

A monad for full ground reference cells

Kammar, Ohad; Levy, Paul Blain; Moss, Sean K.; Staton, Sam

DOI:

[10.1109/LICS.2017.8005109](https://doi.org/10.1109/LICS.2017.8005109)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Kammar, O, Levy, PB, Moss, SK & Staton, S 2017, A monad for full ground reference cells. in *Proceedings of 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS2017)*. IEEE Computer Society Press, 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2017), Reykjavik, Iceland, 20/06/17. <https://doi.org/10.1109/LICS.2017.8005109>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

(c) 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works.
Eligibility for repository: Checked on 5/5/2017

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

A monad for full ground reference cells

Ohad Kammar^{*†§}, Paul B. Levy[‡], Sean K. Moss^{†¶}, and Sam Staton^{*}

^{*}University of Oxford Department of Computer Science

[†]University of Cambridge [§]Computer Laboratory and [¶]Department of Pure Mathematics and Mathematical Statistics

[‡]University of Birmingham School of Computer Science

Abstract—We present a denotational account of dynamic allocation of potentially cyclic memory cells using a monad on a functor category. We identify the collection of heaps as an object in a different functor category equipped with a monad for adding hiding/encapsulation capabilities to the heaps. We derive a monad for full ground references supporting effect masking by applying a state monad transformer to the encapsulation monad. To evaluate the monad, we present a denotational semantics for a call-by-value calculus with full ground references, and validate associated code transformations.

I. INTRODUCTION

Linked lists are a common example of a dynamically allocated, mutable, and potentially cyclic data type involving three sorts of memory cell.

- `linked_list` cells store values of type $1 + \text{ref}_{\text{list_cell}}$, i.e. a variant that is either a `NIL` value or a pointer to a `list_cell`;
- `list_cell` cells store values of type $\text{ref}_{\text{data}} * \text{ref}_{\text{linked_list}}$, i.e. a pair of pointers to a data payload and another list; and
- `data` cells store values of type `bool`, i.e., a single bit.

Such reference cells, that may contain references to other references and create cycles, but may not store functions and thunks, are called *full ground references* [25].

Here we develop a denotational semantics for languages with full ground storage in which types denote sets-with-structure and program terms denote structure-preserving functions. When dynamic allocation is involved, such a semantics typically uses functor categories and is called possible-world semantics, e.g. [27], [28], [30], [33]. To motivate the functorial structure of types, consider the type of `linked_list` and note that unless a memory cell has been previously allocated, the empty list is the only possible value of this type. As we allocate new cells, our programs can access more linked lists, and the type of linked lists is functorial in *worlds*: collections of previously allocated memory locations.

Our work builds on Moggi’s [23] theory of computational effects and monads. Moggi [24] gave monads for dynamic allocation of names and dynamically allocated *ground* storage, where the range of storable values does not change between worlds, such as `ref_data` cells, which only store a bit, regardless of what memory locations have been previously allocated. Plotkin and Power [30] give a different monad for ground storage on the same functor category as Moggi. Their monad is Hilbert-Post complete with respect to the expected first-order program equivalences [34]. Ghica’s masters thesis [10]

pioneered a functor category semantics for Idealised Algol extended with pointers, where dangling pointers played a crucial role in the model. The monad we propose here supports full ground storage for ML-like references without dangling pointers and validates the expected first-order program equivalences.

Possible-world semantics for local storage typically involves *heaps*, which assign a value of the appropriate type to each of a given set of locations. The situation is straightforward in the (ordinary) ground case, as heaps collect into a functor *contravariant* in worlds, with the functorial action given by projection. However, when heaps may contain cyclic data, this functorial action is ill-defined. For example, projecting a heap containing a cyclic list with two cells into a world containing only a single cell reference results in a dangling pointer.

We give a functorial action on heaps by defining the category of *initialisations*. Its objects are worlds, and its morphisms are world morphisms together with *initialisation data*: a specified value for each of the newly added locations in the codomain. The collection of heaps is functorial with respect to initialisations, and moreover, there is a way to transform monads over functors from initialisations into monads over functors from ordinary worlds, equipping the transformed monad with operations for mutating and dereferencing pointers.

We choose a specific monad over functors from initialisation which supports hiding capabilities. Adding this hiding capability to the heaps object gives the generalisation of the contravariant action given by projection in the (ordinary) ground case. The monad for full ground references is given by transforming the hiding monad in the aforementioned way. We define an allocation operation on this resulting monad.

To evaluate the suitability of this monad for modelling reference cells, we use three yardsticks. First, we prove this monad has the *effect masking* property: the global elements of the monad applied to a constant functor factor uniquely through the unit of the monad. We interpret this result as stating that computations that do not leak references are semantically equivalent to pure values. Second, we use the monad to give adequate semantics to a total call-by-value calculus with full ground references. Finally, we show that the axioms for (non-full) ground references [30], [34] hold in this model.

We compare several existing semantic approaches to the denotational semantics of reference cells, to place the possible-worlds semantics in context.

Relational models: In these models [2]–[4], [7], heaps may associate to a location values of any type, and the semantics is given non-functorially. Semantic equivalence in these models does not validate some of the basic intended equations, such as:

$$\text{let } x = \text{new true in true} \equiv \text{true}$$

These models define, in addition, a logical relations interpretation, and being related by this relation implies contextual equivalence [12]. We then validate the equations of interest with respect to this relation. Contrasted with the semantics we present here, relational models are much simpler, and as a consequence they can model richer collections of effects than what we consider here.

Step-indexing models: The key property of full ground storage is that, when initialising a new cell, all the information we need in order to determine its value is in the target world. In contrast, when references can hold functions or thunks, we also need to store the way such higher-order values will behave in future worlds, leading to a potential cyclicity in the semantics. Step-indexing models introduce a well-founded hierarchy on worlds that breaks this circularity. Step-indexing semantics to general references is given syntactically [1], [5], [8], or relationally [5], [6]. We hope that synthesising our technique with step-indexing models or other recursive-domain techniques would allow us to resolve the circularity in possible worlds and extend our model structure to general references.

Games models: Game semantics is especially well-suited for modelling local state in the presence of higher-order functions, and full ground state is no exception [25]. In game semantics, program terms denote *strategies*: dialogues between a Player — the program — and an Opponent — its environment. In such models the semantics of the heap is implicit (though, cf. nominal game semantics [36]), and manifest in the abilities of the Opponent. In contrast, in a sets-with-structure semantics, the heaps are fully explicit. In particular, it might be difficult to semantically decompose a game semantics into a monad over a bi-cartesian closed category for modelling pure languages.

To keep the discussion precise, we give two (general) points of comparison between the two kinds of semantics. The sets-with-structure semantics we present here does not validate some equations at higher types that a game semantics would normally validate, for example:

$$\begin{aligned} &\vdash \lambda_- : \mathbf{1}.\text{true} \\ &\equiv \text{let } x = \text{new true in } \lambda_- : \mathbf{1}.\!x \end{aligned}$$

The reason is that the semantics allows us to inspect, for example, whether a value depends on references [29], [34].

In contrast, a sets-with-structure semantics makes it easier to exploit that some types are uninhabited, for example, to prove that in a total call-by-value language:

$$x : \mathbf{1} \rightarrow \mathbf{0} \vdash \text{true} \equiv \text{false} : \text{bool}$$

where $\mathbf{0}$ is the empty type. There are currently no call-by-value game-semantics for a total language that validate this equation. It might be possible to develop such a model, but it would be a less natural game semantics, as these have partiality wired in.

Parametric models: Using the semantic machinery needed to interpret parametric polymorphism, Reddy and Yang [31] give semantics to full ground storage. Each type denotes a functor that has at every world, in addition to the set of values at that world, a relation between those values. These facilitate semantic universal and existential quantification over worlds both at the level of types and terms. This model also makes use of dangling pointers. We hope further work would clarify how the extra semantic structure in such parametric models relates to our models.

Contribution:

- 1) We give a new monad for full ground storage over the category of functors from worlds and their morphisms into sets and functions.
- 2) We identify the collection of heaps as a functor from worlds and initialisations.
- 3) We decompose our full ground storage monad into a global state transformer applied to a monad for encapsulation.
- 4) We evaluate the monad in three ways:
 - a) We prove the full ground storage monad satisfies the effect masking property.
 - b) We use the monad to give adequate denotational semantics to a total call-by-value calculus for full ground storage.
 - c) We show this model satisfies the usual equations for ground storage.

The rest of the paper is structured as follows. Sec. II defines full ground storage through the syntax and operational semantics of a calculus, and highlights where the semantic structure we expose appears in the operational account. Sec. III reviews the category-theoretic background and concepts we need for our development. Sec. IV defines the category of worlds and the category of initialisations. Sec. V gives an explicit formula for the full ground references monad. Sec. VI decomposes the monad into a state transformed monad for encapsulation, and uses this decomposition to analyse both the encapsulation monad and the full ground storage monad. Sec. VII returns to the calculus of Sec. II, uses our monad to give denotational semantics to this calculus, and uses its adequacy to validate that the program equations for (ordinary) ground state carry over to the full ground setting. Sec. VIII concludes.

II. FULL GROUND STORAGE

Our formalism consists of two parts. First, we fix the description of the storable data structures in a well-founded way. We then define the syntax, type system, and semantics of programs that manipulate data structures involving those types.

The first component is a (typically countable) set \mathbf{S} whose elements c are called *cell sorts*. Given \mathbf{S} , we define the set $\mathbf{G}^{\mathbf{S}}$ of *full ground types* γ given inductively by

$$\gamma ::= \mathbf{0} \mid \gamma_1 + \gamma_2 \mid \mathbf{1} \mid \gamma_1 * \gamma_2 \mid \mathbf{ref}_c$$

We omit the superscript in $\mathbf{G}^{\mathbf{S}}$, and other superscripts and subscripts, wherever possible. The second component is a function $ctype : \mathbf{S} \rightarrow \mathbf{G}$ assigning to each sort its *content type*: the type of values stored in cells of this sort.

Example 1. To capture the example from the introduction, choose $\mathbf{S} := \{\text{linked_list}, \text{list_cell}, \text{data}\}$ and set $ctype\ c, c \in \mathbf{S}$, as in the introduction. \square

A *full ground storage signature* is a pair $\Sigma = \langle \mathbf{S}^{\Sigma}, ctype^{\Sigma} \rangle$ of such a set and a content type assignment for it. We can view a full ground storage signature as an abstract description of a sequence of top-level data declarations. Fix such a signature Σ for the remainder of this manuscript.

A. Syntax

Fig. 1 presents the $\lambda_{\text{ref}}^{\Sigma}$ -calculus, our subject of study. We let x range over a countable set of *identifiers*, and ℓ range over a countably infinite set \mathbb{L} of *locations*. The occurrences of x 's in the following constructs are binding: function abstraction, non-empty pattern matching, and allocation. It is a standard call-by-value Church-style, higher-order calculus. We shade the parts specific to references.

We include reference literals ℓ which will inhabit the reference types \mathbf{ref}_c . The core of high-level languages like ML does not usually have global memory locations. However, in our core calculus we will include values for references for two reasons. First, having values for references makes the operational semantics straightforward and similar in appearance to the natural global state semantics. Second, the resulting calculus enables us to present some program equivalences involving distinct memory locations without introducing additional type constructors. The decision to include memory locations in the base language is common practice in operational and denotational semantics for local state [4], [7], [20], [33].

The assignment and dereferencing constructs are standard. We use a non-standard allocation operation [20] which allows the simultaneous initialisation of a cyclic structure. Each of the initialisation values v_i has access to each of the other newly allocated references x_1, \dots, x_n . We require the initialisation data to be given as values. Allowing computation at this point would be unsound, as an arbitrary computation may try to dereference the yet-uninitialised locations.


We will make use of the following syntactic sugar:

$$\begin{aligned} \mathbf{let}\ (x : \tau) = t\ \mathbf{in}\ s &\equiv (\lambda x : \tau. s)\ t \\ \mathbf{new}_{ct} &\equiv \mathbf{let}\ (x : ctype\ c) = t\ \mathbf{in} \\ &\quad \mathbf{letref}\ (y : \mathbf{ref}_c) := x\ \mathbf{in}\ y \\ (t_1, \dots, t_n) &\equiv (t_1, (\dots, t_n)) \\ \tau_1 * \dots * \tau_n &\equiv \tau_1 * (\dots * \tau_n) \end{aligned}$$

When it is clear from the context, we omit type annotations. With these conventions in place, the examples in the introduction are special cases of full ground storage.

$\tau ::=$	types
\mathbf{ref}_c	reference
$\mathbf{0}$	empty
$\tau_1 + \tau_2$	binary sum
$\mathbf{1}$	unit
$\tau_1 * \tau_2$	binary product
$\tau_1 \rightarrow \tau_2$	function
$v ::=$	values
ℓ	location
x	identifier
$\mathbf{inj}_i^{\tau_1 + \tau_2} v$	sum constructor ($i = 1, 2$)
$()$	unit
(v_1, v_2)	pair
$\lambda x : \tau. t$	function abstraction
$t, s ::=$	terms
ℓ	location
x	identifier
$\mathbf{inj}_i^{\tau_1 + \tau_2} t$	sum constructor ($i = 1, 2$)
$()$	unit
(t, s)	pair
$\lambda x : \tau. t$	function abstraction
$\mathbf{match}\ t\ \mathbf{with}\ \{\}^{\tau}$	pattern matching: empty
$\mathbf{match}\ t\ \mathbf{with}$	binary
$\{\mathbf{inj}_1 x_1 \mapsto s_1$	
$\mid \mathbf{inj}_2 x_2 \mapsto s_2\}$	
$\mathbf{match}\ t\ \mathbf{with}\ (x_1, x_2) \mapsto s$	products
$t\ s$	function application
$t := s$	assignment
$!t$	dereferencing
\mathbf{letref}	allocation
$(x_1 : \mathbf{ref}_{c_1}) := v_1,$	
\vdots	
$(x_n : \mathbf{ref}_{c_n}) := v_n$	
$\mathbf{in}\ t$	

Fig. 1: The types and syntax of $\lambda_{\text{ref}}^{\Sigma}$

Example 2. We allocate a cyclic list:  \mathbf{letref} cyclic_list \bullet

(payload : $\mathbf{ref}_{\text{data}}$) := 42,
(cyclic_list : $\mathbf{ref}_{\text{linked_list}}$) := \mathbf{inj}_2 head,
(head : $\mathbf{ref}_{\text{list_cell}}$) := (payload, cyclic_list) \square

Our type-system needs to associate a sort to each location literal ℓ the program may use. A *heap layout* w is a partial function with finite support $w : \mathbb{L} \rightarrow_{\text{fin}} \mathbf{S}$. We write $\{\ell_1 : c_1, \dots, \ell_n : c_n\}$ for the heap layout whose support is $w := \{\ell_1, \dots, \ell_n\}$ defined by $w(\ell_i) = c_i$. When $w(\ell) = c$, we write $(\ell : c) \in w$. We write $w \leq w'$ when w' extends w . Layout extension is thus a partial order. Heap layouts are a standard abstraction and appear under different names: state-types [4], [7], and location worlds [31].

Typing contexts Γ are partial functions with finite support from the set of identifiers to the set of types. We use the list-

$$\begin{array}{c}
\frac{}{\Gamma \vdash_w \ell : \mathbf{ref}_c} ((\ell : c) \in w) \qquad \frac{}{\Gamma \vdash_w x : \tau} ((x : \tau) \in \Gamma) \\
\\
\frac{\Gamma \vdash_w t : \tau_i}{\Gamma \vdash_w \mathbf{inj}_i^{\tau_1 + \tau_2} t : \tau_1 + \tau_2} (i \in \{1, 2\}) \qquad \frac{}{\Gamma \vdash_w () : \mathbf{1}} \\
\\
\frac{\Gamma \vdash_w t_1 : \tau_1 \quad \Gamma \vdash_w t_2 : \tau_2}{\Gamma \vdash_w (t_1, t_2) : \tau_1 * \tau_2} \qquad \frac{\Gamma, x : \tau \vdash_w t : \tau'}{\Gamma \vdash_w \lambda x : \tau. t : \tau \rightarrow \tau'} \\
\\
\frac{\Gamma \vdash_w t : \mathbf{0}}{\Gamma \vdash_w \mathbf{match} t \mathbf{with} \{ \}^\tau : \tau} \\
\\
\frac{\Gamma \vdash_w t : \tau_1 + \tau_2 \quad \Gamma, x_1 : \tau_1 \vdash_w s_1 : \tau \quad \Gamma, x_2 : \tau_2 \vdash_w s_2 : \tau}{\Gamma \vdash_w \mathbf{match} t \mathbf{with} \{ \mathbf{inj}_1 x_1 \mapsto s_1 \mid \mathbf{inj}_2 x_2 \mapsto s_2 \} : \tau} \\
\\
\frac{\Gamma \vdash_w t : \tau_1 * \tau_2 \quad \Gamma, x_1 : \tau_1, x_2 : \tau_2 \vdash_w s : \tau'}{\Gamma \vdash_w \mathbf{match} t \mathbf{with} (x_1, x_2) \mapsto s : \tau'} \\
\\
\frac{\Gamma \vdash_w t : \tau \rightarrow \tau' \quad \Gamma \vdash_w s : \tau}{\Gamma \vdash_w t s : \tau'} \\
\\
\frac{\Gamma \vdash_w t : \mathbf{ref}_c \quad \Gamma \vdash_w s : \mathit{ctype} c}{\Gamma \vdash_w t := s : \mathbf{1}} \qquad \frac{\Gamma \vdash_w t : \mathbf{ref}_c}{\Gamma \vdash_w !t : \mathit{ctype} c} \\
\\
\text{for } i = 1, \dots, n: \\
\frac{\Gamma, x_1 : \mathbf{ref}_{c_1}, \dots, x_n : \mathbf{ref}_{c_n} \vdash_w v_i : \mathit{ctype} c_i \quad \Gamma, x_1 : \mathbf{ref}_{c_1}, \dots, x_n : \mathbf{ref}_{c_n} \vdash_w t : \tau}{\Gamma \vdash_w \mathbf{letref}(x_1 : \mathbf{ref}_{c_1}) := v_1,} \\
\qquad \vdots \\
\qquad (x_n : \mathbf{ref}_{c_n}) := v_n \mathbf{in} t : \tau
\end{array}$$

Fig. 2: The inductive definition of the typing relation of $\lambda_{\mathbf{ref}}^\Sigma$

like notation $\Gamma, x : \tau$ for the extension of Γ by the assignment $x \mapsto \tau$, and the membership-like notation $(x : \tau) \in \Gamma$ to state that $\Gamma(x) = \tau$. The type system is given in Fig. 2 via an inductively defined quaternary relation $\Gamma \vdash_w t : \tau$ between contexts Γ , layouts w , terms t , and types τ .

Location literals ℓ are limited to the locations in the layout w , which does not change throughout the typing derivation. To assign, the type of the assigned value needs to match the sort of the reference, and similarly the type of the dereferenced value matches that of the reference. Finally, for allocation, the initialisation values may refer to each of the newly allocated references, as does the remainder of the computation.

By construction, every typeable term has a unique type in a given context.

A *value substitution* θ is a partial function with finite support

from the set of identifiers to values. Defining capture avoiding substitution $t[\theta]$ and proving the substitution lemma is standard and straightforward. In the sequel we will hand-wave around the standard issues with α -equivalence, and omit the standard freshness conditions on bound variables.

Finally, the type system is monotone with respect to layout extension: $\Gamma \vdash_w t : \tau$, and $w' \geq w$ implies $\Gamma \vdash_{w'} t : \tau$. In particular, if we define τw to be the set of closed values of type τ assuming layout w , τw is functorial in w in the following sense: for every $w \leq w'$, $\tau w \subseteq \tau w'$.

B. Operational semantics

We present a big-step operational semantics. We expect a small-step semantics or stack-machine semantics to be similarly straightforward.

An *untyped heap* η is a partial function with finite support w_η from \mathbb{L} to values. A *typed heap* η consists of a pair of a layout w_η , and a function from the set of locations in w_η to values that assigns to every $(\ell : c) \in w_\eta$ a well-typed closed value: $\vdash_{w_\eta} \eta(\ell) : \mathit{ctype} c$. We denote the set of heaps with layout w by $\mathbf{H}w$. For every layout w , an untyped heap can be turned into a typed heap from $\mathbf{H}w$ in at most one way. Note that $\mathbf{H}w$ is not functorial with respect to layout extension: there is no obvious way to turn an arbitrary heap in $\mathbf{H}w$ into a heap in $\mathbf{H}w'$ for every $w' \geq w$.

A *configuration* is a pair $\langle t, \eta \rangle$ consisting of a term t and an untyped heap η . A *terminal configuration* is one whose term is a value. We say that a sequence of locations ℓ_1, \dots, ℓ_n is *fresh* for layout w , and write $\#_w \langle \ell_1, \dots, \ell_n \rangle$, when the locations are pairwise distinct and collectively disjoint from w 's support. Fig. 3 defines the evaluation relation $\langle t, \eta \rangle \Downarrow \langle v, \eta' \rangle$ between configurations and terminal configurations.

Locations, as values, are fully evaluated. Assignment updates the heap, and dereferencing retrieves the appropriate value from the heap. The rule for allocation requires the newly allocated locations to be fresh for the current heap, and then extends the heap with the given initialisation values with the new locations substituted in. It then carries the execution in the body of the allocation with those new locations substituted in. As the only requirement of the new location is to be fresh, this semantics is not deterministic, and a phrase might evaluate to several different configurations with different heap layouts.

We prove Felleisen-Wright soundness for $\lambda_{\mathbf{ref}}^\Sigma$.

Theorem 1 (preservation). *Evaluation preserves typeability: for every well-typed closed term $\vdash_w t : \tau$ and for every $w_1 \geq w$ and $\eta_1 \in \mathbf{H}w_1$, if $\langle t, \eta_1 \rangle \Downarrow \langle v, \eta_2 \rangle$, then there is some $w_2 \geq w_1$ such that $\vdash_{w_2} v : \tau$ and $\eta_2 \in \mathbf{H}w_2$.*

The proof is by straightforward induction on the evaluation relation after strengthening the induction hypothesis to closed substitutions in open terms.

Theorem 2 (totality). *All well-typed closed programs fully evaluate: for every $\vdash_w t : \tau$, $w_1 \geq w$ and $\eta_1 \in \mathbf{H}w_1$ there exist some $w_2 \geq w_1$, $\vdash_{w_2} v : \tau$, and $\eta_2 \in \mathbf{H}w_2$ such that $\langle t, \eta_1 \rangle \Downarrow \langle v, \eta_2 \rangle$.*

$$\begin{array}{c}
\frac{\langle \ell, \eta \rangle \Downarrow \langle \ell, \eta \rangle}{\langle \ell, \eta \rangle \Downarrow \langle \ell, \eta \rangle} \\
\frac{\langle t, \eta \rangle \Downarrow \langle v, \eta' \rangle}{\langle \mathbf{inj}_i^{\tau_1 + \tau_2} t, \eta \rangle \Downarrow \langle \mathbf{inj}_i^{\tau_1 + \tau_2} v, \eta' \rangle} \\
\frac{\langle \langle \rangle, \eta \rangle \Downarrow \langle \langle \rangle, \eta \rangle}{\langle \langle t_1, \eta \rangle \Downarrow \langle v_1, \hat{\eta} \rangle \quad \langle t_2, \hat{\eta} \rangle \Downarrow \langle v_2, \eta' \rangle} \quad \langle \langle t_1, t_2 \rangle, \eta \rangle \Downarrow \langle \langle v_1, v_2 \rangle, \eta' \rangle} \\
\frac{\langle \lambda x : \tau.t, \eta \rangle \Downarrow \langle \lambda x : \tau.t, \eta \rangle}{\langle t, \eta \rangle \Downarrow \langle \mathbf{inj}_i^{\tau_1 + \tau_2} \hat{v}, \hat{\eta} \rangle \quad \langle s_i[x_i \mapsto \hat{v}], \hat{\eta} \rangle \Downarrow \langle v, \eta' \rangle} \\
\frac{\langle \mathbf{match} t \text{ with } \{ \mathbf{inj}_1 x_1 \mapsto s_1 \mid \mathbf{inj}_2 x_2 \mapsto s_2 \}, \eta \rangle \Downarrow \langle v, \eta' \rangle}{\langle t, \eta \rangle \Downarrow \langle (v_1, v_2), \eta' \rangle \quad \langle s[x_1 \mapsto v_1, x_2 \mapsto v_2], \eta' \rangle \Downarrow \langle v, \eta'' \rangle} \\
\frac{\langle \mathbf{match} t \text{ with } (x_1, x_2) \mapsto s, \eta \rangle \Downarrow \langle v, \eta'' \rangle}{\langle t, \eta \rangle \Downarrow \langle \lambda x : \tau.t', \eta_1 \rangle \quad \langle s, \eta_1 \rangle \Downarrow \langle v', \eta_2 \rangle} \\
\frac{\langle t', \eta \rangle \Downarrow \langle v', \eta' \rangle}{\langle t s, \eta \rangle \Downarrow \langle v, \eta' \rangle} \\
\frac{\langle t, \eta \rangle \Downarrow \langle \ell, \hat{\eta} \rangle \quad \langle s, \hat{\eta} \rangle \Downarrow \langle v, \eta' \rangle}{\langle t := s, \eta \rangle \Downarrow \langle \langle \rangle, \eta'[\ell \mapsto v] \rangle} \quad (\ell \in \underline{w}_{\eta'}) \\
\frac{\langle t, \eta \rangle \Downarrow \langle \ell, \eta' \rangle}{\langle !t, \eta \rangle \Downarrow \langle \eta'(\ell), \eta' \rangle} \quad (\ell \in \underline{w}_{\eta'})
\end{array}$$

$$\frac{\left\langle t[\theta], \eta \left[\ell_i \mapsto v_i[\theta] \right]_{i=1}^n \right\rangle \Downarrow \langle v, \eta' \rangle}{\text{letref} \quad \left\langle \begin{array}{l} (x_1 : \mathbf{ref}_{c_1}) := v_1, \\ \vdots \\ (x_n : \mathbf{ref}_{c_n}) := v_n \end{array} \text{ in } t \right\rangle \Downarrow \langle v, \eta' \rangle} \quad (\#_{\underline{w}_n} \langle \ell_1, \dots, \ell_n \rangle)$$

Fig. 3: The operational semantics of $\lambda_{\text{ref}}^{\Sigma}$

The proof is standard using Tait's method [35] and Kripke logical predicates, and a w -indexed predicate on $\mathbf{H}w$ that is *not* Kripke, as $\mathbf{H}w$ is not functorial in layout extensions.

We focus on the following aspects from the allocation case in the course of this proof. In that case, the allocation construct is typed with respect to heap layout w , but operates on a heap with an extended layout $w \leq w'$. We can find fresh locations ℓ_1, \dots, ℓ_n , and then form the following square of layout extensions

$$\begin{array}{ccc}
w \leq w \oplus \{ \ell_1 : c_1, \dots, \ell_n : c_n \} & & \\
\lrcorner \wedge & & \lrcorner \wedge \\
w' \leq w' \oplus \{ \ell_1 : c_1, \dots, \ell_n : c_n \} & &
\end{array}$$

where the operation \oplus denotes layout extension by the given locations and sorts. The initialisation data in the allocation construct is given for the top extension. The crucial step in the proof in this case is that we can transform this initialisation data into initialisation data for the extension in the bottom row. Applying this transformed initialisation data on the given heap is precisely the functorial action of the collection of heaps.

C. Observational equivalence

To define observational equivalence, we need a few more technical definitions. Let Γ, Γ' be two typing contexts. We say that Γ' *extends* Γ , and write $\Gamma' \geq \Gamma$ when Γ' extends Γ as a function from identifier names to types. In the sequel we write $\Gamma \vdash_w t, s : \tau$ to indicate that the quintuple $\langle \Gamma, w, t, s, \tau \rangle$ belongs to the quinary relation given by the conjunction $\Gamma \vdash_w t : \tau$ and $\Gamma \vdash_w s : \tau$.

Consider any two terms $\Gamma \vdash_w t, s : \tau$. We define the *set of contexts plugged with* $\Gamma \vdash_w t, s : \tau$, which we denote by $\mathcal{C}[\Gamma \vdash_w t, s : \tau]$, as the smallest quinary relation jointly compatible with the typing rules that contains the quintuples $\langle \Gamma', w', t, s, \tau \rangle$ for every $\Gamma' \vdash_{w'} t, s : \tau$, $\Gamma' \geq \Gamma$, and $w' \geq w$.

We say that two terms $\Gamma \vdash_w t_1, t_2 : \tau$ are *observationally equivalent* when for all closed boolean plugged contexts $\vdash_{w'} s_1, s_2 : \mathbf{bool} \in \mathcal{C}[\Gamma \vdash_w t_1, t_2 : \tau]$, heaps $\eta' \in \mathbf{H}w'$ and boolean values $\vdash v : \mathbf{bool}$, we have:

$$\exists \eta_1 (\langle s_1, \eta' \rangle \Downarrow \langle v, \eta_1 \rangle) \iff \exists \eta_2 (\langle s_2, \eta' \rangle \Downarrow \langle v, \eta_2 \rangle)$$

Note that by the Preservation Theorem 1, if such heaps η_i exist, they can be typed.

III. PRELIMINARIES

We assume familiarity with categories, functors, and natural transformations. We denote by \mathbf{Set} the category of sets and functions. Let \mathbb{C} be a small category. We denote the category of functors $X, Y : \mathbb{C} \rightarrow \mathbf{Set}$ and natural transformations between them by $[\mathbb{C}, \mathbf{Set}]$. As we will interpret types in such a category, we recall its bi-cartesian closed structure: its sums and products are given component-wise, and the exponential is given by an end formula (see below). We assume familiarity with ends and coends over \mathbf{Set} , which we will use in the explicit description of the full ground storage monad. To fix terminology and notation, given a mixed variance functor $P : \mathbb{C}^{\text{op}} \times \mathbb{C} \rightarrow \mathbf{Set}$, we denote its end and its ending wedge as follows, for all $w' \in \mathbb{C}$:

$$\pi_{w'} : \int_{w \in \mathbb{C}} P(w, w) \rightarrow P(w', w')$$

We denote its coend and its coending wedge as follows:

$$q_{w'} : P(w', w') \rightarrow \int^{w \in \mathbb{C}} P(w, w)$$

Consider any functor $u : \mathbb{E} \rightarrow \mathbb{W}$ between two small categories. Precomposition with u induces a functor $u^* : [\mathbb{W}, \mathbf{Set}] \rightarrow [\mathbb{E}, \mathbf{Set}]$. By generalities, this functor has a right adjoint $u_* : [\mathbb{E}, \mathbf{Set}] \rightarrow [\mathbb{W}, \mathbf{Set}]$ called the *right Kan extension along* u , which we will use to establish the

existence of the monoidal strength structure for our monad. Let $w \in \mathbb{W}$ be an object in \mathbb{W} . We will use the following specific kind of comma category in our ends and coends. The comma category $w \downarrow u$ has as objects pairs $\langle e, \rho \rangle$ where $e \in \mathbb{E}$ is an object and $\rho : w \rightarrow ue$ is a morphism in \mathbb{W} . Its morphisms $\varepsilon : \langle e_1, \rho_1 \rangle \rightarrow \langle e_2, \rho_2 \rangle$ are morphisms $\varepsilon : e_1 \rightarrow e_2$ in \mathbb{E} such that $u\varepsilon \circ \rho_1 = \rho_2$. As we assumed \mathbb{E} and \mathbb{W} are small, so is $w \downarrow u$. Precomposition with the projection functor $\pi : (w \downarrow u) \rightarrow \mathbb{E}$ turns every mixed variance functor $P : \mathbb{E}^{\text{op}} \times \mathbb{E} \rightarrow \mathbf{Set}$ into a mixed variance functor $\bar{P} : (w \downarrow u)^{\text{op}} \times (w \downarrow u) \rightarrow \mathbf{Set}$. There is then a canonical isomorphism:

$$\int_{w \downarrow u} \bar{P} = \int_{\langle e, \rho \rangle \in w \downarrow u} P(e, e) \cong \int_{e \in \mathbb{E}} (P(e, e))^{\mathbb{W}(w, ue)}$$

Also recall the end formula for exponentials in $[\mathbb{W}, \mathbf{Set}]$:

$$X^Y w \cong \int_{w' \in \mathbb{W}} (Xw')^{\mathbb{W}(w, w') \times Yw'} \cong \int_{\rho : w \rightarrow w' \in w \downarrow \text{id}_{\mathbb{W}}} (Xw')^{(Yw')}$$

To describe the tensorial strengths for our monads, we assume familiarity with symmetric monoidal closed categories, though concretely we will only use a cartesian structure. The following concept will simplify the presentation of the strengths. Let $\underline{\mathbf{W}} = \langle \mathbf{W}, \otimes, I, a, l, r \rangle$ be a symmetric monoidal closed category. A *monoidal $\underline{\mathbf{W}}$ -action category*, also known as a *$\underline{\mathbf{W}}$ -actegory*, is a tuple $\underline{\mathbf{E}} = \langle \mathbf{E}, \odot, \alpha, \lambda \rangle$ where:

- $\odot : \mathbf{W} \times \mathbf{E} \rightarrow \mathbf{E}$ is a two-argument functor;
- $\alpha_{X,Y,A} : (X \otimes Y) \odot A \rightarrow X \odot (Y \odot A)$ is a natural isomorphism; and
- $\lambda_A : I \odot A \rightarrow A$ is a natural isomorphism

subject to the following coherence axioms:

$$\begin{aligned} \alpha_{X,Y,Z \odot A} \circ \alpha_{X \otimes Y, Z, A} \\ = \text{id}_X \odot \alpha_{Y,Z,A} \circ \alpha_{X,Y \otimes Z, A} \circ \alpha_{X,Y,Z} \odot \text{id}_A \\ \text{id}_X \odot \lambda_A \circ \alpha_{X,I,A} = r_X \odot \text{id}_A \end{aligned}$$

We say that a $\underline{\mathbf{W}}$ -actegory $\underline{\mathbf{E}}$ is *bi-closed* if the following two right adjoints exist:

- $- \odot A \dashv A \multimap - : \mathbf{E} \rightarrow \mathbf{W}$ for every $A \in \mathbf{E}$
- $X \odot - \dashv X \multimap - : \mathbf{E} \rightarrow \mathbf{E}$ for every $X \in \mathbf{W}$.

We denote the natural bijection of the first adjunction by $\text{curry}_{X,A,B}^{\circ} : \mathbf{W}(X \odot A, B) \rightarrow \mathbf{E}(X, A \multimap B)$, its inverse by uncurry° , and its counit by eval° , and similarly for \dashv .

As the name suggests, a $\underline{\mathbf{W}}$ -actegory is a $\underline{\mathbf{W}}$ -category ($\underline{\mathbf{W}}$ -enriched category) with the structure permuted:

Theorem ([11], [13]). *These data are canonically isomorphic:*

- a bi-closed $\underline{\mathbf{W}}$ -actegory $\underline{\mathbf{E}}$; and
- a $\underline{\mathbf{W}}$ -category $\underline{\mathbf{E}}$ with powers and copowers.

and moreover, the adjoints \multimap and \dashv enrich.

Up to isomorphism, the enrichment is given by the adjoint \multimap , the copowers by the monoidal action \odot , and the powers by the adjoint \dashv . The relevance of $\underline{\mathbf{W}}$ -actegories to our situation is that the monoidal action in our setting is much simpler than

the enrichment. Working with monoidal actions thus simplifies many calculations.

Finally, we assume familiarity with monads, formulated as Kleisli triples. Let $\underline{T} = \langle T, \text{return}, \gg \rangle$ be a monad over a $\underline{\mathbf{W}}$ -actegory $\underline{\mathbf{E}}$. A *tensorial strength* for \underline{T} is a natural transformation $\text{str} : X \odot TA \rightarrow T(X \odot A)$ satisfying the following coherence axioms:

$$\begin{aligned} T\lambda_A \circ \text{str}_{T,A} &= \lambda_{TA} \\ \text{str}_{X,Y \odot A} \circ \text{id}_X \odot \text{str}_{Y,A} \circ \alpha_{X,Y,TA} &= T\alpha_{X,Y,A} \circ \text{str}_{X \otimes Y, A} \\ \text{str}_{X,A} \circ \text{id}_X \odot \text{return}_A &= \text{return}_{X \odot A} \\ (\gg \text{id}_{T(X \odot A)}) \circ T\text{str}_{X,A} \circ \text{str}_{X,TA} &= \text{str}_{X,A} \circ \text{id}_X \odot \gg \text{id}_{TA} \end{aligned}$$

Let \underline{T} and \underline{S} be two such strong monads. We say that a monad morphism $m : \underline{T} \rightarrow \underline{S}$ is *strong* when for all $X \in \mathbb{W}, A \in \mathbf{E}$: $m_{X \odot A} \circ \text{str}_{X,A} = \text{str}_{X,SA} \circ \text{id}_X \odot m$. When the actegory is bi-closed, the data for a strong monad and the data for a $\underline{\mathbf{W}}$ -monad are canonically isomorphic (cf. [15]).

We are now ready to describe the state monad transformer that will give our monad for reference cells, as a straightforward consequence of transforming a monad across the adjunction $- \odot A \dashv A \multimap -$ (see also e.g. [22]):

Proposition 3. *Let $\underline{\mathbf{W}}$ be a symmetric monoidal closed category and $\underline{\mathbf{E}}$ be a bi-closed $\underline{\mathbf{W}}$ -actegory. Let A be an object in \mathbf{E} . For every strong monad P over $\underline{\mathbf{E}}$, we have a strong monad $\underline{T} = \{T, \text{return}, \gg, \text{str}\}$ over $\underline{\mathbf{W}}$ given by:*

$$TA := A \multimap P(- \odot A)$$

$$\text{return}_X^T := \text{curry}_{X,A,P(A \odot A)}^{\circ} (\text{return}_{X \odot A}^P)$$

$$\gg_{X,Y}^T f := A \multimap (\gg_{X \odot A, Y \odot A}^P \text{uncurry}_{X,A,P(Y \odot A)}^{\circ} f)$$

$$\text{str}_{X,Y}^T := \text{curry}_{X \otimes TY, A, P((X \otimes Y) \odot A)}^{\circ}$$

$$(P\alpha_{X,Y,A}^{-1} \circ \text{str}_{X,Y \odot A}^P \circ (\text{id}_X \odot \text{eval}_{A,P(Y \odot A)}^{\circ}) \circ \alpha_{X,TY,A})$$

and for every strong monad morphism $m : \underline{P}_1 \rightarrow \underline{P}_2$ we have a strong monad morphism $m^T : T_1 \rightarrow T_2$ between the corresponding monads, given by:

$$m_X^T := A \multimap (m_{X \odot A})$$

$$\begin{array}{ccc} & \mathbb{H} \multimap - & \\ & \curvearrowleft & \curvearrowright \\ \mathbf{W} & \dashv & \mathbf{E} \\ & \curvearrowright & \curvearrowleft \\ & - \odot \mathbb{H} & \end{array}$$

IV. WORLDS AND INITIALISATIONS

In Sec. II, we arranged the terms and values of $\lambda_{\text{ref}}^{\Sigma}$ based on their heap layout assumptions, w , and this arrangement is functorial with respect to layout extension $w \leq w'$. We saw that heaps also divide based on layouts, but not functorially. And finally, we highlighted that the proof of the Totality Theorem 2 makes use of the fact that initialisation data can be promoted along world extension. We will now expose this structure semantically.




Fig. 4: Heaplets, initialisations, and their operations

A. Worlds

The category \mathbb{W} of *worlds* has as objects the heap layouts of Sec. II, i.e., partial functions $w : \mathbb{L} \rightarrow_{\text{fin}} \mathbf{S}$ with finite support $\underline{w} \subseteq \mathbb{L}$. A morphism $\rho : w \rightarrow w'$ is an injection $\rho : \underline{w} \hookrightarrow \underline{w}'$ such that $(\ell : c) \in w$ implies $(\rho(\ell) : c) \in w'$. For brevity, we refer to heap layouts as worlds and to \mathbb{W} -morphisms as (world) injections.

We define several layout-manipulation operations on \mathbb{W} . While we work concretely, these operations can be axiomatised by universal properties using Simpson's *independence structures* [32], and we will use his vocabulary as much as possible.

Let $\# : \mathbb{N} \xrightarrow{\cong} \mathbb{L}$ be an enumeration of the set of locations. Given a world w , define $|w|$ to be the smallest index beyond all the locations in w , i.e. $\min\{n \in \mathbb{N} \mid \forall i \geq n. \# i \notin w\}$. Given two worlds $w_1, w_2 \in \mathbb{W}$, we can embed their supports into the following subset of \mathbb{L} :

$$\underline{w_1 \oplus w_2} := \underline{w_1} \cup \{\#(|w_1| + n) \mid \# n \in w_2\}$$

by setting $\iota_1^\oplus(\ell) := \ell$ and $\iota_2^\oplus(\#n) := \#(|w_1| + n)$. We then have that for every $\ell \in \underline{w_1 \oplus w_2}$ there is exactly one $i \in \{1, 2\}$ and $\ell_i \in w_i$ such that $\ell = \iota_i^\oplus \ell_i$. We define the *independent coproduct* $w_1 \oplus w_2$ whose support is given by $\underline{w_1 \oplus w_2}$ by setting $(w_1 \oplus w_2)(\iota_i^\oplus \ell) := w_i(\ell)$. Then $\iota_i^\oplus : w_i \rightarrow w_1 \oplus w_2$ are world injections which we call the *independent coprojections*. Moreover, the construction \oplus extends to a functor $\oplus : \mathbb{W} \times \mathbb{W} \rightarrow \mathbb{W}$ and each coprojection is a natural transformation. The independent coproduct is *not* a coproduct in \mathbb{W} , for example, there is no codiagonal injection $w \oplus w \rightarrow w$ for $w = \{\ell : c\}$. Independent coproducts are the semantic counterparts to extending a world with fresh locations.

Given an injection $\rho : w_1 \rightarrow w_2$, its *complement* is the injection $\rho^c : w_2 \ominus \rho \rightarrow w_2$ whose domain $\underline{w_2 \ominus \rho} := \underline{w_2} \setminus \text{Im}(\rho)$ are all the locations in w_2 that ρ misses, and the action of ρ^c is given by that of w_2 . There are canonical isomorphisms $w_1 \oplus (w_2 \ominus \rho) \cong w_2$ and $(w_1 \oplus w_2) \ominus \iota_i^\oplus \cong w_{3-i}$. We use complements to define initialisation data below.

Given two injections $\rho_i : w \rightarrow w_i$, $i = 1, 2$, we define their *local independent coproduct* by

$$\rho_1 \oplus_w \rho_2 := w \oplus (w_1 \ominus \rho_1) \oplus (w_2 \ominus \rho_2)$$

We have morphisms $w_1 \xrightarrow{\rho_1^* \rho_2} \rho_1 \oplus_w \rho_2 \xleftarrow{\rho_2^* \rho_1} w_2$ such that:

$$\begin{array}{ccc} w & \xrightarrow{\rho_2} & w_2 \\ \rho_1 \downarrow & = & \downarrow \rho_2^* \rho_1 \\ w_1 & \xrightarrow{\rho_1^* \rho_2} & \rho_1 \oplus_w \rho_2 \end{array}$$

We define the functor category $\mathbf{W} := [\mathbb{W}, \mathbf{Set}]$ in which we will interpret the types of the $\lambda_{\text{ref}}^\Sigma$ -calculus. We interpret the *full ground types*, defining $\llbracket - \rrbracket : \mathbf{G} \rightarrow \mathbf{W}$ by:

$$\llbracket \mathbf{0} \rrbracket := \mathbf{0} \quad \llbracket \gamma_1 + \gamma_2 \rrbracket := \llbracket \gamma_1 \rrbracket + \llbracket \gamma_2 \rrbracket$$

$$\llbracket \mathbf{1} \rrbracket := \mathbf{1} \quad \llbracket \gamma_1 * \gamma_2 \rrbracket := \llbracket \gamma_1 \rrbracket \times \llbracket \gamma_2 \rrbracket$$

$$\llbracket \text{ref}_c \rrbracket w := \{\ell \in w \mid w(\ell) = c\} \quad (\llbracket \text{ref}_c \rrbracket \rho) \ell := \rho(\ell)$$

We can interpret references more compactly by noting that $\llbracket \text{ref}_c \rrbracket \cong \mathbb{W}(\{\ell : c\}, -)$.

B. Initialisations

The account so far has been standard for possible-world semantics of local state. We now turn to defining the semantic counterpart to initialisation data.

Define the mixed-variance functor $\mathbb{H} : \mathbb{W}^{\text{op}} \times \mathbb{W} \rightarrow \mathbf{Set}$:

$$\mathbb{H}(w^-, w^+) := \prod_{(\ell : c) \in w^-} \llbracket \text{ctype } c \rrbracket w^+$$

Its contravariant action is given by projection, and its covariant action is given component-wise by the actions of $\llbracket \text{ctype } c \rrbracket$. Elements of $\mathbb{H}(w^-, w^+)$ are *heaplets* [26] whose layout is given by w^- , and whose values assume the layout w^+ (Fig. 4a). As in separation logic, heaplets are a composable abstraction facilitating local reasoning about the heap. This functor preserves the independent coproducts in the sense that $\mathbb{H}(w_1 \oplus w_2, w)$ and $\mathbb{H}(\iota_i^\oplus, w) : \mathbb{H}(w_1 \oplus w_2, w) \rightarrow \mathbb{H}(w_i, w)$ form the product of $\mathbb{H}(w_1, w)$ and $\mathbb{H}(w_2, w)$, and that $\mathbb{H}(\emptyset, w)$ is the singleton. Consequently, we have canonical isomorphisms $\mathbb{H}^\emptyset : \mathbf{1} \xrightarrow{\cong} \mathbb{H}(\emptyset, w)$ and, depicted in Fig. 4b, $\mathbb{H}^\oplus : \mathbb{H}(w_1, w) \times \mathbb{H}(w_2, w) \xrightarrow{\cong} \mathbb{H}(w_1 \oplus w_2, w)$. Fig. 4c depicts the contravariant action of \mathbb{H} on the canonical isomorphism $w_2 \cong w_1 \oplus (w_2 \ominus \rho)$.

The category \mathbb{E} of *initialisations* has worlds as objects, and as homsets $\mathbb{E}(w_1, w_2) := \sum_{\rho : w_1 \rightarrow w_2} \mathbb{H}(w_2 \ominus \rho, w_2)$ whose elements we call *initialisations*. Explicitly, an initialisation $\varepsilon : w_1 \rightarrow w_2$ is a pair $\langle u\varepsilon, \eta_\varepsilon \rangle$ consisting of an injection

$u\varepsilon : w_1 \rightarrow w_2$ and a heaplet η_ε containing the initialisation data required to transition from heap layout w_1 to w_2 (Fig. 4d). This heaplet may contain cyclic dependencies on the newly added locations, or on locations already present in w_1 . Identities $\text{id}_w^\mathbb{E}$ in \mathbb{E} are given by identities in \mathbb{W} , and formally as $\langle \text{id}_w^\mathbb{W}, \mathbb{H}^\emptyset \rangle$, as no initialisation data is required. The composition of two initialisations is given by composing their underlying injections, and appending their initialisation data, suitably promoted to the later world (Fig. 4e).


The collection of (*semantic*) *heaps* now becomes a representable functor $\mathbb{H} : \mathbb{E} \rightarrow \mathbf{Set}$, given at world w by setting

$$\mathbb{H}w := \mathbb{H}(w, w) \cong \mathbb{E}(\emptyset, w)$$

The latter bijection follows from the canonical isomorphism $w \oplus \text{id}_w \cong \emptyset$ in \mathbb{W} . The functorial action of $\mathbb{E}(\emptyset, -)$ then equips \mathbb{H} with a functorial action over initialisations: given a heap $\eta \in \mathbb{H}w_1$ and an initialisation $\varepsilon : w_1 \rightarrow w_2$, promote η to a heaplet in $\mathbb{H}(w_1, w_2)$, and append the initialisation data to create a heap in $\mathbb{H}w_2$. Given $(\ell : c) \in w$, we use projection to define a look-up operation given for any $\eta \in \mathbb{H}w$ by setting $\eta(\ell) := \pi_\ell \eta$, and an update operation, given for any $\eta \in \mathbb{H}w$ and $x \in \llbracket \text{ctype } c \rrbracket w$ by setting

$$\eta[\ell \mapsto x](\ell') := \begin{cases} x & \ell' = \ell \\ \eta(\ell') & \text{otherwise} \end{cases}$$

Finally, given any injection $\rho : w_1 \rightarrow w'$ and initialisation $\varepsilon : w_1 \rightarrow w_2$, the injection $\rho^* u \varepsilon : w' \rightarrow \rho \oplus_{w_1} u \varepsilon$ in fact has an initialisation structure $\rho^* \varepsilon : w' \rightarrow \rho \oplus_{w_1} u \varepsilon$, where the initialisation data $\eta_{\rho^* \varepsilon}$ is given using the isomorphism



$(h \oplus_{w_1} u \varepsilon) \oplus \rho^* u \varepsilon \cong w_2 \oplus u \varepsilon$ Fig. 5: Promoting init. data

and promotion along $u\varepsilon^* \rho$. We denote $u\varepsilon^* \rho$ by $\varepsilon^* \rho$. This process is the semantic counterpart for the promotion of initialisation data we use in the proof of the Totality Theorem 2.

V. THE MONAD

Consider the functor category $\mathbf{E} := [\mathbb{E}, \mathbf{Set}]$, which contains the heaps functor as an object. As we have a forgetful functor $u : \mathbb{E} \rightarrow \mathbb{W}$ projecting out the underlying injection, we obtain a functor $u^* : \mathbf{W} \rightarrow \mathbf{E}$ given by precomposition. In the following, consider the cartesian closed structure of \mathbf{W} as a symmetric monoidal closed structure.

We equip \mathbf{E} with a bi-closed \mathbf{W} -actegory structure:

$$X \odot A := u^* X \times A := (X \circ u) \times A$$

$$A \multimap B := u_* (B^A) \quad X \multimap A := A^{u^* X}$$

$$\alpha_{X,Y,A} : \langle \langle x, y \rangle, a \rangle \mapsto \langle x, \langle y, a \rangle \rangle \quad \lambda_A : \langle *, a \rangle \mapsto a$$

This structure can be alternatively described as transporting the self-enrichment of \mathbf{E} via the cartesian closed structure along the geometric morphism $\langle u^*, u_* \rangle$ from \mathbf{E} to \mathbf{W} .

We can give an explicit end formula for the enrichment:

$$(A \multimap B)w := \int_{w \rightarrow w' \in w \downarrow u} (Bw')^{Aw'}$$

$$\pi_{\rho'} : w_2 \rightarrow w'_2 (A \multimap B) \left(\begin{smallmatrix} w_1 \\ \rho \downarrow \\ w_2 \end{smallmatrix} \right) (\alpha) := \pi_{\rho' \circ \rho} \alpha$$

$$\pi_{\rho : w \rightarrow w'} (\text{curry}_{X,A,B}^- f(x))(a) = f_{w'}(X \rho x, a)$$

$$(\text{eval}_{X,A}^-)_w (\alpha, a) = \pi_{\text{id}_w} \alpha a$$

We can now give an explicit description of the full ground storage monad $T : \mathbf{W} \rightarrow \mathbf{W}$. The action on worlds is

$$(TX)w := \int_{w \rightarrow w' \in w \downarrow u} \left(\int_{w' \rightarrow w'' \in w \downarrow u} (X \circ u)w'' \times \mathbb{H}w'' \right)^{\mathbb{H}w'}$$

This definition is subtle. First, the argument of the coend is covariant in $w' \rightarrow w'' \in w \downarrow u$, and so this coend is an ordinary colimit. We keep the coend notation for its more convenient presentation. The second subtlety is that, while the inner coend is contravariant in w' , the action with respect to which we define the outer end is *different*, and is in fact *covariant* in the object $w \rightarrow w'$ of the comma category $w \downarrow u$. To describe it explicitly, take any morphism $\varepsilon : \langle w'_1, \rho'_1 \rangle \rightarrow \langle w'_2, \rho'_2 \rangle$ in the comma category, i.e., an initialisation $\varepsilon : w'_1 \rightarrow w'_2$ such that $u\varepsilon \circ \rho'_1 = \rho'_2$. Consider a generic element in the coend $\int_{w'_1 \rightarrow w'' \in w \downarrow u} (X \circ u)w'' \times \mathbb{H}w''$ namely some $q_\rho(x, \eta)$, for some $\rho : w'_1 \rightarrow w''$, $x \in Xw''$ and $\eta \in \mathbb{H}w''$. We promote the initialisation ε to an initialisation $\rho^* \varepsilon : w'' \rightarrow \rho \oplus_{w'_1} u \varepsilon$, and map the generic element as follows:

$$q_\rho(x, \eta) \mapsto q_{u\varepsilon^* \rho : w'_2 \rightarrow \rho \oplus_{w'_1} u \varepsilon} (X(u(\rho^* \varepsilon))x, \mathbb{H}(\rho^* \varepsilon)\eta)$$

This subtlety is the main conceptual reason for the decomposition of this monad we present in the next section. We do indeed use the contravariant action of the coend, implicitly below, and explicitly in the next section, to define the hiding/encapsulation operation. This subtlety also appears in the (ordinary) ground storage monad [30] when defining the functorial action of TX . The end gives the functorial action in the full ground setting:

$$(\pi_{\rho_2 : w_2 \rightarrow w'_2} (TX \left(\begin{smallmatrix} w_1 \\ \rho \downarrow \\ w_2 \end{smallmatrix} \right) \alpha))(\eta_2) = \pi_{\rho_2 \circ \rho} (\alpha)(\eta_2)$$

The monadic unit is given by $(\pi_{\rho : w \rightarrow w'} \circ \text{return}_{w'}^T x)\eta := q_{\text{id}_{w'}}(X \rho x, \eta)$. Given any morphism $f : X \rightarrow TY$ in \mathbf{W} and $\alpha \in TX$, define $(\pi_{\rho : w \rightarrow w'} (\alpha \gg f))(\eta') = q_{\rho' \circ \rho'}(y, \eta''')$ where

$$\begin{aligned} (\pi_\rho \alpha)\eta' &= q_{\rho' : w' \rightarrow w''}(x, \eta'') \\ (\pi_{\text{id}_{w''}} \circ f_{w''})(\eta''') &= q_{\rho'' : w'' \rightarrow w'''}(y, \eta''') \end{aligned}$$

Define the strength for any $x \in Xw$ and $\alpha \in TXw$:

$$(\pi_{\rho : w \rightarrow w'} \circ \text{str}_w^T(x, \alpha))\eta' = q_{\rho'}(\langle X(\rho' \circ \rho)x, y \rangle, \eta''')$$

where $(\pi_\rho \alpha)\eta' = q_{\rho'' : w' \rightarrow w''}(y, \eta''')$. Finally, from the other definitions we calculate the functorial action of T on any morphism $f : X \rightarrow Y$: $(\pi_{\rho : w \rightarrow w'} (Tf \alpha))\eta' = q_{\rho'}(f_{w''}(x), \eta''')$ where $(\pi_\rho \alpha)\eta' = q_{\rho' : w' \rightarrow w''}(x, \eta''')$.

On this monad we define the state manipulation operations by setting, for every $\rho : w \rightarrow w'$, two \mathbf{E} -morphisms:

$$\begin{aligned} \text{get}_c &: [\mathbf{ref}_c] \rightarrow T[\mathit{ctype} \ c] \\ &(\pi_\rho \circ \text{get}_c(\ell))(\eta_1) = q_{\text{id}_{w'}}(\eta(\rho(\ell)), \eta) \\ \text{set}_c &: [\mathbf{ref}_c] \times [\mathit{ctype} \ c] \rightarrow T\mathbb{1} \\ &(\pi_\rho \circ \text{set}_c(\ell, a))(\eta_1) = q_{\text{id}_{w'}}(\star, \eta[\rho(\ell) \mapsto [\mathit{ctype} \ c] \ \rho a]) \end{aligned}$$

To define the allocation operation, first define, for every w_0 in \mathbb{W} the functor $\partial_{w_0} : \mathbf{W} \rightarrow \mathbf{W}$ that evaluates at a later world, namely $\partial_{w_0} X := X(- \oplus w_0)$. Using the isomorphism $\varphi : (w \oplus w_0) \ominus \iota_1^\oplus \cong w_0$, we can then define the \mathbf{W} -morphism that constructs an initialisation from given initialisation data:

$$\text{init}_{w_0, w} : \prod_{(\ell:c) \in w_0} \partial_{w_0} [\mathit{ctype} \ c] \ w \rightarrow \mathbb{E}(w, w \oplus w_0) \\ \langle a_\ell \rangle_{\ell \in w_0} \mapsto \langle \iota_1^\oplus \circ \varphi, \langle a_{\varphi \ell} \rangle_\ell \rangle$$

and define:

$$\text{new}_{w_0} : \prod_{(\ell:c) \in w_0} \partial_{w_0} [\mathit{ctype} \ c] \rightarrow T \prod_{(\ell:c) \in w_0} [\mathbf{ref}_c] \\ (\pi_\rho \circ \text{new}_{w_0} \langle a_\ell \rangle) \eta_1 = q_{\varepsilon \star \rho}(\langle \varepsilon \star \rho(\ell) \rangle_{\ell \in w_0}, \mathbb{H}(\rho \star \varepsilon) \eta_1)$$

where $\varepsilon := \text{init} \langle a_\ell \rangle$.

VI. HIDING AND MASKING

We now analyse the functorial action of the inner coend in T 's definition, which is given by a \mathbf{W} -strong monad \underline{P} on \mathbf{E} .


A. The hiding monad

Consider any $A \in \mathbf{E}$, and define for every w :

$$PAw := \int^{w \rightarrow w' \in w \downarrow u} A$$

Given an extension $\rho : w \rightarrow w'$, we think of locations in $w' \ominus \rho$ as *private* locations, and of locations in w as *public* locations.

Example 3. On the left we depict two representatives for a value in $P\mathbb{H}\{\ell : \text{data}\}$. The left representative has no private locations, whereas the right representative has the two private locations $\{\ell_0 : \text{linked_list}, \ell_1 : \text{list_cell}\}$. As we can initialise the right representative from the left, the two representatives are equivalent.



On the right we depict a representative for a value in $P\mathbb{H}\{\ell : \text{linked_list}\}$, whose private locations are given by $\{\ell_0 : \text{list_cell}, \ell_1 : \text{data}\}$. \square

The contravariant action of the coend gives, for every injection $\rho : w_1 \rightarrow w_2$, a function $\text{hide}_\rho : PAw_2 \rightarrow PAw_1$ defined by $q_{\rho' : w_2 \rightarrow w'}(a) \mapsto q_{\rho' \circ \rho}(a)$ (Fig. 6a). The unit is given by $\text{return}_w^P := q_{\text{id}_w} : A \rightarrow PA$ (Fig. 6b). For every $g : A \rightarrow PB$, define $\gg g : PA \rightarrow PB$ by (Fig. 6c)

$$(q_{\rho : w \rightarrow w'}(a) \gg g) := \text{hide}_\rho(g_{w'}(a))$$

For every initialisation $\varepsilon : w_1 \rightarrow w_2$, we derive the functorial action $PA\varepsilon : PAw_1 \rightarrow PAw_2$ (Fig. 6d):

$$PA\varepsilon(q_{\rho : w_1 \rightarrow w'}(a)) := q_{\varepsilon \star \rho}(A(\rho \star \varepsilon)(a))$$

Finally, for every $X \in \mathbf{W}$ and $A \in \mathbf{E}$, define the strength:

$$\text{str}(x, q_{\rho : w \rightarrow w'}(a)) := q_\rho(\langle X \rho x, a \rangle)$$

Proposition 4. The data $\underline{P} = \langle P, \text{return}, \gg, \text{str} \rangle$ define a strong monad over the \mathbf{W} -actegory \mathbf{E} .

As a consequence of Proposition 3, we obtain a monad over \mathbf{W} , and further calculation using the explicit description of \dashv shows this monad is the monad \underline{T} for full ground storage from the previous section.

B. Effect masking

To evaluate the monad T , we show it can mask hidden effects. First, we define a semantic criterion for not leaking any locations. We say that a functor $X \in \mathbf{W}$ is *constant* when, for every $\rho : w \rightarrow w'$ in \mathbb{W} , the function $X\rho$ is a bijection. We say that a world w is *constant* if, for every $(\ell : c) \in w$, the functor $[\mathit{ctype} \ c]$ is constant. When w is constant, so is every sub-world $\hat{w} \rightarrow w$, and the covariant action of the partially applied functor $\mathbb{H}(w, -)$ is a natural isomorphism. Given $\rho : w \rightarrow w'$, we can then project any heap in $\mathbb{H}w'$ to a heap in $\mathbb{H}w$.

Lemma 5. If w is constant, then the monadic unit is invertible $\text{return}_\mathbb{H}^P : \mathbb{H}w \xrightarrow{\cong} P\mathbb{H}w$. In particular, $P\mathbb{H}\emptyset \cong \mathbb{1}$.

We use this result when we prove the Effect Masking Theorem 7, as well as when working with concrete examples.

Lemma 6. For every constant functor $X \in \mathbf{W}$ and every $A \in \mathbf{E}$, the tensorial strength $\text{str}_{X,A}^P$ is an isomorphism.

While technical, this last result is useful, as it plays the role of the mono requirement [23] in $\lambda_{\text{ref}}^\Sigma$'s adequacy proof.

We can now prove that morphisms that do not leak locations are denotationally equivalent to pure values:

Theorem 7 (effect masking). For every pair of constant functors $\Gamma, X \in \mathbf{W}$, every morphism $f : \Gamma \rightarrow TX$ factors uniquely through the monadic unit:

$$\begin{array}{ccc} \Gamma & \xrightarrow{f} & TX \\ & \searrow & \nearrow \\ & X & \end{array} \quad \begin{array}{c} \text{runST } f \\ = \\ \text{return}^T \end{array}$$

The proof of this theorem is conceptually high-level using our decomposition of T as $\mathbb{H} \dashv P(- \odot \mathbb{H})$:

Proof sketch:

As Γ is constant, it suffices to prove the theorem for $\Gamma = \mathbb{1}$. Calculate as in Fig. 7, chasing a generic morphism upwards. \blacksquare

We named the factored morphism $\text{runST } f$ as we can use it to interpret a monadic metalanguage [23] containing a construct similar to Haskell's runST [17].

As usual in functor categories, two different functors may have the same global elements. Thus, even if TX has the same global elements as X , for any constant X , the two functors might differ, for example, for $X = \mathbb{1}$ and the signature from Example 1. The fact that $T\mathbb{1} \not\cong \mathbb{1}$ for this signature will be an immediate consequence of λ_{ref} 's adequacy (see Example 4




Fig. 6: The hiding monad P

$$\begin{array}{c}
\mathbb{1} \longrightarrow \mathbb{H} \multimap P(X \odot \mathbb{H}) \quad \text{in } \mathbf{W} \\
\hline
\mathbb{H} \longrightarrow P(X \odot \mathbb{H}) \quad \text{in } \mathbf{E} \\
\hline
\mathbb{H} \longrightarrow X \odot P\mathbb{H} \quad \text{in } \mathbf{E}, \text{ by Lemma 6} \\
\hline
\mathbb{E}(\emptyset, -) \longrightarrow X \odot P\mathbb{H} \quad \text{in } \mathbf{E} \\
\hline
\mathbb{1} \longrightarrow (X \odot P\mathbb{H})\emptyset \quad \text{in } \mathbf{Set}, \text{ by Yoneda} \\
\hline
\mathbb{1} \longrightarrow X\emptyset \quad \text{in } \mathbf{Set}, \text{ by Lemma 5} \\
\hline
\mathbb{1} \longrightarrow X \quad \text{in } \mathbf{W}
\end{array}$$

Fig. 7: High-level proof of the Effect Masking Theorem 7

below). However, computations that do not assume anything about the heap nor leak references are pure:

Proposition 8. *For every constant $X \in \mathbf{W}$, we have $\text{return}_X^T : X\emptyset \xrightarrow{\cong} TX\emptyset$.*

To see why it holds, note that the initiality of \emptyset in \mathbf{W} means we can bijectively turn an arbitrary element of $TX\emptyset$ into a global element. We then bijectively apply effect masking to get a global element of X , equivalently an element of $X\emptyset$, and further calculation shows the monadic unit induces it.

VII. SEMANTICS FOR FULL GROUND STORAGE

We now return to the λ_{ref} -calculus.

A. Semantics

Fig. 8 presents the interpretation of $\lambda_{\text{ref}}^{\Sigma}$'s types as functors in \mathbf{W} . It extends the interpretation of full ground types by interpreting function types using the exponentials in \mathbf{W} and the full ground storage monad T .

We define two semantic functions, for values in Fig. 9 and for terms in Fig. 10, by induction on typing judgements. These functions have the following types:

$$\begin{array}{l}
[\Gamma \vdash_w v : \tau]^V : \mathbb{W}(w, -) \times [\Gamma] \rightarrow [\tau] \\
[\Gamma \vdash_w t : \tau] : \mathbb{W}(w, -) \times [\Gamma] \rightarrow T[\tau]
\end{array}$$

The two semantic functions relate by $[v] = \text{return}^T \circ [v]^V$ and consequently we omitted from Fig. 10 the definitions implied by this relationship. The two interpretations take as argument a *location environment* ρ , assigning a location in

$$\begin{array}{l}
[\text{ref}_c] w := \{\ell \in \mathbb{L} \mid (\ell : c) \in w\} \quad [\text{ref}_c] \rho(\ell) := \ell \\
[\mathbf{0}] := \mathbf{0} \quad [\tau_1 + \tau_2] := [\tau_1] + [\tau_2] \quad [\mathbf{1}] := \mathbb{1} \\
[\tau_1 * \tau_2] := [\tau_1] \times [\tau_2] \quad [\tau_1 \rightarrow \tau_2] := (T[\tau_2])^{[\tau_1]} \\
[\Gamma] := \prod_{(x:\tau) \in \Gamma} [\tau]
\end{array}$$

Fig. 8: Type semantics

$$[\ell]^V(\rho, e) := \rho(\ell) \quad [x]^V(\rho, e) := e(x)$$

$$[\text{inj}_i^{\tau_1 + \tau_2} v]^V(\rho, e) := \iota_i([v]^V(\rho, e)) \quad [()]^V(\rho, e) := \star$$

$$[(v_1, v_2)]^V(\rho, e) := \langle [v_1]^V(\rho, e), [v_2]^V(\rho, e) \rangle$$

$$[\lambda x : \tau. t]^V(\rho, e) := \text{curry}[t]^V(\rho, e)$$

Fig. 9: Value semantics

the current world for every location in the heap layout the term assumes, and the more standard (*identifier environment* e , assigning a value of the appropriate type to every identifier in the type context.

The definition makes use of the symmetry, dual strength, and the double strength morphisms:

$$\begin{array}{l}
\text{swap} := \langle \pi_2, \pi_1 \rangle : X \times Y \rightarrow Y \times X \\
\text{str}' := T \text{swap} \circ \text{str} \circ \text{swap} : (TX) \times Y \rightarrow T(X \times Y) \\
\text{dstr} := (\gg \text{str}) \circ \text{str}' : (TX) \times (TY) \rightarrow T(X \times Y)
\end{array}$$

The double strength is given explicitly by

$$(\pi_{\rho_1} \text{dstr}(\alpha, \beta))(\eta_1) = q_{\rho_3 \circ \rho_2}(\langle X\rho_3 x, y \rangle, \eta_3)$$

where $(\pi_{\rho_1} \alpha)\eta_1 = q_{\rho_2}(x, \eta_2)$ and $(\pi_{\rho_2 \circ \rho_1} \beta)\eta_2 = q_{\rho_3}(y, \eta_3)$.

The value semantics is standard, with locations interpreted by the location environment. The term semantics is standard. The interpretation of the empty match construct is given by the empty morphism $[] : \mathbb{0} \rightarrow T[\tau]$, as having a morphism $[t] : \mathbb{W}(w, -) \times [\Gamma] \rightarrow \mathbb{0}$ necessitates $\mathbb{W}(w, -) \times [\Gamma]$ is isomorphic to $\mathbb{0}$. The interpretations of the three storage operations use the corresponding three operations for the monad T from Sec. V. There are two steps in defining the

$$\begin{aligned}
\llbracket \mathbf{inj}_i^{\tau_1 + \tau_2} t \rrbracket (\rho, e) &:= T\iota_i(\llbracket t \rrbracket (\rho, e)) \\
\llbracket (t, s) \rrbracket (\rho, e) &:= \text{dstr} \langle \llbracket t \rrbracket (\rho, e), \llbracket s \rrbracket (\rho, e) \rangle \\
\llbracket \mathbf{match} t \mathbf{with} \{ \}^\tau \rrbracket &= \square \\
\llbracket \mathbf{match} t \mathbf{with} \left[\begin{array}{l} \mathbf{inj}_1 x_1 \mapsto s_1 \\ \mathbf{inj}_2 x_2 \mapsto s_2 \end{array} \right] \rrbracket (\rho, e) &:= \llbracket t \rrbracket (\rho, e) \gg \lambda \iota_i a. \\
&\quad \llbracket t_i \rrbracket (\rho, e[x_i \mapsto a]) \\
\llbracket \mathbf{match} t \mathbf{with} (x_1, x_2) \mapsto s \rrbracket (\rho, e) &:= \\
&\quad \text{str}(\langle \rho, e \rangle, \llbracket t \rrbracket (\rho, e)) \gg \llbracket s \rrbracket \\
\llbracket t s \rrbracket (\rho, e) &:= \text{dstr}(\llbracket t \rrbracket (\rho, e), \llbracket s \rrbracket (\rho, e)) \gg \text{eval} \\
\llbracket t := s \rrbracket (\rho, e) &:= \text{dstr}(\llbracket t \rrbracket (\rho, e), \llbracket s \rrbracket (\rho, e)) \gg \text{set} \\
\llbracket ! t \rrbracket (\rho, e) &:= \llbracket t \rrbracket (\rho, e) \gg \text{get}
\end{aligned}$$

$$\llbracket \mathbf{letref} \left[\begin{array}{l} (x_1 : \mathbf{ref}_{c_1}) := v_1, \\ \vdots \\ (x_n : \mathbf{ref}_{c_n}) := v_n \\ \mathbf{in} t \end{array} \right] \rrbracket (\rho, e) := \text{str}(\langle \rho, e \rangle, \text{new}_{\{\ell_1 : c_1, \dots, \ell_n : c_n\}} \langle \llbracket v_i \rrbracket^v (\rho, e[x_i \mapsto \iota_2^\oplus \ell_i]_{i=1}^n) \rangle_{i=1}^n) \gg \llbracket t \rrbracket$$

Fig. 10: Term semantics

semantics of allocation. First, we interpret the initialisation data in the world extended with w_0 , which gives us the appropriate input to the new morphism from Sec. V. The morphism new then returns the newly allocated locations, which we bind to the remainder of the computation.

The semantics satisfies the usual substitution lemma. It is also uniform with respect to the heap layout in the typing judgement. To phrase it, note that every layout extension $w \leq w'$ denotes the world injection given by inclusion.

Lemma 9. *For every layout extension $w \leq w'$ we have:*

$$\begin{aligned}
\llbracket \Gamma \vdash_{w'} v : \tau \rrbracket_{w'}^v (\text{id}_{w'}, -) &= \llbracket \Gamma \vdash_w v : \tau \rrbracket_{w'}^v (w \leq w', -) \\
\llbracket \Gamma \vdash_{w'} t : \tau \rrbracket_{w'} (\text{id}_{w'}, -) &= \llbracket \Gamma \vdash_w t : \tau \rrbracket_{w'} (w \leq w', -)
\end{aligned}$$

This lemma is the semantic counterpart to the monotonicity of the type system.

B. Soundness and adequacy

To phrase our denotational soundness result, we first extend the semantics to heaps. For brevity's sake, we define the following notation for *closed* program phrases $\llbracket \vdash_w v \rrbracket_\star^v := \llbracket v \rrbracket_w^v (\text{id}_w, \star)$, and $\llbracket \vdash_w t \rrbracket_\star := \llbracket t \rrbracket_w (\text{id}_w, \star)$. Next, for every typed heap $\eta \in \mathbf{H}w$ define:

$$\llbracket \eta \rrbracket := \langle \llbracket \vdash_w \eta(\ell) : \text{ctype } c \rrbracket_\star^v \rangle_{(\ell:c) \in w} \in \mathbb{H}w$$

The semantic heap operations are compatible with the syntactic heap operations, in the sense that for every syntactic heap $\eta \in \mathbf{H}w_1$, location $(\ell : c) \in w_1$, and value $\vdash_{w_1} v : \text{ctype } c$ we have: $\llbracket \eta(\ell) \rrbracket_\star^v = \llbracket \eta \rrbracket(\ell)$ and $\llbracket \eta[\ell \mapsto v] \rrbracket = \llbracket \eta \rrbracket[\ell \mapsto \llbracket v \rrbracket_\star^v]$. For allocation, we need to be more careful. Consider any extension $w \leq w_1$, heap $\eta_1 \in \mathbf{H}w_1$, and fresh locations $\#_{w_1} \langle \ell_1, \dots, \ell_n \rangle$. Then let $w' := w \oplus \{\ell_1 : c_1, \dots, \ell_n : c_n\}$ and $w'_1 := w_1 \oplus \{\ell_1 : c_1, \dots, \ell_n : c_n\}$. Then consider any initialisation data $\langle \vdash_{w'} v_i : \text{ctype } c_i \rangle_{i=1}^n$, and let $\varepsilon := \text{init}(\llbracket v_i \rrbracket_\star^v)_{i=1}^n$ be the corresponding initialisation. We then have that $\llbracket \eta_1[\ell_i \mapsto v_i]_{i=1}^n \rrbracket = \mathbb{H}(\langle w \leq w_1 \rangle^{\varepsilon \downarrow} (\llbracket \eta_1 \rrbracket))$.

The operational and denotational semantics agree:

Theorem 10 (soundness). *The operational and denotational semantics agree: for every closed, well-typed term $\vdash_w t : \tau$, extensions $w \leq w' \leq w''$, value $\vdash_{w''} v : \tau$ and heaps $\eta' \in \mathbf{H}w'$ and $\eta'' \in \mathbf{H}w''$, if $\langle t, \eta' \rangle \Downarrow \langle v, \eta'' \rangle$ then*

$$(\pi_{w \leq w'} \llbracket t \rrbracket_\star) \llbracket \eta' \rrbracket = q_{w' \leq w''}(\llbracket v \rrbracket_\star^v, \llbracket \eta'' \rrbracket)$$

The proof is by induction on typing judgements, using the explicit description of T given in Sec. V.

Given two terms $\Gamma \vdash_w t, s : \tau$, recall the set $\mathcal{C}[\Gamma \vdash_w t, s : \tau]$ of contexts plugged with t and s from Subsec. II-C.

Theorem 11 (compositionality). *For every pair of plugged contexts $\Gamma' \vdash_{w'} s_1, s_2 : \tau' \in \mathcal{C}[\Gamma \vdash_w t_1, t_2 : \tau]$, if $\llbracket t_1 \rrbracket = \llbracket t_2 \rrbracket$ then $\llbracket s_1 \rrbracket = \llbracket s_2 \rrbracket$.*

The proof is by induction on contexts, using the fact that the semantics is given compositionally in terms of sub-terms.

Theorem 12 (adequacy). *For all terms $\Gamma \vdash_w t_1, t_2 : \tau$, if $\llbracket t_1 \rrbracket = \llbracket t_2 \rrbracket$ then $\Gamma \vdash_w t_1 \simeq_{\text{ctx}} t_2 : \tau$.*

The proof is standard using the Compositionality and Soundness theorems. In the final step, where the mono requirement is usually used, use Lemma 6 to project out the shared return value of the contexts.

Example 4. As promised, we show $T\mathbb{1} \not\cong \mathbb{1}$ in the signature from Example 1. Consider the two program phrases:

$$\begin{aligned}
&\vdash_{\{\ell_0, \ell_1 : \text{data}\}}(), \mathbf{let} \mathbf{x} = !\ell_0 \mathbf{in} \\
&\quad \ell_0 := !\ell_1; \\
&\quad \ell_1 := \mathbf{x} \quad : \mathbf{1}
\end{aligned}$$

We can distinguish the two phrases by dereferencing ℓ_0 . Had $T\mathbb{1} \cong \mathbb{1}$, they would have equal denotations, and so the result follows from the Adequacy Theorem 12. \square

C. Program equivalences

There are fourteen program equivalences (ordinary) ground reference cells are expected to satisfy [20], [34], and Staton has shown they are Hilbert-Post complete. While we do not check their Hilbert-Post completeness here, we validate them for full ground references. As some equations require locations to be distinct, we use the heap layout assumption to avoid aliasing:

$$\begin{aligned}
&\mathbf{v}_1 : \text{ctype } c_1, \mathbf{v}_2 : \text{ctype } c_2 \vdash_{\{\ell_1 : c_1, \ell_2 : c_2\}} \\
&\quad \ell_1 := \mathbf{v}_1; \quad \ell_2 := \mathbf{v}_2; \\
&\quad \ell_2 := \mathbf{v}_2 \quad \equiv \quad \ell_1 := \mathbf{v}_1 \quad : \mathbf{1}
\end{aligned}$$

VIII. CONCLUSIONS AND FURTHER WORK

We gave a monad for full ground references. An important ingredient was to view the collection of heaps as functorial on initialisations. Using standard developments in enrichment [9], [19], [34], we decomposed it into a monad for hiding transformed with state capabilities to better account for subtleties in the monad’s definition. We gave evidence that the monad is appropriate for modelling reference cells: we showed it yields adequate semantics for the calculus of full ground references, and also validates the equations expected from a local state monad, as well as the effect masking property.

Further work abounds. We would like to use the Effect Masking Theorem 7 to account for Haskell’s `runST` construct [17] by tying the denotational semantics derived from said theorem with a more operational account. We would also like to use our semantics to investigate the combination of polymorphism and reference cells, as the issues motivating ML’s *value restriction* [37] surface with full ground storage.

We would also like to find monads for general storage, and not just full ground references. As it is possible to tie Landin’s knot [16] with general references and implement full recursion, we expect to need to solve some recursive equation to obtain the category of worlds. It might be possible to do so with a traditional recursive domain equation [18], or using step-indexing methods [5].

We would like to find an algebraic presentation for our monad in the style of Plotkin and Power [30], and investigate its completeness [34]. Doing so would allow us to account for effect-dependent program transformations [14]. We would also like to give a simpler description of the monad’s action at (full) ground types. Our decomposition of the full ground references monad differs from existing decompositions for ground storage [21], [34]. Repeating this decomposition in the ordinary ground case would lead to new insights into existing and new models.

ACKNOWLEDGEMENTS

Supported by the ERC grant ‘events causality and symmetry — the next-generation semantics’, EPSRC grants EP/N007387/1 and EP/N023757/1, an EPSRC Studentship, and a Royal Society University Research Fellowship. The authors would like to thank Bob Atkey, Simon Castellan, Pierre Clairambault, Marcelo Fiore, Martin Hyland, Sam Lindley, James McKinna, Paul-André Melliès, Kayvan Memarian, Dominic Mulligan, Jean Pichon-Pharabod, Gordon Plotkin, Uday Reddy, Alex Simpson, Ian Stark, Kasper Svendsen, and Conrad Watt for fruitful discussions and comments.

REFERENCES

- [1] A. J. Ahmed, “Semantics of types for mutable state,” Ph.D. dissertation, Princeton University, 2004.
- [2] N. Benton, M. Hofmann, and V. Nigam, “Abstract effects and proof-relevant logical relations,” in *Proc. POPL*. ACM, 2014, pp. 619–632.
- [3] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann, “Relational semantics for effect-based program transformations with dynamic allocation,” in *Proc. PPDP*. ACM, 2007, pp. 87–96.
- [4] N. Benton and B. Leperchey, “Relational reasoning in a nominal semantics for storage,” in *TLCA*. Springer, 2005, pp. 86–101.
- [5] L. Birkedal, B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang, “Step-indexed Kripke models over recursive worlds,” in *POPL*. ACM, 2011, pp. 119–132.
- [6] L. Birkedal, K. Støvring, and J. Thamsborg, “Realisability semantics of parametric polymorphism, general references and recursive types,” *Math. Structures Comput. Sci.*, vol. 20, no. 4, pp. 655–703, 2010.
- [7] N. Bohr and L. Birkedal, “Relational reasoning for recursive types and references,” in *Proc. APLAS*. Springer, 2006, pp. 79–96.
- [8] D. Dreyer, G. Neis, and L. Birkedal, “The impact of higher-order state and control effects on local relational reasoning,” *J. Funct. Program.*, vol. 22, no. 4-5, pp. 477–528, 2012.
- [9] J. Egger, R. E. Møgelberg, and A. Simpson, “The enriched effect calculus,” *J. Logic Comput.*, vol. 24, no. 3, p. 615, 2014.
- [10] D. R. Ghica, “Semantics of dynamic variables in Algol-like languages,” Queen’s University, Ontario, Canada, Masters Thesis, March 1997.
- [11] R. Gordon and A. Power, “Enrichment through variation,” *J. Pure Appl. Algebra*, vol. 120, no. 2, pp. 167 – 185, 1997.
- [12] M. Hofmann, “Correctness of effect-based program transformations,” in *Formal Logical Methods for System Security and Correctness*, O. Grumberg, T. Nipkow, and C. Pfaller, Eds. IOS Press, 2008, pp. 149–173.
- [13] G. Janelidze and G. Kelly, “A note on actions of a monoidal category,” *Theory and Applications of Categories*, vol. 9, pp. 61–91, 2001.
- [14] O. Kammar and G. D. Plotkin, “Algebraic foundations for effect-dependent optimisations,” in *Proc. POPL*. ACM, 2012, pp. 349–360.
- [15] A. Kock, “Strong functors and monoidal monads,” *Archiv der Mathematik*, vol. 23, no. 1, pp. 113–120, 1972.
- [16] P. J. Landin, “The mechanical evaluation of expressions,” *The Computer Journal*, vol. 6, no. 4, pp. 308–320, 1964.
- [17] J. Launchbury and S. L. P. Jones, “Lazy functional state threads,” in *PLDI*. ACM, 1994, pp. 24–35.
- [18] P. B. Levy, “Possible world semantics for general storage in call-by-value,” in *Proc. CSL*. Springer, 2002, pp. 232–246.
- [19] —, *Call-By-Push-Value: A Functional/Imperative Synthesis*, ser. Semantics Structures in Computation. Springer, 2004, vol. 2.
- [20] —, “Global state considered helpful,” *ENTCS*, vol. 218, pp. 241 – 259, 2008, MFPS XXIV.
- [21] P. Melliès, “Local states in string diagrams,” in *Proc. RTA-TLCA*, ser. LNCS, G. Dowek, Ed., vol. 8560, 2014, pp. 334–348.
- [22] R. E. Møgelberg and S. Staton, “Linear usage of state,” *Logical Methods in Computer Science*, vol. 10, no. 1, 2014.
- [23] E. Moggi, “Computational lambda-calculus and monads,” in *Proc. LICS*. IEEE Computer Society, 1989, pp. 14–23.
- [24] —, “An Abstract View of Programming Languages,” Edinburgh University, Technical Report, 1989. [Online]. Available: <http://www.lfcs.inf.ed.ac.uk/reports/90/ECS-LFCS-90-113/>
- [25] A. S. Murawski and N. Tzevelekos, “Algorithmic games for full ground references,” in *Proc. ICALP*. Springer, 2012, pp. 312–324.
- [26] P. W. O’Hearn, *Scalable Specification and Reasoning: Challenges for Program Logic*. Springer, 2008, pp. 116–133.
- [27] P. W. O’Hearn and R. D. Tennent, “Semantics of local variables,” *Applications of categories in computer science*, pp. 217–238, 1992.
- [28] F. J. Oles, “A category-theoretic approach to the semantics of programming languages,” Ph.D. dissertation, Syracuse University, Aug. 1982.
- [29] A. M. Pitts and I. D. B. Stark, “Observable properties of higher order functions that dynamically create local names, or what’s new?” in *Proc. MFCS*. Springer, 1993, pp. 122–141.
- [30] G. D. Plotkin and J. Power, “Notions of computation determine monads,” in *Proc. FOSSACS*. Springer, 2002, pp. 342–356.
- [31] U. Reddy and H. Yang, “Correctness of data representations involving heap data structures,” *Sci. Comput. Program.*, vol. 50, no. 1-3, 2004.
- [32] A. K. Simpson, “Category-theoretic structure for independence and conditional independence,” Faculty of Mathematics and Physics, University of Ljubljana, preprint, 2017.
- [33] I. Stark, “Names and higher-order functions,” Ph.D. dissertation, University of Cambridge, Dec. 1994, also available as Technical Report 363, University of Cambridge Computer Laboratory.
- [34] S. Staton, “Completeness for algebraic theories of local state,” in *Proc. FOSSACS*, L. Ong, Ed. Springer, 2010, pp. 48–63.
- [35] W. W. Tait, “Intensional interpretations of functionals of finite type I,” *The journal of symbolic logic*, vol. 32, no. 02, pp. 198–212, 1967.
- [36] N. Tzevelekos, “Nominal game semantics,” Ph.D. dissertation, Brasenose College, University of Oxford, 2008.
- [37] A. K. Wright, “Simple imperative polymorphism,” *Lisp and Symbolic Computation*, vol. 8, no. 4, pp. 343–355, 1995.