

An Algebraic Theory for Shared-State Concurrency*

Yotam Dvir¹[0000-0002-6507-3791], Ohad Kammar²[0000-0002-2071-0929], and
Ori Lahav¹[0000-0003-4305-6998]

¹ Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel
yotamdvir@mail.tau.ac.il, orilahav@tau.ac.il

² School of Informatics, University of Edinburgh, Edinburgh, Scotland
ohad.kammar@ed.ac.uk

Keywords: shared state · concurrency · denotational semantics · monads · equational theory · program refinement · program equivalence · compiler transformations · compiler optimisations

Abstract. We present a monadic denotational semantics for a higher-order programming language with shared-state concurrency, i.e. global-state in the presence of interleaving concurrency. Central to our approach is the use of Plotkin and Power’s algebraic effect methodology: designing an equational theory that captures the intended semantics, and proving a monadic representation theorem for it. We use Hyland et al.’s equational theory of resumptions that extends non-deterministic global-state with an operator for yielding to the environment. The representation is based on Brookes-style traces. Based on this representation we define a denotational semantics that is directionally adequate with respect to a standard operational semantics. We use this semantics to justify compiler transformations of interest: redundant access eliminations, each following from a mundane algebraic calculation; while structural transformations follow from reasoning over the monad’s interface.

1 Introduction

Denotational semantics $\llbracket - \rrbracket$ associates every program M with its meaning, i.e. its denotation, $\llbracket M \rrbracket$. A key feature of a denotational semantics is compositionality: the denotation of a program depends only on the denotations of its constituents.

As a concrete example, consider an imperative language that manipulates a memory store $\sigma \in \mathbb{S}$, and denotational semantics for it that associates with each program M a denotation $\llbracket M \rrbracket : \mathbb{S} \rightarrow \mathbb{S}$ modelling how M transforms the store. For example – denoting by $\sigma[a \mapsto v]$ the store that is the same as σ but with

* Supported by the Israel Science Foundation (grant numbers 1566/18 and 814/22) and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 851811); and by a Royal Society University Research Fellowship and Enhancement Award.

its value at a changed to v – we have $\llbracket a := v \rrbracket \sigma = \sigma [a \mapsto v]$. Compositionality manifests in the semantics of program sequencing: $\llbracket M ; N \rrbracket \sigma = \llbracket N \rrbracket (\llbracket M \rrbracket \sigma)$. Thus $\llbracket a := 0 ; a := 1 \rrbracket \sigma = \llbracket a := 1 \rrbracket (\llbracket a := 0 \rrbracket \sigma) = (\sigma [a \mapsto 0]) [a \mapsto 1] = \sigma [a \mapsto 1]$. Incidentally, we also have $\llbracket a := 1 \rrbracket \sigma = \sigma [a \mapsto 1]$, and so $\llbracket a := 0 ; a := 1 \rrbracket = \llbracket a := 1 \rrbracket$.

A desirable property of a denotational semantics $\llbracket - \rrbracket$ is *adequacy*, meaning that $\llbracket M \rrbracket = \llbracket N \rrbracket$ implies that M and N are contextually equivalent: replacing N with M within some larger program does not affect the possible results of executing that program. Contextual equivalence is useful for optimizations: for example, M could have better runtime performance than N . Adequate denotational semantics can justify optimizations without quantifying over all program contexts, serving in this way as a basis for validating compiler optimizations.

Returning to the example above, although $\llbracket a := 0 ; a := 1 \rrbracket = \llbracket a := 1 \rrbracket$, in the presence of concurrency $a := 0 ; a := 1$ and $a := 1$ are not contextually equivalent. For example, if $b := a?$ (read from a and write the result to b) is executed concurrently, it could write 0 to b only with the first programs. Therefore, the semantics we defined is inadequate for a concurrent programming language; differentiating between these programs requires a more sophisticated denotational semantics.

Moreover, the *transformation* $a := 0 ; a := 1 \rightarrow a := 1$ eliminating the first, redundant memory access is valid in the presence of concurrency, even though the programs are not equivalent. Indeed, a compiler applying this simplification within a program would not introduce any additional possible results (though it may eliminate some), and in particular it would conserve the correctness of the program. We would like our semantics to be able to justify such transformations.

This leads us to the concept of *directional adequacy*, a useful refinement of adequacy. Given a partial order \leq on the set of denotations, the denotational semantics is directionally adequate (w.r.t. \leq) if $\llbracket M \rrbracket \leq \llbracket N \rrbracket$ implies that M contextually refines N : replacing N with M within some larger program does not introduce new possible results of executing that program. Thus, directional adequacy can justify the transformation $N \rightarrow M$ even if it is not an equivalence.

In this paper we define directionally-adequate denotational semantics for a higher-order language, subsuming the imperative language above, that justifies the above transformation along with other standard memory access eliminations:

$$\begin{array}{ll}
\ell := w ; \ell := v \rightarrow \ell := v & \text{(write ; write)} \\
\ell := v ; \ell? \rightarrow \ell := v ; v & \text{(write ; read)} \\
\mathbf{let } x = \ell? \mathbf{ in } \ell := x ; x \rightarrow \ell? & \text{(read ; write)} \\
\mathbf{let } x = \ell? \mathbf{ in } \mathbf{let } y = \ell? \mathbf{ in } \langle x, y \rangle \rightarrow \mathbf{let } x = \ell? \mathbf{ in } \langle x, x \rangle & \text{(read ; read)} \\
\ell? ; M \rightarrow M & \text{(irrelevant read)}
\end{array}$$

Other transformations and equivalences this semantics validates are structural ones, such as $\mathbf{if } M \mathbf{ then } N \mathbf{ else } N \cong M ; N$; and concurrency-related ones, such as $(M \parallel N) ; K \rightarrow M ; N ; K$.

None of these transformations are novel. Rather, the contribution of this paper is in the methodology that is used to justify them, fitting shared-state concurrency semantics into a general, uniform model structure. In particular, each memory access elimination is proven correct via a mundane algebraic calculation; and the structural transformations can be justified using simple structural arguments that abstract away the details of our particular semantics.

Methodology The use of monads to study the denotational semantics of effects [30] has proven fruitful, especially with its recent refinement with algebraic operators and equational theories [9, 10, 23, 35, 36]. We follow the algebraic approach to define denotational semantics for a simple higher-order concurrent programming language, using an equational theory extending non-deterministic global-state with a single unary algebraic operator for yielding computation to a concurrently running program [1]. We find a concrete representation of the monad this theory induces based on sets of traces [4, 6], and use it to define a directionally adequate denotational semantics. From this adequacy result we deduce various program transformations via routine calculations.

The advantages of denotational semantics defined with this approach include:

- Uniformity.** Theories are stated using the same general framework. This uniformity means that many theoretical results can be stated in general terms, applying to all theories. Even if a theorem, like adequacy, must be proven separately for each theory, it is likely that a similar proof technique can be used, and experience can guide the construction of the proof.
- Comparability.** Comparing and contrasting theories is convenient due to uniformity [23]. While our language and semantics is very different from Abadi and Plotkin’s [1], the equational theory we obtain easily compares to theirs.
- Modularity.** Since the theories are stated algebraically, using operators and equations, they are amenable to be combined to form larger theories. Some combinations are the result of general theory-combinators, such as the theory of non-deterministic global-state resulting from combining the theory of global-state [34] with the theory of non-determinism [14]. In this combined theory, equations that are provable in each separate theory remain provable. Even if the combination is bespoke, using an algebraic theory breaks down the problem into smaller components [8].
- Abstraction.** The semantics we define for the fragment of our language without shared-state is identical in form to the standard semantics, by using the monad operations. Therefore, any structural transformation proven using these abstractions remains valid in the language extended with shared-state.
- Implementability.** Monads are ubiquitous as a computational device in functional programming languages, such as Haskell. Thus a theory based on a monad may in the future form a bridge to implementation.

Outline The remaining sections are as follows. The next section provides background to the problem and overviews our results in a simplified manner. Then we

dive into the weeds, starting with a succinct presentation of the equational theory and related notions (§3). We then work our way gradually to define the induced monad’s concrete representation (§4). Next we define the denotations using this representation (§5). With everything in place, we present our metatheoretical results and use them to justify program transformations and equivalences (§6). We conclude with a discussion of related work and future prospects (§7).

2 Overview

Our setting is a simple programming language with state. We fix a finite set of *locations* $\mathbb{L} := \{\mathbb{l}_1, \dots, \mathbb{l}_n\}$ and a finite set of (*storable*) *values* $\mathbb{V} := \{v_1, \dots, v_m\}$. A *store* σ is an element of $\mathbb{S} := \mathbb{L} \rightarrow \mathbb{V}$, where $\sigma[\ell \mapsto v]$ is the store that is equal to σ except (perhaps) at ℓ , where it takes the value v . We use subscripts to apply stores to locations, i.e. we write σ_ℓ instead of $\sigma\ell$. In examples we often assume $\mathbb{L} = \{a, b, c\}$ and $\mathbb{V} = \{0, 1\}$, and write stores in matrix notation, e.g. $\begin{pmatrix} a & b & c \\ 1 & 0 & 1 \end{pmatrix}$.

The language is based on a standard extension of Moggi’s [30] computational lambda calculus with products and variants (labelled sums), further extended with three shared-state constructs. Many other constructs are defined using syntactic sugar, such as if-statements and booleans, let-bindings, and program sequencing. The core syntax is presented below (where $n \geq 0$):

$$\begin{aligned}
G &::= (G_1 * \dots * G_n) \mid \{\iota_1 \text{ of } G_1 \mid \dots \mid \iota_n \text{ of } G_n\} && \text{(Ground types)} \\
A, B &::= (A_1 * \dots * A_n) \mid \{\iota_1 \text{ of } A_1 \mid \dots \mid \iota_n \text{ of } A_n\} \mid A \rightarrow B && \text{(Types)} \\
V, W &::= \langle V_1, \dots, V_n \rangle \mid \iota V \mid \lambda x. M && \text{(Values)} \\
M, N &::= x \mid \langle M_1, \dots, M_n \rangle \mid \iota M \mid \lambda x. M \mid MN && \text{(Terms)} \\
&\mid \text{match } M \text{ with } \langle x_1, \dots, x_n \rangle \rightarrow N \\
&\mid \text{case } M \text{ of } \{\iota_1 x_1 \rightarrow N_1 \mid \dots \mid \iota_n x_n \rightarrow N_n\} \\
&\mid M? \mid M := N \mid M \parallel N
\end{aligned}$$

The typing rules for the shared-state constructs appear at the top of Figure 1, where we define $\mathbf{Loc} := \{\mathbb{l}_1 \text{ of } () \mid \dots \mid \mathbb{l}_n \text{ of } ()\}$ and $\mathbf{Val} := \{v_1 \text{ of } () \mid \dots \mid v_m \text{ of } ()\}$.

The language is equipped with a call-by-value, small-step operational semantics $\sigma, M \rightsquigarrow \rho, N$, meaning that the program M executed from store σ progresses to N with store ρ . The operational semantics for the shared-state constructs appear at the bottom of Figure 1. Parallel execution is interpreted via a standard interleaving semantics, ultimately returning the pair of the results of each side, and synchronizing on their completion. The reflexive-transitive closure of \rightsquigarrow is denoted by \rightsquigarrow^* . The operational semantics can be seen in action in Example 2.

2.1 Global-State (for Sequential Computation)

To make this exposition more accessible, we focus on sequential computation before advancing to denotational semantics of concurrent computation. Sequential computations with global state cause two kinds of side-effects: looking a value up

$$\begin{array}{c}
\boxed{\Gamma \vdash M : A} \\
\frac{\Gamma \vdash M : \mathbf{Loc}}{\Gamma \vdash M? : \mathbf{Val}} \quad \frac{\Gamma \vdash M : \mathbf{Loc} \quad \Gamma \vdash N : \mathbf{Val}}{\Gamma \vdash M := N : ()} \quad \frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash M \parallel N : (A * B)} \\
\boxed{\sigma, M \rightsquigarrow \sigma', M'} \\
\frac{\sigma, M \rightsquigarrow \sigma', M'}{\sigma, M? \rightsquigarrow \sigma', M'?} \quad \frac{}{\sigma, \ell? \rightsquigarrow \sigma, \sigma_\ell} \quad \frac{\sigma, M \rightsquigarrow \sigma', M'}{\sigma, M := N \rightsquigarrow \sigma', M' := N} \\
\frac{\sigma, N \rightsquigarrow \sigma', N'}{\sigma, V := N \rightsquigarrow \sigma', V := N'} \quad \frac{}{\sigma, \ell := v \rightsquigarrow \sigma[\ell \mapsto v], \langle \rangle} \quad \frac{\sigma, M \rightsquigarrow \sigma', M'}{\sigma, M \parallel N \rightsquigarrow \sigma', M' \parallel N} \\
\frac{\sigma, N \rightsquigarrow \sigma', N'}{\sigma, M \parallel N \rightsquigarrow \sigma', M \parallel N'} \quad \frac{}{\sigma, V \parallel W \rightsquigarrow \sigma, \langle V, W \rangle}
\end{array}$$

Fig. 1. Typing and operational semantics of the shared-state constructs.

in the store, and updating a value in the store. Plotkin and Power [34] propose a corresponding equational theory with two operators:

Lookup. Suppose $\ell \in \mathbb{L}$, and $(t_v)_{v \in \mathbb{V}}$ is a \mathbb{V} -indexed sequence of terms. Then $L_\ell(t_v)_{v \in \mathbb{V}}$ is a term representing looking the value in ℓ up, calling it v , and continuing the computation with t_v . We write $L_\ell(v. t_v)$ instead of $L_\ell(t_v)_{v \in \mathbb{V}}$.

Update. Suppose $\ell \in \mathbb{L}$, $v \in \mathbb{V}$, and t is a term. Then $U_{\ell, v} t$ is a term representing updating the value in ℓ to v and continuing the computation with t .

The equations of the theory of global-state are generated by taking the closure of the axioms – listed at the top of Figure 2 – under reflexivity, symmetry, transitivity, and substitution. The grayed-out axioms happen to be derivable, and are included for presentation’s sake. The theory of global-state can be used to define adequate denotational semantics for the sequential fragment of our language, obtained by removing concurrent execution (\parallel).

Example 1. Global-state includes the following equation (1) which, when considering sequential programs, represents the program equivalence (2):

$$L_b(v. U_{a, v} L_c(w. U_{a, w} \langle \rangle)) = L_c(w. U_{a, w} \langle \rangle) \quad (1)$$

$$a := b? ; a := c? \cong a := c? \quad (2)$$

2.2 Shared-State

The equivalence (2) from Example 1 fails in the concurrent setting, since the two program can be differentiated by program contexts with concurrency:

Global-State		
UL-det	$\mathsf{U}_{\ell,w} \mathsf{L}_{\ell}(v. x_v) = \mathsf{U}_{\ell,w} x_w$	
UU-last	$\mathsf{U}_{\ell,v} \mathsf{U}_{\ell,w} x = \mathsf{U}_{\ell,w} x$	
LU-noop	$\mathsf{L}_{\ell}(v. \mathsf{U}_{\ell,v} x) = x$	
LL-diag	$\mathsf{L}_{\ell}(v. \mathsf{L}_{\ell}(w. x_{v,w})) = \mathsf{L}_{\ell}(v. x_{v,v})$	
UU-comm	$\mathsf{U}_{\ell,v} \mathsf{U}_{\ell',w} x = \mathsf{U}_{\ell',w} \mathsf{U}_{\ell,v} x$	$\ell \neq \ell'$
LU-comm	$\mathsf{L}_{\ell}(v. \mathsf{U}_{\ell',w} x_v) = \mathsf{U}_{\ell',w} \mathsf{L}_{\ell}(v. x_v)$	$\ell \neq \ell'$
LL-comm	$\mathsf{L}_{\ell}(v. \mathsf{L}_{\ell'}(w. x_{v,w})) = \mathsf{L}_{\ell'}(w. \mathsf{L}_{\ell}(v. x_{v,w}))$	
Non-Determinism		
ND-return	$\bigvee_{i < 1} x = x$	
ND-epi	$\bigvee_{j < \beta} x_j = \bigvee_{i < \alpha} x_{\varphi i}$	surjective $\varphi : \alpha \rightarrow \beta$
ND-join	$\bigvee_{i < \alpha} \bigvee_{j < \beta_i} x_{i,j} = \bigvee_{j < \sum_{i < \alpha} \beta_i} x_{\varphi j}$	bijjective $\varphi : \sum_{i < \alpha} \beta_i \rightarrow \prod_{i < \alpha} \beta_i$
Interaction with Non-Determinism		
ND-L	$\bigvee_{i < \alpha} \mathsf{L}_{\ell}(v. x_{v,i}) = \mathsf{L}_{\ell}(v. \bigvee_{i < \alpha} x_{v,i})$	
ND-U	$\bigvee_{i < \alpha} \mathsf{U}_{\ell,v} x_i = \mathsf{U}_{\ell,v} \bigvee_{i < \alpha} x_i$	
ND-Y	$\bigvee_{i < \alpha} \mathsf{Y} x_i = \mathsf{Y} \bigvee_{i < \alpha} x_i$	

Fig. 2. The axiomatization of the algebraic theory

Example 2. Consider the program context $\Xi[-] = [-] \parallel \mathbf{a}?$, i.e. executing each program in parallel to a thread dereferencing the location \mathbf{a} . Then there is no execution of $\Xi[\mathbf{a} := \mathbf{c}?$] that starts with the store $\begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 1 & 0 & 1 \end{pmatrix}$ and returns $\langle \rangle, 0$, but there is such a execution of $\Xi[\mathbf{a} := \mathbf{b}? ; \mathbf{a} := \mathbf{c}?$]:

$$\begin{aligned} \begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 1 & 0 & 1 \end{pmatrix}, \mathbf{a} := \mathbf{b}? ; \mathbf{a} := \mathbf{c}? \parallel \mathbf{a}? &\rightsquigarrow^* \begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 0 & 0 & 1 \end{pmatrix}, \mathbf{a} := \mathbf{c}? \parallel \mathbf{a}? \rightsquigarrow \\ \begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 0 & 0 & 1 \end{pmatrix}, \mathbf{a} := \mathbf{c}? \parallel 0 &\rightsquigarrow^* \begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 1 & 0 & 1 \end{pmatrix}, \langle \rangle \parallel 0 \rightsquigarrow \begin{pmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 1 & 0 & 1 \end{pmatrix}, \langle \rangle, 0 \end{aligned}$$

Therefore, the aforementioned denotational semantics defined using the theory of global-state cannot be extended with a denotation for (\parallel) while preserving adequacy. More sophistication is needed in the concurrent setting: the denotations, even of sequential programs, must account for environment effects.

We thus extend the global-state theory with a single unary operator:

Yield. Suppose t is a term. Then $\mathsf{Y}t$ is a term. Its intended meaning is to let the environment read and write and then continue with the computation t .

We also need to account for the non-determinism inherent in parallel executions. We do so by extending the resulting theory with operators for finite non-determinism with their equational theory, and further standard equations axiomatizing commutative interaction with the other operators [14]:

Choice. For every $\alpha \in \mathbb{N}$ there is a respective choice operator. Suppose $(t_i)_{i < \alpha}$ is a sequence of terms $t_0, \dots, t_{\alpha-1}$. Then $\bigvee_{i < \alpha} (t_i)_{i < \alpha}$ is a term. Its intended

meaning is to choose $i < \alpha$ non-deterministically and continue with the computation t_i . We write $\bigvee_{i < \alpha} t_i$ instead of $\bigvee_{i < \alpha} (t_i)_{i < \alpha}$; and when $\alpha = 2$ we use infix notation, i.e. instead of $\bigvee_{i < 2} t_i$ we may write $t_0 \vee t_1$.

The axioms of the resulting theory of resumptions RES [1, 14, 30] are listed in Figure 2. The novelty is not in the equational theory, but rather in the way we use it to define denotations. We compare to related work in §7.

2.3 Denotations

In §5 we define denotations of programs, but for the sake of this discussion we simplify, defining the denotation $\llbracket M \rrbracket$ of a program M to be an equivalence class $|t|$ of a particular term t in RES. Our actual denotations defined in §5 will use the concrete representation of the monad (developed in §4.3), similar to how state transformers represent equivalence classes of terms of global-state.

Dereference & Assignment We use the monadic bind $\gg=$ (defined in §4.3), which we can think of as “chaining operators”; and *possible yields* $Y^? t := t \vee Yt$:

$$\begin{aligned} \llbracket M^? \rrbracket &:= \llbracket M \rrbracket \gg= \lambda \ell. |L_\ell(v. Y^? v)| \\ \llbracket M := N \rrbracket &:= \llbracket M \rrbracket \gg= \lambda \ell. \llbracket N \rrbracket \gg= \lambda v. |U_{\ell, v} Y^? \langle \rangle| \end{aligned}$$

The main idea is to intersperse possible yields to block the use of global-state equations such as in (1) and to allow computations to interleave. Although equation (1) still holds in RES, it does not imply the program equivalence (2) because the programs in (2) do not map to the algebraic terms in (1). Rather:

Example 3. The denotations of the programs in Example 1 are:

$$\llbracket \mathbf{a} := \mathbf{c}^? \rrbracket = |L_c(v. Y^? v)| \gg= \lambda v. |U_{a, v} Y^? \langle \rangle| = |L_c(v. Y^? U_{a, v} Y^? \langle \rangle)| \quad (3)$$

$$\llbracket \mathbf{a} := \mathbf{b}^? ; \mathbf{a} := \mathbf{c}^? \rrbracket = |L_b(w. Y^? U_{a, w} Y^? L_c(v. Y^? U_{a, v} Y^? \langle \rangle))| \quad (4)$$

So the denotation (3) looks \mathbf{c} up finding a value v , then possibly yields, then updates \mathbf{a} to v , then possibly yields, and finally returns the empty tuple. The concrete representation (Theorem 1) immediately proves that (3) and (4) are not equal, in contrast to the situation in Example 1.

Parallel Execution Computations interleave using the yield operator. Interleaving execution of programs suggests, for example, the following calculation:

$$\begin{aligned} \llbracket \ell^? \parallel \ell := 0 \rrbracket &= |L_\ell(v. Y^? \llbracket v \parallel \ell := 0 \rrbracket) \vee U_{\ell, 0} Y^? \llbracket \ell^? \parallel \langle \rangle \rrbracket| \\ &= |L_\ell(v. Y^? U_{\ell, 0} Y^? \llbracket v \parallel \langle \rangle \rrbracket) \vee U_{\ell, 0} Y^? L_\ell(v. Y^? \llbracket v \parallel \langle \rangle \rrbracket)| \\ &= |L_\ell(v. Y^? U_{\ell, 0} Y^? \langle v, \langle \rangle \rangle) \vee U_{\ell, 0} Y^? L_\ell(v. Y^? \langle v, \langle \rangle \rangle)| \end{aligned}$$

The problem with the above is that it lacks compositionality: $\llbracket \ell^? \parallel \ell := 0 \rrbracket$ should be defined in terms of $\llbracket \ell^? \rrbracket$ and $\llbracket \ell := 0 \rrbracket$, without referring to the underlying programs. In §4.6 we define a function $(\parallel\parallel)$ such that $\llbracket M \parallel N \rrbracket = \llbracket M \rrbracket \parallel\parallel \llbracket N \rrbracket$. This definition relies on the concrete representation from §4.3.

2.4 Program Transformations

Our main result is directional adequacy (Theorem 5). Under this simplified view it can be stated, in terms of the partial-order on our denotations generated by $|t| \leq |t \vee s|$, as follows: if $\llbracket M \rrbracket \leq \llbracket N \rrbracket$ then the transformation $N \rightarrow M$ is valid in the concurrent setting. The following example illustrates how directional-adequacy can be used to validate program transformations of interest, of the relatively few that are valid in the strong memory-model we consider here.

Example 4. We validate (write ; read) (see also Example 9):

$$\llbracket \ell := v ; \ell? \rrbracket = |U_{\ell,v} Y^? L_{\ell}(w. Y^? w)| \geq |U_{\ell,v} L_{\ell}(w. Y^? w)| = |U_{\ell,v} Y^? v| = \llbracket \ell := v ; v \rrbracket$$

We can similarly validate the other memory access eliminations from §1. By using $Y^?$ rather than Y in the denotations, the cases where the environment is known to not interleave are taken into account explicitly. The relevant global-state equation can then be exploited to eliminate the redundant memory access.

2.5 Caveats and Limitations

Our goal in this work is to fit concurrency semantics on equal footing with other semantic models of computational effects. As a consequence, the proposed model has its fair share of fine print, which we bring to the front:

Memory Model. We study a very strong memory model: sequential consistency. Modern architectures adhere to much weaker memory-models, where further program transformations are valid.

Concurrency Model. Our semantics involves a simple form of concurrency in which threads interleave their computation without restriction, acting on a shared memory store. This is in contrast to a well-established line of work in which models include a causal partial-order in which incomparable events denote “truly” parallel execution [31]. These causal models are showing promise in modelling sophisticated (i.e. weak) shared-state models [7, 16, 17, 25]. We hope further work would fit these causal models into a reliable semantic footing that easily accommodates higher-order structure.

Features. Our analysis lacks many valuable features that appear in related work, such as recursion [1], higher-order state [4], probabilities [41], and infinitely many locations/values. This simplification is intended: we took to reductionism, finding a minimal model still accounting for core features of shared state. The benefit of the algebraic approach is that this model can be modularly combined with other features, hopefully using standard technology such as sum-and-tensor [14], domain-enrichment [27], and functor categories [21, 32, 33, 38]. For example, to support recursion our model may be integrated with one of the known powerdomain theory-combinators [42]. This requires making a semantic-design choice that is orthogonal to shared-state concurrency, each with different trade-offs. We avoid making such choices.

Semantic Precision. The equational theory and denotational semantics based on it leave much room for improvement in terms of precision and abstraction. For example, our denotational model does not support the *introduction* of irrelevant reads, i.e. it does not justify the valid transformation $M \rightarrow \ell?; M$. Indeed, taking $M = \langle \rangle$, we have $\llbracket \ell?; \langle \rangle \rrbracket = |\mathbb{L}_\ell(v. Y^? \langle \rangle)| = |Y^? \langle \rangle| \not\equiv |\langle \rangle| = \llbracket \langle \rangle \rrbracket$. The problem stems from a “counting” issue: even though the value being looked-up in ℓ is discarded, the additional possible-yield remains. We hope further work could address this semantic inaccuracy.

Full Abstraction. Brookes’s seminal work [1,6] defined denotational semantics for concurrency that is fully-abstract, meaning that the converse of adequacy holds: programs that are replaceable in every context have equal denotations. Our semantics is far from being fully-abstract: there is a first-order valid equivalence, $M \cong \ell?; M$, that our semantics does not support. Moreover, we do not include atomic block executions in our language as Brookes did, which was crucial for the proof of full-abstraction. However, even if our model was precise enough to capture the first-order equivalences, and even if we were to include atomic block executions, we still would not expect to obtain full-abstraction, since this result is famously elusive for higher-order languages (see Abramsky’s recent overview on the full-abstraction problem of PCF [2]).

3 Equational Theory

At the foundation of our approach is the equational theory of resumptions RES [1, 14,30] presented in §2, consisting of operators and equational axioms over them. We succinctly fill-in the formal details below, followed by related definitions.

The signature of RES consists of the following parameterized operators. The notation $O : A \langle P \rangle$ means that the arity of the operator O is the set A and it is parameterized over the set P :

$$\begin{array}{llll} \mathbb{L} : \mathbb{V} \langle \mathbb{L} \rangle & \text{lookup} & \mathbb{Y} : \mathbb{1} \langle \mathbb{1} \rangle & \text{yield} \\ \mathbb{U} : \mathbb{1} \langle \mathbb{L} \times \mathbb{V} \rangle & \text{update} & \bigvee_{\alpha} : \alpha \langle \mathbb{1} \rangle & \text{non-deterministic choice for every } \alpha \in \mathbb{N} \end{array}$$

From now on, whenever we refer to an *operator*, we mean an operator of RES. We denote the set of terms freely generated by the signature over X by $\text{Term}X$.

Figure 2 lists the axioms of RES, classified as follows: an axiomatization of the equational theory of global-state [34]; the standard axiomatization of non-determinism; and an axiomatization of the commutative interaction of non-determinism with the other operators [13] via the tensor [14].

A RES-*algebra* \mathcal{A} consists of a carrier set $\underline{\mathcal{A}}$ together with interpretations $\tilde{O}^{\mathcal{A}} : \underline{\mathcal{A}}^A \times P \rightarrow \underline{\mathcal{A}}$ for each operator $O : A \langle P \rangle$. We elide the superscript if it is clear from context. For a set X , a RES-*algebra on* X consists of a RES-algebra \mathcal{A} and a function $\text{env} : X \rightarrow \underline{\mathcal{A}}$; which extends to $\text{eval} : \text{Term}X \rightarrow \underline{\mathcal{A}}$ homomorphically along the inclusion $X \hookrightarrow \text{Term}X$. A RES-*model on* X is a RES-algebra on X that satisfies each axiom of RES, i.e. the same element of $\underline{\mathcal{A}}$ is obtained by applying eval to either side of the axiom.

In the following, we abbreviate using $\bar{L}(\sigma. t_\sigma) := L_{1_1}(v_1 \dots L_{1_n}(v_n \cdot t_{\lambda_{1_i} \cdot v_i}))$ and $\bar{U}_\sigma t := U_{1_1, \sigma_{1_1}} \dots U_{1_n, \sigma_{1_n}} t$, in addition to $Y^? t := t \vee Y t$ that we saw in §2.3. For example, $\bar{L}\left(\sigma. \bar{U}\left(\begin{smallmatrix} a & b & c \\ 1 & 0 & \sigma_b \end{smallmatrix}\right) Y^? t_{\sigma_c}\right) = L_a(v. L_b(w. L_c(u. U_{a,1} U_{b,0} U_{c,w}(t_u \vee Y t_u))))$. We use similar shorthands with interpretations of operators as well.

4 A Monad for Shared-State

The next part in our denotational semantics is a monad whose elements represent equivalence classes of RES. The monad can be obtained via a universal construction [28] (by quotienting the terms by the equational theory), but a concrete representation is crucial to reason about it; for example, to show that two denotations are different.

4.1 Difficulty of Term Normalization

To motivate the definitions building up to this concrete representation, we first find a representative for each equivalence class in $\text{Term}X/\text{RES}$, by taking an arbitrary $t \in \text{Term}X$ and transforming it via equations in RES to a particular form – a normal form – such that there is only one term of this form equal to t .

Consider an algebraic term $t \in \text{Term}X$. Using LU-noop once for each location, a sequence of LU-comm, and ND-return, we find that $t = \bar{L}(\sigma. \bigvee_{i < 1} \bar{U}_\sigma t)$. Note:

$$\bar{U}_\sigma L_\ell(v. s_v) \stackrel{\text{RES}}{=} \bar{U}_\sigma s_{\sigma_\ell} \quad \bar{U}_\sigma U_{\ell, v} s \stackrel{\text{RES}}{=} \bar{U}_{\sigma[\ell \rightarrow v]} s \quad \bar{U}_\sigma \bigvee_{i < \alpha} s_i \stackrel{\text{RES}}{=} \bigvee_{i < \alpha} \bar{U}_\sigma s_i$$

By applying these equalities left-to-right as long as possible, and applying ND-join and ND-epi to rearrange the sums, we find that t is equal to a term of the form $\bar{L}(\sigma. \bigvee_{i < \alpha_\sigma} \bar{U}_{\rho_{i, \sigma}} s_{i, \sigma})$, where $s_{i, \sigma}$ is either in X or is of the form $Y s'_{i, \sigma}$.

For every σ , we can rearrange the sum according to common prefixes, thus we find that t is equal to a term of the form: $\bar{L}(\sigma. \bigvee_{\rho \in \mathbb{S}} \bar{U}_\rho \bigvee_{j < \alpha_{\rho, \sigma}} s_{j, \rho, \sigma})$ where $s_{i, \rho, \sigma}$ is either in X or is of the form $Y s'_{i, \rho, \sigma}$ (we can take $\alpha_{\rho, \sigma} = 0$ when the prefix \bar{U}_ρ did not appear in t). For every ρ , we can rearrange to obtain the form:

$$\bar{L}\left(\sigma. \bigvee_{\rho \in \mathbb{S}} \bar{U}_\rho \left(Y r_{\rho, \sigma} \vee \bigvee_{j < \beta_{\rho, \sigma}} x_{j, \rho, \sigma} \right) \right) \quad (5)$$

This is not yet a normal form, which to obtain would require recursively applying this procedure to $r_{\rho, \sigma}$ and propagating empty choice operators outward. Were we to continue in this way to find a normal form, we would still need to prove uniqueness and completeness. One standard way to achieve this is to show that this procedure equates the sides of every axiom and respects the deduction rules of equational logic. This requires a careful proof-theoretic analysis of this normalization procedure. Instead, we take a model-theoretic approach, akin to normalization-by-evaluation, constructing for every set a concrete representation of the free RES-model over it. This representation is based on finite sets of traces.

4.2 Traces

Brookes [6] defined a trace to be a non-empty sequence of transitions, where a *transition* is a pairs of stores; e.g. $\langle\langle \begin{smallmatrix} a & b & c \\ 1 & 0 & 1 \end{smallmatrix}, \begin{smallmatrix} a & b & c \\ 0 & 0 & 1 \end{smallmatrix} \rangle\rangle \langle\langle \begin{smallmatrix} a & b & c \\ 0 & 1 & 1 \end{smallmatrix}, \begin{smallmatrix} a & b & c \\ 1 & 0 & 1 \end{smallmatrix} \rangle\rangle$. Brookes used traces to define a denotational semantics for an imperative concurrent programming language. In Brookes's semantics, traces denote interrupted executions, where each transition corresponds to an uninterrupted sequence of computation steps that starts with the first store and end with the second store. The breaks between transitions are where the computation yields to the environment.

The concept was adapted by many, including Benton et al. [4], to define denotational semantics for a functional language, where they have added an additional value at the end of the sequence to refer to the value the computation returns. A *trace* in this paper will refer to this concept: a non-empty sequence of transitions followed by an additional return value. If we wish to specify X as the set of the return values, we will call it an X -trace. For example, if $x \in X$ then $\langle\langle \begin{smallmatrix} a & b & c \\ 1 & 0 & 1 \end{smallmatrix}, \begin{smallmatrix} a & b & c \\ 0 & 0 & 1 \end{smallmatrix} \rangle\rangle \langle\langle \begin{smallmatrix} a & b & c \\ 0 & 1 & 1 \end{smallmatrix}, \begin{smallmatrix} a & b & c \\ 1 & 0 & 1 \end{smallmatrix} \rangle\rangle x$ is an X -trace.

For sets Y, Z , we denote by Y^* the set of sequences over Y , by Y^+ the set of non-empty sequences over Y , and $Y \cdot Z := \{yz \mid y \in Y, z \in Z\}$. That is, (\cdot) is just notation for (\times) which suggests that the elements of the set are written in sequence, eliding the tuple notation. Thus $(\mathbb{S} \times \mathbb{S})^+ \cdot X$ is the set of X -traces.

Following our discussion in §4.1, our representation will use finite sets of traces instead of the algebraic syntax. In particular, the form we have found in (5) suggests the recursive definition:

$$\text{rept} := \{\langle\sigma, \rho\rangle \tau \mid \sigma, \rho \in \mathbb{S}, \tau \in \text{rep } r_{\rho, \sigma}\} \cup \{\langle\sigma, \rho\rangle x_{j, \rho, \sigma} \mid \sigma, \rho \in \mathbb{S}, j < \beta_{\rho, \sigma}\} \quad (6)$$

The model-theoretic approach we use below obviates the need for the syntactic manipulation that leads to the form in (5) as part of finding the representation. In the model definition, eval will play the role of rep .

4.3 Model Definition

We represent elements of $\text{Term}X/\text{RES}$ by $\underline{\mathcal{T}}X := \mathcal{P}_{\text{fin}}((\mathbb{S} \times \mathbb{S})^+ \cdot X)$, i.e. finite sets of X -traces. We equip $\underline{\mathcal{T}}X$ with a RES-algebra structure $\mathcal{F}X$:

$$\begin{aligned} \tilde{\mathcal{L}}_{\ell}(v. P_v) &:= \{\langle\sigma, \rho\rangle \tau \mid \langle\sigma, \rho\rangle \tau \in P_{\sigma_{\ell}}\} & \tilde{\bigvee}_{i < \alpha} P_i &:= \bigcup_{i < \alpha} P_i \\ \tilde{\mathcal{U}}_{\ell, v} P &:= \{\langle\sigma, \rho\rangle \tau \mid \langle\sigma[\ell \mapsto v], \rho\rangle \tau \in P\} & \tilde{\mathcal{Y}}P &:= \{\langle\sigma, \sigma\rangle \tau \mid \sigma \in \mathbb{S}, \tau \in P\} \end{aligned}$$

We further equip it with $\text{env } x := \{\langle\sigma, \sigma\rangle x \mid \sigma \in \mathbb{S}\}$ to make it a RES-algebra over X . We denote $\text{env } x$ by $\text{return } x$, or \tilde{x} for shorthand. This RES-algebra is in fact a RES-model over X by virtue of satisfying the axioms of RES:

Example 5. We verify that $\langle\mathcal{F}X, \text{return}\rangle$ satisfies the axiom LU-noop:

$$\begin{aligned} \text{eval}(\mathbb{L}_{\ell}(v. \mathbb{U}_{\ell, v} x)) &= \tilde{\mathcal{L}}_{\ell}(v. \tilde{\mathcal{U}}_{\ell, v} \tilde{x}) = \tilde{\mathcal{L}}_{\ell}(v. \{\langle\sigma, \sigma[\ell \mapsto v]\rangle x \mid \sigma \in \mathbb{S}\}) \\ &= \{\langle\sigma, \sigma[\ell \mapsto \sigma_{\ell}]\rangle x \mid \sigma \in \mathbb{S}\} = \{\langle\sigma, \sigma\rangle x \mid \sigma \in \mathbb{S}\} = \tilde{x} = \text{eval } x \end{aligned}$$

4.4 Correspondence to Non-Deterministic Global-State

The theory of non-deterministic global-state (the fragment of RES excluding Y) admits a concrete representation using non-deterministic state transformers $\mathbb{S} \rightarrow \mathcal{P}_{\text{fin}}(\mathbb{S}X)$ [14]. This representation corresponds to the one we have defined in an interesting way. Namely, there is a bijection between $\underline{\mathcal{T}}X$ and the set of functions mapping stores to finite sets of X -traces-with-the-first-store-removed:

$$\begin{aligned} \lambda P \in \underline{\mathcal{T}}X. \lambda \sigma \in \mathbb{S}. \{ \rho \tau \in \mathbb{S} \cdot (\mathbb{S} \times \mathbb{S})^* \cdot X \mid \langle \sigma, \rho \rangle \tau \in P \} \\ \lambda \psi \in \mathbb{S} \rightarrow \mathcal{P}_{\text{fin}}(\mathbb{S} \cdot (\mathbb{S} \times \mathbb{S})^* \cdot X). \bigcup_{\sigma \in \mathbb{S}} \{ \langle \sigma, \rho \rangle \tau \in (\mathbb{S} \times \mathbb{S})^+ \cdot X \mid \rho \tau \in \psi \sigma \} \end{aligned}$$

Implicitly identifying the two, the model from §4.3 can be defined using formulas that look exactly like the non-deterministic global-state ones:

$$\begin{aligned} \tilde{L}_\ell(v. P_v) &:= \lambda \sigma. P_{\sigma_\ell} \sigma & \tilde{\bigvee}_{i < \alpha} P_i &:= \lambda \sigma. \bigcup_{i < \alpha} P_i \sigma \\ \tilde{U}_{\ell, v} P &:= \lambda \sigma. P(\sigma[\ell \mapsto v]) & \tilde{x} &:= \lambda \sigma. \{ \sigma x \} \end{aligned}$$

However, these are not the same formulas – they are defined for different elements (sets of traces as opposed to non-deterministic state transformers).

Using this identification for the yield operator, we obtained the definition $\tilde{Y}P := \lambda \sigma. \{ \sigma \tau \mid \tau \in P \}$, which we understand as “the thread does not modify the state, then allows the environment to intervene, and then continues as before.”

4.5 Representation Theorem

The model $\langle \mathcal{F}X, \text{return} \rangle$ defined in §4.3 *represents* $\text{Term}X/\text{RES}$ because – according to the representation theorem – this is a *free* RES-model on X , and therefore equivalent to the model of equivalence classes we used in §2 or the model of syntactic normal forms to which we have alluded in §4.1.

To prove that the model is free we first equip the family of sets $\underline{\mathcal{T}}$ with a monad structure. For every RES-model \mathcal{A} and function $f : X \rightarrow \underline{\mathcal{A}}$, define $\llbracket f \rrbracket : \underline{\mathcal{T}}X \rightarrow \underline{\mathcal{A}}$, the homomorphic extension of f along return , recursively; where $R_P^{(\sigma, \rho)} := \{ \tau \in (\mathbb{S} \times \mathbb{S})^+ \cdot X \mid \langle \sigma, \rho \rangle \tau \in P \}$ and $X_{P, f}^{(\sigma, \rho)} := \tilde{\bigvee}_{\langle \sigma, \rho \rangle x \in P} f x$:

$$\emptyset \llbracket f \rrbracket := \tilde{\bigvee}_0^{\mathcal{A}} \emptyset \quad P \llbracket f \rrbracket := \tilde{L}^{\mathcal{A}} \left(\sigma. \tilde{\bigvee}_{\rho \in \mathbb{S}}^{\mathcal{A}} \tilde{U}_\rho^{\mathcal{A}} \left(\tilde{Y}^{\mathcal{A}} \left(R_P^{(\sigma, \rho)} \llbracket f \rrbracket \right) \tilde{\bigvee}^{\mathcal{A}} X_{P, f}^{(\sigma, \rho)} \right) \right)$$

A simpler definition is available when there exists a set Y such that $\mathcal{A} = \mathcal{F}Y$:

$$\emptyset \llbracket f \rrbracket := \emptyset \quad P \llbracket f \rrbracket := \{ \alpha \langle \sigma, \varsigma \rangle \tau \mid \exists \rho. \alpha \langle \sigma, \rho \rangle x \in P \wedge \langle \rho, \varsigma \rangle \tau \in f x \}$$

The recursion is well-founded since $R_P^{(\sigma, \rho)}$ is smaller than P when measured by the length of the longest trace in the set.

Thus we have our monad structure $\mathcal{T} := \langle \underline{\mathcal{T}}, \text{return}, \llbracket _ \rrbracket \rangle$. We show it is induced by the aforementioned family of free RES-models:

Theorem 1 (Representation for shared-state). *The pair $\langle \mathcal{F}X, \text{return} \rangle$ is a free RES-model on X : for every RES-model \mathcal{A} and $f : X \rightarrow \underline{\mathcal{A}}$, the function $- \gg f : \underline{\mathcal{T}}X \rightarrow \underline{\mathcal{A}}$ is the unique homomorphism g satisfying $f = g \circ \text{return}$.*

As a direct consequence:

Corollary 1 (Model is sound and complete). *Terms over X are equal in RES iff they have the same representation in $\langle \mathcal{F}X, \text{return} \rangle$.*

4.6 Synchronization

To define the denotational semantics of (\parallel) in §5, we will define a corresponding function $(\parallel\parallel)$ on elements of the monad. To this end we first define the *trace synchronization*, an inductively defined relation $\tau_1 \parallel \tau_2 \Rightarrow \tau$ presented below, that relates $\tau_i \in (\mathbb{S} \times \mathbb{S})^+ \cdot X_i$ and $\tau \in (\mathbb{S} \times \mathbb{S})^+ \cdot (X_1 \times X_2)$, representing the fact that τ_1 and τ_2 can synchronize to form τ :

$$\boxed{\tau \parallel \pi \Rightarrow \omega} \quad \frac{}{\langle \sigma, \rho \rangle x \parallel \langle \rho, \varsigma \rangle \beta y \Rightarrow \langle \sigma, \varsigma \rangle \beta \langle x, y \rangle} \text{(VAR-LEFT)}$$

$$\frac{\tau \parallel \pi \Rightarrow \omega}{\langle \sigma, \rho \rangle \tau \parallel \pi \Rightarrow \langle \sigma, \rho \rangle \omega} \text{(BRK-LEFT)} \quad \frac{\tau \parallel \pi \Rightarrow \langle \rho, \varsigma \rangle \omega}{\langle \sigma, \rho \rangle \tau \parallel \pi \Rightarrow \langle \sigma, \varsigma \rangle \omega} \text{(SEQ-LEFT)}$$

Symmetrically: (VAR-RIGHT) (BRK-RIGHT) (SEQ-RIGHT)

One way to understand these rules is to concentrate on the first transition on the left trace $\tau_1 = \langle \sigma, \rho \rangle \tau'_1$; the right-sided rules are treated symmetrically. If the first transition is also the last, i.e. $\tau'_1 \in X$, then ρ must be the initial store when the execution continues (recall that only a break *between transitions* reflects a yield to the environment). This is why VAR-LEFT combines the transitions as it does. The value in τ_3 is the pair of the values in τ_1 and τ_2 , reflecting the operational semantics of (\parallel) returning the pair of the results. If, on the other hand, the first transition is not the last, then we may combine the transition with the continuation of the computation (SEQ-LEFT), or we may not (BRK-LEFT). The first option means the yield was used-up in this synchronization; while in the second option yield remains available to ambient synchronizations.

From this relation we derive the *semantic synchronization* function:

$$(\parallel\parallel) : \underline{\mathcal{T}}X \times \underline{\mathcal{T}}Y \rightarrow \underline{\mathcal{T}}(X \times Y) \quad P \parallel\parallel Q := \left\{ \omega \mid \exists \tau \in P, \pi \in Q. \tau \parallel \pi \Rightarrow \omega \right\}$$

Example 6. For $\sigma, \rho \in \mathbb{S}$, we may synchronize $\langle \sigma, \rho \rangle \langle \rho, \sigma \rangle \langle \rangle$ and $\langle \rho, \rho \rangle 0$ so:

$$\frac{\frac{}{\langle \rho, \sigma \rangle \langle \rangle \parallel \langle \rho, \rho \rangle 0 \Rightarrow \langle \rho, \sigma \rangle \langle \langle \rangle, 0 \rangle} \text{VAR-RIGHT}}{\langle \sigma, \rho \rangle \langle \rho, \sigma \rangle \langle \rangle \parallel \langle \rho, \rho \rangle 0 \Rightarrow \langle \sigma, \sigma \rangle \langle \langle \rangle, 0 \rangle} \text{SEQ-LEFT}}$$

Therefore, if $\langle \sigma, \rho \rangle \langle \rho, \sigma \rangle \langle \rangle \in P$ and $\langle \rho, \rho \rangle 0 \in Q$, then $\langle \sigma, \sigma \rangle \langle \langle \rangle, 0 \rangle \in P \parallel Q$.

The use of SEQ-LEFT was possible since the stores happen to match, resulting in a trace that does not allow the environment to interfere. By using BRK-LEFT we could find a different synchronization, one that does yield to the environment.

5 Denotational Semantics

With the monad in place, denotations of types and contexts are standard [30]:

$$\begin{aligned} \llbracket (A_1 * \dots * A_n) \rrbracket &:= \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket & \llbracket A \multimap B \rrbracket &:= \llbracket A \rrbracket \rightarrow \mathcal{T} \llbracket B \rrbracket \\ \llbracket \{\iota_1 \text{ of } A_1 \mid \dots \mid \iota_n \text{ of } A_n\} \rrbracket &:= \bigcup_{i=1}^n \{\iota_i\} \times \llbracket A_i \rrbracket & \llbracket \Gamma \rrbracket &:= \prod_{(x:A) \in \Gamma} \llbracket A \rrbracket \end{aligned}$$

Define the extension of $\gamma \in \llbracket \Gamma \rrbracket$ to $\gamma[\mathbf{x} \mapsto y] \in \llbracket \Gamma, \mathbf{x} : A \rrbracket$ by $\gamma[\mathbf{x} \mapsto y] \mathbf{x} := y$.

On the above we base two kinds of denotations for programs $\Gamma \vdash M : A$:

Computational. $\llbracket M \rrbracket^c : \llbracket \Gamma \rrbracket \rightarrow \mathcal{T} \llbracket A \rrbracket$. When Γ is empty we may write $\llbracket M \rrbracket^c$ instead of $\llbracket M \rrbracket^c \langle \rangle$. We write $\llbracket M \rrbracket^c \subseteq \llbracket N \rrbracket^c$ for $\forall \gamma \in \llbracket \Gamma \rrbracket. \llbracket M \rrbracket^c \gamma \subseteq \llbracket N \rrbracket^c \gamma$.

Valuational. $\llbracket V \rrbracket^v : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$ defined solely for values, and satisfying $\llbracket V \rrbracket^c \gamma = \text{return}(\llbracket V \rrbracket^v \gamma)$. When Γ is empty we may write $\llbracket V \rrbracket^v$ instead of $\llbracket V \rrbracket^v \langle \rangle$; and if furthermore A is a ground type, we may write V instead of $\llbracket V \rrbracket^v$, noting that the restriction of $\llbracket - \rrbracket^v$ to closed programs of ground type is a bijection.

Most denotations of programs are standard as well, such as:

$$\begin{aligned} \llbracket \langle \rangle \rrbracket^v \gamma &:= \langle \rangle & \llbracket \lambda \mathbf{x}. M \rrbracket^v \gamma &:= \lambda y. \llbracket M \rrbracket^c \gamma[\mathbf{x} \mapsto y] \\ \llbracket \mathbf{x} \rrbracket^v \gamma &:= \gamma \mathbf{x} & \llbracket NM \rrbracket^c \gamma &:= \llbracket N \rrbracket^c \gamma \gg \llbracket M \rrbracket^c \gamma \gg f. \llbracket M \rrbracket^c \gamma \gg f \end{aligned}$$

The denotations of the state effects allow the environment to intervene:

$$\begin{aligned} \llbracket M? \rrbracket^c \gamma &:= \llbracket M \rrbracket^c \gamma \gg \lambda \ell. \tilde{L}_\ell(v. \tilde{Y}^? \tilde{v}) \\ \llbracket M := N \rrbracket^c \gamma &:= \llbracket M \rrbracket^c \gamma \gg \lambda \ell. \llbracket N \rrbracket^c \gamma \gg \lambda v. \tilde{U}_{\ell, v} \tilde{Y}^? \tilde{\langle \rangle} \\ \llbracket M \parallel N \rrbracket^c \gamma &:= \llbracket M \rrbracket^c \gamma \parallel \llbracket N \rrbracket^c \gamma \end{aligned}$$

Example 7. With the definitions above, we can state the denotations from Example 3 precisely. For instance, (4) becomes:

$$\llbracket \mathbf{a} := \mathbf{b}? ; \mathbf{a} := \mathbf{c}? \rrbracket^c = \tilde{L}_b \left(w. \tilde{Y}^? \tilde{U}_{a, w} \tilde{Y}^? \tilde{L}_c \left(v. \tilde{Y}^? \tilde{U}_{a, v} \tilde{Y}^? \tilde{\langle \rangle} \right) \right)$$

Example 8. We can explain the execution of $\mathbf{a} := \mathbf{b}? ; \mathbf{a} := \mathbf{c}? \parallel \mathbf{a}?$ from Example 2 in denotational terms. First we find traces to synchronize:

$$\begin{aligned} \langle \langle \begin{pmatrix} a & b & c \\ 1 & 0 & 1 \end{pmatrix}, \begin{pmatrix} a & b & c \\ 0 & 0 & 1 \end{pmatrix} \rangle \langle \langle \begin{pmatrix} a & b & c \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} a & b & c \\ 1 & 0 & 1 \end{pmatrix} \rangle \rangle \langle \rangle \in \llbracket \mathbf{a} := \mathbf{b}? ; \mathbf{a} := \mathbf{c}? \rrbracket^c \\ \langle \langle \begin{pmatrix} a & b & c \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} a & b & c \\ 0 & 0 & 1 \end{pmatrix} \rangle \rangle 0 \in \llbracket \mathbf{a}? \rrbracket^c \end{aligned}$$

Following from the derivation in Example 6 with $\sigma = \begin{pmatrix} a & b & c \\ 1 & 0 & 1 \end{pmatrix}$ and $\rho = \begin{pmatrix} a & b & c \\ 0 & 0 & 1 \end{pmatrix}$:

$$\langle \langle \begin{pmatrix} a & b & c \\ 1 & 0 & 1 \end{pmatrix}, \begin{pmatrix} a & b & c \\ 1 & 0 & 1 \end{pmatrix} \rangle \langle 0, \langle \rangle \rangle \in \llbracket \mathbf{a} := \mathbf{b}? ; \mathbf{a} := \mathbf{c}? \parallel \mathbf{a}? \rrbracket^c$$

This trace corresponds to the (uninterrupted) execution presented in Example 2.

6 Metatheoretical Results

First we find that the single-transition traces in the denotation of a program account for the possible executions of that program:

Theorem 2 (Soundness). *If $\sigma, M \rightsquigarrow^* \rho, V$ then $\langle \sigma, \rho \rangle \llbracket V \rrbracket^v \in \llbracket M \rrbracket^c$.*

For the proof, omitted for brevity, we instrument the operational-semantics with *actions*, elements of $\{\mathsf{U}_{\ell, v}, \mathsf{L}_{\ell}, \varepsilon\}$ signifying the effect caused by the step, to analyse the change to the denotation of the program as it runs.

Working our way up to the fundamental lemma, we define a unary logical relation: functions $\mathcal{V}(-)$ and $\mathcal{E}(-)$ from types to sets of closed programs by mutual recursion. Specifically, $\mathcal{V}(A)$ is a set of closed values of type A , and $\mathcal{E}(A)$ is a set of closed programs of type A . The definition of $\mathcal{V}(-)$ is standard:

$$\begin{aligned} \mathcal{V}(A \rightarrow B) &:= \{\lambda \mathbf{x}. M \mid \forall V \in \mathcal{V}(A). M[V/\mathbf{x}] \in \mathcal{E}(B)\} \\ \mathcal{V}((A_1 * \dots * A_n)) &:= \{\langle V_1, \dots, V_n \rangle \mid \forall i. V_i \in \mathcal{V}(A_i)\} \\ \mathcal{V}(\{\iota_1 \text{ of } A_1 \mid \dots \mid \iota_n \text{ of } A_n\}) &:= \cup_i \{\iota_i V \mid V \in \mathcal{V}(A_i)\} \end{aligned}$$

The definition of $\mathcal{E}(-)$ is also standard in that it ensures programs in $\mathcal{E}(A)$ compute to values in $\mathcal{V}(A)$, but bespoke in its requirement about *how* they compute. This requirement is based on the way traces specify interrupted executions, a notion we have discussed in §4.2 and now make precise. For a non-empty sequence of transitions $\alpha = \langle \sigma_1, \rho_1 \rangle \dots \langle \sigma_m, \rho_m \rangle$ we write $M \xrightarrow{\alpha} N$ when there exist $M = M_1, M_2 \dots M_m, M_{m+1} = N$ such that $\sigma_i, M_i \rightsquigarrow^* \rho_i, M_{i+1}$ for all $i \in \{1, \dots, m\}$. We write $M \xrightarrow{\alpha x} V$ when $M \xrightarrow{\alpha} V$ and $\llbracket V \rrbracket^v = x$. We now define:

$$\mathcal{E}(A) := \left\{ M \in \cdot \vdash A \mid \forall \tau \in \llbracket M \rrbracket^c \exists V \in \mathcal{V}(A). M \xrightarrow{\tau} V \right\}$$

The last component needed is the function $\mathcal{G}(-)$ from typing contexts to sets of program substitutions: $\mathcal{G}(\Gamma) := \{\Theta \mid \forall (\mathbf{x} : A) \in \Gamma. \Theta \mathbf{x} \in \mathcal{V}(A)\}$. The *semantic typing judgment* $\Gamma \vDash M : A$ is then defined as: $\forall \Theta \in \mathcal{G}(\Gamma). \Theta M \in \mathcal{E}(A)$.

Theorem 3 (Fundamental Lemma). *If $\Gamma \vdash M : A$ then $\Gamma \vDash M : A$.*

This brings us one step closer to proving the theorem of directional adequacy. One piece is still missing: since the theorem assumes set inclusion of denotations rather than equality, we will need a different form of compositionality of the denotations than the one that holds by definition.

To state this form of compositionality we first define the standard notion of a program with holes. A function $\Xi[-] : \Gamma \vdash A \rightarrow \Delta \vdash B$ is a *program context* (or *context* for short) if, in the language extended with a program \bullet and additional axioms $\Gamma' \vdash \bullet : A$ for all $\Gamma' \geq \Gamma$, we have $\Delta \vdash \Xi[\bullet] : B$; and if $\Gamma \vdash M : A$, then $\Xi[M]$ is obtained from $\Xi[\bullet]$ by replacing every occurrence of \bullet with M .

Theorem 4 (Compositionality). *Let $\Xi[-] : \Gamma \vdash A \rightarrow \cdot \vdash G$ be a context for ground G , and $M, N \in \Gamma \vdash A$. If $\llbracket M \rrbracket^c \subseteq \llbracket N \rrbracket^c$ then $\llbracket \Xi[M] \rrbracket^c \subseteq \llbracket \Xi[N] \rrbracket^c$.*

The condition that the context be closed and ground is necessary, so an attempt to prove directly by induction on the structure of the context fails. The proof, omitted for brevity, instead uses a binary logical relation approximating containment that identifies with it on ground types; the main ingredient being:

$$\mathcal{E}^\circ(A) := \{\langle P, Q \rangle \in \mathcal{T}[A] \times \mathcal{T}[A] \mid \forall \alpha x \in P \exists \beta y \in Q. \alpha = \beta \wedge \langle x, y \rangle \in \mathcal{V}^\circ(A)\}$$

With this compositionality in hand we are finally ready to prove the main result of this paper, that we will then use to justify program transformation. To state it we first spell out the standard definition of contextual refinement.

Suppose that $M, N \in \Gamma \vdash A$. We say that M *refines* N , and write $M \sqsubseteq N$, if $\sigma, \Xi[M] \rightsquigarrow^* \rho, V$ implies $\sigma, \Xi[N] \rightsquigarrow^* \rho, V$ whenever $\Xi[-] : \Gamma \vdash A \rightarrow \cdot \vdash G$ is a context for ground G . This justifies the transformation $N \rightarrow M$, since replacing N with M within a larger program introduces no additional behaviours.

Theorem 5 (Directional Adequacy). *If $\llbracket M \rrbracket^c \subseteq \llbracket N \rrbracket^c$ then $M \sqsubseteq N$.*

Proof. Let $\Xi[-] : \Gamma \vdash A \rightarrow \cdot \vdash G$ be a program context for some ground G , and assume $\sigma, \Xi[M] \rightsquigarrow^* \rho, V$ for some V . By soundness, $\langle \sigma, \rho \rangle \llbracket V \rrbracket^v \in \llbracket \Xi[M] \rrbracket^c$. Using compositionality, by assumption $\langle \sigma, \rho \rangle \llbracket V \rrbracket^v \in \llbracket \Xi[N] \rrbracket^c$. By the fundamental lemma, $\Xi[N] \xrightarrow{\langle \sigma, \rho \rangle} W$ for some W such that $\llbracket W \rrbracket^v = \llbracket V \rrbracket^v$. They are of ground type, so $W = V$. Therefore, $\sigma, \Xi[N] \rightsquigarrow^* \rho, V$.

6.1 Example Transformations

Thanks to directional adequacy, we can now justify various transformations and equivalences using rather mundane calculations, requiring no reasoning about the context in which these transformations are to take place.

Example 9. We make the reasoning from Example 4 precise.

Denote $\Gamma := \mathbf{x} : \mathbf{Loc}, \mathbf{y} : \mathbf{Val}$. We have $\Gamma \vdash \mathbf{x} := \mathbf{y}; \mathbf{x} ? : \mathbf{Val}$ and $\Gamma \vdash \mathbf{x} := \mathbf{y}; \mathbf{y} : \mathbf{Val}$. Let $\gamma \in \llbracket \Gamma \rrbracket$, and denote $\ell := \gamma \mathbf{x}$ and $v := \gamma \mathbf{y}$. Calculating, we have:

$$\llbracket \mathbf{x} := \mathbf{y}; \mathbf{x} ? \rrbracket^c \gamma = \tilde{U}_{\ell, v} \tilde{Y}^? \tilde{L}_\ell (w. \tilde{Y}^? \tilde{w}) \supseteq \tilde{U}_{\ell, v} \tilde{Y}^? \tilde{v} = \llbracket \mathbf{x} := \mathbf{y}; \mathbf{y} \rrbracket^c \gamma$$

By directional adequacy, $\mathbf{x} := \mathbf{y}; \mathbf{y} \sqsubseteq \mathbf{x} := \mathbf{y}; \mathbf{x} ?$.

Example 10. We validate elimination of irrelevant reads, i.e. $M \sqsubseteq \mathbf{x} ? ; M$:

$$\llbracket \mathbf{x} ? ; M \rrbracket^c \gamma = \llbracket (\lambda _ . M) \mathbf{x} ? \rrbracket^c \gamma = \llbracket \mathbf{x} ? \rrbracket^c \gamma \gg \lambda v. \llbracket M \rrbracket^c \gamma = \tilde{Y}^? (\llbracket M \rrbracket^c \gamma) \supseteq \llbracket M \rrbracket^c \gamma$$

As mentioned in §2.5, the semantics does not validate *introduction* of irrelevant reads, i.e. we have $\llbracket \mathbf{x} ? ; M \rrbracket^c \not\subseteq \llbracket M \rrbracket^c$ even though $\mathbf{x} ? ; M \sqsubseteq M$.

Example 11. Thanks to our use of standard monad-based semantics, structural transformations and equivalences follow from structural reasoning, avoiding considerations relating to shared-state. For instance:

$$\begin{aligned} \llbracket \mathbf{if} \ \mathbf{y} \ \mathbf{then} \ \lambda \mathbf{x}. K_{\mathbf{true}} \ \mathbf{else} \ \lambda \mathbf{x}. K_{\mathbf{false}} \rrbracket^c \gamma &= \llbracket \lambda \mathbf{x}. K_{\gamma \mathbf{y}} \rrbracket^c \gamma \\ &= \text{return } \lambda z. \llbracket K_{\gamma \mathbf{y}} \rrbracket^c (\gamma [\mathbf{x} \mapsto z]) = \llbracket \lambda \mathbf{x}. \mathbf{if} \ \mathbf{y} \ \mathbf{then} \ K_{\mathbf{true}} \ \mathbf{else} \ K_{\mathbf{false}} \rrbracket^c \gamma \end{aligned}$$

Therefore, $\mathbf{if} \ \mathbf{y} \ \mathbf{then} \ \lambda \mathbf{x}. K_{\mathbf{true}} \ \mathbf{else} \ \lambda \mathbf{x}. K_{\mathbf{false}} \cong \lambda \mathbf{x}. \mathbf{if} \ \mathbf{y} \ \mathbf{then} \ K_{\mathbf{true}} \ \mathbf{else} \ K_{\mathbf{false}}$.

Finally, adequacy can help validate expected transformations involving (\parallel):

Example 12. Defining $\text{map } \psi P := \{\alpha(\psi x) \mid \alpha x \in P\}$ we have:

$$\begin{aligned} \llbracket \langle M, N \rangle \rrbracket^c \gamma &\subseteq \llbracket M \parallel N \rrbracket^c \gamma && \text{(Sequencing)} \\ \llbracket M \parallel V \rrbracket^c \gamma &= \text{map}(\lambda x. \langle x, \llbracket V \rrbracket^v \rangle)(\llbracket M \rrbracket^c \gamma) && \text{(Neutrality)} \\ \llbracket M \parallel N \rrbracket^c \gamma &= \text{map}(\lambda \langle y, x \rangle. \langle x, y \rangle)(\llbracket N \parallel M \rrbracket^c \gamma) && \text{(Symm.)} \\ \llbracket (M \parallel N) \parallel K \rrbracket^c \gamma &= \text{map}(\lambda \langle x, \langle y, z \rangle \rangle. \langle \langle x, y \rangle, z \rangle)(\llbracket M \parallel (N \parallel K) \rrbracket^c \gamma) && \text{(Assoc.)} \end{aligned}$$

Unlike the previous examples, proving the above involves careful reasoning at the level of the traces. We still gain the benefit of justifying equivalences and transformations of programs – even open ones – without resorting to analysis under arbitrary program contexts and substitutions:

$$\begin{aligned} \langle M, N \rangle &\subseteq M \parallel N && \text{(Sequencing)} \\ \langle M, V \rangle &\cong M \parallel V && \text{(Neutrality)} \\ M \parallel N &\cong \mathbf{match} N \parallel M \mathbf{with} \langle y, x \rangle \rightarrow \langle x, y \rangle && \text{(Symm.)} \\ (M \parallel N) \parallel K &\cong \mathbf{match} M \parallel (N \parallel K) \mathbf{with} \langle \langle x, \langle y, z \rangle \rangle \rangle \rightarrow \langle x, y \rangle, z && \text{(Assoc.)} \end{aligned}$$

Coordinating the returned values make these somewhat awkward. More convenient but less informative forms are derivable, such as $M ; N ; K \subseteq (M \parallel N) ; K$ (mentioned as a transformation in §1) which is a consequence of (Sequencing).

7 Conclusion, Related Work, and Future Work

We have defined a monad-based denotational semantics for a language for shared-state providing standard higher-order semantics supporting standard meta-theoretic development. This monad is a representation of the one induced by the equational theory of resumptions, which extends non-deterministic global-state with a delaying/yielding operator [14].

Abadi and Plotkin [1] design a modification for the theory of resumptions to define a denotational semantics for a concurrent imperative programming language with cooperative asynchronous threads. We have shown that the theory of resumptions can be used as-is to define denotational semantics for concurrency, albeit of a different kind. It is interesting to note that they interpret the unary operator analogously to our interpretation of $Y^?$, rather than Y . By decomposing into a sum we were able to validate transformations that are not equivalences.

Benton et al. [4] also define a monad for higher-order shared-state, with additional features such as recursion and abstract locations, using Brookes’s style of semantics. Contrasting, the monad we defined is presented algebraically, and has finite sets of traces, whereas Benton et al.’s denotations are infinite even for recursion-free programs. Although this finiteness makes our definition simpler, we saw in Example 10 that it leads to a resumption-counting issue, thus less abstract semantics. It would be interesting to analyse their semantic model from the algebraic perspective as it may lead to more abstract semantics.

Like in previous work, including those mentioned above, our semantics is based on the sets of traces, originally used by Brookes [6] to define denotational semantics for an imperative concurrent language. Brookes proved that this semantics is not only directionally adequate, but also fully abstract. The proof makes crucial use of atomic execution blocks which we have not included.

Birkedal et al. [5] provide an interesting related model, given by logical relations (step-indexed, Kripke, etc.) over syntactic terms as semantics. Their language is substantially more expressive including higher-order local store, and accounts for a type-and-effect system semantics. A more precise model could lead to a monadic account that reproduces these results less syntactically.

Also of note are process calculi and algebraic laws concerning the structure of programs. Hoare and van Staden [12] give such an account for concurrent programs, unifying previous work. Their laws are much more general, parameterizing over the notions of sequencing programs and running programs in parallel. It would be interesting to discover if and how our semantics is an instance of theirs. There is also a lot of work on semantics of “while” languages where all information flows through the state, which support more advanced features such as probabilistic choice [3, 11, 41]. Others approach the study of concurrency through game semantics, such as Jaber and Murawski’s [15] study of the semantics of a higher-order call-by-value concurrent language. Trace semantics features in their study too, though their traces are quite different, being sequences of player/opponent actions that incrementally transform configurations.

In the future we plan to refine the type system into a type-and-effect system [18, 20, 22, 29, 39, 40], by annotating the typing judgments with the allowed effects. The denotations then depend on the effect annotations, with each annotation having its own associated equational theory. This may allow additional transformations that are currently beyond this model’s reach. For example, the converse of (Sequencing) under certain syntactic and static guarantees would enable compiler parallelism.

Atomic constructs that disallow interference from the environment are a common feature of concurrent languages. Adding such constructs may be a simple matter, since we have a dedicated operator, `yield`, for allowing interference. Nevertheless, in the spirit of reductionism, we leave this investigation to future work.

We would also like to see how well our approach extends to weak-memory models. In particular, we believe that the timestamp-based operational semantics of the release-acquire memory model [19, 24, 26, 37] is amenable to a similar treatment by using more sophisticated traces.

Acknowledgments We thank Andrés Goens for providing his perspective on a previous version of this paper, and the anonymous APLAS reviewers for their helpful feedback.

References

1. Abadi, M., Plotkin, G.D.: A model of cooperative threads. *Log. Methods Comput. Sci.* **6**(4) (2010). [https://doi.org/10.2168/LMCS-6\(4:2\)2010](https://doi.org/10.2168/LMCS-6(4:2)2010), [https://doi.org/10.2168/LMCS-6\(4:2\)2010](https://doi.org/10.2168/LMCS-6(4:2)2010)
2. Abramsky, S.: *Intensionality, Definability and Computation*, pp. 121–142. Springer International Publishing, Cham (2014). https://doi.org/10.1007/978-3-319-06025-5_5, https://doi.org/10.1007/978-3-319-06025-5_5
3. Anderson, C.J., Foster, N., Guha, A., Jeannin, J., Kozen, D., Schlesinger, C., Walker, D.: Netkat: semantic foundations for networks. In: Jagannathan, S., Sewell, P. (eds.) *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14*, San Diego, CA, USA, January 20–21, 2014. pp. 113–126. ACM (2014). <https://doi.org/10.1145/2535838.2535862>, <https://doi.org/10.1145/2535838.2535862>
4. Benton, N., Hofmann, M., Nigam, V.: Effect-dependent transformations for concurrent programs. In: Cheney, J., Vidal, G. (eds.) *Proceedings of the 18th International Symposium on Principles and Practice of Declarative Programming*, Edinburgh, United Kingdom, September 5–7, 2016. pp. 188–201. ACM (2016). <https://doi.org/10.1145/2967973.2968602>, <https://doi.org/10.1145/2967973.2968602>
5. Birkedal, L., Sieczkowski, F., Thamsborg, J.: A Concurrent Logical Relation. In: Cégielski, P., Durand, A. (eds.) *Computer Science Logic (CSL'12) - 26th International Workshop/21st Annual Conference of the EACSL. Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 16, pp. 107–121. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2012). <https://doi.org/10.4230/LIPIcs.CSL.2012.107>, <http://drops.dagstuhl.de/opus/volltexte/2012/3667>
6. Brookes, S.D.: Full abstraction for a shared-variable parallel language. *Inf. Comput.* **127**(2), 145–163 (1996). <https://doi.org/10.1006/inco.1996.0056>, <https://doi.org/10.1006/inco.1996.0056>
7. Castellan, S.: Weak memory models using event structures. In: Signoles, J. (ed.) *Vingt-septièmes Journées Francophones des Langages Applicatifs (JFLA 2016)*. Saint-Malo, France (Jan 2016), <https://hal.inria.fr/hal-01333582>
8. Fiore, M., Saville, P.: List objects with algebraic structure. In: Miller, D. (ed.) *2nd International Conference on Formal Structures for Computation and Deduction, FSCD 2017*, September 3–9, 2017, Oxford, UK. *LIPIcs*, vol. 84, pp. 16:1–16:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017). <https://doi.org/10.4230/LIPIcs.FSCD.2017.16>, <https://doi.org/10.4230/LIPIcs.FSCD.2017.16>
9. Forster, Y., Kammar, O., Lindley, S., Pretnar, M.: On the expressive power of user-defined effects: Effect handlers, monadic reflection, delimited control. *J. Funct. Program.* **29**, e15 (2019). <https://doi.org/10.1017/S0956796819000121>, <https://doi.org/10.1017/S0956796819000121>
10. Gibbons, J., Hinze, R.: Just do it: simple monadic equational reasoning. In: Chakravarty, M.M.T., Hu, Z., Danvy, O. (eds.) *Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011*, Tokyo, Japan, September 19–21, 2011. pp. 2–14. ACM (2011). <https://doi.org/10.1145/2034773.2034777>, <https://doi.org/10.1145/2034773.2034777>
11. Grathwohl, N.B.B., Kozen, D., Mamouras, K.: KAT + b! In: Henzinger, T.A., Miller, D. (eds.) *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14*, Vienna, Austria, July 14 – 18, 2014. pp. 44:1–44:10. ACM (2014). <https://doi.org/10.1145/2603088.2603095>, <https://doi.org/10.1145/2603088.2603095>

12. Hoare, T., van Staden, S.: The laws of programming unify process calculi. *Sci. Comput. Program.* **85**, 102–114 (2014). <https://doi.org/10.1016/j.scico.2013.08.012>, <https://doi.org/10.1016/j.scico.2013.08.012>
13. Hyland, M., Levy, P.B., Plotkin, G.D., Power, J.: Combining algebraic effects with continuations. *Theor. Comput. Sci.* **375**(1-3), 20–40 (2007). <https://doi.org/10.1016/j.tcs.2006.12.026>, <https://doi.org/10.1016/j.tcs.2006.12.026>
14. Hyland, M., Plotkin, G.D., Power, J.: Combining effects: Sum and tensor. *Theor. Comput. Sci.* **357**(1-3), 70–99 (2006). <https://doi.org/10.1016/j.tcs.2006.03.013>, <https://doi.org/10.1016/j.tcs.2006.03.013>
15. Jaber, G., Murawski, A.S.: Complete trace models of state and control. In: Yoshida, N. (ed.) *Programming Languages and Systems - 30th European Symposium on Programming, ESOP 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021, Proceedings. Lecture Notes in Computer Science*, vol. 12648, pp. 348–374. Springer (2021). https://doi.org/10.1007/978-3-030-72019-3_13, https://doi.org/10.1007/978-3-030-72019-3_13
16. Jagadeesan, R., Jeffrey, A., Riely, J.: Pomsets with preconditions: A simple model of relaxed memory. *Proc. ACM Program. Lang.* **4**(OOPSLA) (nov 2020). <https://doi.org/10.1145/3428262>, <https://doi.org/10.1145/3428262>
17. Jeffrey, A., Riely, J., Batty, M., Cooksey, S., Kaysin, I., Podkopaev, A.: The leaky semicolon: Compositional semantic dependencies for relaxed-memory concurrency. *Proc. ACM Program. Lang.* **6**(POPL) (jan 2022). <https://doi.org/10.1145/3498716>, <https://doi.org/10.1145/3498716>
18. Jouvelot, P., Gifford, D.K.: Algebraic reconstruction of types and effects. In: Wise, D.S. (ed.) *Conference Record of the Eighteenth Annual ACM Symposium on Principles of Programming Languages*, Orlando, Florida, USA, January 21–23, 1991. pp. 303–310. ACM Press (1991). <https://doi.org/10.1145/99583.99623>, <https://doi.org/10.1145/99583.99623>
19. Kaiser, J., Dang, H., Dreyer, D., Lahav, O., Vafeiadis, V.: Strong logic for weak memory: Reasoning about release-acquire consistency in iris. In: Müller, P. (ed.) *31st European Conference on Object-Oriented Programming, ECOOP 2017, June 19–23, 2017, Barcelona, Spain. LIPIcs*, vol. 74, pp. 17:1–17:29. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017). <https://doi.org/10.4230/LIPIcs.ECOOP.2017.17>, <https://doi.org/10.4230/LIPIcs.ECOOP.2017.17>
20. Kammar, O.: Algebraic theory of type-and-effect systems. Ph.D. thesis, University of Edinburgh, UK (2014), <http://hdl.handle.net/1842/8910>
21. Kammar, O., Levy, P.B., Moss, S.K., Staton, S.: A monad for full ground reference cells. In: *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20–23, 2017*. pp. 1–12. IEEE Computer Society (2017). <https://doi.org/10.1109/LICS.2017.8005109>, <https://doi.org/10.1109/LICS.2017.8005109>
22. Kammar, O., McDermott, D.: Factorisation systems for logical relations and monadic lifting in type-and-effect system semantics. In: Staton, S. (ed.) *Proceedings of the Thirty-Fourth Conference on the Mathematical Foundations of Programming Semantics, MFPS 2018, Dalhousie University, Halifax, Canada, June 6–9, 2018. Electronic Notes in Theoretical Computer Science*, vol. 341, pp. 239–260. Elsevier (2018). <https://doi.org/10.1016/j.entcs.2018.11.012>, <https://doi.org/10.1016/j.entcs.2018.11.012>
23. Kammar, O., Plotkin, G.D.: Algebraic foundations for effect-dependent optimisations. In: Field, J., Hicks, M. (eds.) *Proceedings of the 39th ACM*

- SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012. pp. 349–360. ACM (2012). <https://doi.org/10.1145/2103656.2103698>, <https://doi.org/10.1145/2103656.2103698>
24. Kang, J., Hur, C., Lahav, O., Vafeiadis, V., Dreyer, D.: A promising semantics for relaxed-memory concurrency. In: Castagna, G., Gordon, A.D. (eds.) Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017. pp. 175–189. ACM (2017). <https://doi.org/10.1145/3009837.3009850>, <https://doi.org/10.1145/3009837.3009850>
 25. Kavanagh, R., Brookes, S.: A denotational semantics for sparc tso. *Electronic Notes in Theoretical Computer Science* **336**, 223–239 (2018). <https://doi.org/https://doi.org/10.1016/j.entcs.2018.03.025>, <https://www.sciencedirect.com/science/article/pii/S1571066118300288>, the Thirty-third Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIII)
 26. Lahav, O., Giannarakis, N., Vafeiadis, V.: Taming release-acquire consistency. In: Bodík, R., Majumdar, R. (eds.) Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016. pp. 649–662. ACM (2016). <https://doi.org/10.1145/2837614.2837643>, <https://doi.org/10.1145/2837614.2837643>
 27. Levy, P.B.: Call-By-Push-Value: A Functional/Imperative Synthesis, *Semantics Structures in Computation*, vol. 2. Springer (2004)
 28. Linton, F.E.J.: An outline of functorial semantics. In: Eckmann, B. (ed.) *Seminar on Triples and Categorical Homology Theory*. pp. 7–52. Springer Berlin Heidelberg, Berlin, Heidelberg (1969)
 29. Lucassen, J.M., Gifford, D.K.: Polymorphic effect systems. In: Ferrante, J., Mager, P. (eds.) *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, San Diego, California, USA, January 10-13, 1988. pp. 47–57. ACM Press (1988). <https://doi.org/10.1145/73560.73564>, <https://doi.org/10.1145/73560.73564>
 30. Moggi, E.: Notions of computation and monads. *Inf. Comput.* **93**(1), 55–92 (1991). [https://doi.org/10.1016/0890-5401\(91\)90052-4](https://doi.org/10.1016/0890-5401(91)90052-4), [https://doi.org/10.1016/0890-5401\(91\)90052-4](https://doi.org/10.1016/0890-5401(91)90052-4)
 31. Nielsen, M., Plotkin, G.D., Winskel, G.: Petri nets, event structures and domains, part I. *Theor. Comput. Sci.* **13**, 85–108 (1981). [https://doi.org/10.1016/0304-3975\(81\)90112-2](https://doi.org/10.1016/0304-3975(81)90112-2), [https://doi.org/10.1016/0304-3975\(81\)90112-2](https://doi.org/10.1016/0304-3975(81)90112-2)
 32. Oles, F.J.: A Category-Theoretic Approach to the Semantics of Programming Languages. Ph.D. thesis (1983)
 33. Oles, F.J.: Type algebras, functor categories, and block structure. *DAIMI Report Series* (156) (1983)
 34. Plotkin, G.D., Power, J.: Notions of computation determine monads. In: Nielsen, M., Engberg, U. (eds.) *Foundations of Software Science and Computation Structures*, 5th International Conference, FOSSACS 2002. Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 8-12, 2002, Proceedings. *Lecture Notes in Computer Science*, vol. 2303, pp. 342–356. Springer (2002). https://doi.org/10.1007/3-540-45931-6_24, https://doi.org/10.1007/3-540-45931-6_24
 35. Plotkin, G.D., Power, J.: Algebraic operations and generic effects. *Appl. Categorical Struct.* **11**(1), 69–94 (2003). <https://doi.org/10.1023/A:1023064908962>, <https://doi.org/10.1023/A:1023064908962>

36. Plotkin, G.D., Pretnar, M.: Handlers of algebraic effects. In: Castagna, G. (ed.) Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5502, pp. 80–94. Springer (2009). https://doi.org/10.1007/978-3-642-00590-9_7, https://doi.org/10.1007/978-3-642-00590-9_7
37. Podkopaev, A., Sergey, I., Nanevski, A.: Operational aspects of C/C++ concurrency. CoRR **abs/1606.01400** (2016), <http://arxiv.org/abs/1606.01400>
38. Reynolds, J.C.: The essence of algol. In: de Bakker, J.W., van Vliet, J.C. (eds.) Algorithmic Languages. pp. 345–372. International Symposium on Algorithmic Languages, Amsterdam; New York: North-Holland Pub. Co. (1981)
39. Talpin, J., Jouvelot, P.: Polymorphic type, region and effect inference. J. Funct. Program. **2**(3), 245–271 (1992). <https://doi.org/10.1017/S095679680000393>, <https://doi.org/10.1017/S095679680000393>
40. Talpin, J., Jouvelot, P.: The type and effect discipline. Inf. Comput. **111**(2), 245–296 (1994). <https://doi.org/10.1006/inco.1994.1046>, <https://doi.org/10.1006/inco.1994.1046>
41. Wagemaker, J., Foster, N., Kappé, T., Kozen, D., Rot, J., Silva, A.: Concurrent netkat: Modeling and analyzing stateful, concurrent networks. CoRR **abs/2201.10485** (2022), <https://arxiv.org/abs/2201.10485>
42. Winskel, G.: On powerdomains and modality. Theor. Comput. Sci. **36**, 127–137 (1985). [https://doi.org/10.1016/0304-3975\(85\)90037-4](https://doi.org/10.1016/0304-3975(85)90037-4), [https://doi.org/10.1016/0304-3975\(85\)90037-4](https://doi.org/10.1016/0304-3975(85)90037-4)