

Shared Data

Consider the **Readers and Writers** problem:

Problem

We have a **large data structure** (i.e. a structure that cannot be updated in one atomic step) that is shared between some number of writers who are updating the data structure and some number of readers who are attempting to retrieve a coherent copy of the data structure.

Desiderata:

- We want **atomicity**, in that each update happens in one go, and updates-in-progress or partial updates are not observable.
- We want **consistency**, in that any reader that starts after an update finishes will see that update.
- We want to minimise **waiting**.

A Crappy Solution

Treat both reads and updates as critical sections — use any old critical section solution (locks, etc.) to sequentialise all reads and writes to the data structure.

Observation

Updates are *atomic* and reads are *consistent* — but reads can't happen concurrently, which leads to unnecessary *contention*.

A Crappy Solution

Treat both reads and updates as critical sections — use any old critical section solution (locks, etc.) to sequentialise all reads and writes to the data structure.

Observation

Updates are *atomic* and reads are *consistent* — but reads can't happen concurrently, which leads to unnecessary *contention*.

A Better Solution

A more elaborate locking mechanism (*condition variables*) could be used to allow multiple readers to read concurrently, but writers are still executed individually and atomically.

Observation

We have atomicity and consistency, and now multiple reads can execute concurrently. Still, we don't allow updates to execute concurrently with reads, to prevent partial updates from being observed by a reader.

A Better Solution

A more elaborate locking mechanism (*condition variables*) could be used to allow multiple readers to read concurrently, but writers are still executed individually and atomically.

Observation

We have atomicity and consistency, and now multiple reads can execute concurrently. Still, we don't allow updates to execute concurrently with reads, to prevent partial updates from being observed by a reader.

Reading and Writing

Complication

Now suppose we don't want readers to wait *(much)* while an update is performed. Instead, we'd rather they get an *older version* of the data structure.

Trick: Rather than update the data structure in place, a writer creates *their own local copy* of the data structure, and then merely updates the (shared) *pointer* to the data structure to point to their copy.

Liam: Draw on the board


Atomicity The only shared write is now just to one pointer.

Consistency Reads that start before the pointer update get the older version, but reads that start after get the latest.

Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.


Example (Binary Search Tree)



Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.


Example (Binary Search Tree)



Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.


Example (Binary Search Tree)



Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.

Example (Binary Search Tree)



Purely Functional Data Structures

Persistent data structures that exclusively make use of copying over mutation are called *purely functional* data structures. They are so called because operations on them are best expressed in the form of mathematical functions that, given an input structure, return a *new* output structure:

$$\begin{aligned}
 \textit{insert } v \text{ Leaf} &= \text{Branch } v \text{ Leaf Leaf} \\
 \textit{insert } v (\text{Branch } x \text{ } l \text{ } r) &= \mathbf{if } v \leq x \mathbf{ then} \\
 &\quad \text{Branch } x (\textit{insert } v \text{ } l) \text{ } r \\
 &\mathbf{else} \\
 &\quad \text{Branch } x \text{ } l (\textit{insert } v \text{ } r)
 \end{aligned}$$

Computing with Functions

We model real processes in Haskell using the IO type. We have previously seen a definition of IO that resembles the state monad:

```
type IO a = (🌐 → (🌐, a))
```

But this is not a good model where concurrency is concerned.

Why?

Instead, we'll just treat IO as an abstract type for now, and give it a *formal semantics* later:

IO τ = A (possibly effectful) process that, when *executed*, produces a result of type τ

Note the semantics of *evaluation* and *execution* are different things.

Computing with Functions

We model real processes in Haskell using the `IO` type. We have previously seen a definition of `IO` that resembles the state monad:

```
type IO a = (State → (State, a))
```

But this is not a good model where concurrency is concerned.

Why?

Instead, we'll just treat `IO` as an abstract type for now, and give it a **formal semantics** later:

`IO τ` = A (possibly effectful) process that, when **executed**, produces a result of type `τ`

Note the semantics of **evaluation** and **execution** are different things.

Building up IO

Recall **monads**:

$$\text{return} :: \forall a. a \rightarrow \text{IO } a$$

$$(\gg) :: \forall a b. \text{IO } a \rightarrow (a \rightarrow \text{IO } b) \rightarrow \text{IO } b$$

$$\text{getChar} :: \text{IO Char}$$

$$\text{putChar} :: \text{Char} \rightarrow \text{IO } ()$$

Example (Echo)

$$\text{echo} :: \text{IO } ()$$

$$\text{echo} = \text{getChar} \gg (\lambda x. \text{putChar } x \gg \lambda y. \text{echo})$$

Or, with **do** notation:

$$\text{echo} :: \text{IO } ()$$

$$\text{echo} = \mathbf{do} \quad x \leftarrow \text{getChar}$$

$$\quad \text{putChar } x$$

$$\quad \text{echo}$$

Adding Concurrency

We can have multiple threads easily enough:

$$\text{forkIO} :: \text{IO } () \rightarrow \text{IO } ()$$

Example (Dueling Printers)

```
let loop c = do putChar c; loop c
in do forkIO (loop 'a'); loop 'z'
```

But what sort of *synchronisation primitives* are available?

MVars

The **MVar** is the simplest synchronisation primitive in Haskell. It can be thought of as a shared box which holds at most one value.

Processes must take the value out of a **full** box to read it, and must put a value into an **empty** box to update it.

MVar Functions

$newMVar :: \forall a. a \rightarrow IO (MVar a)$	Create a new MVar
$takeMVar :: \forall a. MVar a \rightarrow IO a$	Read/remove the value
$putMVar :: \forall a. MVar a \rightarrow a \rightarrow IO ()$	Update/insert a value

Taking from an empty MVar or putting into a full one results in blocking.

An MVar can be thought of as channel containing at most one value.

Readers and Writers

We can treat MVars as shared variables with some definitions:

writeMVar m v = do takeMVar m; putMVar m v

readMVar m = do v ← takeMVar m; putMVar m v; return v

problem :: DB → IO ()

problem initial = do

db ← newMVar initial

wl ← newMVar ()

let *reader = readMVar db* $\gg\gg$ \dots

let *writer = do*

takeMVar wl

d ← readMVar db

let *d' = update d*

evaluate d'

writeMVar db d'

putMVar wl ()

Fairness

Each MVar has an attached FIFO queue, so GHC Haskell can ensure the following fairness property:

No thread can be blocked indefinitely on an MVar unless another thread holds that MVar indefinitely.

Evaluation Semantics

The semantics of Haskell's evaluation are interesting but not particularly relevant for us. We will assume that it happens quietly without a fuss:

$$\beta\text{-equivalence} \quad (\lambda x. M[x]) N \equiv_{\beta} M[N]$$

$$\alpha\text{-equivalence} \quad \lambda x. M[x] \equiv_{\alpha} \lambda y. M[y]$$

$$\eta\text{-equivalence} \quad \lambda x. M x \equiv_{\eta} M$$

Let our ambient congruence relation \equiv be $\equiv_{\alpha\beta\eta}$ enriched with the following extra equations, justified by the *monad laws*:

$$\text{return } N \gg= M \equiv M N$$

$$(X \gg= Y) \gg= Z \equiv X \gg= (\lambda x. Y x \gg= Z)$$

$$X \equiv X \gg= \text{return}$$

Processes

This means that a Haskell expression of type $\text{IO } \tau$ will boil down to either $\text{return } x$ where x is a value of type τ ; or $a \gg= M$ where a is some *primitive IO action* ($\text{forkIO } p$, $\text{readMVar } v$, etc.) and M is some function producing another $\text{IO } \tau$. This is the *head normal form* for IO expressions.

Definition

Define a language of *processes* P , which contains all (head-normal) expressions of type $\text{IO } ()$.

We want to define the semantics of the *execution* of these processes. Let's use *operational semantics*:

$$(\mapsto) \subseteq P \times P$$

Semantics for forkIO

To model *forkIO*, we need to model the parallel execution of multiple processes in our process language. We shall add a *parallel composition* operator to the language of processes:

$$\begin{array}{l}
 P, Q ::= a \gg M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \dots
 \end{array}$$

And the following ambient congruence equations:

$$\begin{array}{l}
 P \parallel Q \equiv Q \parallel P \\
 P \parallel (Q \parallel R) \equiv (P \parallel Q) \parallel R
 \end{array}$$

Semantics for forkIO

To model *forkIO*, we need to model the parallel execution of multiple processes in our process language. We shall add a *parallel composition* operator to the language of processes:

$$\begin{array}{l}
 P, Q ::= a \gg M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \dots
 \end{array}$$

And the following ambient congruence equations:

$$\begin{array}{l}
 P \parallel Q \equiv Q \parallel P \\
 P \parallel (Q \parallel R) \equiv (P \parallel Q) \parallel R
 \end{array}$$

Semantics for forkIO

If we have multiple processes active, pick one of them non-deterministically to move:

$$\frac{P \mapsto P'}{P \parallel Q \mapsto P' \parallel Q}$$

The *forkIO* operation introduces a new process:

$$(forkIO P \gg M) \mapsto P \parallel (return () \gg M)$$

Semantics for forkIO

If we have multiple processes active, pick one of them non-deterministically to move:

$$\frac{P \mapsto P'}{P \parallel Q \mapsto P' \parallel Q}$$

The *forkIO* operation introduces a new process:

$$(\text{forkIO } P \gg\! = M) \mapsto P \parallel (\text{return } () \gg\! = M)$$

Semantics for MVars

MVars are modelled as a special type of *process*, identified by a *unique name*. Values of `MVar` type merely contain the name of the process, so that `putMVar` and friends know where to look.

$$\begin{array}{l}
 P, Q ::= a \gg\equiv M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \langle n \rangle \quad | \langle v \rangle_n \\
 \quad | \dots
 \end{array}$$

$$\langle n \rangle \parallel (\text{putMVar } n \ v \gg\equiv M) \mapsto \langle v \rangle_n \parallel (\text{return } () \gg\equiv M)$$

$$\langle v \rangle_n \parallel (\text{takeMVar } n \gg\equiv M) \mapsto \langle n \rangle \parallel (\text{return } v \gg\equiv M)$$

Semantics for MVars

MVars are modelled as a special type of *process*, identified by a *unique name*. Values of `MVar` type merely contain the name of the process, so that `putMVar` and friends know where to look.

$$\begin{array}{l}
 P, Q ::= a \gg\equiv M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \langle n \rangle \quad | \quad \langle v \rangle_n \\
 \quad | \dots
 \end{array}$$

$$\langle n \rangle \parallel (\text{putMVar } n \ v \gg\equiv M) \quad \mapsto \quad \langle v \rangle_n \parallel (\text{return } () \gg\equiv M)$$

$$\langle v \rangle_n \parallel (\text{takeMVar } n \gg\equiv M) \quad \mapsto \quad \langle n \rangle \parallel (\text{return } v \gg\equiv M)$$

Semantics for MVars

MVars are modelled as a special type of *process*, identified by a *unique name*. Values of `MVar` type merely contain the name of the process, so that `putMVar` and friends know where to look.

$$\begin{array}{l}
 P, Q ::= a \gg\equiv M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \langle n \rangle \quad | \langle v \rangle_n \\
 \quad | \dots
 \end{array}$$

$$\langle n \rangle \parallel (\text{putMVar } n \ v \gg\equiv M) \quad \mapsto \quad \langle v \rangle_n \parallel (\text{return } () \gg\equiv M)$$

$$\langle v \rangle_n \parallel (\text{takeMVar } n \gg\equiv M) \quad \mapsto \quad \langle n \rangle \parallel (\text{return } v \gg\equiv M)$$

Semantics for newMVar

We might think that *newMVar* should have semantics like this:

$$\frac{}{(newMVar\ v \gg\equiv M) \mapsto \langle v \rangle_n \parallel (return\ n \gg\equiv M)}^{(n\ \text{fresh})}$$

But this approach has a number of problems:

- The name *n* is now globally-scoped, without an explicit binder to introduce it.
- It doesn't accurately model the *lifetime* of the MVar, which should be garbage-collected once all processes that can access it finish.
- It makes MVars *global* objects, so our semantics aren't very *abstract*. We would like local communication to be local in our model.

Restriction Operator

We introduce a *restriction operator* ν to our language of processes:

$$\begin{array}{l}
 P, Q ::= a \gg M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \langle \rangle_n \mid \langle v \rangle_n \\
 \quad | (\nu n) P
 \end{array}$$

Writing $(\nu n) P$ says that the MVar name n is *only* available in process P . Mentioning n outside P is not well-formed. We need the following additional congruence equations:

$$\begin{array}{l}
 (\nu n) (\nu m) P \equiv (\nu m) (\nu n) P \\
 (\nu n)(P \parallel Q) \equiv P \parallel (\nu n) Q \quad (\text{if } n \notin P)
 \end{array}$$

Restriction Operator

We introduce a *restriction operator* ν to our language of processes:

$$\begin{array}{l}
 P, Q ::= a \gg M \\
 \quad | \text{return } () \\
 \quad | P \parallel Q \\
 \quad | \langle \rangle_n \quad | \langle v \rangle_n \\
 \quad | (\nu n) P
 \end{array}$$

Writing $(\nu n) P$ says that the MVar name n is *only* available in process P . Mentioning n outside P is not well-formed. We need the following additional congruence equations:

$$\begin{array}{l}
 (\nu n) (\nu m) P \equiv (\nu m) (\nu n) P \\
 (\nu n)(P \parallel Q) \equiv P \parallel (\nu n) Q \quad (\text{if } n \notin P)
 \end{array}$$

Better Semantics for newMVar

The rule for *newMVar* is much the same as before, but now we explicitly restrict the MVar to *M*.

$$\frac{}{(newMVar\ v \gg= M) \mapsto (\nu\ n)(\langle v \rangle_n \parallel (return\ n \gg= M))} \text{ (} n \text{ fresh)}$$

We can always execute under a restriction:

$$\frac{P \mapsto P'}{(\nu\ n)\ P \mapsto (\nu\ n)\ P'}$$

Question

What happens when you put an MVar inside another MVar?

Better Semantics for newMVar

The rule for *newMVar* is much the same as before, but now we explicitly restrict the MVar to *M*.

$$\frac{}{(newMVar\ v \gg\equiv M) \mapsto (\nu\ n)(\langle v \rangle_n \parallel (return\ n \gg\equiv M))} \text{ (} n \text{ fresh)}$$

We can always execute under a restriction:

$$\frac{P \mapsto P'}{(\nu\ n)\ P \mapsto (\nu\ n)\ P'}$$

Question

What happens when you put an MVar inside another MVar?

Better Semantics for newMVar

The rule for *newMVar* is much the same as before, but now we explicitly restrict the MVar to *M*.

$$\frac{}{(newMVar\ v \gg\equiv M) \mapsto (\nu\ n)(\langle v \rangle_n \parallel (return\ n \gg\equiv M))} \quad (n\ \text{fresh})$$

We can always execute under a restriction:

$$\frac{P \mapsto P'}{(\nu\ n)\ P \mapsto (\nu\ n)\ P'}$$

Question

What happens when you put an MVar inside another MVar?

Garbage Collection

If an MVar is no longer used, we just replace it with the do-nothing process:

$$\begin{aligned} (\nu n) \langle \rangle_n &\mapsto \text{return } () \\ (\nu n) \langle v \rangle_n &\mapsto \text{return } () \end{aligned}$$

Extra processes that have outlived their usefulness disappear:

$$\text{return } () \parallel P \mapsto P$$

Garbage Collection

If an MVar is no longer used, we just replace it with the do-nothing process:

$$\begin{aligned} (\nu n) \langle \rangle_n &\mapsto \text{return } () \\ (\nu n) \langle v \rangle_n &\mapsto \text{return } () \end{aligned}$$

Extra processes that have outlived their usefulness disappear:

$$\text{return } () \parallel P \mapsto P$$

Process Algebra

Our language P is called a *process algebra*, a common means of describing semantics for concurrent programs.

Process algebras and calculi of various kinds are covered in great detail in **COMP6752** with Rob van Glabbeek, who is an expert in this field.

The last topic of this course is **Software Transactional Memory**, which we will cover in Week 12.

If there's time!

We can talk about more concurrency topics.

Bibliography



Simon Marlow

Parallel and Concurrent Programming in Haskell

O'Reilly, 2013

<http://chimera.labs.oreilly.com/books/1230000000929>



Simon Peyton Jones, Andrew Gordon and Sigbjorn Finne

Concurrent Haskell

POPL '96

Association for Computer Machinery

<http://microsoft.com/en-us/research/wp-content/uploads/1996/01/concurrent-haskell.pdf>



Simon Marlow (Editor)

Haskell 2010 Language Report

<https://www.haskell.org/onlinereport/haskell2010/>