

PROOFS AS STATEFUL PROGRAMS: A FIRST-ORDER LOGIC WITH ABSTRACT HOARE TRIPLES, AND AN INTERPRETATION INTO AN IMPERATIVE LANGUAGE

THOMAS POWELL 

Department of Computer Science, University of Bath
e-mail address: trjp20@bath.ac.uk

ABSTRACT. We introduce an extension of first-order logic that comes equipped with additional predicates for reasoning about an abstract state. Sequents in the logic comprise a main formula together with pre- and postconditions in the style of Hoare logic, and the axioms and rules of the logic ensure that the assertions about the state compose in the correct way. The main result of the paper is a realizability interpretation of our logic that extracts programs into a mixed functional/imperative language. All programs expressible in this language act on the state in a sequential manner, and we make this intuition precise by interpreting them in a semantic metatheory using the state monad. Our basic framework is very general, and our intention is that it can be instantiated and extended in a variety of different ways. We outline in detail one such extension: A monadic version of Heyting arithmetic with a wellfounded while rule, and conclude by outlining several other directions for future work.

1. INTRODUCTION

The Curry-Howard correspondence lies at the heart of theoretical computer science. Over the years, a multitude of different techniques for extracting programs from proofs have been developed, the majority of which translate formal proof systems into lambda calculi. As such, programs extracted from proofs are typically conceived as pure functional programs.

Everyday programmers, on the other hand, often think and write in an imperative paradigm, in terms of instructions that change some underlying global state. This is reinforced by the fact that many of the most popular programming languages, including C and Python, lean towards this style. Imperative programs are nevertheless highly complex from a mathematical perspective, and while systems such as Hoare logic [Hoa69] or separation logic [Rey02] have been designed to reason about them, the formal extraction of imperative programs from proofs has received comparatively little attention.

In this paper, we propose a new idea in this direction, developing a formal system SL that enriches ordinary first-order logic with Hoare triples for reasoning about an abstract global state. Sequents will have the form $\Gamma \vdash \{\alpha \cdot A \cdot \beta\}$, where A is a formula and α, β assertions about the state, and proofs in the logic will include both ordinary introduction and elimination rules for predicate logic, together with special rules for reasoning about the

Key words and phrases: Program extraction, modified realizability, imperative programs, state monad.

state. We then construct a stateful realizability interpretation (based on Kreisel’s modified realizability [Kre59]) that relates formulas in SL to terms in a mixed functional/imperative language ST. Our main result is a soundness theorem, which confirms that whenever a formula is provable in SL, we can extract a corresponding stateful realizing term in ST. While our initial soundness theorem focuses on pure predicate logic, we subsequently show that it can be extended to arithmetic, where in particular we are then able to extract programs that contain both recursion and controlled while loops.

We are not the first to adapt traditional methods to extract imperative programs: A major achievement in this direction, for example, is the monograph [PCW05], which sets up a variant of intuitionistic Hoare logic alongside a realizability translation into a standard imperative language. Other relevant examples include [Atk09, CMM⁺09, DF89, Fil03, NAMB07, YHB07]. However, these and almost all other prior work in this direction tend to focus on formal verification, with an eye towards using proof interpretations as a method for the synthesis of correct-by-construction software. In concrete terms, this means that the formal systems tend to be quite detailed and oriented towards program analysis, while the starting point is typically a program for which we want to construct a verification proof, rather than a proof from which we hope to extract a potentially unfamiliar program.

Our approach, on the other hand, is much more abstract, with an emphasis on potential applications in logic and proof theory. Our basic system SL makes almost no assumptions about the structure of the state and what we are allowed to do with it. Rather, we focus on producing a general framework for reasoning about ‘stateful formulas’, which can then be instantiated with additional axioms to model concrete scenarios. The simplicity and generality of our framework is its most important feature, and we consider this work to be a first step towards a number of potentially interesting applications. For this reason, we include not only an extension of our system to a monadic theory of arithmetic, but conclude by sketching out some additional ways in which we conjecture that our logic and interpretation could be used and expanded, including the computational semantics of proofs and probabilistic logic.

We take ideas from three main sources. The first is a case study of Berger et al. [BSW14], in which a realizability interpretation is used to extract a version of in-place quicksort, and where the imperative nature of the extracted program is presented in a semantic way using the state monad. While their program behaves imperatively “by-chance”, terms extracted from our logic are forced to be imperative, and thus our framework offers one potential solution to their open problem of designing a proof calculus which only yields imperative programs. Indeed, an implementation of the insert sort algorithm is formally extracted in Section 6 below. Our second source of inspiration is the thesis of Birolo [Bir12], where a general monadic realizability interpretation is defined and then used to give an alternative, semantic presentation of learning-based interactive realizability [Asc11, Bd09]. However, our work goes beyond this in that it also involves a monadic extension of the target logic, whereas Birolo’s applies to standard first-order logic. Finally, a number of ideas are taken from the author’s previous work [Pow18] on extracting stateful programs using the Dialectica interpretation. While there the state is used in a very specific and restricted way, here we use an analogous call-by-value monadic translation on terms.

It is important to stress that we do not claim that our work represents an optimal or complete method for extracting imperative programs from proofs, nor do we claim that it is superior to alternative methods, including the aforementioned works in the direction of verification, or, for instance, techniques based on Krivine’s classical realizability [Kri09],

which could be viewed as imperative in nature. We simply offer what we consider to be a new and interesting perspective that emphasises abstraction and simplicity, and propose that our framework could prove valuable in a number of different contexts.

Overview of the paper. The main technical work that follows involves the design of three different systems, a realizability interpretation that connects them, and an instantiation of this framework in the setting of first-order arithmetic, namely:

- A novel extension SL of predicate logic with abstract Hoare triples, which can be extended with additional axioms for characterising the state (Section 2).
- A standard calculus ST for lambda terms with imperative commands, which can again be extended with additional constants for interacting with the state (Section 3).
- A metalanguage S^ω into which both SL and ST can be embedded (Section 4), which is used to formulate the realizability relation and prove its soundness (Section 5).
- An instantiation of SL as a theory of arithmetic, with programs extracted into an extension of ST with recursion and while loops (Section 6).

Concrete examples are given, and potential applications surveyed in Section 7.

2. THE SYSTEM SL: FIRST-ORDER LOGIC WITH STATE

We begin by introducing our target theory SL from which stateful programs will be extracted. This is an extension of ordinary first-order logic in the sense that the latter can always be embedded into SL (we will make this precise in Proposition 2.1 below). Ultimately, we are interested not so much in SL on its own, but in theories of the form $SL + \Delta_H + \Delta_S$, where Δ_H and Δ_S are collections of (respectively non-computational and computational) axioms that together characterise the state. Several concrete examples will be given to illustrate this, and in Section 6 we present a variant of SL that represents a theory of first-order arithmetic with state.

Before defining SL, we give a standard presentation of first-order intuitionistic predicate logic PL, which serves as an opportunity to fix our basic style of formal reasoning. The language of PL consists of the logical constants $\wedge, \vee, \Rightarrow, \forall, \exists, \top, \perp$, variables x, y, z, \dots , along with function symbols f, g, h, \dots and predicate symbols P, Q, R, \dots , each with a fixed arity. We assume the existence of at least one constant c . Terms are built from variables and function symbols as usual, and formulas are built from prime formulas $P(t_1, \dots, t_n)$, \top and \perp using the logical constants. We use the usual abbreviation $\neg A := A \Rightarrow \perp$. We work in a sequent style natural deduction calculus, where sequents have the form $\Gamma \vdash_I A$ for some context Γ and formula A , and a context is a set of labelled assumptions of the form $A_1^{u_1}, \dots, A_n^{u_n}$ for pairwise distinct labels u_i . The axioms and rules of PL are as in Figure 1.

2.1. Stateful first-order logic. We now define our new logical system SL, which is an extension of ordinary first-order logic with new state propositions. To be more precise, we extend the language of PL with a ternary operation $\{- \cdot - \cdot -\}$, together with special state predicate symbols p, q, r, \dots , which also have a fixed arity. Terms of SL are the same as those of PL. On the other hand, there are two kinds of formulas in SL: state formulas and main formulas. A *state formula* is defined using state predicate symbols and propositional connectives as follows:

Figure 1: Axioms and rules of PL

Propositional logic

$$\Gamma \vdash_I A \quad \text{if } A^u \in \Gamma \text{ for some } u \quad \Gamma \vdash_I \top$$

$$\frac{\Gamma \vdash_I A \quad \Gamma \vdash_I B}{\Gamma \vdash_I A \wedge B} \wedge I \quad \frac{\Gamma \vdash_I A \wedge B}{\Gamma \vdash_I A} \wedge E_L \quad \frac{\Gamma \vdash_I A \wedge B}{\Gamma \vdash_I B} \wedge E_R$$

$$\frac{\Gamma \vdash_I A}{\Gamma \vdash_I A \vee B} \vee I_L \quad \frac{\Gamma \vdash_I B}{\Gamma \vdash_I A \vee B} \vee I_R \quad \frac{\Gamma \vdash_I A \vee B \quad \Gamma, A^u \vdash_I C \quad \Gamma, B^v \vdash_I C}{\Gamma \vdash_I C} \vee E$$

$$\frac{\Gamma, A^u \vdash_I B}{\Gamma \vdash_I A \Rightarrow B} \Rightarrow I \quad \frac{\Gamma \vdash_I A \Rightarrow B \quad \Gamma \vdash_I A}{\Gamma \vdash_I B} \Rightarrow E \quad \frac{\Gamma \vdash_I \perp}{\Gamma \vdash_I A} \perp E$$

Quantifier rules

$$\frac{\Gamma \vdash_I A[y/x]}{\Gamma \vdash_I \forall x A} \forall I \quad \frac{\Gamma \vdash_I \forall x A}{\Gamma \vdash_I A[t/x]} \forall E$$

$$\frac{\Gamma \vdash_I A[t/x]}{\Gamma \vdash_I \exists x A} \exists I \quad \frac{\Gamma \vdash_I \exists x A \quad \Gamma, A[y/x]^u \vdash_I C}{\Gamma \vdash_I C} \exists E$$

for $\forall I$, $y \equiv x$ or y not free in A , and y not free in Γ
for $\exists E$, $y \equiv x$ or y not free in A , and y not free in C or Γ .

- \top and \perp are state formulas,
- if p a state predicate symbol of arity n and t_1, \dots, t_n are terms, then $p(t_1, \dots, t_n)$ is a state formula,
- if α, β are state formulas, so are $\alpha \wedge \beta$, $\alpha \vee \beta$, $\alpha \Rightarrow \beta$.

A *main formula* (or just *formula*) of SL is now defined as:

- \top and \perp are formulas,
- if P is an ordinary predicate symbol of arity n and t_1, \dots, t_n are terms, then $P(t_1, \dots, t_n)$ is a formula,
- if A, B are formulas, so are $A \wedge B$, $A \vee B$ and $\exists x A$,
- if A, B are formulas and α, β state formulas, then $A \Rightarrow \{\alpha \cdot B \cdot \beta\}$ and $\forall x \{\alpha \cdot A \cdot \beta\}$ are formulas.

The notions of free and bound variables, along with substitution $\alpha[t/x]$ and $A[t/x]$ can be easily defined for both state and main formulas.

Analogous to the construction of formulas, our basic proof system uses the auxiliary notion of a *state proof* in order to define a main proof. A *state sequent* has the form $\Gamma \vdash_H \alpha$ where α is a state formula and Γ a set of labelled state formulas. A proof of $\Gamma \vdash_H \alpha$ in SL is built from the axioms and rules of *classical propositional logic* i.e. the propositional axioms

and rules as set out in Figure 1 plus the law of excluded middle $\Gamma \vdash_H \alpha \vee \neg\alpha$, together with a set Δ_H of as yet unspecified state axioms of the form $\Gamma \vdash_H \alpha$.

A main sequent of SL has the form $\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}$, where A is a formula and α, β state formulas, and Γ is a set of labelled main formulas. A proof of $\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}$ in SL uses the axioms and rules given in Figure 2, together with a set Δ_S of additional axioms.

We now make precise what we mean when we characterise SL as an extension of standard first-order logic. The following is provable with an easy induction over derivations in PL:

Proposition 2.1. *For any formula A of PL and state formula α , define the main formula A_α of SL by*

- $Q_\alpha := Q$ for Q atomic,
- $(A \wedge B)_\alpha := A_\alpha \wedge B_\alpha$, $(A \vee B)_\alpha := A_\alpha \vee B_\alpha$ and $(\exists x A)_\alpha := \exists x A_\alpha$,
- $(A \Rightarrow B)_\alpha := A_\alpha \Rightarrow \{\alpha \cdot B_\alpha \cdot \alpha\}$ and $(\forall x A)_\alpha := \forall x \{\alpha \cdot A_\alpha \cdot \alpha\}$.

Then whenever $\Gamma \vdash_I A$ is provable in PL, we have that $\Gamma_\alpha, \Delta \vdash_S \{\alpha \cdot A_\alpha \cdot \alpha\}$ is provable in SL, where Δ is arbitrary and $\Gamma_\alpha := (A_1)_\alpha^{u_1}, \dots, (A_n)_\alpha^{u_n}$ for $\Gamma := A_1^{u_1}, \dots, A_n^{u_n}$.

2.2. The intuition behind SL. The intended semantic meaning of $\Gamma \vdash_H \alpha$ is that α can be inferred from the assumptions Γ for any fixed state. More specifically, if we imagine a semantic variant $[\alpha](\pi)$ of each state formula where now the dependency on an underlying state π is made explicit, the semantics of $\Gamma \vdash_H \alpha$ is just

$$[\Gamma](\pi) \Rightarrow [\alpha](\pi)$$

On the other hand, the intended meaning of $\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}$ is that from assumptions Γ , if α holds with respect to some initial state, then we can infer that A is true and β holds with respect to some modified state, or more precisely:

$$[\Gamma] \Rightarrow (\exists \pi [\alpha](\pi) \Rightarrow ([A] \wedge \exists \pi' [\beta](\pi'))) \quad (2.1)$$

In particular, the computational interpretation of (2.1) above will be a program that takes some input state π satisfying $[\alpha](\pi)$ and returns a realizer-state pair $\langle x, \pi' \rangle$ such that x realizes A and $[\beta](\pi')$ holds.

Our semantic interpretation $[\cdot]$ will be properly defined in Section 4. Crucially, in SL the state is *implicit*, and so there are no variables or terms of state type. The state will rather be made explicit in our metatheory S^ω . The main axioms and rules of SL simply describe how this semantic interpretation propagates in a call-by-value manner through the usual axioms and rules of first-order logic. The state itself is brought into play through the Hoare rules along with the additional axioms Δ_H and Δ_S .

To give a more detailed explanation, consider the introduction rule $\wedge_S I$. Semantically, the idea is that if

$$\exists \pi [\alpha](\pi) \Rightarrow ([A] \wedge \exists \pi_1 [\beta](\pi_1)) \quad \text{and} \quad \exists \pi_1 [\beta](\pi_1) \Rightarrow ([B] \wedge \exists \pi_2 [\gamma](\pi_2))$$

both hold, then we can infer

$$\exists \pi [\alpha](\pi) \Rightarrow ([A] \wedge [B] \wedge \exists \pi_2 [\gamma](\pi_2))$$

and this can be regarded as a ‘stateful’ version of the usual conjunction introduction rule. In particular, we note that this is no longer symmetric: Informally speaking, $[A]$ is ‘proven first’, followed by $[B]$. Under our realizability interpretation, this rule will correspond to realizer for $\{\alpha \cdot A \cdot \beta\}$ being sequentially composed with a realizer for $\{\beta \cdot B \cdot \gamma\}$.

Figure 2: Axioms and rules of SL

Propositional axioms and rules

$$\begin{array}{c}
\Gamma \vdash_S \{\alpha \cdot A \cdot \alpha\} \quad \text{if } A^u \in \Gamma \text{ for some } u \quad \Gamma \vdash_S \{\alpha \cdot \top \cdot \alpha\} \\
\\
\frac{\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\} \quad \Gamma \vdash_S \{\beta \cdot B \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot A \wedge B \cdot \gamma\}} \wedge_S I \\
\\
\frac{\Gamma \vdash_S \{\alpha \cdot A \wedge B \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}} \wedge_S E_L \quad \frac{\Gamma \vdash_S \{\alpha \cdot A \wedge B \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot B \cdot \beta\}} \wedge_S E_R \\
\\
\frac{\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot A \vee B \cdot \beta\}} \vee_S I_L \quad \frac{\Gamma \vdash_S \{\alpha \cdot B \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot A \vee B \cdot \beta\}} \vee_S I_R \\
\\
\frac{\Gamma \vdash_S \{\alpha \cdot A \vee B \cdot \beta\} \quad \Gamma, A^u \vdash_S \{\beta \cdot C \cdot \gamma\} \quad \Gamma, B^v \vdash_S \{\beta \cdot C \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot C \cdot \gamma\}} \vee_S E \\
\\
\frac{\Gamma, A^u \vdash_S \{\alpha \cdot B \cdot \beta\}}{\Gamma \vdash_S \{\gamma \cdot A \Rightarrow \{\alpha \cdot B \cdot \beta\} \cdot \gamma\}} \Rightarrow_S I \quad \frac{\Gamma \vdash_S \{\alpha \cdot A \Rightarrow \{\gamma \cdot B \cdot \delta\} \cdot \beta\} \quad \Gamma \vdash_S \{\beta \cdot A \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot B \cdot \delta\}} \Rightarrow_S E \\
\\
\frac{\Gamma \vdash_S \{\alpha \cdot \perp \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot A \cdot \gamma\}} \perp_S E \quad \text{for } \gamma \text{ arbitrary}
\end{array}$$

Quantifier rules

$$\begin{array}{c}
\frac{\Gamma \vdash_S \{\alpha[y/x] \cdot A[y/x] \cdot \beta[y/x]\}}{\Gamma \vdash_S \{\gamma \cdot \forall x \{\alpha \cdot A \cdot \beta\} \cdot \gamma\}} \forall_S I \quad \frac{\Gamma \vdash_S \{\alpha \cdot \forall x \{\beta \cdot A \cdot \gamma\} \cdot \beta[t/x]\}}{\Gamma \vdash_S \{\alpha \cdot A[t/x] \cdot \gamma[t/x]\}} \forall_S E \\
\\
\frac{\Gamma \vdash_S \{\alpha \cdot A[t/x] \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot \exists x A \cdot \beta\}} \exists_S I \quad \frac{\Gamma \vdash_S \{\alpha \cdot \exists x A \cdot \beta\} \quad \Gamma, A[y/x]^u \vdash_S \{\beta \cdot C \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot C \cdot \gamma\}} \exists_S E
\end{array}$$

for $\forall_S I$, $y \equiv x$ or y not free in A , α , β , and y not free in Γ
for $\exists_S E$, $y \equiv x$ or y not free in A , and y not free in C , α , β , γ or Γ .

Basic Hoare rules

$$\begin{array}{c}
\frac{\alpha \vdash_H \beta \quad \Gamma \vdash_S \{\beta \cdot A \cdot \gamma\} \quad \gamma \vdash_H \delta}{\Gamma \vdash_S \{\alpha \cdot A \cdot \delta\}} \text{cons} \\
\\
\frac{\vdash_H \alpha \vee \beta \quad \Gamma \vdash_S \{\alpha \wedge \gamma \cdot A \cdot \delta\} \quad \Gamma \vdash_S \{\beta \wedge \gamma \cdot A \cdot \delta\}}{\Gamma \vdash_S \{\gamma \cdot A \cdot \delta\}} \text{cond}
\end{array}$$

Additional axioms

state axioms Δ_H of the form $\Gamma \vdash_H \alpha$
main axioms Δ_S of the form $\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}$

Similarly, for the stateful quantifier rule $\forall_S I$, the premise $\{\alpha[y/x] \cdot A[y/x] \cdot \beta[y/x]\}$ corresponds semantically to

$$\exists \pi [\alpha[y/x]](\pi) \Rightarrow [A[y/x]] \wedge \exists \pi' [\beta[y/x]](\pi')$$

and so by the standard quantifier rule of ordinary predicate logic we would have

$$\forall x (\exists \pi [\alpha](\pi) \Rightarrow [A] \wedge \exists \pi' [\beta](\pi'))$$

From this we can directly infer

$$\exists \pi [\gamma](\pi) \Rightarrow \forall x (\exists \pi [\alpha](\pi) \Rightarrow [A] \wedge \exists \pi' [\beta](\pi')) \wedge \exists \pi [\gamma](\pi)$$

for arbitrary γ , and the above now this has the right form, namely the semantic interpretation of $\{\gamma \cdot \forall x \{\alpha \cdot A \cdot \beta\} \cdot \gamma\}$.

We note that the traditional rules of Hoare logic correspond to certain simple cases of our rules. For example, the composition rule could be viewed as special case of $\wedge_S I$ for $A = B = \top$, more specifically the derivation:

$$\frac{\Gamma \vdash_S \{\alpha \cdot \top \cdot \beta\} \quad \Gamma \vdash_S \{\beta \cdot \top \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot \top \wedge \top \cdot \gamma\}} \wedge_S I$$

$$\frac{\Gamma \vdash_S \{\alpha \cdot \top \wedge \top \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot \top \cdot \gamma\}} \wedge_S E_L$$

In a similar spirit, the two Hoare rules of SL correspond to the *consequence* and *conditional* rules of ordinary Hoare logic, with the traditional conditional rule falling out as a special case of ours since we assume $\Gamma \vdash_H \alpha \vee \neg \alpha$. In Section 6 we extend this further by adding a controlled while loop to our logic. But for now, we illustrate our logic with some very straightforward scenarios.

Example 2.2 (Simple query and return). Consider a very simple state, on which we can perform the following actions:

- (1) Store any value from our domain of discourse in some query location.
- (2) For the current value x in the query location, return a suitable answer y such that $P(x, y)$ holds for some fixed binary predicate system of the logic, and store this value.
- (3) Retrieve the computed value y from the state.

We formalise those three actions by including two unary state predicates **query** and **return_P**, where **query**(x) denotes that x is currently stored in the query location, and **return_P**(x) denotes that some y satisfying $P(x, y)$ has been returned. We would then add the following axioms to Δ_S , which intuitively represent each of the above actions:

- (1) $\Gamma \vdash_S \{\alpha \cdot \top \cdot \mathbf{query}(x)\}$ where α ranges over all state formulas,
- (2) $\Gamma \vdash_S \{\mathbf{query}(x) \cdot \top \cdot \mathbf{return}_P(x)\}$
- (3) $\Gamma \vdash_S \{\mathbf{return}_P(x) \cdot \exists y P(x, y) \cdot \top\}$

The second rule should be regarded as representing an abstract updating of the state where the return value is stored somewhere, and the third the act of retrieval, where following the update we can now formally prove $\exists y P(x, y)$. Note that the computational aspects of this interpretation will only become apparent under the realizability interpretation, where a realizer for $\{\mathbf{return}_P(x) \cdot \exists y P(x, y) \cdot \top\}$ will be a function that accesses any state satisfying **return_P**(x) and produces as output the y satisfying $P(x, y)$ (See Example 5.6).

We can now, for example, derive the following theorem in $\text{SL} + \Delta_H + \Delta_S$ (for $\Delta_H = \emptyset$), where α, β are any state formulas:

$$\vdash_S \{ \beta \cdot \forall x \{ \alpha \cdot \exists y P(x, y) \cdot \top \} \cdot \beta \}$$

The derivation below is essentially a proof of $\forall x \exists y P(x, y)$ that, necessarily, utilising properties of the state:

$$\frac{\frac{\frac{\vdash_S \{ \alpha \cdot \top \cdot \text{query}(x) \} \quad \vdash_S \{ \text{query}(x) \cdot \top \cdot \text{return}_P(x) \}}{\vdash_S \{ \alpha \cdot \top \wedge \top \cdot \text{return}_P(x) \}} \wedge_S I}{\vdash_S \{ \alpha \cdot \top \cdot \text{return}_P(x) \}} \wedge_S E_L \quad \vdash_S \{ \text{return}_P(x) \cdot \exists y P(x, y) \cdot \top \}}{\vdash_S \{ \alpha \cdot \top \wedge \exists y P(x, y) \cdot \top \}} \wedge_S I}{\vdash_S \{ \alpha \cdot \exists y P(x, y) \cdot \top \}} \wedge_S E_L}{\vdash_S \{ \beta \cdot \forall x \{ \alpha \cdot \exists y P(x, y) \cdot \top \} \cdot \beta \}} \forall_S I$$

We note that while state formulas and actions are used in the *proof*, if we set $\alpha = \beta = \top$ then the components of the *theorem* itself are just formulas in ordinary first-order logic. A program corresponding to this derivation will be formally extracted in Example 5.6, but referring to the semantic explanations above, we can view each instance of $\wedge_S I$ in the proof as a sequential composition of state actions (with $\wedge_S E$ just cleaning up the central logical formula), and finally $\forall_S I$ is just the stateful version of the usual \forall -introduction rule.

Example 2.3 (Fixed-length array sorting). Let us now consider our state as an array of length three, and elements in that array as having some order structure. We formalise this in SL by introducing 1, 2, 3 as constants representing our three locations, along with two state predicates: a binary predicate \leq for comparing elements at locations l and l' , and a nullary predicate *sorted* that declares that the state is sorted. These can be characterised by adding the following axiom schemes, but to Δ_H rather than Δ_S as they do not represent state *actions*:

$$\Gamma \vdash_H 1 \leq 2 \wedge 2 \leq 3 \Rightarrow \text{sorted}$$

$$\Gamma \vdash_H l \leq l' \vee l' \leq l \quad \text{where } l, l' \text{ range over } \{1, 2, 3\}$$

We then allow a single action on our array, namely the swapping of a pair of elements in the list. Suppose that α is a state formula of the form

$$\alpha \equiv l_1 \leq l'_1 \wedge \dots \wedge l_n \leq l'_n \tag{2.2}$$

where l_i, l'_i range over locations $\{1, 2, 3\}$. Now for $l, l' \in \{1, 2, 3\}$ let $\alpha[l \leftrightarrow l']$ denote α where all instances of l and l' are swapped, so that if e.g. $\alpha = 3 \leq 2 \wedge 1 \leq 2 \wedge 1 \leq 3$ then

$$\alpha[2 \leftrightarrow 3] = 2 \leq 3 \wedge 1 \leq 3 \wedge 1 \leq 2$$

We axiomatise the swapping of the values in locations of some arbitrary pair $l, l' \in \{1, 2, 3\}$ by adding to Δ_S all instances of

$$\Gamma \vdash_S \{ \alpha \cdot \top \cdot \alpha[l \leftrightarrow l'] \}$$

where α ranges over state formulas of the form (2.2). The statement that all arrays of length three can be sorted is then formulated as

$$\vdash_S \{ \top \cdot \top \cdot \text{sorted} \}$$

Let us now give a proof of this statement in $\text{SL} + \Delta_H + \Delta_S$. As with the previous example, we emphasise that the underlying computation represented by this proof will only become visible through the realizability interpretation.

Figure 3: Typing derivations for $ST + \Lambda_S$

$$\begin{array}{c}
\Gamma \vdash f : D^n \rightarrow D \quad \text{where } f \text{ has arity } n \quad \Gamma \vdash c : X \\
\Gamma \vdash x : X \quad \text{if } x : X \text{ in } \Gamma \quad \Gamma \vdash \text{skip} : C \\
\frac{\Gamma \vdash s : X \quad \Gamma \vdash t : Y}{\Gamma \vdash s \circ t : X \times Y} \quad \frac{\Gamma \vdash t : X \times Y}{\Gamma \vdash p_0(t) : X} \quad \frac{\Gamma \vdash t : X \times Y}{\Gamma \vdash p_1(t) : Y} \\
\frac{\Gamma \vdash t : X}{\Gamma \vdash \iota_0(t) : X + Y} \quad \frac{\Gamma \vdash t : Y}{\Gamma \vdash \iota_1(t) : X + Y} \\
\frac{\Gamma \vdash r : X + Y \quad \Gamma \vdash s : X \rightarrow Z \quad \Gamma \vdash t : Y \rightarrow Z}{\Gamma \vdash \text{elim } r \text{ s t} : Z} \\
\frac{\Gamma, x : X \vdash t : Y}{\Gamma \vdash \lambda x. t : X \rightarrow Y} \quad \frac{\Gamma \vdash t : X \rightarrow Y \quad \Gamma \vdash s : X}{\Gamma \vdash t s : Y} \quad \Gamma \vdash \text{default}_X : X \\
\frac{\Gamma \vdash s : X \quad \Gamma \vdash t : X \quad x : D \in \Gamma \text{ for all free variables of } \alpha}{\Gamma \vdash \text{if } \alpha \text{ then } s \text{ else } t : X}
\end{array}$$

is a set of typed variables. Note that the types of constants $c \in \Lambda_S$ are also left unspecified. The type C should be interpreted as a type of commands that act on the state but don't return any values.

A *denotational* semantics of $ST + \Lambda_S$, which is what we require for our realizability interpretation, will be specified in detail in Section 4 below. Operationally, the idea is that terms $t : X$ of $ST + \Lambda_S$ take some input state π and evaluates to some value v and final state π_1 in a call-by-value way. We choose our notation to reflect the underlying stateful computations: For example, $s \circ t$ is used instead of what would normally be a pairing operation, because this plays the role of simulating composition in the stateful setting. Indeed, it will be helpful to consider a derived operator for sequential composition that also incorporates the ‘cleanup’ seen in the examples above, and corresponds to an instance of $\wedge_S I$ followed by an instance of $\wedge_S E_L$:

Definition 3.1. If $\Gamma \vdash s : C$ and $\Gamma \vdash t : X$ then $\Gamma \vdash s * t := p_1(s \circ t) : X$. In particular, if $\Gamma \vdash t : C$ then $\Gamma \vdash s * t : C$.

A full exploration of the operational semantics of $ST + \Lambda_S$ along with correctness with respect to the denotational semantics defined in Section 4.3 is left to future work, but we provide some further insight in Remark 4.8 below, including a more detailed discussion of the relationship between \circ and pairing.

4. A MONADIC EMBEDDING OF SL AND ST INTO A METATHEORY S^ω

We now give a semantic interpretation of both state formulas of $SL + \Delta_H + \Delta_S$ and terms in $ST + \Lambda_S$ into a standard higher-order, many sorted logic $S^\omega + \Lambda_{S^\omega}$.

4.1. The system S^ω . This logic contains typed lambda terms along with equational axioms for reasoning about them, together with the usual axioms and rules of many-sorted predicate logic. Because most aspects of the logic are completely standard, and in any case it is purely a verifying system, we are less detailed in specifying it. Types are defined as follows:

$$X ::= D \mid 1 \mid \text{Bool} \mid S \mid X \times X \mid X \rightarrow X$$

where D represents objects in the domain of SL (just as in ST), Bool a type of booleans, and states are now explicitly represented as objects of type S . Our metatheory is an equational calculus, with an equality symbol $=_X$ for all types. Typed terms include:

- variables x, y, z, \dots for each type, where we denote state variables by π, π_1, π_2, \dots
- a constant $f : D^n \rightarrow D$ for each n -ary function symbol of SL,
- additional, as yet unspecified constant symbols $c : X$ for interpreting objects in Λ_S , along with axioms that characterise them,
- a unit element $() : 1$ along with the axiom $x = ()$,
- boolean constants \mathbf{t} and \mathbf{f} , with the axiom $x =_{\text{Bool}} \mathbf{t} \vee x =_{\text{Bool}} \mathbf{f}$,
- pairing $\langle s, t \rangle$ and projection $\text{proj}_0(t)$, $\text{proj}_1(t)$ operators, with the usual axioms,
- terms formed by lambda abstraction and application, with the rule $(\lambda x.t)s = t[s/x]$,
- for each type X a case operator $\text{case}(b)(s)(t)$ for $b : \text{Bool}$ and $s, t : X$, with axioms $\text{case } \mathbf{f} \ x \ y = x$ and $\text{case } \mathbf{t} \ x \ y = y$.

We sometimes write x^X instead of $x : X$, and we use abbreviations such as $\langle x, y, z \rangle$ for $\langle x, \langle y, z \rangle \rangle$. Atomic formulas of S^ω include all ordinary predicate symbols P, Q, R, \dots of SL as atomic formulas, where an n -ary predicate P in SL takes arguments of type D^n in S^ω , along with predicates p, q, r, \dots for each state predicate symbol of SL, but now, if p is an n -ary state predicate in SL, p takes arguments of type $D^n \times S$ in S^ω . General formulas are built using the usual logical connectives, including quantifiers for all types. The axioms and rules of S^ω include the axioms of rules of predicate logic (now in all finite types), axioms for the terms, along with the usual equality axioms (including full extensionality). Because S^ω acts as a verifying theory, we freely use strong axioms (such as extensionality), without concerning ourselves with the minimal such system that works.

4.2. The embedding $[\cdot]$ on state formulas of SL. The main purpose of our metalanguage is to allow us to reason semantically about SL and ST. To do this, we introduce an embedding of state formulas of SL and terms of ST into S^ω . We use the same notation $[\cdot]$ for both, as there is no danger of ambiguity. An informal explanation of the meaning of $[\cdot]$ on formulas is given in Section 2.2.

An important point to highlight here is that under the semantics, the arity of state formulas change: For any state formula α , the interpreted formula $[\alpha]$ now contains a single additional free variable $\pi : S$ representing the underlying state. As mentioned earlier, state is implicit in our logic and term languages, but needs to be made explicit under the semantics.

Definition 4.1. For each term t of SL, there is a natural interpretation of t as a term of type D in ST, namely $x \mapsto x : D$ and $f(t_1, \dots, t_n) \mapsto f(t_1 \circ \dots \circ t_n) : D$. Similarly, there is a natural interpretation of t into S^ω , this time with $f(t_1, \dots, t_n) \mapsto f(\langle t_1, \dots, t_n \rangle)$. We use the same notation for t in each of the three systems, as there is no risk of ambiguity.

Definition 4.2. For each *state* formula α of SL, we define a formula $[\alpha](\pi)$ of S^ω , whose free variables are the same as those of α (but now typed with type D) with the potential addition of a single state variable π , as follows:

- $[\top](\pi) := \top$ and $[\perp](\pi) := \perp$,
- $[p(t_1, \dots, t_n)](\pi) := p(t_1, \dots, t_n, \pi)$,
- $[\alpha \wedge \beta](\pi) := [\alpha](\pi) \wedge [\beta](\pi)$, and similarly for $\alpha \vee \beta$ and $\alpha \Rightarrow \beta$.

The following Lemma is easily proven using induction over propositional derivations.

Lemma 4.3. *If $\Gamma \vdash_H \alpha$ in SL then $[\alpha](\pi)$ is provable in S^ω from the assumptions $[\Gamma](\pi)$, where $[\Gamma](\pi) := [\alpha_1](\pi), \dots, [\alpha_n](\pi)$ for $\Gamma := \alpha_1, \dots, \alpha_n$. This extends to proofs in $\text{SL} + \Delta_H$ provided that the embedding of any axiom in Δ_H is provable in $S^\omega + \Lambda_{S^\omega}$.*

We are now in a position to make the semantic meaning of main formulas of SL precise. Note that, technically speaking, this is not necessary in what follows, neither to formulate our realizability interpretation nor to prove our soundness theorem. This is because our main realizability relation (Definition 5.2) is of the form $(x \text{ sr } A)$ for main formulas A of SL: This relation is basically equivalent to $(x \text{ mr } [A])$ for $[A]$ as defined below and a standard modified realizability relation mr , but our realizability interpretation essentially acts as a simultaneous inductive definition of both standard realizability and the embedding $[\cdot]$, and so neither of the latter need to be separately defined.

Definition 4.4. *main* formula A of , we define a formula $[A]$ of S^ω , whose free variables are the same as those of A (but now typed with type D), as follows:

- $[\top] := \top$ and $[\perp] := \perp$,
- $[P(t_1, \dots, t_n)] := P(t_1, \dots, t_n)$,
- $[A \wedge B] := [A] \wedge [B]$, $[A \vee B] := [A] \vee [B]$ and $[\exists x A] := \exists x^D [A]$,
- $[A \Rightarrow \{\alpha \cdot B \cdot \beta\}] := [A] \Rightarrow [\{\alpha \cdot B \cdot \beta\}]$ and $[\forall x \{\alpha \cdot A \cdot \beta\}] := \forall x^D [\{\alpha \cdot A \cdot \beta\}]$

where $[\{\alpha \cdot A \cdot \beta\}] := \exists \pi^S [\alpha](\pi) \Rightarrow [A] \wedge \exists \pi' [\beta](\pi')$.

Similarly to Lemma 4.3, we can now prove the following by induction over derivations in SL. We omit the proof, because it is straightforward and in any case not necessary in what follows.

Proposition 4.5. *If $\Gamma \vdash_S \{\alpha \cdot A \cdot \beta\}$ in SL then $[\{\alpha \cdot A \cdot \beta\}]$ is provable in S^ω from the assumptions $[\Gamma]$, where $[\Gamma] := [A_1], \dots, [A_n]$ for $\Gamma := A_1, \dots, A_n$. This extends to proofs in $\text{SL} + \Delta_H + \Delta_S$ provided that the embedding of any axiom in Δ_H and Δ_S is provable in $S^\omega + \Lambda_{S^\omega}$.*

4.3. The embedding $[\cdot]$ on terms of ST. Our translation on terms is a call-by-value monadic translation using the state monad $S \rightarrow X \times S$, which, intuitively speaking, gives a denotational interpretation of a standard call-by-value operational semantics of terms of ST. A full treatment of the corresponding operational semantics and its adequacy with respect to the and denotational semantics will be left to future work, as we only require the latter for the realizability interpretation. However, to help motivate the definitions that follow, and also justify our claim that $[t]$ can be viewed as an imperative program, we give an informal intuition in Remark 4.8 below.

We first define a translation on types of ST as follows:

- $[D] := D$, $[C] := 1$ and $[X \times Y] := [X] \times [Y]$,
- $[X + Y] := \text{Bool} \times [X] \times [Y]$
- $[X \rightarrow Y] := [X] \rightarrow S \rightarrow [Y] \times S$

Lemma 4.6. *For any type X of SL, the type $[X]$ is inhabited, in the sense that we can define a canonical closed term $0_X : [X]$.*

Proof. Induction on types, letting $0_D := c$ for a constant symbol which is assumed to exist in SL. The only other nonstandard case is $0_{X \rightarrow Y}$, which can be defined as $\lambda x, \pi. \langle 0_Y, \pi \rangle$. \square

Finally, before introducing our translation on terms, we need to add characteristic functions to S^ω for all state formulas (analogous to the characteristic functions for quantifier-free formulas in [GK05]). For any state formula $\alpha[x_1, \dots, x_n]$ of SL, where x_1, \dots, x_n are the free variables of α , we introduce constants $\chi_\alpha : D^n \rightarrow S \rightarrow X \rightarrow X \rightarrow X$ satisfying the axioms

$$\begin{aligned} [\alpha][x_1, \dots, x_n](\pi) &\Rightarrow \chi_\alpha \langle x_1, \dots, x_n \rangle \pi y z = y \\ [-\alpha][x_1, \dots, x_n](\pi) &\Rightarrow \chi_\alpha \langle x_1, \dots, x_n \rangle \pi y z = z \end{aligned}$$

Definition 4.7. For each term $\Gamma \vdash t : X$ of ST we define a term $[\Gamma] \vdash [t] : S \rightarrow [X] \times S$ of S^ω as follows, where $[\cdot]$ is defined on contexts as $[x_1 : X_1, \dots, x_n : X_n] := x_1 : [X_1], \dots, x_n : [X_n]$:

- $[x]\pi := \langle x, \pi \rangle$,
- $[\text{skip}]\pi := \langle (), \pi \rangle$,
- $[f]\pi := \langle \lambda x^{D^n}, \pi. \langle fx, \pi \rangle, \pi \rangle$,
- $[c]\pi$ is appropriately defined for each additional constant in Λ_S ,
- $[s \circ t]\pi := \langle a, b, \pi_2 \rangle$ where $\langle a, \pi_1 \rangle := [s]\pi$ and $\langle b, \pi_2 \rangle := [t]\pi_1$,
- $[p_0 t]\pi := \langle a, \pi_1 \rangle$ and $[p_1 t]\pi := \langle b, \pi_1 \rangle$ where $\langle a, b, \pi_1 \rangle := [t]\pi$,
- $[\iota_0 t]\pi := \langle f, a, 0_Y, \pi_1 \rangle$ and $[\iota_1 t]\pi := \langle t, 0_X, b, \pi_1 \rangle$ for $\langle a, \pi_1 \rangle := [t]\pi$,
- $[\text{elim } r \ s \ t]\pi := \text{case } e \ (fa\pi_2) \ (gb\pi_3)$ for $\langle e, a, b, \pi_1 \rangle := [r]\pi$, $\langle f, \pi_2 \rangle := [s]\pi_1$, $\langle g, \pi_3 \rangle := [t]\pi_1$,
- $[\lambda x.t]\pi := \langle \lambda x^{[X]}. [t], \pi \rangle$,
- $[ts]\pi := fa\pi_2$ for $\langle f, \pi_1 \rangle := [t]\pi$ and $\langle a, \pi_2 \rangle := [s]\pi_1$,
- $[\text{default}_X]\pi := \langle 0_X, \pi \rangle$,
- $[\text{if } \alpha[x_1, \dots, x_n] \text{ then } s \text{ else } t]\pi := \chi_\alpha \langle x_1, \dots, x_n \rangle \pi ([s]\pi) ([t]\pi)$ where $\{x_1, \dots, x_n\}$ are the free variables of α .

Remark 4.8. The intuition behind the embedding $[t]$ is to give a denotational semantics for t as a stateful program. Informally, a term $t : X$ of ST would have operational behaviour $\langle t, \pi \rangle \Downarrow \langle v, \pi_1 \rangle$, where we imagine that t in state π evaluates to a value v and returns a state π_1 . We view $[X]$ as the set of denotations of values of type X , and $[t] : S \rightarrow [X] \times S$ accordingly as a function with $[t]\pi = \langle [v], \pi_1 \rangle$. The interpretation of fully typed terms $[\Gamma \vdash t : X]$ correspond to mappings $[\Gamma] \rightarrow S \rightarrow [X] \times S$ where free variables of t can be instantiated by values of the corresponding type.

With this intuition in mind, each component of Definition 4.7 corresponds to a natural *operational* interpretation of the corresponding term forming rule. For example, the intended operational semantics of $s \circ t$ would be

$$\frac{\langle s, \pi \rangle \Downarrow \langle u, \pi_1 \rangle \quad \langle t, \pi_1 \rangle \Downarrow \langle v, \pi_2 \rangle}{\langle s \circ t, \pi \rangle \Downarrow \langle \langle u, v \rangle, \pi_2 \rangle}$$

and this behaviour is embodied by the denotation $[s \circ t]$, which maps π to $\langle a, b, \pi_2 \rangle$ where a can be interpreted as the denotation of the value u , and b as the denotation of v . Similarly, the call-by-value operational semantics of function application would be expressed by the rule

$$\frac{\langle t, \pi \rangle \Downarrow \langle u, \pi_1 \rangle \quad \langle s, \pi_1 \rangle \Downarrow \langle v, \pi_2 \rangle \quad \langle uv, \pi_2 \rangle \Downarrow \langle w, \pi_3 \rangle}{\langle ts, \pi \rangle \Downarrow \langle w, \pi_3 \rangle}$$

i.e. we evaluate first the function t , then the argument s , then the function application itself. This order of evaluation along with the behaviour of the state is represented semantically by $[ts]$. Finally, we note that the main interactions with the state for terms of $\text{ST} + \Lambda_S$ are driven by the constants $c \in \Lambda_S$, whose semantic interpretation as stateful programs will need to be specified in each case.

Our monadic semantics and its relationship with the intended call-by-value operational semantics is very similar in spirit to the monadic denotational semantics used in [DLR15] and related papers (but with the state monad instead of the complexity monad). We leave a formal definition and detailed exploration of the operational semantics of extracted imperative programs to future work, where it is anticipated that existing work on the operational semantics of imperative languages could be useful and revealing (e.g. [CLM13, McC10, Red96]), particularly in extending our realizability interpretation to incorporate richer imperative languages.

The following lemmas will be useful when verifying our realizability interpretation in the next section. The first is by a simple induction on terms.

Lemma 4.9. *For any term t of SL, we have $[t]\pi = \langle t, \pi \rangle$ (cf. Definitions 4.1 and 4.7).*

Lemma 4.10 (Currying in ST). *Suppose that $\Gamma, x : X, y : Y \vdash t : Z$ is a term in ST, and define $\Gamma \vdash \lambda^*v.t : X \times Y \rightarrow Z$ by $\lambda^*v.t := \lambda v.(\lambda x, y.t)(p_0v)(p_1v)$ where v is not free in t . Then for any $s : X \times Y$ we have*

$$[(\lambda^*v.t)s]\pi = [t][a/x, b/y]\pi_1$$

where $\langle a, b, \pi_1 \rangle := [s]\pi$.

Proof. By unwinding the definition of $[\cdot]$. For any variable $v : X \times Y$ we have $[p_0v]\pi = \langle \text{proj}_0v, \pi \rangle$ and $[p_1v]\pi = \langle \text{proj}_1v, \pi \rangle$, and we also have $[\lambda x, y. t]\pi = \langle \lambda x, \pi. \langle \lambda y. [t], \pi \rangle, \pi \rangle$. We therefore calculate

$$[(\lambda x, y.t)(p_0v)]\pi = (\lambda x, \pi. \langle \lambda y. [t], \pi \rangle)(\text{proj}_0v)\pi = \langle \lambda y. [t][\text{proj}_0v/x], \pi \rangle$$

and thus

$$[(\lambda x, y.t)(p_0v)(p_1v)]\pi = (\lambda y. [t][\text{proj}_0v/x])(\text{proj}_1v)\pi = [t][\text{proj}_0v/x, \text{proj}_1v/y]\pi$$

Finally, we can see that if $\langle a, b, \pi_1 \rangle := [s]\pi$ then

$$\begin{aligned} [(\lambda^* v.t)(s)]\pi &= (\lambda v.[(\lambda x, y.t)(p_0 v)(p_1 v)])(\langle a, b \rangle)\pi_1 \\ &= (\lambda v.[t][\mathbf{proj}_0 v/x, \mathbf{proj}_1 v/y])(\langle a, b \rangle)\pi_1 \\ &= [t][\mathbf{proj}_0 v/x, \mathbf{proj}_1 v/y][\langle a, b \rangle/v]\pi_1 \\ &= [t][a/x, b/y]\pi_1 \end{aligned}$$

which completes the proof. \square

5. A REALIZABILITY INTERPRETATION OF SL INTO ST

We now come to the main contribution of the paper, which is the definition of a realizability relation between terms of ST and formulas of SL, along with a soundness theorem that shows us how to extract realizers from proofs. Our metatheory S^ω is used to define the realizability relation and prove the soundness theorem.

Definition 5.1 (Types of realizers). To each main formula A of SL we assign a type $\tau_S(A)$ of ST as follows:

- $\tau_S(\top) = \tau_S(\perp) = \tau_S(P(t_1, \dots, t_n)) := C$,
- $\tau_S(A \wedge B) := \tau_S(A) \times \tau_S(B)$,
- $\tau_S(A \vee B) := \tau_S(A) + \tau_S(B)$,
- $\tau_S(\exists x A) := D \times \tau_S(A)$,
- $\tau_S(A \Rightarrow \{\alpha \cdot B \cdot \beta\}) := \tau_S(A) \rightarrow \tau_S(B)$,
- $\tau_S(\forall x \{\alpha \cdot A \cdot \beta\}) := D \rightarrow \tau_S(A)$.

Definition 5.2 (Realizability relation). For each main formula A of SL we define a formula $x \text{ sr } A$ of S^ω , whose free variables are contained in those of A (now typed with type D) together with a fresh variable $x : [\tau_S(A)]$, by induction on the structure of A as follows:

- $x \text{ sr } Q := Q$ for $Q = \top, \perp$ or $P(t_1, \dots, t_n)$,
- $x \text{ sr } A \wedge B := (\mathbf{proj}_0 x \text{ sr } A) \wedge (\mathbf{proj}_1 x \text{ sr } B)$,
- $x \text{ sr } A \vee B := (\mathbf{proj}_0 x = \mathbf{f} \Rightarrow \mathbf{proj}_0(\mathbf{proj}_1 x) \text{ sr } A) \wedge (\mathbf{proj}_0 x = \mathbf{t} \Rightarrow \mathbf{proj}_1(\mathbf{proj}_1 x) \text{ sr } B)$,
- $x \text{ sr } \exists y A(y) := (\mathbf{proj}_1 x \text{ sr } A)[\mathbf{proj}_0 x/y]$,
- $f \text{ sr } (A \Rightarrow \{\alpha \cdot B \cdot \beta\}) := \forall x^{[\tau_S(A)]} (x \text{ sr } A \Rightarrow f x \text{ sr } \{\alpha \cdot B \cdot \beta\})$,
- $f \text{ sr } (\forall x \{\alpha(x) \cdot A(x) \cdot \beta(x)\}) := \forall x^D (f x \text{ sr } \{\alpha(x) \cdot A(x) \cdot \beta(x)\})$,

where for $x : S \rightarrow [\tau_S(A)] \times S$ we define

- $x \text{ sr } \{\alpha \cdot A \cdot \beta\} := \forall \pi^S ([\alpha](\pi) \Rightarrow \mathbf{proj}_0(x\pi) \text{ sr } A \wedge [\beta](\mathbf{proj}_1(x\pi)))$.

The following substitution lemma is easily proven by induction on formulas of SL.

Lemma 5.3. *For any term t of SL and $s : [\tau_S(A)]$ we have $s \text{ sr } A[t/x] = (s \text{ sr } A)[t/x]$, where x is not free in s and on the right hand side we implicitly mean the natural interpretation of t in S^ω (cf. Definition 4.1).*

Theorem 5.4 (Soundness). *Suppose that*

$$\Gamma := A_1^{u_1}, \dots, A_n^{u_n} \vdash_S \{\alpha \cdot A \cdot \beta\}$$

is provable in SL. Then we can extract from the proof a term $\Delta, \tau_S(\Gamma) \vdash t : \tau_S(A)$ of ST, where Δ contains the free variables of Γ and $\{\alpha \cdot A \cdot \beta\}$ (typed with type D) and $\tau_S(\Gamma) := x_1 : \tau_S(A_1), \dots, x_n : \tau_S(A_n)$ for fresh variables x_1, \dots, x_n , such that the formula

$$[t] \text{ sr } \{\alpha \cdot A \cdot \beta\}$$

is provable in S^ω from the assumptions $(x_1 \text{ sr } A_1)^{u_1}, \dots, (x_n \text{ sr } A_n)^{u_n}$ for $x_i : [\tau_S(A_i)]$. The theorem holds more generally for proofs in $\text{SL} + \Delta_H + \Delta_S$, now provably in $S^\omega + \Lambda_{S^\omega}$, if:

- *for any axiom $\Gamma \vdash_H \alpha$ in Δ_H , the corresponding axiom $[\Gamma](\pi) \Rightarrow [\alpha](\pi)$ is added to Λ_{S^ω} ,*
- *for any axiom in Δ_S there is a term t of $\text{ST} + \Lambda_S$ such that $[t]$ realizes that axiom provably in $S^\omega + \Lambda_{S^\omega}$.*

Proof. Induction on the structure of derivations in SL. In all cases, we assume as global assumptions $(x_1 \text{ sr } A_1)^{u_1}, \dots, (x_n \text{ sr } A_n)^{u_n}$, and our aim is then to produce a term t such that if $[\alpha](\pi)$ holds for some state variable π , then $a \text{ sr } A$ and $[\beta](\pi_1)$ hold for $\langle a, \pi_1 \rangle := [t]\pi$.

- For the axiom $\Gamma \vdash_S \{\alpha \cdot A \cdot \alpha\}$, if $A^u \in \Gamma$ we define $t := x$ for the corresponding variable $x : \tau_S(A)$. Then $[x]\pi := \langle x, \pi \rangle$ for $x \text{ sr } A$ and $[\alpha](\pi)$. For $\Gamma \vdash_S \{\alpha \cdot \top \cdot \alpha\}$ we define $t := \text{skip}$ and the verification is even simpler.
- ($\wedge_S I$) Given terms s, t with $[s] \text{ sr } \{\alpha \cdot A \cdot \beta\}$ and $[t] \text{ sr } \{\beta \cdot B \cdot \gamma\}$, from $[\alpha](\pi)$ we can infer $a \text{ sr } A$ and $[\beta](\pi_1)$ for $\langle a, \pi_1 \rangle := [s]\pi$, and from $[\beta](\pi_1)$ it follows that $b \text{ sr } B$ and $[\gamma](\pi_2)$ for $\langle b, \pi_2 \rangle := [t]\pi_1$, therefore we have shown that $[s \circ t] \text{ sr } \{\alpha \cdot A \wedge B \cdot \gamma\}$.
- ($\wedge_S E_i$) If $[t] \text{ sr } \{\alpha \cdot A \wedge B \cdot \beta\}$ then $\langle a, b \rangle \text{ sr } A \wedge B$ and $[\beta](\pi_1)$ follow from $[\alpha](\pi)$, where $\langle a, b, \pi_1 \rangle := [t]\pi$. But then $[p_0 t] \text{ sr } \{\alpha \cdot A \cdot \beta\}$ and $[p_1 t] \text{ sr } \{\alpha \cdot B \cdot \beta\}$.
- ($\vee_S I_i$) If $[t] \text{ sr } \{\alpha \cdot A \cdot \beta\}$ and $[\alpha](\pi)$ holds, then $a \text{ sr } A$ and $[\beta](\pi_1)$ for $\langle a, \pi_1 \rangle := [t]\pi$, and therefore

$$(b = f \Rightarrow a \text{ sr } A) \wedge (b = t \Rightarrow 0_{\tau_S(B)} \text{ sr } B)$$

for $b := f$. Thus $[\iota_0 t] \text{ sr } A \vee B$. By an entirely analogous argument we can show that $[\iota_1 t] \text{ sr } A \vee B$ whenever $[t] \text{ sr } B$.

- ($\vee_S E$) Suppose that $r, s(x)$ and $t(y)$ are such that $[r] \text{ sr } \{\alpha \cdot A \vee B \cdot \beta\}$, $[s](x) \text{ sr } \{\beta \cdot C \cdot \gamma\}$ assuming $x \text{ sr } A$, and $[t](y) \text{ sr } \{\beta \cdot C \cdot \gamma\}$ assuming $y \text{ sr } B$. We claim that

$$[\text{elim } r (\lambda x. s) (\lambda y. t)] \text{ sr } \{\alpha \cdot C \cdot \gamma\}$$

To prove this, first note that if $[\alpha](\pi)$, we have $\langle e, a, b \rangle \text{ sr } A \vee B$ and $[\beta](\pi_1)$ for $\langle e, a, b, \pi_1 \rangle := [r]\pi$. There are now two possibilities. If $e = f$ then

$$\begin{aligned} [\text{elim } r (\lambda x. s) (\lambda y. t)]\pi &= f a \pi_2 \quad \text{for } \langle f, \pi_2 \rangle := [\lambda x. s]\pi_1 = \langle \lambda x. [s](x), \pi_1 \rangle \\ &= (\lambda x. [s](x)) a \pi_1 \\ &= [s](a) \pi_1 \end{aligned}$$

But since $[\beta](\pi_1)$ holds and $e = f$ also implies that $a \text{ sr } A$, we have $c \text{ sr } C$ and $[\gamma](\pi_2)$ for $\langle c, \pi_2 \rangle := [s](a) \pi_1$, which proves the main claim in the case $e = f$. An analogous argument works for the case $e = t$.

- ($\Rightarrow_S I$) If $t(x)$ is such that $[t](x)$ sr $\{\alpha \cdot B \cdot \beta\}$ whenever x sr A , then by definition we have

$$\lambda x.[t] \text{ sr } A \Rightarrow \{\alpha \cdot B \cdot \beta\}$$

and therefore $[\lambda x.t]$ sr $\{\gamma \cdot A \Rightarrow \{\alpha \cdot B \cdot \beta\} \cdot \gamma\}$ for any γ .

- ($\Rightarrow_S E$) Assume that $[s]$ sr $\{\beta \cdot A \cdot \gamma\}$ and $[t]$ sr $\{\alpha \cdot A \Rightarrow \{\gamma \cdot B \cdot \delta\} \cdot \beta\}$. If $[\alpha](\pi)$ holds then defining $\langle f, \pi_1 \rangle := [t]\pi$ we have $[\beta]\pi_1$ and

$$x \text{ sr } A \Rightarrow f x \text{ sr } \{\gamma \cdot B \cdot \delta\}$$

Similarly, defining $\langle a, \pi_2 \rangle := [s]\pi_1$, it follows that $[\gamma](\pi_2)$ and a sr A . Finally, setting $\langle b, \pi_3 \rangle := f a \pi_2$ it follows that b sr B and $[\delta](\pi_3)$, and we have therefore proven that $[ts]$ sr $\{\alpha \cdot B \cdot \delta\}$.

- ($\perp_S E$) If $[t]$ sr $\{\alpha \cdot \perp \cdot \beta\}$ then from $[\alpha](\pi)$ we can infer a sr \perp and $[\beta](\pi_2)$ for $\langle a, \pi_1 \rangle := [t]\pi$. But a sr $\perp = \perp$, and from \perp we can deduce anything, and in particular $0_{\tau_S(A)}$ sr A and $[\gamma](\pi)$, from which it follows that $[\text{default}_{\tau_S(A)}]$ sr $\{\alpha \cdot A \cdot \gamma\}$.
- ($\forall_S I$) Suppose that $t(x)$ is such that $[t](y)$ sr $\{\alpha[y/x] \cdot A[y/x] \cdot \beta[y/x]\}$, where $y \equiv x$ or y is not free in $\{\alpha \cdot A \cdot \beta\}$, and y is not free in Γ . Then since y is not free in any of the assumptions x_i sr A_i , we can deduce in S^ω that

$$\forall x^D [t](x) \text{ sr } \{\alpha \cdot A \cdot \beta\}$$

and therefore $\lambda x.[t]$ sr $\forall x \{\alpha \cdot A \cdot \beta\}$, and thus (just as for $\Rightarrow_S I$) we have

$$[\lambda x.t] \text{ sr } \{\gamma \cdot \forall x \{\alpha \cdot A \cdot \beta\} \cdot \gamma\}$$

for any γ .

- ($\forall_S E$) Suppose that $[s]$ sr $\{\alpha \cdot \forall x \{\beta \cdot A \cdot \gamma\} \cdot \beta[t/x]\}$ and that $[\alpha](\pi)$ holds. Then we have f sr $\forall x \{\beta \cdot A \cdot \gamma\}$ and $[\beta][t/x](\pi_1)$ for $\langle f, \pi \rangle := [s]\pi$. Now, using Lemma 4.9 we have $[st]\pi = ft\pi_1$ for the natural interpretation of t in S^ω , since we can prove in S^ω that

$$ft \text{ sr } \{\beta[t/x] \cdot A[t/x] \cdot \gamma[t/x]\}$$

it follows that a sr $A[t/x]$ and $[\gamma][t/x](\pi_2)$ for $\langle a, \pi_2 \rangle := ft\pi_1$, and therefore we have shown that $[st]$ sr $\{\alpha \cdot A[t/x] \cdot \gamma[t/x]\}$.

- ($\exists_S I$) If $[s]$ sr $\{\alpha \cdot A[t/x] \cdot \beta\}$ and $[\alpha](\pi)$ then a sr $A[t/x]$ and $[\beta](\pi_1)$ for $\langle a, \pi_1 \rangle := [s]\pi$. By Lemma 5.3 we therefore have $(a \text{ sr } A)[t/x]$, and therefore $\langle t, a \rangle$ sr $\exists x A$. Observing (using Lemma 4.9) that $[t \circ s]\pi = \langle t, a, \pi_1 \rangle$, we have shown that $[t \circ s]$ sr $\{\alpha \cdot \exists x A \cdot \beta\}$.
- ($\exists_S E$) Suppose that s and $t(x, z)$ are such that $[s]$ sr $\{\alpha \cdot \exists x A \cdot \beta\}$ and

$$z \text{ sr } A[y/x] \Rightarrow [t](y, z) \text{ sr } \{\beta \cdot C \cdot \gamma\}$$

where $y \equiv x$ or y is not free in A , and y is also not free in C , α , β , γ or Γ . By Lemma 5.3 that z sr $A[y/x] = (z \text{ sr } A)[y/x] = \langle y, z \rangle$ sr $\exists x A$ we therefore have

$$\langle y, z \rangle \text{ sr } \exists x A \Rightarrow [t](y, z) \text{ sr } \{\beta \cdot C \cdot \gamma\}$$

Now, applying Lemma 4.10 to $\Delta, \Gamma, y : D, z : \tau_S(A) \vdash t : \tau_S(C)$, we have

$$[(\lambda^* v.t)s]\pi = [t](e, a)\pi_1$$

for $\langle e, a, \pi_1 \rangle := [s]\pi$. Now, if $[\alpha](\pi)$ holds, then we have $\langle e, a \rangle$ sr $\exists x A$ and $[\beta](\pi_1)$, and therefore since $[t](e, a)$ sr $\{\beta \cdot C \cdot \gamma\}$, we have c sr C and $[\gamma](\pi_2)$ for $\langle c, \pi_2 \rangle = [t](e, a)\pi_1 = [(\lambda^* v.t)s]\pi$, and thus we have shown that $[(\lambda^* v.t)s]$ sr $\{\alpha \cdot C \cdot \gamma\}$.

- (*cons*) If $\alpha \vdash_H \beta$ and $\gamma \vdash_H \delta$ then by Lemma 4.3 both $[\alpha](\pi) \Rightarrow [\beta](\pi)$ and $[\gamma](\pi) \Rightarrow [\delta](\pi)$ are provable in S^ω (respectively $S^\omega + \Lambda_{S^\omega}$ for the general version of the theorem) for any $\pi : S$. It is then easy to show that if $[t] \text{ sr } \{\beta \cdot A \cdot \gamma\}$ then we also have $[t] \text{ sr } \{\alpha \cdot A \cdot \delta\}$.
- (*cond*) Suppose that $[s] \text{ sr } \{\alpha \wedge \gamma \cdot A \cdot \delta\}$ and $[t] \text{ sr } \{\beta \wedge \gamma \cdot A \cdot \delta\}$. We claim that

$$[\text{if } \alpha \text{ then } s \text{ else } t] \text{ sr } \{\gamma \cdot A \cdot \delta\}$$

To prove this, suppose that $[\gamma](\pi)$ holds. Since $\vdash_H \alpha \vee \beta$ then $[\alpha](\pi) \vee [\beta](\pi)$ is provable in S^ω , and so we consider two cases. Let $\{x_1, \dots, x_n\}$ be the free variables of α . If $[\alpha](\pi)$ holds, then

$$[\text{if } \alpha \text{ then } s \text{ else } t]\pi = \chi_\alpha \langle x_1, \dots, x_n \rangle \pi ([s]\pi) ([t]\pi) = [s]\pi$$

and since then $[\alpha](\pi) \wedge [\gamma](\pi)$ we have $a \text{ sr } A$ and $[\delta](\pi_1)$ for $\langle a, \pi_1 \rangle := [s]\pi$. On the other hand, if $[\beta](\pi)$ holds, then by an analogous argument we can show that $a \text{ sr } A$ and $[\delta](\pi_1)$ for $\langle a, \pi_1 \rangle := [t]\pi = [\text{if } \alpha \text{ then } s \text{ else } t]\pi$, and we are done.

The extension of the soundness theorem to $\text{SL} + \Delta_H + \Delta_S$ is straightforward, as the soundness proof is modular and so any axioms along with their realizers can be added. The first condition is needed so that Lemma 4.3 (needed for the *cons* rule) continues to apply.

For the free variable condition that the free variables of t are contained in those of Γ , $\{\alpha \cdot A \cdot \beta\}$ and $\tau_S(\Gamma)$, if this were not the case, we could simply ground those variables with a canonical constant $c : D$ and we would still have $\tilde{t} \text{ sr } \{\alpha \cdot A \cdot \beta\}$ for the resulting term \tilde{t} . \square

Corollary 5.5 (Program extraction). *Suppose that the sentence*

$$\vdash_S \{\alpha \cdot \forall x \{\beta \cdot \exists y P(x, y) \cdot \gamma(x)\} \cdot \beta\}$$

is provable in $\text{SL} + \Delta_S$. Then we can extract a closed realizing term $t : D \rightarrow D \times C$ in $\text{ST} + \Lambda_S$ such that defining $g : D \rightarrow S \rightarrow D \times S$ by $gx\pi := \langle a, \pi_2 \rangle$ for $\langle f, \pi_1 \rangle := [t]\pi$ and $\langle a, (), \pi_2 \rangle := fx\pi_1$, we have

$$\forall \pi^S([\alpha](\pi) \Rightarrow \forall x^D (P(x, \text{proj}_0(gx\pi)) \wedge [\gamma](x)(\text{proj}_1(gx\pi))))$$

provably in $S^\omega + \Lambda_{S^\omega}$.

5.1. Simplification and removal of unit types. In presentations of modified realizability that use product types instead of type sequences, it is common to introduce the notion of a Harrop formula (a formula that does not contain disjunction or existential quantification in a positive position) and define realizability in a way that all Harrop formulas have unit realizability type, so that e.g. $\tau_S(\forall x (P \wedge Q)) = 1$ for atomic predicates P and Q , rather than $\tau_S(\forall x (P \wedge Q)) = D \rightarrow 1 \times 1$ as for us. We have avoided this simplification earlier on, as it would have added additional cases and bureaucracy to our soundness theorem. However, we can compensate retroactively for this choice by introducing equivalences on types that eliminate unit types, namely the closure under contexts of

$$1 \times X \simeq X \simeq X \times 1 \quad (1 \rightarrow X) \simeq X \quad (X \rightarrow 1) \simeq 1$$

along with corresponding equivalences on terms, also closed under contexts:

$$t^{1 \times X} \simeq \text{proj}_1(t)^X \quad t^{X \times 1} \simeq \text{proj}_0(t)^X \quad t^{1 \rightarrow X} \simeq t() \quad t^X \simeq \lambda x^1. t \quad t^{X \rightarrow 1} \simeq ()$$

For example, in Corollary 5.5 we would then have

$$[t]\pi : (D \rightarrow S \rightarrow D \times 1 \times S) \times S \simeq (D \rightarrow S \rightarrow D \times S) \times S \quad \text{and} \quad gx\pi \simeq fx\pi_1.$$

For us, the equivalence relation \simeq will not play a formal role in the paper, but will be used to provide simplified descriptions of extracted programs. In particular, we do not rely on it in any way for the main results. To incorporate the simplifications formally into our framework, the most obvious route would involve defining a more elaborate realizability relation as indicated above, with the interpretation of Harrop formulas treated separately (as in the recent and related paper [BT21]), which would then also generate many new (but routine) cases in the soundness proof. Given that we are already introducing a new and nonstandard realizability relation, in this article we prefer to work with a simple interpretation and use the equivalences above in an informal way to more concisely describe the meaning of extracted programs.

5.2. Examples of program extraction. We now continue the short illustrative examples we outlined in Section 2.2.

Example 5.6 (Simple read-write). In Example 2.2 we considered a state where three actions were possible (writing to the state, performing a calculation, and reading the output from the state). We can formalise these three actions semantically in the metatheory S^ω by including three constants in Λ_{S^ω} , namely $c_1 : D \rightarrow S \rightarrow S$, $c_2 : S \rightarrow S$ and $c_3 : S \rightarrow D$, along with the characterising axioms:

- (1) $\text{query}(x, c_1 x \pi)$,
- (2) $\text{query}(x, \pi) \Rightarrow \text{return}_P(x, c_2 \pi)$,
- (3) $\text{return}_P(x, \pi) \Rightarrow P(x, c_3 \pi)$.

While we are able to use these constants to form terms in S^ω such as $\lambda \pi, \pi_1, x. \langle c_1 x \pi, c_2 \pi_1 \rangle$, which could be viewed as non-sequential in the sense that we take two input states as arguments, we can force them to be applied in a sequential, call-by-value manner by adding three corresponding constants to our term calculus ST, namely including $\text{write} : D \rightarrow C$, $\text{calc} : C$ and $\text{read} : D \times C$ in Λ_S , along with the embedding rules

- $[\text{write}] \pi := \langle \lambda x, \pi'. \langle (), c_1 x \pi' \rangle, \pi \rangle \simeq \langle c_1, \pi \rangle$,
- $[\text{calc}] \pi := \langle (), c_2 \pi \rangle$ so that $[\text{calc}] \simeq c_2$,
- $[\text{read}] \pi := \langle c_3 \pi, (), \pi \rangle \simeq \langle c_3 \pi, \pi \rangle$.

and then restricting our attention to terms of the form $[t]$ for $t \in \text{ST} + \{\text{write}, \text{calc}, \text{read}\}$. We can then prove the following in S^ω i.e. that all axioms in Δ_S can be realised:

- $[\text{write}(x)] \text{sr } \{\alpha \cdot \top \cdot \text{query}(x)\}$,
- $[\text{calc}] \text{sr } \{\text{query}(x) \cdot \top \cdot \text{return}_P(x)\}$,
- $[\text{read}] \text{sr } \{\text{return}_P(x) \cdot \exists y P(x, y) \cdot \top\}$.

and thus Theorem 5.4 applies to $\text{SL} + \Delta_H + \Delta_S$ for $\Delta_H = \emptyset$. In particular, we have

$$[t] \text{sr } \{\beta \cdot \forall x \{\alpha \cdot \exists y P(x, y) \cdot \top\} \cdot \beta\}$$

for $t := \lambda x. ((\text{write}(x) * \text{calc}) * \text{read})$ where $*$ is sequential composition operator from Definition 3.1. A formal derivation of this term from the corresponding proof given in

Example 2.2 is as follows:

$$\frac{\frac{\frac{x : D \vdash \text{write}(x) : C \quad x : D \vdash \text{calc} : C}{x : D \vdash \text{write}(x) \circ \text{calc} : C \times C} \wedge_{SI} I}{x : D \vdash \text{write}(x) * \text{calc} : C} \wedge_{SE} L \quad \frac{x : D \vdash \text{read} : D \times C}{x : D \vdash (\text{write}(x) * \text{calc}) \circ \text{read} : C \times D \times C} \wedge_{SI} I}{x : D \vdash (\text{write}(x) * \text{calc}) * \text{read} : D \times C} \wedge_{SE} L}{\vdash \lambda x. ((\text{write}(x) * \text{calc}) * \text{read}) : D \rightarrow D \times C} \forall_{SI} I$$

Example 5.7 (Fixed-length array sorting). In Example 2.3 we considered a situation where we are allowed a single action on our state, namely to swap elements. Analogously to the previous example, we can formalise this in our semantic environment S^ω by adding to Λ_{S^ω} constants $c_{l,l'} : S \rightarrow S$ for each pair $l, l' \in \{1, 2, 3\}$ along with the axiom

$$[\alpha](\pi) \Rightarrow [\alpha[l \leftrightarrow l']](c_{l,l'}\pi)$$

ranging over state formulas α of the form (2.2) and locations $l, l' \in \{1, 2, 3\}$ of SL, together with axioms corresponding to those of Δ_H i.e.

$$[1 \leq 2 \wedge 2 \leq 3](\pi) \Rightarrow \text{sorted}(\pi) \quad \text{and} \quad [l \leq l' \vee l' \leq l](\pi)$$

Similarly, for each $l, l' \in \{1, 2, 3\}$ we add a term $\text{swap}_{l,l'} : C$ to Λ_S and define $[\text{swap}_{l,l'}]\pi := \langle(), c_{l,l'}\pi\rangle$ so that

$$\text{swap}_{l,l'} \text{ sr } \{\alpha \cdot \top \cdot \alpha[l \leftrightarrow l']\}$$

A derivation of a closed term $t : C$ of $\text{ST} + \{\text{swap}_{l,l'}\}$ such that $[t] \text{ sr } \{\top \cdot \top \cdot \text{sorted}\}$ is given below. In particular, we can prove in S^ω that $\forall \pi^S \text{sorted}(\text{proj}_1([t]\pi))$, and so the term $\lambda \pi. \text{proj}_1([t]\pi) : S \rightarrow S$ acts as a sorting program for arrays of length three. For an extracted term t corresponding to the proof given in Example 2.3, first we interpret \mathcal{D}_1 as

$$\frac{\frac{\vdash \text{skip} : C}{\vdash \text{skip} : C} \text{cons} \quad \frac{\frac{\vdash \text{swap}_{2,3} : C}{\vdash \text{swap}_{2,3} : C} \text{cons} \quad \frac{\vdash \text{swap}_{2,3} : C}{\vdash \text{swap}_{2,3} : C} \text{cons}}{\vdash \text{skip} : C} \text{cons}}{\vdash t_1 := \text{if } (2 \leq 3) \text{ then } (\text{skip}) \text{ else } (\text{swap}_{2,3}) : C} \text{cond}[2 \leq 3 \vee 3 \leq 2]$$

and define $t_1 := \text{if } (2 \leq 3) \text{ then } (\text{skip}) \text{ else } (\text{swap}_{2,3})$. Now \mathcal{D}_2 is interpreted as

$$\frac{\frac{\frac{\vdash \text{swap}_{1,2} : C}{\vdash \text{swap}_{1,2} : C} \text{cons} \quad \frac{\frac{\mathcal{D}_1}{\vdash t_1 : C}}{\vdash t_1 : C} \wedge_{SI} I}{\vdash \text{swap}_{1,2} \circ t_1 : C \times C} \wedge_{SE} L}{\vdash t_2 := \text{swap}_{1,2} * t_1 : C} \wedge_{SE} L$$

where we define $t_2 := \text{swap}_{1,2} * t_1 : C$. Continuing, \mathcal{D}_3 is interpreted as:

$$\frac{\frac{\mathcal{D}_2}{\vdash t_2 : C} \text{cons} \quad \frac{\vdash \text{skip} : C}{\vdash \text{skip} : C} \text{cons}}{\vdash t_3 := \text{if } (2 \leq 1) \text{ then } t_2 \text{ else } (\text{skip}) : C} \text{cond}[2 \leq 1 \vee 1 \leq 2]$$

where $t_3 := \text{if } (2 \leq 1) \text{ then } t_2 \text{ else } (\text{skip})$, and finally

$$\frac{\frac{\frac{\frac{\vdash \text{skip} : C \quad \overline{\vdash \text{swap}_{2,3} : C}^{2 \leftrightarrow 3}}{\vdash \text{if } (2 \leq 3) \text{ then } (\text{skip}) \text{ else } (\text{swap}_{2,3}) : C}^{\text{cond}[2 \leq 3 \vee 3 \leq 2]}}{\vdash (\text{if } (2 \leq 3) \text{ then } (\text{skip}) \text{ else } (\text{swap}_{2,3})) \circ t_3 : C \times C}^{\wedge_S I}}{\vdash t := (\text{if } (2 \leq 3) \text{ then } (\text{skip}) \text{ else } (\text{swap}_{2,3})) * t_3 : C}^{\wedge_S E_L}}{\mathcal{D}_3} \vdash t_3 : C$$

6. AN EXTENSION TO ARITHMETIC

We now present an extension of our framework to a stateful version of first-order intuitionistic arithmetic. On the logic side, we will add not only a stateful induction rule, but also a Hoare-style while rule for iteration over the natural numbers. On the computational side, these will be interpreted by stateful recursion in all finite types, along with a controlled while loop. The addition of these constants will allow us to extract programs that are more interesting than those obtainable from proofs in pure predicate logic, and which can be clearly compared to well-known stateful algorithms. To exemplify this, we will present a formally synthesised version of insertion sort, and we stress that by further extending our framework with additional rules and terms, we would be able to extract an even richer variety of combined functional/stateful programs.

6.1. The system SA: First-order arithmetic with state. Our system of stateful intuitionistic arithmetic SA builds on SL just as ordinary first-order Heyting arithmetic builds on first-order predicate logic. In both cases, we introduce a constant 0, a unary successor symbol succ , symbols for all primitive recursive functions, and our predicate symbols now include an equality relation $=$. In what follows we write $x + 1$ instead of $\text{succ}(x)$. The axioms and rules of SA are, in turn, analogous to the additional axioms and rules we would require in ordinary first-order arithmetic: They include all axioms and rules of SL (based now on the language of SA), along with a collection of additional axioms and rules. These comprise not only basic axioms and rules for equality and the successor, and an induction rule (all now adapted to incorporate the state), but also a new while rule for stateful iteration, which now exploits our state and, as we will see, allows us to extract programs that contain while loops. These additional axioms and rules are given in Figure 4.

Our formulation of stateful arithmetic follows the same basic idea as the construction of stateful predicate logic, incorporating standard rules but keeping track of an ambient state in a call-by-value manner, and adding new rules that explicitly correspond to stateful constructions. In particular, Proposition 2.1 clearly extends to SA, as the usual axioms and rules of arithmetic can be embedded into those of SA:

Proposition 6.1. *For any formula A of HA and state formula α , define the main formula A_α of SA as in Proposition 2.1. Then whenever $\Gamma \vdash_I A$ is provable in HA, we have that $\Gamma_\alpha, \Delta \vdash_S \{\alpha \cdot A_\alpha \cdot \alpha\}$ is provable in SA, where Δ is arbitrary and $\Gamma_\alpha := (A_1)_\alpha^{u_1}, \dots, (A_n)_\alpha^{u_n}$ for $\Gamma := A_1^{u_1}, \dots, A_n^{u_n}$.*

We can also derive a natural extensionality rule from our stateful equality rules, which assures us that whenever $s = t$ in ordinary Heyting arithmetic, then we can replace s by t for stateful formulas:

Figure 4: Additional axioms and rules of SA

Axioms and rules for equality

$$\Gamma \vdash_S \{\alpha \cdot t = t \cdot \alpha\} \quad \frac{\Gamma \vdash_S \{\alpha \cdot s = t \cdot \beta\}}{\Gamma \vdash_S \{\alpha \cdot t = s \cdot \beta\}} \quad \frac{\Gamma \vdash_S \{\alpha \cdot r = s \cdot \beta\} \quad \Gamma \vdash_S \{\beta \cdot s = t \cdot \gamma\}}{\Gamma \vdash_S \{\alpha \cdot r = t \cdot \gamma\}}$$

$$\frac{\Gamma \vdash_S \{\alpha \cdot s = t \cdot \beta\} \quad \Gamma \vdash_S \{\beta \cdot A(s) \cdot \gamma(s)\}}{\Gamma \vdash_S \{\alpha \cdot A(t) \cdot \gamma(t)\}} \textit{ext}$$

Axioms and rules for arithmetical function symbols

$$\Gamma \vdash_S \{\alpha \cdot \text{succ}(t) \neq 0 \cdot \alpha\} \quad \frac{\{\alpha \cdot \text{succ}(s) = \text{succ}(t) \cdot \beta\}}{\{\alpha \cdot s = t \cdot \beta\}}$$

$\Gamma \vdash_S \{\alpha \cdot l = r \cdot \alpha\}$ where $l = r$ ranges across defining equations for prim. rec. functions

Induction rule

$$\frac{\Gamma \vdash_S \{\alpha \cdot A(0) \cdot \beta(0)\} \quad \Gamma, A(x) \vdash_S \{\beta(x) \cdot A(x+1) \cdot \beta(x+1)\}}{\Gamma \vdash_S \{\gamma \cdot \forall x \{\alpha \cdot A(x) \cdot \beta(x)\} \cdot \gamma\}} \textit{ind}$$

While rule (over natural numbers)

$$\frac{\mathcal{A}_1 \quad \mathcal{A}_2 \quad \mathcal{A}_3}{\Gamma, A(x) \vdash_S \{\alpha(x) \cdot B \cdot \beta\}} \textit{while}$$

$$\mathcal{A}_1 := \Gamma, A(x+1) \vdash_S \{\gamma(x+1) \wedge \alpha(x+1) \cdot A(x) \cdot \alpha(x)\}$$

$$\mathcal{A}_2 := \Gamma, A(x+1) \vdash_S \{\neg \gamma(x+1) \wedge \alpha(x+1) \cdot B \cdot \beta\}$$

$$\mathcal{A}_3 := \Gamma, A(0) \vdash_S \{\alpha(0) \cdot B \cdot \beta\}$$

for *ind* and *while*, x is not free in Γ , and for *while* it is not free in B or β

Proposition 6.2. *Suppose that $\vdash_I s = t$ is provable in HA. Then from $\Gamma \vdash_S \{\alpha(s) \cdot A(s) \cdot \beta(s)\}$ we can derive $\Gamma \vdash_S \{\alpha(t) \cdot A(t) \cdot \beta(t)\}$ in SA.*

Proof. By Proposition 6.1 for $\alpha := \alpha(s)$ we have $\Gamma \vdash_S \{\alpha(s) \cdot s = t \cdot \alpha(s)\}$ and thus using the extensionality rule in SA we can derive

$$\frac{\Gamma \vdash_S \{\alpha(s) \cdot s = t \cdot \alpha(s)\} \quad \Gamma \vdash_S \{\alpha(s) \cdot A(s) \cdot \beta(s)\}}{\Gamma \vdash_S \{\alpha(s) \cdot A(t) \cdot \beta(t)\}} \textit{ext}$$

Since $\vdash_I t = s$ must also be provable in HA, another instance of Proposition 6.1 for $\alpha := \alpha(t)$ along with the true axiom in SA gives us

$$\frac{\Gamma \vdash_S \{\alpha(t) \cdot t = s \cdot \alpha(t)\} \quad \Gamma \vdash_S \{\alpha(t) \cdot \top \cdot \alpha(t)\}}{\Gamma \vdash_S \{\alpha(t) \cdot \top \cdot \alpha(s)\}} \textit{ext}$$

Putting these together we obtain

$$\frac{\frac{\Gamma \vdash_S \{\alpha(t) \cdot \top \cdot \alpha(s)\} \quad \Gamma \vdash_S \{\alpha(s) \cdot A(t) \cdot \beta(t)\}}{\Gamma \vdash_S \{\alpha(t) \cdot \top \wedge A(t) \cdot \beta(t)\}} \wedge_S I}{\Gamma \vdash_S \{\alpha(t) \cdot A(t) \cdot \beta(t)\}} \wedge_S E_L$$

which completes the derivation. \square

6.2. An extended term calculus ST_N . In order to give derivations in SA a computation interpretation, we need to extend our term calculus ST to include a recursor (for induction) and a controlled while loop (for the while rule). The remaining new axioms and rules of SA are dealt with in a straightforward manner.

To be precise: the theory ST_N is defined to be the instance of ST for the case of arithmetic, with function symbols for zero, successor and all primitive recursive functions. Accordingly, we rename the base type D to Nat . In addition to the terms of ST, we add terms $\text{rec } e e$ and $\text{while}_e \gamma[z] e e e$ to our grammar, where $\gamma[z]$ ranges over state formulas of SL with a specified free variable z . The typing rules for these new terms are

$$\frac{\Gamma \vdash s : X \quad \Gamma \vdash t : \text{Nat} \rightarrow X \rightarrow X}{\Gamma \vdash \text{rec } s t : \text{Nat} \rightarrow X}$$

for the recursor, while for the while loop we have

$$\frac{\Gamma \vdash r : \text{Nat} \rightarrow X \rightarrow X \quad \Gamma \vdash s : \text{Nat} \rightarrow X \rightarrow Y \quad \Gamma \vdash t : X \rightarrow Y \quad \Gamma \vdash u : \text{Nat}}{\Gamma \vdash \text{while}_u \gamma[z] r s t : X \rightarrow Y}$$

under the additional variable condition that $z \notin \Gamma$, but $x : \text{Nat} \in \Gamma$ for all free variables of $\gamma[z]$ outside of z . Note that we do not consider z a free variable of $\text{while}_a \gamma[z] r s t$, but rather a placeholder for the loop condition. In order to give the appropriate semantics to our terms, we must add to our metatheory S^ω axioms and rules for arithmetic in all finite types, including the ability to define functions of arbitrary type via recursion over the natural numbers, along the lines of E-HA $^\omega$ [Tro73] (though as before the precise details are not important). We then define:

- $[\text{rec } s t] \pi := \langle R_f, \pi_1 \rangle$ for $\langle f, \pi_1 \rangle := [t] \pi$, where

$$\begin{aligned} R_f 0 \pi &:= [s] \pi \\ R_f (n+1) \pi &:= g a \pi'_2 \text{ for } \langle a, \pi'_1 \rangle := R_f n \pi' \text{ and } \langle g, \pi'_2 \rangle := f n \pi'_1 \end{aligned} \quad (6.1)$$

- $[\text{while}_u \gamma[z] r s t] \pi := \langle L_{f,g,h} m, \pi_4 \rangle$ where $\langle f, \pi_1 \rangle := [r] \pi$, $\langle g, \pi_2 \rangle := [s] \pi_1$, $\langle h, \pi_3 \rangle := [t] \pi_2$ and $\langle m, \pi_4 \rangle := [u] \pi_3$, where

$$\begin{aligned} L_{f,g,h} 0 y \pi' &:= h y \pi' \\ L_{f,g,h} (n+1) y \pi' &:= \begin{cases} L_{f,g,h} n y' \pi_2 \text{ for } \langle a, \pi'_1 \rangle := f n \pi' \text{ and } \langle y', \pi'_2 \rangle := a y \pi'_1 & \text{if } [\gamma][n+1](\pi') \\ b y \pi'_1 \text{ for } \langle b, \pi'_1 \rangle := g n \pi' & \text{if } \neg[\gamma][n+1](\pi') \end{cases} \end{aligned} \quad (6.2)$$

where in the case distinctions, we would technically speaking need to use the characteristic function $\chi_\gamma \langle x_1, \dots, n, \dots, x_k \rangle$ for γ , with n substituted for the special free variable z .

6.3. The soundness theorem for arithmetic. We now need to show that the soundness proof for stateful predicate logic also holds in the extension to arithmetic.

Theorem 6.3. *The statement of Theorem 5.4 remains valid if we replace SL by SA and ST by ST_N .*

Proof. We need to extend the proof of Theorem 6.3 to show that the additional axioms and rules as in Figure 6.1 can be realized by a term of the form $[t]$ for t in ST_N .

- For the non-extensionality equality and arithmetic axioms this is straightforward due to the fact that these are also true in S^ω : For instance, given a realizer $[s]$ sr $\{\alpha \cdot u = v \cdot \beta\}$ and $[t]$ sr $\{\beta \cdot v = w \cdot \gamma\}$, we have that $[s \circ t]$ sr $\{\alpha \cdot u = v \wedge v = w \cdot \gamma\}$, and since from $u = v \wedge v = w$ we can infer $u = w$ in S^ω , it follows that $[p_1(s \circ t)]$ sr $\{\alpha \cdot u = w \cdot \gamma\}$. The other axioms and rules are even simpler.
- (*ext*) Extensionality is similarly simple: If $[s]$ sr $\{\alpha \cdot u = v \cdot \beta\}$ and $[t]$ sr $\{\beta \cdot A(u) \cdot \gamma(u)\}$, then $[\alpha](\pi)$ implies that $u = v$ and $[\beta](\pi_1)$ for $\langle \dots, \pi_1 \rangle := [s]\pi$, and therefore a sr $A(u)$ and $[\gamma](u)(\pi_2)$ for $\langle a, \pi_2 \rangle := [t]\pi_1$. Now applying extensionality in S^ω to the formula $T(x) := a$ sr $A(x) \wedge [\gamma](x)(\pi_2)$, from $u = v$ we have a sr $A(v)$ and $[\gamma](v)(\pi_2)$, and thus $[s \circ t]$ sr $\{\alpha \cdot u = v \wedge A(v) \cdot \gamma(v)\}$ and therefore $[p_2(s \circ t)]$ sr $\{\alpha \cdot A(v) \cdot \gamma(v)\}$.
- (*rec*) Suppose that s and $t(x, y)$ are such that $[s]$ sr $\{\alpha \cdot A(0) \cdot \beta(0)\}$ and

$$[t](x, y) \text{ sr } \{\beta(x) \cdot A(x+1) \cdot \beta(x+1)\}$$

assuming y sr $A(x)$. We show that $[\text{rec } s \lambda x, y. t(x, y)]$ sr $\{\gamma \cdot \forall x \{\alpha \cdot A(x) \cdot \beta(x)\} \cdot \gamma\}$ for any γ . Since $[\text{rec } s \lambda x, y. t(x, y)]\pi = \langle R_f, \pi \rangle$ for $f := \lambda x. [\lambda y. t(x, y)]$ and R_f as in (6.1), it suffices to show that for any $n : \text{Nat}$ we have

$$R_f n \text{ sr } \{\alpha \cdot A(n) \cdot \beta(n)\}$$

We prove this by induction: For the base case, we have $R_f 0 = [s]$ and the claim holds by assumption. For the induction step, let us assume that $[\alpha](\pi')$ holds, and so by the induction hypothesis we have a sr $A(n)$ and $[\beta(n)](\pi'_1)$ for $\langle a, \pi'_1 \rangle := R_f n$. Since $f n \pi'_1 = \langle g, \pi'_1 \rangle$ for $g := \lambda y. [t](n, y)$, we have that $R_f(n+1)\pi' = [t](n, a)\pi'_1$, and since by the property of $[t]$ we then have b sr $A(n+1)$ and $[\beta(n+1)](\pi'_2)$ for $\langle b, \pi'_2 \rangle := [t](n, a)\pi'_1$, we have shown that $R_f(n+1)$ sr $\{\alpha \cdot A(n+1) \cdot \beta(n+1)\}$, which completes the induction.

- (*while*) We suppose that
 - (1) $[r](x, y)$ sr $\{\gamma(x+1) \wedge \alpha(x+1) \cdot A(x) \cdot \alpha(x)\}$ assuming that y sr $A(x+1)$,
 - (2) $[s](x, y)$ sr $\{\neg \gamma(x+1) \wedge \alpha(x+1) \cdot B \cdot \beta\}$ assuming that y sr $A(x+1)$,
 - (3) $[t](y)$ sr $\{\alpha(0) \cdot B \cdot \beta\}$ assuming that y sr $A(0)$.

Our aim is to show that

$$[(\text{while}_x \gamma (\lambda x', y'. r) (\lambda x', y'. s) (\lambda y'. t))y] \text{ sr } \{\alpha(x) \cdot B \cdot \beta\}$$

for any $x, y \in \text{Nat}$ with y sr $A(x)$. We observe, unwinding the definition, that

$$[(\text{while}_x \gamma (\lambda x', y'. r) (\lambda x', y'. s) (\lambda y'. t))y]\pi = L_{f,g,h} x y \pi$$

for $f := \lambda x'. [\lambda y'. r(x', y')]$, $g := \lambda x'. [\lambda y'. s(x', y')]$, $h := \lambda y'. [t](y')$ and $L_{f,g,h}$ as defined in (6.2). We now show by induction on n that if y sr $A(n)$ then

$$L_{f,g,h} n y \text{ sr } \{\alpha(n) \cdot B \cdot \beta\}$$

and then the result follows by setting $n := x$. The base case is straightforward since

$$L_{f,g,y}0y = [t](y)$$

and the claim follows by definition of $[t]$. For the induction step, suppose that y sr $A(n+1)$ and $[\alpha(n+1)](\pi)$. There are two cases. If $\neg[\gamma](n+1)(\pi)$ we have

$$L_{f,g,h}(n+1)y\pi = [s](n,y)\pi$$

and the result holds by the property of $[s]$. On the other hand, if $[\gamma](n+1)(\pi)$ then

$$L_{f,g,h}(n+1)y\pi = L_{f,g,h}ny'\pi'$$

for $\langle y', \pi' \rangle := [r](n,y)\pi$. But by the property of $[r]$ we have y' sr $A(n)$ and $[\alpha(n)](\pi')$, and therefore by the induction hypothesis we have b sr B and $[\beta](\pi'')$ for $\langle b, \pi'' \rangle := L_{f,g,h}ny'\pi' = L_{f,g,h}(n+1)y\pi$, and so the result is proven for $n+1$.

This covers all the additional axioms and rules of SA. □

6.4. Worked example: Insertion sort. We now illustrate our extended system by synthesising a list sorting program that, intuitively, forms an implementation of the insertion sort algorithm. Here our state will represent the structure that is to be sorted, and continuing the spirit of generality that we have adhered to throughout, we characterise this structure through a number of abstract axioms. Instantiating the state as, say, an array of natural numbers, would provide a model for our theory, but our sorting algorithm can be extracted on the more abstract level. Crucially, the proof involves both loop iteration and induction, and the corresponding program combines an imperative while loop with a functional recursor.

We begin by axiomatising our state, just as in previous examples. An intuition here is that states represent an infinite array of elements a_0, a_1, \dots possessing some total order structure \leq , and we seek to extract a program that, for any input n , sorts the first n elements. We use this informal semantics throughout to indicate the intended meaning of our axioms, but stress that none of this plays a formal role in the proof or resulting computational interpretation.

We introduce three state predicates to SA, with the intuition indicated in each case:

- $\text{sort}(N)$ – *Sorted: The first $N+1$ elements of the array i.e. $[a_0, \dots, a_N]$ are sorted*
- $\text{psort}(n, N)$ – *Partially sorted with respect to a_n : if $n < N$ then the list*

$$[a_0, \dots, a_{n-1}, a_{n+1}, \dots, a_N]$$

is sorted and $a_n \leq a_{n+1}$. For the base cases, if $n = N$ then the list $[a_0, \dots, a_{N-1}]$ is sorted, and if $n > N$ then the list $[a_0, \dots, a_N]$ is sorted.

- $\text{comp}(n)$ – *Comparison: true if $a_n \leq a_{n-1}$, and always true if $n = 0$*

We formalise this intuition by adding the following state independent axioms to Δ_H :

- (1) $\Gamma, \text{sort}(N) \vdash_H \text{psort}(N+1, N+1)$ – *If the first $N+1$ elements are sorted, then they are also partially sorted with respect to the next element a_{N+1} .*
- (2) $\Gamma, \neg \text{comp}(n), \text{psort}(n, N) \vdash_H \text{sort}(N)$ – *If $[a_0, \dots, a_{n-1}, a_{n+1}, \dots, a_N]$ is sorted, $a_n \leq a_{n+1}$, but also $a_{n-1} \leq a_n$, then the entire segment $[a_0, \dots, a_N]$ must be sorted.*
- (3) $\Gamma, \text{psort}(0, N) \vdash_H \text{sort}(N)$ – *If $[a_1, \dots, a_N]$ is sorted and $a_0 \leq a_1$, then $[a_0, \dots, a_N]$ is sorted.*

(4) $\Gamma \vdash_H \text{sort}(0)$ – *The singleton array $[a_0]$ is defined to be sorted.*

We complete the axiomatisation by adding a single state-sensitive axiom to Δ_S :

(5) $\Gamma \vdash_S \{\text{comp}(n+1) \wedge \text{psort}(n+1, N) \cdot \top \cdot \text{psort}(n, N)\}$ – *If $[a_0, \dots, a_n, a_{n+2}, \dots, a_N]$ is sorted and $a_{n+1} \leq a_{n+2}$, but $a_{n+1} \leq a_n$, then we can modify the state (i.e. swapping a_n and a_{n+1} by setting $\tilde{a}_n := a_{n+1}$ and $\tilde{a}_{n+1} := a_n$) so that $[a_0, \dots, a_{n-1}, \tilde{a}_{n+1}, \dots, a_N]$ is sorted and $\tilde{a}_n \leq \tilde{a}_{n+1}$. The edge cases for $n \geq N$ are interpreted in a more straightforward way.*

In order to give a realizing term to this axiom, we representing element swapping semantically by adding a constant $c : \text{Nat} \rightarrow S \rightarrow S$ to our metatheory S^ω , which satisfies

$$\text{comp}(n+1, \pi) \wedge \text{psort}(n+1, N, \pi) \Rightarrow \text{psort}(n, N, c n \pi)$$

and a corresponding term $\text{swap} : \text{Nat} \rightarrow C$ to our term calculus, along with the embedding

$$[\text{swap}] \pi := \langle \lambda n, \pi. \langle \cdot \rangle, c n \pi \rangle, \pi \simeq \langle c, \pi \rangle$$

so that we can prove

$$[\text{swap } n] \text{ sr } \{\text{comp}(n+1) \wedge \text{psort}(n+1, N) \cdot \top \cdot \text{psort}(n, N)\}$$

With this in place, we can now prove in SA that the first N elements of the state can be sorted, and extract a corresponding realizing term in ST_N .

6.4.1. *Proof of $\vdash_S \{\gamma \cdot \forall N \{\alpha \cdot \top \cdot \text{sort}(N)\} \cdot \gamma\}$ in SA.* The core of our proof begins with an instance of the *while* rule parametrised by N , with $\Gamma := \emptyset$, $A(n) := \top$, $\alpha(n) := \text{psort}(n, N+1)$, $\beta := \text{sort}(N+1)$ and $\gamma(n) := \text{comp}(n)$:

$$\frac{\frac{\frac{\mathcal{D}_1 \quad \mathcal{D}_2 \quad \mathcal{D}_3}{\top \vdash_S \{\text{psort}(n, N+1) \cdot \top \cdot \text{sort}(N+1)\}} \text{while}}{\top \vdash_S \{\text{psort}(N+1, N+1) \cdot \forall n \{\text{psort}(n, N+1) \cdot \top \cdot \text{sort}(N+1)\} \cdot \text{psort}(N+1, N+1)\}} \forall_S I}{\top \vdash_S \{\text{psort}(N+1, N+1) \cdot \top \cdot \text{sort}(N+1)\}} \forall_S E}{\top \vdash_S \{\text{sort}(N) \cdot \top \cdot \text{sort}(N+1)\}} \text{cons}$$

where the final composition inference makes use of the first state independent axiom. Here \mathcal{D}_1 represents an instance of the state sensitive axiom

$$\top \vdash_S \{\text{comp}(n+1) \wedge \text{psort}(n+1, N+1) \cdot \top \cdot \text{psort}(n, N+1)\}$$

and \mathcal{D}_2 represents the derivation

$$\frac{\top \vdash_S \{\text{sort}(N+1) \cdot \top \cdot \text{sort}(N+1)\}}{\top \vdash_S \{\neg \text{comp}(n+1) \wedge \text{psort}(n+1, N+1) \cdot \top \cdot \text{sort}(N+1)\}} \text{cons}$$

where composition makes use of the second state independent axiom. Finally \mathcal{D}_3 is

$$\frac{\top \vdash_S \{\text{sort}(N+1) \cdot \top \cdot \text{sort}(N+1)\}}{\top \vdash_S \{\text{psort}(0, N+1) \cdot \top \cdot \text{sort}(N+1)\}} \text{cons}$$

this time making use of the third state independent axiom. Finally we can prove that all lists can be sorted with an outer induction as follows:

$$\frac{\frac{\frac{\vdash_S \{\alpha \cdot \top \cdot \alpha\}}{\vdash_S \{\alpha \cdot \top \cdot \text{sort}(0)\}} \text{cons} \quad \mathcal{D}}{\vdash_S \{\gamma \cdot \forall N \{\alpha \cdot \top \cdot \text{sort}(N)\} \cdot \gamma\}} \text{ind}}{\vdash_S \{\alpha \cdot \top \cdot \text{sort}(0)\}} \text{cons} \quad \frac{\mathcal{D}}{\top \vdash_S \{\text{sort}(N) \cdot \top \cdot \text{sort}(N+1)\}} \text{ind}}{\vdash_S \{\gamma \cdot \forall N \{\alpha \cdot \top \cdot \text{sort}(N)\} \cdot \gamma\}} \text{ind}$$

where α is an arbitrary state predicate, the instance of *cons* uses the fourth state independent axiom, and \mathcal{D} represents the derivation above.

6.4.2. *Program extraction.* We now extract a program that corresponds to the above proof. First of all, we note that the three premises of our while rule are realised by *swap* n , *skip* and *skip* respectively, and so our derivation \mathcal{D} corresponds to the following program:

$$\frac{\frac{\frac{\frac{y : C \vdash \text{swap } n : C \quad y : C \vdash \text{skip} : C \quad y : C \vdash \text{skip} : C}{y : C \vdash t(n)y : C} \text{while}}{y : C \vdash \lambda n.t(n)y : C} \forall_S I}{y : C \vdash (\lambda n.t(n)y)(N+1) : C} \forall_S E}{y : C \vdash (\lambda n.t(n)y)(N+1) : C} \text{cons}$$

where

$$\begin{aligned} t(n) &:= \text{while}_n \text{comp}[z] (\lambda x, y. (\text{swap } x)) (\lambda x, y. \text{skip}) (\lambda y. \text{skip}) \\ &\simeq \text{while}_n \text{comp}[z] (\lambda x. (\text{swap } x)) (\text{skip}) (\text{skip}) \end{aligned}$$

Then our final induction generates the following program:

$$\frac{\vdash \text{skip} : C \quad y : C \vdash (\lambda n.t(n)y)(N+1) : C}{\vdash \text{rec} (\text{skip}) (\lambda x, y. ((\lambda n.t(n)y)(x+1))) : \text{Nat} \rightarrow C} \text{ind}$$

Thus our list sorting program is

$$\begin{aligned} &\text{rec} (\text{skip}) (\lambda x, y. ((\lambda n.t(n)y)(x+1))) \\ &\simeq \text{rec} (\text{skip}) (\lambda x. ((\lambda n. (\text{while}_n \text{comp}[z] (\lambda x. (\text{swap } x)) (\text{skip}) (\text{skip}))) (x+1))) \end{aligned}$$

which is essentially an implementation of the insertion sort algorithm, with an outer recursion that sorts initial segments of the list in turn, and an inner loop that inserts new elements into the appropriate place in the current sorted list.

7. DIRECTIONS FOR FUTURE WORK

In this paper we have presented the central ideas behind a new method for extracting stateful programs from proofs, which include an extension of ordinary first-order logic with Hoare triples, a corresponding realizability interpretation, and a soundness theorem. We emphasise once again that our intention has been to offer an alternative approach to connecting proofs with stateful programs, one that seeks to complement rather than improve existing work by embracing simplicity and abstraction, and which might be well suited to a range of applications in proof theory or computability theory. In this spirit, we conclude with a very informal outline of a series interesting directions in which we anticipate that our framework could be applied.

7.1. Further extensions and program synthesis. While our main results have been presented in the neutral setting of first-order predicate logic, it would be straightforward to extend SL to richer logics with more complex data structures and imperative commands. Already, the addition of recursion and loops over natural numbers in Section 6 has allowed us to synthesise a standard in-place sorting algorithm using our abstract axiomatisation of an ordered state, in a similar spirit to [BSW14]. However, further extensions are naturally possible, including the addition of general fixpoint operators and non-controlled while loops, which would then require a S^ω to be replaced by a domain theoretic semantics that allows for partiality.

Looking a step further ahead, by implementing all of this in a proof assistant, we would have at our disposal a new technique for synthesising correct-by-construction imperative programs. While we do not suggest that this pipeline would directly compete with existing techniques for verifying imperative programs, it could be well suited to synthesising and reasoning about programs in very specific domains, where we are interested in algorithms for which interactions with the state have a restricted form that could be suitably axiomatised within our logic. For example, a more detailed axiomatisation our state as an ordered array along the lines of Section 6.4, with a “swap” operation and a few other ways of interacting with the state, might give rise to an interesting theory of in-place sort algorithms. Stateful algorithms on other data structures, such as graphs, could presumably also be formalised within our framework.

7.2. Bar recursion and the semantics of extracted programs. Two of the main starting points for this paper, the monadic realizability of Birolo [Bir12] and the author’s own Dialectica interpretation with state [Pow18], address the broader problem of trying to understand the operational semantics of programs extracted from proofs as stateful procedures (the origins and development of this general idea, from Hilbert’s epsilon calculus onwards, is brilliantly elucidated in Chapter 1 of Aschieri’s thesis [Asc11], who then sets out his own realizability interpretation based on learning). A number of case studies by the author and others [OP15a, OP15b, Pow20, PSW22] have demonstrated that while terms extracted from nontrivial proofs can be extremely complex, they are often much easier to understand if one focuses on the way they interact with the mathematical environment. For example, in understanding a program extracted from a proof using Ramsey’s theorem for pairs [OP15a], it could be illuminating to study the *trace* of the program as it queries a colouring at particular pairs, as this can lead to a simpler characterisation of the *algorithm* ultimately being implemented by the term.

While the aforementioned analysis of programs has always been done in an informal way, our stateful realizability interpretation would in theory allow us to extract programs which store this trace formally in the state, where our abstract characterisation of state would allow us to implement it in whichever way is helpful in a given setting. For example, in the case of the Bolzano-Weierstrass theorem [OP15b], our state might record information of the form $x_n \in I$, collecting information about the location of sequence elements. For applications in algebra [PSW22], one might instead store information about a particular maximal ideal.

The aforementioned theorems are typically proven using some form of choice or comprehension, and that in itself leads to the interesting prospect of introducing both stateful recursors and while-loops that are computationally equivalent to variants of *bar recursion* [Spe62]. In [Pow16], several bar recursive programs that arise from giving a computational

interpretation to arithmetical comprehension principles are formulated as simple while loops, and these could in principle be incorporated into our system with new controlled Hoare rules in the style of update recursion [Ber04], that replace the conditions $n < N$ and $n \geq N$ in the \mathcal{A}_i above with e.g. $n \in \text{dom}(f)$ and $n \notin \text{dom}(f)$, where f is some partial approximation to a comprehension function. An exploration of such while-loops from the perspective of higher-order computability theory might well be of interest in its own right.

7.3. A logic for probabilistic lambda calculi. Probabilistic functional languages are a major topic of research at present. While work in this direction dates back to the late 1970s [JP89, SD78] where it typically had a semantic flavour, a more recent theme [DLZ12, DP95, DPHW05] has been to study simple extensions of the lambda calculus with nondeterministic choice operators \oplus , where $s \oplus t$ evaluates nondeterministically (or probabilistically) to either s or t . While such calculi have been extensively studied, corresponding *logics* that map under some proof interpretation to probabilistic programs are far more rare (although there is some recent work in this direction e.g. [ADLP22]).

We conjecture that our framework offers a bridge between logic and probabilistic computation through incorporating probabilistic disjunctions into our logic SL and taking states to be streams of outcomes of probabilistic events together with a current ‘counter’ that increases each time an event occurs. In a simple setting where only two outcomes are possible with equal probability, we can axiomatise this within SL by adding zero and successor functions (allowing us to create numerals n), along with a unary state predicate $\text{count}(n)$. We can then model probabilistic events by adding the appropriate axioms to Δ_S . Suppose, for example, we add two predicate constants $H(x)$ and $T(x)$ (for *heads* and *tails*), along with constants c_1, c_2, \dots representing coins. Then flipping a coin would be represented by the axiom schema

$$\Gamma \vdash_S \{ \text{count}(n) \cdot H(c_i) \vee T(c_i) \cdot \text{count}(n+1) \}$$

where n ranges over numerals and c_i over coin constants, the counter indicating that a probabilistic event has occurred. The act of reading a probability from the state could be interpreted semantically by introducing a constant $\omega : S \rightarrow \text{Bool} \times S$ to S^ω , with the axiom

$$\text{count}(n, \pi) \Rightarrow (e = \text{f} \Rightarrow H(c_i)) \wedge (e = \text{t} \Rightarrow T(c_i)) \wedge \text{count}(n+1, \pi_1) \text{ for } \langle e, \pi_1 \rangle := \omega\pi$$

(alternatively, we could simply define $S := \text{Nat} \times (\text{Nat} \rightarrow \text{Bool})$ for a type of Nat natural numbers, and define $\omega \langle n, a \rangle := \langle a(n), \langle n+1, a \rangle \rangle$ and $\text{count}(n, \langle m, a \rangle) := m =_{\text{Nat}} n$).

A probabilistic choice operator \oplus can then be added to the language of ST, along with the typing rule $\Gamma \vdash s \oplus t : X + Y$ for $\Gamma \vdash s : X$ and $\Gamma \vdash t : Y$, and the interpretation

$$[s \oplus t]\pi := \text{case } e \text{ } ([\iota_0 s]\pi_1) ([\iota_1 t]\pi_1) \text{ where } \langle e, \pi_1 \rangle := \omega\pi$$

In particular, defining $\text{flip} := \text{skip} \oplus \text{skip} : C + C$ we would have

$$[\text{flip}] \text{sr } \{ \text{count}(n) \cdot H(c_i) \vee T(c_i) \cdot \text{count}(n+1) \}$$

although we stress that the operator \oplus and would allow for much more complex probabilistic disjunctions, potentially involving additional computational content.

Our soundness theorem, extended to these new probabilistic axioms and terms, would then facilitate the extraction of probabilistic programs from proofs. For instance, including

a winner predicate $W(x)$, two player constant symbols p_1, p_2 , and adding axioms

$$\begin{aligned} H(c_1), H(c_2) &\vdash_S \{\alpha \cdot W(p_1) \cdot \alpha\} \\ T(c_1), T(c_2) &\vdash_S \{\alpha \cdot W(p_1) \cdot \alpha\} \\ H(c_1), T(c_2) &\vdash_S \{\alpha \cdot W(p_2) \cdot \alpha\} \\ T(c_1), H(c_2) &\vdash_S \{\alpha \cdot W(p_2) \cdot \alpha\} \end{aligned}$$

for any α , we could prove

$$\vdash_S \{\text{count}(n) \cdot \exists x W(x) \cdot \text{count}(n + 2)\}$$

expressing the fact that a winner can be determined after two flips. We can then extract a corresponding probabilistic term for realizing this statement, which would be isomorphic to the expected program that queries the state twice in order to determine the outcome of those flips, and returns either p_1 or p_2 as a realizer for $\exists x W(x)$ depending on the content of the state.

Of course, the details here need to be worked through carefully in order to properly substantiate the claim that our framework could be used to extract probabilistic programs in a natural and meaningful way. At the very least, it is likely that further additions to SL along with a more intricate state would be needed to incorporate more interesting probabilistic events, such as annotated disjunctions along the lines of [VVB04]. We leave such matters to future work.

ACKNOWLEDGMENTS

The author thanks the anonymous referees who provided comments on earlier versions of this paper, which resulted in a much improved final version. This research was supported by the Engineering and Physical Sciences Research Council grant number EP/W035847/1.

REFERENCES

- [ADLP22] Melissa Antonelli, Ugo Dal Lago, and Paolo Pistone. Curry and Howard meet Borel. In *Proceedings of Logic in Computer Science (LICS '22)*, page 13 pages. ACM, 2022. doi:10.1145/3531130.3533361.
- [Asc11] Federico Aschieri. *Learning, Realizability and Games in Classical Arithmetic*. PhD thesis, Università degli Studi di Torino and Queen Mary, University of London, 2011. doi:10.48550/arXiv.1012.4992.
- [Atk09] Robert Atkey. Parameterised notions of computation. *Journal of Functional Programming*, 19(3&4):335–376, 2009. doi:10.1017/S095679680900728X.
- [Bd09] Stefano Berardi and Ugo de'Liguoro. Toward the interpretation of non-constructive reasoning as non-monotonic learning. *Information and Computation*, 207(1):63–81, 2009. doi:10.1016/j.ic.2008.10.003.
- [Ber04] Ulrich Berger. A computational interpretation of open induction. In *Proceedings of Logic in Computer Science (LICS '04)*, pages 326–334. IEEE, 2004.
- [Bir12] Giovanni Birolò. *Interactive Realizability, Monads and Witness Extraction*. PhD thesis, Università degli Studi di Torino, 2012. doi:10.48550/arXiv.1304.4091.
- [BSW14] Ulrich Berger, Monika Seisenberger, and Gregory J. M. Woods. Extracting imperative programs from proofs: In-place quicksort. In *Proceedings of Types for Proofs and Programs (TYPES'13)*, volume 26 of *LIPICs*, pages 84–106, 2014. doi:10.4230/LIPICs.TYPES.2013.84.
- [BT21] Ulrich Berger and Hideki Tsuiki. Intuitionistic fixed point logic. *Annals of Pure and Applied Logic*, 172:102903, 2021. doi:0.1016/j.apal.2020.102903.

- [CLM13] Martin Churchill, James Laird, and Guy McCusker. Imperative programs as proofs via game semantics. *Annals of Pure and Applied Logic*, 164(11):1038–1078, 2013. doi:10.1016/j.apal.2013.05.005.
- [CMM⁺09] Adam Chlipala, Gregory Malecha, Greg Morrisett, Avraham Shinnar, and Ryan Wisnesky. Effective interactive proofs for higher-order imperative programs. In *Proceedings of International Conference on Functional Programming (ICFP'09)*, pages 79–90. ACM, 2009. doi:10.1145/1596550.1596565.
- [DF89] Olivier Danvy and Andrzej Filinski. A functional abstraction of typed contexts. Technical Report 89/12, BRICS, 1989.
- [DLR15] Norman Danner, Dan Licata, and Ramyaa Ramyaa. Denotational cost semantics for functional languages with inductive types. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*, pages 140–151, 2015. doi:10.1145/2784731.2784749.
- [DLZ12] Ugo Dal Lago and Margherita Zorzi. Probabilistic operational semantics for the lambda calculus. *RAIRO—Theoretical Informatics and Applications*, 46(3):413–450, 2012. doi:10.1051/ita/2012012.
- [DP95] Ugo Deliguoro and Adolfo Piperno. Nondeterministic extensions of untyped λ -calculus. *Information and Computation*, 122(2):149–177, 1995. doi:10.1006/inco.1995.1145.
- [DPHW05] Alessandra Di Pierro, Chris Hankin, and Herbert Wiklicky. Probabilistic lambda-calculus and quantitative program analysis. *Journal of Logic and Computation*, 15(2):159–179, 2005. doi:10.1093/logcom/exi008.
- [Fil03] Jean-Christoph Filliâtre. Verification of non-functional programs using interpretations in type theory. *Journal of Functional Programming*, 13(4):709–745, 2003. doi:10.1017/S095679680200446X.
- [GK05] Philipp Gerhardy and Ulrich Kohlenbach. Extracting Herbrand disjunctions by functional interpretation. *Archive for Mathematical Logic*, 44:633–644, 2005. doi:10.1007/s00153-005-0275-1.
- [Hoa69] Charles Antony Richard Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 1969. doi:10.1145/363235.363259.
- [JP89] Claire Jones and Gordon Plotkin. A probabilistic powerdomain of evaluations. In *Proceedings of Logic in Computer Science (LICS'89)*, pages 186–195. IEEE Press, 1989.
- [Kre59] Georg Kreisel. Interpretation of analysis by means of functionals of finite type. In A. Heyting, editor, *Constructivity in Mathematics*, pages 101–128. North-Holland, Amsterdam, 1959.
- [Kri09] Jean-Louis Krivine. Realizability in classical logic in interactive models of computation and program behaviour. *Panoramas et synthèses*, 27, 2009.
- [McC10] Guy McCusker. A graph model for imperative computation. *Logical Methods in Computer Science*, 6(1:2), 2010. doi:10.2168/LMCS-6(1:2)2010.
- [NAMB07] Aleksandar Nanevski, Amal Ahmed, Greg Morrisett, and Lars Birkedal. Abstract predicates and mutable ADTs in Hoare type theory. In *Proceedings of the European Symposium on Programming (ESOP'07)*, volume 4421 of *LNCS*, pages 189–204, 2007. doi:10.1007/978-3-540-71316-6_14.
- [OP15a] Paulo Oliva and Thomas Powell. A constructive interpretation of Ramsey’s theorem via the product of selection functions. *Mathematical Structure in Computer Science*, 25(8):1755–1778, 2015. doi:10.1017/S0960129513000340.
- [OP15b] Paulo Oliva and Thomas Powell. A game-theoretic computational interpretation of proofs in classical analysis. In Reinhard Kahle and Michael Rathjen, editors, *Gentzen’s Centenary: The Quest for Consistency*, pages 501–531. Springer, 2015. doi:10.1007/978-3-319-10103-3_18.
- [PCW05] Iman Poernomo, John N. Crossley, and Martin Wirsing. *Adapting Proofs-as-Programs*. Monographs in Computer Science. Springer, 2005. doi:10.1007/0-387-28183-5.
- [Pow16] Thomas Powell. Gödel’s functional interpretation and the concept of learning. In *Proceedings of Logic in Computer Science (LICS '16)*, pages 136–145. ACM, 2016. doi:10.1145/2933575.2933605.
- [Pow18] Thomas Powell. A functional interpretation with state. In *Proceedings of Logic in Computer Science (LICS '18)*, pages 839–848. ACM, 2018. doi:10.1145/3209108.3209134.
- [Pow20] Thomas Powell. Well quasi-orders and the functional interpretation. In Peter Schuster, Monika Seisenberger, and Andreas Weiermann, editors, *Well Quasi-Orders in Computation, Logic, Language and Reasoning*, volume 53 of *Trends in Logic*, pages 221–269. Springer, 2020. doi:10.1007/978-3-030-30229-0_9.

- [PSW22] Thomas Powell, Peter Schuster, and Franziskus Wiesnet. A universal algorithm for krull's theorem. *Information and Computation*, 287:104761, 2022. doi:10.1016/j.ic.2021.104761.
- [Red96] Uday Reddy. Global state considered unnecessary. *Lisp and Symbolic Computation*, 9:7–76, 1996.
- [Rey02] John C. Reynolds. Separation logic: a logic for shared mutable data structures. In *Proceedings of Logic in Computer Science (LICS'02)*, pages 55–74. IEEE, 2002. doi:10.1109/LICS.2002.1029817.
- [SD78] Nasser Saheb-Djaromi. Probabilistic LCF. In *Proceedings of International Symposium on Mathematical Foundations of Computer Science (MFCS'78)*, volume 64 of *LNCS*, pages 442–451, 1978.
- [Spe62] Clifford Spector. Provably recursive functionals of analysis: a consistency proof of analysis by an extension of principles in current intuitionistic mathematics. In F. D. E. Dekker, editor, *Recursive Function Theory: Proc. Symposia in Pure Mathematics*, volume 5, pages 1–27. American Mathematical Society, 1962.
- [Tro73] Anne S. Troelstra. *Metamathematical Investigation of Intuitionistic Arithmetic and Analysis*, volume 344 of *Lecture Notes in Mathematics*. Springer-Verlag, 1973. doi:10.1007/BFb0066739.
- [VVB04] Joost Vennekens, Sofie Verbaeten, and Maurice Brunooghe. Logic programs with annotated disjunctions. In *ICLP 2004: Logic Programming*, volume 3132 of *Lecture Notes in Computer Science*, pages 431–445, 2004. doi:10.1007/978-3-540-27775-0_30.
- [YHB07] Nobuko Yoshida, Kohei Honda, and Martin Berger. Logical reasoning for higher-order functions with local state. In *Proceedings of Foundations of Software Science and Computation Structures (FOSSACS'07)*, volume 44213 of *LNCS*, pages 361–377, 2007. doi:10.1007/978-3-540-71389-0_26.