

Relational Semantics for Effect-Based Program Transformations with Dynamic Allocation

Nick Benton Andrew Kennedy

Microsoft Research, Cambridge
{nick,akenn}@microsoft.com

Lennart Beringer Martin Hofmann

LMU, Munich
{beringer,mhofmann}@tcs.ifi.lmu.de

Abstract

We give a denotational semantics to a region-based effect system tracking reading, writing and allocation in a higher-order language with dynamically allocated integer references.

Effects are interpreted in terms of the preservation of certain binary relations on the store, parameterized by region-indexed partial bijections on locations.

The semantics validates a number of effect-dependent program equivalences and can thus serve as a foundation for effect-based compiler transformations.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features – Dynamic storage management; F.3.2 [Logic and Meanings of Programs]: Semantics of Programming Languages – Denotational semantics, Program analysis; F.3.2 [Logic and Meanings of Programs]: Studies of Program Constructs – Type structure

General Terms Languages, Theory

Keywords Type and effect systems, region analysis, logical relations, parametricity, program transformation

1. Introduction

Many analyses and logics for imperative programs are concerned with establishing whether particular mutable variables may be read or written by a phrase. For example, the equivalence of while-programs

$$\begin{aligned} & C ; \text{if } B \text{ then } C' \text{ else } C'' = \\ & \text{if } B \text{ then } (C;C') \text{ else } (C;C'') \end{aligned}$$

is valid when B does not read any variable which C might write. Hoare-style programming logics often have rules with side conditions on possibly-read and possibly-written variable sets, and reasoning about concurrent processes is dramatically simplified if one can establish that none of them may write a variable which another may read.

Effect systems [13, 17] are static analyses that compute upper bounds on the possible side-effects of computations. The literature contains many effect systems that analyse which storage cells may be read and which storage cells may be written (as well as many

other properties), but few satisfactory accounts of the semantics of this information, or of the uses to which it may be put. Note that because effect systems *over-estimate* the possible side-effects of expressions, the information they capture is of the form that particular variables will definitely *not* be read or will definitely *not* be written. But what does that mean?

Thinking operationally, it may seem entirely obvious what is meant by saying that a variable X will not be read (written) by a command C , viz. no execution trace of C contains a read (resp. write) operation to X . But, as we have argued before [4, 7], such intensional interpretations of program properties are over-restrictive, cannot be interpreted in a standard semantics, do not behave well with respect to program equivalence or contextual reasoning and are hard to maintain during transformations. Thus we seek extensional properties that are more liberal than the intensional ones yet still validate the transformations or reasoning principles we wish to apply.

In a previous paper [8] we successfully solved this problem for a simple imperative language with global references. The main idea of that work was that a computation M preserves a store relation R provided this relation is preserved by all the side-effecting operations that M may possibly invoke. In particular, the operation “read from X ” preserves all store relations R with the property that sRs' implies $s(X) = s'(X)$. The operation “write to Y ”, on the other hand, preserves all store relations R with the property that sRs' implies $s[Y:=n]Rs'[Y:=n]$ for all $n \in \mathbb{Z}$. Thus, if an effect system ascribes *only* the effects “read from X ” and “write to Y ” to a computation M then M should preserve all the relations that enjoy both of those properties. This idea extends to a compositional semantics which validates a list of effect-based program transformations, among them lifting semantically pure computations out of functions, and eliminating dead code.

In the present work we substantially extend this idea by allowing dynamic allocation of references. Accordingly, effects can no longer refer to explicit references (“read from X ”) but abstract from individual references using the concept of *regions* as first introduced by Lucassen and Gifford [17]. A major motivation in that work was in parallelizing compilation, though such optimizations were not presented, still less formalized. In our work we validate program equivalences for sequential programs that depend on fine-grained effect information, as in the rule for commuting computations (in C-style notation):

$$\begin{aligned} & x=m1() ; y=m2() ; m3(x, y) ; \equiv \\ & y=m2() ; x=m1() ; m3(x, y) ; \end{aligned}$$

which requires that $m1$ does not read from or write to a region that $m2$ writes to and vice versa. Notice that the values x, y that are computed in different order may themselves be functions possibly embodying (freshly allocated) references.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PPDP'07 July 14–16, 2007, Wrocław, Poland.
Copyright © 2007 ACM 978-1-59593-769-8/07/0007...\$5.00

Notice that the equivalences we prove are generic, in that they hold for all terms with a particular effect-annotated type. Many interesting concrete instances can be shown using existing techniques, but devising schematic equational principles, driven by types rather than terms, is more challenging.

Another important addition to our earlier work is the inclusion of Lucassen-Gifford’s “masking rule” which under certain circumstances allows one to remove manifestly present effect information. Effect masking was used later by Tofte and Talpin in region-based memory management [28] in order to statically infer sound deallocations. Here, we are not interested in such intensional properties of programs; rather, the masking rule merely hides non-observable side-effects, validating more program equivalences than would otherwise hold.

Our approach to the soundness of the effect analysis and to the correctness of effect-dependent program equivalences is to interpret program properties (which may be expressed as points in an abstract domain, or as non-standard types) as binary relations over a standard, non-instrumented (operational or denotational) semantics of the language. Unlike in our earlier work on global store these binary relations are now parameterized over a partially-ordered set of region layouts (“parameters”); a Kripke-style quantification over parameters is used in the definition of the relation for function types.

We stress that the contribution of this paper is not merely methodological; as far as we are aware the program equivalences in Section 7 are genuinely new results. Note that these results are phrased in terms of observational equivalences and effect typing, thus do not refer in their statement to concepts introduced here. Of course, the relational semantics which is the main technical contribution plays a crucial role in the proofs of these statements.

Acknowledgements: Support by MS Research Cambridge, EU grant EmBounded IST-FET 510255 (Hofmann), EU grant MOBIUS IST-FET-15905 (Berlinger, Hofmann) is gratefully acknowledged.

2. Base Language

We study a monadically-typed, normalizing, call-by-value lambda calculus with dynamically allocated integer references. We thus extend the language from [8] but do not include recursion, references to types other than integers, etc. The use of monadic types, making an explicit distinction between values and computations, simplifies the presentation of the effect system and cleans up the equational theory of the language. A more conventionally-typed impure calculus may be translated into the monadic one via the usual ‘call-by-value translation’ [6], and this extends to the usual style of presenting effect systems in which every judgement has an effect, and function arrows are annotated with ‘latent effects’ [29].

We define value types (ranged over by A and B), computation types TA and contexts Γ as follows:

$$\begin{aligned} A, B &::= \text{unit} \mid \text{int} \mid \text{bool} \mid \text{ref} \mid A \times B \mid A \rightarrow TB \\ \Gamma &::= x_1 : A_1, \dots, x_n : A_n \end{aligned}$$

Variables are always given value types, as this is all we shall need to interpret a CBV language. There are two forms of typing judgement: value judgements $\Gamma \vdash V : A$ and computation judgements $\Gamma \vdash M : TA$, defined inductively by the rules in Figure 1. Note that the types on λ -bound variables make typing derivations unique and that addition and comparison are just representative primitive operations. To save space we omit grammars for values V and computations M as they can be inferred from the typing rules.

3. Denotational semantics

Since our simple language has no recursion, we can give it an elementary denotational semantics in the category of sets and functions.

We axiomatise states as follows, assuming a set \mathbb{L} of locations and a set \mathbb{S} of states. There is a constant $\emptyset \in \mathbb{S}$, the empty state. If $s \in \mathbb{S}$ then $\text{dom}(s) \subseteq \mathbb{L}$ and if $\ell \in \text{dom}(s)$ then $s.\ell \in \mathbb{Z}$ is a value; if $v \in \mathbb{Z}, \ell \in \text{dom}(s)$ then $s[\ell \mapsto v] \in \mathbb{S}$; finally $\text{new}(s, v)$ yields a pair (ℓ, s') where $\ell \in \mathbb{L}$ and $s' \in \mathbb{S}$. These operations are subject to the following axioms:

$$\begin{aligned} \text{dom}(\emptyset) &= \emptyset \\ \text{dom}(s[\ell \mapsto v]) &= \text{dom}(s) \\ (s[\ell \mapsto v]).\ell' &= \text{if } \ell = \ell' \text{ then } v \text{ else } s.\ell' \\ \text{new}(s, v) = (\ell, s') &\Rightarrow \text{dom}(s') = \text{dom}(s) \cup \{\ell\} \wedge \\ &\ell \notin \text{dom}(s) \wedge s'.\ell = v \end{aligned}$$

This abstract datatype can be implemented in a number of ways, e.g., as finite maps.¹

If L is a set of locations, typically finite, and $s, s' \in \mathbb{S}$ then we define

$$s \sim_L s' \Leftrightarrow \text{dom}(s) \supseteq L \wedge \text{dom}(s') \supseteq L \wedge \forall \ell \in L. s.\ell = s'.\ell$$

The semantics of types is now given as follows:

$$\begin{aligned} \llbracket \text{unit} \rrbracket &= \{\star\} & \llbracket \text{int} \rrbracket &= \mathbb{Z} & \llbracket \text{bool} \rrbracket &= \mathbb{B} \\ \llbracket \text{ref} \rrbracket &= \mathbb{L} & \llbracket A \times B \rrbracket &= \llbracket A \rrbracket \times \llbracket B \rrbracket \\ \llbracket A \rightarrow TB \rrbracket &= \llbracket A \rrbracket \rightarrow \llbracket TB \rrbracket & \llbracket TA \rrbracket &= \mathbb{S} \rightarrow \mathbb{S} \times \llbracket A \rrbracket \end{aligned}$$

The interpretation of the computation type constructor is the usual state monad. The meaning of contexts is given by $\llbracket \Gamma \rrbracket = \prod_{x \in \text{dom}(\Gamma)} \llbracket \Gamma(x) \rrbracket$, and we can then give the semantics of judgements

$$\begin{aligned} \llbracket \Gamma \vdash V : A \rrbracket &: \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket \\ \text{and } \llbracket \Gamma \vdash M : TA \rrbracket &: \llbracket \Gamma \rrbracket \rightarrow \llbracket TA \rrbracket \end{aligned}$$

inductively, though we omit the completely standard details.

Figure 2 defines sequential composition as syntactic sugar and contains a few example programs with their types. The first, V_{sum} computes the sum of its three arguments using a reference to store an intermediate result. Here int^3 abbreviates $\text{int} \times (\text{int} \times \text{int})$ and we use an obvious pattern-matching notation to avoid accumulation of projections. The second example generates a counter object initialised at zero. Note that each evaluation of M_{cnt} produces a different counter. The next, M_{mem} is a memoised version of the successor function. It should be clear how to extend this example to a memo-functional that can memoise an argument function. The last example, M_{buf} is a function that returns the previous argument it has been called with, or zero on first invocation.

4. Effect system

We now present our effect analysis as a type system that refines the simple type system by annotating the computation type constructor with information about whether a computation may read, write, or allocate within a region.

Formally, we assume a (possibly infinite) set Regs of region identifiers (regions for short) ranged over by r . Types and typing judgements will involve only finitely many regions. Such finite sets are ranged over by Π . Primitive effects are al_r (allocation in region r), wr_r (write access to region r), and rd_r (read access to region r). An effect, ranged over by ε , is a finite set of primitive effects. If ε is an effect, we define the *reads* of ε by $\text{rds}(\varepsilon) = \{r \mid rd_r \in \varepsilon\}$ and similarly its *writes* $\text{wrs}(\varepsilon)$ and its *allocations* $\text{als}(\varepsilon)$.

¹ This is not quite the free implementation because of the *new* operation.

$\Gamma \vdash n : \text{int}$	$\Gamma \vdash b : \text{bool}$	$\Gamma \vdash () : \text{unit}$	$\Gamma, x : A \vdash x : A$
$\Gamma \vdash V_1 : \text{int} \quad \Gamma \vdash V_2 : \text{int}$		$\Gamma \vdash V_1 : \text{int} \quad \Gamma \vdash V_2 : \text{int}$	
$\Gamma \vdash V_1 + V_2 : \text{int}$		$\Gamma \vdash V_1 > V_2 : \text{bool}$	
$\Gamma \vdash V_1 : A \quad \Gamma \vdash V_2 : B$		$\Gamma \vdash V : A_1 \times A_2$	
$\Gamma \vdash (V_1, V_2) : A \times B$		$\Gamma \vdash \pi_i V : A_i$	
$\Gamma, x : A \vdash M : TB$	$\Gamma \vdash V_1 : A \rightarrow TB \quad \Gamma \vdash V_2 : A$	$\Gamma \vdash V : A$	
$\Gamma \vdash \lambda x : A.M : A \rightarrow TB$	$\Gamma \vdash V_1 V_2 : TB$	$\Gamma \vdash \text{val } V : TA$	
$\Gamma \vdash M : TA \quad \Gamma, x : A \vdash N : TB$	$\Gamma \vdash V : \text{bool} \quad \Gamma \vdash M : TA \quad \Gamma \vdash N : TA$		
$\Gamma \vdash \text{let } x \leftarrow M \text{ in } N : TB$	$\Gamma \vdash \text{if } V \text{ then } M \text{ else } N : TA$		
$\Gamma \vdash V : \text{ref}$	$\Gamma \vdash V_1 : \text{ref} \quad V_2 : \text{int}$	$\Gamma \vdash V : \text{int}$	
$\Gamma \vdash \text{read}(V) : T\text{int}$	$\Gamma \vdash \text{write}(V_1, V_2) : T\text{unit}$	$\Gamma \vdash \text{ref}(V) : T\text{ref}$	

Figure 1. Simple computation type system

$M_1; M_2$:= <code>let $_ \leftarrow M_1$ in M_2</code> (sequential composition)
V_{sum}	:= <code>$\lambda(x, y, z) : \text{int}^3. \text{let } x \leftarrow \text{ref}(x + y) \text{ in } \text{read}(x) + z : \text{int}^3 \rightarrow T\text{int}$</code>
M_{cnt}	:= <code>let $x \leftarrow \text{ref}(0)$ in $(\lambda _ : \text{unit}. \text{read}(x), \lambda _ : \text{unit}. \text{write}(x, \text{read}(x) + 1)) :$</code> <code>$T((\text{unit} \rightarrow T\text{int}) \times (\text{unit} \rightarrow T\text{unit}))$</code>
M_{mem}	:= <code>let $x \leftarrow \text{ref}(0)$ in let $y \leftarrow \text{ref}(1)$ in</code> <code>$\lambda a : \text{int}. \text{if } a = \text{read}(x) \text{ then } \text{read}(y) \text{ else } \text{write}(x, a); \text{write}(y, a + 1); \text{read}(y) : T(\text{int} \rightarrow T\text{int})$</code>
M_{buf}	:= <code>let $x \leftarrow \text{ref}(0)$ in $\lambda a : \text{int}. \text{let } o \leftarrow \text{read}(x) \text{ in } \text{write}(x, a); o : T(\text{int} \rightarrow T\text{int})$</code>

Figure 2. Example programs

We define *refined* value types ranged over by X and Y :

$$X, Y := \text{unit} \mid \text{int} \mid \text{bool} \mid \text{ref}_r \mid X \times Y \mid X \rightarrow T_\varepsilon Y$$

A refined computation type takes the form $T_\varepsilon X$ with X a refined value type. A refined typing context Θ is a finite map from variables to refined value types. The well-formedness judgements $\Pi \vdash \varepsilon \text{ ok}$, $\Pi \vdash X \text{ ok}$ and $\Pi \vdash \Theta \text{ ok}$ mean that only regions declared in Π appear in ε , X and Θ . There is a subtyping relation on refined types, axiomatised in Figure 3. The erasure map, $U(\cdot)$, takes refined types to simple types (and contexts) by forgetting the effect annotations. We omit its obvious definition.

The refined type assignment system is shown in Figure 4. Note that the subject terms are the same (we still only have simple types on λ -bound variables).

Lemma 1. *If $\Pi; \Theta \vdash V : X$ then $U(\Theta) \vdash V : U(X)$, and similarly for computations.*

The last rule in Fig. 4 is called the *masking rule* which allows one to remove all reference to a region r in the effect of a computation provided this region is mentioned neither in the result type X of that computation nor in the type of any of its free variables. In the original work on region-based memory management [28] such regions could be safely deallocated upon completion of that computation. We do not consider deallocation in our system, and indeed the refined type system does not affect the semantics of computations in any way. Instead we use the masking rule to enhance applicability of effect-based program equivalences.

Note that any term typable in the original language can be typed in the refined system using a single region r and all effects set, i.e., $\Pi = \{r\}$ and $\varepsilon = \{rd_r, wr_r, al_r\}$ throughout. Of course, in order to maximise applicability of program equivalences and of the masking rule it is in the interest of the programmer or of an automatic type inference (which we do not consider here) to seek refined typings that use more than one region.

4.1 Example programs

Recall the example programs from Figure 2. These can be typed as follows.

$$\begin{aligned}
V_{\text{sum}} & : \text{int}^3 \rightarrow T_{\{al_r, rd_r\}} \text{int} \\
M_{\text{cnt}} & : T_{\{al_r\}} ((\text{unit} \rightarrow T_{\{rd_r\}} \text{int}) \times (\text{unit} \rightarrow T_{\{wr_r\}} \text{unit})) \\
M_{\text{mem}} & : T_{\{al_r\}} (\text{int} \rightarrow T_{\{wr_r, rd_r\}} \text{int}) \\
M_{\text{buf}} & : T_{\{al_r\}} (\text{int} \rightarrow T_{\{wr_r, rd_r\}} \text{int})
\end{aligned}$$

The function V_{sum} can even be typed as

$$V_{\text{sum}} : \text{int}^3 \rightarrow T_\emptyset \text{int}$$

Thus will be (correctly) regarded as pure by our analysis. Unfortunately, the masking rule does not apply to the type of the pure computation M_{mem} . Indeed, if it did, the impure M_{buf} which has the same type as M_{mem} would be falsely given a pure typing.

$$\begin{array}{c}
\frac{}{X \leq X} \qquad \frac{X \leq Y \quad Y \leq Z}{X \leq Z} \qquad \frac{X \leq X' \quad Y \leq Y'}{X \times Y \leq X' \times Y'} \\
\frac{X' \leq X \quad T_\varepsilon Y \leq T_{\varepsilon'} Y'}{(X \rightarrow T_\varepsilon Y) \leq (X' \rightarrow T_{\varepsilon'} Y')} \qquad \frac{\varepsilon \subseteq \varepsilon' \quad X \leq X'}{T_\varepsilon X \leq T_{\varepsilon'} X'}
\end{array}$$

Figure 3. Subtyping refined types

$$\begin{array}{c}
\frac{}{\Pi; \Theta \vdash n : \text{int}} \qquad \frac{}{\Pi; \Theta \vdash b : \text{bool}} \qquad \frac{}{\Pi; \Theta \vdash () : \text{unit}} \qquad \frac{\Pi \vdash X \text{ ok}}{\Pi; \Theta, x : X \vdash x : X} \\
\frac{\Pi; \Theta \vdash V_1 : \text{int} \quad \Pi; \Theta \vdash V_2 : \text{int}}{\Pi; \Theta \vdash V_1 + V_2 : \text{int}} \qquad \frac{\Pi; \Theta \vdash V_1 : \text{int} \quad \Pi; \Theta \vdash V_2 : \text{int}}{\Pi; \Theta \vdash V_1 > V_2 : \text{bool}} \\
\frac{\Pi; \Theta \vdash V_1 : X \quad \Pi; \Theta \vdash V_2 : Y}{\Pi; \Theta \vdash (V_1, V_2) : X \times Y} \qquad \frac{\Pi; \Theta \vdash V : X_1 \times X_2}{\Pi; \Theta \vdash \pi_i V : X_i} \qquad \frac{\Pi; \Theta, x : X \vdash M : T_\varepsilon Y}{\Pi; \Theta \vdash \lambda x : U(X).M : X \rightarrow T_\varepsilon Y} \\
\frac{\Pi; \Theta \vdash V_1 : X \rightarrow T_\varepsilon Y \quad \Pi; \Theta \vdash V_2 : X}{\Pi; \Theta \vdash V_1 V_2 : T_\varepsilon Y} \qquad \frac{\Pi; \Theta \vdash V : X}{\Pi; \Theta \vdash \text{val } V : T_\emptyset X} \\
\frac{\Pi; \Theta \vdash M : T_\varepsilon X \quad \Pi; \Theta, x : X \vdash N : T_{\varepsilon'} Y}{\Pi; \Theta \vdash \text{let } x \leftarrow M \text{ in } N : T_{\varepsilon \cup \varepsilon'} Y} \qquad \frac{\Pi; \Theta \vdash V : \text{bool} \quad \Pi; \Theta \vdash M : T_\varepsilon X \quad \Pi; \Theta \vdash N : T_\varepsilon X}{\Pi; \Theta \vdash \text{if } V \text{ then } M \text{ else } N : T_\varepsilon X} \\
\frac{\Pi; \Theta \vdash V : \text{ref}_r}{\Pi; \Theta \vdash \text{read}(V) : T_{\{rd_r\}}(\text{int})} \qquad \frac{\Pi; \Theta \vdash V_1 : \text{ref}_r \quad \Pi; \Theta \vdash V_2 : \text{int}}{\Pi; \Theta \vdash \text{write}(V_1, V_2) : T_{\{wr_r\}}(\text{unit})} \qquad \frac{\Pi; \Theta \vdash V : \text{int} \quad r \in \Pi}{\Pi; \Theta \vdash \text{ref}(V) : T_{\{al_r\}}(\text{ref}_r)} \\
\frac{\Pi; \Theta \vdash V : X \quad X \leq X'}{\Pi; \Theta \vdash V : X'} \qquad \frac{\Pi; \Theta \vdash M : T_\varepsilon X \quad T_\varepsilon X \leq T_{\varepsilon'} X'}{\Pi; \Theta \vdash M : T_{\varepsilon'} X'} \qquad \frac{\Pi, r; \Theta \vdash M : T_\varepsilon X \quad \Pi \vdash X \text{ ok} \quad \Pi \vdash \Theta \text{ ok}}{\Pi; \Theta \vdash M : T_{\varepsilon \setminus \{rd_r, wr_r, al_r\}} X}
\end{array}$$

Figure 4. Refined type system

A schematic example for the use of the masking rule in the spirit of Haskell’s “runST” [10] is as follows. Suppose that M is a computation of an end result of type int employing on its way a number of counters, i.e., invocations of M_{cnt} (and possibly other stateful operations). We then get a type $T_\varepsilon \text{int}$ for M which can then be improved to $T_\emptyset \text{int}$ assuming that no reference is made to externally visible portions of the store (formally: the regions mentioned in ε are disjoint from regions mentioned in the types of free variables in M).

All these typings only use a single region which, however, is arbitrary and can thus be chosen differently from regions used in the context. For example, we have

$$\text{let } x_1 \leftarrow M_{\text{buf}} \text{ in let } x_2 \leftarrow M_{\text{buf}} \text{ in } (x_1, x_2) : T_{\{al_{r_1}, al_{r_2}\}}(J(r_1) \times J(r_2))$$

where $J(r) = \text{int} \rightarrow T_{\{wr_r, rd_r\}} \text{int}$.

Remark on polymorphism It would be natural to assign a region-polymorphic type such as $\forall r. T_{\{al_r\}} J(r)$ to M_{buf} . In this paper we refrain from considering polymorphism (at the level of regions, effects or types) for the sake of simplicity. Our proof of type soundness can be extended to the effect-polymorphic case quite straightforwardly by interpreting polymorphic types as (possibly infinite) families of monomorphic types, as in [5].

5. Formal preliminaries

In our previous work [8] we approximated contextual equivalence by a partial equivalence relation $\llbracket X \rrbracket$ for each refined type X . In the presence of dynamic allocation such a simple-minded setup will no longer work and we move to a Kripke logical relation indexed by store layouts which we refer to as *parameters*. Parameters introduce (a) a ‘representation independence’ for state, capturing the fact that behaviour is invariant under permutation of locations; and (b) a distinction between observable and non-observable locations, as expressed syntactically by the masking rule.

5.1 Partial bijections

Aspect (a) of parameters is expressed by assigning to each region identifier a partial bijection between locations in the store.

Definition 1 (partial bijection). A partial bijection is a triple (L, L', f) where L, L' are finite subsets of \mathbb{L} and $f \subseteq L \times L'$ such that $(l_1, l'_1) \in f$ and $(l_2, l'_2) \in f$ imply $l_1 = l_2 \Leftrightarrow l'_1 = l'_2$.

If $t = (L, L', f)$ is a partial bijection, we write $\text{dom}(t) = L$, $\text{dom}'(t) = L'$ and refer to f simply by t itself. We let (ℓ, ℓ') denote the partial bijection $\{(\ell), (\ell')\}$, and let \emptyset denote the empty partial bijection.

Two partial bijections t_1, t_2 are *disjoint* if $\text{dom}(t_1) \cap \text{dom}(t_2) = \emptyset$ and $\text{dom}'(t_1) \cap \text{dom}'(t_2) = \emptyset$. In this case, we write $t_1 \otimes t_2$ for

the partial bijection given by

$$t_1 \otimes t_2 = \left(\begin{array}{c} \text{dom}(t_1) \cup \text{dom}(t_2), \\ \text{dom}'(t_1) \cup \text{dom}'(t_2), \\ t_1 \cup t_2 \end{array} \right)$$

Partial bijections are ordered as follows: $t' \geq t$ if and only if $t' = t \otimes t''$ for some (uniquely determined) t'' .

5.2 Parameters

For aspect (b) of parameters we introduce a special symbol $\tau \notin \text{Regs}$ to represent the part of the store arising from regions “masked out” by the masking rule. Commands must not alter this portion of the store at all. We will thus sometimes refer to τ as the *silent* region. This intended meaning will become clear subsequently; for now, τ is just a symbol.

We are now ready to give a formal definition of parameters.

Definition 2 (parameter). A parameter φ is a pair (Π, f) where

- Π is a finite set of regions,
- f is a function assigning to each region $r \in \Pi \cup \{\tau\}$ a partial bijection $f(r)$ such that distinct regions map to disjoint partial bijections.
- $f(\tau) = (L, L', \emptyset)$ for some L, L' .

A parameter is meant to lay out regions in pairwise disjoint parts of the store in two related computations. The partial bijections relate corresponding locations and are used to interpret equality of locations. Note that $\varphi(\tau)$ merely contains two sets of locations and no “links” at all. The silent region represents a store portion that must not be modified at all.

If $\varphi = (\Pi, f)$ is a parameter, we refer to its components as follows:

$$\begin{aligned} \text{regs}(\varphi) &= \Pi \\ \varphi(r) &= \begin{cases} f(r) & \text{when defined} \\ \emptyset & \text{otherwise} \end{cases} \\ \text{dom}(\varphi) &= \bigcup_{r \in \Pi \cup \{\tau\}} \text{dom}(\varphi(r)) \\ \text{dom}'(\varphi) &= \bigcup_{r \in \Pi \cup \{\tau\}} \text{dom}'(\varphi(r)) \end{aligned}$$

Note that $\text{dom}(\varphi)$ and $\text{dom}'(\varphi)$ are both finite. We define $\text{Par}(\Pi) = \{\varphi \mid \text{regs}(\varphi) \subseteq \Pi\}$. We also use the notation $\Pi \vdash \varphi \text{ ok}$ to mean $\varphi \in \text{Par}(\Pi)$. If φ and φ' are parameters such that

$$\text{dom}(\varphi) \cap \text{dom}(\varphi') = \text{dom}'(\varphi) \cap \text{dom}'(\varphi') = \emptyset$$

then φ, φ' are called *disjoint* and we write $\varphi \otimes \varphi'$ for the obvious juxtaposition of φ and φ' . Formally,

$$(\Pi, f) \otimes (\Pi', f') = (\Pi \cup \Pi', \lambda r. (f r) \otimes (f' r)) \in \text{Par}(\Pi \cup \Pi')$$

Whenever we write $\varphi \otimes \psi$ then φ and ψ are presumed to be disjoint from each other so, a statement like $\exists \psi. \dots \varphi \otimes \psi \dots$ is understood as “there exists ψ disjoint from φ such that $\dots \varphi \otimes \psi \dots$ ”.

Each set $\text{Par}(\Pi)$ is partially ordered by $\varphi' \geq_{\Pi} \varphi \iff \varphi' = \varphi \otimes \psi$ for some ψ .

If t is a partial bijection then $[r \mapsto t]$ is the parameter $\varphi \in \text{Par}(\{r\})$ such that $\varphi(r) = t, \varphi(r') = \emptyset$ when $r' \neq r$.

Thus, if $\ell \notin \text{dom}(\psi)$ and $\ell' \notin \text{dom}'(\psi)$ then we can form $\psi \otimes [r \mapsto (\ell, \ell')]$ to add the link (ℓ, ℓ') to r in ψ . Similarly, if $r \notin \Pi$, we can form $\varphi \otimes [r \mapsto \emptyset]$ to initialise a new region r with \emptyset .

If $\Pi, r \vdash \varphi \text{ ok}$ then $\varphi - r$ denotes the parameter such that $\Pi \vdash \varphi - r \text{ ok}$ and

$$\begin{aligned} (\varphi - r)(r') &= \varphi(r') \text{ when } r' \neq r \\ (\varphi - r)(\tau) &= \varphi(\tau) \otimes (\text{dom}(\varphi(r)), \text{dom}'(\varphi(r)), \emptyset) \end{aligned}$$

5.3 Store relations

The main ingredient of our semantics for effects is the preservation of certain sets of relations on stores. We start by introducing the

notion of relations on stores that depend only on particular subsets of locations in their domain and codomain.

Definition 3 (store relations). If L, L' are sets of locations, a store relation on L, L' is defined as a nonempty relation $R \subseteq \mathbb{S} \times \mathbb{S}$ such that whenever $(s, s') \in R$ and $s \sim_L s_1$ and $s' \sim_{L'} s'_1$ then $(s_1, s'_1) \in R$, too. We write $\text{StRel}(L, L')$ for the set of all store relations on L, L' .

Given such a relation, we now formalize what it means to ‘respect’ an effect ε under some parameter φ .

Definition 4 (relations and effects). Suppose $\Pi \vdash \varphi \text{ ok}$ and $\Pi \vdash \varepsilon \text{ ok}$. Let R be a store relation on $\text{dom}(\varphi), \text{dom}'(\varphi)$. We say that R respects ε at φ if it is preserved by all commands that exhibit only ε on the store layout delineated by φ . Formally, we define:

- R respects $\{rd_r\}$ at φ if $(s, s') \in R$ implies $s.\ell = s'.\ell'$ for all $(\ell, \ell') \in \varphi(r)$;
- R respects $\{wr_r\}$ at φ if for all $(s, s') \in R$ and for all $(\ell, \ell') \in \varphi(r)$ and $v \in \mathbb{Z}$, we have $(s[\ell \mapsto v], s'[\ell' \mapsto v]) \in R$;
- R respects $\{al_r\}$ always.

We then define the set $\mathcal{R}_{\varepsilon}^{\Pi}(\varphi)$ of all store relations that respect ε at φ as follows:

$$\mathcal{R}_{\varepsilon}^{\Pi}(\varphi) = \{R \in \text{StRel}(\text{dom}(\varphi), \text{dom}'(\varphi)) \mid \forall e \in \varepsilon, R \text{ resp. } e \text{ at } \varphi\}.$$

Unfortunately, we cannot track the allocation effect with relations; this will be done separately in the definition of the monad.

Finally, we introduce two additional bits of notation. If $s, s' \in \mathbb{S}$ we define

$$s, s' \models \varphi \iff \begin{aligned} &\text{dom}(s) = \text{dom}(\varphi) \wedge \text{dom}(s') = \text{dom}'(\varphi) \end{aligned}$$

We also define the following:

$$s \sim_{\varphi} s' \iff \forall r \in \text{Regs}. \forall (\ell, \ell') \in \varphi(r). s.\ell = s'.\ell'$$

Note that this notation does not constrain the values in the silent region.

6. Logical Relation

This section defines the relational semantics of refined types. It thus contains the main technical contribution of the paper.

Definition 5 (logical relation). For $\Pi \vdash X \text{ ok}$ and $\Pi \vdash \varphi \text{ ok}$ we define a relation $\llbracket X \rrbracket_{\varphi}^{\Pi} \subseteq \llbracket U(X) \rrbracket \times \llbracket U(X) \rrbracket$ by

$$\begin{aligned} \llbracket X \rrbracket_{\varphi}^{\Pi} &\equiv \Delta_{\llbracket U(X) \rrbracket} \text{ when } X \in \{\text{int}, \text{bool}, \text{unit}\} \\ \llbracket \text{ref}_r \rrbracket_{\varphi}^{\Pi} &\equiv \varphi(r) \\ \llbracket X \times Y \rrbracket_{\varphi}^{\Pi} &\equiv \llbracket X \rrbracket_{\varphi}^{\Pi} \times \llbracket Y \rrbracket_{\varphi}^{\Pi} \\ \llbracket X \rightarrow T_{\varepsilon} Y \rrbracket_{\varphi}^{\Pi} &\equiv \{(f, f') \mid \forall \varphi' \geq_{\Pi} \varphi. \forall (x, x') \in \llbracket X \rrbracket_{\varphi'}^{\Pi}. \\ &\quad (f(x), f'(x')) \in \llbracket T_{\varepsilon} Y \rrbracket_{\varphi'}^{\Pi}\} \\ \llbracket T_{\varepsilon} X \rrbracket_{\varphi}^{\Pi} &\equiv \{(f, f') \mid s, s' \models \varphi \Rightarrow \\ &\quad \forall R \in \mathcal{R}_{\varepsilon}^{\Pi}(\varphi). s R s' \Rightarrow s_1 R s'_1 \wedge \\ &\quad \exists \psi \in \text{Par}(\text{als}(\varepsilon)). s_1, s'_1 \models \varphi \otimes \psi \wedge \\ &\quad s_1 \sim_{\psi} s'_1 \wedge (v, v') \in \llbracket X \rrbracket_{\varphi \otimes \psi}^{\Pi} \\ &\quad \text{where } (s_1, v) = f \text{ s and } (s'_1, v') = f' \text{ s'} \} \end{aligned}$$

We define $\llbracket \Theta \rrbracket_{\varphi}^{\Pi}$ by $\llbracket \Theta \rrbracket_{\varphi}^{\Pi} \equiv \{(\gamma, \gamma') \mid \forall i. (\gamma(x_i), \gamma'(x_i)) \in \llbracket X_i \rrbracket_{\varphi}^{\Pi}\}$ where $\Theta = x_1 : X_1, \dots, x_n : X_n$.

The definition of the logical relation on computation types deserves some explanation. First, it says that the store behaviour of

two related computations must respect all relations that are compatible with the declared effect, cf. [8]. Since these relations are completely unconstrained on the silent region τ , this implies in particular that the silent region may neither be read nor modified. The existential quantifier asserts a (disjoint) extension ψ of the current parameter φ which is to hold all newly allocated references, hence $\psi \in \text{Par}(\text{als}(\varepsilon))$, cf. Def. 2. The result values (v, v') are then required to be related with respect to the extended parameter $\varphi \otimes \psi$. Note that if v and v' contain newly allocated references then $(v, v') \in \llbracket X \rrbracket_{\varphi}^{\Pi}$ will in general not hold.

The semantics of value types is monotonic with respect to the ordering on parameters.

Lemma 2 (Monotonicity). *Suppose $\Pi \vdash X$ ok. If $\varphi' \geq_{\Pi} \varphi$ then $\llbracket X \rrbracket_{\varphi'}^{\Pi} \supseteq \llbracket X \rrbracket_{\varphi}^{\Pi}$.*

Proof. By induction on X . \square

Lemma 3 (Masking). *Suppose $\Pi \vdash X$ ok and $\Pi, r \vdash \varphi$ ok. Then $\llbracket X \rrbracket_{\varphi}^{\Pi, r} = \llbracket X \rrbracket_{\varphi-r}^{\Pi}$, and likewise $\llbracket T_{\varepsilon} X \rrbracket_{\varphi}^{\Pi, r} = \llbracket T_{\varepsilon} X \rrbracket_{\varphi-r}^{\Pi}$.*

Proof. By induction on X . The cases where X is a basic type or a product type are trivial and hence omitted. Note, though, that if $X = \text{ref}_r$, then $r' \neq r$.

Case $X = X_1 \rightarrow T_{\varepsilon} X_2$. Suppose that $(f, f') \in \llbracket X \rrbracket_{\varphi}^{\Pi, r}$ and that $\varphi' \geq_{\Pi} \varphi - r$ and that $(x, x') \in \llbracket X_1 \rrbracket_{\varphi'}^{\Pi}$. Write $\varphi' = (\varphi - r) \otimes \theta$ and put $\psi = \varphi \otimes \theta$. We have $\psi \geq_{\Pi, r} \varphi$ and $(x, x') \in \llbracket X_1 \rrbracket_{\psi}^{\Pi, r}$ by the induction hypothesis since $\psi - r = \varphi'$. Note that $\text{als}(\varepsilon) \subseteq \Pi$ and so $\Pi \vdash \theta$ ok. The assumption then gives $(f v, f' v') \in \llbracket T_{\varepsilon} X_2 \rrbracket_{\psi}^{\Pi, r}$ and thence $(f v, f' v') \in \llbracket T_{\varepsilon} X_2 \rrbracket_{\varphi'}^{\Pi}$ as required, again by the induction hypothesis and $\psi - r = \varphi'$.

Conversely, assume that $(f, f') \in \llbracket X \rrbracket_{\varphi-r}^{\Pi}$ and that $\varphi' \geq_{\Pi, r} \varphi$ and $(x, x') \in \llbracket X_1 \rrbracket_{\varphi'}^{\Pi, r}$. We have $\varphi' - r \geq_{\Pi} \varphi - r$ and $(x, x') \in \llbracket X_1 \rrbracket_{\varphi-r}^{\Pi}$ by the induction hypothesis. Therefore, the assumption gives us $(f x, f' x') \in \llbracket T_{\varepsilon} X_2 \rrbracket_{\varphi-r}^{\Pi}$ whence $(f x, f' x') \in \llbracket T_{\varepsilon} X_2 \rrbracket_{\varphi'}^{\Pi, r}$ again by the induction hypothesis and we are done.

Case $T_{\varepsilon}(X)$. Suppose that $(f, f') \in \llbracket T_{\varepsilon}(X) \rrbracket_{\varphi}^{\Pi, r}$ and $R \in \mathcal{R}_{\varepsilon}^{\Pi}(\varphi - r)$ (note that $\Pi \vdash \varepsilon$ ok) and $s, s' \models \varphi - r$ and $(s, s') \in R$. Write $s_1, v = f s$ and $s'_1, v' = f' s'$. We also have $s, s' \models \varphi$ by definition of $\varphi - r$. Now, $R \in \mathcal{R}_{\varepsilon}^{\Pi, r}(\varphi)$ since the masked out region r is not mentioned in ε ; therefore the assumption provides us with $\psi \in \text{Par}(\text{als}(\varepsilon))$ disjoint from s, s' such that $s_1, s'_1 \models \varphi'$ and $(s_1, s'_1) \in R$ and $s_1 \sim_{\psi} s'_1$ and $(v, v') \in \llbracket X \rrbracket_{\varphi'}^{\Pi, r}$ where $\varphi' = \varphi \otimes \psi$. Now put $\varphi'' = (\varphi - r) \otimes \psi = \varphi' - r$. We have $s_1, s'_1 \models \varphi''$ and $(v, v') \in \llbracket X \rrbracket_{\varphi''}^{\Pi}$ by the induction hypothesis and we are done.

Conversely, suppose that $(f, f') \in \llbracket T_{\varepsilon}(X) \rrbracket_{\varphi-r}^{\Pi}$ and $R \in \mathcal{R}_{\varepsilon}^{\Pi, r}(\varphi)$ and $s, s' \models \varphi$ and $(s, s') \in R$. Again, write $s_1, v = f s$ and $s'_1, v' = f' s'$. We also have $s, s' \models \varphi - r$ by definition of $\varphi - r$.

Now, as before, $R \in \mathcal{R}_{\varepsilon}^{\Pi}(\varphi - r)$ and so the assumption provides us with $\psi \in \text{Par}(\text{als}(\varepsilon))$ such that $s_1, s'_1 \models \varphi'$ and $(s_1, s'_1) \in R$ and $(v, v') \in \llbracket X \rrbracket_{\varphi'}^{\Pi, r}$ where $\varphi' = (\varphi - r) \otimes \psi$.

Putting $\varphi'' = \varphi \otimes \psi$ it follows that $(v, v') \in \llbracket X \rrbracket_{\varphi''}^{\Pi, r}$ from the induction hypothesis (note that $\varphi'' - r = \varphi'$). \square

Lemma 4 (Extension).

$$\mathcal{R}_{\varepsilon-r}^{\Pi}(\varphi) = \mathcal{R}_{\varepsilon}^{\Pi, r}(\varphi \otimes [r \mapsto \emptyset])$$

Proof. The \supseteq direction is obvious. Conversely, if R respects $\varepsilon - r$ at φ then it also respects ε at $\varphi \otimes [r \mapsto \emptyset]$ since the additional

restrictions imposed by possible occurrences of r in ε are vacuously true at $\varphi \otimes [r \mapsto \emptyset]$. \square

The following establishes semantic soundness for our subtyping relation.

Lemma 5 (Soundness of subtyping). *If $\Pi \vdash X_i$ ok and $X_1 \leq X_2$ then for all $\varphi \in \text{Par}(\Pi)$ one has $\llbracket X_1 \rrbracket_{\varphi}^{\Pi} \subseteq \llbracket X_2 \rrbracket_{\varphi}^{\Pi}$.*

Proof. By induction on the subtyping derivation. \square

We can now prove the following ‘fundamental theorem’ of logical relations, which states that terms are related to themselves.

Theorem 1 (Fundamental Theorem).

1. *If $\Pi \vdash \varphi$ ok, $\Pi; \Theta \vdash V : X$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$, then*

$$(\llbracket U(\Theta) \vdash V : U(X) \rrbracket_{\gamma}, \llbracket U(\Theta) \vdash V : U(X) \rrbracket_{\gamma'}) \in \llbracket X \rrbracket_{\varphi}^{\Pi}.$$

2. *If $\Pi \vdash \varphi$ ok, $\Pi; \Theta \vdash M : T_{\varepsilon}(X)$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$, then*

$$(\llbracket U(\Theta) \vdash M : T(U(X)) \rrbracket_{\gamma}, \llbracket U(\Theta) \vdash M : T(U(X)) \rrbracket_{\gamma'}) \in \llbracket T_{\varepsilon}(X) \rrbracket_{\varphi}^{\Pi}.$$

Proof. By induction on typing derivations. We give a selection of representative cases.

Case let. From the derivation we have

$$\Pi; \Theta \vdash M_1 : T_{\varepsilon_1}(X_1) \quad (1)$$

$$\Pi; \Theta, x : X_1 \vdash M_2 : T_{\varepsilon_2}(X_2) \quad (2)$$

Assume $\Pi \vdash \varphi$ ok and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$. Suppose $R \in \mathcal{R}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi}(\varphi)$ and $s_0 R s'_0$. Let

$$f_1 := \llbracket U(\Theta) \vdash M_1 : T U X_1 \rrbracket,$$

and $(s_1, v_1) = f_1 \gamma s_0$ and $(s'_1, v'_1) = f_1 \gamma' s'_0$.

$$f_2 := \llbracket U(\Theta), x : U(X_1) \vdash M_2 : T U X_2 \rrbracket,$$

and $(s_2, v_2) = f_2(\gamma, v_1) s_1$ and $(s'_2, v'_2) = f_2(\gamma', v'_1) s'_1$.

We then have $\llbracket U(\Theta) \vdash M : T U X_2 \rrbracket_{\gamma s_0} = (s_2, v_2)$ and $\llbracket U(\Theta) \vdash M : T U X_2 \rrbracket_{\gamma' s'_0} = (s'_2, v'_2)$.

Now $R \in \mathcal{R}_{\varepsilon_i}^{\Pi}(\varphi)$ holds for $i = 1, 2$. The induction hypothesis for (1), applied to $\varphi, (\gamma, \gamma')$ and (s_0, s'_0) therefore yields $s_1 R s'_1$ and also furnishes $\psi_1 \in \text{Par}(\text{als}(\varepsilon_1))$ such that $s_1, s'_1 \models \varphi \otimes \psi_1$ and $s_1 \sim_{\psi_1} s'_1$ and $(v_1, v'_1) \in \llbracket X_1 \rrbracket_{\varphi'}^{\Pi}$, where $\varphi' = \varphi \otimes \psi_1$. We now define a relation R^* by

$$s R^* s' \iff s R s' \wedge s \sim_{\psi_1} s'$$

We then have $R^* \in \mathcal{R}_{\varepsilon_2}^{\Pi}(\varphi')$ and also $s_1 R^* s'_1$. In particular, R^* respects rd_r at ψ_1 in view of the clause $s \sim_{\psi_1} s'$ that has been added for that purpose.

The monotonicity of the interpretation of the value types (Lemma 2) yields $((\gamma, v_1), (\gamma', v'_1)) \in \llbracket \Theta, x : X_1 \rrbracket_{\varphi'}^{\Pi}$. We can therefore apply the induction hypothesis for (2), with $\varphi', ((\gamma, v_1), (\gamma', v'_1))$ and (s_1, s'_1) , and obtain the existence of some $\psi_2 \in \text{Par}(\text{als}(\varepsilon_2))$ such that $s_2, s'_2 \models \varphi' \otimes \psi_2$ and $s_2 R^* s'_2$ and $s_2 \sim_{\psi_2} s'_2$ and $(v_2, v'_2) \in \llbracket X_2 \rrbracket_{\varphi \otimes \psi_1 \otimes \psi_2}^{\Pi}$. From the definition of R^* we then get $s_2 R s'_2$ and $s_2 \sim_{\psi_1 \otimes \psi_2} s'_2$.

Case ref. Suppose $\Pi \vdash \varphi$ ok, $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$, $R \in \mathcal{R}_{al_r}^{\Pi}(\varphi)$, i.e., arbitrary and $s_0 R s'_0$. Writing

$$f := \llbracket U \Theta \vdash \text{ref}(V) : T \text{ref} \rrbracket$$

we have

$$f \gamma s = \text{new}(s, \gamma(x)) = (s_1, \ell) \quad \text{and} \\ f \gamma' s' = \text{new}(s', \gamma'(x)) = (s'_1, \ell').$$

Now, since R is a store relation on $\text{dom}(\varphi), \text{dom}'(\varphi)$ and $s_0, s'_0 \models \varphi$ we find $s_1 R s'_1$. We let $\psi = [r \mapsto (\ell, \ell')]$ and have $s_1, s'_1 \models \varphi \otimes \psi$. We have $(\ell, \ell') \in \llbracket \text{ref}_r \rrbracket_{\varphi \otimes \psi}^{\Pi}$ and since $\llbracket U\Theta \vdash V : \text{int} \rrbracket \gamma = \llbracket U\Theta \vdash V : \text{int} \rrbracket \gamma'$ from $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$, we also have $s_1 \sim_{\psi} s'_1$ and we are done.

Masking Rule Suppose $\Pi \vdash \Theta \text{ ok}$, $\Pi \vdash X \text{ ok}$, $\Pi, r; \Theta \vdash M : T_{\varepsilon} X$, $\Pi \vdash \varphi \text{ ok}$, and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$ and $\Pi \vdash \varphi \text{ ok}$ and $R \in \mathcal{R}_{\varepsilon-r}^{\Pi}(\varphi)$ and $s_0 R s'_0$ and $s_0, s'_0 \models \varphi$.

Write $g := \llbracket U\Theta \vdash M : TUX \rrbracket$ and $(s_1, v_1) = g\gamma s_0$ and $(s'_1, v'_1) = g\gamma s'_0$.

Let us write $\psi = \varphi \otimes [r \mapsto \emptyset]$. The masking lemma 3 shows $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\psi}^{\Pi, r}$ and from the extension lemma 4 we have $R \in \mathcal{R}_{\varepsilon-r}^{\Pi, r}(\psi)$. So the induction hypothesis furnishes $\theta \in \text{Par}(\text{als}(\varepsilon))$ and disjoint from s_0, s'_0 such that $s_1, s'_1 \models \psi'$ and $s_1 \sim_{\theta} s'_1$ and $s_1 R s'_1$ and $(v_1, v'_1) \in \llbracket X \rrbracket_{\psi'}^{\Pi, r}$ where $\psi' = \psi \otimes \theta$.

Now, $\psi' - r = \varphi \otimes (\theta - r)$ and so $(v_1, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes (\theta - r)}^{\Pi}$ again by the masking lemma. The rest is clear from the definitions. \square

Definition 6 (contextual equivalence).

1. Suppose that $\Pi; \emptyset \vdash V_i : X$ for $i = 1, 2$ are two closed values. They are contextually equivalent, written $\Pi; \emptyset \vdash V_1 \equiv V_2 : X$, if for all values $\Pi; \emptyset \vdash V : X \rightarrow T_{\varepsilon} \text{bool}$ (“contexts”) with ε arbitrary it holds that when $(s_1, v_1) = \llbracket V V_1 \rrbracket () \emptyset$ and $(s_2, v_2) = \llbracket V V_2 \rrbracket () \emptyset$ then $v_1 = v_2$. Recall that \emptyset is the initial, empty state.
2. Contextual equivalence is extended to open values and computations by abstracting all free variables using a dummy abstraction of type **unit** in the case of a closed computation. We write $\Pi; \Theta \vdash M_1 \equiv M_2 : A$ to mean that computations M_1 and M_2 are contextually equivalent.

Notice that the examining context V must itself be typable in our effect system. This is important when, e.g., $X = (\text{int} \rightarrow T_{\emptyset} \text{unit}) \rightarrow T_{\varepsilon} \text{unit}$ and the equivalence of V_1 and V_2 relies on their being fed a pure function as input. E.g., in this situation we would want to consider $V_1 = \lambda f. f(0)$ and $V_2 = \lambda f. f(0); f(0)$ as equivalent.

This definition coincides with the more standard yet more complex one involving terms with holes as contexts, cf. [8]. This is due to the fact that we use a monadic metalanguage rather than a language with built-in side effects like ML.

We remark that the restriction to boolean observations is not a severe one. Two contextually equivalent values also agree on all observations of type **int** as is easily seen by wrapping the observation into an appropriate equality test.

Proposition 1. *Contextual equivalence is a congruence validating the equational theory of the monadic metalanguage (in particular beta reduction).*

The following corollary to Theorem 1 now provides a powerful method for establishing stronger contextual equivalences.

Corollary 1. *Suppose that $\Pi; \Theta \vdash M_i : T_{\varepsilon} X$ for $i = 1, 2$ and whenever $\Pi \vdash \varphi \text{ ok}$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$ then*

$$\begin{aligned} & \llbracket [U(\Theta) \vdash M_1 : T(U(X))] \rrbracket \gamma, \llbracket [U(\Theta) \vdash M_2 : T(U(X))] \rrbracket \gamma' \\ & \in \llbracket [T_{\varepsilon} X] \rrbracket_{\varphi}^{\Pi}. \end{aligned}$$

Then M_1 and M_2 are contextually equivalent.

Proof. Let $\lambda(M_1)$ and $\lambda(M_2)$ be the values obtained by lambda abstracting all variables in M_1, M_2 . Let us write $\Theta \rightarrow T_{\varepsilon} A$ for the common type of these values. It is easy to see from the definition

of the logical relation that

$$(\llbracket \lambda(M_1) \rrbracket, \llbracket \lambda(M_2) \rrbracket) \in \llbracket \Theta \rightarrow T_{\varepsilon} X \rrbracket_{\varphi}^{\Pi}$$

for all φ (we have omitted types and contexts here). Now let $V : (\Theta \rightarrow T_{\varepsilon} X) \rightarrow T_{\varepsilon'} \text{bool}$ be a context. By applying the fundamental lemma to V we conclude

$$(\llbracket V \lambda(M_1) \rrbracket, \llbracket V \lambda(M_2) \rrbracket) \in \llbracket [T_{\varepsilon'} \text{bool}] \rrbracket_{\varphi}^{\Pi}$$

Now, let φ be the parameter that assigns to each region the empty partial bijection. We have $\emptyset, \emptyset \models \varphi$ where \emptyset is the empty state. Let $R \in \mathcal{R}_{\varepsilon'}^{\Pi}(\varphi)$ be arbitrary. Clearly, $\emptyset R \emptyset$ no matter what ε' is. Put $(s_i, v_i) = \llbracket V(\lambda(M_i)) \rrbracket$. By the definition of the logical relation we then have $v_1 = v_2$ as required. \square

7. Applications

For $\Pi \vdash \varphi \text{ ok}$ and $\Pi \vdash \varepsilon \text{ ok}$ and $s, s' \models \varphi$ we introduce the notation

$$s \sim_{\text{rds}_{\varphi}(\varepsilon)} s' \iff \forall r \in \text{rds}(\varepsilon). \forall (\ell, \ell') \in \varphi(r). s.\ell = s'.\ell'$$

It expresses that s and s' agree on those locations that are read given effect ε .

We also define

$$\begin{aligned} \text{nwr}_{\varphi}(\varepsilon) &= \text{dom}(\varphi) \setminus \bigcup_{r \in \text{wrs}(\varepsilon)} \text{dom}(\varphi(r)) \\ \text{nwr}'_{\varphi}(\varepsilon) &= \text{dom}'(\varphi) \setminus \bigcup_{r \in \text{wrs}(\varepsilon)} \text{dom}'(\varphi(r)) \end{aligned}$$

Thus $\text{nwr}_{\varphi}(\varepsilon)$ and $\text{nwr}'_{\varphi}(\varepsilon)$ comprise the locations on the left (resp. right) side that are not written to given effect ε . This includes the locations in the silent region.

Lemma 6. *Suppose $\Pi; \Theta \vdash M : T_{\varepsilon} X$ and $\Pi \vdash \varphi \text{ ok}$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_{\varphi}^{\Pi}$ and $s_0, s'_0 \models \varphi$ and $\llbracket M \rrbracket \gamma s_0 = (s_1, x)$ and $\llbracket M \rrbracket \gamma' s'_0 = (s'_1, x')$.*

If $s_0 \sim_{\text{rds}_{\varphi}(\varepsilon)} s'_0$ then there exists $\psi \in \text{Par}(\text{als}(\varepsilon))$ disjoint from s_0, s'_0 such that

1. $s_1, s'_1 \models \varphi \otimes \psi$ and $(x, x') \in \llbracket X \rrbracket_{\varphi \otimes \psi}$ and $s_1 \sim_{\psi} s'_1$.
2. $s_0 \sim_{\text{nwr}_{\varphi}(\varepsilon)} s_1$ and $s'_0 \sim_{\text{nwr}'_{\varphi}(\varepsilon)} s'_1$.
3. For each $(\ell, \ell') \in \varphi(r)$ where $r \in \Pi$ we have either
 - $s_0.\ell = s_1.\ell$ and $s'_0.\ell' = s'_1.\ell'$ (unchanged) or
 - $s_1.\ell = s'_1.\ell'$ (identically written).
4. Suppose that $\ell \in \text{dom}(\varphi)$ but there is no ℓ', r such that $(\ell, \ell') \in \varphi(r)$. Then $s_0.\ell = s_1.\ell$. A symmetric statement holds for s'_0, s'_1 .

Notice that Part 2 asserts in particular that the contents of the silent region do not change from s_0 to s_1 .

Proof. Part 1 is direct from the definition.

For part 2 we define

$$(s, s') \in R \iff s \sim_{\text{rds}_{\varphi}(\varepsilon)} s' \wedge s \sim_{\text{nwr}_{\varphi}(\varepsilon)} s_0$$

We now have $s_0 R s'_0$ (not necessarily $s_0 R s'_0$) and so $s_1 R s'_1$. The claim about s_1 follows. The proof for $\text{nwr}'_{\varphi}(\varepsilon)$ and s'_1 is analogous.

For part 3 we define

$$\begin{aligned} (s, s') \in R & \iff s \sim_{\text{rds}_{\varphi}(\varepsilon)} s' \wedge \\ & \forall r \in \Pi. \forall (\ell, \ell') \in \varphi(r). s.\ell = s_0.\ell \wedge s'.\ell' = s'_0.\ell' \vee s.\ell = s'.\ell' \end{aligned}$$

and note that $R \in \mathcal{R}_{\varepsilon}^{\Pi}(\varphi)$. It follows $s_1, s'_1 \in R$ and the claim follows.

Part 4, finally, follows using the relation

$$s R s' \iff s \sim_{\text{rds}_{\varphi}(\varepsilon)} s' \wedge s.\ell = s_0.\ell$$

\square

Proposition 2 (duplicated computation). *Suppose that $\Pi; \Theta \vdash M : T_{\varepsilon}(X)$ and suppose that $\text{rds}(\varepsilon) \cap \text{wrs}(\varepsilon) = \text{als}(\varepsilon) = \emptyset$. Thus,*

M reads and writes on disjoint portions of the store and makes no allocations except possibly in the silent region. Then the following two terms are contextually equivalent:

$$\begin{aligned} M_1 &:= \text{let } x \leftarrow M \text{ in val } (x, x) \\ M_2 &:= \text{let } x \leftarrow M \text{ in let } y \leftarrow M \text{ in val } (x, y) \end{aligned}$$

Formally, $\Pi; \Theta \vdash M_1 \equiv M_2 : T_\varepsilon (X \times X)$.

Proof. We will use Cor. 1. Let $f(\gamma) = \llbracket U\Theta \vdash M : TUX \rrbracket \gamma$. Suppose that $\Pi \vdash \varphi \text{ ok}$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_\varphi^\Pi$. Suppose that $s_0, s'_0 \models \varphi$ and define

$$\begin{aligned} s_1, v_1 &= f(\gamma)(s_0) \\ s_2, v_2 &= f(\gamma)(s_1) \\ s'_1, v'_1 &= f(\gamma')(s'_0) \end{aligned}$$

Suppose that $s_0 R s'_0$ for some $R \in \mathcal{R}_\varepsilon^\Pi(\varphi)$. We must exhibit some $\psi \in \text{Par}(\emptyset)$ disjoint from s, s' such that $s_2, s'_1 \models \varphi \otimes \psi$ and $(v_2, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes \psi}^\Pi$ and $(v_1, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes \psi}^\Pi$ and $s_2 \sim_\psi s'_1$ and $s_2 R s'_1$. Notice that ψ contains allocations in the silent region only.

From $s_0 R s'_0$ it follows that $s_0 \sim_{\text{rds}_\varphi(\varepsilon)} s'_0$. Thus, Lemma 6 furnishes $\psi_1 \in \text{Par}(\emptyset)$ disjoint from s_0, s'_0 such that $s_1, s'_1 \models \varphi \otimes \psi_1$ and $(v_1, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes \psi_1}^\Pi$. Furthermore, since $\text{rds}(\varphi) \cap \text{wrs}(\varphi) = \emptyset$ we have $s_1 \sim_{\text{rds}_\varphi(\varepsilon)} s'_0$. Decompose ψ_1 as $\psi_1^L \otimes \psi_1^R$ where $\text{dom}'(\psi_1^L) = \emptyset = \text{dom}(\psi_1^R)$.

Now we have $s_1, s'_0 \models \varphi \otimes \psi_1^L$ and thus another application of Lemma 6 furnishes $\psi_2 \in \text{Par}(\emptyset)$ disjoint from s_1, s'_0 such that $s_2, s'_1 \models \varphi \otimes \psi_1^L \otimes \psi_2$ and $(v_2, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes \psi_1^L \otimes \psi_2}^\Pi$.

Again decompose ψ_2 as $\psi_2 = \psi_2^L \otimes \psi_2^R$ where $\text{dom}'(\psi_2^L) = \emptyset = \text{dom}(\psi_2^R)$.

We already know from Part 2 of Lemma 6 that if location l is not written (that includes locations in the silent region) then $s_1.l = s_2.l$. We claim that s_1 agrees with s_2 on the other, written to, locations as well. To see this, we define s_0^* by updating s_0 in such a way that for all $r \in \Pi$ and $(\ell, \ell') \in \varphi(r)$ one has $s_0^*.l' = s_0.l$ and s_0^* agrees with s'_0 elsewhere. Then $s_1, s_0^* \models \varphi \otimes \psi_1^L$ and $s_1 \sim_{\text{rds}_\varphi(\varepsilon)} s_0^*$. Let s_1^* denote the successor state of s_0^* . Lemma 6,3 yields that for each $(\ell, \ell') \in \varphi(r)$ that either $s_2.l = s_1.l$ or $s_2.l = s_1^*.l'$. On the other hand, $s_1^*.l' = s_1.l$ or $s_1.l = s_0.l$ and $s_1^*.l' = s_0^*.l'$. It follows $s_2.l = s_1.l$. We have thus shown $s_1 \sim_{\text{dom}(\varphi)} s_2$. Since $s_1 R s'_1$ and R is a store relation on $\text{dom}(\varphi)$, $\text{dom}'(\varphi)$ we conclude $s_2 R s'_1$ and we are done. \square

Proposition 3 (commuting computations). *Suppose that*

$$\begin{aligned} \Pi; \Theta \vdash M_1 &: T_{\varepsilon_1}(X) \\ \Pi; \Theta \vdash M_2 &: T_{\varepsilon_2}(X) \end{aligned}$$

and suppose that $\text{rds}(\varepsilon_1) \cap \text{wrs}(\varepsilon_2) = \text{rds}(\varepsilon_2) \cap \text{wrs}(\varepsilon_1) = \text{wrs}(\varepsilon_1) \cap \text{wrs}(\varepsilon_2) = \emptyset$.

Put

$$\begin{aligned} N_1 &:= \text{let } y \leftarrow M_1 \text{ in let } x \leftarrow M_2 \text{ in val } (x, y) \\ N_2 &:= \text{let } x \leftarrow M_2 \text{ in let } y \leftarrow M_1 \text{ in val } (x, y) \end{aligned}$$

We have

$$\Pi; \Theta \vdash N_1 \equiv N_2 : T_\varepsilon (X \times X)$$

where $\varepsilon = \varepsilon_1 \cup \varepsilon_2$.

Proof. We will use Cor. 1. Let $f_i(\gamma) = \llbracket U\Theta \vdash M_i : TUX \rrbracket \gamma$ for $i = 1, 2$.

Suppose that $\Pi \vdash \varphi \text{ ok}$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_\varphi^\Pi$. Suppose that $s_0, s'_0 \models \varphi$ and define

$$\begin{aligned} s_1, v_1 &= f_1(\gamma)(s_0) \\ s_2, v_2 &= f_2(\gamma)(s_1) \\ s'_1, v'_1 &= f_2(\gamma')(s'_0) \\ s_2, v_2 &= f_1(\gamma')(s'_1) \end{aligned}$$

Suppose that $s_0 R s'_0$ for some $R \in \mathcal{R}_\varepsilon^\Pi(\varphi)$.

We must exhibit $\psi \in \text{Par}(\text{als}(\varepsilon))$ disjoint from s_0, s'_0 such that $s_2, s'_1 \models \varphi'$ and $(v_1, v'_1) \in \llbracket X \rrbracket_{\varphi'}$ and $(v_2, v'_1) \in \llbracket X \rrbracket_{\varphi'}$ where $\varphi' = \varphi \otimes \psi$ and $s_2 R s'_1$.

Lemma 6 applied to M_2 together with the fact that $\text{rds}(\varepsilon_1) \cap \text{wrs}(\varepsilon_2) = \emptyset$ shows that $s_0 \sim_{\text{rds}_\varphi(\varepsilon_1)} s'_1$. We also have $s_0, s'_1 \models \varphi$ so Lemma 6 applied to M_1 furnishes ψ_1 disjoint from s_0, s'_1 such that $(v_1, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes \psi_1}^\Pi$ and also $s_1, s'_2 \models \varphi \otimes \psi_1$. Arguing symmetrically, we obtain ψ_2 disjoint from s_1, s'_0 such that $(v_2, v'_1) \in \llbracket X \rrbracket_{\varphi \otimes \psi_2}^\Pi$ and also $s_2, s'_1 \models \varphi \otimes \psi_2$. Now, since $\text{dom}(\psi_2)$ is disjoint from $\text{dom}(s_1) \supseteq \text{dom}(\psi_1)$ and analogously on the right-hand side, we find that ψ_1 and ψ_2 are disjoint allowing us to put $\psi := \psi_1 \otimes \psi_2$ and $(v_2, v'_1), (v_1, v'_2) \in \llbracket X \rrbracket_{\varphi \otimes \psi}^\Pi$ by monotonicity. We also have $s_2, s'_1 \models \varphi \otimes \psi$.

Since $s_0 R s'_0$ we can show $s_2 R s'_1$ by induction on the size of $\text{dom}_{le/rw}(\varphi)$ using the definition of $\mathcal{R}_\varepsilon^\Pi(\varphi)$. \square

Proposition 4 (dead computation). *Suppose that*

$$\Pi; \Theta \vdash M : T_\varepsilon(\text{unit})$$

and that $\text{wrs}(\varepsilon) = \emptyset$. Then M is contextually equivalent to $\text{val } ()$.

Proof. Let $\varphi \in \text{Par}(\Pi)$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_\varphi$. Let $f = \llbracket M \rrbracket \gamma$ and suppose $s_0, s'_0 \models \varphi$. Let $v, s_1 = f s_0$. Of course, $v = \star$. The fundamental lemma applied to M, φ, s_0, s'_0 furnishes $\psi \in \text{Par}(\text{als}(\varepsilon))$ disjoint from φ . Construct ψ_1 by taking only the left part of ψ ; formally: $\psi_1(r) = (\text{dom}(\psi(r)), \emptyset, \emptyset)$. Now, $s_1, s'_0 \models \varphi \otimes \psi_1$ and, trivially, $(v, \star) \in \llbracket \text{unit} \rrbracket_{\varphi \otimes \psi_1}$.

If $R \in \mathcal{R}_\varepsilon^\Pi(\varphi)$ and $s_0 R s'_0$ then s_1 agrees with s_0 on all locations that R depends upon, so $s_1 R s'_0$. \square

Proposition 5 (pure lambda hoist). *Suppose that*

$$\begin{aligned} \Pi; \Theta \vdash M &: T_\emptyset Z \\ \Pi; \Theta, x:X, y:Z \vdash N &: T_\varepsilon Y \end{aligned}$$

Put

$$\begin{aligned} M_1 &:= \text{val } (\lambda x:UX. \text{let } y \leftarrow M \text{ in } N) \\ M_2 &:= \text{let } y \leftarrow M \text{ in val } (\lambda x:UX. N) \end{aligned}$$

Then $\Pi; \Theta \vdash M_1 \equiv M_2 : T_\emptyset (X \rightarrow T_\varepsilon)$

Proof. Let $\varphi \in \text{Par}(\Pi)$ and $(\gamma, \gamma') \in \llbracket \Theta \rrbracket_\varphi$ and define

$$\begin{aligned} f &= \llbracket M \rrbracket \gamma \\ f' &= \llbracket M \rrbracket \gamma' \\ g(v_x, v_y) &= \llbracket N \rrbracket \gamma[x \mapsto v_x, y \mapsto v_y] \\ g'(v_x, v_y) &= \llbracket N \rrbracket \gamma'[x \mapsto v_x, y \mapsto v_y] \end{aligned}$$

Suppose that $s_0, s'_0 \models \varphi$. The fundamental theorem applied to M yields $\psi_1 \in \text{Par}(\emptyset)$ such that $s_1, s'_1 \models \varphi \otimes \psi_1$ where $s_1, v = f s_0$ and $s'_1, v' = f'_0$.

Factor ψ_1 as $\psi_1^L \otimes \psi_1^R$ such that $s_1, s'_0 \models \varphi \otimes \psi_1^L$. We claim that

$$(\lambda v_x. g(v_x, v), \lambda v_x. \lambda s'. \text{let } s'_2, v_y = f' s' \text{ in } g'(v_x, v_y) s'_2) \in \llbracket X \rightarrow T_\varepsilon \rrbracket_{\varphi \otimes \psi_1^L}$$

Assume thus $\varphi' \geq_\Pi \varphi \otimes \psi_1^L$ and $(v_x, v'_x) \in \llbracket X \rrbracket_{\varphi'}$. Also assume $s, s' \models \varphi'$ and let $s'_2, v_y = f' s'$. We would like to argue that $(v, v_y) \in \llbracket Z \rrbracket_{\varphi' \otimes \psi_2}$ for some ψ_2 .

Now $s_0, s' \models \varphi$ whence $(v, v_y) \in \llbracket Z \rrbracket_{\varphi \otimes \psi_1^L \otimes \theta}$ where θ is on the right hand side only. We also know from $s, s' \models \varphi'$ that θ is disjoint from φ' . Thus, we have $(v, v_x) \in \llbracket Z \rrbracket_{\varphi' \otimes \theta}$ by monotonicity.

The rest is similar to earlier calculations. \square

Examples Suppose that the let bindings $f_1 \Leftarrow M_{\text{buf}}$ and $f_2 \Leftarrow M_{\text{buf}}$ are in force. If we type them using two different region identifiers Prop. 3 and Prop. 4 and Prop. 1 allow us to conclude that

$$\begin{aligned} & \text{let } x \Leftarrow f_1(5) \text{ in let } y \Leftarrow f_1(4) \text{ in let } z \Leftarrow f_2(6) \text{ in } (x, z) \equiv \\ & \text{let } z \Leftarrow f_2(6) \text{ in let } x \Leftarrow f_1(5) \text{ in } (x, z) \end{aligned}$$

Lambda hoist (Prop. 5) applies to the following code after using the masking rule to give a pure typing to V_{sum} .

$$\lambda x. \text{let } s \Leftarrow V_{\text{sum}}(a, b, c) \text{ in val } (x + s)$$

We remark that the program equivalences we get for pure computations are complete in the following sense:

Proposition 6. *Let C be a cartesian closed category and T be a strong monad on C . Suppose that in the Kleisli category C_T the laws of dead computation, commuting computations, duplicated computations, lambda hoist are valid. Then C_T is cartesian closed.*

In particular, the Kleisli category consisting of computations of type $T_\emptyset(A)$ modulo contextual equivalence is cartesian closed.

8. Conclusion and further work

We have given a relational semantics to a region-based effect type system for a higher-order language with dynamically allocated store. The relational semantics is shown sound for contextual equivalence and thus provides a powerful proof principle for the latter. We have used the semantics to establish the soundness of a collection of useful effect-based program transformations. It would probably be very hard to establish these directly from the definition of contextual equivalence and no such proof appears to exist in the literature.

There has been a great deal of previous work on the soundness of region-based memory management and of its close cousin, encapsulated monadic state, as provided by `runST` in Haskell [15]. We mention some particularly relevant references. Banerjee et al. [3] translate the region calculus into a variant of System F and give a denotational model showing that references in masked regions do not affect the rest of the computation. Moggi and Sabry [20] prove syntactic type soundness for encapsulated lazy state. Fluet and Morrisett [12] bring the two lines of work together by giving a type- and behaviour-preserving translation from a variant of the region calculus into an extension of System F with a region-indexed family of monads. Naumann [21] uses simulation relations to capture a notion of observational purity for boolean-valued specifications that allows mutation of encapsulated state.

The general problem of modelling and proving equivalences in languages with dynamically allocated store and higher order features is a difficult one, with a very long history [27]. The basic techniques we use here, such as partial bijections and parametric logical relations, have been developed and refined over the last 25 years or so [14, 18, 22, 23, 24, 9]. The focus of much of this previous work has been on showing tricky equivalences between particular pairs of terms, such as the well-known Meyer-Sieber examples [18]. One might expect that equivalences justified by simple program analyses, such as those considered here, would generally be much easier to establish than some of the more contorted examples from the literature. Whilst this is broadly true – our relational reasoning technique is far from complete, yet suffices for establishing the interesting equational consequences of the effect system – completely generic reasoning is surprisingly difficult. When proving concrete equivalences one treats the context generically, but has two particular, literal terms in one’s hand, whose denotations one can work with explicitly. In the case of purely type-based equivalences, on the other hand, both the context and the terms are abstract; all one knows are the types, and the semantics of those types has to capture enough information to justify all instances of the transformation.

An alternative approach to proving ‘difficult’ contextual equivalences is to use techniques based on bisimulation. Originally developed for process calculi by Park and Milner [19], bisimulation was adapted for the untyped lambda calculus by Abramsky [2]. Other researchers, particularly Sumii and Pierce, subsequently developed notions of bisimulation for typed lambda calculi that could deal with the kind of encapsulation (data abstraction) given by existential types [26]. These methods have recently been refined by Koutavas and Wand, and applied to an untyped higher-order language with storage [16] and to object-based calculi. It would be extremely worthwhile to investigate whether bisimulation methods can be applied to the typed (and, as discussed above, type-directed) impure equivalences studied here.²

It has been suggested to us that our results might formally be subsumed by our earlier work [8] via a translation that essentially treats a region as a global variable. Reading / writing within a region is then tracked as a read write from/to the corresponding global variable. First, this handles neither masking nor the allocation effect. More importantly such a translation would at best allow one to transfer a type *inference* for global variables to one for the region calculus but it does not seem to help in any way with the soundness proof.

Our base language is deliberately simple so as to allow us to focus on the salient aspects of the semantics. Nevertheless, it would be useful to extend it in various ways. We sketch how this can be done and what difficulties might be faced.

Recursion To accommodate recursive functions, one needs to phrase the denotational semantics in terms of cpos and partial functions. A new ‘possible non-termination’ effect must also be introduced (and attached to potentially-diverging recursive definitions), which prevents “dead computation” from applying, see [5]. The relations interpreting refined types must then be admissible in the sense that suprema of componentwise related ascending chains are related.

To enforce this condition, we must replace the current definition of $\llbracket T_\varepsilon X \rrbracket_\varphi^\Pi$ with the least admissible relation comprising it. This is because the existential quantification in the definition of this relation thwarts an attempt to show admissibility by induction on types. The least admissible relation R^\dagger comprising a given relation R on a cpo D is explicitly given by the set of all (d, d') for which there exist chains $d = \sup_i d_i, d' = \sup_i d'_i$ where $(d_i, d'_i) \in R$ for all i . With this amended definition the cases go through at the expense of a slightly more cluttered notation.

The notational messiness is considerably palliated by the fact that when faced with a situation where one has an assumption of the form $dR^\dagger d'$ and one wants to show some property $\varphi(d, d')$ where φ is itself admissible then one can w.l.o.g. assume that one has dRd' .

An alternative route to admissibility consists of using a semantics phrased in terms of a continuation-based termination judgement [23, 9]. The existential quantifier in the monad then becomes a universal quantifier and admissibility holds from the start.³ This approach, however, yields an apparently coarser relation which breaks our current proof of Prop 2. Writing $M_1 \simeq M_2$ to mean that the denotations of M_1 and M_2 are in the continuation-based logical relation, we are faced with the apparently simpler goal $(x=M; y=M; (x, y)) \simeq (x=M; (x, x))$, however, we have to prove it under the *weaker* assumption $M \simeq M$. This is an example of the difference between proving particular and generic equiv-

² The more general relationship between logical relations and bisimulation still seems slightly murky (at least to us) and clearly demands further study as well.

³ In fact, rather more holds: the continuation-based formulation uses a ‘ \top -closure’ operation that ‘extensionalizes’ the relations [1].

alences discussed above: the problem does not arise in concrete cases in which we know what M actually is.

Purity Certain valid program equivalences are not provable with our proof method. For example, we would like the semantics to justify the pure typing

$$M_{\text{mem}} : T_{\emptyset}(\text{int} \rightarrow T_{\emptyset} \text{int})$$

for M_{mem} and similar memoised functions on the grounds that they behave just as a pure function. With such a typing, an instance of M_{mem} could be hoisted out of an abstraction which it presently cannot. One way to make our semantics believe that M_{mem} is pure consists of including invariants on private portions of the store again as in [23, 9]. This would allow one to justify a strengthened masking rule with a semantically formulated side condition which requires one to “manually” provide an invariant on privately allocated store. In the case of the memo function M_{mem} one could use $\text{read}(y) = \text{read}(x) + 1$ as invariant. One has to prove that this invariant is maintained and moreover that the result value is independent of the store so long as the invariant is satisfied. At present, we cannot see a syntactic approximation to such a semantic masking rule, i.e., a non-contrived typing rule which would recognize M_{mem} as pure yet of course reject M_{buf} as impure.

General references Of course it would be interesting to include references to types other than integers. References to product and reference types which correspond to the usual Java/C heap model should be relatively straightforward and quite useful. Treating references to functions and values of general mixed-variance recursive types ought to be possible using a simultaneous recursive definition of the logical relation and the sets $\mathcal{R}_{\varepsilon}^{\Pi}(\varphi)$. While doing this is technically challenging, the principal feasibility of this approach has been demonstrated in [25, 11].

References

- [1] M. Abadi. $\top\top$ -closed relations and admissibility. *Mathematical Structures in Computer Science*, 10(3), 2000.
- [2] S. Abramsky. The lazy lambda calculus. In D. Turner, editor, *Research Directions in Functional Programming*, chapter 4, pages 65–116. Addison-Wesley, 1988.
- [3] A. Banerjee, N. Heintze, and J. Riecke. Region analysis and the polymorphic lambda calculus. In *Proceedings of the 14th IEEE Symposium on Logic in Computer Science (LICS)*, 1999.
- [4] N. Benton. Simple relational correctness proofs for static analyses and program transformations. In *31st ACM Symposium on Principles of Programming Languages (POPL)*, 2004. Revised version available from <http://research.microsoft.com/~nick/publications.htm>.
- [5] N. Benton and P. Buchlovsky. Semantics of an effect analysis for exceptions. In *3rd ACM-SIGPLAN Workshop on Types in Language Design and Implementation (TLDI)*, 2007.
- [6] N. Benton, J. Hughes, and E. Moggi. Monads and effects. In *Applied Semantics, Advanced Lectures*, volume 2395 of *LNCS*. Springer-Verlag, 2002.
- [7] N. Benton and A. Kennedy. Monads, effects and transformations. In *3rd International Workshop on Higher Order Operational Techniques in Semantics (HOOTS)*, volume 26 of *ENTCS*. Elsevier, 1999.
- [8] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Reading, writing, and relations: Towards extensional semantics for effect analyses. In *4th Asian Symposium on Programming Languages and Systems (APLAS)*, LNCS, 2006.
- [9] N. Benton and B. Leperchey. Relational reasoning in a nominal semantics for storage. In *7th International Conference on Typed Lambda Calculi and Applications (TLCA)*, volume 3461 of *LNCS*, 2005.
- [10] Nick Benton, John Hughes, and Eugenio Moggi. Monads and effects. In *Applied Semantics, International Summer School, Springer LNCS 2395*, pages 42–122, 2000.
- [11] N. Bohr and L. Birkedal. Relational reasoning for recursive types and references. In *APLAS*, 2006.
- [12] M. Fluet and G. Morrisett. Monadic regions. *Journal of Functional Programming*, 2006. to appear.
- [13] D. K. Gifford and J. M. Lucassen. Integrating functional and imperative programming. In *ACM Conference on LISP and Functional Programming*, Cambridge, Massachusetts, August 1986.
- [14] J. Y. Halpern, A. R. Meyer, and B. A. Trakhtenbrot. The semantics of local storage, or what makes the free-list free? In *Proceedings of the 11th ACM Symposium on Principles of Programming Languages (POPL)*, 1984.
- [15] S. Peyton Jones and J. Launchbury. State in Haskell. *Lisp and Symbolic Computation*, 8(4), 1995.
- [16] Vasileios Koutavas and Mitchell Wand. Small bisimulations for reasoning about higher-order imperative programs. In *POPL*, pages 141–152, 2006.
- [17] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *15th ACM Symposium on Principles of Programming Languages (POPL)*, 1988.
- [18] A. R. Meyer and K. Sieber. Towards a fully abstract semantics for local variables: Preliminary report. In *Proceedings of the 15th Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)*, January 1988.
- [19] R. Milner. *Communication and Concurrency*. Prentice-Hall International, 1989.
- [20] E. Moggi and A. Sabry. Monadic encapsulation of effects: A revised approach (extended version). *Journal of Functional Programming*, 11(6), 2001.
- [21] D. Naumann. Observational purity and encapsulation. *Theoretical Computer Science*. To appear.
- [22] P. W. O’Hearn and R. D. Tennent. Parametricity and local variables. *Journal of the ACM*, 42(3):658–709, May 1995.
- [23] A. M. Pitts and I. D. B. Stark. Operational reasoning for functions with local state. In *Higher Order Operational Techniques in Semantics*. CUP, 1998.
- [24] U. S. Reddy and H. Yang. Correctness of data representations involving heap data structures. *Science of Computer Programming*, 50(1–3):129–160, March 2004.
- [25] B. Reus and J. Schwinghammer. Denotational semantics for a program logic of objects. *Mathematical Structures in Computer Science*, 2006. in press.
- [26] E. Sumii and B. C. Pierce. A bisimulation for type abstraction and recursion. In *POPL ’05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM Press, 2005.
- [27] R. D. Tennent and D. R. Ghica. Abstract models of storage. *Higher-Order and Symbolic Computation*, 13(1/2):119–129, 2000.
- [28] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [29] P. Wadler and P. Thiemann. The marriage of effects and monads. *ACM Trans. Comput. Logic*, 4(1):1–32, 2003.