

Omnisemantics: Smoother Handling of Nondeterminism

ARTHUR CHARGUÉRAUD, Inria & Université de Strasbourg, CNRS, ICube, France

ADAM CHLIPALA, MIT CSAIL, USA

ANDRES ERBSEN, MIT CSAIL, USA

SAMUEL GRUETTER, MIT CSAIL, USA

This paper gives an in-depth presentation of the omni-big-step and omni-small-step styles of semantic judgments. These styles describe operational semantics by relating starting states to sets of outcomes rather than to individual outcomes. A single derivation of these semantics for a particular starting state and program describes all possible nondeterministic executions (hence the name *omni*), whereas in traditional small-step and big-step semantics, each derivation only talks about one single execution. This restructuring allows for straightforward modeling of languages featuring both nondeterminism and undefined behavior. Specifically, omnisemantics inherently assert *safety*, i.e. they guarantee that none of the execution branches gets stuck, while traditional semantics need either a separate judgment or additional error markers to specify safety in the presence of nondeterminism.

Omnisemantics can be understood as an inductively defined weakest-precondition semantics (or more generally, predicate-transformer semantics) that does not involve invariants for loops and recursion, but instead uses unrolling rules like in traditional small-step and big-step semantics. Omnisemantics have already been used in the past, but we believe that it has been under-appreciated and that it deserves a well-motivated, extensive and pedagogical presentation of its benefits. We also explore several novel aspects associated with these semantics, in particular their use in type-soundness proofs for lambda calculi, partial-correctness reasoning, and forward proofs of compiler correctness for terminating but potentially nondeterministic programs being compiled to nondeterministic target languages. All results in this paper are formalized in Coq.

1 INTRODUCTION

Today, a typical project in rigorous reasoning about programming languages begins with an operational semantics (or maybe several), with proofs of key lemmas proceeding by induction on derivations of the semantics judgement. An extensive toolbox has been built up for formulating these relations, with common wisdom on the style to choose for each situation. With decades having passed since operational semantics became the standard technique in the 1980s, one might expect that the base of wisdom is sufficient. Yet, omnisemantics have emerged in the recent years as a new, powerful technique with numerous applications.

In short, omnisemantics relate starting states to their sets of possible outcomes, rather than to individual outcomes. The omni-big-step judgment takes the form $t/s \Downarrow Q$ and asserts that every possible evaluation starting from the configuration t/s reaches a final configuration that belongs to the set Q . This set Q is isomorphic to a postcondition from a Hoare triple. The omni-small-step judgment takes the form $t/s \longrightarrow P$. It asserts both that the configuration t/s can take one reduction step and that, for any step it might take, the resulting configuration belongs to the set P . On top of this judgment, one may define the *eventually* judgment $t/s \longrightarrow^\diamond P$, which asserts that every possible evaluation of t/s is safe and eventually reaches a configuration in the set P .

On the one hand, omnisemantics can be viewed as *operational semantics*, because they are not far from traditional operational semantics, and not far from executable interpreters. On the other hand, omnisemantics can be viewed as *axiomatic semantics*, because they are not far from reasoning rules; in particular, they directly give a practical, usable definition of a weakest-precondition judgment,

Authors' addresses: Arthur Charguéraud, Inria & Université de Strasbourg, CNRS, ICube, France, arthur.chargueraud@inria.fr; Adam Chlipala, MIT CSAIL, Cambridge, MA 02139, USA, adamc@csail.mit.edu; Andres Erbsen, MIT CSAIL, Cambridge, MA 02139, USA, andreser@mit.edu; Samuel Gruetter, MIT CSAIL, Cambridge, MA 02139, USA, gruetter@mit.edu.

which can be used for verifying concrete programs. The fact that they are both closely related to operational semantics and to axiomatic semantics is precisely the strength of omnisemantics.

To the best of our knowledge, the ideas of omnisemantics have been studied prior to the writing of this paper by three different groups of researchers. First, Schäfer et al. [2016] present an omni-big-step judgment for a nondeterministic source language of guarded commands, as well as for a deterministic target language with named continuations, using the term *axiomatic semantics* to refer to this style of semantics. They establish the correctness of a function that compiles terminating programs from the source language into the target language. Their proof is by induction on the derivation of an omni-big-step judgment for the *source* program rather than on a derivation for the *target* program, a key insight that we will discuss in Sections 1.3 and 6. They also present characterizations of program equivalence and present a proof of equivalence with traditional small-step semantics, though only in the case of a deterministic semantics. Second, Erbsen et al. [2021] make use of both omni-big-step semantics, applied to a high-level, core imperative language with external calls; and of omni-small-step semantics, applied to a low-level, RISC-V machine language. They call this style of semantics *CPS semantics*. They establish end-to-end compiler-correctness results for terminating programs. They also set up Separation Logic reasoning rules in weakest-precondition style. Third, Charguéraud [2020]’s course notes make use of omni-big-step semantics for the purpose of deriving Separation Logic triples, for both partial and total correctness. The language considered is a nondeterministic, imperative λ -calculus, with a substitution-based semantics. In particular, that work establishes the relationship between omni-big-step semantics and traditional big-step semantics, in the presence of nondeterminism.

Throughout the three pieces of work, the fundamental feature of omnisemantics being exploited is the ability to carry proofs by induction on derivations that follow the flow of program execution, *with smooth handling of nondeterminism*. Indeed, nondeterministic choices result in universally quantified induction hypotheses at steps where nondeterministic choices are made. Before further presenting omnisemantics, we believe that it is useful to begin by presenting in more detail the several important problems that omnisemantics solve.

1.1 Feature #1: Stuck Terms and Nondeterminism

In an impure language, an execution may get stuck, for instance due to a division by zero or an out-of-bounds array access. In a nondeterministic language, some executions may get stuck while others do not. Thus, for an impure, nondeterministic language, the existence of a traditional big-step derivation for a starting configuration is not a proof that getting stuck is impossible.

How to fix the problem? A popular but cumbersome approach is to add errors as explicit outcomes (written *err* in the rules below), so that we can state theorems ruling out stuck terms. For example, if the semantics of an impure functional language includes the rule BIG-LET, it needs to be augmented with two additional rules for propagating errors: BIG-LET-ERR-1 and BIG-LET-ERR-2.

$$\frac{t_1/s \Downarrow v_1/s' \quad ([v_1/x] t_2)/s' \Downarrow v/s''}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow v/s''} \text{BIG-LET}$$

$$\frac{t_1/s \Downarrow \text{err}}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow \text{err}} \text{BIG-LET-ERR-1} \quad \frac{t_1/s \Downarrow v_1/s' \quad ([v_1/x] t_2)/s' \Downarrow \text{err}}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow \text{err}} \text{BIG-LET-ERR-2}$$

The set of inference rules grows significantly, and the very type signature of the relation is complicated. Omni-big-step semantics provide a way to reason, in big-step style, about the absence of stuck terms in nondeterministic languages without introducing error-propagation rules.

1.2 Feature #2: Termination and Nondeterminism

In a nondeterministic language, a total-correctness Hoare triple, written $\text{total}\{H\} t \{Q\}$, asserts that in any state satisfying the precondition H , *any* execution of the term t terminates and reaches a final state satisfying the postcondition Q . In foundational approaches, Hoare triples must be defined in terms of or otherwise formally related to the operational semantics of languages.

When the (nondeterministic) semantics is expressed using the standard small-step relation, there are two classical approaches to defining total-correctness Hoare triples. The first one involves bounding the length of the execution. This approach not only involves tedious manipulation of integer bounds, but it is also restricted to finitely branching forms of nondeterminism. The second approach is to define total correctness as the conjunction of a partial-correctness property (if t terminates, then it satisfies the postcondition) and of a separate, inductively defined termination judgment. With both of these approaches, deriving reasoning rules for total-correctness Hoare triples becomes much more tedious than in the case of partial correctness.

One may hope for simpler proofs using a big-step judgment. Indeed, Hoare triples inherently have a big-step flavor. Moreover, for deterministic, sequential languages, the most direct way to derive reasoning rules for Hoare triples is from the big-step evaluation rules. Yet, when the semantics of a nondeterministic language is expressed using a traditional big-step judgment, we do not know of any direct way to capture the fact that *all* executions terminate. Omni-big-step semantics provide a direct definition of total-correctness Hoare triples with respect to a big-step-style, nondeterministic semantics, in a way that leads to simple proofs of the Hoare-logic rules.

1.3 Feature #3: Simulation Arguments with Nondeterminism and Undefined Behavior

Many compiler transformations map source programs to target programs that require more steps to accomplish the same work, because they must make do with lower-level primitives. Intuitively, we like to think of a compiler transformation being correct in terms of *forward simulation*: the transformation maps each step from the source program to a number of steps in the target program. Yet, in the context of a nondeterministic language, such a result is famously insufficient even in the special case of safely terminating programs. Concretely, compiler correctness requires showing *all* possible behaviors of the target program correspond to possible behaviors of the source program. A tempting approach is to establish a *backward simulation*, by showing that any step in the target program can be matched by some number of steps in the source program. The trouble is that all intermediate target-level states during a single source-level step need to be related to a source-level state, severely complicating the simulation relation.

To avoid that hassle, most compilation phases from CompCert [Leroy 2009] are carried out on *deterministic* intermediate languages, for which forward simulation implies backward simulation. Yet, many realistic languages (C included) are not naturally seen as deterministic. CompCert involves special effort to maintain determinism, through its celebrated memory model. Rather than revealing pointers as integers, CompCert semantics allocate pointers deterministically, taking care to trigger undefined behavior for any coding pattern that would be sensitive to the literal values of pointers. As a result, any compiler transformations that modify allocation order require the complex machinery of memory injections, to connect executions that use different deterministic pointer values. Omnisemantics make it possible to retain the simplicity of forward simulation, while keeping nondeterminism explicit.

1.4 Feature #4: Linear-size Type-Soundness Proofs

Type soundness asserts that if a closed term is well-typed, then none of its possible evaluations gets stuck. A type-soundness proof in the syntactic style [Wright and Felleisen 1994] reduces to a

pair of lemmas: preservation and progress.

$$\begin{array}{l} \text{PRESERVATION: } E \vdash t : T \quad \wedge \quad t \longrightarrow t' \quad \Rightarrow \quad E \vdash t' : T \\ \text{PROGRESS: } \quad \quad \quad \emptyset \vdash t : T \quad \Rightarrow \quad (\text{isvalue } t) \vee (\exists t'. t \longrightarrow t') \end{array}$$

The Wright and Felleisen approach, although widely used, suffers from two limitations that can be problematic at the scale of real-world languages with hundreds of syntactic constructs.

The first limitation is that this approach requires performing two inductions over the typing judgment. Nontrivial language constructs are associated with nontrivial statements of their induction hypotheses, for which the same manual work needs to be performed twice, once in the preservation proof and once in the progress proof. Factoring out the cases makes a huge difference in terms of proof effort and maintainability.

The second limitation is associated with the case inspection involved in the preservation proof. Concretely, for each possible rule that derives the typing judgment ($E \vdash t : T$), one needs to select the applicable rules that can derive the reduction rule ($t \longrightarrow t'$) for that same term t . Typically, only a few reduction rules are applicable. The trouble is that fully rigorous checking of the proof must still inspect all of those cases to confirm their irrelevance. A direct Coq proof, of the form “induction H1; inversion H2”, results in a proof term of size quadratic in the size of the language¹. As we expect to handle each possible transition at most once, a proof that takes only linear work would be more satisfying. It would also avoid potential blow-up in the proof-checking time, for languages involving hundreds of constructs.

Interestingly, in the particular case of a deterministic language, there exists a strategy [Rompf and Amin 2016] for deriving type soundness through a *single* inductive proof, which moreover avoids the quadratic case inspection. The key idea is to carry out an induction over the following statement: a well-typed term is either a value or can step to a term that admits the same type.

$$\emptyset \vdash t : T \quad \Rightarrow \quad (\text{isvalue } t) \vee (\exists t'. (t \longrightarrow t') \wedge (\emptyset \vdash t' : T))$$

Omnisemantics allow to generalize this approach to the case of nondeterministic languages. As we show in one of this paper’s original contributions, practical proofs of type soundness can be carried out with respect to both omni-small-step and omni-big-step semantics.

1.5 Contributions and Contents of the Paper

The contributions of this paper are as follows.

- We present bigstep and smallstep omnisemantics for a standard imperative λ -calculus as well as for a standard imperative while language, which we believe should make the presentation more accessible than in prior publications. Moreover, we accompany this presentation with a Coq formalization of all definitions and proofs.²
- We explain four key beneficial features of omnisemantics: They provide a convenient way to reason about the absence of stuck terms (feature #1) and the absence of diverging terms (feature #2) in nondeterministic languages, they enable forward-simulation-based correctness proofs for compilers with nondeterministic target languages (feature #3), and they enable type soundness proofs that avoid quadratic case inspection even in the case of a nondeterministic language (feature #4).

¹Lean matches Coq, and a proof based on Agda’s flexible dependent pattern matching still takes superlinear time to check.

²The present paper would, in particular, provide a formal publication of the results covered by the chapter on nondeterminism and the chapter on partial correctness from Charguéraud’s *Separation Logic Foundations* course, Volume 6 of the Software Foundations series. These results originally covered only omni-big-step semantics, but have been extended in 2021 to also cover omni-small-step semantics.

- We introduce the coinductive variant of omni-big-step semantics, which yields a partial-correctness judgment. This possibility was left as future work by Schäfer et al. [2016].
- We present numerous properties of omnisemantics, as well as their relationship to traditional operational semantics. Some of these properties were included in Erbsen et al. [2021], but extremely briefly. For example, the connection between traditional and omnisemantics only covered traditional small-step semantics with no undefined behavior, and small-step omnisemantics themselves were devoted one paragraph of space.
- We present in detail the proof techniques from two case studies on compiler-correctness results, adapted from Erbsen et al. [2021]’s prior work.
- We present a new case study illustrating an example of a correctness proof for a compiler transformation that *increases* the amount of nondeterminism. In contrast, work by Schäfer et al. [2016] and Erbsen et al. [2021] only considered transformations that *decrease* the amount of nondeterminism.

The paper is organized as follows.

- In Section 2, we introduce the omni-big-step judgment, which can be defined either inductively, to capture termination of all executions, or coinductively, in *partial-correctness* fashion. We also state and prove properties about the judgment, including the notion of smallest and largest admissible sets of outcomes.
- In Section 3, we introduce the omni-small-step judgment, as well as the *eventually* judgment defined on top of it, and three practical reasoning rules associated with these judgments.
- In Section 4, we present type-soundness proofs carried out with respect to either omni-small-step or omni-big-step semantics. We explain the improvement over the prior state of the art, as suggested in the earlier discussion of features #1 and #4.
- In Section 5, we explain how the omni-big-step judgment or the omni-small-step eventually judgment can be used to define Hoare triples and weakest-precondition predicates. We consider both partial and total correctness, and we show how the associated reasoning rules can be established via one-line proofs (recall feature #2). Moreover, we explain how to derive the frame rule from Separation Logic.
- In Section 6, we demonstrate how omnisemantics can be used to prove that a compiler correctly compiles terminating programs, via forward-simulation proofs (recall feature #3). We illustrate this possibility through two case studies carried out on a while-language. The first one, “heapification” of pairs, increases the amount of nondeterminism; it involves an omni-big-step semantics for both the source and the target language. The second one, introduction of stack allocation, decreases the amount of nondeterminism; it involves an omni-big-step semantics for the source language and an omni-small-step semantics for the target language.

Note that we leave it to future work to investigate how omnisemantics may be exploited to establish *full compiler correctness*, that is, not just the correctness of compilation for terminating programs but also that of programs that may crash, diverge, or perform infinitely many I/O interactions.

2 OMNI-BIG-STEP SEMANTICS

2.1 Definition of the Omni-Big-Step Judgment

Syntax. As a running example, we consider an imperative lambda-calculus, including a random-number generator `rand`. Both this operator and allocation are nondeterministic operations.

The grammar of the language appears next. The metavariable π ranges over primitive operations, v ranges over values, t ranges over terms, and x and f range over program variables. A value can be the unit value $\#$, a Boolean b , a natural number n , a pointer p , a primitive operator, or a closure³.

$$\begin{aligned} \pi &:= \text{add} \mid \text{rand} \mid \text{ref} \mid \text{free} \mid \text{get} \mid \text{set} \\ v &:= \# \mid b \mid n \mid p \mid \pi \mid \mu f. \lambda x. t \\ t &:= v \mid x \mid (t t) \mid \text{let } x = t \text{ in } t \mid \text{if } t \text{ then } t \text{ else } t \end{aligned}$$

For simplicity, we present evaluation rules by focusing first on programs in A-normal form. In an application $(t_1 t_2)$, the two terms must be either variables or values. Similarly, the condition of an if-statement must be either a variable or a value, and likewise for arguments of primitive operations. Restricting code to A-normal form is a common design choice for intermediate representations that appears in compilers and program-verification tools. We discuss in [Appendix B](#) how to state evaluation rules in omnisemantics-style for terms that are not in A-normal form.

Evaluation judgments. The standard big-step-semantics judgment for this language appears in [Figure 1](#). States s are partial maps from pointers p to values v . The evaluation judgment $t/s \Downarrow v/s'$ asserts that the configuration t/s , made of a term t and an initial state s , may evaluate to the final configuration v/s' , made of a value v and a final state s' .

The corresponding omni-big-step semantics appears in [Figure 2](#). Its evaluation judgment, written $t/s \Downarrow Q$, asserts that all possible evaluations starting from the configuration t/s reach final configurations that belong to the set Q . Observe how the standard big-step judgment $t/s \Downarrow v/s'$ describes the behavior of one possible execution of t/s , whereas the omni-big-step judgment describes the behavior of all possible executions of t/s . The set Q that appears in $t/s \Downarrow Q$ corresponds to an overapproximation of the set of final configurations: it may contain configurations that are not actually reachable by executing t/s . We return to that aspect later on, at the end of [§2.2](#).

The set Q contains pairs made of values and states. Such a set can be described equivalently by a predicate of type “ $\text{val} \rightarrow \text{state} \rightarrow \text{Prop}$ ” or by a predicate of type “ $(\text{val} \times \text{state}) \rightarrow \text{Prop}$ ”. In this paper, in order to present definitions in the most idiomatic style, we use set-theoretic notation such as $(v, s) \in Q$ for stating semantics and typing rules, and we use the logic-oriented notation $Q v s$ when discussing program logics. (The type of Q may be generalized for languages that include exceptions; see [Appendix D](#).)

Description of the evaluation rules. The base case is the rule OMNI-BIG-VAL: a final configuration v/s satisfies the postcondition Q if this configuration belongs to the set Q .

The let-binding rule OMNI-BIG-LET ensures that all possible evaluations of an expression $\text{let } x = t_1 \text{ in } t_2$ in state s terminate and satisfy the postcondition Q . First of all, we need all possible evaluations of t_1 to terminate. Let Q_1 denote (an overapproximation of) the set of results that t_1 may reach, as captured by the first premise $t_1/s \Downarrow Q_1$. One can think of Q_1 as the type of t_1 , in a very precise type system where any set of values can be treated as a type. The second premise asserts that, for any configuration v'/s' in that set Q_1 , we need all possible evaluations of the term $[v'/x] t_2$ in state s' to satisfy the postcondition Q .

The evaluation rule OMNI-BIG-ADD for an addition operation is almost like that of a value: it asserts that the evaluation of $\text{add } n_1 n_2$ in state s satisfies the postcondition Q if the pair $((n_1 + n_2), s)$ belongs to the set Q . The nondeterministic rule OMNI-BIG-RAND is more interesting. The term $\text{rand } n$ evaluates safely only if $n > 0$. Under this assumption, its result, named m in the rule, may be any

³ In our Coq formalization, the grammar of values is restricted to *closed* values (i.e., values without free variables). This design choice significantly simplifies the reasoning about substitutions. One minor consequence is that the function construct needs to appear twice: once in the grammar of closed values and once in the grammar of terms.

$$\begin{array}{c}
\text{BIG-APP} \\
\frac{v_1 = (\mu f. \lambda x. t) \quad ([v_2/x] [v_1/f] t)/s \Downarrow v'/s'}{(v_1 v_2)/s \Downarrow v'/s'} \\
\text{BIG-IF-TRUE} \\
\frac{t_1/s \Downarrow v'/s'}{(\text{if true then } t_1 \text{ else } t_2)/s \Downarrow v'/s'} \\
\text{BIG-IF-FALSE} \\
\frac{t_2/s \Downarrow v'/s'}{(\text{if false then } t_1 \text{ else } t_2)/s \Downarrow v'/s'} \\
\text{BIG-LET} \\
\frac{t_1/s \Downarrow v_1/s' \quad ([v_1/x] t_2)/s' \Downarrow v/s''}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow v/s''} \\
\text{BIG-ADD} \\
\frac{}{(\text{add } n_1 \ n_2)/s \Downarrow (n_1 + n_2)/s} \\
\text{BIG-RAND} \\
\frac{0 \leq m < n}{(\text{rand } n)/s \Downarrow m/s} \\
\text{BIG-REF} \\
\frac{p \notin \text{dom } s}{(\text{ref } v)/s \Downarrow p/(s[p := v])} \\
\text{BIG-FREE} \\
\frac{p \in \text{dom } s}{(\text{free } p)/s \Downarrow \#/(s \setminus p)} \\
\text{BIG-GET} \\
\frac{p \in \text{dom } s}{(\text{get } p)/s \Downarrow (s[p])/s} \\
\text{BIG-SET} \\
\frac{p \in \text{dom } s}{(\text{set } p \ v)/s \Downarrow \#/(s[p := v])}
\end{array}$$

Fig. 1. Standard big-step semantics (for terms in A-normal form)

$$\begin{array}{c}
\text{OMNI-BIG-VAL} \\
\frac{(v, s) \in Q}{v/s \Downarrow Q} \\
\text{OMNI-BIG-IF-TRUE} \\
\frac{t_1/s \Downarrow Q}{(\text{if true then } t_1 \text{ else } t_2)/s \Downarrow Q} \\
\text{OMNI-BIG-IF-FALSE} \\
\frac{t_2/s \Downarrow Q}{(\text{if false then } t_1 \text{ else } t_2)/s \Downarrow Q} \\
\text{OMNI-BIG-APP} \\
\frac{v_1 = \mu f. \lambda x. t_1 \quad ([v_1/f] [v_2/x] t_1)/s \Downarrow Q}{(v_1 v_2)/s \Downarrow Q} \\
\text{OMNI-BIG-LET} \\
\frac{t_1/s \Downarrow Q_1 \quad (\forall (v', s') \in Q_1. ([v'/x] t_2)/s' \Downarrow Q)}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow Q} \\
\text{OMNI-BIG-ADD} \\
\frac{(n_1 + n_2, s) \in Q}{(\text{add } n_1 \ n_2)/s \Downarrow Q} \\
\text{OMNI-BIG-RAND} \\
\frac{n > 0 \quad (\forall m. 0 \leq m < n \Rightarrow (m, s) \in Q)}{(\text{rand } n)/s \Downarrow Q} \\
\text{OMNI-BIG-REF} \\
\frac{\forall p \notin \text{dom } s. (p, s[p := v]) \in Q}{(\text{ref } v)/s \Downarrow Q} \\
\text{OMNI-BIG-FREE} \\
\frac{p \in \text{dom } s \quad (\#, s \setminus p) \in Q}{(\text{free } p)/s \Downarrow Q} \\
\text{OMNI-BIG-GET} \\
\frac{p \in \text{dom } s \quad (s[p], s) \in Q}{(\text{get } p)/s \Downarrow Q} \\
\text{OMNI-BIG-SET} \\
\frac{p \in \text{dom } s \quad (\#, s[p := v]) \in Q}{(\text{set } p \ v)/s \Downarrow Q}
\end{array}$$

Fig. 2. Omni-big-step semantics (for terms in A-normal form)

integer in the range $[0, n)$. Thus, to guarantee that every possible evaluation of $\text{rand } n$ in a state s produces a result satisfying the postcondition Q , it must be the case that every pair of the form (m, s) with $m \in [0, n)$ belongs to the set Q .

The evaluation rule **OMNI-BIG-REF**, which describes allocation at a nondeterministically chosen, fresh memory address, follows a similar pattern. For every possible new address p , the pair made of p and the extended state $s[p := v]$ needs to belong to the set Q . The remaining rules, **OMNI-BIG-FREE**, **OMNI-BIG-GET** and **OMNI-BIG-SET**, are deterministic and follow the same pattern as **OMNI-BIG-ADD**, only with a side condition $p \in \text{dom } s$ to ensure that the address being manipulated does belong to the domain of the current state.

2.2 Properties of the Omni-Big-Step Judgment

In this section, we discuss some key properties of the omni-big-step judgment $t/s \Downarrow Q$. Recall that the metavariable Q denotes an overapproximation of the set of possible final configurations.

Total correctness. The predicate $t/s \Downarrow Q$ captures total correctness in the sense that it captures the conjunction of termination (all executions terminate) and partial correctness (if an execution terminates, then its final state satisfies the postcondition Q). Formally, let $t/s \Downarrow v/s'$ denote the standard big-step evaluation judgment, and let $\text{terminates}(t, s)$ be a predicate that captures the fact that all executions of t/s terminate (a formal definition is given in [Appendix F](#)). We prove:

OMNI-BIG-STEP-IFF-TERMINATES-AND-CORRECT :

$$t/s \Downarrow Q \iff \text{terminates}(t, s) \wedge \left(\forall v/s'. (t/s \Downarrow v/s') \implies (v, s') \in Q \right).$$

In particular, if we instantiate the postcondition Q with the *always-true* predicate, we obtain the predicate $t/s \Downarrow \{(v, s') \mid \text{True}\}$, which captures only the termination property.

Consequence rule. The judgment $t/s \Downarrow Q$ still holds when the postcondition Q is replaced with a larger set. In other words, the postcondition can always be weakened, like in Hoare logic.

$$\text{OMNI-BIG-CONSEQUENCE} : \quad t/s \Downarrow Q \wedge Q \subseteq Q' \implies t/s \Downarrow Q'$$

Strongest postcondition. If the omni-big-step judgment holds for at least one set, then there exists a smallest possible set Q for which $t/s \Downarrow Q$ holds. This set corresponds to the strongest possible postcondition Q , in the terminology of Hoare logic. Formally, if $t/s \Downarrow Q$ holds for at least one Q , then $t/s \Downarrow (\text{strongest-post } t/s)$ holds, where the strongest postcondition is equal to the intersection of all valid postconditions.

$$\text{strongest-post } t/s = \bigcap_{Q \mid (t/s \Downarrow Q)} Q = \{(v, s') \mid \forall Q, (t/s \Downarrow Q) \implies (v, s') \in Q\}$$

Outcome set is nonempty. Observe that the judgment $t/s \Downarrow Q$, as defined in [Fig. 2](#), can only hold for a nonempty set Q . This property is a general requirement for an omni-big-step semantics to make sense, and when designing omni-big-step rules for a new language, one has to be careful to make sure it holds. For example, omitting the premise $n > 0$ in the rule [OMNI-BIG-RAND](#) would lead to an ill-formed semantics, allowing $\text{rand } 0$ to be related to any set Q , including the empty set.

No derivations for terms that may get stuck. The fact that $\text{rand } 0$ is a stuck term is captured by the fact that $(\text{rand } 0)/s \Downarrow Q$ does not hold for any Q . More generally, if one or more nondeterministic executions of t may get stuck, then we have: $\forall Q. \neg (t/s \Downarrow Q)$.

Relationship to standard big-step semantics. The standard big-step judgment $t/s \Downarrow v/s'$ relates one input configuration t/s to one single result configuration v/s' . The omni-big-step judgment, which relates inputs to sets of results, thus appears as an immediate generalization of the standard big-step judgment. The following two results formalizes their relationship.

First, if $t/s \Downarrow Q$ holds, then any final configuration for which the standard big-step judgment holds necessarily belongs to the set Q .

$$\text{OMNI-BIG-AND-BIG-INV} : \quad t/s \Downarrow Q \wedge t/s \Downarrow v/s' \implies (v, s') \in Q$$

Second, if $t/s \Downarrow Q$ holds, then there exists at least one evaluation according to the standard big-step judgment whose final configuration belongs to the set Q .

$$\text{OMNI-BIG-TO-ONE-BIG} : \quad t/s \Downarrow Q \implies \exists v/s'. t/s \Downarrow v/s' \wedge (v, s') \in Q$$

$\frac{\text{PRECISE-BIG-VAL}}{v/s \Downarrow' \{(v, s)\}}$	$\frac{\text{PRECISE-BIG-REF}}{(\text{ref } v)/s \Downarrow' \{(p, s[p := v]) \mid p \notin \text{dom } s\}}$
$\frac{\text{PRECISE-BIG-RAND}}{n > 0}}{(\text{rand } n)/s \Downarrow' \{(m, s) \mid 0 \leq m < n\}}$	$\frac{\text{PRECISE-BIG-LET} \quad t_1/s \Downarrow' Q_1 \quad \forall (v', s') \in Q_1. ([v'/x] t_2)/s' \Downarrow Q'_{(v', s')}}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow' \bigcup_{(v', s') \in Q_1} Q'_{(v', s')}}}$

Fig. 3. Selected rules defining a *precise* variant of omni-big-step semantics, written $t/s \Downarrow' Q$.

A corollary asserts that if $t/s \Downarrow Q$ holds with Q being a singleton set made of a unique final configuration v/s' , then the standard big-step judgment holds for that configuration.

$$\text{OMNI-BIG-SINGLETON:} \quad t/s \Downarrow \{(v, s')\} \Rightarrow t/s \Downarrow v/s'$$

Particular case of deterministic languages. In a deterministic language, an input configuration t/s may evaluate to at most one configuration v/s' . In such case, the strongest postcondition is reduced to the singleton set $\{(v, s')\}$.

2.3 About the Overapproximation of the Set of Results

The omni-big-step judgment $t/s \Downarrow Q$ associates an initial configuration t/s with a postcondition Q , which denotes an *overapproximation* of the set of possible final configurations. One may thus wonder: why not associate it with a *precise* set of results? In this section, we show that it is technically possible to define a *precise* judgment, but at the same time argue why that judgment is much less practical to work with than the *overapproximating* omni-big-step judgment.

The precise judgment, written $t/s \Downarrow' Q$, is precise in the sense that it relates a configuration t/s to *at most one* set of results Q . This precise judgment, like the overapproximating omni-big-step judgment, guarantees safety: a judgment $t/s \Downarrow' Q$ can be derived for some Q if and only if none of the possible executions of t/s can get stuck. Thus, the precise judgment relates a *safe* configuration t/s to exactly one Q .

Figure 3 shows selected rules from the definition of the precise judgment, written $t/s \Downarrow' Q$. The rule PRECISE-BIG-VAL relates a value v in a state s to the singleton set made of the pair (v, s) . The rule PRECISE-BIG-REF relates the term $(\text{ref } v)$ in a state s to the set of pairs made of a location p fresh from s and of the state s updated at location p with the value v . Observe how this compares with the rule OMNI-BIG-REF, which only requires that set of pairs to be included in the result set Q . The rule PRECISE-BIG-RAND follows a similar pattern, only with the premise $n > 0$ to ensure that the term is not stuck.

Most interesting is the rule PRECISE-BIG-LET. Its first premise involves an intermediate set Q_1 , which denotes *exactly* the set of results that t_1 can produce when executed in the input state s . The second premise describes, for each result (v', s') from the set Q_1 , the evaluation of $([v'/x] t_2)$ in state s' . The result of the execution is asserted to be exactly a set of configurations written $Q'_{(v', s')}$. Here Q' denotes a (possibly infinite) family of postconditions, indexed by the possible results of t_1 . The final postcondition of the term $(\text{let } x = t_1 \text{ in } t_2)$ is obtained by taking the union over that family of postconditions.⁴

In practice, working with indexed families of postconditions introduces significant overhead, compared with the overapproximating omni-big-step judgment. Moreover, for practical applications

⁴In Coq, we model sets with elements of type A as functions from A to propositions, thus Q_1 is represented as a function that takes a value and a state and returns a proposition, Q' is a function that takes a value, a state, another value, another state and returns a proposition, and the union over the family of results is written $\lambda v'' s''. \exists v' s'. Q_1 v' s' \wedge Q' v' s' v'' s''$.

such as type-checking or program verification (either using weakest precondition or Hoare triples), we are only interested in overapproximations of the semantics. For such applications, building the overapproximation on top of a precise judgment would only introduce a level of indirection. For other situations where a notion of exact set of results might be desirable, typically for metatheoretical results (e.g., completeness results), we can always refer to the *strongest postcondition*, which, as explained earlier, can be formalized as the intersection of all valid postconditions.

In summary, we believe that it is interesting to know that a precise judgment can be defined, as it might be useful in other contexts, but for the applications that we have in mind the overapproximating omni-big-step judgment appears much better suited.

2.4 Coinductive Interpretation of the Omni-Big-Step Judgment

Let $t/s \Downarrow^{\text{co}} Q$ denote the judgment defined by the coinductive interpretation of the same set of rules as for the inductively defined judgment $t/s \Downarrow Q$, i.e., rules from Fig. 2. The coinductive interpretation allows for infinite derivation trees, thus the coinductive omni-big-step judgment can be used to capture properties of nonterminating executions.

More precisely, the judgment $t/s \Downarrow^{\text{co}} Q$ asserts that every possible execution of configuration t/s either diverges or terminates on a final configuration satisfying Q . In particular, this judgment rules out the possibility for an execution of t/s to get stuck, and it can be used to express type soundness, as detailed in §4. The judgment $t/s \Downarrow^{\text{co}} Q$ can also be used to define partial-correctness Hoare triples, as detailed in §5.

Formally, we can relate the meaning of $t/s \Downarrow^{\text{co}} Q$ to the small-step characterization of partial correctness as follows: for every execution prefix, the configuration reached is either a value satisfying the postcondition, or it is a term that can be reduced further. Below, $t/s \longrightarrow t'/s'$ denotes the standard small-step evaluation judgment (defined in Appendix I), and val denotes the constructor that injects values into the grammar of terms.

CO-OMNI-BIG-IFF-SAFE-AND-CORRECT

$$t/s \Downarrow^{\text{co}} Q \iff \forall s't'. (t/s \longrightarrow^* t'/s') \Rightarrow \left(\exists v. t' = \text{val } v \wedge (v, s') \in Q \right) \vee \left(\exists t''s''. t'/s' \longrightarrow t''/s'' \right)$$

The judgment $t/s \Downarrow^{\text{co}} Q$ can also be used to characterize divergence, by instantiating Q as the empty set: the predicate $t/s \Downarrow^{\text{co}} \emptyset$ asserts that every possible execution of t/s diverges. Because the judgment $t/s \Downarrow^{\text{co}} Q$ is covariant in Q , the predicate $t/s \Downarrow^{\text{co}} \emptyset$ holds if and only if the predicate $t/s \Downarrow^{\text{co}} Q$ holds for any Q . In summary, we formally characterize divergence as follows.

$$\text{diverges } t s \equiv (t/s \Downarrow^{\text{co}} \emptyset) \quad \text{diverges } t s \iff \forall Q. (t/s \Downarrow^{\text{co}} Q)$$

3 OMNI-SMALL-STEP SEMANTICS

In this section, we present the omni-small-step judgment, written $t/s \longrightarrow P$. We then present the *eventually* judgment, written $t/s \longrightarrow^\diamond P$. We use these judgments for establishing type-soundness (§4.1) and compiler-verification (§6.6) results.

3.1 The Omni-Small-Step Judgment

The omni-small-step judgment, written $t/s \longrightarrow P$, asserts that the configuration t/s can take one reduction step and that, for any step it might take, the resulting configuration belongs to the set P . It is defined by the rules from Fig. 4. There is one per small-step transition. The interesting rules are those involving nondeterminism, namely OMNI-SMALL-RAND and OMNI-SMALL-REF, which follow a

$$\begin{array}{c}
\text{OMNI-SMALL-APP} \\
\frac{v_1 = (\mu f. \lambda x. t) \quad ([v_2/x] [v_1/f] t, s) \in P}{(v_1 v_2)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-LET-CTX} \\
\frac{t_1/s \longrightarrow P_1 \quad (\forall (t'_1, s') \in P_1. ((\text{let } x = t'_1 \text{ in } t_2), s') \in P)}{(\text{let } x = t_1 \text{ in } t_2)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-LET} \\
\frac{([v_1/x] t_2, s) \in P}{(\text{let } x = v_1 \text{ in } t_2)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-RAND} \\
\frac{n > 0 \quad (\forall m \in [0, n]. (m, s) \in P)}{(\text{rand } n)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-ADD} \\
\frac{(n_1 + n_2, s) \in P}{(\text{add } n_1 \ n_2)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-REF} \\
\frac{(\forall p \notin \text{dom } s. (p, s[p := v]) \in P)}{(\text{ref } v)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-FREE} \\
\frac{p \in \text{dom } s \quad (\#t, s \setminus p) \in P}{(\text{free } p)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-GET} \\
\frac{p \in \text{dom } s \quad (s[p], s) \in P}{(\text{get } p)/s \longrightarrow P} \\
\\
\text{OMNI-SMALL-SET} \\
\frac{p \in \text{dom } s \quad (\#t, s[p := v]) \in P}{(\text{set } p \ v)/s \longrightarrow P}
\end{array}$$

Fig. 4. Omni-small-step semantics (for terms in A-normal form)

pattern similar to the corresponding omni-big-step rules. Observe also how the rule OMNI-SMALL-LET-CTX handles the case of a reduction that takes place in the evaluation context of a let-binding, by quantifying over an intermediate set of results named P_1 .

We prove that the judgment $t/s \longrightarrow P$ captures the expected property w.r.t. the standard small-step judgment: the configuration t/s can make a step, and for every step it might take, it reaches a configuration in P .

OMNI-SMALL-STEP-IFF-PROGRESS-AND-CORRECT

$$t/s \longrightarrow P \iff (\exists t's'. t/s \longrightarrow t'/s') \wedge (\forall t's'. t/s \longrightarrow t'/s' \Rightarrow (t', s') \in P)$$

3.2 The “Eventually” Judgment

The judgment $t/s \longrightarrow^\diamond P$ captures the property that every possible evaluation of t/s is safe and eventually reaches a configuration in the set P . Here, P denotes a set of configurations—it is not limited to being a set of *final* configurations like in the previous section. The judgment $t/s \longrightarrow^\diamond P$ is defined inductively by the following two rules. The first one asserts that the judgment is satisfied if t/s belongs to P . The second one asserts that the judgment is satisfied if t/s is not stuck and that for any configuration t'/s' that it may reduce to, the predicate $t'/s' \longrightarrow^\diamond P$ holds. The latter property is expressed using the omni-small-step judgment $t/s \longrightarrow P'$, where P' denotes an overapproximation of the set of configurations t'/s' to which t/s may reduce.

$$\begin{array}{c}
\text{EVENTUALLY-HERE} \\
\frac{(t, s) \in P}{t/s \longrightarrow^\diamond P} \\
\\
\text{EVENTUALLY-STEP} \\
\frac{t/s \longrightarrow P' \quad (\forall (t', s') \in P'. t'/s' \longrightarrow^\diamond P)}{t/s \longrightarrow^\diamond P}
\end{array}$$

If Q denotes a set of *final* configurations, then the judgment $t/s \longrightarrow^\diamond Q$ can be viewed as a particular case of the judgment $t/s \longrightarrow^\diamond P$, where P denotes a set of configurations. We prove that $t/s \longrightarrow^\diamond Q$ matches our omni-big-step judgment $t/s \Downarrow Q$.

$$\text{EVENTUALLY-IFF-OMNI-BIG-STEP:} \quad t/s \longrightarrow^\diamond Q \iff t/s \Downarrow Q$$

3.3 Chained Rule and Cut Rule for the “Eventually” Judgment

To apply the rule `EVENTUALLY-STEP`, one needs to provide upfront an intermediate postcondition P' . Doing so is not always convenient. It turns out that we can leverage the omni-small-step judgment $t/s \longrightarrow P'$ to provide an introduction rule for $t/s \longrightarrow^\diamond P$ that does not require providing P' upfront. This rule, which we call the *chained* version of `EVENTUALLY-STEP`, admits the statement shown below. It reads as follows: if every possible step of t/s reduces in one step to a configuration that eventually reaches a configuration from the set P , then every possible evaluation of t/s eventually reaches a configuration from the set P .

$$\text{EVENTUALLY-STEP-CHAINED} : \quad t/s \longrightarrow \{(t', s') \mid t'/s' \longrightarrow^\diamond P\} \Rightarrow t/s \longrightarrow^\diamond P$$

Another interesting property of the judgment $t/s \longrightarrow^\diamond P$ is its cut rule, which is derivable. It asserts the following: if every possible evaluation of t/s eventually reaches a configuration in the set P' , and if every configuration from the set P' eventually reaches a configuration from the set P , then every possible evaluation of t/s eventually reaches a configuration from the set P .

$$\text{EVENTUALLY-CUT} : \quad t/s \longrightarrow^\diamond P' \wedge (\forall (t', s') \in P'. t'/s' \longrightarrow^\diamond P) \Rightarrow t/s \longrightarrow^\diamond P$$

This cut rule also admits a *chained* version, which reads as follows: if every possible evaluation of t/s eventually reaches a configuration that itself eventually reaches a configuration from the set P , then every possible evaluation of t/s eventually reaches a configuration from the set P .

$$\text{EVENTUALLY-CUT-CHAINED} : \quad t/s \longrightarrow^\diamond \{(t', s') \mid t'/s' \longrightarrow^\diamond P\} \Rightarrow t/s \longrightarrow^\diamond P$$

The cut rule and the chained rules are particularly handy to work with, as we illustrate in §6.6.

3.4 Coinductive Interpretation of the Omni-Small-Step Judgment

Let $t/s \longrightarrow_{\text{co}}^\diamond P$ denote the coinductive interpretation of the two rules that define $t/s \longrightarrow^\diamond P$. We can relate the coinductive judgment $t/s \longrightarrow_{\text{co}}^\diamond P$ with the coinductive omni-big-step judgment from §2.4. Here again, we view a set Q of final configurations as a set of configurations.

$$t/s \longrightarrow_{\text{co}}^\diamond Q \iff t/s \Downarrow^{\text{co}} Q$$

Divergence can be captured by instantiating P as the empty set. We prove that the judgment $t/s \longrightarrow_{\text{co}}^\diamond \emptyset$ is equivalent to the standard small-step characterization of divergence, which asserts that any execution prefix may be extended with at least one additional step.

$$t/s \longrightarrow_{\text{co}}^\diamond \emptyset \iff \forall s' t'. (t/s \longrightarrow^* t'/s') \Rightarrow (\exists t'' s''. t'/s' \longrightarrow t''/s'')$$

4 TYPE-SOUNDNESS PROOFS USING OMNISEMANTICS

In this section, we show how the omni-small-step and omni-big-step judgments may be used to carry out type-soundness proofs. We illustrate the proof structures using simple types (STLC). As a warm-up, we begin with a presentation of type soundness on the restriction to the state-free fragment of our running-example language.

For this section, we need to consider a different semantics for the random-number generator. Indeed, the current rule `OMNI-BIG-RAND` asserts that the program is stuck if `rand n` is invoked with an argument $n \leq 0$. Since here we are interested in proving that well-typed programs do not get stuck, let us consider a modified semantics, where `rand n` is turned into a total function that returns 0 when $n \leq 0$.

$$\frac{\text{OMNI-BIG-RAND-COMPLETE} \quad \forall m. 0 \leq m < \max(n, 1) \Rightarrow (m, s) \in Q}{(\text{rand } n)/s \Downarrow Q}$$

Additionally, for this section, we also exclude the primitive operation `free`, which is not type-safe. The grammar of types, written T , appears below.

$$T ::= \text{unit} \mid \text{bool} \mid \text{int} \mid T \rightarrow T \mid \text{ref } T$$

A typing environment, written E , maps variable names to types. The judgment $\vdash v : T$ asserts that the closed value v admits the type T . The judgment $E \vdash t : T$ asserts that the term t admits type T in the environment E . We let \mathbb{V} denote the set of terms that are either values or variables—recall that we consider A-normal forms to simplify the presentation. The typing rules (given in appendix G) are essentially standard, apart from the fact that they involve side conditions of the form $t \in \mathbb{V}$ to constrain terms to be in A-normal form.

4.1 Omni-Small-Step Type-Soundness Proof for a State-Free Language

A *stuck term* is a term that is not a value and that cannot take a step. Type soundness asserts that if a closed term t is well-typed, then none of its possible evaluations gets stuck. In other words, if t reduces in a number of steps to t' , then t' is either a value or can further reduce.

TYPE-SOUNDNESS (STATE-FREE LANGUAGE):

$$(\emptyset \vdash t : T) \wedge (t \longrightarrow^* t') \Rightarrow (\text{isvalue } t') \vee (\exists t''. t' \longrightarrow t'')$$

The traditional approach to establishing type soundness is by proving the *preservation* and *progress* properties [Pierce 2002; Wright and Felleisen 1994].

PRESERVATION (STATE-FREE LANGUAGE):

$$E \vdash t : T \wedge t \longrightarrow t' \Rightarrow E \vdash t' : T$$

PROGRESS (STATE-FREE LANGUAGE):

$$\emptyset \vdash t : T \Rightarrow (\text{isvalue } t) \vee (\exists t'. t \longrightarrow t')$$

Each of these proofs is most typically carried out by induction on the typing judgment. One difficulty that might arise in the type-preservation proof for a large language with dozens (if not hundreds) of typing rules is the fact that one needs, for each case of the typing judgment $E \vdash t : T$, to inspect all the potential cases of the reduction judgment $t \longrightarrow t'$. This inspection is not really quadratic in practice, because one can filter out applicable rules based on the shape of the term t . Nevertheless, a typical Coq proof performing “intros HT HR; induction HT; inversion HR” does produce a proof term whose size is quadratic in the number of term constructs. Coq users have experienced performance challenges with quadratic-complexity proof terms when formalizing PL metatheory [Monin and Shi 2013].

Interestingly, in the particular case of a deterministic language, there exists a known strategy (e.g., of Rompf and Amin [2016]) to reformulate the preservation and progress statements in a way that not only factors out the two into a single statement but also can be proved with a linear-size proof term. This combined statement, shown below, asserts that a well-typed term t is either a value or can make a step towards a term t' that admits the same type.

INDUCTION-FOR-TYPE-SOUNDNESS (STATE-FREE, STANDARD SMALL-STEP, DETERMINISTIC CASE)

$$\emptyset \vdash t : T \Rightarrow (\text{isvalue } t) \vee (\exists t'. (t \longrightarrow t') \wedge (\emptyset \vdash t' : T))$$

As we explain next, this approach can be generalized to the case of nondeterministic languages using the omni-small-step judgment. Let us write $t \longrightarrow P$ for the judgment that corresponds to $t/s \longrightarrow P$ without the state argument. We can state type soundness by considering for the postcondition P the set of terms t' that admit the same type as t .

INDUCTION-FOR-TYPE-SOUNDNESS (STATE-FREE, OMNI-SMALL-STEP, GENERAL CASE)

$$\emptyset \vdash t : T \Rightarrow (\text{isvalue } t) \vee (t \longrightarrow \{t' \mid (\emptyset \vdash t' : T)\})$$

The proof is carried out by induction on the typing judgment. For the case where t is a value, the left part of the disjunction applies. For all other cases, the right part needs to be established.

For each term construct, we invoke the relevant constructor from the definition of the judgment $t \longrightarrow P$. The application of such a rule typically results in proof obligations of the form $\emptyset \vdash t' : T$, where t' can be any of the terms to which t reduces. In case of a nondeterministic evaluation rule, these possible terms t' are quantified universally in the premise of the omni-small-step rule, thus they become universally quantified in the proof obligation that arises for the user. In other words, we are required to establish $\emptyset \vdash t' : T$ for each possible t' , as expected.

The statement `INDUCTION-FOR-TYPE-SOUNDNESS` above entails the preservation property (for empty environments) and the progress property. We prove once-and-for-all that the statement of `INDUCTION-FOR-TYPE-SOUNDNESS` entails the `TYPE-SOUNDNESS` property.⁵

4.2 Omni-Small-Step Type-Soundness Proof for an Imperative Language

Let us now generalize the results from the previous section to account for memory operations.

A store-typing environment, written S , is a map from locations to types. The typing judgment for values is extended with a store-typing environment, taking the form $S \vdash v : T$. Likewise, the typing judgment for terms is extended to the form $S; E \vdash t : T$. The store-typing entity S only plays a role in the typing rule for memory locations. The rules for typing memory locations and memory operations are standard; they appear in [Appendix H](#).

The type-soundness property asserts that the execution of any well-typed term, starting from the empty state, does not get stuck. In the statement below, \emptyset denotes an empty state or an empty store typing, whereas \emptyset denotes, as before, the empty typing context.

`TYPE-SOUNDNESS:`

$$(\emptyset; \emptyset \vdash t : T) \wedge (t/\emptyset \longrightarrow^* t'/s') \Rightarrow (\text{isvalue } t') \vee (\exists t''s''. t'/s' \longrightarrow t''/s'')$$

To establish a type-soundness result by induction on a reduction sequence, one needs to introduce a typing judgment for stores. A store s admits type S , written $\vdash s : S$, if and only if s and S have the same domain and, for any location p in the domain, $s[p]$ admits the type $S[p]$. Formally:

$$\vdash s : S \quad \equiv \quad (\text{dom } s = \text{dom } S) \wedge (\forall p \in \text{dom } s. S; \emptyset \vdash s[p] : S[p])$$

The preservation and progress lemmas associated with the traditional approach to proving type soundness are updated as shown below. In particular, the preservation lemma requires the output state to admit a type that extends the store typing associated with the input state ($S' \supseteq S$).

`PRESERVATION:`

$$S; \emptyset \vdash t : T \wedge \vdash s : S \wedge t/s \longrightarrow t'/s' \\ \Rightarrow \exists S' \supseteq S. \vdash s' : S' \wedge S'; \emptyset \vdash t' : T$$

`PROGRESS:`

$$S; \emptyset \vdash t : T \wedge \vdash s : S \Rightarrow (\text{isvalue } t) \vee (\exists t's'. t/s \longrightarrow t'/s')$$

In contrast, using the omni-small-step judgment, we can establish type soundness through a single induction on the typing judgment. To that end, we formulate a lemma in terms of the predicate $t/s \longrightarrow P$, instantiating the set P as the set of configurations t'/s' such that t' admits the same type as t , and such that s' admits a type that extends the type of s .

`INDUCTION-FOR-TYPE-SOUNDNESS (OMNI-SMALL-STEP)`

$$\begin{aligned} & (S; \emptyset \vdash t : T) \wedge (\vdash s : S) \\ \Rightarrow & (\text{isvalue } t) \vee (t/s \longrightarrow \{(t', s') \mid \exists S' \supseteq S. (\vdash s' : S') \wedge (S'; \emptyset \vdash t' : T)\}) \end{aligned}$$

⁵The generic entailment from `INDUCTION-FOR-TYPE-SOUNDNESS` to `TYPE-SOUNDNESS` holds for any typing judgment of the form $\emptyset \vdash t : T$ and for any judgment $t \longrightarrow P$ related to the small-step judgment $t \longrightarrow t'$ in the expected way, that is, satisfying the property `OMNI-SMALL-STEP-IFF-PROGRESS-AND-CORRECT` from [§3.2](#).

4.3 Omni-Big-Step Type-Soundness Proof for an Imperative Language

Traditionally, a big-step soundness proof can only be carried out if the semantics is completed using error-propagation rules. Here, we demonstrate how to establish type soundness with respect to an omni-big-step judgment, without any need for error-propagation rules.

Consider a type T and a store typing S . We define $\llbracket T/S \rrbracket$ as the set of final configurations of the form v/s such that the state s admits a type S' that extends S , and the value v admits type T , under the store typing S' . Formally:

$$\llbracket T/S \rrbracket \equiv \left\{ (v, s) \mid \exists S' \supseteq S. (\vdash s : S') \wedge (S' \vdash v : T) \right\}$$

The coinductive omni-big-step judgment $t/s \Downarrow^{\text{co}} \llbracket T/S \rrbracket$ asserts that any evaluation of t/s executes safely, without ever getting stuck; and that if an evaluation reaches a final configuration v/s' , then this configuration satisfies the postcondition $\llbracket T/S \rrbracket$. Given our definition of $\llbracket T/S \rrbracket$, the judgment $t/\emptyset \Downarrow^{\text{co}} \llbracket T/\emptyset \rrbracket$ thus captures exactly the type-soundness property associated with the typing judgment $\emptyset; \emptyset \vdash t : T$.

Type soundness may be established by proving the following statement by coinduction.

$$\text{COINDUCTION-FOR-TYPE-SOUNDNESS (OMNI-BIG-STEP):} \\ S; \emptyset \vdash t : T \quad \wedge \quad \vdash s : S \quad \Rightarrow \quad t/s \Downarrow^{\text{co}} \llbracket T/S \rrbracket$$

The coinduction hypothesis asserts that we can invoke the result that we are trying to prove, provided that we have already made some progress by applying one of the introduction rules that define the coinductive judgment $t/s \Downarrow^{\text{co}} Q$. The first step of the proof is to perform a case analysis on the typing hypothesis $S; \emptyset \vdash t : T$. We then consider each of the possible typing rules one-by-one. For each typing rule, we exploit the constraints that apply to the term t to show that there exists an omni-big-step evaluation rule that applies to this term. We apply this rule to the current proof obligation, of the form $t/s \Downarrow^{\text{co}} \llbracket T/S \rrbracket$. For each evaluation judgment that appears as a premise of the evaluation rule, we can invoke the coinduction hypothesis.

Like for the small-step settings, we proved once-and-for-all that the statement COINDUCTION-FOR-TYPE-SOUNDNESS entails TYPE-SOUNDNESS.

Our coinductive omni-big-step approach offers, to those who have good reasons to work with a big-step-style semantics, a means to establish type soundness without introducing error rules.

Regarding the comparison with the standard preservation-and-progress approach, at this stage we cannot draw general conclusions on whether omni-big-step and omni-small-step type-soundness proofs are more effective, because we considered a relatively simple language. Nevertheless, it appears that each of the two approaches that we propose results in proof scripts that (1) require only one induction or one coinduction instead of two separate inductions, (2) are no longer than with preservation and progress separated, and (3) avoid the issue of nested inversions requiring a number of cases quadratic in the size of the language.

5 DEFINITION OF TRIPLES

5.1 Semantics of Hoare Triples for Nondeterministic Languages

A Hoare triple, written $\{H\} t \{Q\}$, describes the behavior of the evaluation of the configurations t/s for any s satisfying the precondition H , in terms of the postcondition Q . The exact interpretation of a triple depends on whether it accounts for *total correctness* or *partial correctness*, which differ on how they account for termination. For nondeterministic languages, the key notions of interest for defining a triple $\{H\} t \{Q\}$ are as follows.

- **Safety:** for any s satisfying H , none of the possible evaluations of t/s can get stuck.
- **Correctness:** for any s satisfying H , if t/s can evaluate to v/s' , then $Q v s'$ holds.

- **Termination:** for any s satisfying H , all possible evaluations of t/s are finite.
- **Partial correctness:** safety and correctness hold.
- **Total correctness:** safety, correctness, and termination hold.

5.2 Definition of Hoare Triples w.r.t. Omni-Big-Step Semantics

For a nondeterministic semantics, a total-correctness Hoare triple asserts that every possible execution of t/s terminates and satisfies the postcondition. This property can be captured directly by the omni-big-step judgment:

$$\text{total, nondeterministic } \{H\} t \{Q\} \equiv \forall s. H s \Rightarrow (t/s \Downarrow Q)$$

For a nondeterministic semantics, a partial-correctness Hoare triple asserts that every possible execution of t/s either diverges or terminates and satisfies the postcondition. This property can be captured directly by the coinductive omni-big-step judgment:

$$\text{partial, nondeterministic } \{H\} t \{Q\} \equiv \forall s. H s \Rightarrow (t/s \Downarrow^{\text{co}} Q)$$

Note that instantiating Q with the always-false predicate in the partial-correctness triple yields a characterization of programs whose execution always diverges—and never gets stuck.

Note that, reciprocally, an omni-big-step judgment can be reformulated as a Hoare triple, using a singleton precondition to constrain the input state. For example, in the case of total correctness:

$$(t/s \Downarrow Q) \iff \text{total, nondeterministic } \{(\lambda s'. s' = s)\} t \{Q\}.$$

5.3 Deriving Reasoning Rules for Triples

After defining a notion of Hoare triples, the usual next step is to state and prove the reasoning rules associated with these triples. These rules are proved correct with respect to the semantics. Consider for example the case of a let-binding. Let us compare the semantics rule OMNI-BIG-LET with the Hoare-logic rule HOARE-LET.

Throughout this section, we formulate rules by viewing postconditions as predicates of type $\text{val} \rightarrow \text{state} \rightarrow \text{Prop}$, as this presentation style is more idiomatic in program logics. Also, we here present reasoning rules using the horizontal bar, but keep in mind that the statements are not inductive definitions: they correspond to lemmas proved correct with respect to the semantics.

$$\begin{array}{c} \text{OMNI-BIG-LET} \\ \frac{t_1/s \Downarrow Q_1 \quad (\forall v' s'. Q_1 v' s' \Rightarrow ([v'/x] t_2)/s' \Downarrow Q)}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow Q} \end{array} \qquad \begin{array}{c} \text{HOARE-LET} \\ \frac{\{H\} t_1 \{Q_1\} \quad (\forall v'. \{Q_1 v'\} ([v'/x] t_2) \{Q\})}{\{H\} (\text{let } x = t_1 \text{ in } t_2) \{Q\}} \end{array}$$

The only difference between the two rules is that the first one considers one specific state s , whereas the second one considers a set of possible states satisfying the precondition H . By exploiting the fact that $\{H\} t \{Q\}$ is defined as $\forall s. H s \Rightarrow (t/s \Downarrow Q)$, it is straightforward to establish that HOARE-LET is a consequence of OMNI-BIG-LET. The corresponding Coq proof script witnesses the simplicity of the proof: “intros. apply mbig_let; eauto.”

To give one more example, consider the consequence rule. The Hoare-logic rule is, again, an immediate consequence from the omni-big-step rule.

$$\begin{array}{c} \text{OMNI-BIG-CONSEQUENCE} \\ \frac{t/s \Downarrow Q \quad Q \subseteq Q'}{t/s \Downarrow Q'} \end{array} \qquad \begin{array}{c} \text{HOARE-CONSEQUENCE} \\ \frac{H' \subseteq H \quad \{H\} t \{Q\} \quad Q \subseteq Q'}{\{H'\} t \{Q'\}} \end{array}$$

5.4 Deriving Weakest-Precondition-Style Reasoning Rules

The weakest-precondition operator, written $\text{wp } t \ Q$, computes the weakest predicate H for which the triple $\{H\} t \{Q\}$ holds. Here, “weakest” is interpreted w.r.t. the entailment relation, written $H \vdash H'$ and defined as pointwise predicate implication ($\forall s. H s \Rightarrow H' s$). Weakest reasoning rules are expressed in the form of entailments, e.g., the rule for let-bindings is:

$$\text{WP-LET:} \quad \text{wp } t_1 \left(\lambda v'. \text{wp } ([v'/x] t_2) \ Q \right) \vdash \text{wp } (\text{let } x = t_1 \text{ in } t_2) \ Q.$$

Many proof tools simply axiomatize the weakest-precondition rules. In a *foundational* approach, however, one needs to prove the reasoning rules correct with respect to the formal semantics of the source language.

What is very appealing about describing the semantics of the language using an omni-big-step semantics is that it delivers the weakest-precondition-style reasoning rules almost for free. Indeed, the interpretation of the inductive judgment $t/s \Downarrow Q$ matches, up to the order of arguments, the standard interpretation of the weakest-precondition operator.

$$t/s \Downarrow Q \iff \text{wp } t \ Q$$

Thus, in a foundational approach, we can formally define $\text{wp } t \ Q$ in terms of the omni-big-step semantics as $\lambda t Q s. (t/s \Downarrow Q)$. It remains to describe how the weakest-precondition-style reasoning rules can be derived from the omni-big-step evaluation rules.

Reasoning rules in weakest-precondition style are even easier to derive than the rules for triples. Consider for example the semantics rule and the wp-reasoning rule associated with a let-binding.

$$\frac{t_1/s \Downarrow Q_1 \quad (\forall v' s'. Q_1 v' s' \Rightarrow ([v'/x] t_2)/s' \Downarrow Q)}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow Q} \text{OMNI-BIG-LET}$$

To derive the rule WP-LET from OMNI-BIG-LET, it suffices to instantiate Q_1 as $\lambda v_1. \text{wp } ([v_1/x] t_2) \ Q$.

In many cases, the wp rule is nothing but a copy of the omni-big-step rule with arguments reordered. One interesting case is that of loops—“while” loops have not been discussed so far, but they appear in the language used for the case studies in §6. Typically, standard weakest-precondition rules for while loops are stated using loop invariants. In contrast, an omni-big-step rule essentially unfolds the first iteration of the loop, just like in a standard big-step semantics. From that unfolding rule, one can derive the loop-invariant-based rule by induction, in just a few lines of proof.

In summary, by considering a semantics expressed in omni-big-step style, one can derive practical reasoning rules, both in Hoare-triple style and in weakest-precondition style, in most cases via one-line proofs. The construction of a program logic on top of an omni-big-step semantics is thus a significant improvement, both over the use of a standard big-step semantics, which falls short in the presence of nondeterminism; and over the use of a small-step semantics, which requires much more work for deriving the reasoning rules, especially if one aims for total correctness. Besides, a major benefit of considering an omni-big-step semantics is that, unlike a set of weakest-precondition reasoning rules, it delivers an induction principle for reasoning about program executions.

6 COMPILER-CORRECTNESS PROOFS FOR TERMINATING PROGRAMS

6.1 Motivation: Avoiding Both Backward Simulations and Artificial Determinism

Following CompCert’s terminology [Leroy 2009], one particular evaluation of a program can admit one out of four possible behaviors: *terminate* (with a value, an exception, a fatal error, etc.), trigger *undefined behavior*, *diverge silently* after performing a finite number of I/O operations, or be *reactive* by performing an infinite sequence of I/O operations. Whether an error such as a division by zero is considered as a terminating behavior or as an undefined behavior is a design decision associated

with each programming language. A general-purpose compiler ought to preserve behaviors, except that undefined behaviors can be replaced with anything.

In this paper, we focus on proofs of compiler correctness for programs that always terminate safely. Such a result is sufficient for many practical applications in software verification where source programs are proven to be safe, and often, the only use case for nontermination is a top-level infinite event-handling loop, which can be implemented outside of the compiler [Erbsen et al. 2021]. We leave to future work the application of omnisemantics to the correct compilation of programs that diverge, react, or trigger undefined behavior on some inputs but not others.

In the particular case of a deterministic programming language, compiler correctness for terminating programs can be established via a *forward-simulation* proof. Such a proof consists of showing that each step from the source program *corresponds to* a number of steps in the compiled program. The correspondence involved is captured by a relation between source states and target states. Such forward-simulation proofs work well in practice. The main problem is that they do not generalize to nondeterministic languages.

Indeed, in the presence of nondeterminism, a source program may have several possible executions. As we restrict ourselves to the case of terminating programs, let us assume that all executions of the source program terminate, only possibly with different results. In that setting, a compiler is correct if (1) the compiled program always terminates, and (2) for any result that the compiled program may produce, the source program could have produced that result. It may not be intuitive at first, but the inclusion is indeed *backwards*: the set of behaviors of the target program must be included in the set of behaviors of the source program.

To establish the backward behavior inclusion, one may set up a *backward-simulation* proof. Such a proof consists of showing that each step from the target program corresponds to one or more steps in the source program.⁶ Yet, backward simulations are much more unwieldy to set up than forward simulations. Indeed, in most cases one source program step is implemented by multiple steps in the compiled program, thus a backward-simulation relation typically needs to relate many more pairs than a forward-simulation relation.

This observation motivated the CompCert project [Leroy 2009] to exploit forward simulations as much as possible, at the cost of making the semantics of the intermediate language deterministic even when it is not natural to do so, and even when implementing determinacy requires introducing artificial functions for, e.g., computing fresh memory locations in a deterministic manner.

In this section:

- We explain how omnisemantics sidestep the need for backward simulations, by carrying out forward-simulation proofs of compiler correctness, for nondeterministic terminating programs.
- We show how the idea generalizes to languages including I/O operations and to the case where the source language and target language are different.
- We present two case studies: one transformation that increases the amount of nondeterminism and one that decreases the amount of nondeterminism.
- We comment on the fact that our second case study features an omni-big-step semantics for the source language, whereas it features an omni-small-step semantics for the target language.

⁶ The number of corresponding steps in the source program should not be zero, otherwise the target program could diverge whereas the source program terminates. In practice, however, it is not always easy to find one source-program step that corresponds to a target-program step. In such situations, one may consider a generalized version of backward simulations that allow for zero source program steps, provided that some well-founded measure decreases [Leroy 2009].

6.2 Establishing Correctness via Forward Simulations using Omnisemantics

Consider a compilation function written $C(t)$. For simplicity, we assume that the source and target language are identical, we assume that compilation does not alter the result values, and we assume the language to be state-free—we will generalize the results in the next subsection. In this subsection, $t \Downarrow v$ denotes the standard big-step judgment, $t \Downarrow Q$ denotes the omni-big-step judgment, and $\text{terminates}(t)$ asserts that all executions of t terminate safely, without undefined behavior. The compiler-correctness result for terminating programs captures preservation of termination and backward inclusion for results—points (1) and (2) stated earlier.

CORRECTNESS-FOR-TERMINATING-PROGRAMS:

$$\text{terminates}(t) \quad \Rightarrow \quad \text{terminates}(C(t)) \quad \wedge \quad (\forall v. C(t) \Downarrow v \Rightarrow t \Downarrow v)$$

We claim that this correctness result can be derived from the following statement, which describes forward preservation of specifications.

$$\text{OMNI-FORWARD-PRESERVATION:} \quad \forall Q. \quad t \Downarrow Q \Rightarrow C(t) \Downarrow Q$$

Let us demonstrate the claim. Let us assume that $\text{terminates}(t)$ hold. First of all, recall from §2.2 the equivalence named OMNI-BIG-STEP-IFF-TERMINATES-AND-CORRECT that relates the omni-big-step judgment and the termination judgment.

$$t \Downarrow Q \quad \iff \quad \text{terminates}(t) \quad \wedge \quad (\forall v. (t \Downarrow v) \Rightarrow v \in Q)$$

Exploiting this equivalence, the OMNI-FORWARD-PRESERVATION assumption reformulates as follows.

$$\begin{aligned} \forall Q. \quad & \left(\text{terminates}(t) \quad \wedge \quad (\forall v. (t \Downarrow v) \Rightarrow v \in Q) \right) \\ & \Rightarrow \left(\text{terminates}(C(t)) \quad \wedge \quad (\forall v. (C(t) \Downarrow v) \Rightarrow v \in Q) \right) \end{aligned}$$

The hypothesis $\text{terminates}(t)$ holds by assumption. Let us instantiate Q as the strongest post-condition for t , that is, as the set $\{v \mid (t \Downarrow v)\}$. We obtain:

$$(\forall v. (t \Downarrow v) \Rightarrow (t \Downarrow v)) \quad \Rightarrow \quad \text{terminates}(C(t)) \quad \wedge \quad (\forall v. (C(t) \Downarrow v) \Rightarrow (t \Downarrow v)).$$

The premise is a tautology, and the conclusion proves CORRECTNESS-FOR-TERMINATING-PROGRAMS.

6.3 Generalization to Handle I/O and Cross-Language Compilation

More generally, the behavior of a terminating program consists of the final result and its interactions with the outside world (input and output). These interactions include, e.g., interaction with the standard input and output streams, system calls, etc. Each interaction is called an *event*, and the semantics judgment is extended to collect such events into a *trace* τ . Figure 5 shows three illustrative cases of how the rules from Figure 2 are augmented with traces, making the choice to treat rand calls as observable events while reference-allocation nondeterminism remains internal.

Requiring a compiler to preserve only the nondeterministic choices recorded in the trace enables us to pick and choose which (external) interactions must be preserved by compilations and which (internal) nondeterministic choices the compiler may resolve as it sees fit. As a particularly fine-grained example, the trace might record that `malloc` was called and succeeded but omit the pointer it returned, to allow for optimizations that reduce the amount of allocation. To our knowledge, this level of flexibility is unique to omnisemantics. For a forward-simulation-based compiler-correctness proof, constructing a deterministic model of all internal nondeterminism can be arbitrarily complicated (the CompCert memory model is an example).

We restrict our attention to semantics that only accept terminating commands c that do not go wrong and do not return values, for the rest of this section. For languages of terms (that return values) rather than commands (that do not return values), we would need a representation relation

$$\begin{array}{c}
\text{OMNI-BIG-LET-TRACE} \\
\frac{t_1/s/\tau \Downarrow Q_1 \quad (\forall v', s', \tau' \in Q_1. ([v'/x] t_2)/s'/\tau' \Downarrow Q)}{(\text{let } x = t_1 \text{ in } t_2)/s/\tau \Downarrow Q} \\
\\
\text{OMNI-BIG-RAND-TRACE} \qquad \text{OMNI-BIG-REF} \\
\frac{n > 0 \quad (\forall m. 0 \leq m < n \Rightarrow (m, s, (n, m) :: \tau) \in Q)}{(\text{rand } n)/s/\tau \Downarrow Q} \quad \frac{\forall p \notin \text{dom } s. (p, s[p := v], \tau) \in Q}{(\text{ref } v)/s/\tau \Downarrow Q}
\end{array}$$

Fig. 5. Omni-big-step semantics with traces, selected rules

between source-level and target-level values—we omit one here for brevity, but §6.4 tackles a similar challenge. In the current setting, *behavior inclusion* holds between a source-language program and a target-language program if all traces that the target-language program can produce (according to traditional small-step or big-step semantics) can also be produced by the source-language program. More formally, we define the traces that can be produced from a starting configuration $c/s/\tau$ as

$$\text{traces}(c, s, \tau) := \{\tau' \mid \exists s'. c/s/\tau \Downarrow s'/\tau'\}$$

and say that a compiler $C()$ satisfies behavior inclusion for a command starting from the initial source-level state s_{src} related to the initial target-level state s_{tgt} and initial trace τ if

$$\text{traces}(C(c), s_{\text{tgt}}, \tau) \subseteq \text{traces}(c, s_{\text{src}}, \tau)$$

Assuming omni-big-step semantics \Downarrow_{src} and \Downarrow_{tgt} for the source and target languages, plus a relation R between source- and target-language states, we define *omnisemantics simulation* as follows:

$$\begin{array}{c}
\forall s_{\text{src}} s_{\text{tgt}} \tau Q. R(s_{\text{src}}, s_{\text{tgt}}) \wedge c/s_{\text{src}}/\tau \Downarrow_{\text{src}} Q \implies C(c)/s_{\text{tgt}}/\tau \Downarrow_{\text{tgt}} Q \\
\text{where } Q_R(s'_{\text{tgt}}, \tau') := \exists s'_{\text{src}}. R(s'_{\text{src}}, s'_{\text{tgt}}) \wedge Q(s'_{\text{src}}, \tau')
\end{array}$$

Our goal is to prove that an omnisemantics simulation implies trace inclusion if the source program is safe. We rely on two properties: First, soundness of omni-big-step semantics with respect to traditional big-step semantics:

$$\forall c s s' \tau \tau' Q. c/s/\tau \Downarrow s'/\tau' \wedge c/s \Downarrow Q \implies Q(s', \tau') \quad (1)$$

And conversely, that a program that terminates safely and whose traditional big-step executions all satisfy a postcondition also has an omnisemantics derivation:

$$\forall c s \tau Q. \text{terminates}(c, s, \tau) \wedge (\forall s' \tau'. c/s/\tau \Downarrow s'/\tau' \implies Q(s', \tau')) \implies c/s/\tau \Downarrow Q \quad (2)$$

To show trace inclusion, i.e. $\text{traces}(C(c), s_{\text{tgt}}, \tau) \subseteq \text{traces}(c, s_{\text{src}}, \tau)$, we can assume a target-language execution $C(c)/s_{\text{tgt}}/\tau \Downarrow s'_{\text{tgt}}/\tau'$ producing trace τ' , and we need to show $\tau' \in \text{traces}(c, s_{\text{src}}, \tau)$. By applying (2) to the source program (whose termination we assume) and setting $Q(s'_{\text{src}}, \tau') := c/s_{\text{src}}/\tau \Downarrow s'_{\text{src}}/\tau'$ so that the second premise becomes a tautology, we obtain the source-level omnisemantics derivation $c/s_{\text{src}}/\tau \Downarrow (\lambda s'_{\text{src}} \tau'. c/s_{\text{src}}/\tau \Downarrow s'_{\text{src}}/\tau')$. Passing this fact into the omnisemantics simulation yields $C(c)/s_{\text{tgt}}/\tau \Downarrow (\lambda s'_{\text{tgt}} \tau'. \exists s'_{\text{src}}. R(s'_{\text{src}}, s'_{\text{tgt}}) \wedge c/s_{\text{src}}/\tau \Downarrow s'_{\text{src}}/\tau')$. Now we can apply (1) to this fact and the originally assumed target-level execution and obtain an s'_{src} such that $R(s'_{\text{src}}, s'_{\text{tgt}})$ and $c/s_{\text{src}}/\tau \Downarrow s'_{\text{src}}/\tau'$, which by definition is exactly what needs to hold to have $\tau' \in \text{traces}(c, s_{\text{src}}, \tau)$.

6.4 Case Study: Compiling Immutable Pairs to Heap-Allocated Records

This section describes a simple compiler pass that *increases* the amount of nondeterminism. The source language assumes a primitive notion of tuples, whereas the target language encodes such tuples by means of heap allocation. This case study is formalized with respect to a language based

$$\begin{array}{c}
\text{EVAL-UNOP} \\
\frac{(y, v_y) \in \ell \quad \text{evalunop}(op, v_y, v) \quad Q(m, \ell[x := v], \tau)}{(x = op(y))/m/\ell/\tau \Downarrow Q} \\
\\
\text{EVAL-STORE} \\
\frac{(x, a) \in \ell \quad (a + n) \in \text{dom } m \quad (y, v) \in \ell \quad Q(m[(a + n) := v], \ell, \tau)}{(x[n] = y)/m/\ell/\tau \Downarrow Q} \\
\\
\text{EVAL-ALLOC} \\
\frac{(\forall a \bar{v}. \text{len}(\bar{v}) = n \wedge a, \dots, (a + n - 1) \notin \text{dom } m \implies Q(m[(a, \dots, (a + n - 1)) := \bar{v}], \ell[x := a], \tau))}{(x = \text{alloc}(n))/m/\ell/\tau \Downarrow Q} \\
\\
\text{EVAL-INPUT} \\
\frac{\forall n. Q(m, \ell[x := n], \tau :: \text{IN } n)}{(x = \text{input}())/m/\ell/\tau \Downarrow Q} \\
\\
\text{EVAL-WHILE-AGAIN} \\
\frac{(\forall m' \ell' \tau'. Q_1(m', \ell', \tau') \implies (\text{while } x \text{ do } c)/m'/\ell'/\tau' \Downarrow Q)}{(\text{while } x \text{ do } c)/m/\ell/\tau \Downarrow Q} \\
\\
\text{EVAL-WHILE-DONE} \\
\frac{(x, 0) \in \ell \quad Q(m, \ell, \tau)}{(\text{while } x \text{ do } c)/m/\ell/\tau \Downarrow Q} \\
\\
\text{EVAL-SEQ} \\
\frac{c_1/m/\ell/\tau \Downarrow Q_1 \quad (\forall m' \ell' \tau'. Q_1(m', \ell', \tau') \implies c_2/m'/\ell'/\tau' \Downarrow Q)}{(c_1; c_2)/m/\ell/\tau \Downarrow Q}
\end{array}$$

Fig. 6. Nondeterministic omni-big-step semantics for an imperative language (selected rules)

on commands whose arguments all must be variables. Such a language could be an intermediate language in a compiler pipeline, reached after an expression-flattening phase.

Language syntax. We let c denote a command, x , y , and z denote identifiers, and n denote unbounded natural-number constants. The grammar of the language is as follows.

$$\begin{aligned}
c \quad := \quad & x = \text{unop}(y) \mid x = \text{binop}(y, z) \mid x = \text{input}() \mid \text{output}(x) \mid x = y[n] \mid x[n] = y \mid \\
& x = \text{alloc}(n) \mid x = n \mid x = y \mid c_1; c_2 \mid \text{if } x \text{ then } c_1 \text{ else } c_2 \mid \text{while } x \text{ do } c \mid \text{skip}
\end{aligned}$$

We actually consider two variants of this language, differing only in the types of values and in the available unary operators unop and binary operators binop . The source language features an inductively defined type of values that can be natural numbers n or immutable pairs of values (i.e., the grammar of values is $v := n \mid (v, v)$). It includes as unary operators the projection functions fst and snd (defined only on pairs) and the Boolean negation not (defined only on $\{0, 1\}$). Its binary operators are addition (+) and pair creation mkpair . The target language admits only natural numbers as values. It includes only the negation and addition operators.

Omn-big-step semantics. For both languages, the omni-big-step evaluation judgment takes the form $c/m/\ell/\tau \Downarrow Q$, where c is a command, m is a memory state (a partial map from natural numbers to natural numbers), ℓ is an environment of local variables (a partial map from identifiers to values, whose type differs between the source and target languages as described above), τ denotes the I/O trace made of the events already performed *before* executing c , and the postcondition Q is a predicate over triples of the form (m', ℓ', τ') . A trace consists of a list of I/O events e whose grammar is $e := \text{IN } n \mid \text{OUT } n$.

The rules defining the judgment appear in Figure 6. They are common to both languages—only the set of supported unary and binary operators differs. The semantics of operators are defined by two straightforward auxiliary relations (spelled out in Appendix J), $\text{evalunop}(\text{unop}, v_1, v_2)$ asserting that applying unop to value v_1 results in v_2 , and $\text{evalbinop}(\text{binop}, v_1, v_2, v_3)$ asserting that applying binop to v_1 and v_2 results in v_3 . The *load* command $x = y[n]$ requires that the local variable y

contains a natural number a and stores the value of the memory at address $a + n$ into variable x (and is undefined if $a + n$ is not mapped by the memory). The *store* command $x[n] = y$ stores the natural number contained in the local variable y at memory location $a + n$, where a is the address contained in local variable x , but only if memory at address $a + n$ has already been allocated.

The command $x = \text{input}()$ reads a natural number n , stores it into local variable x , and adds the event $(\text{IN } n)$ to the event trace. The number n is chosen nondeterministically but recorded in the trace, resulting in *external* nondeterminism. The language has a built-in memory allocator but, for simplicity, we do not deal with deallocation. The command $x = \text{alloc}(n)$ nondeterministically picks an address (natural number) a such that a , as well as the $n - 1$ addresses following a , are not yet part of the memory, initializes these addresses with nondeterministically chosen values, and returns a . This rule encodes *internal* nondeterminism, because this action is not recorded in the event trace. Semantics of while loops are given by sequencing the first iteration with the loop itself as long as the loop test succeeds.

In practice, we found it convenient also to derive a chained version EVAL-SEQ-CHAINED of the omni-big-step rule EVAL-SEQ, just like we did for omni-small-step rules in §3.2.

$$\text{EVAL-SEQ-CHAINED} : \quad c_1/m/\ell/\tau \Downarrow (\lambda m' \ell' \tau'. (c_2/m'/\ell'/\tau' \Downarrow Q)) \Rightarrow (c_1; c_2)/m/\ell/\tau \Downarrow Q$$

One might wonder why we did not directly use this rule in the inductively defined judgment, and the reason is related to Coq's *strict positivity* requirement, on which we further elaborate in Appendix A.

Compilation function. The compilation function C lays out the pairs of the source language on the heap memory of the target language. This function is defined recursively on the source program. It maps the source-language operators that are not supported by the target language as follows.

$$\begin{aligned} C(x = \text{fst}(y)) &:= x = y[0] \\ C(x = \text{snd}(y)) &:= x = y[1] \\ C(x = \text{mkpair}(y, z)) &:= \text{tmp} = \text{alloc}(2); \text{tmp}[0] = y; \text{tmp}[1] = z; x = \text{tmp} \end{aligned}$$

Note that to compile `mkpair`, we cannot simply store the address returned by `alloc` directly into x , because if x is the same variable name as y or z , then we would be overwriting the argument. For this reason, we use a temporary variable `tmp` that we declare to be reserved for compiler usage.

Simulation relation. To carry out the proof of correctness of the function $C(c)$, we introduce a simulation relation R for relating a source-language state (m_1, ℓ_1) with a target-language state (m_2, ℓ_2) . To that end, we first define the relation $\text{valuerepr}(v, w, m)$, to relate a source-language value v with the corresponding target-language value w , in a target-language memory m . This relation is implemented as the recursive function shown below—it could equally well consist of an inductive definition. A pair (v_1, v_2) is represented by address w if recursively v_1 is represented by the value at address w and v_2 is represented by the value at address $w + 1$. A natural number n has the same representation on the target-language level, i.e. we just assert $w = n$.

$$\begin{aligned} \text{valuerepr}((v_1, v_2), w, m) &:= (\exists w_1, (w, w_1) \in m \wedge \text{valuerepr}(v_1, w_1, m)) \wedge \\ &\quad (\exists w_2, (w + 1, w_2) \in m \wedge \text{valuerepr}(v_2, w_2, m)) \\ \text{valuerepr}(n, w, m) &:= w = n \end{aligned}$$

The relationship R between source and target states can then be defined using valuerepr . In the definition shown below, we write $m_2 \supseteq m_1$ to mean that memory m_2 extends m_1 , and we write $m_2 \setminus m_1$ to denote the map-subtraction operator that restricts m_2 to contain only addresses not bound in m_1 . The locations bound by m_2 but not by m_1 correspond to the memory addresses of the

pairs allocated on the heap in the target language.

$$R((m_1, \ell_1), (m_2, \ell_2)) := \text{tmp} \notin \text{dom } \ell_1 \wedge m_2 \supseteq m_1 \wedge \\ \forall (x, v) \in \ell_1. \exists w. (x, w) \in \ell_2 \wedge \text{valuerepr}(v, w, m_2 \setminus m_1)$$

Correctness proof. We are now ready to tackle the omni-forward-simulation proof.

THEOREM 6.1 (OMNISIMULATION FOR THE PAIR-HEAPIFICATION COMPILER). $\forall c m_{src} \ell_{src} m_{tgt} \ell_{tgt} \tau Q.$

$$\text{tmp} \notin \text{vars}(c) \wedge R((m_{src}, \ell_{src}), (m_{tgt}, \ell_{tgt})) \wedge c/m_{src}/\ell_{src}/\tau \Downarrow_{src} Q \implies C(c)/m_{tgt}/\ell_{tgt}/\tau \Downarrow_{tgt} Q_R \\ \text{where } Q_R(m'_{tgt}, \ell'_{tgt}, \tau') := \exists m'_{src} \ell'_{src}. R((m'_{src}, \ell'_{src}), (m'_{tgt}, \ell'_{tgt})) \wedge Q(m'_{src}, \ell'_{src}, \tau')$$

PROOF. By induction on the derivation of $c/m_{src}/\ell_{src}/\tau \Downarrow Q$. In each case, the goal to prove is initially of the form $C(c)/m_{tgt}/\ell_{tgt}/\tau \Downarrow Q_R$, where c has some structure that allows us to simplify $C(c)$ into a more concrete program snippet. We consider the resulting simplified goal as an invocation of a weakest-precondition generator on that program snippet, and we view the rules of Figure 6 as weakest-precondition rules that we can apply in order to step through the program snippet, using the hypotheses obtained from inverting the source-level derivation $c/m_{src}/\ell_{src}/\tau \Downarrow Q$ to discharge the side conditions that arise. Whenever we encounter a sequence of commands, we use EVAL-SEQ-CHAINED instead of EVAL-SEQ, so that we do not have to provide an intermediate postcondition. In the cases where commands have subcommands, we use the inductive hypotheses about their execution as if they were previously proven lemmas about these “functions.”

We only present the case where $c = (x = \text{mkpair}(y, z))$ in more detail: We have to prove a goal of the form $C(x = \text{mkpair}(y, z))/m_{tgt}/\ell_{tgt}/\tau \Downarrow Q_R$, which simplifies to

$$(\text{tmp} = \text{alloc}(2); \text{tmp}[0] = y; \text{tmp}[1] = z; x = \text{tmp})/m_{tgt}/\ell_{tgt}/\tau \Downarrow Q_R$$

Applying EVAL-SEQ-CHAINED turns it into:

$$(\text{tmp} = \text{alloc}(2))/m_{tgt}/\ell_{tgt}/\tau \Downarrow (\lambda m'_{tgt} \ell'_{tgt} \tau'. (\text{tmp}[0] = y; \text{tmp}[1] = z; x = \text{tmp})/m'_{tgt}/\ell'_{tgt}/\tau' \Downarrow Q_R)$$

Applying EVAL-ALLOC turns it into:

$$\forall a \bar{v}. \text{len}(\bar{v}) = 2 \implies a, a + 1 \notin \text{dom } m_{tgt} \implies \\ (\text{tmp}[0] = y; \text{tmp}[1] = z; x = \text{tmp})/m_{tgt}[a..(a + 1) := \bar{v}]/\ell_{tgt}[\text{tmp} := a]/\tau \Downarrow Q_R$$

Note how the fact that the address a and the list of initial values \bar{v} are chosen nondeterministically naturally shows up as a universal quantification, and note how the memory and locals appearing in the state to the left of the \Downarrow have been updated by the alloc function. After introducing these universally quantified variables and the hypotheses, we again have a goal of the form “ $\dots \Downarrow \dots$ ” and continue by applying EVAL-SEQ-CHAINED, EVAL-STORE, EVAL-SEQ-CHAINED, EVAL-STORE, EVAL-SET. Finally, we prove Q_R for the locals and memory updated according to the various EVAL-... rules that we applied by using map laws and the initial hypothesis $R((m_{src}, \ell_{src}), (m_{tgt}, \ell_{tgt}))$. \square

6.5 Case Study: Introduction of Stack Allocation

This second case study illustrates the case of a transformation that reduces the amount of nondeterminism. The transformation consists of adding a *stack-allocation* feature to the compiler developed by Erbsen et al. [2021]. Proving this transformation correct using an omni-big-step forward simulation was straightforward and took us only a few days of work—most of the work was *not* concerned with dealing with nondeterminism. This smooth outcome is in stark contrast to the outlook of using traditional evaluation judgments: verifying the same transformation would have required either more complex invariants to set up a backward simulation, or completely rewriting the memory model so that pointers are represented by deterministically generated unobservable identifiers to allow for a compiler-correctness proof by forward simulation. In fact, addressable stack allocation

was initially omitted from the language exactly to avoid these intricacies (based on the experience from CompCert), but switching to omniseantics made its addition local and uncomplicated.

The input language is an imperative command language similar to the one described in §6.4. The memory is modeled as a partial map from machine words (32-bit or 64-bit integers) to bytes. The stack-allocation feature here consists of a command $\text{let } x = \text{stackalloc}(n)$ in c available in the source language. This command assigns an address to variable x at which n bytes of memory will be available during the execution of command c . Our compiler extension implements this command by allocating the requested n bytes on the stack, computing the address at runtime based on the stack pointer.

The key challenge is that the source-language semantics does not feature a stack. The stack gets introduced further down the compilation chain. Thus, the simplest way to assign semantics to the `stackalloc` function in the source language is to pretend that it allocates memory at a *nondeterministically chosen* memory location. This nondeterministic choice is described using a universal quantification in the omni-big-step rule shown below, like in rule OMNI-BIG-REF from §2.

$$\frac{\forall m_{\text{new}} a. \text{dom } m_{\text{new}} \cap \text{dom } m = \emptyset \wedge \text{dom } m_{\text{new}} = [a, a + n) \implies c/m \cup m_{\text{new}}/\ell[x := a]/\tau \Downarrow \lambda m' \ell' \tau'. P(m' \setminus m_{\text{new}}, \ell', \tau')}{(\text{let } x = \text{stackalloc}(n) \text{ in } c)/m/\ell/\tau \Downarrow P}$$

In the source language, the address returned by `stackalloc` is picked nondeterministically, whereas in the target language the address used for the allocation is deterministically computed, as the current stack pointer augmented with some offset. Thus, the compiler phase that compiles away the `stackalloc` command *reduces* the amount of nondeterminism.

The compiler-correctness proof proceeds by induction on the omniseantics derivation for the source language, producing a target-language execution with a related postcondition. The simulation relation R describes the target-language memory as a disjoint union of unallocated stack memory and the source-language memory state. Critically, the case for `stackalloc` has access to a universally quantified induction hypothesis (derived from the rule shown above) about *target-level executions of $C(c)$ for any address a* .

As the address of the stack-allocated memory is not recorded in the trace, we are free to instantiate it with the specific stack-space address, expressed in terms of compile-time stack-layout parameters and the runtime stack pointer. Reestablishing the simulation relation to satisfy the precondition of that induction hypothesis then involves carving out the freshly allocated memory from unused stack space and considering it a part of the source-level memory instead, matching the compiler-chosen memory layout and the preconditions of the `stackalloc` omniseantics rule. It is this last part that made up the vast majority of the verification work in this case study; handling the nondeterminism itself is as straightforward as it gets.

Note that it would not be possible to complete the proof by instantiating a with a compiler-chosen offset from the stack pointer if the semantics recorded the value of a in the trace. The (unremarkable) proof for the input command in the previous section also has access to a universally quantified execution hypothesis, but it *must* directly instantiate its universally quantified induction hypothesis with the variable introduced when applying the target-level omniseantics input rule to the goal, to match the target-language trace to the source-language trace. Either way, reasoning about the reduction of nondeterminism in an omni-forward-preservation proof boils down to instantiating a universal quantifier.

6.6 Compilation from a Language in Omni-Big-Step to One in Omni-Small-Step

If the semantics of the source language of a compiler phase are most naturally expressed in omni-big-step, but the target language's semantics are best expressed in omni-small-step semantics, it is

convenient to prove an omni-forward simulation directly from a big-step source execution to a small-step target execution. For instance, the compiler in the project by Erbsen et al. [2021] includes such a translation, relating a big-step intermediate language to a small-step assembly language. In fact, this translation happens in the same case study described in the previous subsection. In what follows, we attempt to give a flavor of the proof obligations that arise from switching from omni-big-step to omni-small-step during the correctness proof.

Consider one sample omni-small-step rule, for the load-word instruction `lw` that loads the value at the address stored in register r_2 and assigns it to register r_1 :

$$\frac{\text{ASM-LW} \quad (pc, lw \ r_1 \ r_2) \in m \quad (r_2, a) \in rf \quad (a, v) \in m \quad P(m, rf[r_1 := v], pc + 1, \tau)}{m/rf/pc/\tau \longrightarrow P}$$

Here, we model a machine state s_{tgt} as a quadruple of a memory m (that contains both instructions and data), a register file rf mapping register names to machine words, a program counter pc , and a trace τ . One can prove an omni-forward simulation from big-step source semantics directly to small-step target semantics:

$$\forall s_{src} \ s_{tgt} \ P. R(s_{src}, s_{tgt}) \wedge s_{src} \Downarrow P \implies s_{tgt} \longrightarrow^\diamond (\lambda s'_{tgt}. \exists s'_{src}. R(s'_{src}, s'_{tgt}) \wedge P(s'_{src}))$$

where R asserts, among other conditions, that the memory of the target state s_{tgt} contains the compiled program.

Like the proof described in §6.4, this proof also works by stepping through the target-language program by applying target-language rules and discharging their side conditions using the hypotheses obtained by inverting the source-language execution, with the only difference that instead of using the derived big-step rule EVAL-SEQ-CHAINED for chaining, one now uses the following two rules: EVENTUALLY-STEP-CHAINED and EVENTUALLY-CUT.

Applying EVENTUALLY-STEP-CHAINED turns the goal into an omni-single-small-step goal with a given postcondition, which is suitable to discharge using rules with universally quantified postconditions like ASM-LW. Applying EVENTUALLY-CUT, on the other hand, creates two subgoals containing an uninstantiated unification variable for the intermediate postcondition. The unification variable appears as the postcondition in the first subgoal, so an induction hypothesis with the concrete postcondition from the theorem statement can be applied. In the second subgoal, this postcondition becomes the precondition for the remainder of the execution.

7 RELATED AND FUTURE WORK

This work builds on that of Schäfer et al. [2016], Charguéraud [2020], and Erbsen et al. [2021], all which are discussed in the introduction (§1). We now will review some additional connections.

Relationship to coinductive big-step semantics. Leroy and Grall [2009] argue that fairly complex, optimizing compilation passes can be proved correct more easily using big-step semantics than using small-step semantics. These authors propose to reason about diverging executions using *coinductive big-step semantics*, following up on an earlier idea by Cousot and Cousot [1992]. Leroy and Grall's judgment, written $t/s \uparrow^{\text{co}}$, asserts that there exists a diverging execution of t/s . This judgment is defined coinductively, and a number of its rules refer to the standard big-step judgment. For example, consider the two rules associated with divergence of a let-binding.

$$\frac{t_1/s \uparrow^{\text{co}}}{(\text{let } x = t_1 \text{ in } t_2)/s \uparrow^{\text{co}}} \text{DIV-LET-1} \quad \frac{t_1/s \Downarrow v_1/s' \quad ([v_1/x] t_2)/s' \uparrow^{\text{co}}}{(\text{let } x = t_1 \text{ in } t_2)/s \uparrow^{\text{co}}} \text{DIV-LET-2}$$

An expression $\text{let } x = t_1 \text{ in } t_2$ diverges either because t_1 diverges (rule DIV-LET-1) or because t_1 terminates on a value v_1 and the term $[v_1/x] t_2$ diverges (rule DIV-LET-2). In contrast, the coinductive

omni-big-step judgment involves a single rule, namely OMNI-BIG-LET from Fig. 2. In that rule, if Q_1 is instantiated as the empty set, the second premise of OMNI-BIG-LET becomes vacuous, and we recover the rule DIV-LET-1. Otherwise, if Q_1 is nonempty, then it describes the values v_1 that t_1 may evaluate to. For each possible value v_1 , the second premise of OMNI-BIG-LET requires the term $[v_1/x] t_2$ to diverge, just like in the rule DIV-LET-2. In summary, OMNI-BIG-LET captures at once the logic of both DIV-LET-1 and DIV-LET-2.

Semantics of nondeterministic programs. Nondeterminism appears in the early work on semantics, including the language of guarded commands of Dijkstra [1976] that admits nondeterministic choice where guards overlap, and the par construct of Milner [1975]. Plotkin [1976] develops a *powerdomain* construction to give a fully abstract model in which equivalences such as $(p \text{ par } q) = (q \text{ par } p)$ hold. Francez et al. [1979] also present semantics that map each program to a representation of the set of its possible results. In all these presentations, nondeterminism is *bounded*: only a finite number of choices are allowed.

Subsequent work generalizes the powerdomain interpretation to *unbounded nondeterminism*. For example, Back [1983] considers a language construct $x := \epsilon P$ that assigns x to an arbitrary value satisfying the predicate P —the program has undefined behavior if no such value exists. Apt and Plotkin [1986] address the lack of *continuity* in the models presented in earlier work, still leveraging the notion of powerdomains. Their presentation includes a (countable) nondeterministic assignment operator, written $x := ?$, that sets x to an arbitrary integer in \mathbb{Z} .

Coinductive characterization of safety. Wang et al. [2014] define a safety judgment, written $\text{safe}(t, s)$, to assert that all possible executions of the configuration t/s execute safely, i.e., do not get stuck. To reason in big-step style, and to avoid the cumbersome introduction of error-propagation rules, they consider a coinductive definition. We reproduce below the rule for let-bindings, which reads as follows: to establish that $\text{let } x = t_1 \text{ in } t_2$ executes safely, prove that t_1 executes safely and that, for any possible result v_1 produced by t_1 , the result of the substitution $[v_1/x] t_2$ executes safely.

$$\frac{\text{safe}(t_1, s) \quad \left(\forall v_1 s'. (t_1/s \Downarrow v_1/s') \Rightarrow \text{safe}([v_1/x] t_2, s') \right)}{\text{safe}(\text{let } x = t_1 \text{ in } t_2, s)} \text{SAFE-LET (COINDUCTIVE)}$$

Our judgment $t/s \Downarrow^{\text{co}} Q$ generalizes the notion of safety, by baking the postcondition directly into the judgment (§2.4). It asserts not only that t/s cannot get stuck but also that any potential final configuration belongs to Q . We formalized in Coq the following equivalence.

OMNI-CO-BIG-STEP-IFF-SAFE-AND-CORRECT :

$$t/s \Downarrow^{\text{co}} Q \iff \text{safe}(t, s) \wedge \left(\forall v s'. (t/s \Downarrow v/s') \Rightarrow (v, s') \in Q \right)$$

Our rule OMNI-BIG-LET extends SAFE-LET not just by adding the postcondition Q to the judgment, but also by changing the quantification over v_1/s' . In the rule SAFE-LET, the quantification is constrained by $t_1/s \Downarrow v_1/s'$, whereas in the rule OMNI-BIG-LET, it is constrained by $(v_1, s') \in Q_1$, where Q_1 denotes the postcondition of t_1/s . The key innovation here is that, thanks to the introduction of postconditions, we no longer need to refer to the standard big-step judgment—the judgment $t/s \Downarrow Q$ gives a stand-alone definition of the semantics of the language.

Semantics of reactive programs. One key question is how much of a program’s internal nondeterminism should be reflected in its *execution trace*. At one extreme, one could include a *delay* event, a.k.a. a *tick*, to reflect in the trace each transition performed by the program, following the approaches of Danielsson [2012]. More recent work on interaction trees [Koh et al. 2019; Xia et al. 2019] maps each program to a coinductive structure featuring ticks in addition to I/O steps. Yet, these approaches come at the cost of reasoning “up to removal of a finite number of ticks.”

A promising route to avoiding ticks is the *mixed inductive-coinductive* approach of Nakata and Uustalu [2010], for distinguishing between *reactive* programs that always eventually perform I/O operations and *silently diverging* programs that eventually continue executing forever without performing any I/O. So far, the mixed inductive-coinductive approach has not been proved to scale up to realistic compiler proofs. Yet, perhaps such a result could be achieved by combining the mixed inductive-coinductive approach with the omnisemantics approach.

Semantics of concurrent programs. Concurrency increases the amount of nondeterminism, due to interleavings, and generally increases the sources of undefined behaviors, due in particular to data races. The work on CompCertTSO [Ševčík et al. 2013] shows how to deal with this additional complexity in a compiler-correctness proof. A direction for future work is to investigate the extent to which omni-small-step semantics would help simplifying proofs from CompCertTSO.

Semantics of probabilistic programs. Probabilistic semantics aim to describe not just which executions are possible but also to describe with what probability each execution may happen. A probabilistic small-step execution relation assigns a probability to every transition. One caveat is that probabilities do not suffice to describe all nondeterminism: in particular, memory is allocated at nondeterministically chosen addresses. We refer to Batz et al. [2019] for a solution to this challenge. In the context of program logics, McIver and Morgan [2005] introduce a *weakest preexpectation calculus*. Batz et al. [2019] generalize this notion to set up a *Quantitative Separation Logic*.

Additionally, there is a long line of work aiming at providing denotational models for probabilistic programs—e.g., Staton et al. [2016]; Wang et al. [2019]. Denotational and operational semantics serve different purposes; one important practical benefit of omnisemantics is that it is grounded in inductive definitions, with respect to which proofs by induction can be carried out easily in a proof assistant. An interesting question is whether omnisemantics could be adapted to provide an inductively defined operational semantics that accounts for probabilities, by relating configurations not to sets of outcomes but instead to probability distributions of outcomes.

The problem of termination of probabilistic programs is particularly subtle. On the one hand, one may be interested in capturing that any execution terminates. For example, Staton et al. [2016] define termination as “there exists n , such that termination occurs in n steps.” However, this approach does not apply to infinitely branching nondeterminism. On the other hand, one may design rules to establish *almost-sure termination* or *positive-almost-sure termination* [Chakarov and Sankaranarayanan 2013; Ferrer Fioriti and Hermanns 2015; Kaminski et al. 2016; McIver et al. 2017].

Dijkstra monads. Dijkstra monads [Ahman et al. 2017; Maillard et al. 2019] target code written in monadic form and specified using dependent types. The type-checking process essentially applies weakest-precondition rules and results in the production of proof obligations. In practice, specifications are expressed in first-order logic, so that proof obligations can be discharged using SMT solvers. Dijkstra monads encourage metareasoning using object-language dependent types only; they do not appear to have been designed for, or demonstrated capable of, carrying out inductions over program executions. Dijkstra monads can be instantiated in particular with the nondeterminism monad (NDet). In the current presentation [Ahman et al. 2017], the monad models sets of possible outcomes using finite sets, which rules out infinitely branching nondeterminism and does not allow for abstraction in intermediate postconditions (e.g., asserting that a subterm t_1 returns an even integer).

REFERENCES

- Danel Ahman, Cătălin Hrițcu, Kenji Maillard, Guido Martínez, Gordon Plotkin, Jonathan Protzenko, Aseem Rastogi, and Nikhil Swamy. 2017. Dijkstra Monads for Free. *ACM SIGPLAN Notices* 52, 1 (Jan. 2017), 515–529. <https://doi.org/10.1145/3093333.3009878> <https://doi.org/10.1145/3093333.3009878>.
- K. R. Apt and G. D. Plotkin. 1986. Countable Nondeterminism and Random Assignment. *J. ACM* 33, 4 (Aug. 1986), 724–767. <https://doi.org/10.1145/6490.6494>
- R.J.R. Back. 1983. A continuous semantics for unbounded nondeterminism. *Theoretical Computer Science* 23, 2 (1983), 187–210. [https://doi.org/10.1016/0304-3975\(83\)90055-5](https://doi.org/10.1016/0304-3975(83)90055-5)
- Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Thomas Noll. 2019. Quantitative Separation Logic: A Logic for Reasoning about Probabilistic Pointer Programs. *Proc. ACM Program. Lang.* 3, POPL, Article 34 (Jan. 2019), 29 pages. <https://doi.org/10.1145/3290347>
- Aleksandar Chakarov and Sriram Sankaranarayanan. 2013. Probabilistic Program Analysis with Martingales. In *Computer Aided Verification*, Natasha Sharygina and Helmut Veith (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 511–526. [10.1007/978-3-642-39799-8_34](https://doi.org/10.1007/978-3-642-39799-8_34)
- Arthur Charguéraud. 2013. Pretty-Big-Step Semantics. In *Proceedings of the 22nd European Conference on Programming Languages and Systems (ESOP'13)*. Springer-Verlag, Rome, Italy, 41–60. https://doi.org/10.1007/978-3-642-37036-6_3 https://doi.org/10.1007/978-3-642-37036-6_3.
- Arthur Charguéraud. 2020. Separation Logic for Sequential Programs (Functional Pearl). *Proc. ACM Program. Lang.* 4, ICFP, Article 116 (Aug. 2020), 34 pages. <https://doi.org/10.1145/3408998>
- Patrick Cousot and Radhia Cousot. 1992. Inductive Definitions, Semantics and Abstract Interpretations. In *Proceedings of the 19th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Albuquerque, New Mexico, USA) (POPL '92)*. Association for Computing Machinery, New York, NY, USA, 83–94. <https://doi.org/10.1145/143165.143184>
- Nils Anders Danielsson. 2012. Operational Semantics Using the Partiality Monad. *SIGPLAN Not.* 47, 9 (Sept. 2012), 127–138. <https://doi.org/10.1145/2398856.2364546>
- Edsger W. Dijkstra. 1976. *A Discipline of Programming*. Prentice-Hall. I–XVII, 1–217 pages.
- Andres Erbsen, Samuel Gruetter, Joonwon Choi, Clark Wood, and Adam Chlipala. 2021. Integration Verification Across Software and Hardware for a Simple Embedded System. *PLDI'21 (2021)*. <https://doi.org/10.1145/3453483.3454065>.
- Luis María Ferrer Fioriti and Holger Hermanns. 2015. Probabilistic Termination: Soundness, Completeness, and Compositionality. *SIGPLAN Not.* 50, 1 (Jan. 2015), 489–501. <https://doi.org/10.1145/2775051.2677001>
- Nissim Francez, C. A. R. Hoare, Daniel J. Lehmann, and Willem P. De Roeer. 1979. Semantics of Nondeterminism, Concurrency, and Communication. *J. Comput. System Sci.* 19, 3 (Dec. 1979), 290–308. [https://doi.org/10.1016/0022-0000\(79\)90006-0](https://doi.org/10.1016/0022-0000(79)90006-0) <http://www.sciencedirect.com/science/article/pii/0022000079900060>.
- Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. 2016. Weakest Precondition Reasoning for Expected Run–Times of Probabilistic Programs. In *Programming Languages and Systems*, Peter Thiemann (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 364–389. https://doi.org/10.1007/978-3-662-49498-1_15
- Nicolas Koh, Yao Li, Yishuai Li, Li-yao Xia, Lennart Beringer, Wolf Honoré, William Mansky, Benjamin C. Pierce, and Steve Zdancewicz. 2019. From C to Interaction Trees: Specifying, Verifying, and Testing a Networked Server. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2019)*. Association for Computing Machinery, Cascais, Portugal, 234–248. <https://doi.org/10.1145/3293880.3294106> <https://doi.org/10.1145/3293880.3294106>.
- Xavier Leroy. 2009. A Formally Verified Compiler Back-end. *Journal of Automated Reasoning* 43, 4 (Dec. 2009), 363–446. <https://doi.org/10.1007/s10817-009-9155-4> <http://link.springer.com/10.1007/s10817-009-9155-4>.
- Xavier Leroy and Hervé Grall. 2009. Coinductive big-step operational semantics. *Information and Computation* 207, 2 (2009), 284–304. <https://doi.org/10.1016/j.ic.2007.12.004> Special issue on Structural Operational Semantics (SOS).
- Kenji Maillard, Danel Ahman, Robert Atkey, Guido Martínez, Cătălin Hrițcu, Exequiel Rivas, and Éric Tanter. 2019. Dijkstra Monads for All. *Proceedings of the ACM on Programming Languages* 3, ICFP (July 2019), 104:1–104:29. <https://doi.org/10.1145/3341708> <https://doi.org/10.1145/3341708>.
- Annabelle McIver and Carroll Morgan. 2005. *Abstraction, Refinement And Proof For Probabilistic Systems*. Springer.
- Annabelle McIver, Carroll Morgan, Benjamin Lucien Kaminski, and Joost-Pieter Katoen. 2017. A New Proof Rule for Almost-Sure Termination. *Proc. ACM Program. Lang.* 2, POPL, Article 33 (Dec. 2017), 28 pages. <https://doi.org/10.1145/3158121>
- Robin Milner. 1975. Processes: a mathematical model of computing agents. In *Studies in Logic and the Foundations of Mathematics*. Vol. 80. Elsevier, 157–173.
- Jean-François Monin and Xiaomu Shi. 2013. Handcrafted Inversions Made Operational on Operational Semantics. In *Interactive Theorem Proving*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie (Eds.). Vol. 7998. Springer Berlin Heidelberg, Berlin, Heidelberg, 338–353. https://doi.org/10.1007/978-3-642-39634-2_25 http://link.springer.com/10.1007/978-3-642-39634-2_25.

- Keiko Nakata and Tarmo Uustalu. 2010. Mixed Induction–Coinduction at Work for Coq. *2nd Workshop of Coq users, developers, and contributors* (2010). <http://www.cs.ioc.ee/~keiko/papers/Coq2.pdf>.
- Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press.
- G. D. Plotkin. 1976. A Powerdomain Construction. *Siam J. of Computing* (1976).
- Tiark Rompf and Nada Amin. 2016. Type Soundness for Dependent Object Types (DOT). In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. ACM, Amsterdam Netherlands, 624–641. <https://doi.org/10.1145/2983990.2984008> <https://dl.acm.org/doi/10.1145/2983990.2984008>.
- Steven Schäfer, Sigurd Schneider, and Gert Smolka. 2016. Axiomatic Semantics for Compiler Verification. In *Proceedings of the 5th ACM SIGPLAN Conference on Certified Programs and Proofs*. ACM, St. Petersburg FL USA, 188–196. <https://doi.org/10.1145/2854065.2854083> <https://dl.acm.org/doi/10.1145/2854065.2854083>.
- Jaroslav Ševčík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2013. CompCertTSO: A Verified Compiler for Relaxed-Memory Concurrency. *J. ACM* 60, 3 (June 2013), 1–50. <https://doi.org/10.1145/2487241.2487248> <http://dl.acm.org/citation.cfm?doid=2487241.2487248>.
- Sam Staton, Hongseok Yang, Chris Heunen, Ohad Kammar, and Frank Wood. 2016. Semantics for Probabilistic Programming: Higher-Order Functions, Continuous Distributions, and Soft Constraints. *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science* (July 2016), 525–534. <https://doi.org/10.1145/2933575.2935313> arXiv:1601.04943 <http://arxiv.org/abs/1601.04943>.
- Di Wang, Jan Hoffmann, and Thomas Reps. 2019. A Denotational Semantics for Low-Level Probabilistic Programs with Nondeterminism. *Electronic Notes in Theoretical Computer Science* 347 (2019), 303–324. <https://doi.org/10.1016/j.entcs.2019.09.016> Proceedings of the Thirty-Fifth Conference on the Mathematical Foundations of Programming Semantics.
- Peng Wang, Santiago Cuellar, and Adam Chlipala. 2014. Compiler Verification Meets Cross-Language Linking via Data Abstraction. In *OOPSLA*. ACM Press, 675–690. <https://doi.org/10.1145/2660193.2660201>
- A.K. Wright and M. Felleisen. 1994. A Syntactic Approach to Type Soundness. *Information and Computation* 115, 1 (Nov. 1994), 38–94. <https://doi.org/10.1006/inco.1994.1093> <https://doi.org/10.1006/inco.1994.1093>.
- Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2019. Interaction Trees: Representing Recursive and Impure Programs in Coq. *Proceedings of the ACM on Programming Languages* 4, POPL (Dec. 2019), 51:1–51:32. <https://doi.org/10.1145/3371119> <https://doi.org/10.1145/3371119>.

A ON THE CHALLENGE OF DEFINING WP INDUCTIVELY

The weakest-precondition style reasoning rule for let-bindings is traditionally stated as follows.

$$\text{WP-LET:} \quad \text{wp } t_1 (\lambda v'. \text{wp } ([v'/x] t_2) Q) \vdash \text{wp } (\text{let } x = t_1 \text{ in } t_2) Q.$$

Translating it to a big-step omnisemantics rule results in the following rule.

$$\frac{t_1/s \Downarrow \{(v', s') \mid ([v'/x] t_2)/s' \Downarrow Q\}}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow Q} \text{ OMNI-BIG-LET-CHAINED}$$

The rule OMNI-BIG-LET-CHAINED can be useful for reasoning when one does not want to specify an explicit postcondition that needs to hold between t_1 and t_2 . This *chained* rule can be straightforwardly derived from the OMNI-BIG-LET rule part of the definition of the omni-big-step semantics, by instantiating Q_1 as $\{(v', s') \mid ([v'/x] t_2)/s' \Downarrow Q\}$ in the first premise, then checking the tautology associated with the second premise.

$$\frac{t_1/s \Downarrow Q_1 \quad (\forall (v', s') \in Q_1. ([v'/x] t_2)/s' \Downarrow Q)}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow Q} \text{ OMNI-BIG-LET}$$

One might wonder why we do not use OMNI-BIG-LET-CHAINED directly in the inductively defined rules. The reason is that Coq's *strict positivity* requirement on the well-formedness of inductive definitions does not allow it.

To elaborate on this point, consider the four candidate Coq rules stated below.

Notation "H1 \vdash H2" := ($\forall s, H1 \ s \rightarrow H2 \ s$). (* notation for entailment *)

Inductive wp : trm \rightarrow (val \rightarrow state \rightarrow Prop) \rightarrow (state \rightarrow Prop) :=

| wp_let_invalid : $\forall x \ t_1 \ t_2 \ Q, (* \text{ non strictly positive occurrence of } [wp]. *)$

wp t1 (fun v1 \Rightarrow wp (subst x v1 t2) Q)

\vdash wp (trm_let x t1 t2) Q

| wp_let_invalid' : $\forall Q_1 \ x \ t_1 \ t_2 \ Q \ s, (* \text{ non strictly positive occurrence of } [wp]. *)$

wp t1 Q1 s \rightarrow

Q1 = (fun v1 s2 \Rightarrow wp (subst x v1 t2) Q s2) \rightarrow

wp (trm_let x t1 t2) Q s

| wp_let_valid : $\forall x \ t_1 \ t_2 \ Q, (* \text{ accepted, but with useless induction principle } *)$

(fun s \Rightarrow $\exists Q_1, \text{wp } t_1 \ Q_1 \ s \wedge (\forall v_1, Q_1 \ v_1 \vdash \text{wp } (\text{subst } x \ v_1 \ t_2) \ Q)$)

\vdash wp (trm_let x t1 t2) Q

| wp_let_valid' : $\forall x \ t_1 \ t_2 \ Q_1 \ Q, (* \text{ accepted, with useful induction principle } *)$

wp t1 Q1 s \rightarrow

($\forall v_1 \ s_2, Q_1 \ v_1 \ s_2 \rightarrow \text{wp } (\text{subst } x \ v_1 \ t_2) \ Q \ s_2$) \rightarrow

wp (trm_let x t1 t2) Q s.

The first rule directly translates WP-LET. It is rejected by Coq because it includes a non-strictly-positive occurrence of the predicate wp.

The second rule attempts a reformulation by expanding the definition of entailment, and by introducing a variable name Q1 for the intermediate postcondition, together with an equality constraint on Q1. Yet, Coq rejects this rule just like the previous.

The third rule modifies the first rule by introducing an existentially quantified intermediate postcondition Q1, and quantifies over the items that belong to it. This rule is accepted by Coq. Yet, in that form, Coq (v8.14) generates a useless induction principle, which provides no induction hypothesis for the nested occurrence of wp.

The fourth rule corresponds to OMNI-BIG-LET. It adapts the previous rule by quantifying Q_1 universally at the level of the constructor. This presentation is properly recognized by the induction principle generator of Coq.

B OMNI-BIG-STEP RULES BEYOND A-NORMAL FORM

Throughout the paper, we have presented evaluation rules for terms in A-normal form. We believe that, in the context of program verification and compiler verification, it makes very much sense to consider an intermediate language that imposes the A-normal form restriction. Indeed, doing so considerably decreases the proof effort involved. In this section, we nevertheless discuss how the omni-big-step semantics may be generalized to languages that do not impose A-normal form.

In other words, what would be the statement of an omni-big-step reasoning rule for a general application of the form $(t_1 t_2)$? How would we quantify over the two intermediate postconditions involved, in a similar way as we quantify over the intermediate postcondition Q_1 in the rule OMNI-BIG-LET? In what follows, we present several possible approaches, assuming a left-to-right evaluation order. We discuss the treatment of unspecified evaluation order further on, in [Appendix C](#).

The pretty-big-step approach. One possible solution is to avoid the problem altogether by applying the technique of the pretty-big-step semantics [[Charguéraud 2013](#)], which decomposes rules in such a way that no rule contains more than two evaluation premises, and thus no more than one intermediate postcondition.

A solution for languages featuring pairs. The omni-big-step judgment $t/s \Downarrow Q$ may include as part of its definition the following evaluation rule for pairs.

$$\frac{t_1/s \Downarrow Q_1 \quad (\forall (v_1, s') \in Q_1. t_2/s' \Downarrow \{(v_2, s'') \mid ((v_1, v_2), s'') \in Q\})}{(t_1, t_2)/s \Downarrow Q} \text{ OMNI-BIG-PAIR}$$

For languages that do not include pairs, it is often an option to artificially extend the grammar of terms to include a pair construct.

Assuming the existence of pairs, the rule for applications may be stated as follows.

OMNI-BIG-APP-TERMS-VIA-PAIRS

$$\frac{(t_1, t_2)/s \Downarrow Q_0 \quad (\forall ((v_1, v_2), s') \in Q_0. \exists f x t_3. v_1 = \mu f. \lambda x. t_3 \wedge ([v_1/f] [v_2/x] t_3)/s' \Downarrow Q)}{(t_1 t_2)/s \Downarrow Q}$$

One caveat, thought, is that the implementation of Coq (as of 2021) does not automatically generate appropriate induction principles for premises that involve existential quantifiers and conjunctions. A possible work-around is to add a premise to constrain Q_0 to only include closures in its first component; this way, in the last premise we can replace the conjunction with an implication. This transformation leads to the generation of a useful induction principle.

OMNI-BIG-APP-TERMS-VIA-PAIRS-WITH-INDUCTION-PRINCIPLE

$$\frac{(t_1, t_2)/s \Downarrow Q_0 \quad (\forall ((v_1, v_2), s') \in Q_0. \exists f x t_3. v_1 = \mu f. \lambda x. t_3) \quad (\forall ((v_1, v_2), s') \in Q_0. \forall f x t_3. v_1 = \mu f. \lambda x. t_3 \Rightarrow ([v_1/f] [v_2/x] t_3)/s' \Downarrow Q)}{(t_1 t_2)/s \Downarrow Q}$$

A solution without pairs in the grammar of terms. If the language does not feature pairs and, for some reasons, cannot be extended with pairs, it suffices to consider an auxiliary judgment of the form $t_1, t_2/s \Downarrow Q$. This judgment, which is defined in mutual induction with $t/s \Downarrow Q$, describes the sequential evaluation of two terms t_1 and t_2 , starting in a state s . The postcondition Q describes a set of results, each result being of the form (v_1, v_2, s') . The rules OMNI-BIG-APP-TERMS-VIA-PAIRS and

OMNI-BIG-APP-TERMS-VIA-PAIRS-WITH-INDUCTION-PRINCIPLE can be trivially adapted by using this auxiliary judgment in the first premise. The main drawback of this approach is that the introduction of an auxiliary judgment means that all proofs needs to be carried out by mutual induction.

A solution without pairs and without auxiliary judgment. If we want to avoid the introduction of pairs or of an auxiliary judgment, then we need the statement of the evaluation rule for an application to quantify over intermediate postconditions. The rule shown below describes the evaluation of such an application $(t_1 t_2)$.

OMNI-BIG-APP-TERMS

$$\frac{t_1/s \Downarrow Q_1 \quad \left(\forall (v_1, s') \in Q_1. \exists Q_2. t_2/s' \Downarrow Q_2 \wedge \left(\forall (v_2, s'') \in Q_2. \exists f x t_3. v_1 = \mu f. \lambda x. t_3 \wedge ([v_1/f][v_2/x] t_3)/s'' \Downarrow Q \right) \right)}{(t_1 t_2)/s \Downarrow Q}$$

For the same reasons as before, Coq does not generate a usable induction principle. A work-around that may be applied here consists of pulling the existential quantifiers on Q_2 at the top of the rule, by making Q_2 be a function of s' . Doing so allows us to split the conjuncts, although at the cost of duplicating a few quantifiers.

OMNI-BIG-APP-TERMS-WITH-INDUCTION-PRINCIPLE

$$\frac{t_1/s \Downarrow Q_1 \quad \left(\forall (v_1, s') \in Q_1. t_2/s' \Downarrow Q_2(s') \right) \quad \left(\forall s'. \forall (v_1, s'') \in Q_2(s'). \exists f x t_3. v_1 = \mu f. \lambda x. t_3 \right) \quad \left(\forall (v_1, s') \in Q_1. \forall (v_2, s'') \in Q_2(s'). \forall f x t_3. v_1 = \mu f. \lambda x. t_3 \Rightarrow ([v_1/f][v_2/x] t_3)/s'' \Downarrow Q \right)}{(t_1 t_2)/s \Downarrow Q}$$

Remark: to prove the reformulated rule OMNI-BIG-APP-TERMS-WITH-INDUCTION-PRINCIPLE equivalent to the original rule OMNI-BIG-APP-TERMS, we need to exploit classical logic and the axiom of choice.

The encodings at play in OMNI-BIG-APP-TERMS-WITH-INDUCTION-PRINCIPLE appear a little bit involved. For this reason, we would recommend in practice to follow any of the other possible approaches: (1) assume A-normal form, possibly by A-normalizing programs once and for all in the front-end of the verified tool chain; or (2) exploit the pretty-big-step technique to evaluate subterms one at a time; or (3) assume a language with pairs, or extend the language with pairs; or (4) introduce an auxiliary judgment for evaluating a pair of terms.

C UNSPECIFIED EVALUATION ORDER

For a language that does not specify the order of evaluation for arguments of, e.g., pairs, or applications, we can consider a generalized version of the rule OMNI-BIG-PAIR from the previous section. Essentially, we duplicate the premises to account for the two possible evaluation orders.

OMNI-BIG-PAIR-UNSPECIFIED-ORDER

$$\frac{t_1/s \Downarrow Q_1 \quad \left(\forall (v_1, s') \in Q_1. t_2/s' \Downarrow \{(v_2, s'') \mid ((v_1, v_2), s'') \in Q\} \right) \quad t_2/s \Downarrow Q_2 \quad \left(\forall (v_2, s') \in Q_2. t_1/s' \Downarrow \{(v_1, s'') \mid ((v_1, v_2), s'') \in Q\} \right)}{(t_1 t_2)/s \Downarrow Q}$$

To avoid the duplication in the premises, one can follow the approach described in §5.5 of the paper on the pretty-big-step semantics [Charguéraud 2013], which presents a general rule for evaluating a list of subterms in arbitrary order.

D OMNISEMANTICS RULES IN THE PRESENCE OF EXCEPTIONS

For a programming language that features exceptions, the reasoning rule for let-bindings needs to be adapted in two ways. Indeed, if the body of the let-binding raises an exception, then the continuation should not be evaluated. Moreover, the exception raised should be included in the set of results that the let-binding can produce.

There are two ways to extend the grammar of results with exceptions. The first possibility is to add a constructor to the grammar of values. In this case, the postcondition Q remains a predicate over pairs of values and states. The second possibility is to introduce a type, to capture the sum of the type of values and of the type of exceptions. In that case, the postcondition Q becomes a predicate over pairs of results and states.

For simplicity, let us assume in what follows that the grammar of values includes a constant exception construct, written exn . In that setting, the omni-big-step evaluation rule for a let-binding of the form $(\text{let } x = t_1 \text{ in } t_2)$ can be stated as follows. The first premise describes the evaluation of t_1 . The second premise handles the case where t_1 produces a normal value. The third premise handles the case where t_1 produces an exception.

OMNI-BIG-LET-WITH-EXCEPTIONS

$$\frac{t_1/s \Downarrow Q_1 \quad \left(\forall (v', s') \in Q_1. v' \neq \text{exn} \Rightarrow ([v'/x] t_2)/s' \Downarrow Q \right) \quad \left(\forall s'. Q_1 \text{ exn } s' \Rightarrow Q \text{ exn } s' \right)}{(\text{let } x = t_1 \text{ in } t_2)/s \Downarrow Q}$$

We proved in Coq the equivalence of this treatment of exceptions with the formalization of exceptions expressed both in standard small-step and in standard big-step semantics.

E AN ALTERNATIVE DEFINITION FOR THE EVENTUALLY JUDGMENT

In §3.2 of the paper, we define the *eventually judgment* in terms of the omni-small-step judgment. This judgment, written $t/s \longrightarrow^\diamond P$, captures the property that every possible evaluation of t/s is safe and eventually reaches a configuration in the set P . Here, we present an alternative definition, expressed in terms of the standard small-step judgment.

This alternative definition involves the two rules shown below. The first rule asserts, as before, that the judgment $t/s \longrightarrow^\diamond P$ is satisfied if the configuration t/s belongs to P . The second rule asserts that the judgment is satisfied if the configuration t/s is not stuck and if, for any configuration t'/s' that it may reduce to—with respect to the small-step evaluation judgments—, the predicate $t'/s' \longrightarrow^\diamond P$ holds.

EVENTUALLY-HERE

$$\frac{(t, s) \in P}{t/s \longrightarrow^\diamond P}$$

EVENTUALLY-STEP-USING-STANDARD-SMALL-STEP

$$\frac{\left(\exists t's'. t/s \longrightarrow t'/s' \right) \quad \left(\forall t's'. (t/s \longrightarrow t'/s') \Rightarrow (t'/s' \longrightarrow^\diamond P) \right)}{t/s \longrightarrow^\diamond P}$$

This alternative definition might be of interest if one already has a small-step judgment at hand and does not wish to state an omni-small-step judgment. We argue, however, that the original definition presented in the paper is both more concise and more practical for carrying out proofs.

F DEFINITION OF THE TERMINATION JUDGMENT

We introduced the termination judgment to formalize the interpretation of the omni-big-step judgment (§2.2, OMNI-BIG-STEP-IFF-TERMINATES-AND-CORRECT). The predicate $\text{terminates}(t, s)$ asserts that all executions of configuration t/s terminate. In this section, we present two formal definitions of this predicate, one in small-step style and one in big-step style.

The small-step version is inductively defined by the two rules show below. This definition corresponds to a specialized version of the alternative definition for the judgment EVENTUALLY,

presented in [Appendix E](#), with the postcondition instantiated as the always-true predicate.

$$\begin{array}{c}
 \text{SMALL-TERMINATES-HERE} \\
 \hline
 \text{terminates}(v, s)
 \end{array}
 \qquad
 \begin{array}{c}
 \text{SMALL-TERMINATES-STEP} \\
 (\exists t's'. t/s \longrightarrow t'/s') \\
 (\forall t's'. (t/s \longrightarrow t'/s') \Rightarrow \text{terminates}(t', s')) \\
 \hline
 \text{terminates}(t, s)
 \end{array}$$

The big-step version is inductively defined using one rule per language construct. We show below the rules for values and for let-bindings. This definition corresponds to an inductive version of the coinductive judgment safe from [Wang et al. \[2014\]](#), described in [§7](#).

$$\begin{array}{c}
 \text{BIG-TERMINATES-VAL} \\
 \hline
 \text{terminates}(v, s)
 \end{array}
 \qquad
 \begin{array}{c}
 \text{BIG-TERMINATES-LET} \\
 \text{terminates}(t_1, s) \\
 (\forall v_1s'. (t_1/s \Downarrow v_1/s') \Rightarrow \text{terminates}(\llbracket v_1/x \rrbracket t_2, s')) \\
 \hline
 \text{terminates}(\text{let } x = t_1 \text{ in } t_2, s)
 \end{array}$$

G DEFINITION OF THE TYPING JUDGMENT

This section states the typing rules for the state-free language considered in [§4.1](#). The typing rules are given for terms in A-normal form. The judgment $\vdash v : T$ asserts that the closed value t admits the type T . The judgment $E \vdash t : T$ asserts that the term t admits type T in the environment E . Finally, \mathbb{V} denotes the set of terms that are either values or variables.

$$\begin{array}{c}
 \text{VTYP-UNIT} \\
 \hline
 \vdash \# : \text{unit}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VTYP-BOOL} \\
 \hline
 \vdash b : \text{bool}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VTYP-INT} \\
 \hline
 \vdash n : \text{int}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VTYP-FIX} \\
 f : (T_1 \rightarrow T_2), x : T_1 \vdash t : T_2 \\
 \hline
 \vdash ((\mu f. \lambda x. t)) : (T_1 \rightarrow T_2)
 \end{array}$$

$$\begin{array}{c}
 \text{TYP-VAL} \\
 \vdash v : T \\
 \hline
 E \vdash v : T
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TYP-VAR} \\
 x \in \text{dom } E \quad E[x] = T \\
 \hline
 E \vdash x : T
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TYP-FIX} \\
 E, f : (T_1 \rightarrow T_2), x : T_1 \vdash t : T_2 \\
 \hline
 E \vdash (\mu f. \lambda x. t) : T
 \end{array}$$

$$\begin{array}{c}
 \text{TYP-APP} \\
 E \vdash t_1 : (T_1 \rightarrow T_2) \quad E \vdash t_2 : T_1 \quad t_1, t_2 \in \mathbb{V} \\
 \hline
 E \vdash (t_1 t_2) : T_2
 \end{array}$$

$$\begin{array}{c}
 \text{TYP-IF} \\
 E \vdash t_0 : \text{bool} \quad E \vdash t_1 : T \quad E \vdash t_2 : T \quad t_0 \in \mathbb{V} \\
 \hline
 E \vdash (\text{if } t_0 \text{ then } t_1 \text{ else } t_2) : T
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TYP-LET} \\
 E \vdash t_1 : T_1 \quad E, x : T_1 \vdash t_2 : T_2 \\
 \hline
 E \vdash (\text{let } x = t_1 \text{ in } t_2) : T_2
 \end{array}$$

$$\begin{array}{c}
 \text{TYP-ADD} \\
 E \vdash t_1 : \text{int} \quad E \vdash t_2 : \text{int} \quad t_1, t_2 \in \mathbb{V} \\
 \hline
 E \vdash (\text{add } t_1 t_2) : \text{int}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TYP-RAND} \\
 E \vdash t_1 : \text{int} \quad t_1 \in \mathbb{V} \\
 \hline
 E \vdash (\text{rand } t_1) : \text{int}
 \end{array}$$

H EXTENSION OF THE TYPING JUDGMENT FOR STATE

This section states the typing rules for the imperative language considered in [§4.2](#). There, the typing judgment for terms takes the form $S; E \vdash t : T$ and the typing judgment for closed values takes the form $S \vdash v : T$, where the store typing S maps locations to types. The rules from the previous appendix are extended to simply thread S throughout the judgment. The new rules include the rule

for typing locations, and the rules for memory operations. They are shown next.

$$\begin{array}{c}
\text{VTYP-LOC} \\
\frac{p \in \text{dom } S \quad S[p] = T}{S \vdash p : (\text{ref } T)}
\end{array}
\qquad
\begin{array}{c}
\text{TYP-REF} \\
\frac{S; E \vdash t_1 : T \quad t_1 \in \mathbb{V}}{S; E \vdash (\text{ref } t_1) : (\text{ref } T)}
\end{array}$$

$$\begin{array}{c}
\text{TYP-GET} \\
\frac{S; E \vdash t_1 : (\text{ref } T) \quad t_1 \in \mathbb{V}}{S; E \vdash (\text{get } t_1) : T}
\end{array}
\qquad
\begin{array}{c}
\text{TYP-SET} \\
\frac{S; E \vdash t_1 : (\text{ref } T) \quad S; E \vdash t_2 : T \quad t_1, t_2 \in \mathbb{V}}{S; E \vdash (\text{set } t_1 t_2) : \text{unit}}
\end{array}$$

I DEFINITION OF THE STANDARD SMALL-STEP JUDGMENT

In §2.4, we gave a characterization of coinductive omni-big-step semantics in terms of the standard small-step semantics, written $t/s \longrightarrow t'/s'$. For reference, we give below the rules that define the standard small-step judgment:

$$\begin{array}{c}
\text{SMALL-APP} \\
\frac{v_1 = (\mu f. \lambda x. t)}{(v_1 v_2)/s \longrightarrow ([v_2/x] [v_1/f] t)/s}
\end{array}
\qquad
\begin{array}{c}
\text{SMALL-IF-TRUE} \\
\frac{}{(\text{if true then } t_1 \text{ else } t_2)/s \longrightarrow t_1/s}
\end{array}$$

$$\begin{array}{c}
\text{SMALL-IF-FALSE} \\
\frac{}{(\text{if false then } t_1 \text{ else } t_2)/s \longrightarrow t_2/s}
\end{array}
\qquad
\begin{array}{c}
\text{SMALL-LET-CTX} \\
\frac{t_1/s \longrightarrow t'_1/s'}{(\text{let } x = t_1 \text{ in } t_2)/s \longrightarrow (\text{let } x = t'_1 \text{ in } t_2)/s'}
\end{array}$$

$$\begin{array}{c}
\text{SMALL-LET-VAL} \\
\frac{}{(\text{let } x = v_1 \text{ in } t_2)/s \longrightarrow ([v_1/x] t_2)/s}
\end{array}
\qquad
\begin{array}{c}
\text{SMALL-ADD} \\
\frac{}{(\text{add } n_1 n_2)/s \longrightarrow (n_1 + n_2)/s}
\end{array}
\qquad
\begin{array}{c}
\text{SMALL-RAND} \\
\frac{0 \leq m < n}{(\text{rand } n)/s \longrightarrow m/s}
\end{array}$$

$$\begin{array}{c}
\text{SMALL-REF} \\
\frac{p \notin \text{dom } s}{(\text{ref } v)/s \longrightarrow (s[p := v])/s}
\end{array}
\qquad
\begin{array}{c}
\text{SMALL-FREE} \\
\frac{p \in \text{dom } s}{(\text{free } p)/s \longrightarrow t/(s \setminus p)}
\end{array}$$

$$\begin{array}{c}
\text{SMALL-GET} \\
\frac{p \in \text{dom } s}{(\text{get } p)/s \longrightarrow (s[p])/s}
\end{array}
\qquad
\begin{array}{c}
\text{SMALL-SET} \\
\frac{p \in \text{dom } s}{(\text{set } p v)/s \longrightarrow t/(s[p := v])}
\end{array}$$

J EVALUATION OF UNARY AND BINARY OPERATORS

The following definitions complete the semantics described in the case study “compiling immutable pairs to heap-allocated records” (§6.4).

$$\frac{}{\text{evalunop}(\text{fst}, (v_1, v_2), v_1)}
\qquad
\frac{}{\text{evalunop}(\text{snd}, (v_1, v_2), v_2)}
\qquad
\frac{}{\text{evalunop}(\text{not}, 1, 0)}$$

$$\frac{}{\text{evalunop}(\text{not}, 0, 1)}
\qquad
\frac{}{\text{evalbinop}(+, n_1, n_2, n_1 + n_2)}
\qquad
\frac{}{\text{evalbinop}(\text{mkpair}, v_1, v_2, (v_1, v_2))}$$