

Effect-Dependent Transformations for Concurrent Programs

Nick Benton

Facebook, London, UK

Martin Hofmann

LMU, Munich, Germany

Vivek Nigam

fortiss, Munich

Abstract

We describe a denotational semantics for an abstract effect system for a higher-order, shared-variable concurrent programming language. We prove the soundness of a number of general effect-based program equivalences, including a parallelization equation that specifies sufficient conditions for replacing sequential composition with parallel composition. Effect annotations are relative to abstract locations specified by contracts rather than physical footprints allowing us in particular to show the soundness of some transformations involving fine-grained concurrent data structures, such as Michael-Scott queues, that allow concurrent access to different parts of mutable data structures.

Our semantics is based on refining a trace-based semantics for first-order programs due to Brookes. By moving from concrete to abstract locations, and adding type refinements that capture the possible side-effects of both expressions and their concurrent environments, we are able to validate many equivalences that do not hold in an unrefined model. The meanings of types are expressed using a game-based logical relation over sets of traces. Two programs e_1 and e_2 are logically related if one is able to solve a two-player game: for any trace with result value v_1 in the semantics of e_1 (challenge) that the player presents, the opponent can present an (response) equivalent trace in the semantics of e_2 with a logically related result value v_2 .

1. Introduction

Type-and-effect systems refine conventional types with extra static information capturing a safe upper bound on the possible side-effects of expression evaluation. Since their introduction by Gifford and Lucassen [19], effect systems have been used for

Email addresses: nick.benton@gmail.com (Nick Benton), hofmann@ifi.lmu.de (Martin Hofmann), vivek.nigam@gmail.com (Vivek Nigam)

5 many purposes, including region-based memory management [13], tracking exceptions
6 [27, 26], communication behaviour [5] and atomicity [18] for concurrent programs, and
7 information flow [14].

8 A major reason for tracking effects is to justify program transformations, most ob-
9 viously in optimizing compilation [11]. For example, one may remove computations
10 whose results are unused, *provided* that they are sufficiently pure, or commute two
11 state-manipulating computations, *provided* that the locations they may read and write
12 are suitably disjoint. Several groups have recently studied the semantics of effect sys-
13 tems, with a focus on formally justifying such effect-dependent equational reasoning
14 [21, 9, 6, 12, 29]. A common approach, which we follow here, is to interpret effect-
15 refined types using a logical relation over the (denotational or operational) semantics
16 of the unrefined (or untyped) language, simultaneously identifying both the subset of
17 computations that have a particular effect type and a coarser notion of equivalence (or
18 approximation) on that subset. Such a semantic approach decouples the meaning of
19 effect-refined types from particular syntactic rules: one may establish that a term has a
20 type using various more or less approximate inference systems, or by detailed semantic
21 reasoning.

22 For sequential computations with global state, denotational models already provide
23 significant abstraction. For example, the denotations of `skip` and `X++;X--` are typi-
24 cally equal, so it is immediate that the second is semantically pure. More generally,
25 the meaning of a judgement $\Gamma \vdash e : \tau \& \varepsilon$ guarantees that the result of evaluating e will
26 be of type τ with side-effects at most ε , under assumptions Γ (a ‘rely’ condition), on
27 the behaviour of e ’s free variables. The possible interaction points between e and its
28 environment are restricted to initial states and parameter values, and final states and
29 results, of e itself and its explicitly-listed free variables. Furthermore, all those interac-
30 tion points are visible in the term and are governed by specific annotations appearing
31 in the typing judgement.

32 For shared-variable concurrency, there are many more possible interactions. An ex-
33 pression’s environment now also includes anything that may be running concurrently
34 and, moreover, atomic steps of e and its concurrent environment may be arbitrarily in-
35 terleaved, so it is no longer sufficient to just consider initial and final states. A priori,
36 this leads to far fewer equations between programs. For example, `X++;X--` may be dis-
37 tinguished from `skip` by being run concurrently with a command that reads or writes
38 `X`. But few programs do anything useful in the presence of unconstrained interference,
39 so we need ways to describe and control it. Fine-grained, optimistic algorithms, which
40 rely on custom protocols being followed by multiple threads with concurrent access
41 to a shared data structure, can significantly outperform ones based on coarse-grained
42 locking, but are notoriously challenging to write and verify.

43 There is a huge literature on shared-variable concurrency, from type systems en-
44 suring race-freedom of programs with locks [1] to sophisticated semantic models for
45 reasoning about refinement of fine-grained concurrent datastructures [31]. This paper
46 explores effect types as a straightforward, lightweight interface language for modular
47 reasoning about equivalence and refinement, e.g. for safely transforming sequential
48 composition into parallelism. We show how the semantics of a simple effect system
49 scales smoothly to the concurrent setting, allowing us to control interference and prove
50 non-trivial equivalences, extending (somewhat to our surprise) to the correctness of

51 some fine-grained algorithms.

52 We build on a trace semantics for concurrent programs, due to Brookes [15], which
53 explicitly describes possible interference by the environment. We extend Brookes’s
54 semantics to a higher-order language and then refine it by a semantically-formulated
55 effect system that separately tracks: (1) the store effects of an expression during eval-
56 uation; (2) the assumed effects of transitions by the environment; and (3) the overall
57 end-to-end effect. Rather than tracking effects at the level of individual concrete heap
58 cells, we view the heap as a set of abstract data structures, each of which may span
59 several locations, or parts of locations [6]. Each abstract location has its own notion of
60 equality, and its own notion of legal mutation. Write effects, for example, need only
61 be flagged when the equivalence class of an abstract location may change. Both typing
62 and refinement judgements may be established by a combination of generic type-based
63 rules and semantic reasoning in the model.

64 This paper is an extended archival version of [10] which has been presented at
65 PPDP 2016. In addition to the conference version this paper has more detail about
66 the higher-order version of Brookes’ trace semantics (Section 3), more examples, in
67 particular the one on loop parallelization, and detailed proofs of the main results on
68 soundness of the logical relation and general reasoning principles (Theorem 7.7) and
69 on canonical program equivalences (Theorem 9.1).

70 We begin with some motivating examples.

71 *Equivalence modulo non-interference.* Our semantics justifies the equation ($X :=$
72 $!X + 1; X := !X + 1) = (X := !X + 2)$ at the effect type $\text{unit} \ \& \ \{co_X\} \mid \varepsilon \mid \varepsilon \cup \{rd_X, wr_X\}$,
73 provided that the effect, ε , of the concurrent environment does not involve X . This says
74 that the two commands are equivalent with return type unit ,¹ exhibit the effect co_X ,
75 signifying concurrent or ‘chaotic’ access to X along the way, and have an overall end-
76 to-end effect of ε plus reading and writing X .

77 *Overlapping References.* Let p, p^{-1} implement a bijection $\mathbb{Z} \rightarrow \mathbb{Z} \times \mathbb{Z}$, and consider
78 the following functions:

```
readFst () = p(!X).1
readSnd () = p(!X).2
wrtFst n = (rec try _ = let m = !X in
             if cas(X, m, p^{-1}(n, p(m).2))
             then () else try () ())
wrtSnd n = (rec try _ = let m = !X in
             if cas(X, m, p^{-1}(p(m).1, n))
             then () else try () ())
```

79 which multiplex two abstract integer references onto a single concrete one. Note that
80 the write functions, `wrtFst` and `wrtSnd`, use compare-and-swap, `cas`, to atomically

¹Being equal at a type means being may-indistinguishable for any observations which use the terms at that type.

81 update the value of the reference. More precisely as follows

```
cas(X, v1, v2) = atomic(if !X = v1 then X := v2; true else false)
```

82 where `atomic` enforces the atomic evaluation of the argument expression.

83 Our generic rules (Figure 5) then say that a program, e_1 , that only reads and/or
84 writes one abstract reference can be commuted, or executed in parallel, with another
85 program, e_2 , that only reads and/or writes into a different reference. This lets one use
86 types to, say, justify parallelizing a call to `wrtFst` followed by one to `wrtSnd`, even
87 though they read and write the same concrete location, which looks like a race.

88 *Version numbers*:. One can isolate a transaction that reads and then writes a piece of
89 state simply by enclosing the whole thing in `atomic(·)`. A more concurrent alternative
90 adds a monotonic version number to the data. A transaction then works on a private
91 copy, only committing its changes back (and incrementing the version) if the current
92 version number is the same as that of the original copy. We can define an abstract integer
93 reference \mathfrak{X} in terms of two concrete ones, X_{ver} and X_{val} , governed by a specification
94 that says $!X_{\text{val}}$ may only change when $!X_{\text{ver}}$ increases. We define

```
transact f = let rec try() =
  let (val, ver) = atomic(!Xval, !Xver)
  in let res = f(val) in if atomic(if !Xver = ver then
    Xver := ver + 1; Xval := res; true else false)
  then () else try()
in try()
```

95 Under the assumption that f is a pure function (has effect type $\text{int} \xrightarrow{\emptyset|\varepsilon} \text{int}$ for any
96 ε), we can show

```
transact f = atomic(Xval := f(!Xval); Xver := !Xver + 1)
```

97 at type $\text{unit} \&\{rd_{\mathfrak{X}}, wr_{\mathfrak{X}}\} \mid \varepsilon \mid \varepsilon \cup \{rd_{\mathfrak{X}}, wr_{\mathfrak{X}}\}$ for any ε not including chaotic access, $co_{\mathfrak{X}}$,
98 to \mathfrak{X} . The environment effect ε here *may* include reading and writing \mathfrak{X} , so concurrent
99 calls to `transact` are linearizable.

100 *Loop Parallelization*:. Our next example is inspired by a loop unrolling optimiza-
101 tion [30]. Assume given a linked list of integers pointed by *head*. Consider the fol-




Figure 1: Illustration of a Michael-Scott Queue. The list resulting from the pointer to the element n_0 (the *head* pointer with the continuous arrow in black) contains the list of elements $[n_1, \dots, n_j]$. The enqueueing operation is illustrated by the dotted arrow and the box with the element n_{j+1} (in blue), while the dequeueing operation is illustrated by the dot dashed head pointer (in red).

102 lowering functions:

```

map f = let rec applyf n =
        n.ele := f(n.ele); if n.next = null then ()
        else applyf (n.next)
      in if !head = null then () else applyf (!head)

map2Par f = let rec applyf2 n =
        n.ele := f(n.ele) || n.next.ele := f(n.next.ele);
        if n.next.next = null then ()
        else if n.next.next.next = null then
            n.next.next.ele = f(n.next.next.ele)
        else applyf2 (n.next.next)
      in if !head = null then ()
        else if !head.next = null then
            !head.next.ele := f(!head.next.ele)
        else applyf2 (!head)

```

103 The function `map` simply applies a pure function f to each element of the list, each
 104 element per iteration. The function `map2Par`, on the other hand, applies f to two
 105 consecutive elements of the list in parallel, potentially allowing one to exploit multiple
 106 cores. Our effect-based reasoning will soundly transform `map` into `map2Par` (under
 107 the assumption that the environment does not interfere with the list).

108 *Michael-Scott queue.* The Michael-Scott Queue [25] (MSQ) is a fine grained concurr-
 109 ent data structure, allowing threads to access and modify different parts of a queue
 110 safely and simultaneously. We present an idealized version like that of Turon et al [31],
 111 which omits a tail pointer.

112 An MSQ maintains a pointer *head* to a non-empty linked list as depicted in Fig-
 113 ure 1. The first node, that containing the element n_0 in the figure, is not an element of
 114 the queue, but is a “sentinel”. Hence the queue in the figure holds $[n_1, \dots, n_j]$.

115 The enqueue and dequeue operations are defined in Figure 2 and illustrated in the
 116 diagram to the right. Elements are dequeued from the beginning of the list, and en-

```

dequeue () = (rec try () = let n0 =!head in
              if !n0.next = null then null
              else let n1 =!n0.next in
                   if cas(!head, n0, n1) then !n1.ele
                   else try () ())
enqueue(x) = (rec try (p) =
              if !p.next = null then
                if atomic(if !p.next = null then
                           !p.next := ref(x, null); true else false)
                then () else try (!p.next)
              else try (!p.next)) !head)
mem x =      (rec find l =
              if l = null then false else
                if !l.ele = x then true else
                find !l.next) !head.next)
reset () =   (rec deqAll () =
              if dequeue () = null then ()
              else deqAll () ())

```

Figure 2: Enqueue, Dequeue, Membership, and Reset programs for a Michael-Scott Queue at location *head*.

117 queued at the end, involving a traversal that is done without locking. Once the end, *p*,
 118 of the list is found, the program atomically attempts to insert the new element. This
 119 operation has to be atomic because other programs may have enqueued elements to the
 120 end of the list, meaning that *p* is no longer the end of the list.

121 We prove that the enqueue and dequeue of Figure 2 are equivalent to their atomic
 122 versions `atomic(enqueue)` and `atomic(dequeue)`, which perform all operations in a
 123 single step, at a type that allows the environment to be concurrently reading and writing
 124 the queue. So the fine-grained MSQ behaves like a synchronized queue, as might also
 125 be implemented using locks.

126 We can also show that `mem` is equivalent to its atomic version `atomic(mem)` at
 127 type $\text{int} \xrightarrow[\varepsilon_2]{\emptyset \mid \varepsilon_2, rd_{MSQ}} \text{bool}$ provided the environment does not access the MSQ chaot-
 128 ically, *i.e.*, $co_{MSQ} \notin \varepsilon_2$. This typing denotes that `mem` has the effect of reading the
 129 MSQ, both during execution and as overall effect. With more assumptions on the en-
 130 vironment effects ε_2 , namely, that it does not enqueue nor dequeue MSQ, `mem` may
 131 participate in many of the equations we prove sound, *e.g.*, commuting, deadcode.

132 Similarly, `reset` is equivalent to `atomic(reset)` at the type $\text{unit} \xrightarrow[\varepsilon_2]{rd_{MSQ} wr_{MSQ} \mid \varepsilon_2, wr_{MSQ}}$
 133 `unit`. During execution, `reset` both reads and writes the MSQ, but we can show se-
 134 mantically that its overall effect is only the environmental effect ε_2 plus writing the
 135 MSQ; there is no overall read effect. Again, from the typing (and assumptions on ε_2),
 136 one obtains equations involving `reset` without further semantic reasoning.

137 **2. Syntax**

138 In this section we define the syntax of a metalanguage for concurrent, stateful com-
 139 putations and higher-order functions. Communication between parallel computations
 140 is via a shared heap mapping dynamically allocated locations to structured values,
 141 which include pointers. To keep the model simple, we do not allow functions to be
 142 stored in the heap (no higher-order store).

143 *Memory model.* We assume a countably infinite set \mathbb{L} of physical locations X_1, \dots, X_n, \dots
 144 and a set \mathbb{VB} of “R-values” that can be stored in those references including integers,
 145 booleans, locations, and tuples of R-values, written (v_1, \dots, v_n) . We assume that it is
 146 possible to tell of which form a value is and to retrieve its components in case it is a tu-
 147 ple. A heap h , then, is a *finite map* from \mathbb{L} to \mathbb{VB} , written $\{(X_1, c_1), (X_2, c_2), \dots, (X_n, c_n)\}$,
 148 specifying that the value stored in location X_i is c_i . We write $\text{dom}(h)$ for the domain of
 149 h and write $h[X \mapsto c]$ for the heap that agrees with h except that it gives the variable X
 150 the value c . The set of heaps is denoted by \mathbb{H} . We also assume that $\text{new}(h, v)$ yields a
 151 pair (X, h') where $X \in \mathbb{L}$ is a fresh location and $h' \in \mathbb{H}$ is $h[X \mapsto v]$.

152 *Syntax of expressions.* The syntax of untyped values and computations is:

$$\begin{aligned} v &::= x \mid (v_1, v_2) \mid v_r \mid c \mid \mathbf{rec} \ f \ x = t \\ e &::= v \mid \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \mid v_1 \ v_2 \mid \mathbf{if} \ v \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 \\ &\quad \mid !v \mid v_1 := v_2 \mid \mathbf{ref}(v) \mid e_1 \parallel e_2 \mid \mathbf{atomic}(e) \end{aligned}$$

153 Here, x ranges over variables, v_r over R-values, and c over built-in functions, which
 154 include arithmetic, testing whether a value is an integer, function, pair or reference,
 155 equality on simple values, etc. Each c has a corresponding semantic partial function
 156 F_c , so for example $F_+(n, n') = n + n'$ for integers n, n' .

157 The construct $\mathbf{rec} \ f \ x = e$ defines a recursive function with body e and recursive
 158 calls made via f ; we use $\lambda x.e$ as syntactic sugar in the case when f is not free in e .
 159 Next, $!v$ (reading) returns the contents of location v , $v_1 := v_2$ (writing) updates location
 160 v_1 with value v_2 , and $\mathbf{ref}(v)$ (allocating) returns a fresh location initialized with v . The
 161 metatheory is simplified by using “let-normal form”, in which the only elimination
 162 for computations is \mathbf{let} , though we sometimes nest computations as shorthand for let-
 163 expanded versions in examples. We emphasize that the use of let-normal form is merely
 164 a convenience not reducing expressivity in any way. For example, we view $v_1 := v_2$ as
 165 short-hand for $\mathbf{let} \ x_1 = v_1 \ \mathbf{in} \ \mathbf{let} \ x_2 = v_2 \ \mathbf{in} \ x_1 := x_2$ thus writing the evaluation order
 166 explicitly.

167 The construct $e_1 \parallel e_2$ is evaluated by arbitrarily interleaving evaluation steps of e_1
 168 and e_2 until each has produced a value, say v_1 and v_2 ; the result is then (v_1, v_2) . Assign-
 169 ment, dereferencing and allocation are atomic, but evaluation of nested expressions is
 170 generally not. To enforce atomicity, $\mathbf{atomic}(e)$ evaluates an arbitrary e in one step,
 171 without any environmental interference.

172 We define the free variables, $FV(e)$, of a term, closed terms, and the substitution
 173 $e[v/x]$ of v for x in e , in the usual way. Locations may occur in terms, but the type
 174 system will constrain their use.

175 **3. Denotational Model**

176 We now describe a denotational semantics for our metalanguage based on Brookes’
 177 trace semantics [15]. In the technical report [7] we give some more detail and in partic-
 178 ular a proof of adequacy with respect to an interleaving operational semantics, which
 179 we elide here since it is not germane to the topic of this article.

180 *3.1. Preliminaries*

181 A *predomain* is an ω -cpo, *i.e.*, a partial order with suprema of ascending chains. A
 182 *domain* is a predomain with a least element, \perp . Recall that $f : A \rightarrow A'$ is *continuous* if
 183 it is monotone $x \leq y \Rightarrow f(x) \leq f(y)$ and preserves suprema of chains, *i.e.*, $f(\sup_i x_i) =$
 184 $\sup_i f(x_i)$. Any set is a predomain with the discrete order (flat predomain). If X is a set
 185 and A a predomain then any $f : X \rightarrow A$ is continuous. We denote a partial (continuous)
 186 function from set (predomain) A to set (predomain) B by $f : A \rightarrow B$. If A, B are
 187 predomains the cartesian product $A \times B$ and the set of continuous functions $A \rightarrow B$ form
 188 themselves predomains (with the obvious componentwise and pointwise orders) and
 189 make the category of predomains cartesian closed. Likewise, the partial continuous
 190 functions $A \rightarrow B$ between predomains A, B form a domain.

191 If $P \subseteq A$ and $Q \subseteq B$ are subsets of predomains A and B we define $P \times Q \subseteq A \times B$
 192 and $P \rightarrow Q \subseteq A \rightarrow B$ in the usual way. We may write $f : P \rightarrow Q$ for $f \in P \rightarrow Q$.

193 A subset $U \subseteq A$ is *admissible* if whenever $(a_i)_i$ is an ascending chain in A such that
 194 $a_i \in U$ for all i , then $\sup_i a_i \in U$, too. If $f : X \times A \rightarrow A$ is continuous and A is a
 195 domain then one defines $f^\ddagger(x) = \sup_i f_x^i(\perp)$ with $f_x(a) = f(x, a)$. As usual, f_x^i is the i -th
 196 iteration of f_x . One has, $f(x, f^\ddagger(x)) = f^\ddagger(x)$ and if $U \subseteq A$ is admissible and contains \perp
 197 and $f : X \times U \rightarrow U$ then $f^\ddagger : X \rightarrow U$, too. Thinking of U as a predicate on the elements
 198 of A , we have that $f^\ddagger(x)$ satisfies U provided that f_x preserves and U is admissible and
 199 an $\perp \in U$. This principle is known as Scott induction. An element d of a predomain A
 200 is *compact* if whenever $d \leq \sup_i a_i$ then $d \leq a_i$ for some i . E.g. in the domain of partial
 201 functions from \mathbb{N} to \mathbb{N} the compact elements are precisely the finite ones. A continuous
 202 partial function $f : A \rightarrow A$ is a *retract* if $f(a) \leq a$ and $f(f(a)) = f(a)$ hold for all $a \in A$.
 203 In short: $f \leq id_A$ and $f \circ f \leq f$. If, in addition, f has a finite image then f is called a
 204 *deflation* [3]. Note that if f is a retract then $\text{dom}(f) = \text{Img}(f)$ and if $a \in \text{Img}(f)$ then
 205 $a = f(a)$. We also note that if a is in the image of a deflation then a is compact.

206 We define the usual state monad on predomains, by taking $SA = \mathbb{H} \rightarrow \mathbb{H} \times A$.
 207 As we seen, Scott induction applies to admissible predicates only which motivates the
 208 following definition:

209 **Definition 3.1.** *Let P be a subset of a predomain A . Then $\text{Adm}(P)$ is the least admissi-*
 210 *ble superset of P . Concretely, $a \in \text{Adm}(P)$ iff there exists a chain $(a_i)_i$ such that $a_i \in P$*
 211 *for all i and $a = \sup_i a_i$.*

212 We will often find ourselves in the situation of wanting to show some property P ,
 213 but (since we want to use Scott induction) are only able to prove $\text{Adm}(P)$. The following
 214 lemma says intuitively that if we know $x_1 \in \text{Adm}(P_1) \dots x_n \in \text{Adm}(P_n)$ then we can
 215 actually assume $x_1 \in P_1 \dots x_n \in P_n$ so long as the end result (“ Q ”) is admissible and
 216 the x_i s are used in a continuous fashion.

217 **Lemma 3.2.** *If $f : A_1 \times \dots \times A_n$ is continuous; $P_i \subseteq A_i$ are arbitrary subsets and $Q \subseteq B$*
 218 *is admissible then $f : P_1 \times \dots \times P_n \rightarrow Q$ implies $f : Adm(P_1) \times \dots \times Adm(P_n) \rightarrow Q$.*

219 The following lemma has a similar purpose. It asserts that under mild condition on
 220 the pre-domains involved, in order to show that some continuous function $Adm(P \rightarrow Q)$
 221 it suffices to show that it is in $P \rightarrow Adm(Q)$.

222 **Lemma 3.3.** *Let A, B be predomains and let $(p_i)_i$ be a chain of retracts on B such that*
 223 *$p_i(b)$ is compact for each i and $\sup_i p_i = id_B$ and $b \in Q$ implies $p_i(b) \in Q$ for all i .*
 224 *Then $P \rightarrow Adm(Q) = Adm(P \rightarrow Q)$.*

225 3.2. Traces

226 A trace models a terminating run of a concurrent computation as a sequence of
 227 pairs of heaps, each representing pre- and post-state of one or more atomic actions.
 228 The semantics of a program then is a (typically large) set of traces (and final values),
 229 accounting for all possible environment interactions.

230 **Definition 3.4** (Traces). *A trace is a finite sequence of the form $(h_1, k_1)(h_2, k_2) \dots (h_n, k_n)$*
 231 *where for $1 \leq j \leq i \leq n$, we have $h_i, k_i \in \mathbb{H}$ and $\text{dom}(h_j) \subseteq \text{dom}(h_i), \text{dom}(h_j) \subseteq$*
 232 *$\text{dom}(k_i), \text{dom}(k_j) \subseteq \text{dom}(h_i), \text{dom}(k_j) \subseteq \text{dom}(k_i)$. We write Tr for the set of traces.*

233 Let t be a trace. A trace of the form $u(h, h)v$ where $t = uv$ is said to arise from t by
 234 stuttering. A trace of the form $u(h, k)v$ where $t = u(h, q)(q, k)v$ is said to arise from t by
 235 mumbling. For example, if $t = (h_1, k_1)(h_2, k_2)(h_3, k_3)$ then $(h_1, k_1)(h, h)(h_2, k_2)(h_3, k_3)$
 236 arises from t by stuttering. In the case where $k_1 = h_2$ the trace $(h_1, k_2)(h_3, k_3)$ arises
 237 from t by mumbling. A set of traces U is closed under stuttering and mumbling if
 238 whenever t' arises from t by stuttering or mumbling and $t \in U$ then $t' \in U$, too.

239 Brookes [15] gives a fully-abstract semantics for while-programs with parallel
 240 composition using sets of traces closed under stuttering and mumbling. We here extend
 241 his semantics to higher-order functions and general recursion.

242 **Definition 3.5** (Trace Monad). *Let A be a predomain. Elements of the domain TA are*
 243 *sets U of pairs (t, a) where t is a trace and $a \in A$ such that the following properties are*
 244 *satisfied:*

- 245 • [S&M]: *if t' arises from t by stuttering or mumbling and $(t, a) \in U$ then $(t', a) \in$*
 246 *U .*
- 247 • [Down]: *if $(t, a_1) \in U$ and $a_2 \leq a_1$ then $(t, a_2) \in U$.*
- 248 • [Sup]: *if $(a_i)_i$ is a chain in A and $(t, a_i) \in U$ for all i then $(t, \sup_i a_i) \in U$.*

249 *The elements of TA are partially ordered by inclusion.*

250 **Lemma 3.6.** *If A is a predomain then TA is a domain.*

251 An element U of TA represents the possible outcomes of a nondeterministic, inter-
 252 active computation with final result in A . Thus, if $(t, a) \in U$ for $t = (h_1, k_1) \dots (h_n, k_n)$
 253 then there could be n interactions with the environment with heaps h_1, \dots, h_n being

254 “played” by the environment and “answered” with heaps k_1, \dots, k_n by the computa-
 255 tion. After that, this particular computation ends and a is the final result value.

256 For example, the semantics of $X := !X + 1; X := !X + 1; !X$ contains many traces,
 257 including the following, where we write $[n]$ for the heap in which X has value n :

(([10], [12]), 12),
 (([10], [11])([15], [16]), 16),
 258 (([10], [11])([15], [16])([17], [17]), 17),
 (([10], [11])([15], [16])([17], [17]), 16),
 (([10], [11])([17], [17])([15], [16]), 16), ...

259 Axiom [S&M] is taken from Brookes. It ensures that the semantics does not distin-
 260 guish between late and early choice [31] and related phenomena which are reflected,
 261 e.g., in resumption semantics [28], but do not affect observational equivalence. Note
 262 that non-termination is modelled by the empty set, so we are working with an ‘an-
 263 gelic’ notion of equivalence (‘may semantics’ [17]). For example, the semantics of
 264 $X := 0; \text{if } X=0 \text{ then } 0 \text{ else diverge}$ is the same as that of $X := 0; 0$ and contains,
 265 for example $(([10], [0]), 0)$ but also (stuttering) $((([10], [0]), ([34], [34])), 0)$. Note that
 266 it is not possible to tell from a trace whether an external update of X has happened
 267 before or after the reading of X .

268 Let us also illustrate how traces iron out some intensional differences that show up
 269 when concurrency is modelled using transition systems or resumptions. Consider the
 270 following two programs where $?$ denotes a nondeterministically chosen boolean value.

$$\begin{aligned}
 e_1 &\equiv \text{if } ? \text{ then } X := 0; \text{true} \text{ else } X := 0; \text{false} \\
 e_2 &\equiv X := 0; ?
 \end{aligned}$$

271 Both e_1 and e_2 admit the same traces, namely $(([x], [0]), \text{true})$ and $(([x], [0]), \text{false})$
 272 and stuttering variants thereof. In semantic models based on transition systems or
 273 resumptions and bisimulation, these are distinguished, which necessitates the use of
 274 special mechanisms such as history and prophecy variables [2], forward-backward sim-
 275 ulation [24], or speculation [31] in reasoning.

276 Axioms [Down] and [Sup] are known from the Hoare powerdomain [28]. Re-
 277 call that the Hoare powerdomain PA contains the subsets of A which are downclosed
 278 ([Down]) and closed under suprema of chains ([Sup]). Such subsets are also known
 279 as Scott-closed sets. Thus, TA is the restriction of $P(\text{Tr} \times A)$ to the sets closed under
 280 stuttering and mumbling. Axiom [Down] ensures that the ordering is indeed a partial
 281 order and not merely a preorder. Additional nondeterministic outcomes that are less
 282 defined than existing ones are not recorded in the semantics.

283 **Definition 3.7.** *If $U \subseteq \text{Tr} \times A$ then U^\dagger is the least subset of TA containing U , i.e. U^\dagger is*
 284 *the closure of U under [S&M], [Down], [Sup].*

285 **Definition 3.8.** *Let A, B be a predomains. We define the continuous functions $\text{rtn} :$*
 286 *$A \rightarrow TA$ and $\text{bnd} : (A \rightarrow TB) \times TA \rightarrow TB$ by:*

$$\begin{aligned}
 \text{rtn}(a) &:= (\{(\text{h}, a) \mid \text{h} \in \mathbb{H}\})^\dagger \\
 \text{bnd}(f, g) &:= (\{(uv, b) \mid (u, a) \in g \wedge (v, b) \in f(a)\})^\dagger
 \end{aligned}$$

287 These endow TA with the structure of a strong monad. The continuous function
 288 $fromstate : SA \rightarrow TA$ is defined by:

$$fromstate(c) := \{((h, k), a) \mid c(h) = (k, a)\}^\dagger$$

289 If t_1, t_2, t_3 are traces, we write $inter(t_1, t_2, t_3)$ to mean that t_3 can be obtained by inter-
 290 leaving t_1 and t_2 in some way, i.e., t_3 is contained in the shuffle of t_1 and t_2 . In order to
 291 model parallel composition we introduce the following helper function

$$\begin{aligned} | & : TA \times TB \rightarrow T(A \times B) \\ U | V & := \{(t_3, (a, b)) \mid inter(t_1, t_2, t_3), (t_1, a) \in U, (t_2, b) \in V\}^\dagger \end{aligned}$$

292 The continuous map $at : TA \rightarrow TA$ is defined by:

$$at(U) := \{((h, k), v) \mid ((h, k), v) \in U\}^\dagger$$

293 Notice that due to mumbling $((h, k), v) \in U$ iff there exists an element

$$((h_1, h_2)(h_2, h_3)(h_{n-2}, h_{n-1})(h_{n-1}, h_n), v) \in U$$

294 where $h = h_1$ and $h_n = k$. The presence of such an element, however, models an atomic
 295 execution of the computation represented by U .

296 3.3. Semantic values

297 The predomain \mathbb{V} of untyped values is the least solution of the following domain
 298 equation:

$$\mathbb{V} \simeq \mathbb{V}\mathbb{B} + (\mathbb{V} \rightarrow T\mathbb{V}) + \mathbb{V}^*$$

299 That is, values are either R-values, continuous functions from values to computations
 300 ($T\mathbb{V}$), or tuples of values. We tend to identify the summands of the right hand side with
 301 subsets of \mathbb{V} but may use tags like $fun(f) \in \mathbb{V}$ when $f : \mathbb{V} \rightarrow T\mathbb{V}$ to avoid ambiguity.

302 We have families of deflations $p_i : \mathbb{V} \rightarrow \mathbb{V}$ and $q_i : T\mathbb{V} \rightarrow T\mathbb{V}$, referred to
 303 as canonical deflations, so that $(p_i)_i$ and $(q_i)_i$ are ascending chains converging to the
 304 identity. The definition is entirely standard and may be found in the technical report
 305 [7]. It shows in particular that \mathbb{V} and $T\mathbb{V}$ are *bifinite* (equivalently SFP) (pre-)domains
 306 [3] and as such also Scott (pre-) domains. The presence of these deflations allows us to
 307 apply Lemma 3.3 and simplifies reasoning in general.

308 The semantics of values $\llbracket v \rrbracket \in \mathbb{V} \rightarrow \mathbb{V}$ and terms $\llbracket t \rrbracket \in \mathbb{V} \rightarrow T\mathbb{V}$ are given by the
 309 recursive clauses in Figure 3. Environments, ρ , are properly tuples of values; we abuse
 310 notation slightly by treating them as maps from variables, x , to values, v , (and write
 311 $\rho[x \mapsto v]$ for functional update) to avoid mentioning an explicit context in which untyped
 312 terms are well-formed. The last clause applies to semantically ill-typed programs, for
 313 example:

$$\llbracket \text{if } v \text{ then } e_1 \text{ else } e_2 \rrbracket \rho$$

314 when $\llbracket v \rrbracket \rho$ does not return a boolean value, but, e.g., a number or a location.

$$\begin{aligned}
\llbracket x \rrbracket \rho &= \rho(x) \\
\llbracket v_r \rrbracket \rho &= v_r \\
\llbracket (v_1, v_2) \rrbracket \rho &= (\llbracket v_1 \rrbracket \rho, \llbracket v_2 \rrbracket \rho) \\
\llbracket v.i \rrbracket \rho &= d_i \text{ if } i = 1, 2, \llbracket v \rrbracket \rho = (d_1, d_2) \\
\llbracket c \rrbracket \rho &= \text{fun}(f) \\
&\quad \text{where } f(v) = \text{rtn}(F_c(v)) \text{ if } F_c(v) \text{ is defined} \\
&\quad \text{and } f(v) = \emptyset, \text{ otherwise.} \\
\llbracket \text{rec } f \ x = e \rrbracket \rho &= \text{fun}(g^\ddagger(\rho)) \\
&\quad \text{where } g(\rho, u) = \lambda d. \llbracket e \rrbracket \rho[f \mapsto u, x \mapsto d] \\
\llbracket v \rrbracket \rho &= 0, \text{ otherwise} \\
\llbracket v \rrbracket \rho &= \text{rtn}(\llbracket v \rrbracket \rho) \\
\llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket \rho &= \text{bnd}(\lambda d. \llbracket e_2 \rrbracket \rho[x \mapsto d], \llbracket e_1 \rrbracket \rho) \\
\llbracket v_1 \ v_2 \rrbracket \rho &= \llbracket v_1 \rrbracket \rho(\llbracket v_2 \rrbracket \rho) \\
\llbracket \text{if } v \text{ then } e_1 \text{ else } e_2 \rrbracket \rho &= \llbracket e_1 \rrbracket \rho, \text{ if } \llbracket v \rrbracket \rho = \text{true} \\
\llbracket \text{if } v \text{ then } e_1 \text{ else } e_2 \rrbracket \rho &= \llbracket e_2 \rrbracket \rho, \text{ if } \llbracket v \rrbracket \rho = \text{false} \\
\llbracket !v \rrbracket \rho &= \text{fromstate}(\lambda h. (h, h(X))), \text{ when } \llbracket v \rrbracket \rho = X \\
\llbracket v_1 := v_2 \rrbracket \rho &= \text{fromstate}(\lambda h. (h[X \mapsto \llbracket v_2 \rrbracket \rho], ())), \text{ if } \llbracket v_1 \rrbracket \rho = X \\
\llbracket \text{ref}(v) \rrbracket \rho &= \text{fromstate}(\lambda h. \text{new}(h, \llbracket v \rrbracket \rho)) \\
\llbracket \text{atomic}(e) \rrbracket \rho &= \text{at}(\llbracket e \rrbracket \rho) \\
\llbracket e_1 \parallel e_2 \rrbracket \rho &= \llbracket e_1 \rrbracket \rho \mid \llbracket e_2 \rrbracket \rho \\
\llbracket e \rrbracket \rho &= \emptyset, \text{ otherwise}
\end{aligned}$$

Figure 3: Denotational semantics

315 **4. Abstract Locations**

316 We build on the concept of abstract locations defined by Benton et al [6]. These
 317 allow complicated data structures that span several concrete locations, or only parts
 318 of them, to be regarded as a single “location” that can be written to and read from.
 319 Essentially, an abstract location is given by a partial equivalence relation on heaps
 320 modelling well-formedness and equality together with a transitive relation modelling
 321 allowed modifications of the abstract location. Abstract locations then allow certain
 322 commands that modify the physical heap to be treated as read-only or even pure if they
 323 respect the contracts. Abstract locations are related to *islands* [4] which also allow one
 324 to specify heap allocated data structures and use transition systems for that purpose.
 325 An important difference is that abstract locations do not require physical footprints in
 326 the form of sets of concrete locations.

327 Due to the absence of dynamic allocation at the level of abstract locations in the
 328 present paper, we can slightly simplify the original definition [6], dropping those ax-
 329 ioms that involve the interaction with dynamic allocation.² On the other hand, in the
 330 presence of concurrency, we need *two* partial equivalence relations: one that models
 331 semantic equivalence and well-formedness and a finer one that constrains the heap
 332 modifications that other concurrent computations that are independent of the given ab-
 333 stract locations are allowed to do *while* an operation on the abstract location is ongoing,
 334 but temporarily preempted.

335 **Definition 4.1** (Concurrent Abstract Location). *A concurrent abstract location \mathfrak{l} con-*
 336 *sists of the following data:*

337 (1) *a partial equivalence relation $\overset{\mathfrak{l}}{\sim}$ on \mathbb{H} modeling the “semantic equivalence” on*
 338 *the bits of the store that \mathfrak{l} uses. If $h \overset{\mathfrak{l}}{\sim} h'$ then the same computation started on h and*
 339 *h' , respectively, will yield related or even equal results.*

340 (2) *a partial equivalence relation $\overset{\mathfrak{l}}{=}$ on \mathbb{H} refining $\overset{\mathfrak{l}}{\sim}$ and modeling the “strict equiv-*
 341 *alence” on the bits of the store that \mathfrak{l} uses. If a concurrent computation on \mathfrak{l} has reached*
 342 *h and is preempted, then another computation may replace h with h' where $h \overset{\mathfrak{l}}{=} h'$ and*
 343 *then the original computation on \mathfrak{l} may resume on h' without the final result being*
 344 *compromised.*

345 (3) *a transitive (and reflexive on the support of $\overset{\mathfrak{l}}{\sim}$) relation $\overset{\mathfrak{l}}{\rightarrow}$ modeling how exactly*
 346 *the heap may change upon writing the abstract location and in particular what bits of*
 347 *the store such writes leave intact. In other words, if $h \overset{\mathfrak{l}}{\rightarrow} h_1$ then h_1 might arise by*
 348 *writing to \mathfrak{l} in h and all possible writes are specified by $\overset{\mathfrak{l}}{\rightarrow}$. We call $\overset{\mathfrak{l}}{\rightarrow}$ the step relation*
 349 *of \mathfrak{l} .*

350 *In addition, we require the following conditions where $h : \mathfrak{l}$ stands for $h \overset{\mathfrak{l}}{\sim} h$.*

- 351 1. *If $h : \mathfrak{l}$ then $h \overset{\mathfrak{l}}{=} h$;*

²Though our examples do all satisfy these axioms, leaving the way open to a future extension with dynamically allocation of abstract locations and concurrency.

352 2. if $h \xrightarrow{\dagger} h_1$ then $h : \dagger$ and $h_1 : \dagger$.

353 If $h \xrightarrow{\dagger} h_1$ and at the same time $h \stackrel{\dagger}{=} h_1$, then we say that h_1 arises from h by a silent
354 move in \dagger . Our semantic framework will permit silent moves at all times.

355 We now introduce some examples of abstract locations.

356 *Single Integer.* For our simplest example, consider the following abstract location
357 parametric with respect to concrete location X as follows:

$$\begin{aligned} h \stackrel{\text{int}(X)}{\sim} h' &\iff \exists n. h(X) = \text{int}(n) \wedge h'(X) = \text{int}(n) \\ h \stackrel{\text{int}(X)}{=} h' &\iff h \stackrel{\text{int}(X)}{\sim} h' \\ h \xrightarrow{\text{int}(X)} h_1 &\iff \\ &h : \text{int}(X), h_1 : \text{int}(X) \text{ and } \forall X' \in \mathbb{L}. X' \neq X \Rightarrow h(X') = h_1(X') \end{aligned}$$

358 Two heaps are semantically equivalent (w.r.t. $\text{int}(X)$ that is) if the values stored in X
359 are integers and equal; the step relation requires all other concrete locations to be un-
360 changed.

361 We will sometimes abuse notation and write rd_X, wr_X, co_X for $rd_{\text{int}(X)}, wr_{\text{int}(X)}, co_{\text{int}(X)}$.

362 *Overlapping references.* Let X be a concrete location encoding a pair of integer values
363 using a bijection p . We define the abstract location $\dagger\text{st}(X)$ as below. We omit $\text{snd}(X)$
364 which is similar, but only looks at the second projection, instead of the first.

$$\begin{aligned} h \stackrel{\dagger\text{st}(X)}{\sim} h' &\iff \exists a_1 a_2 a'_1 a'_2 \in \mathbb{Z}. h(X) = p^{-1}(a_1, a_2) \wedge \\ &h'(X) = p^{-1}(a'_1, a'_2) \wedge a_1 = a'_1 \\ h \stackrel{\dagger\text{st}(X)}{=} h' &\iff h \stackrel{\dagger\text{st}(X)}{\sim} h' \\ h \xrightarrow{\dagger\text{st}(X)} h_1 &\iff h : \dagger\text{st}(X), h_1 : \dagger\text{st}(X) \text{ and} \\ &(\forall X' \neq X. h(X') = h_1(X')) \wedge (\forall a_1 a_2 a'_1 a'_2 \in \mathbb{Z}. h(X) = p^{-1}(a_1, a_2) \wedge \\ &h_1(X) = p^{-1}(a'_1, a'_2) \Rightarrow a_2 = a'_2) \end{aligned}$$

365 The semantic (and strict) equivalence of $\dagger\text{st}(X)$ (respectively, $\text{snd}(X)$) specifies that two
366 heaps h and h' are equivalent whenever they both store a pair of values in X and the
367 first projections (respectively, second projection) of these pairs are the same. The step
368 relation of $\dagger\text{st}(X)$ (respectively, $\text{snd}(X)$) specifies that it keeps all other locations alone
369 and does not change the second projection (respectively, first projection) of the pair
370 stored at location X .

371 *Version Numbers.* The abstract location \mathfrak{X} consists of two concrete locations X_{Val} and
372 X_{Ver} and its relations are specified as follows:

$$\begin{aligned} h \stackrel{\mathfrak{X}}{\sim} h' &\iff h(X_{Val}) = h'(X_{Val}) \\ h \stackrel{\mathfrak{X}}{=} h' &\iff h \stackrel{\mathfrak{X}}{\sim} h' \\ h \xrightarrow{\mathfrak{X}} h_1 &\iff \forall X' \notin \{X_{Ver}, X_{Val}\}. h(X') = h_1(X') \wedge \\ &h : \mathfrak{X} \wedge h_1 : \mathfrak{X} \wedge h(X_{Ver}) \leq h_1(X_{Ver}) \wedge \\ &[h(X_{Val}) \neq h_1(X_{Val}) \Rightarrow h(X_{Ver}) < h_1(X_{Ver})] \end{aligned}$$

373 Two heaps are semantically equivalent if they have the same value (independent of the
 374 version number). The step relation specifies that the version number does not decrease
 375 and it increases if the value changes.

376 *Loop Parallelization.* For a concrete location X , we introduce two concurrent abstract
 377 locations $\text{listeven}(X)$ and $\text{listodd}(X)$, which only look, respectively, at the elements in
 378 the even and odd positions of the linked list pointed to by X . Formally, let $L(X, h)$
 379 denote that $h(X)$ points to a well formed linked list of integers of length $L(X, h).len$ and
 380 locations $L(X, h).locs$ and that $L(X, h)[i]$ is the i^{th} node of the list for $1 \leq i \leq L(X, h).len$.
 381 The relations for $\text{listeven}(X)$ are as below. We omit the relations for $\text{listodd}(X)$, which
 382 are similar.

$$\begin{aligned}
 h \stackrel{\text{listeven}(X)}{\sim} h' &\iff L(X, h) \wedge L(X, h') \wedge L(X, h).len = L(X, h').len \wedge \\
 &L(X, h)[2i] = L(X, h')[2i] \\
 &\text{for } 0 \leq i \leq \lfloor L(X, h).len/2 \rfloor \\
 h \stackrel{\text{listeven}(X)}{=} h' &\iff h \stackrel{\text{listeven}(X)}{\sim} h' \\
 h \xrightarrow{\text{listeven}(X)} h_1 &\iff h : \text{listeven}(X) \wedge h_1 : \text{listeven}(X) \wedge \\
 &L(X, h).len = L(X, h_1).len \\
 &\text{for } 0 \leq i \leq \lfloor L(X, h).len/2 \rfloor \\
 &L(X, h)[2i + 1] = L(X, h_1)[2i + 1] \wedge \\
 &L(X, h)[2i].next = L(X, h_1)[2i].next \wedge \\
 &\forall X' \notin L(X, h).locs. h(X') = h_1(X')
 \end{aligned}$$

383 The step relation $h \xrightarrow{\text{listeven}(X)} h_1$ specifies that $h : \text{listeven}(X)$ and that h_1 arises from h
 384 by possibly modifying the list entries at even positions leaving everything else alone.

385 *Michael-Scott queue.* For concrete location X we introduce a concurrent abstract lo-
 386 cation $\text{msq}(X)$ first informally as follows: we have $h \stackrel{\text{msq}(X)}{\sim} h'$ if both h and h' contain
 387 a well-formed MSQ rooted at X and these queues contain the same entries in the same
 388 order. They may, however, use different locations for the nodes and also have different
 389 garbage tails.

390 The relation $h \stackrel{\text{msq}(X)}{=} h'$ asserts that h and h' are identical on the part reachable
 391 and co-reachable from X via *next* pointers. This means that while an MSQ operation is
 392 working on the queue no concurrent operation working elsewhere is allowed to relocate
 393 the queue or remove the garbage trail which would be the case if we merely required
 394 that such operations do not change the $\text{MSQ}^{\sim}(X)$ -class.

395 The relation $\xrightarrow{\text{msq}(X)}$, finally, is defined as the transitive closure of the actions of
 396 operations on the MSQ: adding nodes at the tail and moving nodes from the head to
 397 the garbage tail.

398 We now give a formal definition. We represent pointers *head*, *next*, *elem* using
 399 some layout convention, e.g. $v.head = v.1$, etc. We then define

$$h, X \xrightarrow{\text{next}} X' \iff X' \text{ can be reached from } X \text{ in } h \\
 \text{by following a chain of next pointers}$$

400 We use $List(X, h, (X_0, \dots, X_n), (v_1 \dots, v_n))$ to signal that $h(X)$ points to a linked list with
 401 nodes X_0, \dots, X_n and entries v_1, \dots, v_n . Note that the first node X_0 acts as a sentinel and
 402 its *elem* component is ignored. Formally:

$$\begin{aligned} h(X).head &= X_0 & h(X_i).elem &= v_i \text{ for } i = 1, \dots, n \\ h(X_i).next &= X_{i+1} \text{ for } i = 0, \dots, n-1 & h(X_n).next &= null \end{aligned}$$

403 We define $fp(X, h)$ as the set of locations reachable and co-reachable from X via *next*,
 404 formally:

$$fp(X, h) = \{X' \mid X \xrightarrow{next} X' \vee X' \xrightarrow{next} X\}$$

405 Finally, we define $snoc(h, h', X, v)$ to mean that h' arises from h by attaching a new
 406 node containing v at the end of the list pointed to by X in h . Thus, in particular,
 407 $List(X, h, (X_0, \dots, X_n), (v_1 \dots, v_n))$ implies $List(X, h', (X_0, \dots, X_n, X_{n+1}), (v_1 \dots, v_n, v))$ for
 408 some $X_{n+1} \notin \text{dom}(h)$. We omit the obvious frame conditions. We now define

$$\begin{aligned} h \stackrel{\text{msg}(X)}{\sim} h' &\iff \exists \vec{X} \vec{X}' \exists \vec{v}. List(X, h, \vec{X}, \vec{v}) \wedge List(X, h', \vec{X}', \vec{v}) \\ h \stackrel{\text{msg}(X)}{=} h' &\iff h \stackrel{\text{msg}(X)}{\sim} h' \wedge \forall X' \in fp(X, h). h(X') = h'(X') \\ h \xrightarrow{\text{msg}(X)} h_1 &\iff h : \text{msg}(X) \wedge h_1 : \text{msg}(X) \wedge \text{step}^*(h, h_1) \\ \text{step}(h, h_1) &\iff \forall X' \neq X. h(X') = h_1(X') \wedge \\ & [h_1(X) = h(X).next \vee \exists v. snoc(h, h_1, X, v)] \end{aligned}$$

409 In all of these examples, the only silent moves are identity moves. This is not so
 410 in the examples from [6] which contained data-structures that would reorganize during
 411 lookups and also patterns like late initialisation.

412 4.1. Worlds

413 We will group the abstract locations used to describe a program into a *world*. In
 414 this paper we do not model dynamic evolution of worlds; all abstract locations ever
 415 used must be set up upfront. While allocation of concrete locations may happen to
 416 increase a data structure modelled by an abstract location, e.g. in the Michael-Scott
 417 Queue example, no new such datastructures can appear. It is possible, however, to
 418 extend our work in this direction by using (proof-relevant) Kripke logical relations
 419 [6, 4].

420 **Definition 4.2** (world). *A world is a set of abstract locations.*

421 *The relation $h \models w$ (heap h satisfies world w) is defined as the largest relation such*
 422 *that $h \models w$ implies*

- 423 • $h : \dagger$ for all $\dagger \in w$;
- 424 • if $\dagger \in w$ and $h \xrightarrow{\dagger} h_1$ then $h \stackrel{\dagger}{=} h_1$ holds for all $Y \in w$ with $Y \neq \dagger$ and $h_1 \models w$.

425 The original account of abstract locations [6] also has a notion of independence
 426 of locations which facilitates reasoning in the presence of dynamic allocation, and in
 427 particular permitted relocation of abstract locations. Since we are not currently treating
 428 dynamic allocation of abstract locations, we can avoid this notion here.

429 We remark that if our world w contains two obviously “dependent” abstract lo-
 430 cations, e.g. has both an integer location and a boolean location placed at the same
 431 physical location, then there will be no heap h such that $h \models w$.

432 We assume a fixed *current* world w which may appear in definitions without being
 433 notationally reflected. See also Assumption 1.

434 5. Effects

435 For each abstract location \mathfrak{l} we have three elementary effects $rd_{\mathfrak{l}}$ (reading from \mathfrak{l}),
 436 $wr_{\mathfrak{l}}$ (writing to \mathfrak{l}), and $co_{\mathfrak{l}}$ (chaotic or concurrent access). The chaotic access is similar
 437 to writing, but allows writes that are not in sync. For example, $e_1 = X := 1$ and
 438 $e_2 = X := 2$ both have individually the wr_X effect, but e_1 and e_2 are distinguishable
 439 with a context that assumes the wr_X -effect. Thus, e_1 and e_2 are not equal “at type” wr_X .
 440 At type co_X they are, however, equal, because a context that copes with this effect may
 441 not assume that both produce equal results.

442 We use the $co_{\mathfrak{l}}$ effect to tell the environment not to look at a particular location
 443 during a concurrent computation. For example, we will be able to show that $X :=$
 444 $!X + 1; X := !X + 1$ is equivalent to $X := !X + 2$ “at type” $\text{unit} \ \& \ co_X \mid \varepsilon \mid \varepsilon \cup$
 445 $\{rd_X, wr_X\}$ whenever $X \notin \text{locs}(\varepsilon)$. This means that the two computations are indistin-
 446 guishable by environments that do not read, let alone modify X during the computation
 447 and assume regular read-write access once it is completed. It would alternatively be
 448 possible to replace the *co*-effect using a special set of private locations akin to the
 449 private regions from [12].

450 We use the notation $\text{rds}(\varepsilon)$, $\text{wrs}(\varepsilon)$, $\text{cos}(\varepsilon)$ to refer to the abstract locations \mathfrak{l} for
 451 which ε contains $rd_{\mathfrak{l}}$, $wr_{\mathfrak{l}}$, and $co_{\mathfrak{l}}$, respectively. We write $\text{locs}(\varepsilon) := \text{rds}(\varepsilon) \cup \text{wrs}(\varepsilon) \cup$
 452 $\text{cos}(\varepsilon)$. We also write ε^C for ε with all read effects removed and each $wr_{\mathfrak{l}}$ in ε replaced
 453 by $co_{\mathfrak{l}}$.

454 **Definition 5.1.** *An effect ε is well-formed (with respect to the current world) if $\text{locs}(\varepsilon) \subseteq$
 455 w and $\text{rds}(\varepsilon) \cap \text{cos}(\varepsilon) = \emptyset$ and $\text{cos}(\varepsilon) \subseteq \text{wrs}(\varepsilon)$. An effect specification is a triple
 456 $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ of well-formed effects such that $\varepsilon_2 \subseteq \varepsilon_3$.*

457 An effect specification $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ approximates the behaviour of a computation e
 458 in the following way: the effect ε_1 summarizes side effects that may occur during the
 459 execution of e (corresponding to a guarantee condition in the rely-guarantee formalism
 460 [16]); the effect ε_2 summarizes effects of the interacting environment that e can tolerate
 461 while still functioning as expected (corresponding to a rely condition). Finally, ε_3
 462 summarizes the side effects that may occur between start and completion of e . All
 463 the effects that the environment might introduce must be recorded in ε_3 because they
 464 are not under “our” control and might happen at any time even as the very last thing
 465 before the final result is returned. The effects flagged in ε_1 , on the other hand, do
 466 not necessarily show up in ε_3 , for a computation might be able to clean up those effects
 467 prior to returning the final result. The requirement that $\text{rds}(\varepsilon) \cap \text{cos}(\varepsilon) = \emptyset$ is owed to the
 468 fact that all effects should preserve their own precondition, however the precondition
 469 of $rd_{\mathfrak{l}}$ is agreement on \mathfrak{l} which is not preserved by $co_{\mathfrak{l}}$. The requirement $\text{cos}(\varepsilon) \subseteq \text{wrs}(\varepsilon)$
 470 reflects the fact that $\text{cos}(\mathfrak{l})$ includes $wr_{\mathfrak{l}}$ as a special case.

471 Note that if $\varepsilon^C \cup \varepsilon_1$ is a (well-formed) effect, then it is the case that $\text{rds}(\varepsilon_1) \cap$
 472 $(\text{wrs}(\varepsilon) \cup \text{cos}(\varepsilon)) = \emptyset$. We will use this observation to simplify some side conditions.

473 In our concrete examples, we abbreviate $\{co_1\} \cup \{wr_1\}$ by just co_1 , in other words,
 474 the chaotic effect silently implies the write effect.

475 Consider the computations $e_1 = X := !X + 1; X := !X + 1$ and $e_2 = X := !X + 2$.
 476 Let ε_X stand for $\{rd_X, wr_X\}$ and analogously ε_Y . Each of the two computations can be
 477 assigned the effect $(\varepsilon_X, \varepsilon_Y, \varepsilon_X \cup \varepsilon_Y)$, but they are distinguishable at that effect typing.
 478 Under the looser specification $(\{co_{\varepsilon_X}\}, \varepsilon_Y, \varepsilon_X \cup \varepsilon_Y)$, however, they are indistinguishable,
 479 and our semantics is able to validate this equivalence, see Example 7.5.

480 Finally, consider the program $e = !X$ that simply reads a location storing an integer.
 481 We can show that this program has type $\mathbb{Z} \ \& \ \emptyset \mid \varepsilon \mid \varepsilon, rd_X$, where the read effect on X
 482 is only in the global effects.

483 *Notations..* For any well-formed effects $\varepsilon, \varepsilon'$ we use the notation $\varepsilon \perp \varepsilon'$ to mean that
 484 $\text{rds}(\varepsilon) \cap \text{wrs}(\varepsilon') = \text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon) = \text{wrs}(\varepsilon) \cap \text{wrs}(\varepsilon') = \emptyset$. Note that this implies in
 485 particular $\text{cos}(\varepsilon) \cap \text{rds}(\varepsilon') = \emptyset$, etc. Intuitively, two programs exhibiting effects ε and
 486 ε' , respectively, commute with each other. We write $h \stackrel{\text{rds}(\varepsilon)}{\sim} h'$ to mean $h \stackrel{\perp}{\sim} h'$ for each
 487 $l \in \text{rds}(\varepsilon)$. We write $\xrightarrow{\varepsilon}$ for the transitive closure of $\bigcup_{l \in \text{wrs}(\varepsilon)} \xrightarrow{l} \cup \bigcup_{l \in \text{w}} \xrightarrow{l} \cap \stackrel{\perp}{\rightarrow}$. Thus,
 488 $\xrightarrow{\varepsilon}$ allows steps by locations recorded as writing in ε and silent steps by all locations in
 489 the current world.

490 We define the notation $\varepsilon_1 \sqcup \varepsilon_2$ which appears in the parallel congruence rule by

$$\varepsilon_1 \sqcup \varepsilon_2 = \varepsilon_1 \cup \varepsilon_2 \setminus \{wr_\ell \mid wr_\ell \notin \varepsilon_1 \cap \varepsilon_2\} \setminus \{co_\ell \mid co_\ell \notin \varepsilon_1 \cap \varepsilon_2\}$$

491 6. Typing and congruence rules

492 Types are given by the grammar

$$\tau ::= \text{unit} \mid \text{int} \mid \text{bool} \mid A \mid \tau_1 \times \tau_2 \mid \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2$$

493 where A ranges over user-specified abstract types. They will typically include reference
 494 types such as `intref` and also types like lists, sets, and even objects. In $\tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2$
 495 the triple of effects $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ must be an effect specification.

496 We use two judgments:

- 497 • $\Gamma \vdash v \leq v' : \tau$ specifying that values v and v' have type τ and that v approximates
 498 v' ,
- 499 • $\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ specifying that the programs e and e' under the
 500 context Γ have type τ , with the effect specification $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ specifying, respec-
 501 tively, the effects during execution, the effects of the interacting environment and
 502 the start and completion effects. Moreover, e approximates e' at this specifica-
 503 tion.

504 We assume an ambient set of *axioms* each having the form (v, v', τ) where v, v' are
 505 values in the metalanguage and τ is a type meaning that v and v' are claimed to be
 506 of type τ and that v approximates v' . This must then be proved “manually” using the
 507 semantics rather than using the rules. We assume that whenever (v, v', τ) an axiom,
 508 then so are (v, v, τ) and (v', v', τ) .

509 We also define typing judgements $\Gamma \vdash v : \tau$ and $\Gamma \vdash e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ which
 510 denote the special case when $\Gamma \vdash v \leq v : \tau$ and $\Gamma \vdash e \leq e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ can be
 511 derived from the rules from Figure 6. We do not formulate explicit typing rules to save
 512 space.

513 The plan is to justify all the rules semantically using a logical relation (Section 7)
 514 and to then conclude their soundness w.r.t. typed observational approximation and equiv-
 515 alence (Section 8).

516 The parallel composition rule states that two programs e_1 and e_2 can be composed
 517 when their internal effects are not conflicting in the sense that the internal effects of
 518 one program appear as environment interaction effects of the other program. Note the
 519 relationship to the parallel composition rule of the rely-guarantee formalism [16]. Also
 520 note that the effects of computations e_1 and e_2 are not required to be independent from
 521 each other as we do in the parallization rule further down.

522 The appearance of the \sqcup -operation deserves special mention. It might be, for ex-
 523 ample, that e_1 modifies X on the way, thus $wr_X \in \varepsilon_1$ but cleans up this modification by
 524 eventually restoring the old value of X . This would be reflected by $wr_X \notin \varepsilon \cup \varepsilon' \cup \varepsilon_2$.
 525 In that case, we would not expect to see wr_X in the end-to-end effect of the parallel
 526 composition and that is precisely what \sqcup achieves.

527 The rules labelled (Sem) make available all kinds of program transformations that
 528 are valid on the level of the *untyped* denotational semantics, including commuting con-
 529 versions for let and if, fixpoint unrolling, and beta and eta equalities.

530 Finally, we have several effect-dependent (in)equalities: the parallelization rule
 531 generalises a similar rule from [12]. The other ones are concurrent version of analogous
 532 rules for sequential computation that have been analysed in previous work [9, 8, 29, 6]
 533 and are at the basis of all kinds of compiler optimizations. The side conditions on the
 534 effects are rather subtle and much less obvious than those found in a sequential setting.
 535 The parallelization rule is similar to the parallel congruence rule in that it requires the
 536 participating computations to mutually tolerate each other. This time, however, since
 537 the two computations being compared will do rather different things temporarily they
 538 must be oblivious against chaotic access, hence the $(-)^C$ strengthenings in the premise.

539 The reason for the appearance of $(-)^C$ in the other rules is similar. The rule for
 540 pure lambda hoist seems unusual and will thus be explained in more detail. First, the
 541 computation e_1 to be hoisted may indeed have side effects ε_1 so long as they are cleaned
 542 up by the time e_1 completes and the intervening environment does not notice (modelled
 543 by the conditions $\varepsilon_1 \perp \varepsilon$ and final effect $\varepsilon^C = \varepsilon^C \cup \emptyset$). In the conclusion the transient
 544 effect ε_1 shows up again, but $(-)^C$ -ed since it only appears in different sides. Also in
 545 the other rules like commuting etc. it is the case that the familiar side conditions on
 546 applicability only affect the end-to-end effects whereas the transient effects are merely
 547 required not to interfere with the environment.

$$\begin{array}{c}
\overline{\Gamma \vdash \text{true} \leq \text{true} : \text{bool}} \quad \overline{\Gamma \vdash \text{false} \leq \text{false} : \text{bool}} \quad \overline{\Gamma \vdash n \leq n : \text{int}} \\
\\
\overline{\Gamma, x : \tau \vdash x \leq x : \tau} \quad \frac{\Gamma \vdash v \leq v' : \tau}{\Gamma \vdash v \leq v' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \quad \frac{\Gamma \vdash v \leq v' : \tau_1 \times \tau_2}{\Gamma \vdash v.i \leq v'.i : \tau_i} \\
\frac{\Gamma \vdash e_1 \leq e_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \Gamma \vdash e_1 \leq e_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash e_1 \leq e_3 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \\
\frac{\Gamma \vdash v_i \leq v'_i : \tau_i \ i = 1, 2}{\Gamma \vdash (v_1, v_2) \leq (v'_1, v'_2) : \tau_1 \times \tau_2} \\
\frac{\Gamma \vdash v_1 \leq v'_1 : \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2 \quad \Gamma \vdash v_2 \leq v'_2 : \tau_1}{\Gamma \vdash v_1 \ v_2 \leq v'_1 \ v'_2 : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \\
\frac{\Gamma \vdash v \leq v' : \text{bool} \quad \Gamma \vdash e_1 \leq e'_1 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \Gamma \vdash e_2 \leq e'_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 \leq \text{if } v' \text{ then } e'_1 \text{ else } e'_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \\
\frac{\Gamma \vdash e_1 \leq e'_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \Gamma, x : \tau_1 \vdash e_2 \leq e'_2 : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \leq \text{let } x = e'_1 \text{ in } e'_2 : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \\
\frac{\Gamma \vdash e_1 \leq e'_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon \cup \varepsilon_2 \mid \varepsilon \cup \varepsilon_2 \cup \varepsilon' \quad \Gamma \vdash e_2 \leq e'_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon \cup \varepsilon_1 \mid \varepsilon \cup \varepsilon_1 \cup \varepsilon'}{\Gamma \vdash e_1 \parallel e_2 \leq e'_1 \parallel e'_2 : \tau_1 \times \tau_2 \ \& \ \varepsilon_1 \cup \varepsilon_2 \mid \varepsilon \mid \varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2)} \\
\frac{\Gamma \vdash e_1 \leq e_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \llbracket e_1 \rrbracket = \llbracket e'_1 \rrbracket \quad \llbracket e_2 \rrbracket = \llbracket e'_2 \rrbracket}{\Gamma \vdash e'_1 \leq e'_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \text{ Sem} \\
\frac{\Gamma, f : \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2, x : \tau_1 \vdash e \leq e' : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad (v, v', \tau) \text{ an axiom}}{\Gamma \vdash \text{rec } f \ x = e \leq \text{rec } f \ x = e' : \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2 \quad \Gamma \vdash v \leq v' : \tau} \text{ Ax} \\
\frac{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \varepsilon_1 \subseteq \varepsilon'_1 \quad \varepsilon'_2 \subseteq \varepsilon_2 \quad \varepsilon_3 \subseteq \varepsilon'_3}{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon'_1 \mid \varepsilon'_2 \mid \varepsilon'_3} \\
\frac{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \emptyset \mid \varepsilon_3}{\Gamma \vdash \text{atomic}(e) \leq \text{atomic}(e') : \tau \ \& \ \varepsilon_3 \mid \varepsilon_2 \mid \varepsilon_2 \cup \varepsilon_3} \text{ Atom}
\end{array}$$

Figure 4: Typing and congruence rules

$$\begin{array}{c}
\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \cup \varepsilon_2^C \mid \varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1 \\
\Gamma \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon^C \cup \varepsilon_1^C \mid \varepsilon^C \cup \varepsilon_1^C \cup \varepsilon'_2 \\
\varepsilon_1 \perp \varepsilon_2 \quad \varepsilon_1 \perp \varepsilon \quad \varepsilon_2 \perp \varepsilon \\
\hline
\Gamma \vdash e_1 \parallel e_2 \leq (\text{let } x=e_1 \text{ in let } y=e_2 \text{ in } (x,y)) : \tau_1 \times \tau_2 \ \& \ \varepsilon_1^C \cup \varepsilon_2^C \mid \varepsilon \mid \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2 \quad \text{Parallelization}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \mid \varepsilon^C \cup \varepsilon'_1 \\
\Gamma \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon^C \mid \varepsilon^C \cup \varepsilon'_2 \\
\varepsilon'_1 \perp \varepsilon'_2 \quad \varepsilon_1 \perp \varepsilon \quad \varepsilon_2 \perp \varepsilon \\
\hline
\Gamma \vdash (\text{let } x=e_1 \text{ in let } y=e_2 \text{ in } (x,y)) = \\
(\text{let } y=e_2 \text{ in let } x=e_1 \text{ in } (x,y)) : \tau_1 \times \tau_2 \ \& \ \varepsilon_1^C \cup \varepsilon_2^C \mid \varepsilon \mid \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2 \quad \text{Commuting}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2^C \mid \varepsilon_2^C \cup \varepsilon' \quad \text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon') = \emptyset \quad \varepsilon_2 \perp \varepsilon_1 \\
\hline
\Gamma \vdash (\text{let } x=e \text{ in } (x,x)) \leq \\
(\text{let } x=e \text{ in let } y=e \text{ in } (x,y)) : \tau \times \tau \ \& \ \varepsilon_1^C \mid \varepsilon_2 \mid \varepsilon_2 \cup \varepsilon' \quad \text{Duplicated}
\end{array}$$

$$\begin{array}{c}
\frac{(v, v', \tau) \text{ an axiom}}{\Gamma \vdash v \leq v' : \tau} \text{ Ax} \\
\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \mid \varepsilon^C \\
\Gamma, x : \tau_3, y : \tau_1 \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon \mid \varepsilon \cup \varepsilon_2 \\
\varepsilon \perp \varepsilon_1 \\
\hline
\Gamma \vdash \text{let } y=e_1 \text{ in } \lambda x. e_2 \leq \lambda x. \text{let } y=e_1 \text{ in } e_2 : \tau_3 \xrightarrow[\varepsilon]{\varepsilon_1^C \cup \varepsilon_2 \mid \varepsilon \cup \varepsilon_3} \tau_2 \ \& \ \varepsilon_1^C \mid \varepsilon \mid \varepsilon \quad \text{Lambda Hoist}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \mid \varepsilon^C \cup \varepsilon'_1 \\
\Gamma \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon \mid \varepsilon'_2 \\
\varepsilon_1 \perp \varepsilon \\
\text{wrs}(\varepsilon'_1) = \emptyset \\
\hline
\Gamma \vdash e_2 \leq (\text{let } x=e_1 \text{ in } e_2) : \tau_2 \ \& \ \varepsilon_1^C \cup \varepsilon_2 \mid \varepsilon \mid \varepsilon \cup \varepsilon'_2 \quad \text{Deadcode}
\end{array}$$

Figure 5: Effect-dependent transformations.

548 The following definitions provide the semantics of our effect annotations.

549 **Definition 6.1** (Tiling). *Let $w \vdash \varepsilon$. We write $[\varepsilon](h, h', h_1, h'_1)$ to mean that (i) $h \models w \Rightarrow$
550 $h \xrightarrow{\varepsilon} h_1$ and (ii) $h' \models w \Rightarrow h' \xrightarrow{\varepsilon} h'_1$ and (iii) $h \stackrel{\text{rds}(\varepsilon)}{\sim} h'$ and $\dagger \in \text{wrs}(\varepsilon) \setminus \text{cos}(\varepsilon)$ imply
551 $(h \stackrel{\dagger}{=} h_1 \wedge h' \stackrel{\dagger}{=} h'_1) \vee h_1 \stackrel{\dagger}{\sim} h'_1$.*

552 Thus, assuming semantic consistency of heaps, h and h' evolve to h_1 and h'_1 ac-
553 cording to the modifying (writing or chaotic) locations in ε , and if h, h' agree on the
554 reads of ε then written locations will either be identically modified or left alone.

555 If the step relations of all abstract locations commute with each other then tiling
556 admits an alternative characterisation in terms of preservation of binary relations [9].
557 The present more operational version is inspired by the treatment of effects in [12].

558 **Lemma 6.2.** *Suppose that $w \vdash \varepsilon$, $w \vdash \varepsilon_1$, $w \vdash \varepsilon_2$. The following hold whenever*
559 *well-formed.*

- 560 1. *If $[\varepsilon](h, h', h_1, h'_1)$ and $[\varepsilon](h_1, h'_1, h_2, h'_2)$ then $[\varepsilon](h, h', h_2, h'_2)$;*
- 561 2. *$[\varepsilon](h, h', h, h')$*
- 562 3. *If $\varepsilon_1 \subseteq \varepsilon_2$ then $[\varepsilon_1](h, h', h_1, h'_1) \Rightarrow [\varepsilon_2](h, h', h_1, h'_1)$*
- 563 4. *$[\varepsilon](h, h', h_1, h'_1) \Rightarrow [\varepsilon^C](h, h', h_1, h'_1)$*
- 564 5. *If $[\varepsilon](h, h', k, k')$ and $h \stackrel{\text{rds}(\varepsilon)}{\sim} h'$ then $k \stackrel{\text{rds}(\varepsilon)}{\sim} k'$. (this relies on $\text{rds}(\varepsilon) \cap \text{cos}(\varepsilon) = \emptyset$.)*
- 565 6. *Suppose $[\varepsilon](h, h', h_1, h'_1)$. If $h \models w$ then $h_1 \models w$; if $h' \models w$ then $h'_1 \models w$.*

566 7. Logical Relation

567 **Definition 7.1** (Specifications). *A value specification is a relation $E \subseteq \mathbb{V} \times \mathbb{V}$ such that*

- 568 • *if $x_1 \leq x$ and $y \leq y_1$ and $x E y$ then $x_1 E y_1$ (in short thus $\leq; E; \leq \subseteq E$);*
- 569 • *if $(x_i)_i$ and $(y_i)_i$ are chains such that $x_i E y_i$ then $\sup_i x_i E \sup_i y_i$, i.e., E is an*
570 *admissible subset of $\mathbb{V} \times \mathbb{V}$;*
- 571 • *if $x E y$ then $p_i(x) E p_i(y)$ for each i , i.e. E is closed under the canonical defla-*
572 *tions.*

573 *Similarly, a computation specification is an admissible subset of $T\mathbb{V} \times T\mathbb{V}$ such that*
574 *the relation $Q \subseteq T\mathbb{V} \times T\mathbb{V}$, $\leq; Q; \leq \subseteq Q$ and Q is closed under the canonical deflations*
575 *q_i .*

576 The requirement $\leq; E; \leq \subseteq E$ ensures smooth interaction with the down-closure
577 built into our trace monad. Admissibility is needed for the soundness of recursion
578 and closure under the canonical deflations, finally is needed so that Lemma 3.3 can be
579 applied.

580 **Definition 7.2.** *If $E \subseteq \mathbb{V} \times \mathbb{V}$ and $Q \subseteq T\mathbb{V} \times T\mathbb{V}$ then the relation $E \rightarrow Q \subseteq \mathbb{V} \times \mathbb{V}$ is*
581 *defined by*

$$fE \rightarrow Qf' \iff \forall x x'. (x E x') \Rightarrow (f(x) Q f'(x'))$$

582 *In particular, for $fE \rightarrow Qf'$ to hold, both f, f' must be functions (and not elements of*
583 *base type or tuples).*

584 **Lemma 7.3.** *If E and Q are specifications so is $E \rightarrow Q$.*

585 The following is the crucial definition of this paper; it gives a semantic counterpart
 586 to observational approximation and, due to its game-theoretic flavour, allows for very
 587 intuitive proofs.

588 **Definition 7.4.** *Let $E \subseteq \mathbb{V} \times \mathbb{V}$ be a value specification and $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ an effect spec-*
 589 *ification. We define the relations $T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ and $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ between sets of*
 590 *trace-value pairs, i.e. on $\mathcal{P}(\text{Tr} \times \text{Values})$:*

591 $(U, U') \in T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ if and only if

$$\left[\begin{array}{l} \forall ((h_1, k_1) \dots (h_n, k_n), a) \in U. h_1 \models w \Rightarrow \\ \forall h'_1. h'_1 \models w \Rightarrow h_1 \overset{\text{rds}(\varepsilon_3)}{\sim} h'_1 \Rightarrow \\ \exists k'_1. [\varepsilon_1](h_1, h'_1, k_1, k'_1) \wedge \forall h'_2. [\varepsilon_2](k_1, k'_1, h_2, h'_2) \Rightarrow \\ \exists k'_2. [\varepsilon_1](h_2, h'_2, k_2, k'_2) \wedge \forall h'_3. [\varepsilon_2](k_2, k'_2, h_3, h'_3) \Rightarrow \\ \dots \\ \exists k'_n. [\varepsilon_1](h_n, k_n, h'_n, k'_n) \wedge [\varepsilon_3](h_1, h'_1, k_n, k'_n) \wedge \\ \exists a' \in \mathbb{V}. (a, a') \in E \wedge ((h'_1, k'_1) \dots (h'_n, k'_n), a') \in U' \end{array} \right]$$

592 We define the relation $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3) \subseteq T\mathbb{V} \times T\mathbb{V}$ as the admissible closure of T_0 , i.e.
 593 $\text{Adm}(T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3))$.

594 The game-theoretic view of $T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ may be understood as follows. Given
 595 $U, U' \in T\mathbb{V}$ we can consider a game between a proponent (who believes $(U, U') \in T\mathbb{V}$)
 596 and an opponent who believes otherwise. The game begins by the opponent selecting
 597 an element $((h_1, k_1) \dots (h_n, k_n), a) \in U$ and $h_1 \models w$, the *pilot trace* and a start heap
 598 $h'_1 \models w$ such that $h_1 \overset{\text{rds}(\varepsilon_3)}{\sim} h'_1$ to begin a trace in U' . Then, the proponent answers with
 599 a matching heap k'_1 so that $[\varepsilon_1](h_1, h'_1, k_1, k'_1)$. If $h_1 \overset{\text{rds}(\varepsilon_1)}{\sim} h'_1$ does not hold, proponent
 600 does not need to ensure that writes are in sync. The opponent then plays a heap h'_2 so
 601 that $[\varepsilon_2](k_1, k'_1, h_2, h'_2)$. At this point, it is in the proponent's interest to make sure that
 602 $k_1 \overset{\text{rds}(\varepsilon_2)}{\sim} k'_1$ for otherwise opponent may make “funny” moves.

603 Then, again, proponent plays a heap k'_2 such that $[\varepsilon_1](h_2, h'_2, k_2, k'_2)$ and so on
 604 until, proponent has played k'_n so that $[\varepsilon_1](h_n, h'_n, k_n, k'_n)$. After that final heap has
 605 been played, it is checked that $[\varepsilon_3](h, h', k_n, k'_n)$ holds. If not, proponent loses. If
 606 yes, then proponent must also play a value a' and it is then checked whether or not
 607 $((h'_1, k'_1) \dots (h'_n, k'_n), a') \in U'$ and $(a E a')$. If this is the case or if at any one point in
 608 the game the opponent was unable to move because there exists no appropriate heap
 609 then the proponent has won the game. Otherwise the opponent wins and we have
 610 $(U, U') \in T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ iff the proponent has a winning strategy for that game.

611 We notice that by Lemma 6.2(6) well-formedness of heaps w.r.t. the ambient world
 612 is a global invariant which allows us to refrain from explicitly assuming and asserting
 613 it in subsequent proofs and statements.

614 We now illustrate the game with a few examples.

615 **Example 7.5.** Consider the following programs:

$$e_1 = (X := !X + 1; X := !X + 1) \quad \text{and} \quad e_2 = (X := !X + 2).$$

616 Let $\dagger = \text{int}(X)$ be the abstract location for a single integer stored at X (see Section 4).
 617 Let $E = \llbracket \text{unit} \rrbracket = \{(\dagger, \dagger)\}$ be the value specification for the unit type.

618 We show that $(\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket) \in T(E, \{co_1\}, \varepsilon, \varepsilon \cup \{rd_\dagger, wr_1\})$ under the assumption that
 619 $\{co_1\} \perp \varepsilon$, that is, when the environment does not read nor write X . This condition is
 620 clearly necessary, for e_1 and e_2 can be distinguished by an environment allowed to read
 621 or write X .

622 Let us now prove the claim when $\{co_1\} \perp \varepsilon$. The opponent picks a pilot trace in
 623 the semantics of e_1 , for example, $((h_1, k_1)(h_2, k_2), \dagger)$ where $h_1(X) = n$ and $k_1(X) =$
 624 $n + 1$ and $h_2(X) = n'$ and $k_2(X) = n' + 1$. The other possible traces are stuttering or
 625 mumbling variants of this one and do not present additional difficulties. The opponent
 626 also chooses a heap h'_1 such that $h_1 \stackrel{\dagger}{\sim} h'_1$, i.e., $h'_1(X) = n$. Now the proponent will
 627 choose to stutter for the time being and thus selects $k'_1 := h'_1$. Indeed, $[co_1](h_1, h'_1, k_1, k'_1)$
 628 holds, so this is legal. The opponent now presents h'_2 such that $[\varepsilon](k_1, k'_1, h_2, h'_2)$. By the
 629 assumption on ε we know that $n' = h_2(X) = k_1(X) = n + 1$ and also $h'_2(X) = k'_1(X) = n$.
 630 The proponent now answers with $k'_2 := h'_2[X \mapsto n + 2]$. It follows that $[co_1](h_2, h'_2, k_2, k'_2)$
 631 and also $[rd_1, wr_1](h_1, h'_1, k_2, k'_2)$. Finally, by stuttering $(h'_1, h'_1)(h'_2, h'_2[X \mapsto n + 2]) \in \llbracket e_2 \rrbracket$
 632 so that proponent wins the game.

633 **Example 7.6.** Consider the following programs e_1 and e_2 :

634 $(X := !X + 1 \parallel Y := !Y + 1)$ and $(X := !X + 1; Y := !Y + 1)$.

635 We show $(\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket) \in T(E, \{co_X, co_Y\}, \varepsilon, \varepsilon \cup \{rd_X, rd_Y, wr_X, wr_Y\})$, provided ε does
 636 not read nor modify X and Y . This equivalence could be deduced syntactically using
 637 our parallelization equation shown in Figure 5. For illustrative purpose, however, we
 638 describe its semantic proof using a game.

639 The opponent picks a pilot trace in $\llbracket e_1 \rrbracket$, for example, the trace $([n_1|n_2], [n_1|n_2 +$
 640 $1])([n_1|n_2 + 1], [n_1 + 1|n_2 + 1])(\dagger, \dagger)$, where $[n_X|n_Y]$ denotes a heap where X and Y store
 641 n_X and n_Y , respectively. Notice that in this trace, Y is incremented before X and since
 642 ε does not read nor modify X and Y , the environment move does not change the values
 643 in X nor Y . We are also given an initial heap h'_1 that agrees with the initial heap $[n_1|n_2]$
 644 on the reads of $\varepsilon \cup \{rd_X, rd_Y, wr_X, wr_Y\}$. Thus, h'_1 should be of the form $[n_1|n_2]$.

645 We now play the move $([n_1|n_2], [n_1 + 1|n_2])$. This is a valid move in the game as
 646 $[co_X, co_Y]([n_1|n_2], [n_1|n_2], [n_1|n_2 + 1], [n_1 + 1|n_2])$. The environment moves returning
 647 $[n_1 + 1|n_2]$ as it does not read nor modify X and Y . We can now match the trace above
 648 by playing $([n_1 + 1|n_2], [n_1 + 1|n_2 + 1])$ and returning (\dagger, \dagger) , winning the game.

649 The following is one of the main technical result of our paper and shows that the
 650 computation specifications $T(\dots)$ can indeed serve as the basis for a logical relation.

651 **Theorem 7.7.** *The following hold whenever well-formed.*

- 652 1. *If $(U, U') \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ then $(q_i(U), q_i(U')) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.*
- 653 2. *$T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ is a computation specification.*
- 654 3. *If $(U, U') \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ then $(U^\dagger, U'^\dagger) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.*
- 655 4. *If $(a, a') \in E$ then $(\text{rtn}(a), \text{rtn}(a'))$ is in $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.*
- 656 5. *Suppose that $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ is an effect specification where $\varepsilon_1 \cup \varepsilon_2 \subseteq \varepsilon_3$. Suppose*
 657 *that whenever $h \stackrel{\text{rds}(\varepsilon_1)}{\sim} h'$ and $c(h) = (h_1, a)$ then there exist (h'_1, a') such*
 658 *that $c'(h') = (h'_1, a')$ and $[\varepsilon_1](h, h', h_1, h'_1)$ and aEa' . We then have for any ε_2 ,*
 659 *$(\text{fromstate}(c), \text{fromstate}(c')) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.*

660 6. If $(f, f') \in E_1 \rightarrow T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ and $(U, U') \in T(E_1, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ then

$$(bnd(f, U), bnd(f', U')) \in T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$$

661 7. If $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon \cup \varepsilon_2, \varepsilon \cup \varepsilon_2 \cup \varepsilon')$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon \cup \varepsilon_1, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$
662 then

$$(U_1 \mid U'_1, U_2 \mid U'_2) \in T(E_1 \times E_2, \varepsilon_1 \cup \varepsilon_2, \varepsilon, \varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2))$$

663 8. $(U, U') \in T(E, \varepsilon_1, \emptyset, \varepsilon_3) \Rightarrow (at(U), at(U')) \in T(\varepsilon_3, \varepsilon_2, \varepsilon_2 \cup \varepsilon_3)$.

664 *Proof.* In each case, using Corollary 3.2 and Lemma 3.3 (for case 6), we can in fact
665 assume w.l.o.g. that the assumed pairs are in $T_0(\dots)$ rather than $T(\dots)$.

666 Ad 1. Let $(t, a) \in q_i(U)$, i.e. $a = p_i(a_0)$ where $(t, a_0) \in U$. By down-closure
667 ([Down]) we also have $(t, a) \in U$. We can now play the strategy guaranteed by the
668 assumption $(U, U') \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ which will yield (depending on the opponent's
669 moves) a trace t' and a value a' such that $(t', a') \in U'$ and $(p_i(a), a') \in E$. Now, since E
670 is a specification we get $(p_i(a), p_i(a')) \in E$ noting that p_i is idempotent. So, we modify
671 the strategy so as to return $p_i(a')$ rather than a' and thus obtain a winning strategy
672 asserting the desired conclusion.

673 Ad 2 This is an easy consequence from 1.

674 Ad 3 Pick $(U, U') \in T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$. Since $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ is closed under suprema
675 it suffices to show that $(q_j(U^\dagger), q_j(U'^\dagger)) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ for each j . Fix such j and
676 pick $(t, p_j(a)) \in q_j(U^\dagger)$, thus $(t, a) \in U^\dagger$.

677 By induction on the closure process we can assume w.l.o.g. that (t, a) arises from
678 $(t_1, a) \in U$ by a single mumbling or stuttering step or that $(t, a_1) \in U$ for some $a_1 \geq a$
679 or else that $(t, a_i) \in U$ where $\sup_i a_i = a$.

680 In the former two cases fix a strategy for the original element of U . We will use
681 this strategy to build a new one demonstrating that $(t, a) \in U'$, hence $(t, p_j(a)) \in q_j(U')$
682 as required.

683 If (t, a) arises by stuttering, so $t = u(h, h)v$ and $t_1 = uv$ we play the strategy until u
684 is worked off. If the opponent then produces a heap h' to match h we answer h' .

685 Now $[\varepsilon_1](h, h', h, h')$ is always true (Lemma 6.2) so this is a legal move. There-
686 after, we continue just as in the original strategy. In the special case where v is
687 empty, we must also show that $[\varepsilon_3](h_1, h'_1, h, h')$ knowing $[\varepsilon_3](h_1, h'_1, k_n, k'_n)$ where
688 $u = (h_1, k_1) \dots (h_n, k_n)$ and $u' = (h'_1, k'_1) \dots (h'_n, k'_n)$ is the matching trace. We have
689 $[\varepsilon_2](k_n, k'_n, h, h')$ for otherwise opponent's playing h' would have been illegal. Since,
690 by assumption $\varepsilon_2 \subseteq \varepsilon_3$, we can conclude $[\varepsilon_3](k_n, k'_n, h, h')$ and then $[\varepsilon_3](h_1, h'_1, h, h')$
691 by Lemma 6.2(3&1).

692 If (t, a) arises by mumbling then we must have $t = u(h_1, h_3)v$ and $t_1 = u(h_1, h_2)(h_2, h_3)v$.
693 We play until the strategy has produced a match h'_2 for h_2 . So far, the play has produced
694 a trace u' matching u , and a state h'_1 so that $[\varepsilon_1](h_1, h'_1, h_2, h'_2)$. Now, we can ask what
695 the original strategy would produce if we gave it (temporarily assuming opponent's
696 role) the state h'_2 as a match for h_2 . Note that this is legal because $[\varepsilon_2](h_2, h'_2, h_2, h'_2)$.
697 The strategy will then produce h'_3 such that $[\varepsilon_1](h_2, h'_2, h_3, h'_3)$ and our answer in the
698 play on the new trace against the challenge h'_1 will be this very h'_3 . Indeed, by compos-
699 ing tiles (Lemma 6.2) we have $[\varepsilon_1](h_1, h'_1, h_3, h'_3)$ as required. Thereafter, the play
700 continues according to the original strategy.

701 For down-closure, we play the strategy against (t, a_1) yielding a match $(t', a'_1) \in U'$
 702 where $a_1 E a'_1$. That same strategy also wins against (t, a) because $a E a'_1$ since E is a
 703 value specification.

704 For closure under [Sup], finally, pick i so that $a_i \geq p_j(a)$ recalling that $a = \sup_i a_i$.
 705 Since we have a winning strategy for (t, a_i) , we also have one (by down-closure which
 706 was already proved) for $(t, p_j(a))$ as required.

707 Ad 4. Suppose $a E a'$. By 3 which we have just proved we only need to match
 708 elements of the form $((h, h)a)$. The opponent plays h' where $h \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'$ and we answer
 709 with h' itself and a' . This is always a legal move (Lemma 6.2) and $a E a'$, so we win the
 710 game.

711 Ad 5. Again, we only need to match traces of the form $((h, h_1), a)$ where $c(h) =$
 712 (h_1, a) . In this case, suppose that the opponent plays h' where $h \stackrel{\varepsilon_3}{\approx} h'$. The assumption
 713 gives (h'_1, a') such that $c'(h') = (h'_1, a')$ and $[\varepsilon_1](h, h', h_1, h'_1)$ and $a E a'$. We thus play
 714 h'_1 and a' and indeed $[\varepsilon_{1/3}](h, h', h_1, h'_1)$ and $a E a'$ hold so this is a winning move.

715 Ad 6. Suppose $(f, f') \in E_1 \rightarrow T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ and $(U, U') \in T(E_1, \varepsilon_1, \varepsilon_2, \varepsilon_3)$. Sup-
 716 pose that $(uv, b) \in ap(f, U)$ where $(u, a) \in U$ and (v, b) in $f(a)$ (note that we can ignore
 717 the \dagger -closure). We need to produce a trace $(u'v', b') \in ap(f', U')$ such that $(u', a') \in U'$
 718 and (v', b') in $f'(a')$ and $b E_2 b'$. Assume that:

$$u = (h_1, k_1) \cdots (h_n, k_n) \text{ and } v = (h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m})$$

719 We are given a heap h'_1 , such that $h_1 \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'_1$. We can use the strategy S_1 from
 720 $(U, U') \in T(E_1, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ for (u, a) . We play according to S_1 to work off the u -part.
 721 This results in a matching trace $u' \in U'$:

$$u' = (h'_1, k'_1) \cdots (h'_n, k'_n)$$

722 where $[\varepsilon_3](h_1, h'_1, k_n, k'_n)$ and $(a, a') \in E_2$. We get $(f(a), f(a')) \in T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$. Now,
 723 we are given a heap h'_{n+1} that is an environment move forming the tile

$$[\varepsilon_2](k_n, k'_n, h_{n+1} h'_{n+1})$$

724 From the fact that $\varepsilon_2 \subseteq \varepsilon_3$ and Lemma 6.2(5) we can conclude $h_{n+1} \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'_{n+1}$.

725 Thus we can continue our play by using the strategy S_2 from

$$(f(a), f(a')) \in T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$$

726 which yields a continuation v' of our trace and a final answer b' . It is then clear that
 727 $(u'v', b') \in bnd(f', U')$ so this combination of strategies does indeed win.

728 Ad 7. Suppose that $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon \cup \varepsilon_2, \varepsilon \cup \varepsilon_2 \cup \varepsilon')$ and $(U_2, U'_2) \in$
 729 $T(E_2, \varepsilon_2, \varepsilon \cup \varepsilon_1, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$ and let $(t, (a, b)) \in U_1 \mid U_2$, thus $inter(t_1, t_2, t)$ (ignor-
 730 ing \dagger by item 3) where $(t_1, a) \in U_1$ and $(t_2, b) \in U_2$. Let S_1, S_2 be corresponding
 731 winning strategies. The idea is to use S_1 when we are in t_1 and to use S_2 when we are
 732 in t_2 . Supposing that t starts with a t_1 fragment we begin by playing according to S_1 .
 733 Let t be of the form:

$$t = (h_1, k_1) \cdots (h_n, k_n)(h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) \\ (h_{n+m+1}, k_{n+m+1}) \cdots (h_{n+m+k}, k_{n+m+k}) \cdots (h_p, k_p)$$

734 composed of pieces of the traces t_1 and t_2 . Assume w.l.o.g. that the first piece $(h_1, k_1) \cdots$
 735 $\cdots (h_n, k_n)$ is a part of t_1 . We are given a initial heap h'_1 such that $h \stackrel{\text{rds}(\varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2))}{\sim} h'_1$.
 736 Since $\text{rds}(\varepsilon_1 \sqcup \varepsilon_2) = \text{rds}(\varepsilon_1) \cup \text{rds}(\varepsilon_2)$, we can apply strategy S_1 to guide us through the
 737 first part of the game, obtaining:

$$(h'_1, k'_1) \cdots (h'_n, k'_n)$$

738 Moreover, we have an environment move which forms the tile $[\varepsilon](k_n, k'_n, h_{n+1}, h'_{n+1})$.
 739 Thus, we have the tile $[\varepsilon \cup \varepsilon_1](h_1, h'_1, h_{n+1}, h'_{n+1})$ which can be seen as an environment
 740 move for t_2 . Therefore, we can use strategy S_2 for the U' and continue the game,
 741 obtaining the trace piece:

$$(h'_{n+1}, k'_{n+1}) \cdots (h'_{n+m}, k'_{n+m})$$

742 Now, we can return to the S_1 game as the trace above is seen as an environment move
 743 for U . Alternating these strategies, we get a trace t which is in $(U \mid U')$. Let (a', b') be
 744 the final values reached at the end. It is clear that $[\varepsilon \cup \varepsilon' \cup \varepsilon_1 \cup \varepsilon_2](h, h', h_p, h'_p)$ and
 745 also aE_1a' and bE_2b' .

746 It remains to assert the stronger statement $[\varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2)](h, h', h_p, h'_p)$. To
 747 see this suppose that $wr_1 \in \varepsilon_1 \setminus \varepsilon_2 \setminus \varepsilon \setminus \varepsilon'$. Since the entire game can be viewed as an
 748 instance of the game U_1 vs U'_1 with interventions by U_2 vs. U'_2 regarded as environment
 749 interactions we have $[\varepsilon \cup \varepsilon_2 \cup \varepsilon'](h, h', h_p, h'_p)$ so that in fact $h \stackrel{!}{=} h_p$ and $h' \stackrel{!}{=} h'_p$. The
 750 case of co_1 and $\varepsilon_1, \varepsilon_2$ interchanged is analogous.

751 Ad 8. This is direct from the definition of atomic and appealing on the fact that
 752 $(U, U') \in T(E, \varepsilon_1, \emptyset, \varepsilon_3)$. \square

753 We assign a value specification $\llbracket \tau \rrbracket$ to each refined type by

- $\llbracket \text{int} \rrbracket = \{(v, v') \mid v = v' \in \mathbb{Z}\}$ • $\llbracket \tau_1 \times \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket$
- $\llbracket \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \rightarrow T(\llbracket \tau_2 \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)$

754 We omit the obvious definition of the other basic types and assume value specifications
 755 for user-specified types as given.

756 **Assumption 1.** *We henceforth adopt the following soundness assumption which must*
 757 *be established concretely for every concrete instance of our framework.*

- 758 • *The initial heap satisfies the current world: $h_{\text{init}} \models w$.*
- 759 • *Each axiom is type sound: whenever (v, v', τ) is an axiom then $(v, v) \in \llbracket \tau \rrbracket$ and*
 760 *$(v', v') \in \llbracket \tau \rrbracket$.*
- 761 • *Each axiom is inequationally sound: whenever (v, v', τ) is an axiom then $(v, v') \in$*
 762 *$\llbracket \tau \rrbracket$.*

763 **Theorem 7.8.** *Suppose that $\Gamma \vdash v : \tau$ and $\Gamma \vdash e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$. Then $(\eta, \eta') \in$*
 764 *$\llbracket \Gamma \rrbracket$ (interpreting a context as a cartesian product) implies $(\llbracket v \rrbracket \eta, \llbracket v \rrbracket \eta') \in \llbracket \tau \rrbracket$ and*
 765 *$(\llbracket e \rrbracket \eta, \llbracket e \rrbracket \eta') \in T(\llbracket \tau \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.*

766 *Proof.* By induction on derivations. Most cases are already subsumed by Theorem 7.7.
 767 The typing rules regarding functions and recursion follow from the definitions and from
 768 the fact that all specifications are admissible. \square

769 **8. Typed observational approximation**

770 **Definition 8.1** (Observational approximation). *Let v, v' be value expressions where*
 771 *$\vdash v : \tau$ and $\vdash v' : \tau$. We say that v observationally approximates v' at type τ if for all f*
 772 *such that $\vdash f : \tau \xrightarrow[\varepsilon]{\varepsilon_1 | \varepsilon_3} \mathbf{int}$ (“observations”) it is the case that if $((h_{init}, k), n) \in \llbracket f v \rrbracket$*
 773 *for $n \in \mathbb{Z}$ and starting from h_{init} then $((h_{init}, k'), n) \in \llbracket f v' \rrbracket$ for some k' . We write*
 774 *$\vdash v \leq_{obs} v'$ in this case. We say that v and v' are observationally equivalent at type τ ,*
 775 *written $\vdash v =_{obs} v'$ if both $\vdash v \leq_{obs} v' : \tau$ and $\vdash v' \leq_{obs} v : \tau$.*

776 This means that for every test harness f we build around v and v' , no matter how
 777 complicated it is and whatever environments it sets up to run concurrently with v and
 778 v' it is the case that each terminating computation of v (in the environment installed by
 779 f) can be matched by a terminating computation with the same result by v' in the same
 780 environment. It is important, however, that the environment be well typed, thus will
 781 respect the contracts set up by the type τ . E.g. if τ is a functional type expecting, say,
 782 a pure function as argument then, by the typing restriction, the environment f cannot
 783 suddenly feed v and v' a side-effecting function as input.

784 We remark that observational approximation extends canonically to open terms by
 785 lambda abstracting free variables (and adding a dummy abstraction in the case of closed
 786 terms) [6].

787 As usual, the logical relation is sound with respect to typed observational approx-
 788 imation and thus can be used to deduce nontrivial observational approximation rela-
 789 tions. We state and prove the precise formulation of this result.

790 **Theorem 8.2.** *Let v, v' be closed values and suppose that $(\llbracket v \rrbracket, \llbracket v' \rrbracket) \in \llbracket \tau \rrbracket$. Then*
 791 *$\vdash v \leq_{obs} v' : \tau$.*

792 *Proof.* If $\vdash f : \tau \xrightarrow[\varepsilon_2]{\varepsilon_1 | \varepsilon_3} \mathbf{int}$ then by Thm 7.8 we have $(\llbracket f \rrbracket, \llbracket f \rrbracket) \in \llbracket \tau \xrightarrow[\varepsilon_2]{\varepsilon_1 | \varepsilon_3} \mathbf{int} \rrbracket$, so
 793 $(\llbracket f v \rrbracket, \llbracket f v' \rrbracket) \in T(\llbracket \mathbf{int} \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)^+$.

794 Let $((h_{init}, k), n) \in \llbracket f v \rrbracket$. We have $h_{init} \models \mathbf{w}$ and thus in particular $h_{init} \stackrel{\text{rds}(\varepsilon_3) \cup \text{rds}(\varepsilon_1)}{\sim} h_{init}$
 795 h_{init} . There must therefore exist a matching heap k' and a value n' such that

$$((h_{init}, k'), n') \in \llbracket f v' \rrbracket \text{ and } n = n' \in \mathbb{Z}$$

796 □

797 This means that the examples from earlier on give rise to valid transformations in
 798 the sense of observational approximation. For instance, for e_1 and e_2 from Example 7.5
 799 we find that $\lambda \dots e_1 =_{obs} \lambda \dots e_2$ at type $\mathbf{unit} \xrightarrow[\varepsilon]{\{co_1\} | \varepsilon \cup \{rd_1, wr_1\}} \mathbf{unit}$ whenever X does not
 800 appear in ε .

801 **9. Effect-dependent transformations**

802 We will now establish the semantic soundness of the inequational theory of effect-
 803 dependent program transformations given in Figure 5. It includes concurrent versions

804 of the effect-dependent equations from [9, 29], but the side conditions on the environ-
805 mental interaction are by no means obvious. We also note that some equations now
806 only hold in one direction thus become inequations. This is in particular the case for
807 duplicated computations. Suppose that $?$ is a computation that nondeterministically
808 chooses a boolean value and let $e := \text{let } x = ? \text{ in } (x, x)$. Then, even though $?$ does
809 not read nor write any location we only have $e \leq (?, ?)$, but not $(?, ?) \leq e$ for $(?, ?)$
810 admits the result $(\text{true}, \text{false})$ but e does not. Furthermore, due to presence of non-
811 termination the equations for dead code elimination and pure lambda hoist also hold
812 in one direction only. It might be possible to restore both directions of said equations
813 by introducing special effects for nondeterminism and nontermination; we have not ex-
814 plored this avenue. We concentrate on the individual effect-dependent transformations
815 before summarising the foregoing results in the general soundness Theorem 9.2.

816 In many of the equations, co-effects play an important role. For example, in the
817 commuting and parallelization equations, the internal effects ε_1 and ε_2 in the premises
818 are replaced by ε_1^C and ε_2^C in the internal effects of the conclusion. This makes sense
819 intuitively because the computations are run in a different order, so for the internal
820 moves, the locations in ε_1 and ε_2 can be modified in any way (see Example 7.6). How-
821 ever, in the global effect, we can still guarantee the effects ε'_1 and ε'_2 because of the
822 \perp -conditions. This intuition appears directly in the soundness proofs.

823 The following thus constitutes the second main technical result of our paper.

824 **Theorem 9.1.** *The following hold whenever well-formed.*

- 825 • **Commuting** If $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon_1^C, \varepsilon^C \cup \varepsilon'_1)$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon_2^C, \varepsilon^C \cup$
826 $\varepsilon'_2)$ and $\varepsilon_1 \perp \varepsilon$ and $\varepsilon_2 \perp \varepsilon$ and $\varepsilon'_1 \perp \varepsilon'_2$ then

$$\begin{aligned} & \{(t_1 t_2, (v_1, v_2)) \mid (t_1, v_1) \in U_1, (t_2, v_2) \in U_2\}^\dagger, \\ & \{(t'_1 t'_2, (v'_1, v'_2)) \mid (t'_1, v'_1) \in U'_1, (t'_2, v'_2) \in U'_2\}^\dagger \\ & \in T(E_1 \times E_2, (\varepsilon_1 \cup \varepsilon_2)^C, \varepsilon, \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2) \end{aligned}$$

- 827 • **Duplicated** If $(U, U') \in T(E, \varepsilon_1, \varepsilon_2^C, \varepsilon_2^C \cup \varepsilon')$ and $\text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon') = \emptyset$ and $\varepsilon_2 \perp \varepsilon_1$,
828 then

$$\begin{aligned} & \{(t, (v, v)) \mid (t, v) \in U\}^\dagger, \{(t'_1 t'_2, (v'_1, v'_2)) \mid (t'_1, v'_1) \in U', \\ & (t'_2, v'_2) \in U'\}^\dagger \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_2 \cup \varepsilon') \end{aligned}$$

- 829 • **Pure Let** Let $(U, U') \in T(E, \varepsilon_1, \varepsilon_2^C, \varepsilon_2^C)$, such that $\varepsilon_1 \perp \varepsilon_2$. If $((q_1, k_1) \dots (q_n, k_n), v) \in$
830 U for some arbitrary trace $t = (q_1, k_1) \dots (q_n, k_n)$ (with $q_1 \models \mathbf{w}$) and value v , then
831 $(\text{rtn}(v), U') \in T(E, \varepsilon_1^C, \varepsilon_2, \varepsilon_2)$;

- 832 • **Dead** Suppose that $(U, U') \in T(\text{unit}, \varepsilon_1, \varepsilon_2, \varepsilon_2 \cup \varepsilon'_1)$, where $\text{wrs}(\varepsilon'_1) = \emptyset$ and
833 $\varepsilon_1 \perp \varepsilon_2$. Then $(U, \text{rtn}(())) \in T(\text{unit}, \varepsilon_1^C, \varepsilon_2, \varepsilon_2 \cup \varepsilon'_1)$.

- 834 • **Parallelization** If $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon_1^C \cup \varepsilon_2^C, \varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1)$ and $(U_2, U'_2) \in$
835 $T(E_2, \varepsilon_2, \varepsilon_2^C \cup \varepsilon_1^C, \varepsilon^C \cup \varepsilon_1^C \cup \varepsilon'_2)$ and $\varepsilon_1 \perp \varepsilon_2$ and $\varepsilon_1 \perp \varepsilon$ and $\varepsilon_2 \perp \varepsilon$, then

$$\begin{aligned} & (U_1 \parallel U_2, \{(t'_1 t'_2, (v'_1, v'_2)) \mid (t'_1, v'_1) \in U'_1, (t'_2, v'_2) \in U'_2\}^\dagger) \in \\ & T(E_1 \times E_2, \varepsilon_1^C \cup \varepsilon_2^C, \varepsilon, \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2) \end{aligned}$$

836 *Proof. Commuting.* By Theorem 7.7(3) we can assume our pilot trace t to be of the
837 form:

$$(h_1, k_1)(h_2, k_2) \cdots (h_n, k_n) (h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) (a, b)$$

838 where

$$\begin{aligned} t_1 &= (h_1, k_1)(h_2, k_2) \cdots (h_n, k_n) v_1 \in U_1 \\ t_2 &= (h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) v_2 \in U_2 \end{aligned}$$

839 We make similar use of Theorem 7.7(3) in the subsequent cases without explicit men-
840 tion.

841 We are also given a heap h'_1 such that

$$h_1 \stackrel{\text{rds}(\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2)}{\sim} h'_1$$

842 Because $\varepsilon'_1 \perp \varepsilon'_2$, h_1 and h_{n+1} agree on the reads of ε'_2 . Thus we can start a game U_2
843 vs. U'_2 using h'_1 and t_2 . We forward all environment's moves from the main game to
844 the side game and use the responses from the side game to answer in the main game.
845 Suppose that the side game leads to the valid U_2 -trace

$$(h'_1, k'_1)(h'_2, k'_2) \cdots (h'_m, k'_m) v'_2$$

846 where $v_2 E_2 v'_2$ and (1) $[\varepsilon^C \cup \varepsilon'_2](h_{n+1}, h'_1, k_{n+m}, k'_m)$. Notice that in the global game these
847 are legal responses as $[\varepsilon^C_i \cup \varepsilon^C_2](h_i, h'_i, k_i, k'_i)$ for $1 \leq i \leq m$.

848 We now have an environment move $[\varepsilon](k_m, k'_m, h_{m+1}, h'_{m+1})$. Since $\varepsilon'_1 \perp \varepsilon$ and $\varepsilon'_2 \perp$
849 ε'_1 , the heaps h'_1 and h'_{m+1} agree in the reads of ε'_1 . Therefore, we can run a game U_1
850 vs. U'_1 using h'_{m+1} and t_1 , obtaining the trace:

$$(h'_{m+1}, k'_{m+1})(h'_{m+2}, k'_{m+2}) \cdots (h'_{m+n}, k'_{m+n}) v'_1$$

851 where $v_1 E_1 v'_1$ and (2) $[\varepsilon^C \cup \varepsilon'_1](h_1, h'_{m+1}, k_n, k'_{m+n})$. The reasoning is similar to the use
852 of the previous game.

853 Thus we have that $(v_1, v_2)(E_1 \times E_2)(v'_1, v'_2)$.

854 Now, we need to conclude that $[\varepsilon^C \cup \varepsilon'_1 \cup \varepsilon'_2](h_1, h'_1, k_{n+m}, k'_{m+n})$. This follows from
855 the fact that $\varepsilon'_1 \perp \varepsilon'_2$ and (1) and (2). In particular, from (1) and $\varepsilon'_1 \perp \varepsilon'_2$, we get that
856 k_{m+n} and k'_{m+n} agree on the locations in ε'_2 , while from (2), we get that k_{m+n} and k'_{m+n}
857 agree on the locations in ε'_1 . This finishes the proof.

858 **Duplicated.** Assume given a trace in U :

$$t = (h_1, k_1) \cdots (h_n, k_n) v$$

859 and a heap h'_1 such that $h_1 \stackrel{\text{rds}(\varepsilon_2 \cup \varepsilon')}{\sim} h'_1$. Recall that $\text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon') = \emptyset$ and moreover,
860 since $\varepsilon^C_2 \cup \varepsilon'$ is well formed, we also have $\text{rds}(\varepsilon') \cap (\text{wrs}(\varepsilon_2) \cup \text{cos}(\varepsilon_2)) = \emptyset$. Thus h_1
861 and k_n agree on the reads of $\varepsilon' \cup \varepsilon^C_2$, i.e., the reads of ε' .

862 We start by simply stuttering:

$$t' = (h'_1, h'_1)(h'_2, h'_2) \cdots (h'_n, ??).$$

863 leaving the final heap ?? yet to be determined. So far, this is a legal play in the
864 main game because for $1 \leq i \leq n-1$, we have $[\varepsilon^C_i](h_i, h'_i, k_i, h'_i)$ and a chaotic ef-
865 fect on a location allows any changes to that location. Moreover, we may assume

866 $[\varepsilon_2](k_i, h_{i+1}, h'_i, h'_{i+1})$ for otherwise we would have won immediately. As a result, since
 867 $\varepsilon_1 \perp \varepsilon_2$, we inductively get $h_i \stackrel{\text{rds}(\varepsilon_2)}{\sim} h'_i$ and, of course, $h_i \stackrel{\text{rds}(\varepsilon_2)}{\sim} k_i$.
 868 We will now play two side-games U vs. U' with pilot trace t so as to construct the
 869 missing heap “??”. We first run a game starting at h'_n , where the environment moves are
 870 simply stutter moves. Recall that $h_1 \stackrel{\text{rds}(\varepsilon' \cup \text{cos}(\varepsilon_2^C))}{\sim} h'_n$ has already been asserted above.
 871 We thus obtain the following trace $t_1 \in U'$

$$t_1 = (h'_n, q_1)(q_1, q_2) \cdots (q_{n-1}, q_n) v'_1$$

872 where vEv'_1 and $[\varepsilon_2^C \cup \varepsilon'](h_1, h'_n, k_n, q_n)$. Notice that using stuttering environment moves
 873 is valid as $[\varepsilon_2^C](k_i, q_i, h_{i+1}, q_i)$ for $1 \leq i \leq n-1$.

874 Since h_1 and k_n agree on the reads of ε' and q_n and k_n agree on $\text{rds}(\varepsilon')$ from
 875 $[\varepsilon_2^C \cup \varepsilon'](h_1, h'_n, k_n, q_n)$, we can run the game U vs. U' again on q_n and t with stut-
 876 ter environment moves:

$$(q_n, q_{n+1})(q_{n+1}, q_{n+2}) \cdots (q_{n+n-1}, q_{n+n}) v'_2$$

877 where vEv'_2 and $[\varepsilon_2^C \cup \varepsilon'](h_1, q_n, k_n, q_{n+n})$. Thus, $(v, v)(E \times E)(v'_1, v'_2)$. This trace is
 878 again valid for the same reasons above, namely ε_1^C allows any internal moves, and
 879 since $\varepsilon_1 \perp \varepsilon_2$, the environment moves are also legal.

880 We now put $?? := q_{n+n}$ which leads to a valid trace due to repeated mumbling.
 881 Finally, we shall show that $[\varepsilon_2 \cup \varepsilon'](h_1, h'_1, k_n, q_{n+n})$ that is k_n and q_{n+n} agree on the
 882 reads of ε_2 and of ε' :

- 883 • They agree on the reads of ε' because $[\varepsilon_2^C \cup \varepsilon'](h_1, q_n, k_n, q_{n+n})$ obtained from
 884 the game above;
- 885 • They agree on the reads of ε_2 because $\varepsilon_1 \perp \varepsilon_2$. The internal moves did not affect
 886 the locations read by ε_2 .

887 **Duplicated for result value unit:** We can show that equality holds and not just \leq
 888 when the result type is `unit`. The reverse direction is proved as follows: For a given
 889 pilot trace t of U , where e is executed twice, we can construct a trace t' in U' by first
 890 stuttering and then mimicking the second execution of e . Since the resulting type is
 891 `unit`, there values obtained in t are necessarily $()$ which is also necessarily the same
 892 value obtained in the trace t' .

893 **Pure.** We start with a trace from $\text{rtn}(v)$, for example $(h_1, h_1), v$ and an arbitrary
 894 heap h'_1 . We now consider the game involving U vs. U' on t, v and h'_1 :

$$\begin{aligned} t &= (q_1, k_1)(q_2, k_2) \cdots (q_n, k_n), v \\ t' &= (h'_1, k'_1)(k'_1, k'_2) \cdots (k'_{n-1}, k'_n), v' \end{aligned}$$

895 We have that vEv' and $[\varepsilon_3](q_1, h'_1, k_n, k'_n)$. By mumbling, $(h'_1, k'_n) \in U'$. We can reply
 896 with k'_n in the main game.

897 **Dead.** Assume given a trace of the form:

$$(h_1, k_1) \cdots (h_n, k_n) v$$

898 and h'_1 such that $h_1 \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'_1$. We now initiate a side game U vs. U' on this trace and
 899 respond in the main game by stuttering. Thus, we obtain traces $(h'_1, h'_1) \cdots (h'_n, h'_n)$ ()
 900 in the main game and $(h'_1, k'_1) \cdots (h'_n, k'_n)$ v' in the side game.

901 The main trace is in $\text{rtn}()$. The side game tells us that $v = ()$ and that $h_i \xrightarrow{\varepsilon_1} k_i$ and
 902 therefore $[\varepsilon_1^C](h_i, h'_i, k_i, h'_i)$. It remains to show that $[\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2](h_1, h'_1, k_n, k'_n)$. This
 903 follows from the fact that ε_1 has only reads as h_i and k_i agree on all locations.

904 *Parallelization.* We start with a trace in $U_1 \parallel U_2$. Assume that the trace is of the fol-
 905 lowing form:

$$t_{1,1}t_{2,1}t_{1,2}t_{2,2} \dots t_{1,n}t_{2,n} (v_1, v_2)$$

906 where each $t_{i,j}$ is a possibly empty sequence of moves of the form $(h_{i,j}^1, k_{i,j}^1) \cdots (h_{i,j}^{m_{i,j}}, k_{i,j}^{m_{i,j}})$
 907 and

$$\begin{aligned} t_1 &= t_{1,1} \cdots t_{1,n} \quad v_1 \in U_1 \\ t_2 &= t_{2,1} \cdots t_{2,n} \quad v_2 \in U_2 \end{aligned}$$

908 are traces from U_1 and U_2 , respectively. We are also given a heap h'_1 such that $h_{1,1}^1 \stackrel{\text{rds}(\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2)}{\sim} h'_1$
 909 h'_1 . We also have $h_{1,1}^1 \stackrel{\text{rds}(\varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1)}{\sim} h'_1$. We run a side game U_1 vs. U'_1 using h'_1 and t_1 ,
 910 yielding:

$$t'_{1,1} \cdots t'_{1,n} \quad v'_1$$

911 Assume that (h'_1, k'_1) and (h'_o, k'_o) are, respectively, the first and last moves of this trace.
 912 We have $v_1 E_1 v'_1$ and (1) $[\varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1](h_{1,1}^1, h'_1, k_{1,n}^m, k'_o)$. Notice that these are legal
 913 moves in the global game as we have $[\varepsilon_1^C \cup \varepsilon_2^C]$ tiles for the player moves and $[\varepsilon]$ times
 914 for the environment moves.

915 Now, assume there is an environment move (k_o, h'_{o+1}) . Since $\varepsilon_1 \perp \varepsilon_2$ and $\varepsilon \perp \varepsilon_2$,
 916 the heaps $h_{1,1}^1$ and $h_{2,1}^1$ agree on the reads of ε_2 and h'_1 and h'_{o+1} also agree on the reads
 917 of ε_2 . (Notice as well that $\text{wrs}(\varepsilon_1) \cap \text{rds}(\varepsilon_2) = \emptyset$ as $\varepsilon^C \cup \varepsilon_1^C \cup \varepsilon_2$ is a valid effect.)
 918 Therefore, we can invoke an U_2 game using h'_{o+1} and t_2 , obtaining the trace:

$$t'_{2,1} \cdots t'_{2,n} \quad v'_2$$

919 Assume that (h'_{o+1}, k'_{o+1}) and (h'_{o+p}, k'_{o+p}) are, respectively, the first and last moves of
 920 this trace. We have $v_2 E_2 v'_2$ and (2) $[\varepsilon^C \cup \varepsilon_1^C \cup \varepsilon'_2](h_{2,1}^1, h'_{o+1}, k_{2,n}^m, k'_{o+p})$. For the same
 921 reasons as above, these are legal moves in the global game.

922 Therefore $(v_1, v_2)(E_1 \times E_2)(v'_1, v'_2)$.

923 We need now to prove that $[\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2](h_{1,1}^1, h'_1, k_{2,n}^m, k_{o+p})$. From (1) and $\varepsilon_1 \perp \varepsilon_2$
 924 and $\varepsilon \perp \varepsilon_1$, we have that $k_{2,n}^m$ and k_{o+p} agree on the locations of ε_1 . Similarly, $k_{2,n}^m$ and
 925 k_{o+p} agree on the locations of ε_2 . Since there are only ε tiles and $\varepsilon \perp \varepsilon_1$ and $\varepsilon \perp \varepsilon_2$,
 926 $k_{2,n}^m$ and k_{o+p} agree on the locations of ε . This finishes the proof.
 927 □

928 **Theorem 9.2.** Suppose that $\Gamma \vdash v \leq v' : \tau$ and $\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$
 929 and assume that for each axiom (v, v', τ) it holds that $(v, v') \in \llbracket \tau \rrbracket^+$. Then $(\eta, \eta') \in$
 930 $\llbracket \Gamma \rrbracket^+$ (interpreting a context as a cartesian product) implies $(\llbracket v \rrbracket \eta, \llbracket v' \rrbracket \eta') \in \llbracket \tau \rrbracket^+$ and
 931 $(\llbracket e \rrbracket \eta, \llbracket e' \rrbracket \eta') \in T(\llbracket \tau \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)^+$.

932 *Sketch.* In essence the proof is by induction on derivations of inequalities. However,
 933 we need to slightly strengthen the induction hypothesis as follows:

934 Define

$$\begin{aligned} \llbracket \Gamma \vdash \tau \rrbracket &= \{(f, f') \mid \forall (\eta, \eta') \in \llbracket \Gamma \rrbracket. (f(\eta), f'(\eta')) \in \llbracket \tau \rrbracket\} \\ \llbracket \Gamma \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket &= \{(f, f') \mid \forall (\eta, \eta') \in \llbracket \Gamma \rrbracket. \\ &\quad (f(\eta), f'(\eta')) \in T(\llbracket \tau \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)\} \end{aligned}$$

935 We now show by induction on derivations that $\Gamma \vdash v \leq v' : \tau$ implies $(\llbracket v \rrbracket, \llbracket v' \rrbracket) \in$
 936 $\llbracket \Gamma \vdash \tau \rrbracket^+$ and that $\Gamma \vdash e \leq e' : \tau \& \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ implies $(\llbracket e \rrbracket, \llbracket e' \rrbracket) \in \llbracket \Gamma \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^+$.

937 The various cases now follow from earlier results in a straightforward manner.
 938 Namely, we use Theorem 7.7 for the congruence rules and Theorem 9.1 for the effect-
 939 dependent transformations.

940 As a representative case we show the case where $e \equiv \text{let } x = e_1 \text{ in } e_2$ and
 941 $e' \equiv \text{let } x = e'_1 \text{ in } e'_2$. Inductively, we know $(\llbracket e_1 \rrbracket, \llbracket e'_1 \rrbracket) \in \llbracket \Gamma \vdash \tau_1 \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^{n_1}$ and
 942 $(\llbracket e_1 \rrbracket, \llbracket e'_1 \rrbracket) \in \llbracket \Gamma, x:\tau_1 \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^{n_2}$ for some $n_1, n_2 > 0$. By Theorem 7.8, we
 943 also have $(\llbracket e_1 \rrbracket, \llbracket e_1 \rrbracket) \in \llbracket \Gamma \vdash \tau_1 \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket$ and analogous statements for e'_1, e_2, e'_2 .
 944 We can, therefore, assume, w.l.o.g. that $n_1 = n_2$ and then use Theorem 7.7 (6) repeat-
 945 edly (n_1 times) so as to conclude $(\llbracket e \rrbracket, \llbracket e' \rrbracket) \in \llbracket \Gamma \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^{n_1}$.

946 The rules for dead code and pure lambda hoist rely on the cases “Dead” and “Pure”
 947 of Thm 9.1 in a slightly indirect way. We sketch the argument for pure lambda hoist.
 948 The pilot trace begins with a trace belonging to e_1 and yielding a value v for x . We can
 949 then invoke case “Pure” on subsequent occurrences of e_1 in the right hand side. \square

950 **Theorem 9.3.** Suppose that $\vdash v : \tau$ and $\vdash v' : \tau$ and that $(\llbracket v \rrbracket, \llbracket v' \rrbracket) \in \llbracket \tau \rrbracket^+$ where $(-)^+$
 951 denotes transitive closure. Then $\vdash v \leq_{\text{obs}} v' : \tau$.

952 *Proof.* If $\vdash f : \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \text{int}$ then by Thm 7.8 we have $(\llbracket f \rrbracket, \llbracket f \rrbracket) \in \llbracket \tau \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \text{int} \rrbracket$, so
 953 $(\llbracket f v \rrbracket, \llbracket f v' \rrbracket) \in T(\llbracket \text{int} \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)^+$.

954 Let $(h_{\text{init}}, k, v) \in \llbracket f v \rrbracket$. We have $h_{\text{init}} \models w$ and thus in particular $h_{\text{init}} \stackrel{\text{rds}(\varepsilon_3) \cup \text{rds}(\varepsilon_1)}{\sim} h_{\text{init}}$
 955 h_{init} . There must therefore exist a matching heap k' and a value v' such that

$$((h_{\text{init}}, k'), v') \in \llbracket f v' \rrbracket \text{ and } v = v' \in \mathbb{Z}$$

956 \square

957 We now return to the examples that we discussed in Section 1 and demonstrate
 958 how to prove using our denotational semantics the properties that have been discussed
 959 informally.

960 *Overlapping References.* With this example, we illustrate the parallelization rule. In
 961 particular, the functions declared in Section 1 have the following type, where ε does
 962 not read nor write X :

$$\begin{aligned} \text{readFst} : \text{unit} &\xrightarrow[\varepsilon^C, \text{co_snd}(X)]{\emptyset \mid \varepsilon^C, \text{co_snd}(X), \text{rd}_{\text{fst}}(X)} \text{int} \\ \text{writeFst} : \text{int} &\xrightarrow[\varepsilon^C, \text{co_snd}(X)]{\text{wr}_{\text{fst}}(X) \mid \varepsilon^C, \text{co_snd}(X), \text{wr}_{\text{fst}}(X)} \text{unit} \end{aligned}$$

964 The obvious and analogous typings for `readSnd` and `writeSnd` are elided. We
 965 justify this typing semantically as described in Theorem 7.7. To illustrate how this is
 966 done, consider the function (`writeSnd 17`). We show how the game is played against
 967 itself using the typing shown above. We start with a “pilot trace”, say:

$$([2|3], [2|3]), ([2|17], [2|17]), ((),$$

968 where $[x|y]$ denotes a store with $X = p(x, y)$ and other components left out for simplic-
 969 ity. The first step corresponds to our reading of X and in the second step – since there
 970 was no environment intervention – we write 17 into the first component.

971 We now start to play: Say that we start at the heap $[13|12]$. We answer $[13|12]$. If
 972 the environment does not change X , then we write 17 to its first component resulting in
 973 the following trace, which is possible for `writeFst(17)`.

$$974 ([13|12], [13|12]), ([13|12], [17|12]), ((),$$

975 If, however, the environment plays $[18|21]$ (a modification of both components of X
 976 has occurred), then we answer $[17|21]$. Again,

$$977 ([13|12], [13|12]), ([18|21], [17|21]), ((),$$

978 is a possible trace for `writeFst(17)`. It is easy to check that there is a strategy that
 979 justifies the typing given above.

980 Now, consider a program, e_1 , that only calls `readFst`, `writeFst`, and another program,
 981 e_2 , that only calls `readSnd`, `writeSnd`. Since the former functions have disjoint effects
 982 to the latter ones, e_1 and e_2 will have effect specifications, respectively, of the form
 983 $(\varepsilon_1, \varepsilon^C \cup \varepsilon_2^C, \varepsilon^C \cup \varepsilon_2^C \cup \varepsilon_1)$ and $(\varepsilon_2, \varepsilon^C \cup \varepsilon_1^C, \varepsilon^C \cup \varepsilon_1^C \cup \varepsilon_2)$, where $\varepsilon_1 \cap \varepsilon_2 = \varepsilon_1 \cap \varepsilon = \varepsilon_2 \cap \varepsilon = \emptyset$.
 984 Thus we can use the parallelization rule shown in Figure 5 to conclude that the behavior
 985 of $e_1 \parallel e_2$ is the same as executing these programs sequentially, although they read and
 986 write to the same concrete location.

987 *Loop Parallelization.* We show that the function `map` is equivalent to `map2Par`. It is
 988 easy to see that the function `map` is equivalent to the program `map2Seq`, which is the
 989 program obtained from `map2Par` by replacing the underlined parallel operator ‘ \parallel ’ in
 990 `map2Par` by a sequential operator ‘;’. The proof goes simply by unfolding `map`.

991 We then proceed by showing `map2Seq` and `map2Par` are equivalent using our
 992 equations and the abstract locations `listodd(X)` and `listeven(X)` defined above. The piece
 993 of code that applies f first, namely $e_1 = n.ele := f(n.ele)$, has global effects $\varepsilon'_1 =$
 994 $rd_{listodd(X)}, wr_{listodd(X)}$, while the second application, namely,

$$e_2 = n.next.ele := f(n.next.ele)$$

995 has effects $\varepsilon'_2 = rd_{listeven(X)}, wr_{listeven(X)}$. Notice that $\varepsilon'_1 \perp \varepsilon'_2$. Therefore, provided that
 996 the environment does not read nor modify the list, we can apply the parallelization
 997 equation to justify running e_1 and e_2 parallel is equivalent to running them in sequence.

998 *Michael-Scott Queue.* We now show that the `enqueue` and `dequeue` functions de-
 999 scribed in Section 1 for the Michael-Scott Queue have the same behavior as their atomic
 1000 versions. We only show the case for `dequeue`, as the case for `enqueue` is similar. More
 1001 precisely, we now justify the axiom

$$(\text{dequeue}, \text{atomic}(\text{dequeue}), \text{unit} \xrightarrow[\text{MSQ}]{\text{MSQ}|\text{MSQ}} \text{int})$$

1002 where $MSQ = \{rd_{msq(X)}, wr_{msq(X)}\}$. That is, they approximate each other at a type where
 1003 the environment is allowed to operate on the queue as well. We also note that the
 1004 converse of the axiom is obvious by stuttering and mumbling. After consuming a
 1005 dummy argument () let the resulting pilot trace be $(h_1, k_1) \dots (h_i, k_i) \dots (h_n, k_n)a$ and h'_1
 1006 be the start heap to match. We can now assume that the passages from k_i to h_{i+1} are
 1007 according to the protocol, i.e. $k_i \xrightarrow{msq(X)} h_{i+1}$. Namely, should this not be the case we
 1008 are free to make arbitrary moves and still win the game by default of the environment
 1009 player. Therefore, there must exist i such that in the move (h_i, k_i) the element a is
 1010 dequeued and $h_j = k_j$ holds for $j \neq i$. We can thus match this trace by a trace in the
 1011 semantics of `atomic(dequeue ())` by stuttering until i :

1012 $(h'_1, h'_1) \dots (h'_i, \dots$
 1013 where h_j and h'_j have the same content, but not necessarily the exact same layout.
 1014 Given the environment's allowed effects it is then clear that also h_i and h'_i have the
 1015 same content, but not necessarily the same as h_1 and h'_1 because in the meantime other
 1016 operations on the queue might have succeeded. We then dequeue the corresponding
 1017 element from h'_i leading to k'_i and continue by stuttering.

1018 $\dots, k'_i)(h'_{i+1}, h'_{i+1}) \dots (h'_n, h'_n)a'$
 1019 It is now clear that this is a matching trace and that $a = a'$ so we are done.

1020 Notice that the congruence rules now allow us to deduce the equivalence of $op_1 \parallel$
 1021 $\dots \parallel op_n$ and `atomic(op_1)` $\parallel \dots \parallel$ `atomic(op_n)` for op_i being enqueues or dequeues,
 1022 which effectively amounts to linearizability.

1023 10. Discussion

1024 We have shown how a simple effect system for stateful computation and its rela-
 1025 tional semantics, combined with the notion of abstract locations, scales to a concurrent
 1026 setting. The resulting type system provides a natural and useful degree of control over
 1027 the otherwise anarchic possibilities for interference in shared variable languages, as
 1028 demonstrated by the fact that we can delineate and prove the conditions for non-trivial
 1029 contextual equivalences, including fine-grained data structures.

1030 The primary goal of this line of work is not so much to find reasoning principles
 1031 that support the most subtle equivalence arguments for particular programs, but rather
 1032 to capture more generic properties of modules, expressed in terms of abstract locations
 1033 and relatively simple effect annotations, that can be exploited by clients (including op-
 1034 timizing compilers) in external reasoning and transformations. But there are of course,
 1035 particularly in view of the fact that we allow deeper reasoning to be used to establish
 1036 that expressions can be assigned particular effect-refined types, very close connections
 1037 with other work on richer program logics and models.

1038 Rely-guarantee reasoning is widely used in program logics for concurrency, in-
 1039 cluding relational ones [23], whilst our abstract locations are very like the *islands* of
 1040 Ahmed et al [4]. Recent work of Turon et al [31] on relational models for fine-grained
 1041 concurrency introduces richer abstractions, notably state transition systems expressing
 1042 inter-thread protocols that can involve ownership transfer. These certainly allow the
 1043 verification of more complex fine-grained algorithms than can be dealt with in our set-
 1044 ting, and it would be natural to try defining an effect semantics over such a model.

1045 Indeed, one might reasonably hope that effects could provide something of a ‘sim-
1046 plifying lens’, with refined types capturing things that would otherwise be extra model
1047 structure or more complex invariants, such that the combination does not lead to further
1048 complexity. The use of Brookes’s trace model (also used by, for example, Turon and
1049 Wand [32]) already seems to bring some simplification compared to transition systems
1050 or resumptions.

1051 Birkedal et al [12] have also studied relational semantics for effects in a concur-
1052 rent language. The language considered there has dynamic allocation via regions and
1053 higher-order store, neither of which we have here. On the other hand, their invariants
1054 are based on simply-typed concrete locations and thus do not allow to capture effects
1055 at the level of whole datastructures as abstract locations do. As a result, the examples
1056 in [12] are of a simpler nature than ours. Furthermore, we offer a subtler parallelization
1057 rule, distinguish transient and end-to-end effects, and validate other effect-dependent
1058 equivalences like commuting, lambda hoist, deadcode and duplication. Our use of
1059 denotational methods and in particular the extension of Brookes’ trace semantics to
1060 higher-order functions does result in a rather simpler and more intuitive definition of
1061 the logical relation by comparison with [12]. While some of the complications are due
1062 to the dynamic allocation and typed locations, others like the explicit step counting,
1063 the need for effect-instrumented operational semantics, and the separation of branches
1064 in the definition of safety are not. We thus see our work also as a proof-of-concept for
1065 denotational semantics in the realm of higher-order concurrent programming.

1066 The ‘RGSim’ relation proposed by Liang *et al.* for proving concurrent refinements
1067 under contextual assumptions also has many similarities with our logical relation [23,
1068 Def.4]. The focus of that work is on proving particular equivalences and refinements,
1069 whereas we encapsulate general patterns of behaviour in a refined type system and can
1070 show the soundness of generic program transformations relying only on effect types
1071 (which combine smoothly with hand proofs of particular equivalences).

1072 Since this work has been presented at PPDP 2016, Krogh-Jespersen et al [22] have
1073 proposed a system with similar goals as ours. It features higher-order store, i.e., the
1074 possibility of storing computations in the heap and not only flat values and pointer
1075 structures. In [10] we argued how our semantics-based approach can be extended to
1076 higher-order store as well, however, since the issue is mostly orthogonal we refrained
1077 from elaborating this path here in the context of concurrency. On the other hand, [22]
1078 has weaker rules than ours. Parallelization relies on essentially complete separation
1079 and it is even argued explicitly that parallelization comes down to “framing”. In our
1080 work, and in [23], closer interaction is possible provided one establishes appropriate
1081 invariants in the style of rely-guarantee. Also, presumably due to lack of space, the
1082 classical effect-dependent rules such as duplication are not treated in [22] and few
1083 examples are given. A more detailed comparison should thus await an extended journal
1084 version of [22]. From a methodological point of view, [22] is rather different from the
1085 work presented here. Namely, equivalences are justified by a translation into the *unary*
1086 program logic Iris [20]. This approach has become popular in the last couple of years.
1087 Essentially, the idea is to compare the behaviour of two programs, i.e., both sides of
1088 an equivalence, by proving a statement in Hoare logic about one of them. The Hoare
1089 logic must for this purpose be augmented with special assertions allowing one to speak
1090 about steps of the other program. The big advantage of this approach is that the difficult

1091 soundness proof needs to be carried out only once and for a unary Hoare logic which is
1092 easier. Moreover, the unary Hoare logic, Iris, has been formalised in Coq. A possible
1093 disadvantage is that the encoding via a unary Hoare logic might be complicated and
1094 unwieldy. It is, however, a very interesting and potentially promising proposal. It
1095 would be interesting to see whether it can be used to justify the exact equational theory
1096 given in this paper. This would allow one to compare the approaches in a more direct
1097 way.

1098 Besides that, there are many other directions for further work. Most importantly,
1099 we would like to add dynamic allocation of abstract locations following [6]. In addition
1100 to relieving us from having to set up all data structures in the initial heap this would, as
1101 we believe, also allow us to model and reason about lock-based protocols in an elegant
1102 way. Other possible extension include higher-order store as mentioned above and weak
1103 concurrency models. Somewhat further afield, it would be interesting to study ways of
1104 automatically inferring opportunities for applications of our equivalences to optimize
1105 programs and, relatedly, to use our theory to justify concrete compiler optimisations.

1106 References

- 1107 [1] M. Abadi, C. Flanagan, and S. N. Freund. Types for safe locking: Static race
1108 detection for java. *ACM Trans. Program. Lang. Syst.*, 28(2):207–255, 2006.
- 1109 [2] M. Abadi and L. Lamport. The existence of refinement mappings. *Theor. Comput.*
1110 *Sci.*, 82(2):253–284, 1991.
- 1111 [3] S. Abramsky and A. Jung. Domain theory, 1994. Online Lecture Notes, available
1112 from CiteSeerX.
- 1113 [4] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation inde-
1114 pendence. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on*
1115 *Principles of Programming Languages, POPL 2009, Savannah, GA, USA, Jan-*
1116 *uary 21-23, 2009*, pages 340–353, 2009.
- 1117 [5] T. Amtoft, F. Nielson, and H. R. Nielson. *Type and Effect Systems: Behaviours*
1118 *for Concurrency*. World Scientific, 1999.
- 1119 [6] N. Benton, M. Hofmann, and V. Nigam. Abstract effects and proof-relevant log-
1120 ical relations. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Prin-*
1121 *ciples of Programming Languages, POPL '14, San Diego, CA, USA, January*
1122 *20-21, 2014*, pages 619–632, 2014.
- 1123 [7] N. Benton, M. Hofmann, and V. Nigam. Effect-dependent transformations for
1124 concurrent programs. *CoRR*, abs/1510.02419, 2015.
- 1125 [8] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics
1126 for effect-based program transformations: higher-order store. In *Proceedings of*
1127 *the 11th International ACM SIGPLAN Conference on Principles and Practice of*
1128 *Declarative Programming, September 7-9, 2009, Coimbra, Portugal*, pages 301–
1129 312, 2009.

- 1130 [9] N. Benton, A. Kennedy, M. Hofmann, and L. Beringer. Reading, writing and re-
1131 lations. In *Programming Languages and Systems, 4th Asian Symposium, APLAS*
1132 *2006, Sydney, Australia, November 8-10, 2006, Proceedings*, pages 114–130,
1133 2006.
- 1134 [10] N. Benton, A. Kennedy, M. Hofmann, and V. Nigam. Counting successes: Effects
1135 and transformations for non-deterministic programs. In S. Lindley, C. McBride,
1136 P. W. Trinder, and D. Sannella, editors, *A List of Successes That Can Change the*
1137 *World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*,
1138 volume 9600 of *Lecture Notes in Computer Science*, pages 56–72. Springer, 2016.
- 1139 [11] N. Benton, A. Kennedy, and G. Russell. Compiling standard ML to java byte-
1140 codes. In *Proceedings of the third ACM SIGPLAN International Conference on*
1141 *Functional Programming (ICFP '98), Baltimore, Maryland, USA, September 27-*
1142 *29, 1998.*, pages 129–140, 1998.
- 1143 [12] L. Birkedal, F. Sieczkowski, and J. Thamsborg. A concurrent logical relation.
1144 In *Computer Science Logic (CSL'12) - 26th International Workshop/21st Annual*
1145 *Conference of the EACSL, CSL 2012, September 3-6, 2012, Fontainebleau,*
1146 *France*, pages 107–121, 2012.
- 1147 [13] L. Birkedal, M. Tofte, and M. Vejlstrup. From region inference to von Neumann
1148 machines via region representation inference. In *Proceedings of the 23rd ACM*
1149 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL*
1150 *'96)*, 1996.
- 1151 [14] N. Broberg and D. Sands. Flow locks: Towards a core calculus for dynamic flow
1152 policies. In *15th European Symposium on Programming (ESOP '06)*, volume
1153 3924 of *LNCS*. Springer, 2006.
- 1154 [15] S. D. Brookes. Full abstraction for a shared-variable parallel language. *Inf. Com-*
1155 *put.*, 127(2):145–163, 1996.
- 1156 [16] J. W. Coleman and C. B. Jones. A structural proof of the soundness of
1157 rely/guarantee rules. *J. Log. Comput.*, 17(4):807–841, 2007.
- 1158 [17] R. De Nicola and M. Hennessy. Testing equivalence for processes. In *Automata,*
1159 *Languages and Programming, 10th Colloquium, Barcelona, Spain, July 18-22,*
1160 *1983, Proceedings*, pages 548–560, 1983.
- 1161 [18] C. Flanagan and S. Qadeer. A type and effect system for atomicity. In *Proceed-*
1162 *ings of the ACM SIGPLAN Conference on Programming Language Design and*
1163 *Implementation (PLDI '03)*, 2003.
- 1164 [19] D. K. Gifford and J. M. Lucassen. Integrating functional and imperative program-
1165 ming. In *LISP and Functional Programming*, 1986.
- 1166 [20] R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and
1167 D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent rea-
1168 soning. In S. K. Rajamani and D. Walker, editors, *Proceedings of the 42nd Annual*

- 1169 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages,*
1170 *POPL 2015, Mumbai, India, January 15-17, 2015*, pages 637–650. ACM, 2015.
- 1171 [21] O. Kammar and G. D. Plotkin. Algebraic foundations for effect-dependent op-
1172 timisations. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on*
1173 *Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania,*
1174 *USA, January 22-28, 2012*, pages 349–360, 2012.
- 1175 [22] M. Krogh-Jespersen, K. Svendsen, and L. Birkedal. A relational model of types-
1176 and-effects in higher-order concurrent separation logic. In G. Castagna and A. D.
1177 Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Prin-*
1178 *ciples of Programming Languages, POPL 2017, Paris, France, January 18-20,*
1179 *2017*, pages 218–231. ACM, 2017.
- 1180 [23] H. Liang, X. Feng, and M. Fu. A rely-guarantee-based simulation for verifying
1181 concurrent program transformations. In J. Field and M. Hicks, editors, *Proceed-*
1182 *ings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Program-*
1183 *ming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28,*
1184 *2012*, pages 455–468. ACM, 2012.
- 1185 [24] N. A. Lynch and F. W. Vaandrager. Forward and backward simulations, ii:
1186 Timing-based systems. *Inf. Comput.*, pages 1–25, 1996.
- 1187 [25] M. M. Michael and M. L. Scott. Nonblocking algorithms and preemption-safe
1188 locking on multiprogrammed shared memory multiprocessors. *J. Parallel Distrib.*
1189 *Comput.*, 51(1):1–26, May 1998.
- 1190 [26] N. Benton and P. Buchlovsky. Semantics of an effect analysis for exceptions. In
1191 *3rd ACM Workshop on Types in Language Design and Implementation (TLDI*
1192 *'07)*, 2007.
- 1193 [27] F. Pessaux and X. Leroy. Type-based analysis of uncaught exceptions. In *Pro-*
1194 *ceedings of the 26 ACM Symposium on Principles of Programming Languages*
1195 *(POPL '99)*, 1999.
- 1196 [28] G. D. Plotkin. A powerdomain construction. *SIAM J. Comput.*, 5(3):452–487,
1197 1976.
- 1198 [29] J. Thamsborg and L. Birkedal. A Kripke logical relation for effect-based program
1199 transformations. In M. M. T. Chakravarty, Z. Hu, and O. Danvy, editors, *Proceed-*
1200 *ing of the 16th ACM SIGPLAN international conference on Functional Program-*
1201 *ming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, pages 445–456. ACM,
1202 2011.
- 1203 [30] J.-B. Tristan and X. Leroy. A simple, verified validator for software pipelining.
1204 In *POPL*, 2010.
- 1205 [31] A. J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical rela-
1206 tions for fine-grained concurrency. In *The 40th Annual ACM SIGPLAN-SIGACT*
1207 *Symposium on Principles of Programming Languages, POPL '13, Rome, Italy -*
1208 *January 23 - 25, 2013*, pages 343–356, 2013.

- 1209 [32] A. J. Turon and M. Wand. A separation logic for refining concurrent objects. In
1210 T. Ball and M. Sagiv, editors, *Proceedings of the 38th ACM SIGPLAN-SIGACT*
1211 *Symposium on Principles of Programming Languages, POPL 2011, Austin, TX,*
1212 *USA, January 26-28, 2011*, pages 247–258. ACM, 2011.