

Modelling Deterministic Concurrent I/O

Malcolm Dowse*

Trinity College Dublin, Ireland
dowsem@tcd.ie

Andrew Butterfield

Trinity College Dublin, Ireland
Andrew.Butterfield@cs.tcd.ie

Abstract

The problem of expressing I/O and side effects in functional languages is a well-established one. This paper addresses this problem from a general semantic viewpoint by giving a unified framework for describing shared state, I/O and deterministic concurrency. We develop a modified state transformer which lets us mathematically model the API, then investigate and machine verify some broad conditions under which confluence holds. This semantics is used as the basis for a small deterministic Haskell language extension called CURIO, which enforces determinism using runtime checks.

Our confluence condition is first shown to hold for a variety of small components, such as individual shared variables, 1-to-1 communication channels, and I-structures. We then show how models of substantial APIs (like a modification of Haskell’s file I/O API which permits inter-process communication) may be constructed from these smaller components using “combinators” in such a way that determinism is always preserved. We describe combinators for product, name-indexing and dynamic allocation, the last of which requires some small extensions to cater for process locality.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; D.3.3 [Programming Languages]: Language Constructs and Features

General Terms Design, Languages, Theory

Keywords Monads, effects, I/O, concurrency, determinism

1. Introduction

Both expressing and giving a semantics to I/O has been a long-standing difficulty in functional languages. This is because I/O represents a side effect that changes the state of a unique, uncopyable external world. Given that pure functional languages rely on copying as a means of ensuring referential transparency, the integration of I/O into such languages is problematic.

An early solution was to treat I/O as a special case which resulted in visible side effects—the approach adopted by ML [20] and Miranda [32]. More recent approaches have focussed on sophisticated type and effect systems [21], or have looked at ways and means for restricting programs to a form where I/O accesses are single-threaded. Two key approaches in this are the use of uniqueness types in the language Clean [1, 3] to statically ensure such

single-threaded access, and the use of monads in Haskell [25, 33], which makes it structurally impossible to write I/O in other than a single-threaded manner (if we ignore the unsafe I/O operations such as `unsafePerformIO`).

This paper presents a minimal language called CURIO, which outlines a more flexible approach to both modelling I/O and expressing I/O in pure functional languages such as Haskell and Clean. The motivation for this is deterministic concurrency—relaxing single-threadedness to allow concurrent threads access to separate I/O resources. The benefits of this limited form of concurrency are: (i) the improved facility it makes available for structuring programs which conceptually consist of largely independent threads; and (ii) the ability to allow implementors to use concurrency to optimise performance, safe in the knowledge that the language semantics will be unaffected. Previously there had been few, if any, attempts to formally model this sort of behaviour, especially in the presence of inter-process communication.

CURIO is a small monadic language designed primarily for the theoretical study of I/O, concurrency and determinism. There are five key constructs, `>>=`, `return`, `action`, `par` and `test`, and determinism is enforced using run-time checks. We don’t claim that the use of run-time checks is particularly novel in itself—static checks would indeed be better, and this could be promising future work. It is instead our rigorous semantics for describing the runtime checks that is the focus of this paper. CURIO’s semantics is

- **expressive** One explicitly describes the API and the effects of each individual action. Programs could therefore be subject to formal correctness proofs.
- **general** A wide variety of different side effecting systems and APIs can be modelled under the same broad framework. This includes basic stores, I-structures, inter-process communication constructs, actions which dynamically allocate new structures, as well as traditional I/O, such as file system access.
- **deterministic** There is a formal condition which guarantees confluence in the presence of concurrency.
- **modular** The semantics of complicated, deterministic APIs may be given directly in terms of their smaller components using “I/O model” combinators. I/O models define both an API and the concurrency which it permits. The I/O model we give for Haskell, called `io`, is defined as:

```
term * (smap file * dmap [] chan)
```

Here `term`, `file` and `chan` are models of terminal I/O, single file I/O and inter-process channel communication respectively. The `smap` combinator adds a mapping from filenames to files, the `dmap` combinator provides a facility for dynamic allocation, and the “*” combinator combines the APIs of two I/O models.

The paper is structured as follows: Section 2 presents an informal introduction to a possible extension to Haskell’s I/O API. This is used as motivation for the more general CURIO system.

* Supported by Enterprise Ireland Research Grant SC/2002/0283

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

ICFP’06 September 16–21, 2006, Portland, Oregon, USA.
Copyright © 2006 ACM 1-59593-309-3/06/0009...\$5.00.

Sections 3 and 4 give a formal semantics for CURIO based on general I/O models, constructed from modified state transformers, and a (machine verified) proof of a condition under which confluence holds. Section 5 describes an extension to the I/O model concept, namely location-based models. These facilitate the modelling of dynamic allocation, and they are proved to be interconvertible with the original I/O models. Section 6 gives a description of combinators that allow us to build complex I/O models from simpler building blocks. Section 7 shows how this core language may provide a transparent, semantic back-end to a subset of Haskell’s I/O API, extended to permit deterministic concurrency as described in Section 2. In Sections 8 and 9 we discuss related and future work, and Appendix A contains additional technical material.

The work in this paper significantly extends that of [10]. This earlier work did not mention inter-process communication, combinators or dynamic allocation, and the confluence proof was weaker.

2. A programmer-oriented view of CURIO

We begin in this section by giving an example of a functional I/O programming problem, and then informally describing a language and API extension to solve this.

2.1 Why deterministic concurrency?

I/O is sequenced in Haskell using monads [25, 33]. This is an elegant solution to the problem of I/O in pure functional languages, and monads obey a small set of useful algebraic laws. Monadic I/O is performed using the Haskell functions `>>=` and `return`. `return v` is a program which immediately terminates, yielding value `v`, and `m >>= f` is a program which first performs `m`, and then, if `m` terminates yielding `v`, performs `f v`, the application of function `f` to `v`. In Haskell, a program performing I/O which returns a value of type `β` is of type `IO β`. Example I/O operations are as follows:

```
return  :: ∀β. β → IO β
(>>=)  :: ∀β. ∀γ. IO β → (β → IO γ) → IO γ
openFile :: FilePath → IOMode → IO Handle
hGetChar :: Handle → IO Char
```

A recognised difficulty with using monads to enforce single-threaded access to the I/O world is that the programmer is forced to over-sequentialize their I/O code: I/O operations that are independent and which could be performed in any order still need to be given an explicit ordering in the program. In Peyton Jones’ 15 year retrospective on the Haskell programming language [22] he mentioned that a more fine-grained way of partitioning I/O effects was one of two significant open challenges associated with the use of monads. What CURIO seeks to provide is a limited form of concurrency that reduces the need to over-sequentialize I/O code, while still retaining the determinism that allows us to use referential transparency when reasoning about I/O effects.

Concurrency, even when deterministic, and therefore limited, will still improve efficiency and user response time for I/O-bound programs, increase program reusability, and generally provide a better specification of program behaviour. As an example, consider a file compressor program to which we wish to give a terminal-based user interface. The file compressor and the interface should probably be understood as two relatively distinct pieces of code—one reads and writes to files, and the other communicates with the user. Yet if actions are sequenced explicitly then either user input will have to halt while a file is being encoded:

```
compressor fs = do
  cmd <- interact_with_user
  ...
  compress_whole_file fsource ftarg
```

or the two applications will have to be entangled together: (here, and throughout this paper, we use Haskell’s convenient “do-notation”)

```
compressor fs = do
  info <- compress_for_100ms_seconds fs
  ...
  interact_with_user
  compressor fs
```

Neither solution is particularly desirable. What we need is (i) a way of spawning processes, giving each certain permissions which don’t conflict with one another—such as exclusive access to different files—and (ii) a way of creating 1-to-1 communication channels on the fly so that two processes can communicate with one another.

2.2 Concurrency and communication—`parIO` and `newChannel`

When performing concurrency, to retain determinism we must split the permissions of the original process so that they are distributed among the two new processes. When these processes terminate, the permissions should then be merged once again. Since the termination of the child processes must be synchronised, we adopt a symmetric concurrency construct in the same style as [18], in contrast to Concurrent Haskell’s UNIX-like `forkIO` [24].

The question remains: how does the programmer specify the permissions, if available, that he or she would like each process to receive? Ideally the type system should handle these sorts of constraints, but for the moment we just give a basic solution: the programmer must specify the permissions directly. In our case, we may request that a process be given access to individual file handles. The call to create two parallel I/O processes is `parIO`, and has the following type, polymorphic in `β` and `γ`:

$$\text{parIO} :: \forall\beta. \forall\gamma. [\text{Handle}] \rightarrow \text{IO } \beta \rightarrow [\text{Handle}] \rightarrow \text{IO } \gamma \rightarrow \text{IO } (\beta, \gamma)$$

`parIO hsl m1 hsr mr` concurrently executes programs `m1` and `mr`, attempting to give `m1` access to the handles in list `hsl`, and `mr` access to the handles in list `hsr`. It terminates when `m1` and `mr` have both terminated, returning their respective return values in a tuple. If there is any ambiguity or overlap in the permissions which we wish each process to have, by default the right-hand side receives access to the resource.

Our solution to inter-process communication is to only permit characters to be communicated, and to allow Haskell’s `Handle` type to also refer to either the read or write handle of a communication channel. This means that when using `parIO` we can distribute read or write permissions to a channel as well. In order to guarantee determinism, we can only have one reader and one writer for any given channel, and, furthermore, the reading process may not check if data is available or perform time-outs while waiting for data.

`newChannel` creates a new communication channel, returning the new read and write handle.

$$\text{newChannel} :: \text{IO } (\text{Handle}, \text{Handle})$$

The outline of a file compressor is given in Figure 1 using `parIO` and `newChannel`. The functions

$$\text{wholeFile}, \text{readAccess} :: \text{FilePath} \rightarrow \text{IO Handle}$$

construct dummy handles which are used to indicate that a process would like complete access or read access to a file.

2.3 Process capabilities

We enforce determinism at runtime, not statically. If a process tries to perform an action which is forbidden, it fails, causing the whole

```

compressor :: [(FilePath,FilePath,Handle)] -> IO ()
compressor pdata = do
  cmd <- interact_with_user
  case cmd of
    CompressCmd fsource ftarg -> do
      (cr,cw) <- newChannel
      (retv,_) <- parIO
        [wholeFile ftarg, readAccess fsource, cw]
        (compress_and_communicate ftarg fsource cw)
        [cr, stdin, stdout]
        (compressor ((fsource,ftarg,cr):pdata))
    ExitCmd -> return ()
  ... other commands

```

Figure 1. Fragments of a file compressor

program to fail. To give a safer way for a program to deal with this, we must provide a few extra API calls to allow a process to check what actions are permitted.

```

hAllowed :: Handle -> IO Bool
fAllowedW :: FilePath -> IO Bool

```

`hAllowed` checks to see if we can use the handle in question, be it a read handle or write handle. `fAllowedW` checks if a process has exclusive access to a file. Therefore, if we were to write defensive code, before attempting file access we would first check that our process had access to that file.

```

compress_and_communicate ftarg fsource cw =
  do b1 <- fAllowedW ftarg
  if b1 then do
    ht <- openFile ftarg ReadWriteMode
    .... compress file 'fsource'
    hClose ht
  else (return ())

```

`parIO` defines a strict, almost syntactic binary tree of permissions (only the leaf processes are active at any one time). There is therefore no way for a process to exchange its permissions with that of another process at runtime.

3. Formal Semantics for CURIO

The API extension outlined in the previous section is a plausible one, but it is only an informal description, and the solution appears ad hoc and highly specific. In this section we describe a more general-purpose language called CURIO.

3.1 Generalising an API

Formally describing the API extension in Section 2 presents some immediate difficulties. We need a semantics for I/O actions (how do we usefully describe the meaning of file access, inter-processes communication and terminal I/O?), a formal definition of the non-interference of these actions, and a rigorous guarantee that our whole language is deterministic.

We begin by abstracting away from our API. Since we generalise over the API, we must generalise also over the monadic type constructor `IO`. The type $\text{Prog}_s \beta$ denotes a CURIO program which returns a value of type β , and whose API is defined by I/O model s . In our general framework, we define

$$\text{IO } \beta \triangleq \text{Prog}_{\text{io}} \beta$$

where `io` is the I/O model defined, later, in Section 7.

Let us make the observation that the API calls of the previous section can be partitioned into three general classes:

- Those which perform actions, like `hPutChar`, `newChannel` and `openFile`.
- `parIO`, which is used to initiate concurrency.
- Those which check if actions can be performed, like `hAllowed`.

In place of these three classes we substitute three general primitives: `action`, `par` and `test`.

```

action ::  $\alpha \rightarrow \text{Prog}_s \nu$ 
par    ::  $\forall \beta. \forall \gamma. \rho \rightarrow \text{Prog}_s \beta \rightarrow \text{Prog}_s \gamma \rightarrow \text{Prog}_s (\beta, \gamma)$ 
test   ::  $\alpha \rightarrow \text{Prog}_s \text{Bool}$ 

```

α is the type of primitive actions. `action` a performs an (atomic) action a , and returns a value of type ν . In Haskell, `getChar`, which reads a character from `stdin`, is just defined as primitive and at runtime executes the corresponding system call.

```
primitive getChar :: IO Char
```

Our aim, with the design of CURIO, was to retain this high level I/O interface but replace the primitive I/O actions with a semantics which tries to usefully model its behaviour. In CURIO, `getChar` is a library function implemented internally using `action` (the actual semantics will be given much later).

```
getChar = do (Left c) <- action (Left GetCh)
            return c
```

The “Left” constructors occur because we are selecting a variant of a sum type, produced by the I/O combinators (Section 6).

`par` is a generalisation of `parIO`. In Section 2, when performing concurrency one had to supply two lists of handles. With `par`, the parameter of type ρ abstracts away from this. `parIO` is implemented as:

```
parIO hsl ml hsr mr = par (encodeP hsl hsr) ml mr
  where encodeP hsl hsr = .....
```

where the function `encodeP` encodes the handle-lists as a single parameter of type ρ . I/O models, as defined later, monomorphically bind the three types α , ν and ρ .

`test` a returns `True` if action a is currently permitted, and can be used to implement `hAllowed`.

CURIO is monadic, so we also add `return` and `>>=`, bringing our total number of primitives to five.

```
return ::  $\forall \beta. \beta \rightarrow \text{Prog}_s \beta$ 
(>>=) ::  $\forall \beta. \forall \gamma. \text{Prog}_s \beta \rightarrow (\beta \rightarrow \text{Prog}_s \gamma) \rightarrow \text{Prog}_s \gamma$ 
```

It is worth noting that the second argument of `>>=` is a function taking an arbitrary type β and returning a CURIO program of type $\text{Prog}_s \gamma$. This function argument can be expressed using the full range of constructs available in a lazy functional language, and not just the five CURIO primitives.

Finally, generalising over an API appears to be quite unusual. A reasonable question is, what can we expect to gain from this generality? There are three main benefits: (i) it shall become apparent that our language extension is applicable to many other situations; (ii) practically speaking it is easier to prove meta-properties for a small language; and (iii), we shall show later how small APIs which are unrealistic in their own right can be combined to form powerful and highly practical ones.

3.2 Semantics of actions

We now require a framework for describing the behaviour of our atomic actions. A state transformer is often satisfactory for this:

$$\text{st} :: \alpha \rightarrow \omega \rightarrow (\omega, \nu)$$

Here, a mathematical function `st` defines the effect of each action of type α on a global state of type ω , and the value of type ν

which it returns. So, if we were describing the effects of actions on a store which was a mapping from variables identified by integers, to integer values, the function might be something like

$$\begin{array}{l} \text{st}_{\text{store}} \quad (\text{Write } i \ v) \quad m \quad = \quad (m[i \mapsto v], 0) \\ \text{st}_{\text{store}} \quad (\text{Read } i) \quad m \quad = \quad (m, m(i)) \end{array}$$

However, we want inter-process communication, too. A single communication channel could possibly be described in a state-like manner using a buffer implemented as a list. The sender and receiver independently manipulate the structure—the sender places a value at the end of the buffer, and the receiver takes the first item from the list.

$$\begin{array}{l} \text{st}_{\text{buffer}} \quad (\text{Send } v) \quad vs \quad = \quad (vs \# [v], 0) \\ \text{st}_{\text{buffer}} \quad (\text{Rcve}) \quad (v:vs) \quad = \quad (vs, v) \end{array}$$

The difficulty here is that we have not specified what happens when “receiving” from an empty buffer. The receive action could return a token value indicating an error, but this would go against our desire for determinism—depending on scheduling, a receiving process may behave differently depending on whether a sending process has or has not got around to sending data.

What we need is a way of synchronising the access to shared state, and we modify our state transformer accordingly:

$$\text{st} :: \alpha \rightarrow \omega \rightarrow (\omega, \nu) + \$$$

If $\text{st } a \ w = \$$ then action a is said to be stalled (in world state w), and this indicates that it should not be performed just yet. If all actions are single-threaded, stalling makes no sense, but when other concurrent processes are also making arbitrary changes to the state, stalling allows one action to wait for another. Now, attempting to receive from an empty buffer will stall until another process has written to that buffer:

$$\begin{array}{l} \text{st}_{\text{buffer}} \quad (\text{Rcve}) \quad (v:vs) \quad = \quad (vs, v) \\ \text{st}_{\text{buffer}} \quad (\text{Rcve}) \quad [] \quad = \quad \$ \end{array}$$

We make one final modification to state transformers. We add the possibility that an action can fail.

$$\text{st} :: \alpha \rightarrow \omega \rightarrow (\omega, \nu) + \$ + \perp$$

If $\text{st } a \ w = \perp$ then action a will fail in world state w . If an action fails, then so does the whole program.

3.3 I/O contexts and I/O models

We have defined the semantics of actions of type α . Now we need to give the behaviour of ρ , which describes how process permissions are split. Each process’ permissions are a set of actions, or an element of $\mathcal{P}\alpha$, the set of all subsets of α ¹. We refer to these permission sets as “I/O contexts.” When initiating concurrency, the type ρ indicates how the current set of actions is to be distributed among the two new child processes. The function pf models this:

$$\text{pf} :: \rho \rightarrow \mathcal{P}\alpha \rightarrow (\mathcal{P}\alpha, \mathcal{P}\alpha)$$

We have no need for a corresponding function which “combines” I/O contexts once two concurrent processes have terminated. This is because the parent process never goes away. It is merely suspended while it awaits the termination of both child processes.

We may now give the formal definition for I/O models. An I/O model is a tuple containing the functions st and pf , and it binds the types ν, α, ρ and ω :

$$\text{IOModel } \nu \ \alpha \ \rho \ \omega \quad \triangleq \quad (\text{st} :: \alpha \rightarrow \omega \rightarrow (\omega, \nu) + \$ + \perp, \text{pf} :: \rho \rightarrow \mathcal{P}\alpha \rightarrow (\mathcal{P}\alpha, \mathcal{P}\alpha))$$

¹ Often these sets of actions will be infinite, which makes them unworkable as a data structure for a real implementation, but viewing I/O contexts as sets gives a simpler language definition. The metalanguage encoding, shown in Appendix A.1, uses a more practical solution.

This describes an API along with the potential for deterministic concurrency which it exhibits. Complete I/O models will be given in the next section, along with example programs.

3.4 Semantics of CURIO

The nondeterministic small-step semantics for CURIO is defined in Figure 2. Appendix A.1 briefly describes the meta-encoding of this language in the metalanguage, Core-Clean.

All the reduction rules describe the behaviour of what we call a world/program pair, written $w \Vdash m$, where w is the world and m is the program. This allows one to describe how a program and world state interact over time. The context $C \in \mathcal{P}\alpha$ in which a program is run also affects how the program behaves, so reduction rules are annotated with the current context. We say a CURIO program is a value if it is of the form $\text{return } v$ for some v .

There are three reduction relations, \longrightarrow^C , \uparrow^C and \downarrow^C .

$$\begin{array}{l} w \Vdash m \longrightarrow^C w' \Vdash m' \quad \triangleq \quad \text{“}w \Vdash m \text{ can reduce to } w' \Vdash m' \\ \quad \text{in context } C\text{”} \\ w \Vdash m \uparrow^C \quad \triangleq \quad \text{“}w \Vdash m \text{ can fail in context } C\text{”} \\ w \Vdash m \downarrow^C \quad \triangleq \quad \text{“}w \Vdash m \text{ is in normal form in } C\text{”} \end{array}$$

CURIO is defined in a similar fashion to Haskell’s I/O extensions in [23]. The behaviour of the I/O primitives is defined as an operational semantics layered on top of an implicit denotational semantics of the host language. This explains why it is necessary to mention failure, and the behaviour of an undefined CURIO program, \perp . A failed small-step reduction may denote a never-ending sequence of reduction steps or a runtime error in the metalanguage.

A world/program pair is in normal form if it is either a single value, or all concurrent actions it is attempting to perform are stalled. We say a world/program pair is in normal form, rather than can be in normal form, because despite nondeterminism it can be shown that a world/program pair is in normal form if and only if it cannot either fail or successfully reduce.

The runtime checks occur with each `action`, `test` and `par`. The behaviour of the first two is dependent on whether the action is permitted. `par` requires `pf` to be called and this in practice need only happen once, when the two new processes are first spawned.

4. A precondition for confluence

As is typical for concurrent languages, CURIO’s small-step semantics is nondeterministic. In this section we describe determinism in CURIO, define a precondition PRE_s for I/O models which is formally proved to guarantee this, and give a few example I/O models.

4.1 Defining PRE_s

To find a precondition for guaranteeing confluence we must first examine what it means for two actions to be order independent.

Let’s begin with the basic case where actions cannot become stalled. By a standard definition, actions a_l and a_r are order independent if, for all w, w_2, v_l and v_r ,

$$\begin{array}{c} (\exists w_1 \in \omega. \text{st } a_l \ w = (w_1, v_l) \wedge \text{st } a_r \ w_1 = (w_2, v_r)) \\ \iff \\ (\exists w'_1 \in \omega. \text{st } a_r \ w = (w'_1, v_r) \wedge \text{st } a_l \ w'_1 = (w_2, v_l)) \end{array}$$

Under this definition, for every initial world state, if performing a_l and a_r is successful then swapping the execution order gives rise to the same resultant world state and the same return values—even though the intermediate world state may differ. So, for example, modifying two different locations in a store is order independent. This definition also guarantees that if failure occurs for one ordering, it will also occur for the other.

$$\begin{array}{c}
\frac{w \Vdash m \uparrow^C}{w \Vdash m \gg= f \uparrow^C} \quad \frac{w \Vdash m \downarrow^C}{w \Vdash m \gg= f \downarrow^C} \quad (m \text{ not a value}) \quad \frac{w \Vdash m \longrightarrow^C w' \Vdash m'}{w \Vdash m \gg= f \longrightarrow^C w' \Vdash m' \gg= f} \\
w \Vdash \text{return } v \gg= f \longrightarrow^C w \Vdash f v \quad w \Vdash \text{return } v \downarrow^C \quad w \Vdash \perp \uparrow^C \\
\frac{a \in C \quad \text{st } a w = (w', v)}{w \Vdash \text{action } a \longrightarrow^C w' \Vdash \text{return } v} \quad \frac{a \in C \quad \text{st } a w = \$}{w \Vdash \text{action } a \downarrow^C} \quad \frac{a \in C \quad \text{st } a w = \perp}{w \Vdash \text{action } a \uparrow^C} \\
\frac{a \notin C}{w \Vdash \text{action } a \uparrow^C} \quad \frac{a \in C}{w \Vdash \text{test } a \longrightarrow^C w \Vdash \text{return True}} \quad \frac{a \notin C}{w \Vdash \text{test } a \longrightarrow^C w \Vdash \text{return False}} \\
\text{pf } p C = (C_l, C_r) \left\{ \begin{array}{l}
w \Vdash \text{par } p (\text{return } v_l) (\text{return } v_r) \longrightarrow^C w \Vdash \text{return } (v_l, v_r) \\
\frac{w \Vdash m_l \longrightarrow^{C_l} w' \Vdash m'_l}{w \Vdash \text{par } p m_l m_r \longrightarrow^C w' \Vdash \text{par } p m'_l m_r} \quad (m_r \neq \perp) \quad \frac{w \Vdash m_l \uparrow^{C_l}}{w \Vdash \text{par } p m_l m_r \uparrow^C} \\
\frac{w \Vdash m_r \longrightarrow^{C_r} w' \Vdash m'_r}{w \Vdash \text{par } p m_l m_r \longrightarrow^C w' \Vdash \text{par } p m_l m'_r} \quad (m_l \neq \perp) \quad \frac{w \Vdash m_r \uparrow^{C_r}}{w \Vdash \text{par } p m_l m_r \uparrow^C} \\
\frac{w \Vdash m_l \downarrow^{C_l} \quad w \Vdash m_r \downarrow^{C_r}}{w \Vdash \text{par } p m_l m_r \downarrow^C} \quad (m_l, m_r \text{ not both values})
\end{array} \right.
\end{array}$$

Figure 2. Nondeterministic small-step semantics for CURIO

When actions can become stalled, this definition must be extended. The crucial (and, as far as we know, original) observation is that if one action succeeds and the other is then stalled, *the stalled action must already have been stalled prior to the successful execution of the first action*. Put another way: two actions are order independent only if one cannot cause the other, if initially unstalled, to become stalled. For example, with model `bfft`, appending to the buffer cannot cause a receive action to become stalled. So these actions are order independent, since, when they are both unstalled, they also modify different parts of the same buffer.

We define two relations on actions to capture these ideas, \parallel_s and ally_s . Both are defined in Figure 3.

- $a_l \parallel_s a_r$: for any two actions a_l and a_r , if neither are stalled then the order in which they are executed is irrelevant—both with regard to their effect on world state and their return values.
- $\text{ally}_s(a_l, a_r)$: if after performing a_l action a_r is stalled then it must have been stalled beforehand. The word “ally” hints at the fact that action a_l will not obstruct or hinder action a_r —they are in effect working with each other. This rules out competition for a limited resource.

We define a_l and a_r to be order independent if $a_l \parallel_s a_r$, $\text{ally}_s(a_l, a_r)$, and $\text{ally}_s(a_r, a_l)$, and using these, $C_l \diamond_s C_r$ is defined to mean that all actions in context C_l are order independent with respect to all actions in context C_r . If, when performing concurrency with `par`, the parent process’ I/O context C is split by `pf` into left- and right-hand processes’ contexts C_l and C_r , and we wish to retain determinism, then:

- the child processes must not be able to perform actions forbidden by the parent: $C_l \subseteq C$ and $C_r \subseteq C$.
- the child processes should not interfere with one another: $C_l \diamond_s C_r$

These observations motivate our definition of precondition PRE_s , shown also in Figure 3.

4.2 PRE_s implies determinism

First, define \longrightarrow^C to be the reflexive, transitive closure of the reduction relation \longrightarrow^C . Then define $w \Vdash m \Downarrow^C w' \Vdash m'$ as meaning that both $w \Vdash m \longrightarrow^C w' \Vdash m'$ and $w' \Vdash m' \downarrow^C$ —that it is possible to reduce to a normal form in zero or more steps. We use \Downarrow^C , convergence, when this normal form is unique regardless of reduction order. \uparrow^C , divergence, indicates that there is no combination of steps that lead to a normal form. The formal definition of these relations requires some details of our meta-encoding, and these can be found in Appendix A.1.

$w \Vdash m \Downarrow^C w' \Vdash m'$ implies $w \Vdash m \Downarrow^C w' \Vdash m'$ trivially, and the purpose of our confluence proof is to show that, if PRE_s holds, then $w \Vdash m \Downarrow^C w' \Vdash m'$ implies $w \Vdash m \Downarrow^C w' \Vdash m'$. The full proof [9] required many smaller lemmas, and only a brief sketch is given. For all results below we assume that PRE_s holds.

Lemma 4.1. *If $w \Vdash m \longrightarrow^C w' \Vdash m'$ then either $w = w'$ and for all w_1 , $w_1 \Vdash m \longrightarrow^C w_1 \Vdash m'$, or for some action $a \in C$, $\text{st } a w = (w', v')$, for some v' .*

Proof. Induction on m . □

Lemma 4.2. *If $C_l \diamond_s C_r$, then if $w \Vdash m_l \longrightarrow^{C_l} w_1 \Vdash m'_l$ and $w \Vdash m_r \longrightarrow^{C_r} w_r \Vdash m'_r$, then either*

- *there exists a w_2 such that both $w_r \Vdash m_l \longrightarrow^{C_l} w_2 \Vdash m'_l$ and $w_l \Vdash m_r \longrightarrow^{C_r} w_2 \Vdash m'_r$*
- *or $w_r \Vdash m_l \uparrow^{C_l}$ and $w_l \Vdash m_r \uparrow^{C_r}$.*

Proof. Case analysis, using Lemma 4.1, on whether one reduction performed an action. □

Lemma 4.3. *$w \Vdash m \uparrow^C$ implies $w \Vdash m \uparrow^C$.*

Proof. Proof by induction over the number of reduction steps that program m cannot “escape” failure. □

$$\begin{aligned}
a_l \parallel_s a_r &\triangleq \forall w \in \omega. \forall w_2 \in \omega. \forall v_l \in \nu. \forall v_r \in \nu. \text{st } a_l w \neq \$ \wedge \text{st } a_r w \neq \$ \implies \\
&\quad (\exists w_1 \in \omega. \text{st } a_l w = (w_1, v_l) \wedge \text{st } a_r w_1 = (w_2, v_r)) \\
&\iff \\
&\quad (\exists w'_1 \in \omega. \text{st } a_r w = (w'_1, v_r) \wedge \text{st } a_l w'_1 = (w_2, v_l)) \\
\text{ally}_s(a_l, a_r) &\triangleq \forall w \in \omega. \forall w_1 \in \omega. \forall v \in \nu. \text{st } a_l w = (w_1, v) \wedge \text{st } a_r w_1 = \$ \implies \text{st } a_r w = \$ \\
C_l \diamond_s C_r &\triangleq \forall a_l \in C_l. \forall a_r \in C_r. a_l \parallel_s a_r \wedge \text{ally}_s(a_l, a_r) \wedge \text{ally}_s(a_r, a_l) \\
\text{PRE}_s &\triangleq \forall p \in \rho. \forall C \in \mathcal{P}\alpha. \forall C_l \in \mathcal{P}\alpha. \forall C_r \in \mathcal{P}\alpha. \text{pf } p C = (C_l, C_r) \implies C_l \subseteq C \wedge C_r \subseteq C \wedge C_l \diamond_s C_r
\end{aligned}$$

Figure 3. Formal definition of PRE_s and associated relations

Lemma 4.4. *If both $w \Vdash m \longrightarrow^C w_1 \Vdash m_1$ and $w_1 \Vdash m_1 \Downarrow^C$ then $w \Vdash m \Downarrow^C w_1 \Vdash m_1$.*

Proof. Induction on m , using Lemmas 4.3 and 4.2. This proves that if there is one small-step reduction before reaching normal form, then the redex must be unique. \square

Lemma 4.5. *Diamond Property. If $w \Vdash m \longrightarrow^C w_1 \Vdash m_1$ and $w \Vdash m \longrightarrow^C w_2 \Vdash m_2$ then either*

- *The same redex was reduced: $w_1 = w_2$ and $m_1 = m_2$.*
- *There is a common reduct: for some w_3 and m_3 , $w_1 \Vdash m_1 \longrightarrow^C w_3 \Vdash m_3$ and $w_2 \Vdash m_2 \longrightarrow^C w_3 \Vdash m_3$.*
- *Both sides will diverge: $w_1 \Vdash m_1 \Uparrow^C$ and $w_2 \Vdash m_2 \Uparrow^C$.*

Proof. Induction on m using Lemmas 4.4 and 4.3, and Lemma 4.2 as one of the base cases. \square

Theorem 4.1. *Confluence. If $w \Vdash m \Downarrow^C w_1 \Vdash m_1$ then $w \Vdash m \Downarrow^C w_1 \Vdash m_1$.*

Proof. Induction on the number of reduction steps, using the diamond property and Lemma 4.3. \square

4.3 Small Examples

We now give three small, and, on their own, largely impractical example I/O models, all of which obey PRE_s . These serve as easy examples of the relevant concepts and will also be used when assembling more practical I/O models using the combinators described, later, in Section 6.

For the store example given in the previous section, one possible context splitter pf_{store} is that given in Figure 4. The context-splitting parameter ρ_{store} is defined to be a list of pairs, the first component of which identifies the store location, whilst the second denotes the required permission: read (WantR) or write (WantW). When performing concurrency with par one lists the permissions we would like to go to the left-hand process, if possible, and all the rest go the right-hand process.

```

storeExample :: Progstore (Int,Int)
storeExample = par [(0,WantW),(1,WantR)]
  -- LHS wants to read int 1 and write to int 0
  (do {i1 <- action (Read 1); i0 <- action (Read 0)
      ; action (Write 0 i1); return i0})
  -- RHS wants to read int 1 and write to int 2
  (do {i1 <- action (Read 1); i2 <- action (Read 2)
      ; action (Write 2 i1); return i2})

```

The proof of $\text{PRE}_{\text{store}}$ relies on the fact that (i) all reads or writes to different integer variables are order independent, (ii) reads from the same integer variable are order independent, and (iii) that

pf_{store} guarantees that if one child process can write to a particular variable, the other must not be able to access it at all.

For an individual communication buffer, bfft , a receive may be performed concurrently with a send, but that is really the only opportunity for deterministic concurrency. The idea for pf_{bfft} is similar but on a smaller scale:

$$\begin{aligned}
\rho_{\text{bfft}} &\triangleq \text{WantS} + \text{WantR} + \text{WantRS} + \text{None} \\
\text{pf}_{\text{bfft}} p C &= \begin{cases} (C \setminus \{\text{Rcve}\}, C \setminus \{\text{Write } v \mid v \in \text{Int}\}) & , p = \text{WantS} \\ (C \setminus \{\text{Write } v \mid v \in \text{Int}\}, C \setminus \{\text{Rcve}\}) & , p = \text{WantR} \\ (C, \emptyset) & , p = \text{WantRS} \\ (\emptyset, C) & , p = \text{None} \end{cases}
\end{aligned}$$

Here we note that we can request permissions to perform zero (None), one (WantR or WantS) or both (WantRS) of the two actions available on a buffer: send or receive. In all cases, each permission can only be allocated to one process.

PRE_{bfft} holds because the only two actions for which we must prove a non-interference property are those of the form Rcve and $\text{Send } i$ (two concurrent Sends are forbidden, as are two concurrent Rcves). Neither can cause the other to become stalled, and when both are unstalled they don't interfere with one another.

Our final example is model istr , a single integer I-structure [2]. I-structures are concurrent, deterministic, write once data structures which have been used to develop elegant algorithms for computing matrices efficiently.

$$\begin{aligned}
\text{st}_{\text{istr}} \text{ (ReadI)} \quad \text{EMPTY} &= \$ \\
\text{st}_{\text{istr}} \text{ (ReadI)} \quad \text{FULL } i &= (\text{FULL } i, i) \\
\text{st}_{\text{istr}} \text{ (WriteI } i_1) \quad \text{EMPTY} &= (\text{FULL } i_1, 0) \\
\text{st}_{\text{istr}} \text{ (WriteI } i_1) \quad \text{FULL } i &= \perp \\
\text{pf}_{\text{istr}} \text{ ()} \quad C &= (C, C)
\end{aligned}$$

If the structure is EMPTY a ReadI will wait until a concurrent WriteI has written to it. Writing to an already full I-structure causes failure. We have no need for I/O contexts in model istr —all actions can be performed concurrently with one another, and ρ_{istr} is Haskell's unit type, $()$.

To prove PRE_{istr} one first shows that no action can cause another to become stalled. Then one shows that the only occasion in which two actions are unstalled is in the case of two consecutive WriteIs (which fail, regardless of ordering), a WriteI and a ReadI to a full I-structure (in which the write always causes failure, regardless of ordering), or two ReadIs to a full I-structure (which are order independent).

Since model istr is so small, we can use it to give some example CURIO reductions. Two concurrent ReadIs attempting to read from an empty buffer is in normal form:

$$\text{EMPTY} \Vdash \text{par } () \text{ (action ReadI) (action ReadI)} \Downarrow^{\alpha_{\text{istr}}}$$

$$\begin{aligned}
\rho_{\text{store}} &\triangleq [(\text{Int}, \text{WantW} + \text{WantR})] \\
\text{pf}_{\text{store}} \text{ } ps \ C &= (C \setminus \{\text{Write } i \ v \mid v \in \text{Int}, (i, \text{WantW}) \notin ps \vee (i, \text{WantR}) \in ps\} \setminus \{\text{Read } i \mid (i, \text{WantW}) \notin ps\}, \\
&\quad C \setminus \{\text{Write } i \ v \mid v \in \text{Int}, (i, \text{WantW}) \in ps \vee (i, \text{WantR}) \in ps\} \setminus \{\text{Read } i \mid (i, \text{WantW}) \in ps\})
\end{aligned}$$

Figure 4. Context splitting function for model *store*

A `ReadI` performed concurrently with a `WriteI` cannot proceed until the I-structure is full:

$$\begin{aligned}
\text{EMPTY} \Vdash \text{par } () \ (\text{action } (\text{WriteI } 7)) \ (\text{action } \text{ReadI}) \\
\quad \longrightarrow^{\alpha_{\text{istr}}} \\
\text{FULL } 7 \Vdash \text{par } () \ (\text{return } 0) \ (\text{action } \text{ReadI}) \\
\quad \longrightarrow^{\alpha_{\text{istr}}} \\
\text{FULL } 7 \Vdash \text{par } () \ (\text{return } 0) \ (\text{return } 7) \\
\quad \longrightarrow^{\alpha_{\text{istr}}} \\
\text{FULL } 7 \Vdash \text{return } (0, 7)
\end{aligned}$$

which is in normal form.

It should be noted that our I/O models make no mention of initial world states or initial, outermost contexts. We do not mention initial world states because a programmer, in general, has no control over the contents, say, of the file system. The natural choice for an outermost context is “every action”. This is not made explicit anywhere because, from the point of view of proving general language properties, it is just one special case.

5. Location-based models

5.1 Summary of the problem and solution

This section tackles a specific technical difficulty which requires our attention—how to model situations where an API call returns a pointer or handle to a brand new structure. Examples of such actions from Section 2 are `newChannel` and `openFile`. The problem with these API calls is that any naive implementation just allocates the “next” handle or structure available in a linear fashion (for example, by appending an item to a list and returning the index). This immediately means that two allocations become order dependent. In practice it usually does not matter, since the handle returned by these APIs is only a token whose literal, numeric value should not be examined. But, without a major overhaul of our semantics, there is no way for us to enforce this or state it formally.

The solution is to modify our semantics, I/O model and precondition very slightly so that actions can distinguish different calling processes. In these new, “location-based” I/O models, a process now has an associated “location” as well as a context. The location identifies the process at runtime and is given as an extra hidden parameter to both `st` and `pf`. This, on its own, allows allocation to be order independent.

This is a small, often transparent extension only for convenience, and it does not make I/O models any more powerful. To show this we give the semantics of location-based models in terms of location-free ones, and prove that confluence is still maintained. In the “location-free” I/O models an action could not directly determine a unique identity for a calling process. However, if we encode the location in the context, the process itself could still do so using the `test` command. When converting back to a normal model, a process, using `test`, determines the location in which it is running, and passes this as an extra explicit argument.

5.2 Location-based I/O models

A location must uniquely identify processes. The `par` primitive in effect creates a binary tree of processes, so we define locations, \mathbb{L} ,

to be lists of L/R values. $l_1 \preceq l_2$ means l_1 is a prefix of l_2 .

$$\mathbb{L} \triangleq [\text{L} + \text{R}] \quad l_1 \preceq l_2 \Leftrightarrow \exists l_3 \in \mathbb{L}. l_2 = l_1 + l_3$$

Location-based models are defined below. The old functions `st` and `pf` now have an extra \mathbb{L} parameter:

$$\begin{aligned}
s : \text{IOModel}' \ \nu \ \alpha \ \rho \ \omega &\triangleq (\\
\text{st}' :: (\mathbb{L}, \alpha) &\rightarrow \omega \rightarrow (\omega, \nu) + \$ + \perp, \\
\text{pf}' :: (\mathbb{L}, \rho) &\rightarrow \mathcal{P}\alpha \rightarrow (\mathcal{P}\alpha, \mathcal{P}\alpha)
\end{aligned}$$

Since location-based models are a small modification to location-free ones, despite appearances, we shall indicate them simply by adding a prime ($'$), so `IOModel'` and `st'`, for example, refer to the location-based counterparts of `IOModel` and `st`, respectively.

The types of the five CURIO primitives are unchanged, and there are only small modifications to be made to the small-step semantics. The three reduction relations are now annotated with locations as well (\longrightarrow^C now becomes $\longrightarrow^{l;C}$). This location is supplied to `st'` implicitly for all action rules. For example:

$$\frac{a \in C \quad \text{st}'(l, a) \ w = (w_1, v)}{w \Vdash \text{action } a \longrightarrow^{l;C} w_1 \Vdash \text{return } v}$$

The `par` rules are all modified so that there is also a left- and right-hand location l_l and l_r :

$$\text{pf}'(l, p) \ C = (C_l, C_r) \quad l_l = l + [\text{L}] \quad l_r = l + [\text{R}]$$

and in the individual rules, C , C_l and C_r are changed to $l; C$, $l_l; C_l$ and $l_r; C_r$ respectively. For example:

$$\frac{w \Vdash m_l \downarrow^{l_l; C_l} \quad w \Vdash m_r \downarrow^{l_r; C_r}}{w \Vdash \text{par } p \ m_l \ m_r \downarrow^{l; C}} \quad (m_l, m_r \text{ not both values})$$

The new precondition PRE'_s is textually almost identical. The only difference is the types of a_l , a_r and p :

$$\begin{aligned}
C_l \diamond'_s C_r &\triangleq \forall a_l \in (\mathbb{L}, C_l). \forall a_r \in (\mathbb{L}, C_r). \dots \\
\text{PRE}'_s &\triangleq \forall p \in (\mathbb{L}, \rho). \dots
\end{aligned}$$

5.3 Converting between models

The function `toM'` converts a location-free model to a location-based one. In the new model all extra location information is ignored. It is not difficult to show that PRE_s implies $\text{PRE}'_{\text{toM}' \ s}$.

$$\begin{aligned}
\text{toM}' &:: \text{IOModel } \nu \ \alpha \ \rho \ \omega \rightarrow \text{IOModel}' \ \nu \ \alpha \ \rho \ \omega \\
\text{toM}'(\text{st}, \text{pf}) &\triangleq (\lambda(l, a). \text{st } a, \lambda(l, p). \text{pf } p)
\end{aligned}$$

From now we shall work entirely with location-based models. Since the conversion to location-based models is so simple, we will just implicitly assume the above conversion.

To show that PRE'_s guarantees confluence for the new, modified language, we define a function `toM` which gives the semantics of location-based models by giving the equivalent location-free ones:

$$\text{toM} :: \text{IOModel}' \ \nu \ \alpha \ \rho \ \omega \rightarrow \text{IOModel } \nu \ (\mathbb{L}, \text{Act } \alpha + \text{Probe}) \ \rho \ \omega$$

In the resultant model the location argument becomes explicit in all actions. The rough idea is that if a process' context is C , its location is l , and $a \in C$, then, when we convert everything back to location-free I/O models, action $(l_1, \text{Act } a)$ is allowed for all locations l_1 such that $l \preceq l_1$. The extra action (l_1, Probe) always

$$\begin{aligned}
& (\text{st}'_1, \text{pf}'_1) * (\text{st}'_2, \text{pf}'_2) \triangleq (\text{st}'_3, \text{pf}'_3) \\
& \text{where} \\
& \text{st}'_3 (l, \text{Left } a_1) (w_1, w_2) = \\
& \quad \begin{cases} ((w'_1, w'_2), \text{Left } v) & , \text{st}'_1 (l, a_1) w_1 = (w'_1, v) \\ \$ & , \text{st}'_1 (l, a_1) w_1 = \$ \\ \perp & , \text{st}'_1 (l, a_1) w_1 = \perp \end{cases} \\
& \text{st}'_3 (l, \text{Right } a_2) (w_1, w_2) = \\
& \quad \begin{cases} ((w_1, w'_2), \text{Right } v) & , \text{st}'_2 (l, a_2) w_2 = (w'_2, v) \\ \$ & , \text{st}'_2 (l, a_2) w_2 = \$ \\ \perp & , \text{st}'_2 (l, a_2) w_2 = \perp \end{cases} \\
& \text{pf}'_3 (l, (p_1, p_2)) (C_1, C_2) = ((C_{l_1}, C_{l_2}), (C_{r_1}, C_{r_2})) \\
& \text{where} \\
& (C_{l_1}, C_{r_1}) = \text{pf}'_1 (l, p_1) C_1 \\
& (C_{l_2}, C_{r_2}) = \text{pf}'_2 (l, p_2) C_2
\end{aligned}$$

Figure 6. Definition of cartesian product combinator

fails when executed, but is *permitted* if the process' location is l , where $l \preceq l_1$. This allows a process in the old model, using `test`, to calculate its current location correctly. Further details of the proof that PRE'_s implies $\text{PRE}_{\text{toM } s}$ can be found in Appendix A.2.

6. I/O model combinators

This section defines general-purpose “glue” for combining smaller deterministic I/O models to form larger deterministic ones.

6.1 Introduction

Let us recall the model `store`, in which the world state is a mapping from integers to integers. We could also easily imagine a store of I-structures, or a store of communication buffers. Yet the definition and proof of PRE'_s for each would contain the same boilerplate—actions on different parts of the store are order independent, regardless of the individual store elements. The idea behind combinators is that one first proves small API models confluent, such as `istr` and `bfft`, and one then develops mathematical functions which, given some confluent I/O model s , implements the API “store of s ”, preserving both confluence and the opportunities for determinism which each individual s allows.

We present three combinators in this section, and their types can be found in Figure 5. The idea of combining API models is perhaps slightly surprising. We should emphasise that combinators were initially developed as a practical solution to the problem of machine verifying PRE'_s for large I/O models. An I/O model is just a data structure in our meta-encoding, and combinators are just functions which examine this structure, returning a new, modified one. The advantages to using combinators are (i) they are more modular, rendering unnecessary much boilerplate, (ii) they make clear the *logical* make-up of the API and (iii) since the semantics of I/O is somewhat subjective and open to debate, they let us focus on general properties of APIs, as opposed to specifics.

The important (machine verified) technical results are that, for all I/O models s_1 and s_2 ,

- PRE'_{s_1} and PRE'_{s_2} together imply $\text{PRE}'_{s_1 * s_2}$.
- PRE'_{s_1} implies $\text{PRE}'_{\text{smap } s_1}$.
- for all w_0 , PRE'_{s_1} implies $\text{PRE}'_{\text{dmap } w_0 s_1}$.

6.2 Cartesian product combinator

This combinator is defined in full in Figure 6. $s_1 * s_2$ represents the joining of the APIs of s_1 and s_2 . The world state becomes the cartesian product of the world states of s_1 and s_2 , and an action can be either an action from s_1 or one from s_2 . When performing

$$\begin{aligned}
& \text{dmap } w_0 (\text{st}', \text{pf}') \triangleq (\text{st}'_{\text{dmap}}, \text{pf}'_{\text{dmap}}) \\
& \text{where} \\
& \text{st}'_{\text{dmap}} (l, \text{DAct } h a) p = \\
& \quad \begin{cases} (p[h \mapsto w_1], \text{Left } v) & , \text{st}' (l, a) p(h) = (w_1, v) \\ \$ & , \text{st}' (l, a) p(h) = \$ \\ \perp & , \text{st}' (l, a) p(h) = \perp \end{cases} \\
& \text{st}'_{\text{dmap}} (l, \text{Next}) p = (p, \text{Right } (l, \text{len } p l)) \\
& \text{st}'_{\text{dmap}} (l, \text{Alloc } (l_1, i_1)) p = \\
& \quad \begin{cases} \perp & , l \neq l_1 \text{ or } i_1 \neq (\text{len } p l) \\ (p[(l_1, i_1) \mapsto w_0], \text{Right } ([], 0)), & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 7. Partial definition of `dmap`

concurrency, one must specify the permissions given to each sub-process for both sides of the world state.

For example, model `bfft * istr` is an API which combines access to a buffer and an I-structure.

```

example :: Prog_bfft_istr ()
example = do (Left v) <- action (Left Rcve)
             (r1, r2) <- par (WantS, ())
                           (action (Right (WriteI 1)))
             ...

```

6.3 Static map combinator

The `smap` combinator turns an I/O model s into a “store of s .” The world state is a map from any type δ to the world state of s . This is similar to cartesian product, except there are a (potentially infinite) number of s models, each indexed by a different δ . An action (d, a) performs action a from s on the world state associated with name d . When performing concurrency, one must supply a list of how the permissions for each $d :: \delta$ are to be distributed between the left- and right-hand sides. If some d isn't mentioned in the list, then all the current permissions will be given to the right side, preventing the left side from doing anything at all with d .

For example, imagine an I/O model `ivart`, modelling a single integer variable.

```
ivart :: IOModel Int ((Write Int)+Read) (WantR+WantS) Int
```

The aforementioned `store` model could have been defined instead as `smap ivart`. This would be completely identical except that the actions `(i, Write v)` and `(i, Read)` would replace `Write i v` and `Read i`. This looks less natural, but it is just as powerful.

6.4 Dynamic map combinator

`dmap w_0 s` is an I/O model which allows processes to dynamically allocate individual s models of initial world state w_0 . Each action either performs an action a on a particular ω identified by handle h , `DAct h a`, or it is one of two separate steps required to create a new ω . A partial definition of `dmap` can be found in Figure 7. This is the only combinator to specifically require the location-based constructs of the previous section.

The basic idea is that our world state is partitioned into individual locations. Since each process is identified uniquely by a location, this allows each process to allocate “locally” without the danger of it interfering with another process doing the same at that very moment.

$$\text{Pool } \omega \triangleq \mathbb{L} \rightarrow [\omega] \quad \text{HndP} \triangleq (\mathbb{L}, \text{Int})$$

The resultant world state of `dmap` is of the form `Pool \omega`. The structure `Pool \omega` can be understood as a map from locations to lists of ω (it is implemented, actually, using a strict binary tree of

```

(*) :: IOModel'  $\nu_1$   $\alpha_1$   $\rho_1$   $\omega_1$   $\rightarrow$  IOModel'  $\nu_2$   $\alpha_2$   $\rho_2$   $\omega_2$   $\rightarrow$  IOModel' (Either  $\nu_1$   $\nu_2$ ) (Either  $\alpha_1$   $\alpha_2$ ) ( $\rho_1, \rho_2$ ) ( $\omega_1, \omega_2$ )
smap :: IOModel'  $\nu$   $\alpha$   $\rho$   $\omega$   $\rightarrow$  IOModel'  $\nu$  ( $\delta, \alpha$ ) [( $\delta, \rho$ )] ( $\delta \rightarrow \omega$ )
dmap ::  $\omega \rightarrow$  IOModel'  $\nu$   $\alpha$   $\rho$   $\omega \rightarrow$  IOModel' (Either  $\nu$  HndP) (DynAction HndP  $\alpha$ ) [(HndP,  $\rho$ )] (Pool  $\omega$ )

Either  $\beta$   $\gamma \triangleq$  Left  $\beta$  + Right  $\gamma$     DynAction  $\iota$   $\alpha \triangleq$  DAct  $\iota$   $\alpha$  + Next + Alloc  $\iota$ 

```

Figure 5. Type signatures for combinators

lists of ω). Since at runtime each process has a unique location, this gives each process its own unique list to which it can append new elements. The data allocated is still completely global, and may be accessed long after the creating process has gone—this approach merely ensures that the *act* of allocating data will not affect other processes. In Figure 7 we use a simple shorthand for examining and modifying pools: $p(h)$ looks up pool p at handle h , $p[h \mapsto w]$ is pool p modified so handle h maps to w , and $\text{len } p$ returns the next free index at location l in pool p .

Each element of type ω in this structure is identified uniquely by a handle of type HndP. This handle is a tuple containing a location, and an Int which identifies an individual ω for that particular location. To create a new ω a process first determines, using Next, the next free handle h at its location, and then allocates it with Alloc h . The technical reason for needing two steps is that a process must have complete access to a handle that it allocates, and I/O contexts cannot forbid actions based on their return values, only their arguments.

```

newBffr :: Progdmap [] bffr HndP
newBffr = do (Right h) <- action Next
             action (Alloc h)
             return h

```

Contexts are split in a style similar to that of smap except

- A sub-process should have access to *all* possible future handles that it or its child processes may end up using. Therefore, by default, when a process in location l splits into two processes then all permissions to access the ω identified by handle (l_h, l_h) will go the left-hand side if and only if $l+[L] \preceq l_h$.
- The action Alloc h must have complete, single-threaded access to the ω identified by handle h . Therefore, when doing concurrency, if we distribute permissions to this handle, immediately Alloc h will be forbidden by both children. This is quite reasonable: why would you want to distribute permissions to the structure identified by an as yet unused handle?
- There is no reason to forbid action Next.

We are not interested in deallocation at the moment. Ideally a clever runtime system would look after this behind the scenes.

7. A real world I/O model

Finally, we can define model io, and it is

```
term * (smap file * dmap [] chan)
```

We now briefly define term, file and chan. Figure 8 shows how some of the front-end API calls from Section 2 are implemented.

Our io model ignores quite a few aspects of Haskell's I/O interface. These include: those actions which are semantically transparent, and included for efficiency reasons (i.e. those to do with buffering); verbose error messages; actions which do not give us any further means of exploiting concurrency, for example, those which relate to directory structure; lazy file I/O functions such as

```

data Handle = StdInHnd | StdOutHnd | ChnWrHnd HndP
             | ChnRdHnd HndP | FileHnd String FHand

```

```
(stdin, stdout) = (StdInHnd, StdOutHnd)
```

```
openFile :: String  $\rightarrow$  IOMode  $\rightarrow$  Progio Handle
```

```
openFile n ReadMode = do
  Right (Left (RHndP rh)) <-
    action (Right (Left (n, FNextRdHnd)))
  action (Right (Left (n, HRdOpen rh)))
  return (FileHnd n (RHnd rh))
```

```
openFile n ReadWriteMode = do
  action (Right (Left (n, FWrOpen)))
  return (FileHnd n RWHnd)
```

```
hPutChar :: Handle  $\rightarrow$  Char  $\rightarrow$  Progio ()
```

```
hPutChar hnd c = do
  case hnd of
    StdOutHnd      -> action (Left (PutCh c))
    FileHnd n RWHnd ->
      action (Right (Left (n, FWrite c)))
    ChnWrHnd h     ->
      action (Right (Right (DAct h (Send c))))
  return ()
```

```
newChannel :: Progio (Handle, Handle)
```

```
newChannel = do
  Right (Right (Right h)) <-
    action (Right (Right Next))
  action (Right (Right (Alloc h)))
  return (ChnRdHnd h, ChnWrHnd h)
```

```
hAllowed :: Handle  $\rightarrow$  Progio Bool
```

```
hAllowed hnd = case hnd of
  StdInHnd      -> test (Left GetCh)
  StdOutHnd     -> test (Left (PutCh 'X'))
  FileHnd n RWHnd ->
    test (Right (Left (n, FWrite 'X')))
  FileHnd n h   -> test (Right (Left (n, HGetC h)))
  ChnRdHnd h    ->
    test (Right (Right (DAct h Rcve)))
  ChnWrHnd h    ->
    test (Right (Right (DAct h (Send 'X'))))
```

```
wholeFile :: String  $\rightarrow$  Handle
```

```
wholeFile n = FileHnd n RWHnd
```

```
encodeP :: [Handle]  $\rightarrow$  [Handle]  $\rightarrow$ 
```

```
  ( $\rho_{\text{term}}, [(String, \rho_{\text{file}})], [(HndP, \rho_{\text{chan}})]])$ 
encodeP hsl hsr = (rhoT th, (rhoFS fsh, rhoChs ch))
  where (th, fsh, ch) = splitHandles hsl hsr
        splitHandles .. -- separate 3 handle types
        rhoFS .. -- isolate 'rho' for each file
        ...
```

Figure 8. Semantics of some high-level API calls

file	$:: \text{IOModel}' \nu_{\text{file}} \alpha_{\text{file}} \rho_{\text{file}} \omega_{\text{file}}$
ω_{file}	$\triangleq \text{NOFILE} + \text{FILE} [\text{Char}] (\text{Either} (\text{Pool FPtr}) \text{Int})$
FPtr	$\triangleq \text{Active Int} + \text{Stale}$
ν_{file}	$\triangleq \text{RChar Char} + \text{RHndP HndP} + \text{RNull} + \dots$
FHnd	$\triangleq \text{RWHnd} + \text{RHnd HndP}$
α_{file}	$\triangleq \text{FDelete} + \text{FWrite Char} + \text{FwrOpen} \dots$
	$ \text{HGetC FHnd} + \text{HClose FHnd} + \text{HIsEOF FHnd} \dots$
	$ \text{FNextRdHnd} + \text{HRdOpen HndP} + \text{FDoesExist} \dots$
ρ_{file}	$\triangleq ([\text{FHnd}], [\text{FHnd}])$

Figure 9. Fragments of model file

`hGetContents`. Extensions of the model to cater for these features is an obvious candidate for future work.

7.1 A file I/O model

Figure 9 contains the outline of a model of a single file. A file either doesn't exist, `NOFILE`, or it exists and contains both the file content (a sequence of characters) and either a single write pointer or a pool of read pointers, some of which may be closed (`Stale`).

The primary difficulty when modelling a file is the dynamic allocation of read handles. For this reason, many aspects of file borrow directly from the `dmap` combinator definition. The pool is necessary because we want the opening of a file for reading to be order independent. When splitting a context one must supply the specific handles which each side requires access to (by giving two lists of `FHnd`, one for each process). One may use this to request complete access to the file or access to individual read handles, and like with `dmap`, we must ensure that a process by default has access to all the potential read handles it may need, if it has shared access to a file. The procedure for allocating a new read handle is likewise very similar to the two-step allocation process of `dmap`.

The concurrent allocation and deallocation of read and write pointers required a few subtle design decisions. All deallocated pointers, `Stale`, must be left within the pool because reusing old read handles could cause nondeterminism. For similar reasons, there is no explicit representation of a closed file: a file is closed if it only contains inactive read pointers. We found it useful to separate file actions into three categories: (i) actions, such as writing, which affect the entire file and must be explicitly sequenced; (ii) actions, such as reading a character, which only affect a single read pointer and may examine file contents but not change them; (iii) anomalous actions which must be treated separately, such as the two actions required to allocate a new read handle.

7.2 Models term and chan

Our model of terminal I/O, `term`, is a basic one (see Figure 10). Each character outputted to `stdout` by the program using `PutCh` gives rise instantaneously to one or more characters ready to be read from `stdin`. These are buffered, and will be consumed by subsequent calls to `GetCh`.

Model `term` has the exact same opportunities for deterministic concurrency as `chan/bfft`. In fact, `term` is a complete generalisation of `chan`. If the world state of `term` is $([], \text{tchan})$,

$$\text{tchan} \triangleq T (\lambda c. ([c], \text{tchan}))$$

this models a communication channel by giving the “semantics” of a user who re-inputs every character outputted to him or her.

TermIO	$\triangleq T (\text{Char} \rightarrow ([\text{Char}], \text{TermIO}))$
ω_{term}	$\triangleq ([\text{Char}], \text{TermIO})$
α_{term}	$\triangleq \text{GetCh} + (\text{PutCh Char})$
ν_{term}	$\triangleq \text{Char}$
$\text{st}_{\text{term}} \text{ GetCh}$	$([], t) = \$$
$\text{st}_{\text{term}} \text{ GetCh}$	$((c:cs), t) = ((cs, t'), c)$
$\text{st}_{\text{term}} (\text{PutCh } c)$	$(cs, T f) = ((cs + cs', t'), 'X')$
where	$(cs', t') = f c$

Figure 10. Partial definition of term

We omit the definition of `chan`. It is mostly identical to `bfft`, with the exception that characters are sent, not integers².

8. Related Work

CURIO's main contributions are in the field of deterministic concurrency and the problems with giving a semantics to/reasoning about I/O (or global state), especially in the presence of concurrency.

8.1 Deterministic concurrency

The Clean language uses uniqueness types [3, 1] to allow actions to be sequenced with respect to individual files. But, unlike CURIO, there is no way for two separate program fragments performing I/O on separate pieces of world state to communicate with one another—all communication must be performed at a synchronisation point at which the two fragments of code become one. Furthermore, Clean does not give a semantics to I/O as such. The I/O interface is constructed in an ad hoc style based on Clean's graph rewriting semantics [26].

The Brisk Haskell Compiler uses rank-2 polymorphism to guarantee deterministic access to multiple global sub-states [17], including individual files. In her Ph.D. thesis [28], Eleni Spiliopoulou describes how concurrent threads may communicate deterministically. However, the language description is largely implementation-driven, with no semantics. Shared file reads are performed using lazy file I/O with an informal argument justifying why this is deterministic. Although lazy file I/O is not mentioned in CURIO, we can still describe shared file reads in a rigorous fashion.

Terauchi and Aiken [29] give a means of structuring side effects under lazy evaluation via “witnesses.” An ordering is enforced on side effects by making one “see” a witness of the other. The downside is that confluence is only statically decidable for some reduction strategies (call-by-need but not call-by-name, for example), and it has only been applied to simple read and write cells, not full I/O. Programs in CURIO are always confluent regardless of reduction strategy, and can cater for a variety of I/O mechanisms.

Type and effect systems [21] are a standard approach to merging imperative and functional features, where pure and side effecting terms are distinguished by the type system. However, these rely on an *implicit* program evaluation strategy and thus are not directly applicable. The language Vault [8] makes use of region-based [31, 6] type and effect systems. Vault is a modification of C and has been used to provide a more secure interface to Windows 2000 device drivers. Although its type system is sensitive to specific aspects of the I/O API, it does not go quite so far as to give a semantics to I/O. Boyland [4] gives a statically decidable means of splitting permissions on mutable state, but this work is focussed on imperative

²One possible extension would be to allow the sender process to close the buffer. At a high level, this would allow a writer to perform `hClose` on the channel, letting the reader determine when communication has ended.

languages with implicit reduction orders, and the permissions do not seem to be general enough to handle communication channels.

8.2 Semantics for I/O

There are a number of elegant techniques for expressing state manipulation in functional languages which also happen to be applicable to I/O. This is true of Clean's uniqueness types, and Brisk uses a modification of lazy functional state threads [19] to structure global state. Composable Memory Transactions [16] use a form of memory transaction to permit concurrent data structures, and may be applicable to certain aspects of everyday I/O. These ad hoc techniques are useful and elegant, but we would argue that CURIO offers a more direct description of I/O and APIs.

Haskell's I/O semantics is largely derived from the work in Gordon's Ph.D. thesis [12]. In this document he develops a full operational semantics for a pure functional language using bisimulation, and extends this notion of bisimulation so that transitions, or observations, may include I/O actions. All actions are opaque, observable events and two different sequences of actions are always distinguishable. A similar approach was used to give the semantics to nondeterministic Concurrent Haskell [24] and, later, to describe Haskell's I/O, exception and foreign interface mechanisms [23]. This technique gives an elegant model of user interaction and non-termination. The main flaw to this bisimulation approach is its usefulness since it "explicitly represents the instructions issued by a program, rather than their observable effect" [24]. This style of semantics was used to prove the monad laws in [13]. However, even in Gordon's Ph.D. thesis, in order to prove useful properties about the *effects* of I/O actions, such as the order independence we ourselves want, a state transformer semantics had to be used.

There have been many attempts at reasoning about purely sequential I/O. The House functional operating system [15] allows small propositions concerning sequential monadic code to be proved correct using Programatica. Butterfield and Strong [5] performed a case study comparing the ease of formal reasoning about I/O in C, Haskell and Clean using a semantic model of the file system. This was then followed with a larger proof of a simplified version of the UNIX make utility [11]. Thompson [30] gives a trace-based semantics to lazy stream I/O in Miranda. Hall and Hammond's draft dynamic semantics for Haskell 1.3 [14] describes the effect of I/O actions with respect to a single system state, including a semantics for file system actions. Yet CURIO is more powerful than these in that it handles concurrency.

Sewell [27] addresses I/O in the concurrent language Pict. He develops a simplified notion of a sequential C program and the UNIX X-Windows request buffer, and proves an abstract machine correct with respect to the semantics. This is based on the experience of implementing a real language but it is mostly the same as the bisimulation approach, and does not appear to easily generalise.

9. Future Work and Conclusions

We defined the CURIO core language, gave a proof of a general confluence condition, and showed how it is powerful enough to describe the semantics of Haskell I/O and a basic language extension for deterministic concurrency and inter-process communication. The most noteworthy feature is perhaps our use of a state transformer semantics to model concurrency and communication.

Future work will include an investigation into whether the runtime checks required by CURIO could instead be performed statically, by proving a type system sound with respect to CURIO's semantics. Ideally this would make `test` unnecessary. I/O contexts are quite type-like, since a program's enclosing context does not change over time. Another limitation to our system is that communication channels currently only permit the sending of characters, and fully polymorphic channels would be far more desirable. Fur-

thermore we also plan to use CURIO as the basis for actual proofs about the behaviour of I/O, and develop general notions of program equivalence for CURIO programs. Some progress in this area has already been reported in [9]. Extending the `io` model to cover more of the Haskell API, and the APIs of other languages would also be worthwhile.

References

- [1] Peter Achten and Rinus Plasmeijer. The Ins and Outs of Clean I/O. *Journal of Functional Programming*, 5(1):81–110, January 1995.
- [2] Arvind, Rishiyur S. Nikhil, and Keshav K. Pingali. I-structures: Data Structures for Parallel Computing. *ACM Transactions on Programming Language Systems*, 11(4):598–632, 1989.
- [3] Erik Barendsen and Sjaak Smetsers. Uniqueness Typing for Functional Languages with Graph Rewriting Semantics. *Mathematical Structures in Computer Science*, 6(6):579–612, 1996.
- [4] John Boyland. Checking Interference with Fractional Permissions. In Radhia Cousot, editor, *SAS*, volume LNCS2694, pages 55–72. Springer, 2003.
- [5] Andrew Butterfield and Glenn Strong. Proving Correctness of Programs with I/O — a Paradigm Comparison. In Thomas Arts and Markus Mohnen, editors, *Proceedings of IFL2001*, volume LNCS2312, pages 72–87, 2001.
- [6] Karl Crary, David Walker, and Greg Morrisett. Typed Memory Management in a Calculus of Capabilities. In *POPL '99: Proceedings of the 26th Symposium on Principles of Programming Languages*, pages 262–275, New York, NY, USA, 1999. ACM Press.
- [7] Maarten de Mol, Marko van Eekelen, and Rinus Plasmeijer. Theorem Proving for Functional Programmers. In Thomas Arts and Markus Mohnen, editors, *Proceedings of IFL2001*, volume LNCS2312, pages 55–71. Springer-Verlag, 2001.
- [8] Robert DeLine and Manuel Fähndrich. Enforcing High-Level Protocols in Low-Level Software. In *PLDI '01: Proceedings of the 2001 Conference on Programming Language Design and Implementation*, pages 59–69, New York, NY, USA, 2001. ACM Press.
- [9] Malcolm Dowse. *A Semantic Framework for Deterministic Functional Input/Output*. PhD thesis, Department of Computer Science, University of Dublin, 2006.
- [10] Malcolm Dowse, Andrew Butterfield, and Marko van Eekelen. Reasoning about Deterministic Concurrent Functional I/O. In Clemens Grellck and Frank Huch, editors, *Proceedings of IFL 2004*, volume LNCS3474, pages 177–194. Springer-Verlag, 2005.
- [11] Malcolm Dowse, Glenn Strong, and Andrew Butterfield. Proving 'make' Correct — I/O Proofs in Haskell and Clean. In Ricardo Peña and Thomas Arts, editors, *Proceedings of IFL 2002*, volume LNCS2670, pages 68–83, 2002.
- [12] Andrew Gordon. *Functional Programming and Input/Output*. Distinguished Dissertations in Computer Science. Cambridge University Press, 1994.
- [13] Andrew Gordon. Bisimilarity as a Theory of Functional Programming: Mini-course. Notes Series BRICS-NS-95-3, BRICS, Department of Computer Science, University of Aarhus, July 1995.
- [14] Cordelia Hall and Kevin Hammond. A Dynamic Semantics for Haskell (draft), May 20 1993.
- [15] Thomas Hallgren, Mark Jones, Rebekah Leslie, and Andrew Tolmach. A Principled Approach to Operating System Construction in Haskell. In *ICFP '05: Proceedings of the 10th International Conference on Functional Programming*, pages 116–128, New York, NY, USA, 2005. ACM Press.
- [16] Tim Harris, Simon Marlow, Simon Peyton Jones, and Maurice Herlihy. Composable Memory Transactions. In *Proceedings of the 10th Symposium on Principles and Practice of Parallel Programming*, pages 48–60, New York, NY, USA, 2005. ACM Press.

- [17] Ian Holyer and Eleni Spiliopoulou. Concurrent Monadic Interfacing. In Kevin Hammond, Antony Davie, and Chris Clack, editors, *Proceedings of IFL '98*, volume LNCS1595, pages 73–89. Springer-Verlag, 1999.
- [18] Mark Jones and Paul Hudak. Implicit and Explicit Parallel Programming in Haskell. Technical Report YALEU/DCS/RR-982, Department of Computer Science, Yale University, August 1993.
- [19] John Launchbury and Simon Peyton Jones. Lazy Functional State Threads. In *PLDI '94: Proceedings of the 1994 Conference on Programming Language Design and Implementation*, pages 24–35, Orlando, Florida, June 20–24, 1994.
- [20] Robin Milner, Mads Tofte, and David MacQueen. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1997.
- [21] Flemming Nielson and Hanne Riis Nielson. Type and Effect Systems. In *Correct System Design, Recent Insight and Advances*, pages 114–136, London, UK, 1999. Springer-Verlag.
- [22] Simon Peyton Jones. Wearing the Hair Shirt: A retrospective on Haskell. Invited talk at POPL 2003.
- [23] Simon Peyton Jones. Tackling the Awkward Squad — monadic input/output, concurrency, exceptions, and foreign language calls in Haskell. In CAR Hoare, M Broy, and R Stein-brueggen, editors, *Engineering theories of software construction, Marktoberdorf Summer School 2000*, pages 47–96. IOS Press, 2001.
- [24] Simon Peyton Jones, Andrew Gordon, and Sigbjorn Finne. Concurrent Haskell. In *POPL '96: Proceedings of the 23rd Symposium on Principles of Programming Languages*, pages 295–308, New York, NY, USA, 1996. ACM Press.
- [25] Simon Peyton Jones and Philip Wadler. Imperative Functional Programming. In *POPL '93: Proceedings of the 20th Symposium on Principles of Programming Languages*, pages 71–84, New York, NY, USA, 1993. ACM Press.
- [26] Rinus Plasmeijer and Marko van Eekelen. *Functional Programming and Parallel Graph Rewriting*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1993.
- [27] Peter Sewell. On Implementations and Semantics of a Concurrent Programming Language. In *CONCUR '97: Proceedings of the 8th International Conference on Concurrency Theory*, pages 391–405, London, UK, 1997. Springer-Verlag.
- [28] Eleni Spiliopoulou. *Concurrent and Distributed Functional Systems*. PhD thesis, University of Bristol, Department of Computing Science, August 1999.
- [29] Tachio Terauchi and Alex Aiken. Witnessing Side-Effects. In *ICFP '05: Proceedings of the 10th International Conference on Functional Programming*, pages 105–115, New York, NY, USA, 2005. ACM Press.
- [30] Simon Thompson. Interactive Functional Programs: a method and a formal semantics. Technical Report 48*, Computing Laboratory, University of Kent., November 1987.
- [31] Mads Tofte and Jean-Pierre Talpin. Region-Based Memory Management. *Information and Computation*, 132(2):109–176, 1997.
- [32] David Turner. An overview of Miranda. *ACM SIGPLAN Notices*, 21(12):158–166, December 1986.
- [33] Philip Wadler. The Essence of Functional Programming. In *POPL '92: Proceedings of the 19th Symposium on Principles of Programming Languages*, pages 1–14, New York, NY, USA, 1992. ACM Press.

A. Additional Material

The proofs were machine verified using the LCF-based Sparkle [7] proof-assistant, and using Core-Clean as a metalanguage. We use Haskell syntax instead as it will be more familiar to the average reader. Full details concerning these proofs, including all results in this paper, may be found in [9].

A.1 Meta-encoding

I/O models are defined as follows.

```
data IOModel v a p w c = IOModel {
  af :: a -> w -> (w,v), wa :: a -> w -> Bool,
  ap :: c -> a -> Bool, pf :: p -> c -> (c,c)}
```

The functions `af` and `wa` form `st`, the type `c` denotes contexts, `ap` gives the actions each context permits, and `pf` is `pf`. `CURIO`'s five primitives are encoded with the following algebraic type:

```
data Prog v a p = Ret v | Action a
  | Bind (Prog v a p) (v -> Prog v a p)
  | Test a (Prog v a p) (Prog v a p)
  | Par p (Prog v a p) (Prog v a p) (v -> v -> v)
```

Small-step reduction is implemented with the function `next`, and `run` continually applies `next`, only terminating when a normal form is reached:

```
next :: IOModel v a p w c -> c -> Guess ->
  (w,Prog v a p) -> Maybe (w,Prog v a p)
run :: IOModel v a p w c -> c -> [Guess] ->
  (w,Prog v a p) -> (w,Prog v a p)
```

The `Guess` parameter is used to implement nondeterminism, and the reduction relations are defined in terms of these functions:

$$\begin{aligned}
 w \Vdash m \downarrow^c &\iff \exists_{g \in \text{Guess}}. \text{next}_s \, g \, c \, (w, m) = \text{Nothing} \\
 w \Vdash m \uparrow^c &\iff \exists_{g \in \text{Guess}}. \text{next}_s \, g \, c \, (w, m) = \perp \\
 w \Vdash m \uparrow^c &\iff \forall_{g \in [\text{Guess}]} . \text{run}_s \, g \, c \, (w, m) = \perp \\
 w \Vdash m \xrightarrow{c} w_1 \Vdash m_1 &\iff \exists_{g \in \text{Guess}}. \text{next}_s \, g \, c \, (w, m) = \text{Just} \, (w_1, m_1) \\
 w \Vdash m \downarrow^c w_1 \Vdash m_1 &\iff \exists_{g \in [\text{Guess}]} . \text{run}_s \, g \, c \, (w, m) = (w_1, m_1) \\
 w \Vdash m \downarrow^c w_1 \Vdash m_1 &\iff \forall_{g \in [\text{Guess}]} . \text{run}_s \, g \, c \, (w, m) = (w_1, m_1)
 \end{aligned}$$

A.2 Encoding location-based models

The trick to the encoding is being able to convert $(\mathbb{L}, \mathcal{P}\alpha)$ to and from the context $\mathcal{P}(\mathbb{L}, \text{Act } \alpha + \text{Probe})$. This encoding and decoding is as follows.

$$\begin{aligned}
 \text{encode} &:: (\mathbb{L}, \mathcal{P}\alpha) \rightarrow \mathcal{P}(\mathbb{L}, \text{Act } \alpha + \text{Probe}) \\
 \text{encode} \, (l, C) &\triangleq \{(l_1, \text{Act } a) \mid l \preceq l_1, a \in C\} \\
 &\quad \cup \{(l_1, \text{Probe}) \mid l \preceq l_1\} \\
 \text{decode} &:: \mathcal{P}(\mathbb{L}, \text{Act } \alpha + \text{Probe}) \rightarrow (\mathbb{L}, \mathcal{P}\alpha) \\
 \text{decode} \, C' &\triangleq \left(\bigcap \{l \mid (l, \text{Probe}) \in C'\}, \right. \\
 &\quad \left. \{a \mid (l, \text{Act } a) \in C'\} \right)
 \end{aligned}$$

$\bigcap L$, where L is a set of locations, returns the largest list $l \in \mathbb{L}$ such that, for all $l_1 \in L$, $l \preceq l_1$, or, if $L = \emptyset$, it returns $[\]$.

The proof that PRE'_s implies $\text{PRE}_{\text{toM } s}$ relies on the correctness of the above encoding and decoding.

The definition of `toM` is

$$\begin{aligned}
 \text{toM} \, (st', pf') &\triangleq (st_0, pf_0) \\
 \text{where} & \\
 st_0 \, (l, \text{Act } a) \, w &= st' \, (l, a) \, w \\
 st_0 \, (l, \text{Probe}) \, w &= \perp \\
 pf_0 \, p \, C' &= (C' \cap C'_l, C' \cap C'_r) \\
 \text{where} & \\
 (l, C) &= \text{decode} \, C' \\
 (C'_l, C'_r) &= pf' \, (l, p) \, C \\
 C'_l &= \text{encode} \, (l \# [L], C_l) \\
 C'_r &= \text{encode} \, (l \# [R], C_r)
 \end{aligned}$$