

Asynchronous Liquid Separation Types

Johannes Kloos Rupak Majumdar Viktor Vafeiadis

Max Planck Institute for Software Systems

ECOOP 2015

Introducing asynchronous concurrency

Example of asynchronous execution:

Task A


Introducing asynchronous concurrency

Example of asynchronous execution:

Task A —————

Introducing asynchronous concurrency

Example of asynchronous execution:



Introducing asynchronous concurrency

Example of asynchronous execution:

Task B

.....


Post B

Task A




Introducing asynchronous concurrency

Example of asynchronous execution:




Introducing asynchronous concurrency

Example of asynchronous execution:




Introducing asynchronous concurrency

Example of asynchronous execution:




Introducing asynchronous concurrency

Example of asynchronous execution:




Introducing asynchronous concurrency

Example of asynchronous execution:




Introducing asynchronous concurrency

Example of asynchronous execution:



Introducing asynchronous concurrency

Example of asynchronous execution:



Used in:

- Web applications, AJAX
- Android programming model
- Operation systems: Interrupt handling

Asynchronous concurrency in OCaml

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
  copy ins outs
```

Asynchronous concurrency in OCaml

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
  copy ins outs
```

- 1 Create task to read from input stream.

Asynchronous concurrency in OCaml

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
  copy ins outs
```

- 1 Create task to read from input stream.
- 2 Wait for results of read task.

Asynchronous concurrency in OCaml

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
  copy ins outs
```

- 1 Create task to read from input stream.
- 2 Wait for results of read task.
- 3 Use task to write data to output stream.

Asynchronous concurrency in OCaml

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
copy ins outs
```

- 1 Create task to read from input stream.
- 2 Wait for results of read task.
- 3 Use task to write data to output stream.
- 4 Copy more data.

Asynchronous concurrency in OCaml

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
  copy ins outs
```

- 1 Create task to read from input stream.
- 2 Wait for results of read task.
- 3 Use task to write data to output stream.
- 4 Copy more data.
- 5 If end-of-file, terminate.

```
let rec copy ins outs =  
  if eof ins then () else  
  let fill = post (read ins) in  
  let buf = wait fill in  
  let drain = post (write outs buf) in  
  wait drain;  
  copy ins outs
```

How can we reason about such programs?

Challenge: Typing asynchronous code

Consider the code fragment `post` (read ins).

read: *stream* \rightarrow *buffer*

Challenge: Typing asynchronous code

Consider the code fragment `post` (read ins).

`read`: *stream* \rightarrow *buffer*

`post` (read ins): *promise buffer*

Challenge: Typing asynchronous code

Consider the code fragment `post (read ins)`.

`read: stream → buffer`

`post (read ins): promise buffer`

Introduce subtyping. Suppose `initbuffer <: buffer`.

Challenge: Typing asynchronous code

Consider the code fragment `post (read ins)`.

`read: stream → buffer`

`post (read ins): promise buffer`

Introduce subtyping. Suppose `initbuffer <: buffer`.

`read: stream → initbuffer`

Challenge: Typing asynchronous code

Consider the code fragment `post (read ins)`.

```
read: stream → buffer  
post (read ins): promise buffer
```

Introduce subtyping. Suppose `initbuffer <: buffer`.

```
read: stream → initbuffer  
post (read ins): promise initbuffer
```


Challenge: Typing asynchronous code

Consider the code fragment `post (read ins)`.

```
read: stream → buffer  
post (read ins): promise buffer
```

Introduce subtyping. Suppose `initbuffer <: buffer`.

```
read: stream → initbuffer  
post (read ins): promise initbuffer
```

In the paper: refinement typing, using liquid types

Challenge: Typing with mutable state

Another copying loop, with buffers on the heap (no concurrency):

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```

Challenge: Typing with mutable state

Another copying loop, with buffers on the heap (no concurrency):

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```

- 1 Main copying loop gets a pre-allocated buffer

Challenge: Typing with mutable state

Another copying loop, with buffers on the heap (no concurrency):

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```

- 1 Main copying loop gets a pre-allocated buffer
- 2 Read fills this buffer, write drains it.

Challenge: Typing with mutable state

Another copying loop, with buffers on the heap (no concurrency):

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```

- 1 Main copying loop gets a pre-allocated buffer
- 2 Read fills this buffer, write drains it.
- 3 The copying loop is called with a fresh buffer.

Challenge: Typing with mutable state

Another copying loop, with buffers on the heap (no concurrency):

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```

Can we make use subtyping to ensure: Only initialized buffers are written?

Subtyping and mutation don't mix well

Easier example:

```
let r = ref 0 in
assert (!r = 0); (* The content of r is 0 *)
r := 1;
assert (!r = 1) (* The content of r is 1 *)
```

Subtyping and mutation don't mix well

Easier example:

```
let r = ref 0 in
assert (!r = 0); (* The content of r is 0 *)
r := 1;
assert (!r = 1) (* The content of r is 1 *)
```

Can the type system prove that the assertions hold?

Subtyping and mutation don't mix well

Easier example:

```
let r = ref 0 in
assert (!r = 0); (* The content of r is 0 *)
r := 1;
assert (!r = 1) (* The content of r is 1 *)
```

Can the type system prove that the assertions hold?

Suppose $int_{=0} <: int$ and $int_{=1} <: int$.

Subtyping and mutation don't mix well

Easier example:

```
let r = ref 0 in
assert (!r = 0); (* The content of r is 0 *)
r := 1;
assert (!r = 1) (* The content of r is 1 *)
```

Can the type system prove that the assertions hold?

Suppose $int_{=0} <: int$ and $int_{=1} <: int$.

- $r:ref\ int$: Can't prove anything
- $r:ref\ int_{=1}$: Not true – initially, !r is 0
- $r:ref\ int_{=0}$: Not true

Subtyping and mutation don't mix well

Easier example:

```
let r = ref 0 in
assert (!r = 0); (* The content of r is 0 *)
r := 1;
assert (!r = 1) (* The content of r is 1 *)
```

Can the type system prove that the assertions hold?

Suppose $int_{=0} <: int$ and $int_{=1} <: int$.

- $r:ref\ int$: Can't prove anything
- $r:ref\ int_{=1}$: Not true – initially, !r is 0
- $r:ref\ int_{=0}$: Not true

We need some form of **strong updates!**

Program logics work well with updating state

```
{ emp }  
let r = ref 0 in  
{ r ↦ 0 }  
assert (!r = 0);  
{ r ↦ 0 }  
r := 1;  
{ r ↦ 1 }  
assert (!r = 1);  
{ r ↦ 1 }
```

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.
- **But:** Don't deal with mutable state out-of-the box.

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.
- **But:** Don't deal with mutable state out-of-the box.

Program logics:

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.
- **But:** Don't deal with mutable state out-of-the box.

Program logics:

- Are good at dealing with mutable state

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.
- **But:** Don't deal with mutable state out-of-the box.

Program logics:

- Are good at dealing with mutable state
- Also work well with concurrency

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.
- **But:** Don't deal with mutable state out-of-the box.

Program logics:

- Are good at dealing with mutable state
- Also work well with concurrency
- **But:** Automation is much harder.

The situation

Subtyping, refinement types:

- Expressive
- In the guise of liquid types: Provide strong automation.
- **But:** Don't deal with mutable state out-of-the box.

Program logics:

- Are good at dealing with mutable state
- Also work well with concurrency
- **But:** Automation is much harder.

Can we combine both?

Our Contribution

Combine **type systems** and **program logics** to reason about OCaml programs with mutable state and asynchronous concurrency.

Combine **type systems** and **program logics** to reason about OCaml programs with mutable state and asynchronous concurrency.

$$\begin{aligned} &\text{Liquid type} + \text{Concurrent Separation Logic} \\ &= \\ &\text{Asynchronous Liquid Separation Types (ALST)} \end{aligned}$$

- The type system
- Type inference
- Some larger examples

- The type system
- Type inference
- Some larger examples

Typing the copying loop with buffers

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```


Typing the copying loop with buffers

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      read ins buf;  
      write outs buf;  
      do_copy buf  
  in do_copy (make_buffer ())
```


Traditional typing judgment

Γ
↑
Environment

$\vdash e :$
↑
Expression

τ
↑
Type

Extended typing judgment



Typing the example

ins: *stream*, *buf*: *ref buffer* < >
⊢ read *ins* *buf*: *unit* < >

Typing the example

ins: *stream*, *buf*: *ref buffer* $\langle \mu \mapsto \textit{buffer} \rangle$

\vdash read *ins* *buf*: *unit* $\langle \mu \mapsto \textit{initbuffer} \rangle$

- Precondition: Content of heap cell μ has type *buffer*
Postcondition: Content of heap cell μ has type *initbuffer*

Typing the example

ins: *stream*, *buf*: $\text{ref}_\mu \text{buffer}$ $\langle \mu \mapsto \text{buffer} \rangle$

\vdash read *ins* *buf*: *unit* $\langle \mu \mapsto \text{initbuffer} \rangle$

- Precondition: Content of heap cell μ has type *buffer*
Postcondition: Content of heap cell μ has type *initbuffer*
- *buf* is a reference to μ

Typing the example

ins: *stream*, *buf*: $\text{ref}_\mu \text{buffer}$ $\langle \mu \mapsto \text{buffer} \rangle$

\vdash read *ins* *buf*: *unit* $\langle \mu \mapsto \text{initbuffer} \rangle$

- Precondition: Content of heap cell μ has type *buffer*
Postcondition: Content of heap cell μ has type *initbuffer*
- *buf* is a reference to μ
- **Note:** Memory cells have types here, subtyping allowed!

We can even handle concurrency

Consider modified code fragment:

```
post (read ins buf)
```


We can even handle concurrency

Consider modified code fragment:

```
post (read ins buf)
```

Known: Type of read ins buf

ins: *stream*, *buf*: *ref_μ buffer* $\langle \mu \mapsto \textit{buffer} \rangle$

\vdash read ins buf: *unit* $\langle \mu \mapsto \textit{initbuffer} \rangle$

We can even handle concurrency

Consider modified code fragment:

```
post (read ins buf)
```

Known: Type of read ins buf

ins: *stream*, *buf*: $ref_{\mu} \textit{buffer}$ $\langle \mu \mapsto \textit{buffer} \rangle$

\vdash read ins buf: *unit* $\langle \mu \mapsto \textit{initbuffer} \rangle$

The post expression:

ins: *stream*, *buf*: $ref_{\mu} \textit{buffer}$ $\langle \mu \mapsto \textit{buffer} \rangle$

\vdash post (read ins buf):

$\langle \quad \quad \quad \rangle$

We can even handle concurrency

Consider modified code fragment:

```
post (read ins buf)
```

Known: Type of read ins buf

$ins: stream, buf: ref_{\mu} buffer \quad \langle \mu \mapsto buffer \rangle$
 $\vdash \text{read ins buf}: \quad unit \quad \langle \mu \mapsto initbuffer \rangle$

The post expression:

$ins: stream, buf: ref_{\mu} buffer \quad \langle \mu \mapsto buffer \rangle$
 $\vdash \text{post (read ins buf)}:$
 $\quad promise_{\pi} unit \quad \langle \quad \quad \quad \rangle$

We can even handle concurrency

Consider modified code fragment:

```
post (read ins buf)
```

Known: Type of read ins buf

$ins: stream, buf: ref_{\mu} buffer \quad \langle \mu \mapsto buffer \rangle$
 $\vdash \text{read ins buf}: \quad unit \quad \langle \mu \mapsto initbuffer \rangle$

The post expression:

$ins: stream, buf: ref_{\mu} buffer \quad \langle \mu \mapsto buffer \rangle$
 $\vdash \text{post (read ins buf)}:$
 $promise_{\pi} unit \quad \langle Wait(\pi, \mu \mapsto initbuffer) \rangle$

Package resources produced by task in *wait permission*

A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));  
    do_copy buf  
  in do_copy (make_buffer ())
```

A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));
```

do
in do

Given read: $stream \rightarrow ref_{\mu} buffer \langle \mu \mapsto buffer \rangle \rightarrow$
 $unit \langle \mu \mapsto initbuffer \rangle$

... $\langle \mu \mapsto buffer \rangle \vdash \text{read ins buf} : unit \langle \mu \mapsto initbuffer \rangle$

A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));  
  do_copy buf
```

```
in do_copy (n ...  $\langle \mu \mapsto \text{buffer} \rangle \vdash \text{post (read ins buf)}:$   
 $\text{promise}_{\pi} \text{ unit } \langle \text{Wait}(\pi, \mu \mapsto \text{initbuffer}) \rangle$ )
```

A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));  
  do_copy buf
```

```
in do_copy ...  $\langle \mu \mapsto \text{buffer} \rangle \vdash$  wait (post (read ins buf)):  
  unit  $\langle \mu \mapsto \text{initbuffer} \rangle$ 
```


A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));  
  do_copy buf
```

in do_copy (make ... $\langle \mu \mapsto \text{buffer} \rangle \vdash \text{read, then write:}$
 $\text{unit} \langle \mu \mapsto \text{initbuffer} \rangle$

A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));  
    do_copy buf  
  in do_copy (make_buffer ())
```

... $\langle \mu \mapsto \text{buffer} \rangle \vdash$ loop body:
unit $\langle \mu \mapsto \text{buffer} \rangle$

A non-trivial example

Earlier example, slightly extended.

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
    wait (post (read ins buf));  
    wait (post (write outs buf));  
    do_copy buf  
  in do_copy (make_buffer ())
```

$\dots \langle emp \rangle \vdash \text{function body: } unit \langle \mu \mapsto buffer \rangle$

- The type system
- **Type inference**
- Some larger examples

What type inference does

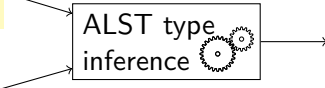
OCaml program

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof ins then () else  
      wait (post (read ins buf));  
      wait (post (write outs buf));  
      do_copy buf  
  in do_copy (make_buffer ())
```

Types of external functions

```
read:  $stream \rightarrow ref_{\mu} buffer \langle \mu \mapsto buffer \rangle \rightarrow$   
       $unit \langle \mu \mapsto initbuffer \rangle$   
write:  $stream \rightarrow ref_{\mu} buffer \langle \mu \mapsto initbuffer \rangle \rightarrow$   
        $unit \langle \mu \mapsto initbuffer \rangle$   
make_buffer:  $unit \langle emp \rangle \rightarrow$   
             $ref_{\mu} buffer \langle \mu \mapsto buffer \rangle$ 
```

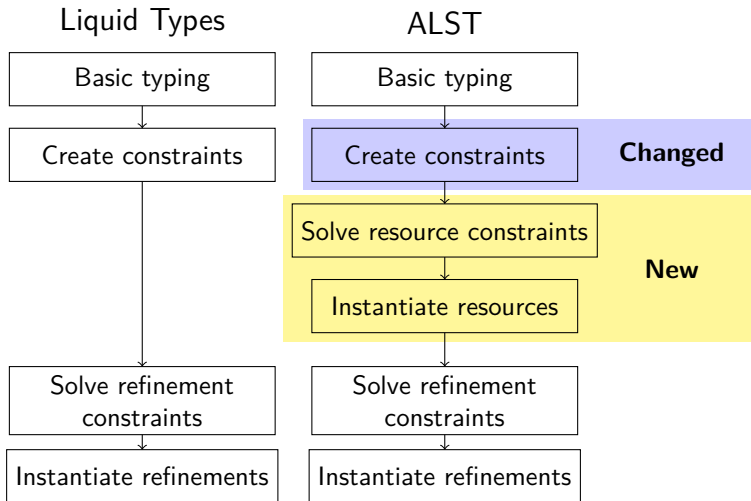
ALST type
inference



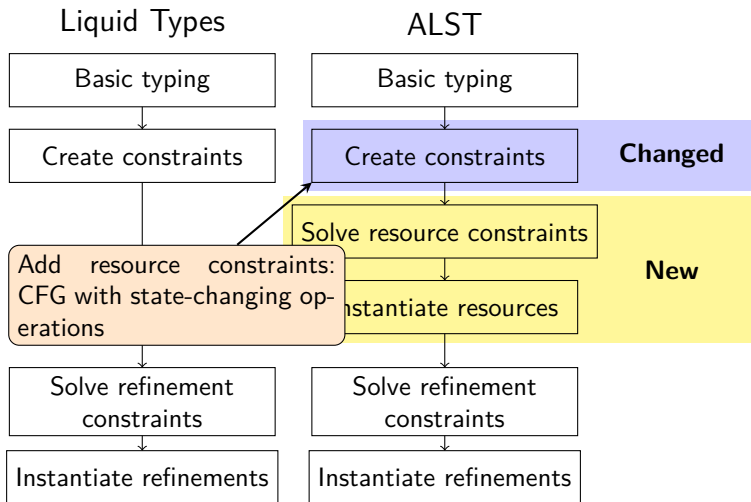
Typed program

```
let copy_buf ins outs =  
  let rec do_copy buf =  
    if eof  $\dots \langle emp \rangle \vdash$  function body:  
      wait  $unit \langle \mu \mapsto buffer \rangle$   
      wait (post (write outs buf));  
      do_copy buf  
  in do_copy (make_buffer ())
```


From liquid types to ALST



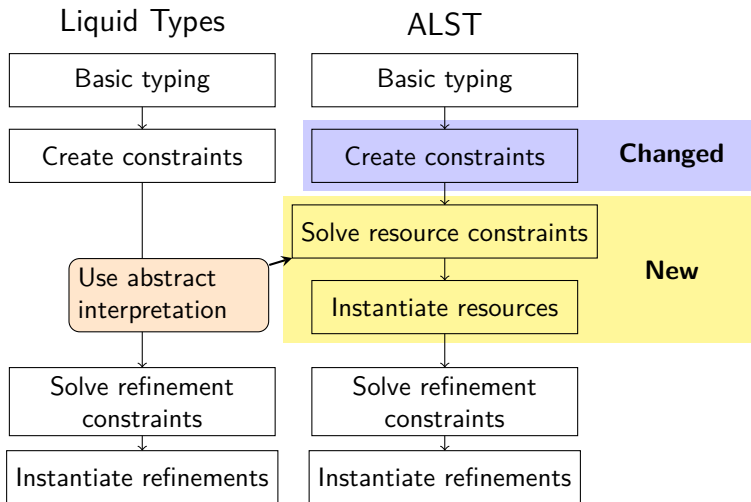
From liquid types to ALST



From liquid types to ALST



From liquid types to ALST



Fundamental limitations

- No way to express pre-/post-condition polymorphism.
Cannot provide a single useful type for this:

```
let rec iter f l = match l with  
  | x::l -> f x; iter f l  
  | [] -> ()
```

Fundamental limitations

- No way to express pre-/post-condition polymorphism.
Cannot provide a single useful type for this:

```
let rec iter f l = match l with  
  | x::l -> f x; iter f l  
  | [] -> ()
```

This is much harder than it seems: Pre-/post-condition of `iter` would depend on those of `f`.

Fundamental limitations

- No way to express pre-/post-condition polymorphism.
Cannot provide a single useful type for this:

```
let rec iter f l = match l with  
  | x::l -> f x; iter f l  
  | [] -> ()
```

This is much harder than it seems: Pre-/post-condition of `iter` would depend on those of `f`.

Workaround: Consider partially-applied versions of `iter`.

Fundamental limitations

- No way to express pre-/post-condition polymorphism.
Cannot provide a single useful type for this:

```
let rec iter f l = match l with  
  | x::l -> f x; iter f l  
  | [] -> ()
```

This is much harder than it seems: Pre-/post-condition of `iter` would depend on those of `f`.

Workaround: Consider partially-applied versions of `iter`.

- No way to handle unbounded task creation.

```
let rec make_tasks l = match l with  
  | x::l -> post (f x) :: make_tasks l  
  | [] -> []
```

- The type system
- Type inference
- Some larger examples

Outline: Larger examples

- The type system
- Type inference
- **Some larger examples**
 - Several variants of copying loops.
 - Coordination code for a parallel SAT solver.
 - Code from the MirageOS FAT file system.

Outline: Larger examples

- The type system
- Type inference
- Some larger examples
 - Several variants of copying loops.
 - Coordination code for a parallel SAT solver.
 - Code from the MirageOS FAT file system.

The MirageOS FAT example

Code taken from an old version of MirageOS:

```
type state = { format: ...; mutable fat: fat_type }

let update_directory_containing x path =
  post (let c = Fat_entry.follow_chain x.format x.fat in ...)

let update x = ...
  update_directory_containing x path;
  x.fat <- List.fold_left update_allocations x.fat fat_allocations
```

The MirageOS FAT example

Code taken from an old version of MirageOS:

```
type state = { format: ...; mutable fat: fat_type }

let update_directory_containing x path =
  post (let c = Fat_entry.follow_chain x.format x.fat in ...)

let update x = ...
  update_directory_containing x path;
  x.fat <- List.fold_left update_allocations x.fat fat_allocations
```

ALST output: Access to resource ξ_{42} that is wrapped in a wait permission.

The MirageOS FAT example

Code taken from an old version of MirageOS:

```
type state = { format: ...; mutable fat: fat_type }

let update_directory_containing x path =
  post (let c = Fat_entry.follow_chain x.format x.fat in ...)

let update x = ...
  update_directory_containing x path;
  x.fat <- List.fold_left update_allocations x.fat fat_allocations
```

ALST output: Access to resource ξ_{42} that is wrapped in a wait permission.

This code has a race condition!

Summary

Asynchronous Liquid Separation Types allow reasoning about OCaml programs with asynchronous concurrency and mutable state. They are an extension of liquid types with concurrent separation logic and allow type inference.

<http://plv.mpi-sws.org/ALSTypes>

Contact: Johannes Kloos, jkloos@mpi-sws.org

Other approaches to the same problem:

- 1 Low-level liquid types (Rondon, Kawaguchi, Jhala; POPL 2010)
Applies the liquid types machinery to C programs.
No reasoning about first-class functions, no concurrency
- 2 Hoare Type Theory (Nanevski, Morrisett, Birkedal; JFP 2007)
A more ambitious way of combining types and program logics.
Geared towards type-checking and interactive theorem proving, little automated inference.
- 3 Mezzo (Protzenko; PhD thesis, INRIA 2014)
Takes a radically different approach based on permissions.
Does not provide type inference.