

Abstract Effects and Proof-Relevant Logical Relations

Nick Benton

Microsoft Research, Cambridge, UK
nick@microsoft.com

Martin Hofmann

LMU, Munich, Germany
hofmann@ifi.lmu.de

Vivek Nigam

UFPB, João Pessoa, Brazil
vivek.nigam@gmail.com

Abstract

We give a denotational semantics for a region-based effect system that supports type abstraction in the sense that only externally visible effects need to be tracked: non-observable internal modifications, such as the reorganisation of a search tree or lazy initialisation, can count as ‘pure’ or ‘read only’. This ‘fictional purity’ allows clients of a module to validate soundly more effect-based program equivalences than would be possible with previous semantics. Our semantics uses a novel variant of logical relations that maps types not merely to partial equivalence relations on values, as is commonly done, but rather to a proof-relevant generalisation thereof, namely setoids. The objects of a setoid establish that values inhabit semantic types, whilst its morphisms are understood as proofs of semantic equivalence. The transition to proof-relevance solves two awkward problems caused by naïve use of existential quantification in Kripke logical relations, namely failure of admissibility and spurious functional dependencies.

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

The last decade has witnessed significant progress in modelling and reasoning about the tricky combination of effects and higher-order language features (first-class functions, modules, classes). The object of study may be ML-like, Java-like, or assembly-like, but the common source of trickiness is the way effectful operations may be *partially* encapsulated behind higher-order abstractions. Several closely-related problems have been addressed using common techniques that include separation, Kripke logical relations and step-indexing. One is to devise models and reasoning principles for establishing contextual (in)equivalences [18, 36]. A second is to establish equivalence between high-level and low-level

code fragments, *e.g.*, for compiler correctness [5, 23]. A third is to define Hoare-style logics for showing programs satisfy assertions [38]. A fourth, which we address here, is to study type systems and analyses that can characterize particular *classes* of behavior (such as purity) and be used to justify equivalences more generically.

Effect systems [22] refine conventional types by adding information capturing an upper bound on the effects expressions may have. Several recent papers [3, 6, 7, 26, 41], have explored the semantics of effect systems, with a focus not merely on showing correctness of analyses, but on providing a rigorous account of the conditions under which effect-dependent optimizations and refactorings are sound. An example is the commutation of stateful computations M and N , subject to the condition that the sets of storage locations potentially written by M and N are disjoint, and that neither potentially reads a location that the other writes.

We seek interpretations of effect-refined types (over an unrefined model) that can justify such equivalences. Indeed, it is the interpretations, rather than rules for assigning such types to terms, that we regard as primary. Types provide a common interface language that can be used by clients in modular reasoning about rewrites; types can be assigned to particular terms by a mixture of more or less sophisticated inference systems, or by deeper semantic reasoning about particular implementations.

A key notion in reasoning about encapsulated state is that of *separation*: invariants depending upon mutually disjoint parts of the store. Intuitively, if each function with direct access to a part preserves the corresponding invariant, then all the invariants will be preserved by any composition of functions. Disjointness is naïvely understood in terms of sets of heap locations. A memory allocator, for example, guarantees that its own private data structures, memory previously handed out to clients, and any freshly-allocated block inhabit mutually disjoint sets of locations. Since the introduction of fractional permissions, work on separation logic often goes beyond this simple model, introducing resources that are combined with a separating conjunction, but which are not literally interpreted as predicates on disjoint locations. Research on ‘domain-specific’ [28], ‘fictional’ [17, 25], ‘subjective’ [30], or ‘superficial’ [29] separation logics and type theories aims to allow custom notions of separable resource to be introduced and combined modularly. This paper presents a semantics for effect systems supporting fictional, or ‘abstract’, notions of both effects and separation.

We have previously interpreted effect-refined types for stateful computations as binary relations, defined via preservation of particular sets of store relations [7]. This already provides some abstraction. The semantics is extensional so, for example, a function that reads a reference cell but doesn’t produce observably different results depending on the value read can soundly be counted as pure (contrasting with, for example, models of permissions that instrument the concrete semantics). Such a semantics can also interpret the ‘masking’ rule, allowing certain non-observable effects not to appear in annotations [6]. But here we go further, generalizing the interpretation of regions to, intuitively, partial equivalence

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

POPL ’14, January 22–24, 2014, San Diego, CA, USA.
Copyright © 2014 ACM 978-1-4503-2544-8/14/01...\$15.00.
<http://dx.doi.org/10.1145/2535838.2535869>

lence relations. This allows, for example, a lookup function for a set ADT to be assigned a read-but-not-write effect, even though the concrete implementation may involve non-observable writes to rebalance an internal datastructure. Moving to PERs in this way requires us to revisit the notion of separation, allowing types to involve distinct regions whose concrete implementations overlap, albeit non-observably, in memory.

Earlier models of dynamic allocation [6] have used Kripke logical relations in which worlds are finite partial bijections between locations, with region-colored links. Two computations $c, c' : \mathbb{H} \rightarrow \mathbb{H} \times \mathbb{V}$, where \mathbb{H}, \mathbb{V} are sets of (partial) heaps and values, respectively, are in the computation relation $T_\varepsilon \llbracket A \rrbracket \mathbf{w}$, for world \mathbf{w} , effect ε and result type A when $\forall h, h', h_1, h'_1. h, h' \models \mathbf{w} \Rightarrow \exists \mathbf{w}_1 \geq \mathbf{w}. h_1, h'_1 \models \mathbf{w}_1 \wedge (h, h_1, h', h'_1, \mathbf{w}, \mathbf{w}_1) \in \llbracket \varepsilon \rrbracket \wedge (v, v') \in \llbracket A \rrbracket \mathbf{w}_1$ where $c h = (h_1, v)$ and $c' h' = (h'_1, v')$. Here, $\llbracket A \rrbracket \mathbf{w}_1$ is the (world-dependent) logical relation at the result type and $\llbracket \varepsilon \rrbracket$ the interpretation of the effect, which can be given in various ways, e.g., in terms of binary relations on stores to be preserved. The notation $h, h' \models \mathbf{w}$ means that for each link $(l, l') \in \mathbf{w}$, $l \in \text{dom}(h), l' \in \text{dom}(h')$.

However, generalizing such a logical relation from bijections to PERs is surprisingly difficult. The problematic part is the existential quantification over world extensions: $\exists \mathbf{w}_1 \geq \mathbf{w}$. This says that c and c' are related at world \mathbf{w} if there is *some* future world \mathbf{w}_1 , allowing for the computations to perform allocations, at which the respective final states and values are related. This pattern of quantification occurs in many accounts of generativity, but the dependence of \mathbf{w}_1 on both h and h' creates a difficulty when one generalizes from bijections to PERs and tries to prove equivalences. Roughly, one has to consider varying the initial heap in which one computation, say c' , is started; the existential then supplies one with a *different* extension that is not at all related, even on the side of c where the heap stays the same, to the one with which one started. In the case of bijections, the fact that h_1 depends only on h (not on h') allows one to deduce sufficient information about the domain of \mathbf{w}_1 from the clause $h_1, h'_1 \models \mathbf{w}_1$, but this breaks down in the more abstract setting.

To fix this problem, we will essentially replace the existential quantifier in the logical relation by appropriate Skolem functions, explicitly enforcing the correct dependencies. In the language of type theory, this amounts to replacing the existential with a Σ -type. A statement like $(c, c') \in T_\varepsilon \llbracket A \rrbracket$ is no longer just a proposition, but we rather have a “set of proofs” $(p : c \sim c') \in T_\varepsilon \llbracket A \rrbracket$, and our constructions will explicitly map proofs to proofs.

We have previously shown [4] how the formalism of *setoids* can be used to make such intuitions both rigorous and more general. That work, dealing with a version of Pitts and Stark’s ν -calculus, used proof-relevant setoids to solve another problem associated with the use of existential quantification: that it fails to preserve admissibility of predicates and relations and so interacts badly with general recursion. Here we show how that technology scales to a much richer language and type system.

2. Motivating Examples

We will assign effect-refined types to terms in a fairly conventional monadically-structured untyped metalanguage with higher-order functions and dynamically allocated references to flat data. The language is defined formally later on, but the following examples of the kind of equivalences and typings we will be able to justify should be comprehensible.

Memoisation Let *memo* be the function

$$\lambda f. \text{let } x \leftarrow \text{ref}(0) \text{ in let } y \leftarrow \text{ref}(f \ 0) \text{ in} \\ \lambda a. \text{if } a = !x \text{ then } !y \text{ else let } r \leftarrow f \ a \text{ in } x := a; y := r; r$$

where $t_1; t_2 = \text{let } _ \leftarrow t_1 \text{ in } t_2$ is sequential composition and $=$ is equality on storable values. So, *memo* returns a memoized

version of its argument. Unlike previous models of effects, our model justifies the typing $\text{memo} : (int \xrightarrow{0} int) \xrightarrow{0} (int \xrightarrow{0} int)$, saying that if f is observationally pure, $\text{memo } f$, is too, and so can participate in any program equivalence relying on purity. For example, a client can deduce the observational equivalence of the programs below purely on the basis of typing:

$$\lambda y. \text{let } g \leftarrow \text{memo} (\lambda x. x + 1) \text{ in } e \equiv \text{let } g \leftarrow \text{memo} (\lambda x. x + 1) \text{ in } \lambda y. e$$

Overlapping references Let p, p^{-1} implement a bijection $\mathbb{Z} \rightarrow \mathbb{Z} \times \mathbb{Z}$, then

$$v_{or} := \lambda _ . \text{let } r \leftarrow \text{ref}(0) \text{ in } (\lambda _ . (p(!r)).1, \lambda _ . (p(!r)).2), \\ \lambda n. \text{let } (x, y) \leftarrow p(!r) \text{ in } r := p^{-1}(n, y), \\ \lambda n. \text{let } (x, y) \leftarrow p(!r) \text{ in } r := p^{-1}(x, n))$$

which multiplexes two abstract integer references onto a single concrete one, can be given the type

$$\tau_{or} := \text{unit} \xrightarrow{alr_1, alr_2} (\text{unit} \xrightarrow{rd_1} int) \times (\text{unit} \xrightarrow{rd_2} int) \times \\ (\text{int} \xrightarrow{wr_1} \text{unit}) \times (\text{int} \xrightarrow{wr_2} \text{unit})$$

expressing that it allocates in two regions, r_1 and r_2 , and returns a tuple of functions, each of which reads or writes just one region. This typing justifies, for example, permuting writes to the two abstract references.

Set factory A more realistic example, for which we do not give concrete code, is a function *setfactory* that generates mutable integer sets. Calling *setfactory* allocates a new reference cell pointing to a linked list of integers (initially empty) and returns a triple of functions (*mem*, *add*, *rem*) for testing membership of the represented set, adding a new integer to the set and removing an integer from the set. The implementation of *mem* searches for its integer argument in the linked list, but at the same time mutates the list by removing duplicates and, just for fun, potentially relocating some of the nodes. The other operations can potentially perform similar optimizations.

We can justify the following semantic typing for *setfactory*:

$$\text{setfactory} : \forall r. \text{unit} \xrightarrow{alr} (\text{int} \xrightarrow{rd_r} \text{bool}) \times (\text{int} \xrightarrow{rd_r, wr_r} \text{unit}) \times (\text{int} \xrightarrow{rd_r, wr_r} \text{unit})$$

which expresses that *setfactory* allocates in some (possibly fresh) region r and returns operations that, observably, only read r (the first one) and/or write in r (the second and third one) even though, physically, all three functions may read, write, and allocate concrete locations.

Thus, these functions can participate in corresponding effect-dependent program equivalences. For example, *mem* operations may be swapped and duplicated.

Lazy Initialization The following definitions illustrate an ‘allocate-on-write’ form of lazy initialization:

$$\begin{aligned} \text{make} &= \lambda _ . \text{ref}(0) \\ \text{getx} &= \lambda p. \text{if } !p = 0 \text{ then } 0 \text{ else } !(p.1) \\ \text{gety} &= \lambda p. \text{if } !p = 0 \text{ then } 0 \text{ else } !(p.2) \\ \text{setx} &= \lambda p. \lambda n. \text{if } !p = 0 \text{ then } p := (\text{ref}(n), \text{ref}(0)) \text{ else } !p.1 := n \\ \text{sety} &= \lambda p. \lambda n. \text{if } !p = 0 \text{ then } p := (\text{ref}(0), \text{ref}(n)) \text{ else } !p.2 := n \end{aligned}$$

We can introduce a new abstract type point_{r_1, r_2} for which the following typings can be semantically justified:

$$\begin{aligned} \text{make} : \text{unit} \xrightarrow{alr_2} \text{point}_{r_1, r_2} \\ \text{getx} : \text{point}_{r_1, r_2} \xrightarrow{rd_{r_1}} int \quad \text{gety} : \text{point}_{r_1, r_2} \xrightarrow{rd_{r_2}} int \\ \text{setx} : \text{point}_{r_1, r_2} \xrightarrow{0} int \xrightarrow{wr_{r_1}} \text{unit} \quad \text{sety} : \text{point}_{r_1, r_2} \xrightarrow{0} int \xrightarrow{wr_{r_2}} \text{unit} \end{aligned}$$

Observe that the allocation effects are flagged in the *make* function, even though the physical allocations happen on demand, in the *set*

functions. Note also that semantic reasoning is necessary to justify that these definitions are well-typed at all, as a *point* holds different kinds of value at different times.

3. Syntax and Semantics

In this section we define the syntax and denotational semantics of our untyped metalanguage for stateful computations. We also give an effect-refined type system for this language that is parameterized by semantically-justified axioms. We omit the standard details of interpreting CBV languages via the metalanguage and of adequacy, relating the operationally induced observational equivalence to equality in the model.

Denotational Model A *predomain* is an ω -cpo, i.e., a partial order with suprema of ascending chains. A *domain* is a predomain with a least element, \perp . Recall that $f : A \rightarrow A'$ is *continuous* if it is monotone $x \leq y \Rightarrow f(x) \leq f(y)$ and preserves suprema of chains, i.e., $f(\sup_i x_i) = \sup_i f(x_i)$. Any set is a predomain with the discrete order. If X is a set and A a predomain then any $f : X \rightarrow A$ is continuous. A subset U of a predomain A is *admissible* if whenever $(a_i)_i$ is an ascending chain in A such that $a_i \in U$ for all i , then $\sup_i a_i \in U$, too. If $f : X \times A \rightarrow A$ is continuous and A is a domain then one defines $f^\dagger(x) = \sup_i f_x^i(\perp)$ with $f_x(a) = f(x, a)$. One has, $f(x, f^\dagger(x)) = f^\dagger(x)$ and if $U \subseteq A$ is admissible and $f : X \times U \rightarrow U$ then $f^\dagger : X \rightarrow U$, too. We denote a partial (continuous) function from set (predomain) A to set (predomain) B by $f : A \rightarrow B$.

We assume two sets \mathbb{L} and \mathbb{V}_b modelling concrete locations, and the R-values that can be stored in locations. We assume that R-values include integers, written $int(n)$ for some $n \in \mathbb{Z}$, booleans, written $bool(b)$ for $b \in \mathbb{B}$, locations $loc(l)$ for some $l \in \mathbb{L}$, and tuples of R-values, written (v_1, \dots, v_n) . We assume that it is possible to tell whether a value is of that form and in this case to retrieve the components. Heaps $h \in \mathbb{H}$ are *finite maps* from $\mathbb{L} \rightarrow \mathbb{V}_b$. The domain of a heap h is $\text{dom}(h) \subseteq \mathbb{L}$. We write $\emptyset \in \mathbb{H}$ for the empty heap. If $v \in \mathbb{V}_b, l \in \text{dom}(h)$ then $h[l \mapsto v]$ is the heap that updates l to contain v ; $new(h, v)$ yields a pair (l, h') where $l \in \mathbb{L}$ and $h' \in \mathbb{H}$.

The following properties hold: $\text{dom}(\emptyset) = \emptyset$, $\text{dom}(h[l \mapsto v]) = \text{dom}(h)$, $(h[l \mapsto v])(l') = \text{if } l = l' \text{ then } v \text{ else } h(l')$, and if $new(h, v) = (l, h')$ then $\text{dom}(h') = \text{dom}(h) \cup \{l\}$ and $l \notin \text{dom}(h)$ and $h'(l) = v$.

Note that heaps and R-values are discrete predomains; we do not here model higher-order store (storing objects with nontrivial order, such as functions). Our stores are ‘flat’, like those of object-oriented languages.

We define the predomain of values \mathbb{V} and the domain of computations \mathbb{C} simultaneously as follows: $\mathbb{C} = \mathbb{H} \rightarrow \mathbb{H} \times \mathbb{V}$ are partial continuous functions from \mathbb{H} to $\mathbb{H} \times \mathbb{V}$, the bottom element being the everywhere undefined function. Values are defined as either R-values, tuples of values or continuous functions from values to computations: $\mathbb{V} \simeq \mathbb{V}_b + \text{fun}(\mathbb{V} \rightarrow \mathbb{C}) + \mathbb{V}^*$. Such domain equations can be solved by standard methods.

Syntax The syntax of untyped values and computations is:

$$\begin{aligned} v &::= x \mid () \mid c \mid (v_1, v_2) \mid v.1 \mid v.2 \mid \text{rec } f \ x = t \\ t &::= v \mid \text{let } x \leftarrow t_1 \text{ in } t_2 \mid v_1 \ v_2 \mid \text{if } v \text{ then } t_1 \text{ else } t_2 \\ &\quad !v \mid v_1 := v_2 \mid \text{ref}(v) \end{aligned}$$

Here, x ranges over variables and constants c over constant symbols, each of which has an associated interpretation $\llbracket c \rrbracket \in \mathbb{V}$; these include numerals \underline{n} with $\llbracket \underline{n} \rrbracket = int(n)$, booleans, arithmetic operations, test functions to tell whether a value is an integer, a function, a pair, or a reference, equality test for R-values, etc. $\text{rec } f \ x = t$ defines a recursive function with body e and recursive calls made via f ; we use $\lambda x.t$ as syntactic sugar in the case when $f \notin \text{fv}(t)$. Finally, $!v$ (reading) returns the contents of location v , $v_1 := v_2$ (writing) updates location v_1 with value v_2 , and $\text{ref}(v)$ (allocating) returns a fresh location initialized with v . The metatheory is sim-

plified by using “let-normal form”, in which the only elimination for computations is let , though we sometimes nest computations as shorthand for let-expanded versions in examples.

The untyped denotational semantics of values $\llbracket v \rrbracket \in \mathbb{V} \rightarrow \mathbb{V}$ and terms $\llbracket t \rrbracket \in \mathbb{V} \rightarrow \mathbb{C}$ is defined in usual way by inductive clauses, which we omit here.

Types are given by the grammar

$$\tau ::= \text{unit} \mid \text{int} \mid \text{bool} \mid A \mid \tau_1 \times \tau_2 \mid \tau_1 \xrightarrow{\varepsilon} \tau_2$$

where A ranges over user-specified abstract types. They will typically include region-indexed reference types such as intref , and also types like lists, sets, and even objects, again possibly refined by regions. The metavariable ε represents an *effect*, that is a subset of some fixed set of elementary effects about which we say more later. The core typing rules for values and computations are shown in Figure 1. The side condition on rule DEAD requires us to ascertain that t_1 terminates for all inputs and heaps satisfying the contracts specified in Γ, ε .

We assume an ambient set of *typing axioms* each having the form (v, τ) where v is a value in the metalanguage and τ is a type meaning that v is claimed to be of type τ and that this will be proved “manually” using the semantics rather than using the typing rules.

We also assume an ambient set of *equality axioms* each having the form (v, v', τ) and asserting semantic equality of v and v' at type τ , again to be justified “manually”. We assume that whenever (v, v', τ) is an axiom so are (v, τ) and (v', τ) .

The equational theory can be classified in three different categories: basic, congruence and effect-dependent. An extract of the basic and congruence equations are depicted in Figure 2. The theory also includes all the usual beta and eta laws and commuting conversions for conditionals as well as for let . We will give a semantic interpretation of typed equality judgments which is sound for observational equivalence.

As with typings, further equations involving effectful computations may be justified semantically in a particular model and added to the theory. In particular, we can justify the soundness of four effect-based equations, namely, dead, duplicated and commuting computation, and pure lambda hoist also shown in Figure 2. The core theory then allows one to deduce new semantic equalities from already proven ones.

4. Abstract Locations

We now define the concept of an *abstract location* which generalizes physical locations in that it models a portion of the store that can be read from and updated. Such portion may comprise a fixed set of physical locations or a varying such set (as in the case of a linked list with some given root). It may also reside in just a part of a physical location, e.g., comprise the two low order bits of an integer value stored in a physical location. Furthermore, the equality on such abstract location may be coarser than physical equality, e.g., two linked lists might be considered equal when they hold the same set of elements, and there may be an invariant, e.g., the linked list should contain integer entries and be neither circular nor aliased with other parts of the heap. This then prompts us to model an abstract location as a partial equivalence relation (PER) on heaps together with two more components that describe how modifications of the abstract location interact with the heap as a whole. Thus, next to a PER, an abstract location also contains a reflexive transitive relation (“guarantee”) modelling its evolution by way of modifying actions.

We will define a notion of when abstract locations are independent, generalizing the traditional notion of separation. However, there is some extra complexity due to the fact that whatever custom notions of equality and separation we introduce, we would like them all to interact well with the underlying built-in allocator. A

$\overline{\Gamma \vdash \text{true} : \text{bool}}$	$\overline{\Gamma \vdash \text{false} : \text{bool}}$	$\overline{\Gamma \vdash \underline{n} : \text{int}}$	$\frac{(v, \tau) \text{ a type axiom}}{\Gamma \vdash v : \tau}$	$\overline{\Gamma, x : \tau \vdash x : \tau}$	$\frac{\Gamma \vdash v : \tau}{\Gamma \vdash v : \tau \& \emptyset}$
$\frac{\Gamma \vdash v : \tau_1 \times \tau_2}{\Gamma \vdash v.i : \tau_i}$	$\frac{\Gamma \vdash v_1 : \tau_1 \quad \Gamma \vdash v_2 : \tau_2}{\Gamma \vdash (v_1, v_2) : \tau_1 \times \tau_2}$	$\frac{\Gamma \vdash v_1 : \tau_1 \xrightarrow{\varepsilon} \tau_2 \quad \Gamma \vdash v_2 : \tau_1}{\Gamma \vdash v_1 \ v_2 : \tau_2 \& \varepsilon}$	$\frac{\Gamma, f : \tau_1 \xrightarrow{\varepsilon} \tau_2, x : \tau_1 \vdash e : \tau_2 \& \varepsilon}{\Gamma \vdash \text{rec } f \ x = e : \tau_1 \xrightarrow{\varepsilon} \tau_2}$	$\frac{\Gamma \vdash e : \tau \& \varepsilon_1 \quad \varepsilon_1 \subseteq \varepsilon_2}{\Gamma \vdash e : \tau \& \varepsilon_2}$	
$\frac{\Gamma \vdash v : \text{int} \quad \Gamma \vdash e_1 : \tau \& \varepsilon \quad \Gamma \vdash e_2 : \tau \& \varepsilon}{\Gamma \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 : \tau \& \varepsilon}$		$\frac{\Gamma \vdash e_1 : \tau_1 \& \varepsilon \quad \Gamma, x : \tau_1 \vdash e_2 : \tau_2 \& \varepsilon}{\Gamma \vdash \text{let } x \leftarrow e_1 \text{ in } e_2 : \tau_2 \& \varepsilon}$	$\frac{\Gamma \vdash t : \tau \& \varepsilon \quad r \notin \text{regs}(\Gamma) \cup \text{regs}(\tau)}{\Gamma \vdash t : \tau \& \varepsilon \setminus \{rd_r, wr_r, al_r\}}$ Masking		

Figure 1. Core rules for effect typing

BASIC EQUATIONS (EXTRACT)

$\frac{\Gamma \vdash t : \tau \& \varepsilon}{\Gamma \vdash t = t : \tau \& \varepsilon}$	$\frac{\Gamma \vdash t = t' : \tau \& \varepsilon}{\Gamma \vdash t' = t : \tau \& \varepsilon}$	$\frac{\Gamma \vdash t = t' : \tau \& \varepsilon \quad \Gamma \vdash t' = t'' : \tau \& \varepsilon}{\Gamma \vdash t = t'' : \tau \& \varepsilon}$	$\frac{\Gamma \vdash v = v' : \tau}{\Gamma \vdash v = v' : \tau \& \emptyset}$	$\frac{\Gamma, \vdash v : \tau_1 \times \tau_2}{\Gamma \vdash v = (v.1, v.2) : \tau_1 \times \tau_2}$
$\frac{\Gamma \vdash v_1 : \tau_1 \quad \Gamma \vdash v_2 : \tau_2}{\Gamma \vdash (v_1, v_2).i = v_i : \tau_i}$	$\frac{\Gamma, f : \tau_1 \xrightarrow{\varepsilon} \tau_2, x : \tau_1 \vdash t = t' : \tau_2 \& \varepsilon}{\Gamma \vdash (\text{rec } f \ x = t) = (\text{rec } f \ x = t') : \tau_1 \xrightarrow{\varepsilon} \tau_2}$		$\frac{\Gamma, f : \tau_1 \xrightarrow{\varepsilon} \tau_2, x : \tau_1 \vdash t : \tau_2 \& \varepsilon \quad \Gamma \vdash v : \tau_1}{\Gamma \vdash (\text{rec } f \ x = t) \ v = t[v/x, (\text{rec } f \ x = t)/f] : \tau_2 \& \varepsilon}$	
$\frac{\Gamma \vdash v : \tau_1 \& \varepsilon \quad \Gamma, x : \tau_1 \vdash t : \tau_2 \& \varepsilon}{\Gamma \vdash \text{let } x \leftarrow v \text{ in } t = t[v/x] : \tau_2 \& \varepsilon}$	$\frac{\Gamma \vdash t_1 : \tau_1 \& \varepsilon \quad \Gamma \vdash t_2 : \tau_2 \& \varepsilon \quad \Gamma, x : \tau_2, y : \tau_1 \vdash t_3 : \tau_3 \& \varepsilon}{\Gamma \vdash \text{let } x \leftarrow (\text{let } y \leftarrow t_1 \text{ in } t_2) \text{ in } t_3 = \text{let } y \leftarrow t_1 \text{ in let } x \leftarrow t_2 \text{ in } t_3 : \tau_3 \& \varepsilon}$		$\frac{(v, v', \tau) \text{ an equality axiom}}{\Gamma \vdash v = v' : \tau}$	

CONGRUENCE EQUATIONS (EXTRACT)

$\frac{\Gamma \vdash v = v' : \tau_1 \times \tau_2}{\Gamma \vdash v.i = v'.i : \tau_i}$	$\frac{\Gamma \vdash v = v' : \text{bool} \& \emptyset \quad \Gamma \vdash t_1 = t'_1 : \tau \& \varepsilon \quad \Gamma \vdash t_1 = t'_1 : \tau \& \varepsilon}{\Gamma \vdash \text{if } v \text{ then } t_1 \text{ else } t_2 = \text{if } v' \text{ then } t'_1 \text{ else } t'_2 : \tau \& \varepsilon}$
---	--

EFFECT BASED EQUATIONS

$\frac{\Gamma \vdash t_1 : \tau_1 \& \varepsilon_1 \quad \Gamma \vdash t : \tau \& \varepsilon \quad \text{wrs}(\varepsilon_1) = \emptyset}{\Gamma \vdash (\text{let } x \leftarrow t_1 \text{ in } t) = t : \tau \& \varepsilon}$ Dead, provided t_1 terminates
$\frac{\Gamma \vdash t_1 : \tau_1 \& \emptyset \quad \Gamma, x : X, y : \tau_1 \vdash t : \tau \& \varepsilon}{\Gamma \vdash (\lambda x. \text{let } y \leftarrow t_1 \text{ in } t) = (\text{let } y \leftarrow t_1 \text{ in } \lambda x. t) : (X \xrightarrow{\varepsilon} \tau) \& \emptyset}$ Pure Lambda Hoist
$\frac{\Gamma \vdash t_1 : \tau_1 \& \varepsilon_1 \quad \Gamma, x : \tau_1, y : \tau_1 \vdash t : \tau \& \varepsilon \quad \text{wrs}(\varepsilon_1) \cap \text{rds}(\varepsilon_1) = \text{als}(\varepsilon_1) = \emptyset}{\Gamma \vdash (\text{let } x \leftarrow t_1 \text{ in let } y \leftarrow t_1 \text{ in } t) = (\text{let } x \leftarrow t_1 \text{ in } t) : \tau \& \varepsilon \cup \varepsilon_1}$ Duplicated
$\frac{\Gamma \vdash t_1 : \tau_1 \& \varepsilon_1 \quad \Gamma \vdash t_2 : \tau_2 \& \varepsilon_2 \quad \Gamma, x : \tau_1, y : \tau_2 \vdash t : \tau \& \varepsilon \quad \text{rds}(\varepsilon_1) \cap \text{wrs}(\varepsilon_2) = \text{rds}(\varepsilon_2) \cap \text{wrs}(\varepsilon_1) = \text{wrs}(\varepsilon_1) \cap \text{wrs}(\varepsilon_2) = \emptyset}{\Gamma \vdash (\text{let } x \leftarrow t_1 \text{ in let } y \leftarrow t_2 \text{ in } t) = (\text{let } y \leftarrow t_2 \text{ in let } x \leftarrow t_1 \text{ in } t) : \tau \& \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon}$ Commuting

Figure 2. Basic, congruence, and effect based equational theory

$\pi \cdot \mathbb{1} = \pi(\mathbb{1})$	$\pi \cdot (a_1, a_2) = (\pi \cdot a_1, \pi \cdot a_2)$
$\pi \cdot (L \subseteq \mathbb{L}) = \{\pi(l) \mid l \in L\}$	$\text{dom}(\pi \cdot h) = \pi \cdot (\text{dom}(h))$
$(\pi \cdot h)(l) = \pi \cdot (h(\pi^{-1}(l)))$	$\pi \cdot R = \{\tilde{h} \mid \pi^{-1} \cdot \tilde{h} \in R\}$
$\pi \cdot (h_1, \dots, h_n) = (\pi \cdot h_1, \dots, \pi \cdot h_n)$	

Figure 3. Action of permutations

freshly allocated concrete cell should be independent of any existing abstract locations, for example, and abstract locations should not be sensitive to the behaviour of the underlying allocator. Were we working in a lower-level language (in which the allocator is just another piece of code to be verified), we would hope that concrete locations could be treated as a special case of abstract ones, but this seems to require a notion of when one notion of location is an abstraction of another (Jensen and Birkedal have taken some steps along such a path [25]). Instead, we here bake in compatibility with the underlying allocator using nominal ideas

Following Gabbay and Pitts [21], we write $\text{perm}(\mathbb{L})$ for the group of permutations of \mathbb{L} and define an action $\pi \cdot -$ of permutations $\pi \in \text{perm}(\mathbb{L})$ for values, stores, sets of locations, and relation on stores. Table 3 contains the most important cases: Note that permutations act differently on sets of locations and on sets / relations on heaps. In the sequel we call “object” anything that the

permutations act on, *i.e.* locations, values, stores, etc. For example, a list of heaps \vec{h} is in the relation $\pi \cdot R$ iff the list of heaps obtained by undoing the permutation ($\pi^{-1} \cdot \vec{h}$) is in the relation R . We define the set of fixpoints of a permutation π as $\text{fix}(\pi) = \{\mathbb{1} \mid \pi(\mathbb{1}) = \mathbb{1}\}$. For an object x , a support of x is a set of locations L such that $\text{fix}(\pi) \supseteq L \Rightarrow \pi \cdot x = x$. Gabbay and Pitts show [21] that if an object has a finite support then it has a (unique) *least* support, which we write $\text{supp}(x)$.

It is clear that stores, values, and finite sets of locations always have a finite support. A store relation having a finite support is called finitely supported. For example, the predicate P on stores asserting that there is a linked list starting at location X has support $\{X\}$. Note that although a linked list in h witnessing that $h \in P$ may have a large “footprint”, the support still is just $\{X\}$ because the internal nodes of such a linked list will be consistently renamed by the permutation action.

Definition 4.1 (Abstract Location). *An abstract location $\mathbb{1}$ (on the chosen set \mathbb{H}) consists of the following data:*

- a non-empty, finitely supported, partial equivalence relation (PER) $\mathbb{1}^R$ on \mathbb{H} modelling the “semantic equality” on the bits of the store that $\mathbb{1}$ uses (a “rely-condition”);
- a finitely supported transitive relation $\mathbb{1}^G$ modelling what bits of the store “writes to $\mathbb{1}$ ” leave intact (a “guarantee condition”).

That is, if $(h, h_1) \in \mathcal{I}^G$ then h_1 might arise by writing to l in h and all possible writes are specified by \mathcal{I}^G ;

subject to the conditions, where π is a permutation and $h : \mathcal{I}^R$ is shorthand for $(h, h) \in \mathcal{I}^R$:

1. if $\text{fix}(\pi) \supseteq \text{supp}(\mathcal{I}^R)$ then $h : \mathcal{I}^R \Rightarrow (h, \pi \cdot h) \in \mathcal{I}^R$.
2. if $h : \mathcal{I}^R$ then $(h, h) \in \mathcal{I}^G$;
3. if $(h, h_1) \in \mathcal{I}^G$ then $h : \mathcal{I}^R$ and $h_1 : \mathcal{I}^R$;
4. $\mathcal{I}^R; \mathcal{I}^G \subseteq \mathcal{I}^G; \mathcal{I}^R$, that is, if $(h, h') \in \mathcal{I}^R$ and $(h, h_1) \in \mathcal{I}^G$, then there exists h'_1 such that $(h', h'_1) \in \mathcal{I}^G$ and $(h_1, h'_1) \in \mathcal{I}^R$.
5. if $h : \mathcal{I}^R$ and $h_1 \supseteq h$ then $(h_1, h) \in \mathcal{I}^R$ thus \mathcal{I}^R “looks” no further than the currently allocated portion of the heap.
6. if $h : \mathcal{I}^R$ and $(h \uplus h', h'_1) \in \mathcal{I}^G$, then $\exists h_1. (h, h_1) \in \mathcal{I}^G \wedge h'_1 = h_1 \uplus h'$.

Condition 1 asserts that semantic equality is closed under relocation. It rules out relations that stipulate equality of internal pointers (not to fixed ones which would enlarge the support, but between h and h'). Conditions 5 and 6 ensure that relies and guarantees interact well with future extensions of the heap, preserving and being independent of fresh locations.

We write $h \stackrel{\sim}{\sim} h'$ to denote $(h, h') \in \mathcal{I}^R$; $h \xrightarrow{l} h_1$ to denote $(h, h_1) \in \mathcal{I}^G$. We also use the notation $\text{supp}(l)$ for $\text{supp}(\mathcal{I}^R) \cup \text{supp}(\mathcal{I}^G)$.

If π is a permutation we define $\pi \cdot l = (\pi \cdot \mathcal{I}^R, \pi \cdot \mathcal{I}^G)$; it is easy to see that $\pi \cdot l$ is an abstract location and $\text{fix}(\pi) \supseteq \text{supp}(l) \Rightarrow \pi \cdot l = l$.

Furthermore, for any π and π' such that $\pi|_{\text{supp}(l)} = \pi'|_{\text{supp}(l)}$, one has $\pi \cdot l = \pi' \cdot l$. This makes the following definition possible:

Definition 4.2 (relocation of abstract locations). *Let l be an abstract location and $u : \text{supp}(l) \rightarrow \mathbb{L}$ be injective. Then $u \cdot l$ is defined as $\pi \cdot l$ for some π extending u . We have $\text{supp}(u \cdot l) = u(\text{supp}(l))$.*

For an example, consider the following abstract location set_X , with support $\{X\}$, implementing the set example introduced above:

$$\begin{aligned} \text{set}_X^R &= \{(h, h') \mid \text{If } h \text{ and } h' \text{ contain a linked list starting from } X \\ &\quad \text{and these lists contain the same set of integers.}\} \\ \text{set}_X^G &= \{(h, h_1) \mid h : \mathcal{I}_X^R, h_1 : \mathcal{I}_X^R \text{ and } \forall l. l \in \text{dom}(h) \setminus F(h, X) \Rightarrow h(l) = h_1(l)\} \end{aligned}$$

where $F(h, X)$ is the set of the locations in h reached by the linked list starting from X . The rely set, set_X^R , specifies that two heaps are equivalent when linked lists denote the same set of integers. The guarantee, set_X^G , specifies that a write to this abstract location may not change concrete locations not in the linked list.

Note, in particular, how this example validates condition Definition 4.1(1): We can relocate one of the two linked lists without compromising relatedness in set_X^R provided we leave the entry point X fixed. If one tried to modify our abstract location to relate linked lists with the exact same footprint, then property 1 would fail.

Definition 4.3 (Separated Heap). *Let l_1, l_2, \dots, l_n be abstract locations. A heap h is separated with respect to l_1, l_2, \dots, l_n , written $\text{sep}(h, \{l_1, l_2, \dots, l_n\})$, if for all $1 \leq i \leq n$. $h : \mathcal{I}_i^R$ and for all $j \neq i$ and for any heap h_1 such that $h \xrightarrow{l_i} h_1$, we have $h \stackrel{\sim}{\sim} h_1$ and, coinductively, $\text{sep}(h_1, \{l_1, l_2, \dots, l_n\})$.*

Definition 4.4 (Independence). *Abstract locations l_1, l_2, \dots, l_n are independent, written $\perp(\{l_1, l_2, \dots, l_n\})$, if $h_1 : l_1, \dots, h_n : l_n$, then there exists h , where $\text{sep}(h, \{l_1, l_2, \dots, l_n\})$ and $h \stackrel{\sim}{\sim} h_i$ for $1 \leq i \leq n$.*

The definition of independence implies that quotienting by the intersection of the PERs yields a product of the individual quotients. As the separated heap h as defined above is unique up to \sim (with respect to l_1, \dots, l_n), we use the notation $h_1 \oplus h_2 \oplus \dots \oplus h_n$ to denote a canonical separated heap.

Returning to the *setfactory* example, consider two sets set_X and set_Y , where X and Y are the concrete locations where the linked lists implementing, respectively, set_X and set_Y start. If X and Y are different, we use coinduction to show set_X and set_Y are independent, using the following set as the witnessing co-inductive hypothesis

for the separated heaps requirement:

$$\begin{aligned} P_{\text{set}} &= \{h \mid \exists h_1, h_2. (h = h_1 \uplus h_2) \wedge \\ &\quad ([X \mapsto \text{nil}] \xrightarrow{\text{set}_X} h_1) \wedge ([Y \mapsto \text{nil}] \xrightarrow{\text{set}_Y} h_2)\}. \end{aligned}$$

This set contains the disjoint unions of any evolution of the linked lists pointed by the concrete locations X and Y .

The following lemma will be helpful for relocating abstract locations and asserting independence.

Lemma 4.5. *If $\text{supp}(l_1) \cap \text{supp}(l_2) = \emptyset$, then $l_1 \perp l_2$.*

Proof Let $h_1 : l_1$ and $h_2 : l_2$. Let $L_1 = \text{dom}(h_1)$ and $L_2 = \text{dom}(h_2)$. Choose permutations π_1 and π_2 such that $\pi_1 \cdot L_1 \cap \pi_2 \cdot L_2 = \emptyset$ and $\text{fix}(\pi_i) \supseteq \text{supp}(l_i^R)$. Define the pasted heap $h_1 \oplus h_2$ by

$$(h_1 \oplus h_2)(l) = \begin{cases} (\pi_1 \cdot h_1)(l), & \text{if } l \in \pi_1 \cdot L_1 \\ (\pi_2 \cdot h_2)(l), & \text{if } l \in \pi_2 \cdot L_2 \\ \text{undefined otherwise} \end{cases}$$

Thus, in particular, $\text{dom}(h_1 \oplus h_2) = \pi_1 \cdot L_1 \uplus \pi_2 \cdot L_2$.

From Def. 4.1(1) we know $(h_1, \pi_1 \cdot h_1) \in \mathcal{I}_1^R$. We also have $(h_1, \pi_1^{-1} \cdot (h_1 \oplus h_2)) \in \mathcal{I}_1^R$ by Def. 4.1(5), thus $(\pi_1 \cdot h_1, h_1 \oplus h_2) \in \mathcal{I}_1^R$ so we can conclude by transitivity. The remaining properties are left to the reader. \square

Definition 4.6. *If l_1, l_2 are independent, we form a joint location $l_1 \otimes l_2$ specified as follows:*

- $(l_1 \otimes l_2)^R = \{(h, h') \mid (h, h') \in \mathcal{I}_1^R \cap \mathcal{I}_2^R \text{ and } \text{sep}(h, \{l_1, l_2\}) \text{ and } \text{sep}(h', \{l_1, l_2\})\}$
- $(l_1 \otimes l_2)^G = (\mathcal{I}_1^G \cup \mathcal{I}_2^G)^*$ restricted to $(l_1 \otimes l_2)^R$;

Lemma 4.7. *If l_1 and l_2 are independent abstract locations, then the joint location $l_1 \otimes l_2$ is an abstract location.*

Lemma 4.8. *Let l_1 and l_2 be independent abstract locations. Then $(l_1 \otimes l_2)^G \subseteq \mathcal{I}_1^G; \mathcal{I}_2^G; (l_1 \otimes l_2)^R$.*

Lemma 4.9. $\perp(\{l_1, l_2, l_3\})$ if and only if $l_2 \perp l_3$ and $l_1 \perp (l_2 \otimes l_3)$.

Definition 4.10. *Let $\mathbf{1}$ be the abstract location as follows, $\mathbf{1}^R = \mathbb{H} \times \mathbb{H}$ and $(h, h') \in \mathbf{1}^G$ if and only if $h = h'$. Notice that $\text{supp}(\mathbf{1}) = \emptyset$.*

Lemma 4.11. $l_1 \perp l_2$ if and only if and $l_2 \perp l_1$. Moreover, $\mathbf{1} \otimes l = l$.

Remark 4.12. *Motivated by the axioms of separation algebras [13] one might conjecture that \otimes is cancellative, i.e. that $l_1 \otimes l = l_2 \otimes l$ implies $l_1 = l_2$. However, for the above definitions, this does not generally hold.*

4.1 Examples of Abstract Locations

Besides abstraction locations for sets of integers used above, we illustrate other instances of abstract locations:

Single Integer For our simplest example, consider an abstract location l_X , parametric with respect to X as follows:

$$\begin{aligned} \text{int}_X^R &= \{(h, h') \mid h(X) = n \wedge h'(X) = n' \Rightarrow \text{int}(n) = \text{int}(n')\} \\ \text{int}_X^G &= \{(h, h_1) \mid h : \text{int}_X^R, h_1 : \text{int}_X^R \text{ and } \forall l \in \mathbb{L}. l \neq X \Rightarrow h(l) = h_1(l)\} \end{aligned}$$

Two heaps are in its rely relation if the values stored in X are the same; and its guarantee is to leave all other concrete locations alone.

It is also the case that int_X and int_Y are independent, i.e., $\text{int}_X \perp \text{int}_Y$. This is witnessed by the following co-inductive hypothesis to show the separateness condition of heaps with respect to these locations, similar to the one used before for the showing the independence of set_X and set_Y

$$P_{\text{int}} = \{h \mid \exists h_1, h_2. (h = h_1 \uplus h_2) \wedge ([X \mapsto 0] \xrightarrow{\text{int}_X} h_1) \wedge ([Y \mapsto 0] \xrightarrow{\text{int}_Y} h_2)\}.$$

Overlapping references Recall the overlapping references example introduced earlier. Let X be the concrete location encoding a

pair of values. We define abstract locations \mathfrak{fst}_X and \mathfrak{snd}_X by:

$$\begin{aligned}\mathfrak{fst}_X^R &= \{(h, h') \mid h(X) = (a_1, a_2) \wedge h'(X) = (a'_1, a'_2) \wedge a_1 = a'_1\} \\ \mathfrak{snd}_X^R &= \{(h, h') \mid h(X) = (a_1, a_2) \wedge h'(X) = (a'_1, a'_2) \wedge a_2 = a'_2\} \\ \mathfrak{fst}_X^G &= \{(h, h_1) \mid h : \mathfrak{fst}_X^R, h_1 : \mathfrak{fst}_X^R \text{ and} \\ &\quad \forall l \in \mathbb{L}. l \neq X \wedge h(X) = (a_1, a_2) \wedge h_1(X) = (a'_1, a'_2) \Rightarrow \\ &\quad h(l) = h_1(l) \wedge a_2 = a'_2\} \\ \mathfrak{snd}_X^G &= \{(h, h_1) \mid h : \mathfrak{snd}_X^R, h_1 : \mathfrak{snd}_X^R \text{ and} \\ &\quad \forall l \in \mathbb{L}. l \neq X \wedge h(X) = (a_1, a_2) \wedge h_1(X) = (a'_1, a'_2) \Rightarrow \\ &\quad h(l) = h_1(l) \wedge a_1 = a'_1\}\end{aligned}$$

The rely of \mathfrak{fst}_X (respectively, \mathfrak{snd}_X) specifies that two heaps h and h' are equivalent whenever they both store a pair of values in X and the first projections (respectively, second projection) of these pairs are the same. The guarantee of \mathfrak{fst}_X (respectively, \mathfrak{snd}_X) specifies that it keeps all other locations alone and does not change the second projection (respectively, first projection) of the pair stored at location X .

By using a coinduction hypothesis for separated heaps similar to the ones used above, it is straightforward to verify that \mathfrak{fst}_X and \mathfrak{snd}_X are independent, although they share the concrete location X .

Lazy Initialization Recall the Lazy Initialization example introduced earlier. Let X be the concrete location used to store the data that is initialized. We define the following abstract location parametric on $i \in \{1, 2\}$:

$$\begin{aligned}\mathfrak{ia}_{3\mathfrak{v}}(i)_X^R &= \{(h, h') \mid h(X) = 0 \Rightarrow \mathit{initial}(h', X) \text{ and} \\ &\quad h'(X) = 0 \Rightarrow \mathit{initial}(h, X) \text{ and} \\ &\quad h(X) = (a_1, a_2) \wedge h'(X) = (a'_1, a'_2) \Rightarrow h(a_i) = h'(a'_i)\} \\ \mathfrak{ia}_{3\mathfrak{v}}(i)_X^G &= \{(h, h_1) \mid h : \mathfrak{fst}_X^R, h_1 : \mathfrak{fst}_X^R \text{ and } h_1(X) = 0 \Rightarrow h(X) = 0 \text{ and} \\ &\quad h(X) = (a_1, a_2) \wedge h_1(X) = (a'_1, a'_2) \Rightarrow \\ &\quad \forall l. l \notin \{X, a_i, a'_i\} \Rightarrow h(l) = h_1(l)\}\end{aligned}$$

where $\mathit{initial}(h, X)$ is the predicate $h(X) = 0 \vee (h(X) = (a_1, a_2) \wedge h(a_1) = 0 = h(a_2))$.

The rely specifies that two heaps are equivalent when either X is in the initial, uninitialized, state in both of them, as specified by the predicate $\mathit{initial}$, or X in both heaps points to a pair of references, the i -th of which also contain the same values. The guarantee condition specifies that X in an updated heap is 0 only if it was so before, and moreover all locations not used in the rely are left unchanged.

5. Setoids

Abstract locations capture the invariants that we will use to build our model. The ‘big picture’ is that types are modelled as functors from a category of worlds into a category of setoids, with terms being interpreted as natural transformations. Worlds will comprise independent abstract locations, whilst setoids are predomains augmented with a proof-relevant notion of equality. In this section, we define the category of setoids more precisely. Section 6 then explains the structure of the category of worlds and setoid-valued functors.

While we assume that the reader is familiar with some notions of category theory, such as functor, natural transformation, and cartesian closure [31], we review some basic concepts and notation here. We compose morphisms in the usual applicative order, thus, if $f : A \rightarrow B$ and $g : B \rightarrow C$ then $gf : A \rightarrow C$. A morphism u in a category C is a monomorphism if $ux = ux'$ implies $x = x'$ for all morphisms x, x' .

A commuting square $xu = x'u'$ of morphisms is a pullback if whenever $xv = x'v'$ there is unique t such that $v = ut$ and $v' = u't$. A pair of morphisms u, u' with common domain is a span, a pair of morphisms x, x' with common codomain is a co-span. A category has pullbacks if every co-span can be completed to a pullback square.

We define the *category of setoids* as the exact completion of the category of predomains, see [10, 14]. We recall here the elementary description using the language of dependent types given in an earlier paper *et al.* [4].

Definition 5.1. A setoid A consists of a predomain $|A|$ and for any two $x, y \in |A|$ a set $A(x, y)$ of “proofs” (that x and y are equal). The set of triples $\{(x, y, p) \mid p \in A(x, y)\}$ must itself be a predomain and the first and second projections must be continuous. Furthermore, there are continuous functions $r_A : \prod x \in |A|. A(x, x)$ and $s_A : \prod x, y \in |A|. A(x, y) \rightarrow A(y, x)$ and $t_A : \prod x, y, z. A(x, y) \times A(y, z) \rightarrow A(x, z)$, witnessing reflexivity, symmetry and transitivity; note that no equations between these are imposed.

In the above definition, continuity of a dependent function such as $t(-, -)$ means the following: If $(x_i)_i$ and $(y_i)_i$ and $(z_i)_i$ are ascending chains in A with suprema x, y, z , and $p_i \in A(x_i, y_i)$ and $q_i \in A(y_i, z_i)$ are proofs such that (x_i, y_i, p_i) ; and (y_i, z_i, q_i) ; are also ascending chains, with suprema (x, y, p) and (y, z, q) , then $(x_i, z_i, t(p_i, q_i))$ is an ascending chain of proofs (by monotonicity of $t(-, -)$) and its supremum is $(x, z, t(p, q))$.

If $p \in A(x, y)$ we may write $p : x \sim y$ or simply $x \sim y$. We also omit $|-|$ wherever appropriate.

Setoids based on types rather than predomains have been used to provide extensional equality in intuitionistic type theory [2].

Definition 5.2. A morphism from setoid A to setoid B is an equivalence class of pairs $f = (f_0, f_1)$ of continuous functions where $f_0 : |A| \rightarrow |B|$ and $f_1 : \prod x, y \in |A|. A(x, y) \rightarrow B(f_0(x), f_0(y))$. Two such pairs $f, g : A \rightarrow B$ are identified if there exists a continuous function $\mu : \prod a \in |A|. B(f(a), g(a))$.

Morphisms are composed in the obvious way to form the category of setoids. In the sequel we tend to omit $|-|$ and 0, 1-subscripts where appropriate, thus writing $x(a)$ for $x_0(a)$ or $a \in A$ or even $a : A$ for $a \in |A|$ etc.

We remark that one could also choose not to identify \sim -equal morphisms. We decided to make the identification because standard category-theoretic concepts like cartesian closure, pullbacks, etc. then apply directly, rather than only holding up to \sim .

The following is folklore, see also [10].

Proposition 5.3. The category of setoids is cartesian-closed. The cartesian product $A \times B$ of setoids A, B has components $|A \times B| = |A| \times |B|$ and $(A \times B)((a, b), (a', b')) = A(a, a') \times B(b, b')$. The ordering is given component-wise.

The underlying predomain of the function space of $A \Rightarrow B$ of setoids A, B is given by the set of representatives (f_0, f_1) of morphisms from A to B ordered component- and point-wise. Given two such elements $f = (f_0, f_1)$ and $g = (g_0, g_1)$ a proof $\mu : f \sim g$ in $A \Rightarrow B$ is a continuous function establishing equality of f and g qua morphisms from A to B as in Definition 5.2.

Proof The verification of the cartesian product is straightforward. The application morphism ap from $(A \Rightarrow B) \times A$ to B has components given by $ap_0((f_0, f_1), a) = f_0(a)$ and $ap_1(\mu, p) = f_1(p); \mu(a')$ where $\mu : f \sim g$ and $p : a \sim a'$. Note that we have $f_1(p) : f_0(a) \sim f_0(a')$ and $\mu(a') : f_0(a') \sim g_0(a')$.

Conversely, if $h : C \times A \rightarrow B$ is a morphism represented by (h_0, h_1) we define $\lambda(h) : C \rightarrow A \Rightarrow B$ by $(\lambda h)_0(c) = (f_0, f_1)$ where $f_0(a) = h_0(c, a)$ and $f_1(p) = h_1(r_C(c), p)$. We also put $(\lambda h)_1(p : c \sim c')(a) = h_1(p, r_A(a))$. If (h'_0, h'_1) is another representative of h witnessed by a proof μ then $\lambda c. \lambda a. \mu(c, a)$ witnesses equality of the so constructed abstraction. The remaining verifications are direct. \square

Proposition 5.4. The category of setoids has pullbacks.

Proof To construct a pullback of $x : A \rightarrow B$ and $x' : A' \rightarrow B$ pick representatives (x_0, x_1) and (x'_0, x'_1) and take the low

point of the pullback U to be given by $|U| = \{(a, a', p) \mid a \in A, a' \in A', p \in B(x(a), x'(a'))\}$ and then $U((a, a', p), (a_1, a'_1, p_1)) = A(a, a_1) \times A(a', a'_1)$. The ordering is component wise. The morphism $u : U \rightarrow A$ is defined by $u_0(a, a', q) = a$ and $u_1(p, p') = p$. The remaining definitions and verifications are straightforward. Note, however, that when proving the uniqueness of fill-ins to establish the pullback property we make use of the fact that pointwise \sim -equal morphisms are identified. \square

Definition 5.5. A setoid D has a least element if $|D|$ has a least element \perp and there is also a least proof $\perp \in D(\perp, \perp)$.

Proposition 5.6. For each D with least element there is a morphism $Y_D : (D \Rightarrow D) \rightarrow D$ such that the following hold:

- $Y_D = ap(id_{D \Rightarrow D}, Y_D)$ (so $Y_D(f)$ is a fixpoint of f).
- Suppose that both D, D' have least elements and $u : D \rightarrow D'$ satisfy $u_0(\perp) = \perp$. Let $D \Rightarrow D \leftarrow^s X \rightarrow^s D' \Rightarrow D'$ be the pullback of $u(-) : D \Rightarrow D \rightarrow D' \Rightarrow D'$ and $u(-) : D' \Rightarrow D' \rightarrow D \Rightarrow D'$. So, intuitively, X contains pairs (f, f') with $f : D \Rightarrow D, f' : D' \Rightarrow D'$ and $uf = f'u$. Then $uY_D g = Y_{D'} g'$. Intuitively, this means that if $uf = f'u$ and $u(\perp) = \perp$ then $u(Y_D(f)) = Y_{D'}(f')$.

Proof Given $f \in D \Rightarrow D$ construct the least fixpoint of f_0 as $\sup_i f_0^i(\perp)$. The passage from f to the least fixpoint being continuous this defines the first part of Y_D , i.e., $(Y_D)_0(f) = \sup_i f_0^i(\perp)$. As for the proof component assume that $p : f \sim f'$. By repeatedly invoking p we then get a family of proofs $p_i : f_0^i(\perp) \sim f_0'^i(\perp)$. Using the fact that $r(\perp)$ is the least proof we can then show by induction on i that this family of proofs is increasing and its supremum satisfies $\sup_i p_i : Y_D(f) \sim Y_{D'}(f')$. Thus, we put $(Y_D)_1(p) = \sup_i p_i$ which again is continuous.

It is obvious from the construction as a least fixpoint that the required fixpoint equation is indeed satisfied. As for the second one pick $u : D \Rightarrow D'$ and assume $f : D \Rightarrow D$ and $f' : D' \Rightarrow D'$ and a proof $p : uf \sim f'u$. Now, by repeatedly invoking p we get an increasing family of proofs $p_i : u f^i(\perp) \sim f'^i(u(\perp)) = f'^i(\perp)$ whose supremum witnesses $u(Y_D(f)) \sim Y_{D'}(f')$. This establishes the claim. \square

6. Worlds and Setoid-Valued Functors

We now organize abstract locations into a category that will index the meanings of types and terms in a functor category semantics [33, 39]. Our category is morally similar to that of finite sets and injections, which has been used in other work [6, 39]. The difference is that objects are now sets of mutually independent abstract locations and morphisms are slightly more complicated in that they allow relocation of abstract locations. Furthermore, locations are grouped into regions allowing one to approximate them statically.

We begin with some more notational preliminaries. Write $x \diamond_u^{x'} u'$ or $w_u^x \diamond_u^{x'} w'$ (when $w^{(')} = \text{dom}(x^{(')})$) for a pullback square with morphisms x, u, x', u' . We call the common codomain of x and x' the apex of the pullback written \bar{w} , while the common domain of u, u' , the low point of the square, is written \underline{w} . A pair of morphisms u, u' with common domain is a span, a pair of morphisms x, x' with common codomain is a co-span. A category has pullbacks if every co-span can be completed to a pullback square. A pullback square $w_u^x \diamond_u^{x'} w'$ with apex \bar{w} is minimal if whenever there is another pullback $w_u^x \diamond_u^{x'} w'$ over the same span and with apex \bar{w}_1 , then there is a unique morphism $t : \bar{w} \rightarrow \bar{w}_1$ such that $x_1 = tx$ and $x'_1 = tx'$.

Lemma 6.1. If a category \mathbf{C} has pullbacks and $w_u^x \diamond_u^{x'} w'$ with apex \bar{w} is a minimal pullback then the morphisms x and x' are jointly epic, that is, if $fx = gx$ and $fx' = gx'$ then $f = g$.

For illustration purposes, we notice that the category of finite sets and injections has pullbacks and in it every span can be completed to a minimal pullback. Indeed, if $u : U \rightarrow A$ and

$u' : U \rightarrow A'$ then the apex of the minimal pullback can be chosen as $A \setminus \text{Img}(u) + A'$ with $+$ denoting disjoint union.

Definition 6.2 (Category of worlds). Assume given an infinite set Regs of region names. The category of worlds \mathbf{W} is defined as follows:

- An object (written w and called a world) comprises a finite set of mutually independent abstract locations written $\text{dom}(w)$ and a function tagging each abstract location in w with a region from Regs ; we write $w(r)$ for the set of abstract locations in w tagged with r . We define $\text{supp}(w) = \bigcup_{l \in \text{dom}(w)} \text{supp}(l)$.
- Let w and w' be two worlds. A morphism from w to w' is given by an injective function $u : \text{supp}(w) \rightarrow \text{supp}(w')$ such that for each $l \in \text{dom}(w)$ we have $u \cdot l \in \text{dom}(w')$ (cf. Def. 4.2) and moreover, if $l \in w(r)$ then $u \cdot l \in w'(r)$.
- We write $I(w, w')$ for the set of morphisms $u : w \rightarrow w'$ that are set-theoretic inclusions. Note that $u \cdot l = l$ in this case.
- The empty world is denoted \emptyset , i.e. $\text{dom}(\emptyset) = \emptyset$.

For example, a morphism between $\{\text{int}_X\} \rightarrow \{\text{int}_Y, \text{set}_Z\}$ (region tags omitted) is given by the injective function that simply maps the concrete location X to the concrete location Y .

The abstract locations grouped together in one region represent all the stateful components, be they physical locations or more complex updatable data structures, that belong to that region. Static effect annotations amalgamate possible side effects on any of those components as effects on the region.

The notion of morphism presented here is the simplest possible one that allows us to justify interesting effect-dependent equivalences. There are more general notions of morphism that allow one to justify equivalences relying on representation independence, e.g. ones resulting from replacing int_X by snd_Y in a program.

It is also possible to define the category of worlds differently by combining all the abstract locations contained in a given region using the \otimes -operation. A world then becomes a function assigning to each region a single ‘large’ abstract location, possibly the trivial one. The presence of a morphism from one world to another then has to imply that the abstract locations attached to a region grow properly, i.e., go from l to an abstract region isomorphic to $l \otimes l'$, which can be elegantly captured by adapting Oles’s [33] idea of morphisms between store shapes, which has later been popularized as *lenses* in the context of bidirectional synchronization [20].

We reluctantly decided, for reasons of space and simplicity, against using these more general notions of worlds and morphisms here, but plan to discuss them in an extended version of the paper.

The following is proved just as in the case of finite sets and injections.

Proposition 6.3. In the category \mathbf{W} all morphisms are monomorphisms; \mathbf{W} has pullbacks and all spans in \mathbf{W} can be completed to minimal pullbacks.

Moreover, these pullbacks and minimal pullbacks can be chosen in such a way that if one of the two given maps is an inclusion then the one parallel to it is also an inclusion.

If $u : w \rightarrow w_1$ and $x : w_1 \rightarrow w_2$ then we can find w'_1 and $u' : w \rightarrow w'_1$ and $x' : w'_1 \rightarrow w_2$ so that $w_u^x \diamond_u^{x'} w'_1$ is a minimal pullback and if one of u, x is an inclusion then the opposite morphism can be chosen as an inclusion as well.

The world w'_1 in the last clause comprises the locations in w and those in w_2 that are not in w_1 . Intuitively, if $w = A, w_1 = A + B, w_2 = A + B + C$ then $w'_1 = A + C$.

6.1 Setoid-valued functors

A functor A from a category of worlds \mathbf{W} to the category of setoids comprises as usual for each $w \in \mathbf{W}$ a setoid Aw and for each $u : w \rightarrow w'$ a morphism of setoids $Au : Aw \rightarrow Aw'$ preserving identities and composition. If $u : w \rightarrow w'$ and $a \in Aw$ we may

write $u.a$ or even ua for $Au(a)$ and likewise for proofs in Aw . Note that $(uv).a = u.(v.a)$.

Definition 6.4. We call a functor A pullback-preserving (p.p.f.) if for every pullback square $w_u^x \diamond_u^x w'$ with apex \bar{w} and low point \underline{w} the diagram below is a pullback in the category of setoids.

$$\begin{array}{ccc} & A\bar{w} & \\ & \swarrow & \searrow \\ Aw & & Aw' \\ & \nwarrow & \nearrow \\ & A\underline{w} & \end{array}$$

Given our characterization of pullbacks in the category of setoids (Prop. 5.4), this condition is equivalent to the existence of a continuous function of type

$$\begin{aligned} \Pi a \in Aw. \Pi a' \in Aw'. A\bar{w}(x.a, x'.a') \rightarrow \\ \Sigma \underline{a} \in A\underline{w}. A\underline{w}(u.\underline{a}, a) \times Aw'(u'.\underline{a}, a') \end{aligned}$$

which amounts to saying that if two values $a \in Aw$ and $a' \in Aw'$ are equal in a common world \bar{w} , then this can only be the case because there is a value in the “intersection world” \underline{w} from which both a and a' arise. The witness \underline{a} is unique up to \sim . More importantly, even, two values cannot suddenly become equal just because we move to another world. Even though pullback preservation is not formally needed in the proofs of any of the subsequent results it simplifies the intuition considerably.

Lemma 6.5. If A is a p.p.f. and $u : w \rightarrow w'$ and $a, a' \in Aw$ then there is a continuous function $Aw'(u.a, u.a') \rightarrow Aw(a, a')$. (There is also a continuous function in the other direction just by functoriality.)

Note that the ordering on worlds and world morphisms is discrete so that continuity only refers to the $Aw'(u.a, u.a')$ argument.

Definition 6.6 (Morphism of functors). If A, B are p.p.f.s., a morphism from A to B is a natural transformation from A to B . Thus, concretely, a morphism is an equivalence class of pairs $e = (e_0, e_1)$ of continuous functions where $e_0 : \Pi w. Aw \rightarrow Bw$ and

$$\begin{aligned} e_1 : \Pi w. \Pi w'. \Pi x : w \rightarrow w'. \Pi a \in Aw. \Pi a' \in Aw'. \\ Aw'(x.a, a') \rightarrow Bw'(x.e_0(a), e_0(a')). \end{aligned}$$

Two such morphisms e, e' are equal if there exists a continuous function $\mu : \Pi w. \Pi a \in Aw. Bw(e(a), e'(a))$.

These morphisms compose in the obvious way and so the p.p.f.s and morphisms between them form a category.

Theorem 6.7. The category of p.p.f.s is cartesian-closed and has pullbacks. If A, B are p.p.f. then $(A \times B)w = Aw \times Bw$ and $|A \Rightarrow B|w$ contains pairs (f_0, f_1) as follows: $f_0(u) \in |Aw_1 \Rightarrow Bw_1|$ for each w_1 and $u : w \rightarrow w_1$. If $u : w \rightarrow w_1$ and $v : w_1 \rightarrow w_2$ then

$$f_1(u, v) \in (Aw_1 \Rightarrow Bw_2)([Av \Rightarrow Bw_2] f_0(vu), [Aw_1 \Rightarrow Bv] f_0(u))$$

where

$$\begin{aligned} [Av \Rightarrow Bw_2] : (Aw_2 \Rightarrow Bw_2) \rightarrow (Aw_1 \Rightarrow Bw_2) \\ [Aw_1 \Rightarrow Bv] : (Aw_1 \Rightarrow Bw_1) \rightarrow (Aw_1 \Rightarrow Bw_2) \end{aligned}$$

are the obvious composition morphisms.

A proof in $(A \Rightarrow B)w((f_0, f_1), (f'_0, f'_1))$ is a function g that for each $u : w \rightarrow w_1$ yields a proof $g(u) \in (Aw_1 \Rightarrow Bw_1)(f_0(u), f'_0(u))$.

The order on objects and proofs is pointwise as usual.

Moreover, if D is a p.p.f. with the property that Dw has a least element \perp_w for each w and whenever $u : w \rightarrow w_1$ then $u.\perp_w = \perp_{w_1}$ then there is a morphism $Y_D : (D \Rightarrow D) \rightarrow D$ satisfying the same properties as those asserted in Prop. 5.6.

Proof This is a special case of the construction of a functor category with respect to an internal category (here the category

of worlds viewed as discrete setoids, thus internal to the category of setoids). One mimics the standard constructions for set-valued functors taking care to keep representatives of morphisms distinct as we did in the construction of the function space of setoids.

The existence of the fixpoint operator does not seem to be known but follows by applying the same method. The rather fine-grained specification of the fixpoint operator on the level of setoids is needed here so that we can assure naturality of Y_D and $Y_D(f)$. \square

Heaps as a setoid We will now equip heaps with a world-dependent setoid structure singling out the separated ones (for a particular world) and setting equality to be induced by the ambient rely relations.

Definition 6.8. If w is a world then the setoid $\mathfrak{S}w$ is defined by $|\mathfrak{S}w| = \{h \in \mathbb{H} \mid \forall l \in w.h : l^R \wedge \text{sep}(h, \text{dom}(w))\}$ and $\mathfrak{S}w(h, h') = \{\star\} \iff \forall l \in w.h \stackrel{\perp}{\sim} h'$ and $\mathfrak{S}w(h, h') = \emptyset$ otherwise, where $\text{sep}(h, \text{dom}(w))$ is the separated heap with respect to the locations in $\text{dom}(w)$ (see Definition 4.3).

If $u : w \rightarrow w'$ and $h \in \mathfrak{S}w'$, then we define $h.u \in \mathfrak{S}w$ to be $\pi^{-1} \cdot h$ for some π extending u . This is well-defined up to \sim .

In this way, \mathfrak{S} becomes a contravariant functor from the category of worlds to the category of setoids.

Lemma 6.9 (Pasting Lemma). Let $w_u^x \diamond_u^x w'$ be a minimal pullback with apex \bar{w} . Let $h \in \mathfrak{S}w$ and $h' \in \mathfrak{S}w'$ such that $h.u \sim h'.u'$. Then there exists a heap $h_1 \in \mathfrak{S}w$, unique up to \sim , such that $h_1.x \sim h$ and $h_1.x' \sim h'$.

Corollary 6.10. Let $u : w \rightarrow w'$ be a morphism and $h \in \mathfrak{S}w$ and $h' \in \mathfrak{S}w'$. There exists h_1 unique up to \sim such that $h_1.u \sim h$ and $h_1.u' \sim h'.u'$ where $w_0^u \diamond_u^0 w_1$ is a minimal pullback with low point \emptyset , the empty world.

Intuitively, h_1 is obtained by overwriting the w -part of h' according to h .

Heap Relations

Definition 6.11. A relation R on \mathfrak{S} consists of a subset $Rw \subseteq \mathfrak{S}w \times \mathfrak{S}w$ for each w , such that if $(h, h') \in Rw$ and $u : w_0 \rightarrow w$ then $(h.u, h'.u) \in Rw_0$, and if $p : h \sim h_1$ and $p' : h' \sim h'_1$ then $(h_1, h'_1) \in Rw$, as well.

The elementary effects track reading, writing, and allocating at the level of regions and are as follows: wr_r (writing within region r), rd_r (reading from within region r), and al_r (allocating within region r). Each elementary effect ϵ is associated with a set $\mathcal{R}(\epsilon)$ of relations on \mathfrak{S} as follows:

$$\begin{aligned} R \in \mathcal{R}(rd_r) &\iff (h, h') \in Rw \implies \forall l \in w(r). h \stackrel{\perp}{\sim} h' \\ R \in \mathcal{R}(wr_r) &\iff (h, h') \in Rw \implies \forall l \in w(r). \forall h_1, h'_1. h_1 \stackrel{\perp}{\sim} h'_1 \implies \\ &\quad h \stackrel{\perp}{\sim} h_1 \wedge h' \stackrel{\perp}{\sim} h'_1 \implies (h_1, h'_1) \in Rw \\ R \in \mathcal{R}(al_r) &\iff (h, h') \in Rw \implies \\ &\quad \forall w_1. \forall u \in \mathcal{I}(w, w_1). (w_1 \setminus w) \subseteq w_1(r) \implies \\ &\quad \forall h_1, h'_1 \in \mathfrak{S}w_1. [h_1.u \sim h \wedge h'_1.u \sim h' \wedge \\ &\quad (h_1, h'_1) \in \bigcap_{l \in w_1 \setminus w} l^R] \implies (h_1, h'_1) \in Rw_1 \end{aligned}$$

Thus, $\mathcal{R}(rd_r)$ is the set of relations R for which R -related heaps contain “equal” (in the sense of l^R) values for all abstract locations in region r ; a relation $R \in \mathcal{R}(wr_r)$ is oblivious to (preserved by) “equal” writes to any abstract location in r ; and a relation $R \in \mathcal{R}(al_r)$ is oblivious to extensions of the current world provided that they only add abstract locations in region r , that the initial contents of these newly allocated locations are “equal” in the sense of $(-)^R$, and that nothing else is changed.

Finally, $\mathcal{R}(\epsilon)$ for a set of effects ϵ is defined as follows:

$$\mathcal{R}(\epsilon) = \bigcap_{\epsilon \in \epsilon} \mathcal{R}(\epsilon).$$

As in our earlier work [7], the key idea is that computations with effect ε will preserve all the heap relations that are preserved by all the operations allowed by ε . The larger the set of operations, the fewer relations need be preserved.

Definition 6.12 (Monad). *Let A be a p.p.f. and ε an effect. A p.p.f. $T_\varepsilon A$ is defined as follows:*

- (Objects) Elements of $(T_\varepsilon A)\mathbf{W}$ are pairs (c_0, c_1) of partial functions where

$$c_0 : \mathfrak{S}\mathbf{W} \rightarrow \Sigma \mathbf{W}_1. \mathcal{I}(\mathbf{w}, \mathbf{w}_1) \times \mathfrak{S}\mathbf{W}_1 \times A\mathbf{W}_1$$

and c_1 is as follows. If $R \in \mathcal{R}(\varepsilon)$ and $(h, h') \in R\mathbf{w}$ then $c_1(R, h, h')$ either is undefined and $c_0(h)$ and $c_0(h')$ are both undefined or else $c_1(R, h, h')$ is defined and then $c_0(h)$ and $c_0(h')$ are both defined, say $c_0(h) = (\mathbf{w}_1, u, h_1, \mathbf{a})$ and $c_0(h') = (\mathbf{w}'_1, u', h'_1, \mathbf{a}')$. In this case, $c_1(R, h, h')$ returns a pair $(x \diamond_{v'}^{x'}, p)$ where $\mathbf{w}_1 \overset{x}{\diamond}_{v'}^{x'} \mathbf{w}'_1$ such that $xu = x'u'$. Furthermore, $p \in A\overline{\mathbf{w}}(x.\mathbf{a}, x'.\mathbf{a}')$ and, finally, $(h_1.v, h'_1.v') \in R\overline{\mathbf{w}}$ where $\overline{\mathbf{w}}$ and $\underline{\mathbf{w}}$ are low point and apex of $x \diamond_{v'}^{x'}$.

- (Proofs) Proofs only look at the $(-)_0$ components. Let $\mathbf{c} = (c_0, _)$ $\in T_\varepsilon A\mathbf{W}$ and $\mathbf{c}' = (c'_0, _)$ $\in T_\varepsilon A\mathbf{W}$, then a proof in $(T_\varepsilon A)\mathbf{W}(\mathbf{c}, \mathbf{c}')$ is a partial (continuous) function μ , such that for a given $h \in \mathfrak{S}\mathbf{W}$:

1. if $\mu(h)$ is undefined, then $c_0(h)$ and $c'_0(h)$ are both undefined;
2. if $\mu(h)$ is defined, then $c_0(h) = (\mathbf{w}_1, u, h_1, \mathbf{a})$ and $c'_0(h) = (\mathbf{w}'_1, u', h'_1, \mathbf{a}')$ are both defined. In this case, $\mu(h)$ returns a tuple $(x \diamond_{v'}^{x'}, q)$ satisfying $xu = x'u'$, $q \in A\overline{\mathbf{w}}(x.\mathbf{a}, x'.\mathbf{a}')$ and $h_1.v \sim h'_1.v'$ in $\mathfrak{S}\overline{\mathbf{w}}$, with $\overline{\mathbf{w}}$ and $\underline{\mathbf{w}}$ apex and low point of $x \diamond_{v'}^{x'}$.

- (Order) The order between computations is given point-wise: $(c_0, c_1) \leq (c'_0, c'_1)$ if and only if for any $h \in \mathfrak{S}\mathbf{W}$ we have $c_0(h) \leq c'_0(h)$, which is defined as follows: for all $h \in \mathfrak{S}\mathbf{W}$ if $c_0(h)$ is defined so is $c'_0(h)$. Moreover, if both $c_0(h)$ and $c'_0(h)$ are defined as $c_0(h) = (\mathbf{w}_1, u_1, h_1, \mathbf{a}_1)$ and $c'_0(h) = (\mathbf{w}'_1, u'_1, h'_1, \mathbf{a}'_1)$, then $\mathbf{w}_1 = \mathbf{w}'_1$, $u_1 = u'_1$, $h_1 = h'_1$ and $\mathbf{a}_1 \leq \mathbf{a}'_1$. Finally, if $c_1(R, h, h')$ and $c'_1(R, h, h')$ are defined returning $(x_1 \diamond_{v_1}^{x'_1}, q)$ and $(x_2 \diamond_{v_2}^{x'_2}, q')$, then $x_1 \diamond_{v_1}^{x'_1} = x_2 \diamond_{v_2}^{x'_2}$ and $q \leq q'$.¹

- (Morphism part) Let $x : \mathbf{w} \rightarrow \mathbf{w}'$ be a morphism and $\mathbf{c} = (c_0, c_1) \in T_\varepsilon A\mathbf{W}$. We are going to describe the first component of $(c'_0, _)$ $\in T_\varepsilon A\mathbf{W}(\mathbf{c})$. Given $h \in \mathfrak{S}\mathbf{w}'$ write $c_0(h.x) = (\mathbf{w}_1, u, h_1, \mathbf{a})$. Should this application be undefined so is $c'_0(h)$. Now, using Proposition 6.3, complete u, x to a minimal pullback square $\overset{x}{\underset{u}{\diamond}}_{\mathbf{w}}^{\mathbf{w}'}$ with $u' : \mathbf{w}' \rightarrow \mathbf{w}'_1$ where \mathbf{w}'_1 is the apex of this pullback.

Now, obtain $h' \in \mathfrak{S}\mathbf{w}'$ from h by overwriting with $h_1.u$ using Corollary 6.10. Then paste h_1 and h' (note that $h_1.u \sim h'.u$ by construction) using Lemma 6.9 to yield the desired heap $h'_1 \in \mathfrak{S}\mathbf{w}'_1$. We then put $c'_0(h) = (\mathbf{w}'_1, u', h'_1, \mathbf{a})$.

We omit the remaining components and verifications.

The proof of the following is tedious but straightforward and will be given in an extended version of this paper.

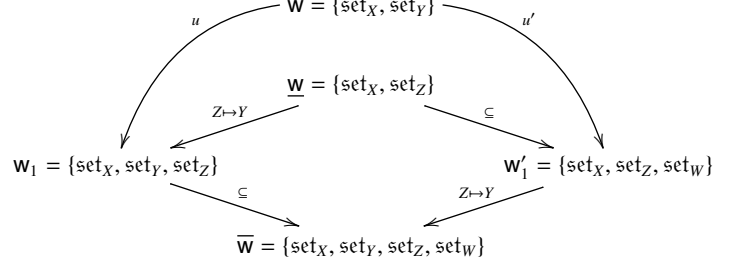
Theorem 6.13. T_ε is the functor part of an effect-indexed strong monad [8, 43].

This theorem means in particular, that each T_ε is a monad on the category of p.p.f. i.e., for each A there are natural transformation $\eta_A : A \rightarrow T_\varepsilon A$ and $\mu_A : T_\varepsilon T_\varepsilon A \rightarrow T_\varepsilon A$ subject to the usual laws and whenever $\varepsilon_1 \subseteq \varepsilon_2$ then there is a morphism $T_{\varepsilon_1} A \rightarrow T_{\varepsilon_2} A$ which interacts with the monad data in the expected sense.

Illustration Assume an initial heap h , where $h(X) \mapsto [1, 2, 3]$ and $h(Y) \mapsto [3, 2, 1]$, that is, h points to the lists $[1, 2, 3]$ and $[3, 2, 1]$ from concrete locations X and Y , respectively. Moreover, let $\mathbf{c}(h) = (\mathbf{w}_1, u, h_1, (3, Y))$ and $\mathbf{c}'(h) = (\mathbf{w}'_1, u', h'_1, (3, Z))$, where

¹ Strictly speaking, $q \leq q'$ needs to be a tuple with the objects they are proving the equality of.

$h_1(X) \mapsto [1, 2, 3, 4]$; $h_1(Y) \mapsto [3, 2, 1, 0]$; and $h_1(Z) = [4, 2]$, while $h'_1(X) \mapsto [4, 3, 2, 1]$; $h'_1(Z) \mapsto [3, 2, 1, 3, 2, 1, 0]$; and $h'_1(W) = [0, 0]$. We can show that \mathbf{c} and \mathbf{c}' are equal when given h . The proof is illustrated by the diagram below, where an arrow labeled with \subseteq denotes an inclusion and an arrow labeled with a mapping, e.g., $Z \mapsto Y$, denotes the corresponding renaming.



Here, it is easy to check that h_1 and h'_1 are equal when taken to the world $\underline{\mathbf{w}}$, as the concrete locations X in h_1 and h'_1 and Y in h_1 and Z in h'_1 point to lists containing the same set of elements: $\{1, 2, 3, 4\}$ and $\{0, 1, 2, 3\}$, respectively. Notice that the abstract locations $\text{set}_Z \in \mathbf{w}_1$ and $\text{set}_W \in \mathbf{w}'_1$ are not taken into account. Moreover, the value $(3, Y)$ is also equivalent to $(3, Z)$ when taken to the world $\overline{\mathbf{w}}$, as the morphism $\mathbf{w}'_1 \rightarrow \overline{\mathbf{w}}$ renames Z to Y , while the morphism $\mathbf{w}_1 \rightarrow \overline{\mathbf{w}}$ is an inclusion.

We see here that the use of setoids rather than plain sets or predomains is crucial. The proofs in $T_\varepsilon A$ are clearly relevant as they explain in what way the abstract locations correspond to each other. Likewise, the elements of $[T_\varepsilon A]$ contain more information than the mere computation, i.e., the evolution of the heap. Namely, the world extension \mathbf{w}_1 will typically contain new abstract locations with new contracts that can be arbitrarily complicated. Thus, it would not be possible to replace $T_\varepsilon A$ with a simple logical predicate on actual computations.

7. Proof-relevant Logical Relations

The structure identified so far allows us to interpret our effect type system and the equational theory in the category of p.p.f.

Definition 7.1 (type interpretation). *A type interpretation consists of an assignment of a p.p.f. $\llbracket A \rrbracket$ to each basic type A .*

Given a type interpretation we can then assign a p.p.f. $\llbracket \tau \rrbracket$ to each type the essential clause being $\llbracket \tau_1 \xrightarrow{\varepsilon} \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \Rightarrow T_\varepsilon \llbracket \tau_2 \rrbracket$. We also interpret a typing context Γ as the cartesian product over its bindings: $\llbracket \Gamma \rrbracket = \prod_{x \in \text{dom}(\Gamma)} \llbracket \Gamma(x) \rrbracket$.

Lemma 7.2 (masking). *Let r be a region. For world \mathbf{w} define $\mathbf{w}-r$ by removing all locations in r . If type τ does not contain r then for each world \mathbf{w} the morphism $u.(-) : \llbracket \tau \rrbracket \mathbf{w}-r \rightarrow \llbracket \tau \rrbracket \mathbf{w}$ where $u : \mathbf{w}-r \rightarrow \mathbf{w}$ is the obvious inclusion is an isomorphism.*

Definition 7.3 (axiom interpretation). *An axiom interpretation consists of an element $x_{(v,\tau)} \in \llbracket \tau \rrbracket \emptyset$ (empty world) for each type axiom (v, τ) and of an element $p_{(v,v',\tau)} \in \llbracket \tau \rrbracket \emptyset(x_{(v,\tau)}, x_{(v',\tau)})$ for each equality axiom (v, v', τ) .*

Given an axiom interpretation (and a type interpretation) and a well-typed term $\Gamma \vdash e : \tau \ \& \ \varepsilon$ we define a morphism

$$\llbracket \Gamma \vdash e : \tau \ \& \ \varepsilon \rrbracket : \llbracket \Gamma \rrbracket \rightarrow T_\varepsilon \llbracket \tau \rrbracket$$

using the fact that all the typing rules are reflected in the semantics: cartesian closure takes care of application and abstraction and variable management. The fixpoint morphism Y_D takes care of recursive definitions; the indexed monadicity of T_ε takes care of the rules associated with “let” and effect propagation. Lemma 7.2, fi-

nally, takes care of the masking rule. If $\Gamma \vdash v : \tau$ because of the axiom (v, τ) then we put $\llbracket \Gamma \vdash v : \tau \rrbracket(\gamma) = x_{(v, \tau)}$ (constant morphism).


Effects Given a set of effects ε , we write $\text{rds}(\varepsilon) = \{r \mid rd_r \in \varepsilon\}$, $\text{wrs}(\varepsilon) = \{r \mid wr_r \in \varepsilon\}$, $\text{als}(\varepsilon) = \{r \mid al_r \in \varepsilon\}$ and $\text{regs}(\varepsilon) = \text{rds}(\varepsilon) \cup \text{wrs}(\varepsilon) \cup \text{als}(\varepsilon)$. Moreover, the set $\text{nwrs}(\varepsilon) = \text{regs}(\varepsilon) \setminus \text{wrs}(\varepsilon)$. We also introduce the following piece of notation for $h, h' \in \mathfrak{S}\mathbf{w}$:

$$h \sim_{\text{rds}(\varepsilon, \mathbf{w})} h' \iff \forall l \in \mathbf{w}(\text{rds}(\varepsilon)). h \stackrel{l}{\sim} h'$$

$$h \sim_{\text{nwrs}(\varepsilon, \mathbf{w})} h' \iff \forall l \in \mathbf{w}(\text{nwrs}(\varepsilon)). h \stackrel{l}{\sim} h'$$

which specifies that the heaps h and h' are equivalent on all the abstract locations l in regions associated to read effects in ε and on the not write locations, respectively.

Lemma 7.4. *Let $\Gamma \vdash e : \tau \ \& \ \varepsilon$. Let $\mathbf{w}_v^x \diamond_{v'}^x \mathbf{w}'$ be a pullback square, $\gamma \in \llbracket \Gamma \rrbracket \mathbf{w}, \gamma' \in \llbracket \Gamma \rrbracket \mathbf{w}'$ be contexts, such that $\mu_\gamma : x.\gamma \sim x.\gamma'$ and $\mathbf{c} = \llbracket \Gamma \vdash e : \tau \ \& \ \varepsilon \rrbracket \mathbf{w}(\gamma) \in T_\varepsilon \mathbf{A}\mathbf{w}$ and $\mathbf{c}' = \llbracket \Gamma \vdash e : \tau \ \& \ \varepsilon \rrbracket \mathbf{w}'(\gamma') \in T_\varepsilon \mathbf{A}\mathbf{w}'$ be computations, such that $\mu : x.\mathbf{c} \sim x'.\mathbf{c}'$. Let $h \in \mathfrak{S}\mathbf{w}, h' \in \mathfrak{S}\mathbf{w}'$ be heaps, such that $h.v \sim_{\text{rds}(\varepsilon, \mathbf{w})} h'.v'$. Then $\mathbf{c}(h)$ and $\mathbf{c}'(h')$ co-terminate. Moreover, if they do terminate and $\mathbf{c}(h) = (\mathbf{w}_1, u_1, h_1, \mathbf{a}_1)$ and $\mathbf{c}'(h') = (\mathbf{w}'_1, u'_1, h'_1, \mathbf{a}'_1)$, then the following diagram exists, in particular, the dashed arrows, where the following is satisfied*



1. $x_1.a \sim x'_1.a'$, $\mathbf{w} \cong \mathbf{w} \otimes \mathbf{q}$, $\mathbf{w}' \cong \mathbf{w} \otimes \mathbf{q}'$, $\mathbf{w}_1 \cong \mathbf{w} \otimes \mathbf{q} \otimes \mathbf{q}_1$, $\mathbf{w}'_1 \cong \mathbf{w} \otimes \mathbf{q}' \otimes \mathbf{q}'_1$, for some \mathbf{q}_1 and \mathbf{q}'_1 ;
2. for all $l \in \mathbf{w}$, we have either: $h.v \stackrel{l}{\sim} h_1.u_1.v$ and $h'.v' \stackrel{l}{\sim} h'_1.u'_1.v'$ (remain equivalent) or $h_1.u_1.v_1 \stackrel{l}{\sim} h'_1.u'_1.v'_1$ (equally modified);
3. if $l \in \mathbf{w}(\text{nwrs}(\varepsilon))$, then $h.v \stackrel{l}{\sim} h_1.u_1.v$ and $h'.v' \stackrel{l}{\sim} h'_1.u'_1.v'$;
4. if $\text{als}(\varepsilon) = \emptyset$, then there exists a morphism $\mathbf{c}_* \in \llbracket \Gamma \rrbracket \rightarrow T_\varepsilon \llbracket \tau \rrbracket$, such that $\mathbf{c}_* \sim \mathbf{c}$ and if $\mathbf{c}_*(\mathbf{w})(\gamma)h = (\mathbf{w}_*, u_*, h_*, \mathbf{a}_*)$ then $\mathbf{w}_* \cong \mathbf{w}$ (no world extension is needed);
5. $\mathbf{w}\mathbf{w}_1 \mathbf{w}_1$ and $\mathbf{w}\mathbf{w}'_1 \mathbf{w}'_1$ are pullbacks.

Proof The proof that the values are equal in \mathbf{w}_1 follows directly from the definition of computations and effects.

For the item 2, we use the following relation R defined as follows for all worlds \mathbf{w}_1 , such that $u : \mathbf{w} \rightarrow \mathbf{w}_1$:

$$\{(h_1, h'_1) \mid h_1.u \sim_{\text{rds}(\varepsilon, \mathbf{w})} h'_1.u \wedge \forall l \in \mathbf{w}. (h_1.u \stackrel{l}{\sim} h \wedge h'_1.u \stackrel{l}{\sim} h') \vee h_1.u \stackrel{l}{\sim} h'_1.u\}$$

Otherwise, for the worlds \mathbf{w}_2 not reachable from \mathbf{w} , the relation $R\mathbf{w}_2$ is the trivial set. Notice that $R \in \mathcal{R}(\varepsilon)$ and it is contravariant. The claim then follows directly by using the morphism from $\mathbf{w} \rightarrow \mathbf{w}_1$.

The proof of item 3 follows in a similar fashion, but we use the following relation:

$$\{(h_1, h'_1) \mid h_1.u \sim_{\text{rds}(\varepsilon, \mathbf{w})} h'_1.u \wedge h_1.u \sim_{\text{nwrs}(\varepsilon, \mathbf{w})} h'_1.u\}$$

And we use a similar relation for showing that h' and $h'_1.u'$ agree on the not written locations $\mathbf{w}(\text{nwrs}(\varepsilon))$.

For item 4, first, we show that $\mathbf{w}(r)$ and $\mathbf{w}'(r)$ are equal up to renaming for all regions $r \notin \text{als}(r)$ by using the following relation $R\mathbf{q}$ for some world \mathbf{q} :

$$\{(h_1, h'_1) \mid h_1 \sim h'_1 \wedge \forall r \notin \text{als}(\varepsilon). \#_r(\mathbf{q}) \leq \#_r(\mathbf{w})\}$$

where $\#_r$ denotes the number of abstract locations colored with r . Clearly, $R \in \mathcal{R}(\varepsilon)$ as ε does not contain any allocation effects. Thus, the number of locations in \mathbf{w}_* is the same as the locations in \mathbf{w} , by using the pullback $\mathbf{w}_1 \diamond_{v'}^x \mathbf{w}_*$. This gives us one direction, while the other direction is obtained by the universal property of pullbacks. Given this property, one can easily construct the function \mathbf{c}_* . \square

Theorem 7.5 (Equational soundness). *Assume an ambient interpretation of types and axioms. If $\Gamma \vdash e = e' : \tau \ \& \ \varepsilon$ then the morphisms $\llbracket \Gamma \vdash e : \tau \ \& \ \varepsilon \rrbracket$ and $\llbracket \Gamma \vdash e' : \tau \ \& \ \varepsilon \rrbracket$ are equal morphisms.*

Proof By induction on derivations. The core rules are direct from the category-theoretic structure identified so far. The effect-specific rules are nontrivial; however given Lemma 7.4, their proofs follow essentially the same reasoning as in existing literature [6, 41]. As for the axiom rule we argue as follows.

If $\Gamma \vdash v = v' : \tau$ using the axiom (v, v', τ) then we have $\llbracket \Gamma \vdash v : \tau \rrbracket(\gamma) = x_{(v, \tau)}$ and $\llbracket \Gamma \vdash v' : \tau \rrbracket(\gamma) = x_{(v', \tau)}$ and the constant function $\lambda\gamma.p_{(v, v', \tau)}$ is continuous and thus establishes the desired equality. \square

8. Observational equivalence

We will now relate semantic equality to observational equivalence. So far, a trivial model that identifies everything could also have been used to “justify” our rules. Doing this, however, requires a loose relationship between the setoid interpretation and the computational meanings of raw terms which will be given by a realizability relation.

Definition 8.1. *For each type τ , effect ε , and world \mathbf{w} we define admissible relations $\Vdash_{\mathbf{w}}^{\tau, \varepsilon} \subseteq \mathbb{V} \times \llbracket \tau \rrbracket \mathbf{w}$ and $\Vdash_{\mathbf{w}}^{\tau, \varepsilon, \mathbf{c}} \subseteq \mathbb{C} \times |T_\varepsilon \llbracket \tau \rrbracket \mathbf{w}|$ as follows:*

- If $c \in \mathbb{C}$ then $c \Vdash_{\mathbf{w}}^{\tau, \varepsilon, \mathbf{c}} (c_0, c_1)$ means that for each $h \in \mathfrak{S}\mathbf{w}$ one has $c(h)$ defined iff $c_0(h)$ defined and if $c(h) = (h_1, a)$ then $c_0(h) = (\mathbf{w}_1, u, h_1, \mathbf{a})$ (same h_1) and $h_1 \in \mathfrak{S}\mathbf{w}_1$ and $a \Vdash_{\mathbf{w}_1}^{\tau} \mathbf{a}$.
- If $\text{fun}(f) \in \mathbb{V}$ and $(f_0, f_1) \in \llbracket \tau_1 \xrightarrow{\varepsilon} \tau_2 \rrbracket \mathbf{w}$ then $f \Vdash_{\mathbf{w}}^{\tau_1 \xrightarrow{\varepsilon} \tau_2} (f_0, f_1)$ iff whenever $u : \mathbf{w} \rightarrow \mathbf{w}_1$ is an inclusion and $a \Vdash_{\mathbf{w}_1}^{\tau_1} x$ then $f(a) \Vdash_{\mathbf{w}_1}^{\tau_2, \varepsilon} f_0(x)$.
- $\eta \Vdash_{\mathbf{w}}^{\Gamma} \gamma \iff \eta(x) \Vdash_{\mathbf{w}}^{\Gamma(x)} \gamma(x)$ for all $x \in \text{dom}(\Gamma)$.

Lemma 8.2. *The \Vdash relations are indeed admissible and are stable under inclusion; that is whenever $v \Vdash_{\mathbf{w}}^X x$ and $u : \mathbf{w} \rightarrow \mathbf{w}'$ is an inclusion then $v \Vdash_{\mathbf{w}_1}^X u.x$ holds.*

The following soundness property of \Vdash is proved by straightforward induction on typing derivations.

Theorem 8.3 (Fundamental Lemma). *Assume that the interpretation of types and axioms is chosen such that for all type axioms (v, τ) one has $v \Vdash_{\emptyset}^{\tau} x_{(v, \tau)}$. Whenever $\Gamma \vdash e : \tau \ \& \ \varepsilon$ and $\eta \Vdash_{\mathbf{w}}^{\Gamma} \gamma$ then $\llbracket e \rrbracket \eta \Vdash_{\mathbf{w}}^{\tau, \varepsilon} \llbracket \Gamma \vdash e : \tau \rrbracket \mathbf{w}(\gamma)$.*

Definition 8.4 (Observational equivalence). *Two value expressions v, v' where $\vdash v : \tau$ and $\vdash v' : \tau$ are observationally equivalent at type τ if for all f such that $\vdash f : \tau \xrightarrow{\varepsilon} \text{unit}$ (“observations”) it is the case that $\llbracket f v \rrbracket \emptyset$ is defined iff $\llbracket f v' \rrbracket \emptyset$ is defined. (\emptyset stands for empty heap and environment).*

It is easy to extend this definition to open terms and computation expressions by closing up with lambda abstractions. Intuitively, observationally equivalent expressions, can be replaced by one another in any context without compromising observable behavior. To see this, one can build observations f from such a context by first λ -abstracting its hole and then adding at the end a wrapper function that tests for some observable feature of the end result signaling failure of the test by nontermination.

Theorem 8.5 (main result). *If $\vdash v = v' : \tau$ then v and v' are observationally equivalent.*

Proof If f is an observation then by the congruence rules we get $\vdash f v = f v' : \tau$ so, by Theorem 7.5 we find $\llbracket \vdash f v \rrbracket = \llbracket \vdash f v' \rrbracket$, thus in particular $\llbracket \vdash f v \rrbracket \emptyset$ (empty world, empty heap) is defined iff $\llbracket \vdash f v' \rrbracket \emptyset$ is defined. Now, if $\llbracket \vdash f v \rrbracket \emptyset$ is defined then by Theorem 8.3 and the definition of \Vdash we can conclude that $\llbracket \vdash f v \rrbracket \emptyset$

is defined so, using Theorem 8.3 again we get that $\llbracket f v' \rrbracket 0$ is defined. \square

9. Applications

In this section we give some concrete instantiations of our framework in the form of axioms and their semantic justifications. We do so in more detail for the Overlapping References example shown in Section 2, and in less detail for the other examples.

Overlapping References Recall value v_{or} and type τ_{or} from Section 2 (two overlapping references). In order to justify the axiom (v_{or}, τ_{or}) we have to construct a semantic object $x_{or} \in \llbracket \tau_{or} \rrbracket$ such that $v_{or} \Vdash_0^{\tau_{or}} x_{or}$ which we will now do.

Given world w we put $x_{or}(w)(\star) = (c_0, c_1)$, where c_0 and c_1 are defined as follows $(\star \in \llbracket \text{unit} \rrbracket w)$:

$$c_0(h) = (u : w \rightarrow w_1, h_1, (g_1, g_2, s_1, s_2))$$

where $new(h, 0) = (X, h_1)$, that is, h_1 is obtained from h by allocating a concrete location X ; w_1 is the world obtained by extending w with the abstract locations fst_X and snd_X defined in Section 4.1, where the former is marked with region $r1$ and the latter with region $r2$; finally, (g_1, g_2, s_1, s_2) are the corresponding semantic objects that get and set values to the projections of pairs stored in X , e.g., $g_1 \in \llbracket (\text{unit} \xrightarrow{rd_1} \text{int}) \rrbracket w_1$ is the semantic counter part of the function getting the first projection of the pair stored in X , defined in a similar way as we do for x_{or} .

Note that $v_{or} \Vdash_0^{\tau_{or}} x_{or}$ in particular requires that the new heap returned by c_0 agrees with the one returned by v_{or} and, more generally, we are forced to follow the computational behavior of v_{or} in x_{or} . We are free, though to choose the semantic components such as world extensions and, later, pullbacks.

The object c_1 , now, is defined as follows: Assume that $R \in \mathcal{R}(\{al_{r1}, al_{r2}\})$ is a heap relation and two heaps $h, h' \in \mathcal{S}w$ that are related, i.e., $(h, h') \in R$. Moreover, assume that $c_0(h) = (u : w \rightarrow w \otimes \text{fst}_X \otimes \text{snd}_X, h_1, (g_1, g_2, s_1, s_2))$ and $c_0(h') = (u : w \rightarrow w \otimes \text{fst}_{X'} \otimes \text{snd}_{X'}, h'_1, (g'_1, g'_2, s'_1, s'_2))$, where the former allocates X , while the latter allocates X' . Note that everything terminates in this example. Now $c_1(R, h, h')$ returns the pullback $w_1 \overset{X \rightarrow X'}{\diamond} \overset{1}{\diamond} \overset{1}{X \rightarrow X'} w'_1$, where X is mapped to X' . It is easy to check that h_1 and h'_1 are related in R when taken to the pullback's low point and that (g_1, g_2, s_1, s_2) and (g'_1, g'_2, s'_1, s'_2) are equal in its apex.

Note that once this typing has been justified all equations derivable from it are automatically sound for observational equivalence. Thus, in particular, writes to the two ‘‘overlapping reference’’ commute.

Set factory Let w be a world and $h \in \mathcal{S}w$. Suppose that h_1 arises from h by allocating a fresh set data structure, e.g., a linked list, with entry point X . Let set_X be the abstract location describing this fresh data structure.

Now for any chosen region r we can add set_X to r to yield a new world w_1 . The function $\text{setfactory}_0(w)(h)$ then returns w_1 and a tuple of semantic functions for reading, membership, removal of which we only sketch reading here: If $u : w_1 \rightarrow w_2$ and $h_1 \in \mathcal{S}w_1$ and $i \in \mathbb{Z}$ then the reading function looks up i in the data structure starting at the entry points X in h_1 (note that $h_1 \in \mathcal{S}w$ asserts that this data structure exists and is well-formed). The returned (abstract) store h_2 might not be the same as h because internal reorganizations, e.g., removal of duplicates, might have occurred, but no world extension is needed and $h_1 \sim h_2$ holds. This together with the fact that the outcome only depends on the set_X^R equivalence class (in the proof-relevant sense) justifies a read-only typing for reading.

Memoization For the simple *memo* functional from Section 2 we produce just as in the previous example a fresh abstract location l

whose specification looks only at the two newly allocated concrete locations, say l_x, l_y , where the first stores the argument value i and the second the integer value $f(i)$ where f is the pure function to be memoised. The location's l rely specifies that two heaps are equivalent when the values stored in l_x and l_y are the same: $(h, h') \in l^R \iff h(l_x), h'(l_x) \wedge h(l_y), h'(l_y)$. We see in Lemma 7.4 that if a function is semantically pure (empty effect) then there is a world- and heap-independent function describing its action. Thus the *memo* functional may be considered pure.

9.1 State Dependent Abstract Data Types

Although our focus is on effect-dependent equivalences that can be derived from refined typings using our equational rules, we are able to justify some of the well-known tricky examples from the literature on proving equivalences in ML-like languages without effect annotations. Our guarantee conditions here act like the transition system components of the worlds used by Ahmed *et al.* [1, 19].

Awkward Example The first example is the classic *awkward example* [34]. Consider the following two programs:

$$e_1 = \text{let } x \leftarrow \text{ref}(0) \text{ in } \lambda f. x := 1; f(); !x \quad \text{and} \quad e_2 = \lambda f. f(); 1.$$

Intuitively, e_1 and e_2 are equivalent as they both return the value 1, although e_1 uses a fresh location to do so. We can formally prove the equivalence as follows: Assign the region where x is allocated as r . If f has the type $\text{unit} \xrightarrow{\varepsilon} \text{unit}$ with effects ε , then e_1 has type $(\text{unit} \xrightarrow{\varepsilon} \text{unit}) \xrightarrow{rd_r, wr_r} \text{int} \ \& \ \varepsilon, al_r$, while e_2 has type $(\text{unit} \xrightarrow{\varepsilon} \text{unit}) \rightarrow \text{int} \ \& \ \varepsilon$. Notice that ε may contain rd_r or wr_r or both. Moreover, assume that an abstract location in region r looks at a single concrete location l :

$$l^R = \{(h, h') \mid h(l) = h'(l)\}$$

$$l^G = \{(h, h_1) \mid h_1(l) = 1 \wedge \forall l' \neq l. h(l') = h_1(l')\}$$

Notice that it is correct to assign e_1 a write effect in region r , as it writes the value 1 to the location assigned to x .

For proving the equivalence of e_1 and e_2 , assume a world w and a heap h . Let $\llbracket e_1 \rrbracket wh = (w \uplus w_1 \uplus w_r, u_1, h_1, a_1)$ and $\llbracket e_2 \rrbracket wh = (w \uplus w_1, u_1, h_2, a_2)$. We construct a pullback square $w \uplus w_1 \uplus w_r \diamond w \uplus w_1$ such that the values a_1 and a_2 are equal in its apex and h_1 and h_2 are equal in its low point. Since wr_r is in the effects of e_1 , we have that $a_1 = 1$. We also have $a_2 = 1$ trivially. Hence a_1 and a_2 are equal in the apex of the pullback square $w \uplus w_1 \uplus w_r \diamond w \uplus w_1$. Similarly, h_1 when taken to the low point of the square, that is, where the locations in w_r are forgotten, the resulting heap is equivalent to h_2 .

Modified Awkward Example Consider now the following variant [1, 19] of the awkward example:

$$e_1 = \text{let } x \leftarrow \text{ref}(0) \text{ in } \lambda f. x := 0; f(); x := 1; f(); !x$$

$$e_2 = \lambda f. f(); f(); 1.$$

The difference is that in the first program x is first assigned 0 and the call-back function is used twice. Interestingly, however, our solution for the Awkward example still works just fine. We can prove semantically that the type of the program e_1 has the same type as before in the Awkward example, where the only writes allowed on the abstract location assigned for x is to write the concrete location to one. Then the reasoning follows in a similar way.

Callback with Lock Example We prove the equivalence of the following programs [1, 19]:

$$e_1 = \text{let } b \leftarrow \text{ref}(\text{true}) \text{ in } \text{let } x \leftarrow \text{ref}(0) \text{ in}$$

$$\langle \lambda f. \text{if } !b \text{ then } (b := \text{false};$$

$$f(); x := !x + 1; b := \text{true}) \text{ else } (), \lambda _ . !x \rangle$$

$$e_2 = \text{let } b \leftarrow \text{ref}(\text{true}) \text{ in } \text{let } x \leftarrow \text{ref}(0) \text{ in}$$

$$\langle \lambda f. \text{if } !b \text{ then } (b := \text{false};$$

$$\text{let } n \leftarrow !x \text{ in } f(); x := n + 1; b := \text{true}) \text{ else } (), \lambda _ . !x \rangle.$$

Both programs produce a pair of functions, one that increments the value stored in x and the second that returns the value stored in x . They use the boolean reference b as lock in the incrementing

function. In particular, once this function is called the value in b is set to `false` and only after the call-back is called and the value in x is incremented is b set again to `true`. However, the implementation of the increment function is different. While e_1 invokes the call-back $f()$ and then increments the value in x , e_2 remembers (in n) the value of x before the call-back and then uses n to update the value stored in x .

Assume that x and b are allocated in the same location (l) in the region r . We show that these programs are equivalent in the type

$((\text{unit} \xrightarrow{\varepsilon} \text{unit}) \xrightarrow{\varepsilon, wr_r, rd_r} \text{unit}) \times (\text{unit} \xrightarrow{rd_r} \text{unit}) \ \& \ al_r$, where ε may contain the effects wr_r, rd_r . In particular, the location l looks at two concrete locations l_b and l_x (storing x and b):

$$\begin{aligned} \mathcal{I}^R &= \{(h, h') \mid h(l_b) = h'(l_b) \wedge h(l_x) = h'(l_x)\} \\ \mathcal{I}^G &= \{(h, h_1) \mid h(l_b) = \text{false} \Rightarrow h_1 = h \text{ and} \\ &\quad (h(l_b) = \text{true} \wedge h_1(l_x) = i \wedge h_1(l_x) = j) \Rightarrow \\ &\quad [(h_1(l_b) = \text{true}) \wedge (j \leq i) \wedge (\forall l' \notin \{l_x, l_b\}. h(l') = h_1(l'))]\} \end{aligned}$$

First, notice that indeed the two programs above have the effect wr_r . The increment of x is allowed by \mathcal{I}^G , as b is initially `true`. To show that the two programs above are equivalent, we show that the value stored in x before and after the call back is called remains the same. This is the case, as even if the callback function has a write effect in the region r , *i.e.*, $wr_r \in \varepsilon$, it cannot change the value of x . This is because when the function is called, the value stored in b is `false`, meaning that the heap remains the same.

Despite the above, the model as presented here does not directly justify equivalences involving representation independence, *e.g.*, that our overlapping references v_{or} are observationally equivalent to an implementation that really allocates two integer references. Generalizing world morphisms as discussed after Def. 6.2 does allow many such equivalences to be proved, though a parametric variant, along the lines of that described by Stark [39], would be more powerful still.

10. Related and Future Work

We have shown Proof-Relevant Logical Relations, introduced in our previous work [4], can be used to justify nontrivial effect-dependent program equivalences. For the first time it was possible to combine effect-dependent program equivalences with hidden invariants allowing “silent modifications” that do not count towards the ascription of an effect. Earlier accounts of effect-dependent program equivalences [3, 6, 7, 26, 41] do not provide such possibilities.

Bisimulation An alternative approach to proving ‘difficult’ contextual equivalences is to use techniques based on bisimulation [27]. While bisimulation has been applied to typed calculi [24, 40] the strength of the method lies in being able to deal with fancy computation rules and fairly simple types and contracts rather than well-understood computation rules (call-by-value lambda calculus in our case) and fancy types and contracts. Indeed, we believe that something like our abstract locations and the proof-relevant world extensions would also suggest itself in a bisimulation-based approach to the equivalences studied here.

Model Variables Proof-relevant logical relations or rather the sets $[AW]$ where A is a semantic type bear a vague relationship with the *model variables* [15] from “design by contract” [32] and more generally data refinement [16]. The commonality is that we track the semantic behavior of a program part with abstract functions on some abstracted set of data that may contain additional information (the “model”). The difference is that we do not focus on particular proof methods or specification formalisms but that we provide a general, sound semantic model for observational equivalence and program transformation and not merely for functional correctness. This is possible by the additional, also proof-relevant part of the semantic equality proofs between the elements of the models. We

also note that our account rigorously supports higher-order functions, recursion, and dynamic allocation.

Separation Logic Our abstract locations draw upon several ideas from separation logic [37], in particular the conditions on rely/guarantee assumptions from [42]. Intriguingly, we did not need something resembling the “frame rule” although perhaps the Π -quantification over larger worlds in function spaces plays its role.

Our use of pullback-preserving functors is motivated by FM-sets [21] or rather the *Schanuel topos* to which they are equivalent. Pitts’s [35] gives a comprehensive account. The way in which we work with both permutation actions and functors has some precedent [9], but does feel slightly awkward: it really ought to be possible to combine the two structures into one.

We would like to define a stylized format that allows one to discharge semantic proof obligations in some cases without having to go down to the low-level semantic definitions. We noted that *private transitions* [1] serve a similar purpose and might perhaps suggest a possible approach in our case, too.

We would also like to be able to store values with proof-relevant equality, which would allow a stratified form of higher-order store [12]. Unrestricted higher-order store requires circular definitions of worlds, which has previously been addressed using, for example, metric spaces [11].

Acknowledgments We thank Lennart Beringer and Andrew Kennedy for fruitful discussions. Nigam was supported by CAPES / CNPq and DAAD.

References

- [1] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *POPL*, 2009.
- [2] G. Barthe, V. Capretta, and O. Pons. Setoids in type theory. *J. Funct. Program.*, 13(2):261–293, 2003.
- [3] N. Benton, L. Beringer, M. Hofmann, and A. Kennedy. Relational semantics for effect-based program transformations: higher-order store. In *PPDP*, 2009.
- [4] N. Benton, M. Hofmann, and V. Nigam. Proof-relevant logical relations for name generation. In *TLCA*, 2013.
- [5] N. Benton and C.-K. Hur. Biorthogonality, step-indexing and compiler correctness. In *ICFP*, 2009.
- [6] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations with dynamic allocation. In *PPDP*, 2007.
- [7] N. Benton, A. Kennedy, M. Hofmann, and L. Beringer. Reading, writing and relations. In *APLAS*, volume 4279 of *LNCS*, 2006.
- [8] N. Benton, A. Kennedy, and G. Russell. Compiling Standard ML to Java bytecodes. In *ICFP*, 1998.
- [9] N. Benton and B. Leperchey. Relational reasoning in a nominal semantics for storage. In *TLCA*, volume 3461 of *LNCS*, 2005.
- [10] L. Birkedal, A. Carboni, G. Rosolini, and D. S. Scott. Type theory via exact categories. In *LICS*, 1998.
- [11] L. Birkedal, K. Stovring, and J. Thamsborg. The category-theoretic solution of recursive metric-space equations. *Theor. Comp. Sci.*, 411:4102–4122, 2010.
- [12] G. Boudol. Typing termination in a higher-order concurrent imperative language. *Inf. Comput.*, 208(6), 2010.
- [13] C. Calcagno, P. W. O’Hearn, and H. Yang. Local action and abstract separation logic. In *LICS*, pages 366–378, 2007.
- [14] A. Carboni, P. J. Freyd, and A. Scedrov. A categorical approach to realizability and polymorphic types. In *MFPS, LNCS 298*, 1987.
- [15] Y. Cheon, G. T. Leavens, M. Sitaraman, and S. H. Edwards. Model variables: cleanly supporting abstraction in design by contract. *Softw. Pract. Exper.*, 35(6):583–599, 2005.

- [16] W. P. de Roeper and K. Engelhardt. *Data Refinement: Model-oriented Proof Theories and their Comparison*. Cambridge University Press, 1998.
- [17] T. Dinsdale-Young, P. Gardner, and M. J. Wheelhouse. Abstraction and refinement for local reasoning. In *VSTTE*, volume 6217 of *LNCS*, 2010.
- [18] D. Dreyer, A. Ahmed, and L. Birkedal. Logical step-indexed logical relations. *Logical Methods in Computer Science*, 7(2), 2011.
- [19] D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. In *Proc. ICFP, ACM*, pages 143–156, 2010.
- [20] J. N. Foster, M. B. Greenwald, J. T. Moore, B. C. Pierce, and A. Schmitt. Combinators for bi-directional tree transformations: a linguistic approach to the view update problem. In *POPL*, pages 233–246, 2005.
- [21] M. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding. *Formal Asp. Comput.*, 13(3-5):341–363, 2002.
- [22] D. K. Gifford and J. M. Lucassen. Integrating functional and imperative programming. In *LISP and Functional Programming*, 1986.
- [23] C.-K. Hur and D. Dreyer. A Kripke logical relation between ML and assembly. In *POPL*, 2011.
- [24] C.-K. Hur, D. Dreyer, G. Neis, and V. Vafeiadis. The marriage of bisimulations and kripke logical relations. In J. Field and M. Hicks, editors, *POPL*, pages 59–72. ACM, 2012.
- [25] J. B. Jensen and L. Birkedal. Fictional separation logic. In *ESOP*, volume 7211 of *LNCS*, 2012.
- [26] O. Kammar and G. D. Plotkin. Algebraic foundations for effect-dependent optimisations. In *POPL*, 2012.
- [27] V. Koutavas and M. Wand. Small bisimulations for reasoning about higher-order imperative programs. In *POPL*, 2006.
- [28] N. Krishnaswami, L. Birkedal, and J. Aldrich. Verifying event-driven programs using ramified frame properties. In *TLDI*, 2010.
- [29] N. Krishnaswami, A. Turon, D. Dreyer, and D. Garg. Superficially substructural types. In *ICFP*, 2012.
- [30] R. Ley-Wild and A. Nanevski. Subjective auxiliary state for coarse-grained concurrency. In *POPL*, 2013.
- [31] S. Mac Lane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics. Springer, 2nd edition, Sept. 1998.
- [32] B. Meyer. Applying "design by contract". *IEEE Computer*, 25(10):40–51, 1992.
- [33] F. Oles. *A Category-Theoretic Approach to the Semantics of Programming Languages*. PhD thesis, CMU, 1982.
- [34] A. Pitts and I. Stark. Operational reasoning for functions with local state. In *Higher order operational techniques in semantics*, pages 227–273. Cambridge University Press, 1998.
- [35] A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*, volume 57 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2013.
- [36] A. M. Pitts and I. D. B. Stark. Observable properties of higher-order functions that dynamically create local names, or what's new? In *MFCs*, volume 711 of *LNCS*, 1993.
- [37] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.
- [38] J. Schwinghammer, L. Birkedal, B. Reus, and H. Yang. Nested Hoare triples and frame rules for higher-order store. *Logical Methods in Computer Science*, 7(3), 2011.
- [39] I. Stark. *Names and Higher-Order Functions*. PhD thesis, U. Cambridge, 1994.
- [40] E. Sumii and B. C. Pierce. A bisimulation for type abstraction and recursion. In *POPL*, 2005.
- [41] J. Thamsborg and L. Birkedal. A Kripke logical relation for effect-based program transformations. In *ICFP*, 2011.
- [42] V. Vafeiadis and M. J. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, volume 4703 of *LNCS*, 2007.
- [43] P. Wadler and P. Thiemann. The marriage of effects and monads. *ACM Trans. Comput. Log.*, 4(1):1–32, 2003.