual formulas, these predicates are lifted in order to deal with gradual formulas in a consistent way. We propose the following definitions of consistent formula evaluation and implication.

Definition 9 (Consistent Formula Evaluation).

Let $\cdot \models \cdot \subseteq \text{ENV} \times \tilde{\text{FORMULA}}$ be defined as $\rho \models \tilde{\phi} \iff \rho \models \text{static}(\tilde{\phi})$

Definition 10 (Consistent Formula Implication).

Let $\cdot \rightsquigarrow \cdot \subseteq \tilde{\text{FORMULA}} \times \tilde{\text{FORMULA}}$ be defined inductively as

$$\frac{\phi_1 \Rightarrow \text{static}(\tilde{\phi}_2)}{\phi_1 \rightsquigarrow \tilde{\phi}_2} \tilde{\text{IMPLSTATIC}} \quad \frac{\phi \in \text{SATFORMULA} \quad \phi \Rightarrow \text{static}(\tilde{\phi}_2)}{\phi_1 \wedge ? \rightsquigarrow \tilde{\phi}_2} \tilde{\text{IMPLGRAD}}$$

In rule $\tilde{\text{IMPLGRAD}}$, ϕ represents a *plausible* formula represented by $\phi_1 \wedge ?$.

Abstract interpretation. Garcia et al. [7] define consistent liftings of predicates as their *existential liftings*:

Definition 11 (Consistent Predicate Lifting). The consistent lifting $\tilde{P} \subseteq \tilde{\text{FORMULA}} \times \tilde{\text{FORMULA}}$ of a predicate $P \subseteq \text{FORMULA} \times \text{FORMULA}$ is defined as:

$$\tilde{P}(\tilde{\phi}_1, \tilde{\phi}_2) \stackrel{\text{def}}{\iff} \exists \phi_1 \in \gamma(\tilde{\phi}_1), \phi_2 \in \gamma(\tilde{\phi}_2). P(\phi_1, \phi_2)$$

Our definitions above are proper predicate liftings.

Lemma 1 (Consistent Formula Evaluation and Implication).

$\cdot \models \cdot$ (Def. 9) is a consistent lifting of $\cdot \models \cdot$ and $\cdot \rightsquigarrow \cdot$ (Def. 10) is a consistent lifting of $\cdot \Rightarrow \cdot$.

3.3 Lifting Functions

Deriving gradual semantics from SVL also involves lifting *functions* that operate on formulas, most importantly WLP (Definition 3). Figure 5 gives the definition of $\widetilde{\text{WLP}} : \text{STMT} \times \widetilde{\text{FORMULA}} \rightarrow \widetilde{\text{FORMULA}}$, the consistent lifting of WLP. For

$$\begin{aligned}
 \widetilde{\text{WLP}}(\text{skip}, \tilde{\phi}) &= \tilde{\phi} & \widetilde{\text{WLP}}(s_1; s_2, \tilde{\phi}) &= \widetilde{\text{WLP}}(s_1, \widetilde{\text{WLP}}(s_2, \tilde{\phi})) \\
 \widetilde{\text{WLP}}(x := e, \tilde{\phi}) &= \tilde{\phi}[e/x] & \widetilde{\text{WLP}}(\text{assert } \phi_a, \tilde{\phi}) &= \phi_a \wedge \tilde{\phi} \\
 \widetilde{\text{WLP}}(y := m(x), \tilde{\phi}) &= \begin{cases} \phi' & \text{if } \tilde{\phi}, \text{mpre}(m), \text{mpost}(m) \in \text{FORMULA} \\ \phi' \wedge ? & \text{otherwise} \end{cases} \\
 &\text{where } \phi' = \max_{\Rightarrow} \{ \phi'' \mid y \notin \text{FV}(\phi'') \wedge (\phi'' \Rightarrow \text{mpre}(m)[x/\text{mparam}(m)]) \wedge \\
 &\quad (\phi'' \wedge \text{mpost}(m)[x, y/\text{old}(\text{mparam}(m)), \text{result}]) \Rightarrow \tilde{\phi} \}
 \end{aligned}$$

Fig. 5. GVL: Weakest precondition (selected rules)

most statements $\widetilde{\text{WLP}}$ is defined almost identical to WLP, however, calls are more complex. Note that for calls, $\widetilde{\text{WLP}}$ not only has to deal with the fact that $\tilde{\phi}$ is a gradual formula, but also that procedure m may now have imprecise contracts. In a sense, the function is lifted w.r.t. three formula parameters, two of them referenced through the procedure's name. To accomplish this, we first determine the static part ϕ' of the result which is analogous to the WLP, but resorting to lifted predicates. Next, we determine whether it would be sufficient to return ϕ' unmodified, or whether it is plausible that the precondition must be stronger. If all three influencing formulas are precise $\widetilde{\text{WLP}}$ should coincide with WLP, so ϕ' is returned. Otherwise, ϕ' might have been chosen too weak, which is counteracted by making it imprecise.

Abstract interpretation. Again, AGT [7] formalizes the notion of consistent functions using an *abstraction* function α that maps a set of static formulas back to the most precise gradual formula that represents this set, such that $\langle \gamma, \alpha \rangle$ forms a Galois connection.

Definition 12 (Abstraction of formulas). Let $\alpha : \mathcal{P}^{\text{Sat}}\text{FORMULA} \rightarrow \widetilde{\text{FORMULA}}$ be defined as $\alpha(\bar{\phi}) = \min_{\sqsubseteq} \{ \tilde{\phi} \in \widetilde{\text{FORMULA}} \mid \bar{\phi} \subseteq \gamma(\tilde{\phi}) \}$

α is partial since \min does not necessarily exist, e.g. $\alpha(\{(x \neq x), (x = x)\})$ is undefined. Using concretization for gradual parameters and abstraction for return values one can consistently lift (partial) functions:

Definition 13 (Consistent Function Lifting). Given a partial function $f : \text{FORMULA} \rightarrow \text{FORMULA}$, its consistent lifting $\tilde{f} : \widetilde{\text{FORMULA}} \rightarrow \widetilde{\text{FORMULA}}$ is defined as $\tilde{f}(\tilde{\phi}) = \alpha(\{ f(\phi) \mid \phi \in \gamma(\tilde{\phi}) \})$

Lemma 2 (Consistent WLP). $\widetilde{\text{WLP}}$ is a consistent lifting of WLP.

3.4 Lifting the Verification Judgment

Gradual verification in GVL must deal with imprecise contracts. The static verifier of SVL uses WLP and implication to determine whether contracts and the overall program are valid (Def. 4, 5). Because contracts in GVL may be imprecise, we have to resort to $\widetilde{\text{WLP}}$ (Fig. 5) and consistent implication (Def. 10).

Example 3 (Static Checker of GVL). Static semantics of SVL and GVL coincide for precise contracts, so example 2 applies to GVL without modification. We extend the example with imprecise contracts:

requires (`balance` \geq `amount`) **ensures** (`result` \geq 0) \wedge ?

Note the similarity to the precise contract in example 1 which causes GVL's static checker to reject the main procedure. However, with the imprecise postcondition we now have (`balance` \geq 0) \wedge ? $\widetilde{\Rightarrow}$ (`balance` \geq 40).

As a result, the static checker optimistically accepts the program. At the same time, it is not *guaranteed* that the precondition is satisfied at runtime without additional checks. We expect GVL's runtime semantics (Section 4) to add such checks as appropriate. These runtime checks should succeed for the main procedure of example 1, however they should fail if we modify the main program as follows, withdrawing more money than available:

```
int b := 100; b := withdraw(b, 30); b := withdraw(b, 80);
```

Static checking succeeds since (`b` \geq 0) \wedge ? $\widetilde{\Rightarrow}$ (`b` \geq 80), but `b`'s value at runtime will not satisfy the formula. Note that the presence of imprecision does not necessarily imply success of static checking:

```
int b := 100; b := withdraw(b, 30); assert (b < 0);
```

It is *implausible* that this program is valid since (`b` \geq 0) \wedge ? $\widetilde{\Rightarrow}$ (`b` < 0) does not hold. However, further weakening `withdraw`'s postcondition to ? would again result in static acceptance but dynamic rejection.

requires ? **ensures** (`result` = `old(balance)` - `old(amount`)) \wedge ?

This contract demonstrates that imprecision must not necessarily result in runtime checks. The body's $\widetilde{\text{WLP}}$ is ?, which is implied by the annotated precondition ? without having to be optimistic (i.e. resort to the plausibility interpretation). We expect that an efficient runtime semantics, like the one we discuss in Section 4.3, adds no runtime overhead through checks here.

4 GVL: Dynamic Semantics

Accepting a gradually-verified program means that it is *plausible* that the program remains valid during each step of its execution, precisely as it is guaranteed by soundness of SVL. To prevent a GVL program from stepping from a valid

state into an invalid state, we extend the dynamic semantics of SVL with (desirably minimal) runtime checks. As soon as the validity of the execution is no longer plausible, these checks cause the program to step into a dedicated error state. Example 3 motivates this behavior.

Soundness. We revise the soundness definition of SVL to the gradual setting which will guide the upcoming efforts to achieve soundness via runtime assertion checks. Validity of states (Def. 6) relies on sWLP (Fig. 4) which itself consumes postconditions of procedures. Hence, GVL uses a consistent lifting of sWLP which we define analogous to Fig. 4, but based on WLP . Save for using sWLP instead of WLP , valid states of GVL are defined just like those of SVL.

We expect there to be error derivations $\pi \xrightarrow{\sim} \mathbf{error}$ whenever it becomes implausible that the remaining program can be run safely. Note that we do not extend STATE , but instead define $\cdot \xrightarrow{\sim} \cdot \subseteq \text{STATE} \times (\text{STATE} \cup \{\mathbf{error}\})$. As a result, we can leave Prop. 2 (Preservation) untouched.

In Section 4.1 we derive a naive runtime semantics driven by the soundness criteria of GVL. We then examine the properties of the resulting gradually verified language. In Section 4.3 we discuss optimizing this approach by combining WLP with strongest preconditions SP in order to determine statically-guaranteed information that can be used to minimize the runtime checks ahead of time.

4.1 Naive semantics

We start with a trivially correct but expensive strategy of adding runtime assertions to each execution step, checking whether the new state would be valid (preservation), right before actually transitioning into that state (progress).⁶

Let $\rho'_{1..m}, \rho_{1..n} \in \text{ENV}$, $s'_{1..m}, s_{1..n} \in \text{STMT}$

If $\langle \rho'_m, s'_m \rangle \cdot \dots \cdot \langle \rho'_1, s'_1 \rangle \cdot \text{nil} \longrightarrow \langle \rho_n, s_n \rangle \cdot \dots \cdot \langle \rho_1, s_1 \rangle \cdot \text{nil}$ holds⁷, then

$$\langle \rho'_m, s'_m \rangle \cdot \dots \cdot \langle \rho'_1, s'_1 \rangle \cdot \text{nil} \xrightarrow{\sim} \begin{cases} \langle \rho_n, s_n \rangle \cdot \dots \cdot \langle \rho_1, s_1 \rangle \cdot \text{nil} & \text{if } (\rho_n \widetilde{\models} \widetilde{\phi}_n) \wedge \dots \wedge (\rho_1 \widetilde{\models} \widetilde{\phi}_1) \\ & \text{where } \widetilde{\phi}_n \cdot \dots \cdot \widetilde{\phi}_1 \cdot \text{nil} = \text{sWLP}(s_n \cdot \dots \cdot s_1 \cdot \text{nil}, \mathbf{true}) \\ \mathbf{error} & \text{otherwise} \end{cases}$$

Before showing how to implement the above semantics, we confirm its soundness: Progress of GVL follows from progress of SVL. The same is true for preservation: in the first reduction case, validity of the resulting state follows from preservation of SVL.

⁶ Note the difference between runtime assertions and the **assert** statement. The former checks assertions at runtime, transitioning into a dedicated exceptional state on failure. The latter is a construct of a statically verified language, and is hence implementable as a no-operation.

⁷ SSCALL and SSCALLEXIT as defined in Fig. 2 are not defined for gradual formulas. Thus, we adjust those rules to use consistent evaluation $\widetilde{\models}$ instead of \models . Since $\widetilde{\models}$ coincides with \models for precise formulas, this is a conservative extension of SVL.

While we can draw the implementation of $\cdot \longrightarrow \cdot$ from SVL, implementing the case condition $(\rho_n \widetilde{\models} \widetilde{\phi}_n) \wedge \dots \wedge (\rho_1 \widetilde{\models} \widetilde{\phi}_1)$ results in overhead. As a first step, we can heavily simplify this check using inductive properties of our language: Stack frames besides the topmost one are not changed by a single reduction, *i.e.* $\rho_{n-1}, \dots, \rho_1, s_{n-1}, \dots, s_1$ stay untouched. It follows that $\widetilde{\phi}_i$ for $1 \leq i < n$ remains unchanged since changes in s_n do not affect lower components of $\widetilde{\text{sWLP}}$ (see Fig. 4). As a result, it is sufficient to check $\rho_n \widetilde{\models} \widetilde{\text{sWLP}}_n(s_n \cdot \dots \cdot s_1 \cdot \text{nil}, \text{true})$.

Recall how we argued that a weakest precondition approach is more suited for the dynamic semantics of GVL than Hoare logic. Due to the syntax-directed sequence rule, all potentially occurring $\widetilde{\text{sWLP}}_n$ are partial results of statically *precomputed* preconditions. Contrast this with a gradual sequence rule of Hoare logic: $\{?\}\text{skip}; \text{skip}\{?\}$ could be accepted statically by, say, instantiating the existential with $(x = 3)$, which is allowed if both premises of the rule are lifted independently. However, the partial result $\{(x = 3)\}\text{skip}\{?\}$ has no (guaranteed) relationship with the next program state since the existential was chosen too strong. Any attempt to fix the gradual sequence rule by imposing additional restrictions on the existential must necessarily involve a weakest precondition calculus, applied to the suffix of the sequence.

4.2 Properties of GVL

Before discussing practical aspects of GVL, we turn to its formal properties: GVL is a *sound, gradual* language. The following three properties are formalized and proven in Coq.

Soundness. Our notion of soundness for GVL coincides with that of SVL, save for the possibility of runtime errors. Indeed, it is up to the dynamic semantics of GVL to make up for the imprecisions that weaken the statics of GVL.

Lemma 3 (Soundness of GVL). *GVL is sound:*

Progress *If $\pi \in \text{STATE}$ is a valid state and $\pi \notin \{\langle \rho, \text{skip} \rangle \cdot \text{nil} \mid \rho \in \text{ENV}\}$ then $\pi \longrightarrow \pi'$ for some $\pi' \in \text{STATE}$ or $\pi \longrightarrow \text{error}$.*

Preservation *If π is a valid state and $\pi \longrightarrow \pi'$ for some $\pi' \in \text{STATE}$ then π' is a valid state.*

We call the state $\langle \rho_n, s_n \rangle \cdot \dots \cdot \langle \rho_1, s_1 \rangle \cdot \text{nil}$ valid if $\rho_i \widetilde{\models} \widetilde{\text{sWLP}}_i(s_n \cdot \dots \cdot s_1 \cdot \text{nil}, \text{true})$ for all $1 \leq i \leq n$.

Conservative extension. GVL is a conservative extension of SVL, meaning that both languages coincide on fully-precise programs. This property is true *by construction*. Indeed, the definition of concretization and consistent lifting captures this property, which thus percolates to the entire verification logic. In order for the dynamic semantics to be a conservative extension, GVL must progress whenever SVL does, yielding the same continuation. This is the case since the reduction rules of GVL coincide with those of SVL for fully-precise annotations (the runtime checks succeed due to preservation of SVL, so we do not step to **error**).

Gradual guarantees. Siek *et al.* formalize a number of criteria for gradually-typed languages [19], which we can adapt to the setting of program verification. In particular, the *gradual guarantee* captures the smooth continuum between static and dynamic verification. More precisely, it states that typeability (here, verifiability) and reducibility are *monotone* with respect to precision. We say a program p_1 is more precise than program p_2 ($p_1 \sqsubseteq p_2$) if p_1 and p_2 are equivalent except in terms of contracts and if p_1 's contracts are more precise than p_2 's contracts. A contract **requires** ϕ_p^1 **ensures** ϕ_q^1 is more precise than contract **requires** ϕ_p^2 **ensures** ϕ_q^2 if $\phi_p^1 \sqsubseteq \phi_p^2$ and $\phi_q^1 \sqsubseteq \phi_q^2$.

In particular, the static gradual guarantee for verification states that a valid gradual program is still valid when we reduce the precision of contracts.

Proposition 3 (Static gradual guarantee of verification).

Let $p_1, p_2 \in \text{PROGRAM}$ such that $p_1 \sqsubseteq p_2$. If p_1 is valid then p_2 is valid.

The dynamic gradual guarantee states that a valid program that takes a step still takes the same step if we reduce the precision of contracts.

Proposition 4 (Dynamic gradual guarantee of verification).

Let $p_1, p_2 \in \text{PROGRAM}$ such that $p_1 \sqsubseteq p_2$ and $\pi \in \text{STATE}$.

If $\pi \xrightarrow{p_1} \pi'$ for some $\pi' \in \text{STATE}$ then $\pi \xrightarrow{p_2} \pi'$.

This also means that if a gradual program fails at runtime, then making its contracts more precise will *not* eliminate the error. In fact, doing so may only make the error manifest *earlier* during runtime or manifest *statically*. This is a fundamental property of gradual verification: a runtime verification error reveals a fundamental mismatch between the gradual program and the underlying verification discipline.

4.3 Practical aspects

Residual checks. Compared to SVL, the naive semantics adds a runtime assertion to every single reduction. Assuming that the cost of checking an assertion is proportional to the formula size, *i.e.* proportional to the size of the WLP of the remaining statement, this is highly unsatisfying. The situation is even worse if the entire GVL program has fully-precise annotations, because then the checks are performed even though they are not necessary for safety.

We can reduce formula sizes given static information, expecting formulas to vanish (reduce to **true**) in the presence of fully-precise contracts and gradually grow with the amount of imprecision introduced, yielding a pay-as-you-go cost model. Example 4 illustrates this relationship.

Example 4 (Residual checks). Consider the following variation of **withdraw**:

```
int withdraw(int balance, int amount)
  requires ? ensures result ≥ 0 {
    // WLP: (balance - amount ≥ 0) ∧ (amount > 0)
    assert amount > 0;
```

```

// WLP: balance - amount ≥ 0
balance = balance - amount;
// WLP: balance ≥ 0
return balance;
// WLP: result ≥ 0
}

```

Precomputed preconditions are annotated. Following the naive semantics (Section 4.1), these are to be checked *before* entering the corresponding state, to ensure soundness. However, many of these checks are redundant. When entering the procedure (*i.e.* stepping to the state prior to the assertion statement), we must first check $\phi_1 = (\text{balance} - \text{amount} \geq 0) \wedge (\text{amount} > 0)$. According to the naive semantics, in order to execute the assertion statement, we would then check $\phi_2 = (\text{balance} - \text{amount} \geq 0)$. Fortunately, it is derivable statically that this check must definitely succeed: The strongest postcondition of `assert amount > 0` given ϕ_1 is again ϕ_1 . Since $\phi_1 \Rightarrow \phi_2$ there is no point in checking ϕ_2 at runtime. Similar reasoning applies to both remaining statements, making all remaining checks redundant. Only the initial check remains, which is itself directly dependent on the imprecision of the precondition. Preconditions $(\text{balance} - \text{amount} \geq 0) \wedge ?$ or $(\text{amount} > 0) \wedge ?$ would allow dropping the corresponding term of the formula, the conjunction of both (with or without a $?$) allows dropping the entire check.

The above example motivates the formalization of a strongest postcondition function $\widetilde{\text{SP}}$ and function `diff` which takes a formula $\tilde{\phi}_a$ to check, a formula $\tilde{\phi}_k$ known to be true and calculates a residual formula $\text{diff}(\tilde{\phi}_a, \tilde{\phi}_k)$ sufficient to check in order to guarantee that $\tilde{\phi}_a$ holds.

Definition 14 (Strongest postcondition). Let $\text{SP} : \text{STMT} \times \text{FORMULA} \rightarrow \text{FORMULA}$ be defined as: $\text{SP}(s, \phi) = \min_{\Rightarrow} \{ \phi' \in \text{FORMULA} \mid \phi \Rightarrow \text{WLP}(s, \phi') \}$

Furthermore let $\widetilde{\text{SP}} : \text{STMT} \times \widetilde{\text{FORMULA}} \rightarrow \widetilde{\text{FORMULA}}$ be the consistent lifting (Def. 13) of SP .

Definition 15 (Reducing formulas).

Let $\text{diff} : \text{FORMULA} \times \text{FORMULA} \rightarrow \text{FORMULA}$ be defined as:

$$\text{diff}(\phi_a, \phi_k) = \max_{\Rightarrow} \{ \phi \in \text{FORMULA} \mid (\phi \wedge \phi_k \Rightarrow \phi_a) \wedge (\phi \wedge \phi_k \in \text{SatFormula}) \}$$

Furthermore let $\widetilde{\text{diff}} : \widetilde{\text{FORMULA}} \times \widetilde{\text{FORMULA}} \rightarrow \widetilde{\text{FORMULA}}$ be defined as:

$$\widetilde{\text{diff}}(\phi_a, \tilde{\phi}_k) = \text{diff}(\phi_a, \text{static}(\tilde{\phi}_k)) \quad \widetilde{\text{diff}}(\phi_a \wedge ?, \tilde{\phi}_k) = \text{diff}(\phi_a, \text{static}(\tilde{\phi}_k)) \wedge ?$$

Both SP and diff can be implemented approximately by only approximating min/max. Likewise, an implementation of $\widetilde{\text{SP}}$ may err towards imprecision. As a result, the residual formulas may be larger than necessary.⁸

⁸ Even a worst case implementation of $\widetilde{\text{SP}}(s, \tilde{\phi})$ as $?$ will only result in no reduction of the checks, but not affect soundness.

$$\begin{array}{c}
 \frac{\langle \rho'_n, (s; s_n) \rangle \cdot \dots \longrightarrow \langle \rho_n, s_n \rangle \cdot \dots}{\langle \rho'_n, (s; s_n) \rangle \cdot \dots \widetilde{\longrightarrow} \langle \rho_n, s_n \rangle \cdot \dots} \widetilde{\text{SSLOCAL}} \\
 \\
 \frac{\langle \rho_{n-1}, (y := m(x); s_{n-1}) \rangle \cdot \dots \longrightarrow \langle \rho_n, \text{mbody}(m) \rangle \cdot \dots \quad \rho_n \models \text{diff}(\widetilde{\text{WLP}}_n(\text{mbody}(m), \text{mpost}(m)), \text{mpre}(m))}{\langle \rho_{n-1}, (y := m(x); s_{n-1}) \rangle \cdot \dots \widetilde{\longrightarrow} \langle \rho_n, \text{mbody}(m) \rangle \cdot \dots} \widetilde{\text{SSCALL}} \\
 \\
 \frac{\rho_n \models \text{diff}(\widetilde{\text{sWLP}}_n(s_n \cdot \dots, \text{true}), \widetilde{\text{SP}}(y := m(x), \widetilde{\text{sWLP}}_n((y := m(x); s_n) \cdot \dots, \text{true}))) \quad \langle \rho'_{n+1}, \text{skip} \rangle \cdot \langle \rho'_n, (y := m(x); s_n) \rangle \cdot \dots \longrightarrow \langle \rho_n, s_n \rangle \cdot \dots}{\langle \rho'_{n+1}, \text{skip} \rangle \cdot \langle \rho'_n, (y := m(x); s_n) \rangle \cdot \dots \widetilde{\longrightarrow} \langle \rho_n, s_n \rangle \cdot \dots} \widetilde{\text{SSCALEXIT}}
 \end{array}$$

Fig. 6. Dynamic semantics with reduced checks.

Definition 16 (Approximate algorithm for $\widetilde{\text{diff}}$).

```

FORMULA  $\widetilde{\text{diff}}$  (FORMULA  $\widetilde{\phi}_a$ , FORMULA  $\widetilde{\phi}_b$ )
  let  $\widetilde{\phi}_1 \wedge \widetilde{\phi}_2 \wedge \dots \wedge \widetilde{\phi}_n := \widetilde{\phi}_a$  (such that all  $\widetilde{\phi}_i$  are atomic)
  for  $i$  from 1 to  $n$ 
    if  $\llbracket \widetilde{\phi}_b \rrbracket \cap \llbracket \widetilde{\phi}_1 \wedge \dots \wedge \widetilde{\phi}_{i-1} \wedge \widetilde{\phi}_{i+1} \dots \wedge \widetilde{\phi}_n \rrbracket \subseteq \llbracket \widetilde{\phi}_a \rrbracket$ 
       $\widetilde{\phi}_i := \text{true}$ 
  return  $\widetilde{\phi}_1 \wedge \dots \wedge \widetilde{\phi}_n$  // one may drop the true terms
    
```

Figure 6 shows the dynamic semantics using residual checks (omitting error reductions). Runtime checks of reductions operating on a single stack frame vanish completely as there exists no source of imprecision in that subset of GVL. This property can be formalized as: *For all $s \in \text{STMT}$ that do not contain a call, $\text{diff}(\widetilde{\text{sWLP}}_n(s_n \cdot \dots, \text{true}), \widetilde{\text{SP}}(s, \widetilde{\text{sWLP}}_n((s; s_n) \cdot \dots, \text{true}))) \in \{\text{true}, ?\}$* The check in $\widetilde{\text{SSCALL}}$ vanishes if $\text{mpre}(m)$ is precise. The check in $\widetilde{\text{SSCALEXIT}}$ vanishes if $\text{mpost}(m)$ is precise. Recall that Example 4 demonstrates this effect: We concluded that only the initial check is necessary and derived that it also vanishes if the precondition is precise.

If $\text{mpost}(m)$ is imprecise, the assertion is equivalent to

$$\rho_n \models \text{diff}(\text{diff}(\widetilde{\text{sWLP}}_n(s_n \cdot \dots, \text{true}), \text{mpre}(m)[x/\text{mparam}(m)]), \text{mpost}(m)[x, y/\text{old}(\text{mparam}(m)), \text{result}])$$

which exemplifies the pay-as-you-go relationship between level of imprecision and run-time overhead: both $\text{mpre}(m)$ and $\text{mpost}(m)$ contribute to the reduction of $\widetilde{\text{sWLP}}_n(s_n \cdot \dots, \text{true})$, *i.e.* increasing their precision results in smaller residual checks.

Pay-as-you-go. To formally establish the pay-as-you-go characteristic of gradual verification, we introduce a simple cost model for runtime contract checking.

Let $\text{size}(\tilde{\phi})$ be the number of conjunctive terms of (the static part of) $\tilde{\phi}$, e.g. $\text{size}(\mathbf{x} = 3) \wedge (\mathbf{y} \neq \mathbf{z}) \wedge ?) = 2$. We assume that measure to be proportional to the time needed to evaluate a given formula. Let $\text{checksize}(p)$ be the sum of the size of all residual checks in program p .

Proposition 5 (Pay-as-you-go overhead).

- a) *Given programs p_1, p_2 such that $p_1 \sqsubseteq p_2$, then $\text{checksize}(p_1) \leq \text{checksize}(p_2)$.*
- b) *If a program p has only precise contracts, then $\text{checksize}(p) = 0$.*

5 Scaling up to Implicit Dynamic Frames

We applied our approach to a richer program logic, namely *implicit dynamic frames* (IDF) [20], which enables reasoning about shared mutable state. The starting point is an extended statically verified language similar to Chalice [10], called SVL_{IDF} . Compared to SVL, the language includes classes with publicly-accessible fields and instance methods. We add field assignments and object creation. Formulas may also contain field *accessibility predicates* **acc** from IDF and use the separating conjunction $*$ instead of regular (non-separating) conjunction \wedge . An accessibility predicate **acc**($e.f$) denotes exclusive access to the field $e.f$. It justifies accessing $e.f$ both in the source code (e.g. $\mathbf{x.f} := 3$ or $\mathbf{y} := \mathbf{x.f}$) and in the formula itself (e.g. $\mathbf{acc}(\mathbf{x.f}) * (\mathbf{x.f} \neq 4)$), which is called *framing*. Compared to SVL, the main challenge in gradualizing SVL_{IDF} is framing.

The linear logic ensures that accessibility predicates cannot be duplicated, hence entitling them to represent *exclusive* access to a field. SVL_{IDF} can *statically* prove that any field access during execution is justified. To formalize and prove soundness, SVL_{IDF} has a reference dynamic semantics that tracks, for each stack frame, the set of fields that it has exclusive access to; deciding at runtime whether a formula holds depends on this information. Of course, thanks to soundness, an implementation of SVL_{IDF} needs no runtime tracking.

Recall that our approach to derive a gradual language includes the insertion of runtime checks, which requires that formulas *can* be evaluated at runtime. Therefore, the overhead of the reference semantics of SVL_{IDF} carries over to a naive implementation semantics. Additionally, it is no longer clear how accessibility is split between stack frames in case of a call: SVL_{IDF} transfers exclusive access to fields that are mentioned in the precondition of a procedure from the call site to the callee’s (fresh) stack frame. As we allow $?$ to also plausibly represent accessibility predicates, an imprecise precondition poses a challenge.

A valid strategy is to conservatively forward *all* accesses from caller to callee. As for GVL, we can devise an efficient implementation strategy for accessibility tracking that results in pay-as-you-go overhead. The fact that reducing the precision of contracts may now result in a divergence of program states (specifically, the accessible fields) asks for an adjustment of the dynamic part of the gradual guarantee, which originally requires lock-step reduction behavior. We carefully adjust the definition, preserving the core idea that reducing precision of a valid program does not alter the *observable* behavior of that program. The formalization of SVL_{IDF} and GVL_{IDF} are available in a companion

report available online, along with an interactive prototype implementation at <http://olydis.github.io/GradVer/impl/HTML5wp/>.

The prototype implementation of GVL_{IDF} displays the static and dynamic semantics of code snippets interactively, indicating the location of inserted runtime checks where necessary. It also displays the stack and heap at any point of execution. A number of predefined examples are available, along with an editable scratchpad. In particular, Example 2 demonstrates how imprecision enables safe reasoning about a recursive data structure that was not possible in SVL_{IDF} , because SVL_{IDF} lacks recursive predicates. This illustrates that, similarly to gradual types, imprecision can be used to augment the expressiveness of the static verification logic. In this case, the example does not even require a single runtime check.

6 Related Work

We have already related our work to the most-closely related research, including work on the underlying logics [8,4,15,20], the theory of gradual typing [18,17,19,7], closely related approaches to static [10] and dynamic [6,12,11,3] verification, as well as hybrid verification approaches [2,14]. Additional related work includes gradual type systems that include notions of *ownership* or *linearity*; one can think of the **acc** predicate as representing ownership of a piece of the heap, and **acc** predicates are treated linearly in the implicit dynamic frames methodology [20]. [21] developed a gradual typestate checking system, in which the state of objects is tracked in a linear type system. Similar to **acc** predicates, permissions to objects with state are passed linearly from one function to another, without being duplicated; if a strong permission is lost (*e.g.* due to a gradual specification), it can be regained with a runtime check, as long as no conflicting permission exist.

The gradual ownership approach of [16] is also related in that, like implicit dynamic frames, it aids developers in reasoning (gradually) about heap data structures. In that work, developers can specify containment relationships between objects, such that an *owned* object cannot be accessed from outside its owner. These access constraints can be checked either statically using a standard ownership type system, but if the developer leaves out ownership annotations from part of the program, dynamic checks are inserted.

Typestate and ownership are finite-state and topological properties, respectively, whereas in this work we explore gradual specification of logical contracts for the first time. Neither of these prior efforts benefited from the Abstracting Gradual Typing (AGT) methodology [7], which guided more principled design choices in our present work. Additionally, it is unclear whether the gradual guarantee of Siek *et al.* [19] holds in these proposals, which were developed prior to the formulation of the gradual guarantee.

One contrasting effort, which was also a stepping-stone to our current paper, is recent work on gradual refinement types [9]. In that approach, the AGT methodology is applied to a functional language in which types can be refined by

logical predicates drawn from a decidable logic. The present work is in a different context, namely first-order imperative programs as opposed to higher-order pure functional programs. This difference has a strong impact on the technical development. The present work is simpler in one respect, because formulas do not depend on their lexical context as in the gradual refinement setting. However, we had to reformulate gradual verification based on a weakest precondition calculus, whereas the prior work could simply extend the type system setting used when the AGT methodology was proposed. In addition, we provide a runtime semantics directly designed for the gradual verification setting, rather than adapting the evidence-tracking approach set forth by the AGT methodology and used for gradual refinement types. In fact, we investigated using the evidence-based approach for the runtime semantics of gradual verification, and found that it was semantically equivalent to what we present here but introduces unnecessary complexities. Overall, by adapting the AGT methodology to the verification setting, this work shows that the abstract interpretation approach to designing gradual languages can scale beyond type systems.

7 Conclusion

We have developed the formal foundations of gradual program verification, taking as starting point a simple program logic. This work is the first adaptation of recent fundamental techniques for gradual typing to the context of program verification. We have shown how to exploit the Abstracting Gradual Typing methodology [7] to systematically derive a gradual version of a weakest precondition calculus. Gradual verification provides a continuum between static and dynamic verification techniques, controlled by the (im)precision of program annotations; soundness is mediated through runtime checks.

Later, we briefly discuss how we applied our strategy to a more advanced logic using implicit dynamic frames (IDF) [20] in order to reason about mutable state. The use of IDF presents an additional challenge for obtaining a full pay-as-you-go model for gradual verification, because the footprint has to be tracked. We point to our prototype implementation which also references a formalization of gradualizing SVL_{IDF} . Further optimization of this runtime tracking is an interesting direction of future work. Another interesting challenge is to exercise our approach with other, richer program logics, as well as to study the gradualization of type systems that embed logical information, such as Hoare Type Theory [13], establishing a bridge between this work and gradual refinement types [9].

References

1. Arlt, S., Rubio-González, C., Rümmer, P., Schäf, M., Shankar, N.: The gradual verifier. In: *NASA Formal Methods Symposium*. pp. 313–327. Springer (2014)
2. Burdy, L., Cheon, Y., Cok, D.R., Ernst, M.D., Kiniry, J.R., Leavens, G.T., Leino, K.R.M., Poll, E.: An overview of jml tools and applications. *International Journal on Software Tools for Technology Transfer* 7(3), 212–232 (2005), <http://dx.doi.org/10.1007/s10009-004-0167-4>
3. Cheon, Y., Leavens, G.T.: A runtime assertion checker for the java modeling language (jml) (2002)
4. Dijkstra, E.W.: Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM* 18(8), 453–457 (Aug 1975), <http://doi.acm.org/10.1145/360933.360975>
5. Fahndrich, M., Barnett, M., Logozzo, F.: Embedded contract languages. In: *ACM SAC - OOPS*. Association for Computing Machinery, Inc. (March 2010), <https://www.microsoft.com/en-us/research/publication/embedded-contract-languages/>
6. Findler, R.B., Felleisen, M.: Contracts for higher-order functions. In: *Proceedings of the 7th ACM SIGPLAN Conference on Functional Programming (ICFP 2002)*. pp. 48–59. Pittsburgh, PA, USA (Sep 2002)
7. Garcia, R., Clark, A.M., Tanter, E.: Abstracting gradual typing. In: *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. pp. 429–442. POPL '16, ACM, New York, NY, USA (2016), <http://doi.acm.org/10.1145/2837614.2837670>
8. Hoare, C.A.R.: An axiomatic basis for computer programming. *Communications of the ACM* 12(10), 576–580 (1969)
9. Lehmann, N., Tanter, É.: Gradual refinement types. In: *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. pp. 775–788. Paris, France (Jan 2017)
10. Leino, K.R.M., Müller, P., Smans, J.: Verification of concurrent programs with chalice. In: *Foundations of Security Analysis and Design V*, pp. 195–222. Springer (2009)
11. Meyer, B.: Eiffel: A language and environment for software engineering. *Journal of Systems and Software* 8(3), 199–246 (1988)
12. Meyer, B.: *Object-Oriented Software Construction*. Prentice Hall (1988)
13. Nanevski, A., Morrisset, G., Birkedal, L.: Hoare type theory, polymorphism and separation. *Journal of Functional Programming* 5-6, 865–911 (2008)
14. Nguyen, H.H., Kuncak, V., Chin, W.N.: Runtime checking for separation logic. In: *International Workshop on Verification, Model Checking, and Abstract Interpretation*. pp. 203–217. Springer (2008)
15. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: *Logic in Computer Science, 2002. Proceedings. 17th Annual IEEE Symposium on*. pp. 55–74. IEEE (2002)
16. Sergey, I., Clarke, D.: Gradual ownership types. In: *Proceedings of the 21st European Conference on Programming Languages and Systems*. pp. 579–599. ESOP'12, Springer-Verlag, Berlin, Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-28869-2_29
17. Siek, J., Taha, W.: Gradual typing for objects. In: *European Conference on Object-Oriented Programming*. pp. 2–27. Springer (2007)

18. Siek, J.G., Taha, W.: Gradual typing for functional languages. In: Scheme and Functional Programming Workshop. vol. 6, pp. 81–92 (2006)
19. Siek, J.G., Vitousek, M.M., Cimini, M., Boyland, J.T.: Refined criteria for gradual typing. In: LIPIcs-Leibniz International Proceedings in Informatics. vol. 32. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2015)
20. Smans, J., Jacobs, B., Piessens, F.: Implicit dynamic frames: Combining dynamic frames and separation logic. In: European Conference on Object-Oriented Programming. pp. 148–172. Springer (2009)
21. Wolff, R., Garcia, R., Tanter, É., Aldrich, J.: Gradual typestate. In: European Conference on Object-Oriented Programming. pp. 459–483. Springer (2011)
22. Wright, A., Felleisen, M.: A syntactic approach to type soundness. *Inf. Comput.* 115(1), 38–94 (Nov 1994), <http://dx.doi.org/10.1006/inco.1994.1093>

Verification, Model Checking, and Abstract
Interpretation

19th International Conference, VMCAI 2018, Los
Angeles, CA, USA, January 7-9, 2018, Proceedings

Dillig, I.; Palsberg, J. (Eds.)

2018, XVIII, 540 p. 103 illus., Softcover

ISBN: 978-3-319-73720-1