

Proof methods for concurrent programs

2. concurrent separation logic

Francesco Zappa Nardelli

INRIA Paris-Rocquencourt, MOSCOVA project-team

francesco.zappa_nardelli@inria.fr

<http://moscova.inria.fr/~zappa/teaching/mpri/2010/>

Warm-up

Hoare logic:

- Commands operate on the state: $C / s \rightarrow C' / s'$;
- statements P are assertions on the state: $s \models P$;
- a triple $\{P\} c \{Q\}$ states that whenever c is executed in a state satisfying P and the execution of c terminates, the state in which c 's execution terminates satisfies Q ;
- a logic system allows us to prove $\vdash \{P\} c \{Q\}$. The logic system is sound.

Separation logic. All of the above plus:

- Special assertions, $P * Q$, $E_1 \mapsto E_2$, empty, to describe the heap part of the state.

Warm-up

Special assertions, $P * Q$, $E_1 \mapsto E_2$, empty , to describe the heap part of the state.

Three axioms to reason about separation:

- *write*: $\{ E \mapsto _ \} [E] = E' \{ E \mapsto E' \}$
- *dispose*: $\{ E \mapsto _ \} \text{dispose}(E) \{ \text{empty} \}$
- *alloc*: $\{ \text{empty} \} x = \text{cons}(E_1, \dots, E_n) \{ x \mapsto E_1 * x+1 \mapsto E_2 * \dots * x+(n-1) \mapsto E_n \}$

where $E \mapsto _$ is a shorthand for $\exists x. E \mapsto x$.

Exercise: prove that $\{ i \mapsto v \} x := [i] \{ i \mapsto v \wedge x = v \}$.

Pure assertions

Remark: some assertions are independent of the heap, e.g. $x = v$.

Definition: an assertion P is *pure*, iff for all stores s and heaps h_1 and h_2 , it holds

$$(s, h_1) \vdash P \quad \text{iff} \quad (s, h_2) \vdash P .$$

Some key properties of pure assertions:

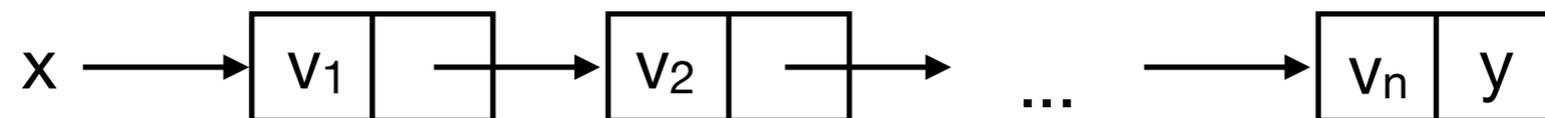
$$P \wedge Q \Rightarrow P * Q \quad \text{when } P \text{ or } Q \text{ is pure;}$$

$$P * Q \Rightarrow P \wedge Q \quad \text{when } P \text{ and } Q \text{ are pure;}$$

$$(P \wedge Q) * R \Rightarrow (P * R) \wedge Q \quad \text{when } Q \text{ is pure.}$$

Warm-up: list segments

The `lseg` predicate denotes *list segments*:



$$\text{lseg } [] (x, y) \equiv \text{empty} \wedge x = y$$

$$\text{lseg } v::\alpha (x, y) \equiv \exists j. x \mapsto v * (x+1 \mapsto j) * \text{lseg } \alpha (j, y)$$

Exercise: prove that the triple below holds.

$$\{ \text{lseg } a \cdot \alpha (i, k) \} r := [i+1]; \text{dispose } i; \text{dispose } i+1; i := r \{ \text{lseg } \alpha (i, k) \}$$

Remark: it is important to be able to reason on the assertion. Prove, by structural induction on α , that:

$$\text{lseg } \alpha \cdot \beta (x, y) \Leftrightarrow \exists j. \text{lseg } \alpha (x, j) \wedge \text{lseg } \beta (j, y)$$

Warm-up: a cyclic buffer

We implement a *cyclic buffer* using:

- an active list segment $\text{lseg } \alpha (i, j)$ (where α is the content of the buffer);
- an inactive list segment $\text{lseg } \beta (j, i)$ (where β is arbitrary);
- an unchanged variable n records the combined length of the two lists.

When $i=j$ the buffer is empty or full:

- a variable m records the length of the active list segment.

Inserting and deleting elements on the buffer must preserve the invariant:

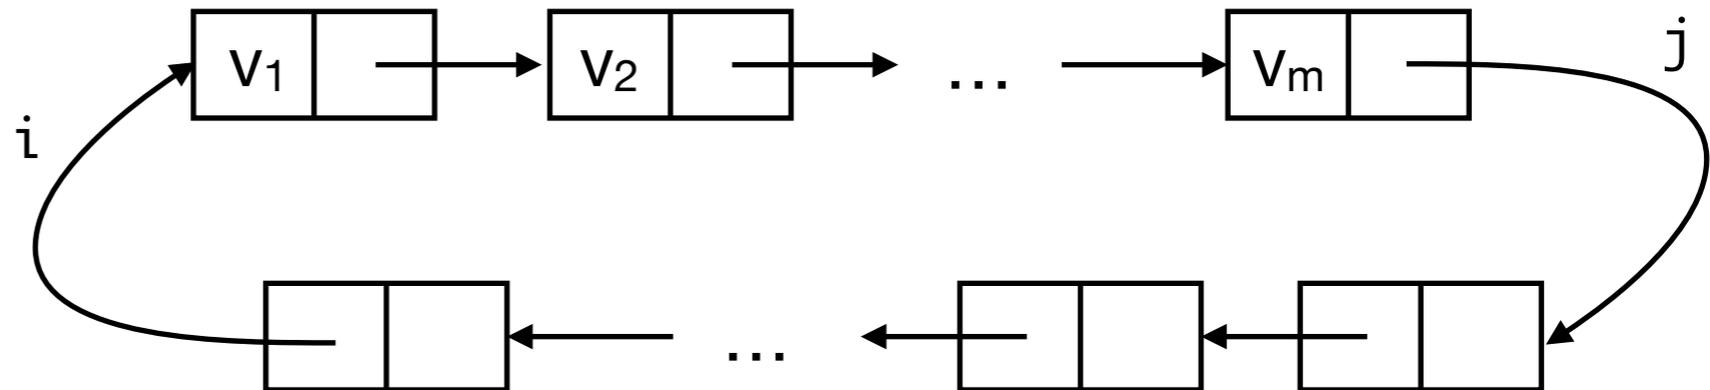
$$\exists \beta. (\text{lseg } \alpha (i, j) * \text{lseg } \beta (j, i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta$$

(where $\#$ computes the length of a sequence).

Warm-up: a cyclic buffer

Adding x to the buffer can be done by the code below (under the hypothesis that $n-m > 0$):

```
[j] := x;  
j := [j+1];  
m := m+1;
```

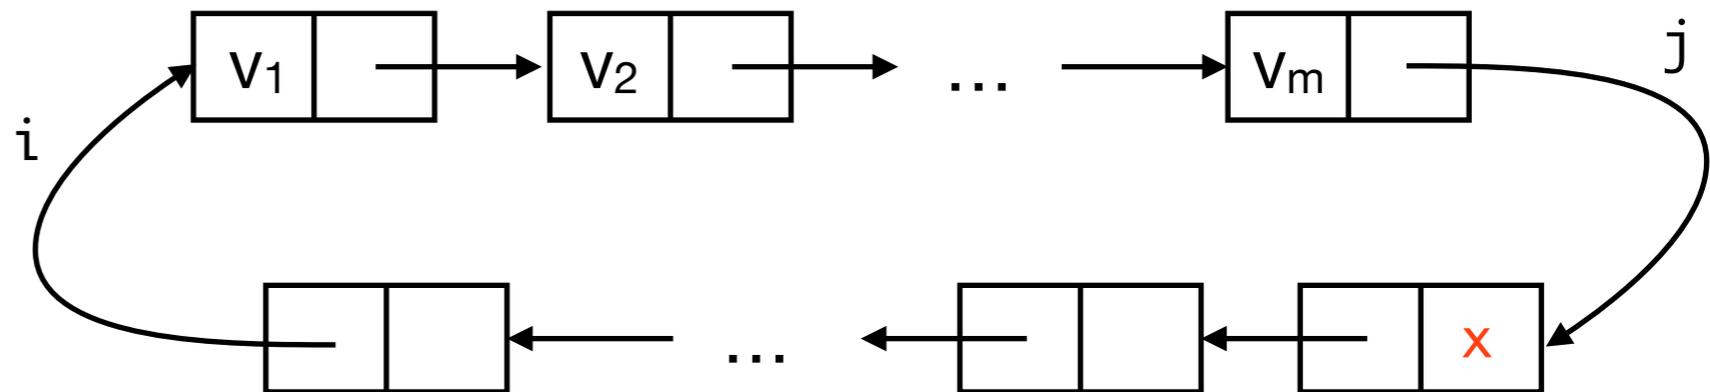


For reference: $\exists \beta. (\text{lseg } \alpha (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta$

Warm-up: a cyclic buffer

Adding x to the buffer can be done by the code below (under the hypothesis that $n-m > 0$):

```
[j] := x;  
j := [j+1];  
m := m+1;
```

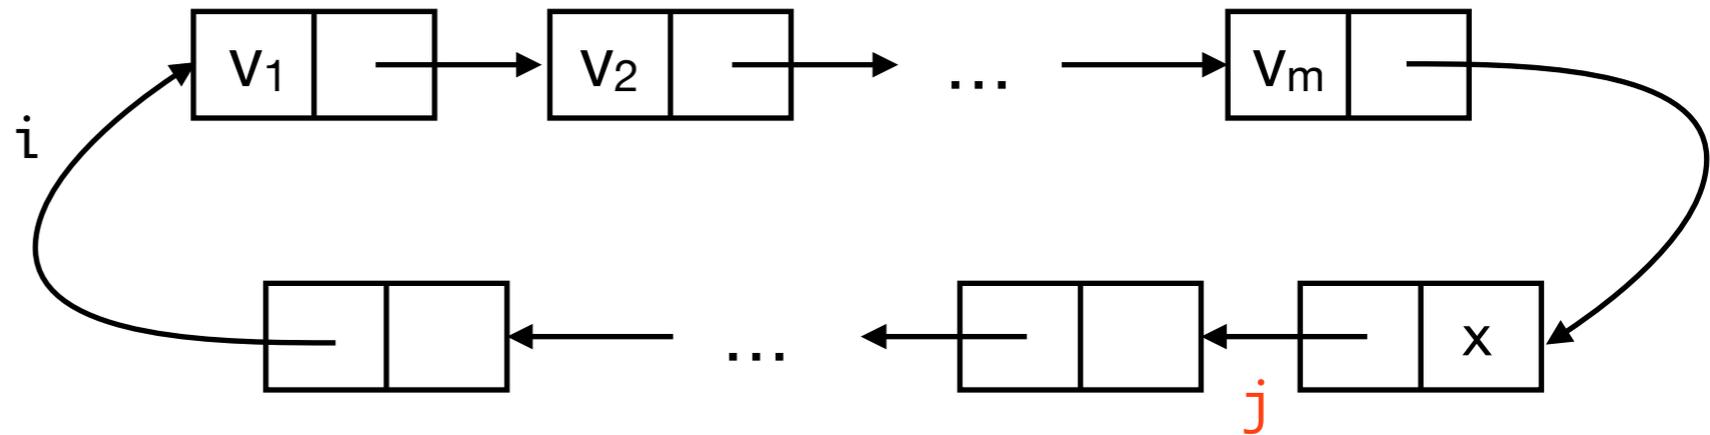


For reference: $\exists \beta. (\text{lseg } \alpha (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta$

Warm-up: a cyclic buffer

Adding x to the buffer can be done by the code below (under the hypothesis that $n-m > 0$):

```
[j] := x;  
j := [j+1];  
m := m+1;
```

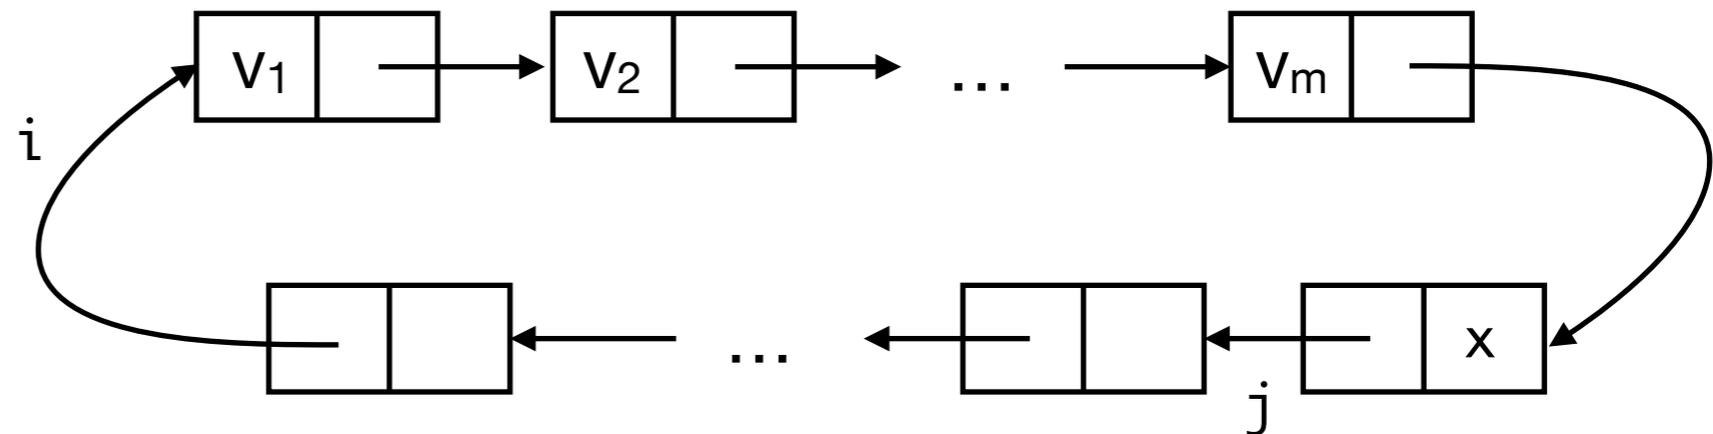


For reference: $\exists \beta. (\text{lseg } \alpha (i, j) * \text{lseg } \beta (j, i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta$

Warm-up: a cyclic buffer

Adding x to the buffer can be done by the code below (under the hypothesis that $n-m > 0$):

```
[j] := x;  
j := [j+1];  
m := m+1;
```



For reference: $\exists \beta. (\text{lseg } \alpha (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta$

Warm-up: a cyclic buffer

Exercise: we prove that the code below inserts x in the buffer.

$$\{ \exists \beta. (\text{lseg } \alpha (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta \wedge n-m > 0 \}$$

$[j] := x;$

$j := [j+1];$

$m := m+1;$

$$\{ \exists \alpha, \beta. (\text{lseg } \alpha \cdot x (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \cdot x \wedge n = \#\alpha \cdot x + \#\beta \}$$

Warm-up: a cyclic buffer

Exercise: we prove that the code below inserts x in the buffer.

$$\{ \exists \beta. (\text{lseg } \alpha (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#\beta \wedge n-m > 0 \}$$

$$\{ \exists b,\beta. (\text{lseg } \alpha (i,j) * \text{lseg } b \cdot \beta (j,i)) \wedge m = \#\alpha \wedge n = \#\alpha + \#b \cdot \beta \}$$

$$\{ \exists k,\beta. (\text{lseg } \alpha (i,j) * j \mapsto _,k * \text{lseg } \beta (k,i)) \wedge m = \#\alpha \wedge n-1 = \#\alpha + \#\beta \}$$

$[j] := x;$

$$\{ \exists k,\beta. (\text{lseg } \alpha (i,j) * j \mapsto x,k * \text{lseg } \beta (k,i)) \wedge m = \#\alpha \wedge n-1 = \#\alpha + \#\beta \}$$

$$\{ \exists k,\beta. j+1 \mapsto k * (\text{lseg } \alpha (i,j) * j \mapsto x * \text{lseg } \beta (k,i)) \wedge m = \#\alpha \wedge n-1 = \#\alpha + \#\beta \}$$

$j := [j+1];$

$$\{ \exists l,\beta. l+1 \mapsto j * (\text{lseg } \alpha (i,l) * l \mapsto x * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n-1 = \#\alpha + \#\beta \}$$

$$\{ \exists l,\beta. (\text{lseg } \alpha (i,l) * l \mapsto x,j * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \wedge n-1 = \#\alpha + \#\beta \}$$

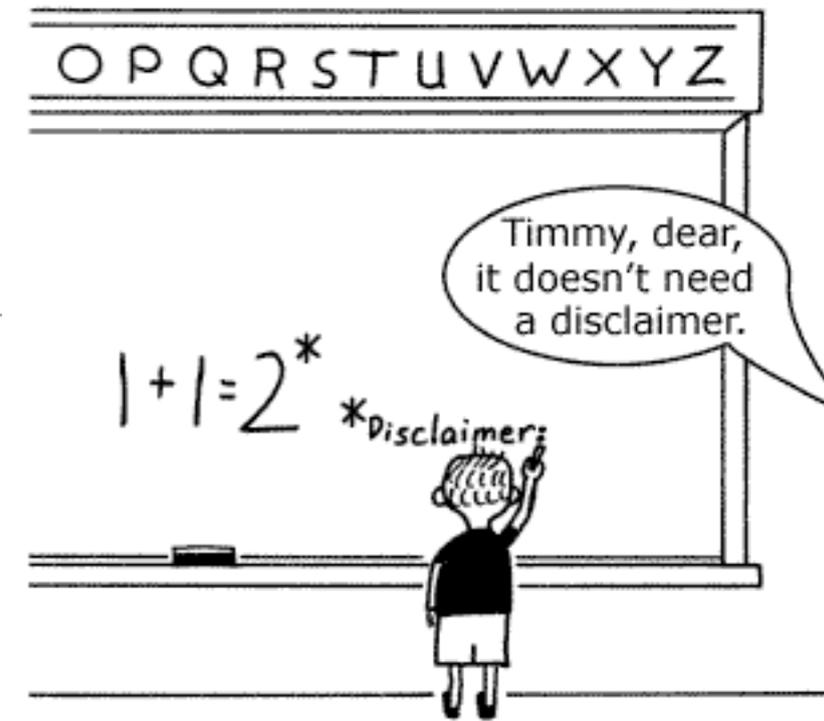
$$\{ \exists l,\beta. (\text{lseg } \alpha \cdot x (i,j) * \text{lseg } \beta (j,i)) \wedge m+1 = \#\alpha \cdot x \wedge n = \#\alpha \cdot x + \#\beta \}$$

$m := m+1;$

$$\{ \exists l,\beta. (\text{lseg } \alpha \cdot x (i,j) * \text{lseg } \beta (j,i)) \wedge m = \#\alpha \cdot x \wedge n = \#\alpha \cdot x + \#\beta \}$$

Be careful

Despite the appearances...

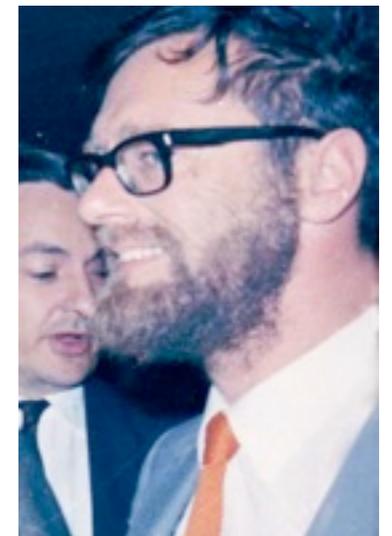


...mastering separation logics takes time...

(after all, we are reasoning about the heap!)

Concurrent separation logic

1. threads that mind their own business



Threads that mind their own business

Imagine a program composed by two threads, one updates $[x]$, the other $[y]$:

$[x] := 4 \quad \parallel \quad [y] := 5$

What can we prove about it?

Threads that mind their own business

Imagine a program composed by two threads, one updates $[x]$, the other $[y]$:

$$\begin{array}{ccc} \{x \mapsto _ \} & & \{y \mapsto _ \} \\ [x] := 4 & \parallel & [y] := 5 \\ \{x \mapsto 4 \} & & \{y \mapsto 5 \} \end{array}$$

1) We can give a (sequential) specification to each thread.

Threads that mind their own business

Imagine a program composed by two threads, one updates [x], the other [y]:

$$\begin{array}{ccc} & \{ x \mapsto - * y \mapsto - \} & \\ \{ x \mapsto - \} & & \{ y \mapsto - \} \\ [x] := 4 & \parallel & [y] := 5 \\ \{ x \mapsto 4 \} & & \{ y \mapsto 5 \} \\ & \{ x \mapsto 4 * y \mapsto 5 \} & \end{array}$$

- 1) We can give a (sequential) specification to each thread.
- 2) If x and y do not point to the same location, then we can guarantee that the final state satisfies $(x \mapsto 4 * y \mapsto 5)$.

Parallel composition of non-interfering threads

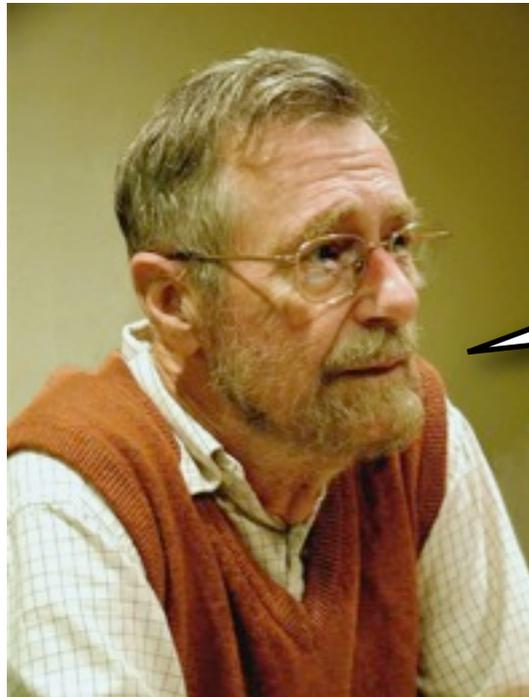
$$\frac{\{ P_1 \} C_1 \{ Q_1 \} \quad \{ P_2 \} C_2 \{ Q_2 \}}{\{ P_1 * P_2 \} C_1 \parallel C_2 \{ Q_1 * Q_2 \}}$$

if $\text{modifies}(C_1) \cap \text{fv}(P_2) = \text{modifies}(C_2) \cap \text{fv}(P_1) = \emptyset$, and
 $\text{modifies}(C_1) \cap \text{fv}(C_2) = \text{modifies}(C_2) \cap \text{fv}(C_1) = \emptyset$.

Parallel composition of non-interfering threads

$$\frac{\{ P_1 \} C_1 \{ Q_1 \} \quad \{ P_2 \} C_2 \{ Q_2 \}}{\{ P_1 * P_2 \} C_1 \parallel C_2 \{ Q_1 * Q_2 \}}$$

if $\text{modifies}(C_1) \cap \text{fv}(P_2) = \text{modifies}(C_2) \cap \text{fv}(P_1) = \emptyset$, and
 $\text{modifies}(C_1) \cap \text{fv}(C_2) = \text{modifies}(C_2) \cap \text{fv}(C_1) = \emptyset$.



... apart from the (rare) moments of explicit communication, *processes are to be regarded as completely independent of each other...*

Parallel composition of non-interfering threads

Remark: the "proof figure" below

$$\begin{array}{ccc} & \{ x \mapsto _ * y \mapsto _ \} & \\ \{ x \mapsto _ \} & & \{ y \mapsto _ \} \\ [x] := 4 & \parallel & [y] := 5 \\ \{ x \mapsto 4 \} & & \{ y \mapsto 5 \} \\ & \{ x \mapsto 4 * y \mapsto 5 \} & \end{array}$$

is an annotation form for

$$\frac{\{ x \mapsto _ \} [x] := 4 \{ x \mapsto 4 \} \quad \{ y \mapsto _ \} [y] := 5 \{ y \mapsto 5 \}}{\{ x \mapsto _ * y \mapsto _ \} [x] := 4 \parallel [y] := 5 \{ x \mapsto 4 * y \mapsto 5 \}}$$

Example: parallel disposal of a tree

$\text{tree } p \equiv (p = \text{nil} \wedge \text{empty}) \vee (\exists j, k. p \mapsto j * p+1 \mapsto k * \text{tree } j * \text{tree } k)$

```
procedure DispTree(p) {  
  if p != nil then {  
    a := [p];  
    b := [p+1];  
    DispTree(a) || DispTree(b);  
    dispose(p+1);  
    dispose(p);  
  }  
}
```

This is an example of a **shape predicate**: it only describes the memory layout of the data structure, not the actual content.

This is a recursive procedure: to prove its correctness you can assume that the specification holds for the recursive calls. Cheating: in these lectures we won't formalise procedure calls...

This is a bad implementation of parallel disposal of a tree, why?

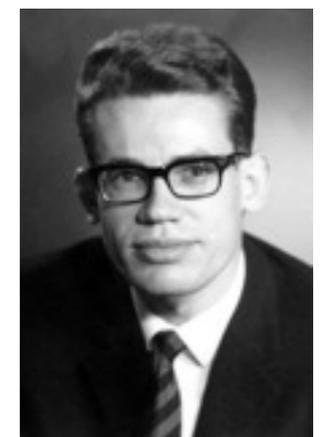
Exercise: assume that $\{ \text{tree } p \} \text{DispTree}(p) \{ \text{empty} \}$ holds.

Prove that the body of `DispTree` satisfies $\{ \text{tree } p \} \text{body} \{ \text{empty} \}$.

Concurrent separation logic



2. synchronising racy threads



Racy programs

In current practice, most programs of interest are racy, e.g.:

$$\begin{array}{ccc} \{ x \mapsto _ \} & & \{ x \mapsto _ \} \\ [x] := 10 & \parallel & [x] := 20 \\ \{ x \mapsto 10 \} & & \{ x \mapsto 20 \} \end{array}$$

But we cannot send $x \mapsto _$ to both threads:

$$(x \mapsto _ * x \mapsto _) \Leftrightarrow F$$

Racy programs

In current practice, most programs of interest are racy, e.g.:

$$\begin{array}{ccc} \{ x \mapsto _ \} & & \{ x \mapsto _ \} \\ [x] := 10 & \parallel & [x] := 20 \\ \{ x \mapsto 10 \} & & \{ x \mapsto 20 \} \end{array}$$

But we cannot send $x \mapsto _$ to both threads:

$$(x \mapsto _ * x \mapsto _) \Leftrightarrow F$$

This program does not have a race-free start state... to reason about such programs we must be explicit about the granularity of the interactions.

Racy programs

In current practice, most programs of interest are racy, e.g.:

$\{x \mapsto _ \}$ $\{x \mapsto _ \}$

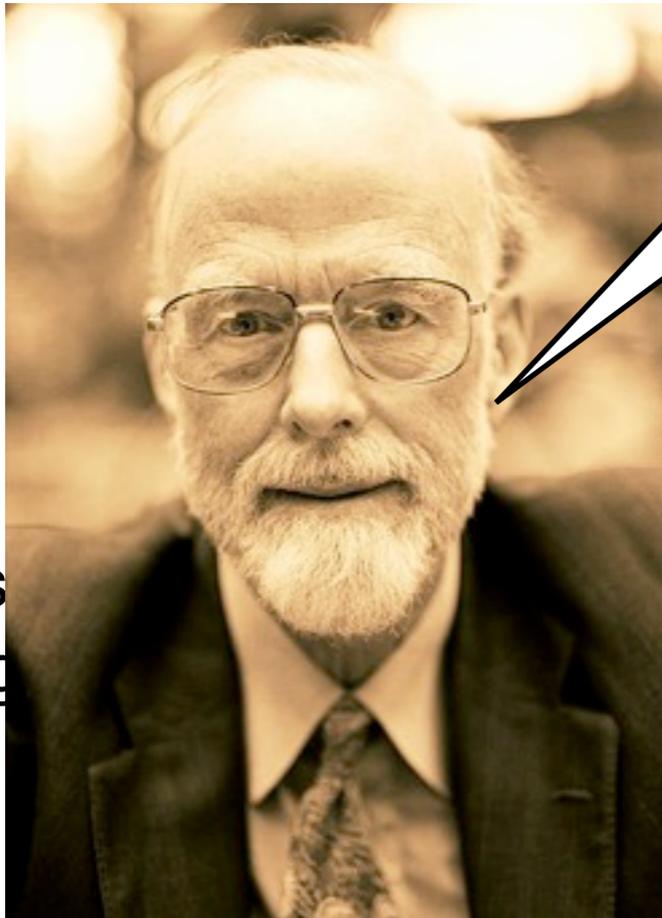
...designing a program to control the fantastic number of combinations involved in arbitrary interleaving...

But

$_$ to both threads:

$$(x \mapsto _ * x \mapsto _) \Leftrightarrow F$$

This program does not have a race-free start state... to reason about such programs we need to be explicit about the granularity of the interactions.



Conditional critical regions (Hoare, 72)

A program is a collection of resources shared by concurrent threads:

init

resource r1 (list of variables) ... resource rn (list of variables)

C1 || ... || Cm

A thread can obtain an exclusive access to a resource:

with r when B do C

Constraints:

- if a variable belongs to a resource, it cannot appear in a parallel process except in a critical section for that resource;
- if a variable is changed in one process, it cannot appear in another unless it belongs to a resource.

Examples of racy programs

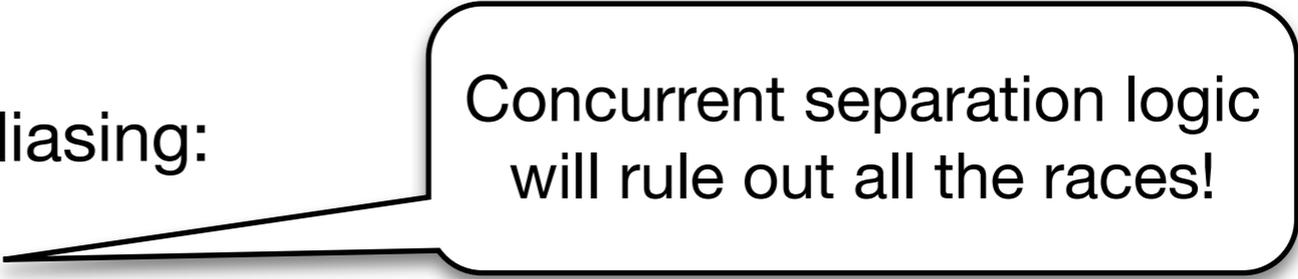
- if a variable belongs to a resource, it cannot appear in a parallel process except in a critical section for that resource;
- if a variable is changed in one process, it cannot appear in another unless it belongs to a resource.

These programs do not respect the constraints above:

- $x := 3 \parallel x := x + 1$
- $x := 3 \parallel \text{with } r \text{ when true do } x := x + 1$

In general, races depend on aliasing:

- $[x] := 3 \parallel [y] := 4$



Concurrent separation logic
will rule out all the races!

(Informal) semantics of CCRs

- The `init` command is executed first (and allocates some resources).

- A declaration

`resource r (x1, ..., xn)`

states that the variables x_1, \dots, x_n can only be accessed while holding the resource r .

- The command

`with r when B do C`

executes C while holding the resource r : no other thread can access the variables x_1, \dots, x_n while C executes. However the execution of C is postponed until the statement B is true.

Programming a bounded buffer with CCRs

Remark: to simplify notations we use arrays instead of pointers.

Buffer space and pointers are encapsulated in the *buffer* resource:

```
resource buffer (  
    item pool[n];  
    int count, in, out;  
)
```

and here the producer and consumer code:

Producer:

```
with buffer when (count < n) {  
    pool[in] = nextp;  
    in = (in+1) % n;  
    count++;  
}
```

Consumer:

```
with buffer when (count > 0) {  
    nextc = pool[out];  
    out = (out+1) % n;  
    count--;  
}
```

Example: semaphores

We can program binary semaphores with CCRs:

```
s := 1;
```

```
resource s (s)
```

```
P(s) = with s when s = 1 do s := 0
```

```
V(s) = with s when s = 0 do s := 1
```

Remark: usually CCRs are implemented using semaphores, not the other way round. However CCRs are simpler from a logical point of view.

Can we associate some property (some invariant?) to a semaphore s ?

Example: semaphores

Typical use of a semaphore s :

$$\begin{array}{ccc} P(s) & & P(s) \\ [10] := 43 & \parallel & [10] := 57 \\ V(s) & & V(s) \end{array}$$

The location $[10]$ is protected by $/$ associated to $/$ (*owned by*) the semaphore.

Idea:

$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge [10] \mapsto _)$$

when the semaphore is held, the location $[10]$ is owned by the thread that holds the semaphore.

When the semaphore is not held, no thread can access the location $[10]$. (because no thread can have $[10] \mapsto _$ in its precondition)

Axioms for CCRs

$$\frac{\{ P \} \text{ init } \{ RI_1 * \dots * RI_n * P' \} \quad \{ P' \} C_1 \parallel \dots \parallel C_n \{ Q \}}{\{ P \} \text{ init}; \text{ resource } r_1 (\dots) \dots \text{ resource } r_n (\dots); C_1 \parallel \dots \parallel C_n \{ RI_1 * \dots * RI_n * Q \}}$$

The init code allocates the resources stored in the resource invariants; the threads are then executed. Threads grab control of the resource invariants when entering the CCRs:

$$\frac{\{ (P * RI_r) \wedge S \} C \{ Q * RI_r \}}{\{ P \} \text{ with } r \text{ when } S \text{ do } C \{ Q \}}$$

Idea: inside the critical region, the threads has visibility of the state associated to (protected by) the resource.

Example: semaphores

Exercise: suppose that

$$Rl_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

Can we prove that the following holds?

$$\begin{array}{ccc} & \{ \text{empty} \} & \\ P(s) & & P(s) \\ [10] := 43 & \parallel & [10] := 57 \\ V(s) & & V(s) \\ & \{ \text{empty} \} & \end{array}$$

Example: semaphores

Reminder: $RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$

Zoom on the proof of thread 1:

{ empty }

P(S)

{ 10 \mapsto _ }

[10] := 43

{ 10 \mapsto _ }

V(s)

{ empty }

$$\frac{\{ \text{empty} * s = 1 * RI_s \} s := 0 \{ s = 0 * 10 \mapsto _ * RI_s \}}{\{ \text{empty} \} \text{ with } s \text{ when } s=1 \text{ do } s := 0 \{ 10 \mapsto _ \}}$$

$$\frac{\{ 10 \mapsto _ * s = 0 * RI_s \} s := 1 \{ s = 1 * RI_s \}}{\{ 10 \mapsto _ \} \text{ with } s \text{ when } s = 0 \text{ do } s := 1 \{ \text{empty} \}}$$

Key observation: the resource $10 \mapsto _$ "flows" from the RI to the thread and back!

Example: semaphores

Reminder: $RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$

Zoom on the proof of thread 1:

Since $s = 0$ the RI cannot hold any resource (empty).

$\{ \text{empty} \}$
 $P(S)$
 $\{ 10 \mapsto _ \}$
 $[10] := 43$
 $\{ 10 \mapsto _ \}$
 $V(s)$
 $\{ \text{empty} \}$

$$\frac{\{ \text{empty} * s = 1 * RI_s \} s := 0 \{ s = 0 * 10 \mapsto _ * RI_s \}}{\{ \text{empty} \} \text{ with } s \text{ when } s=1 \text{ do } s := 0 \{ 10 \mapsto _ \}}$$

The resource $10 \mapsto _$ gets into the scope of the thread.

$$\frac{\{ 10 \mapsto _ * RI_s \}}{\{ 10 \mapsto _ \} \text{ with } s \text{ when } s = 0 \text{ do } s := 1 \{ \text{empty} \}}$$

Key observation: the resource $10 \mapsto _$ "flows" from the RI to the thread and back!

Example: semaphores

Reminder: $RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$

Zoom on the proof of thread 1:

{ empty }

P(S)

{ 10 \mapsto _ }

[10] := 43

{ 10 \mapsto _ }

V(s)

{ empty }

$$\frac{\{ \text{empty} * s = 1 * RI_s \} s := 0 \{ s = 0 * 10 \mapsto _ * RI_s \}}{\{ \text{empty} \} \text{ with } s \text{ when } s=1 \text{ do } s := 0 \{ 10 \mapsto _ \}}$$

$$\frac{\{ 10 \mapsto _ * s = 0 * RI_s \} s := 1 \{ s = 1 * RI_s \}}{\{ 10 \mapsto _ \} \text{ with } s \text{ when } s = 0 \text{ do } s := 1 \{ \text{empty} \}}$$

Key observation: the resource $10 \mapsto _$ "flows" from the RI to the thread and back!

Example: semaphores

Exercise: can you prove the triple below under the stronger invariant?

$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto 57)$$

	{ empty }	
P(s)		P(s)
[10] := 43		[10] := 57
V(s)		V(s)
	{ empty }	

With the stronger invariant we cannot prove that the invariant holds after executing $V(s)$!

Remark: a dull specification?

$\vdash \{ \text{empty} \} P(s); [10] := 43; V(s) \parallel P(s); [10] := 57; V(s) \{ \text{empty} \}$

holds under the invariant $RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$.

This guarantees that if the program is executed in a state that satisfies $\{ \text{empty} * RI_s \}$, if the program terminates, it ends in a state that satisfies $\{ \text{empty} * RI_s \}$.

Even if the precondition and the postcondition does not look very interesting, the triple (also) guarantees that:

- the program is race-free;
all the accesses to the shared resource were correctly protected by locks
- the resource invariant was preserved by all the threads;
- no memory leaks occurred.

Producer/consumer via semaphores

Two semaphores. Initially $free = 1$ and $busy = 0$.

{ empty }

P($free$)		P($busy$)
[10] := 43		x := [10]
V($busy$)		V($free$)

{ empty }

For s being either $free$ or $busy$, the semaphore invariant is:

$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

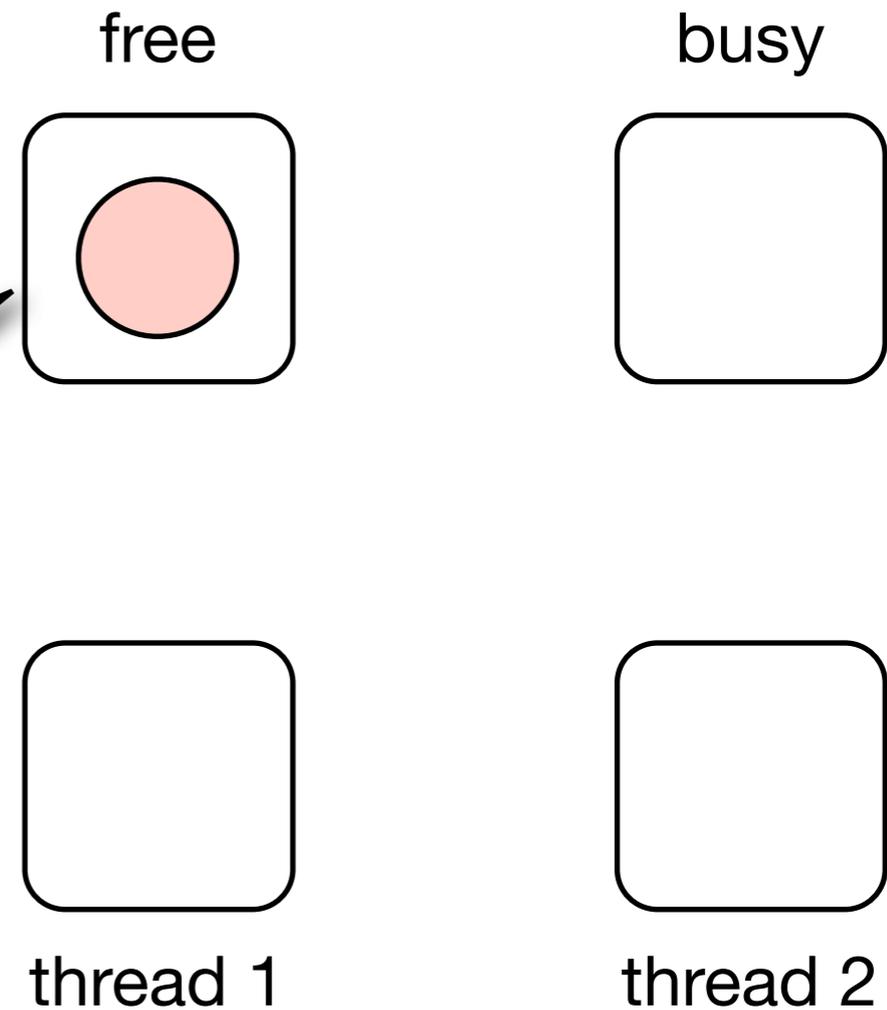
Exercise: prove the triple above.

Producer/consumer via semaphores

Initially free = 1 and busy = 0.

$\{ \text{empty} \}$
 $\{ \text{empty} \}$
 $\{ \text{empty} \}$
 $\{ 10 \mapsto _ \}$
 $[10] :=$
 $\{ 10 \mapsto _ \}$
 $\{ 10 \mapsto _ \}$
 $\{ \text{empty} \}$
 $\{ \text{empty} \}$
 $\{ \text{empty} \}$

This diagram records the current owner of the resource [10].



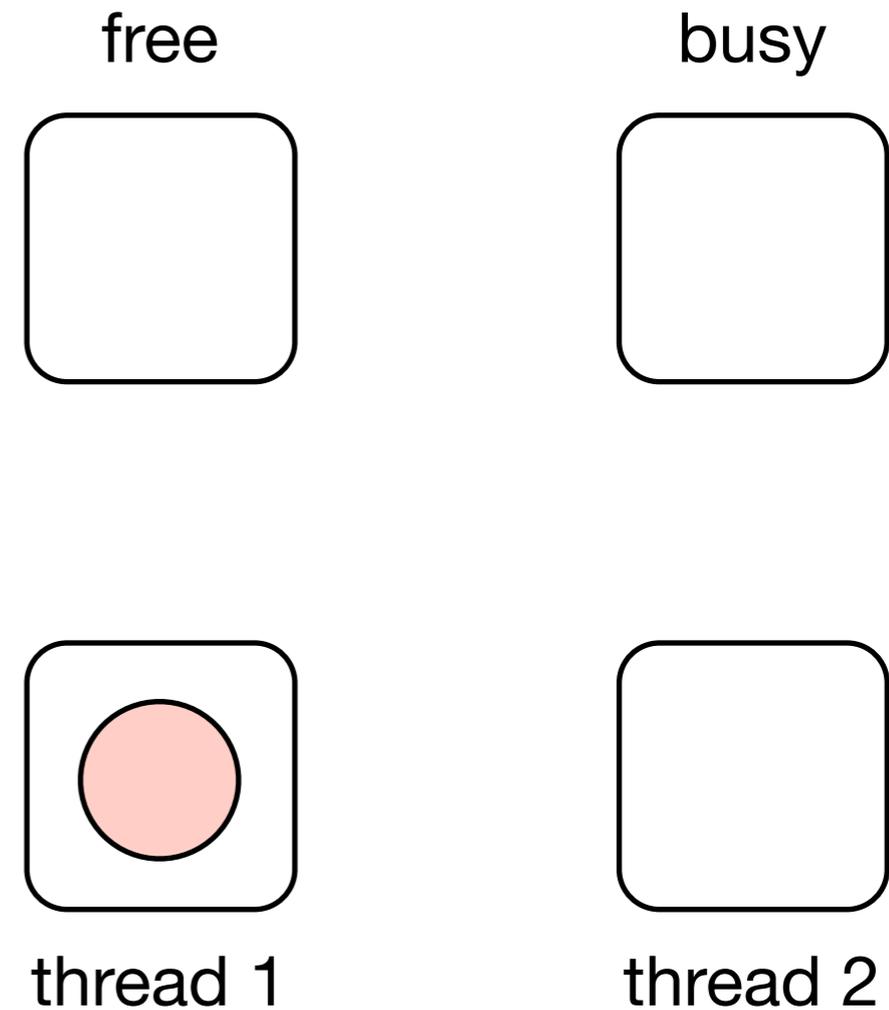
$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

Producer/consumer via semaphores

Initially $free = 1$ and $busy = 0$.

```

    { empty }
    { empty }
    P(free)   P(busy)
    { 10 ↦ _ } { 10 ↦ _ }
    [10] := 43 || x := [10]
    { 10 ↦ _ } { 10 ↦ _ }
    V(busy)   V(free)
    { empty } { empty }
    { empty }
  
```

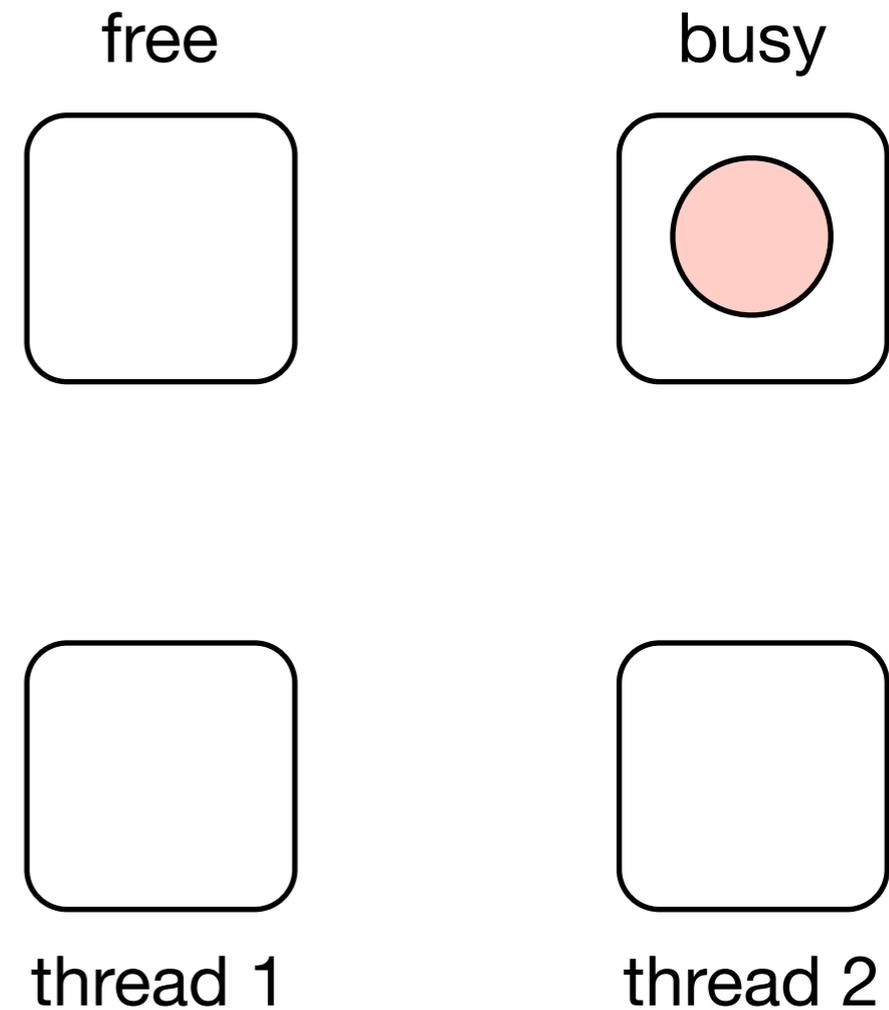


$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

Producer/consumer via semaphores

Initially $free = 1$ and $busy = 0$.

{ empty }		{ empty }
P(free)		P(busy)
{ 10 \mapsto _ }		{ 10 \mapsto _ }
[10] := 43		x := [10]
{ 10 \mapsto _ }		{ 10 \mapsto _ }
V(busy)		V(free)
{ empty }		{ empty }
	{ empty }	



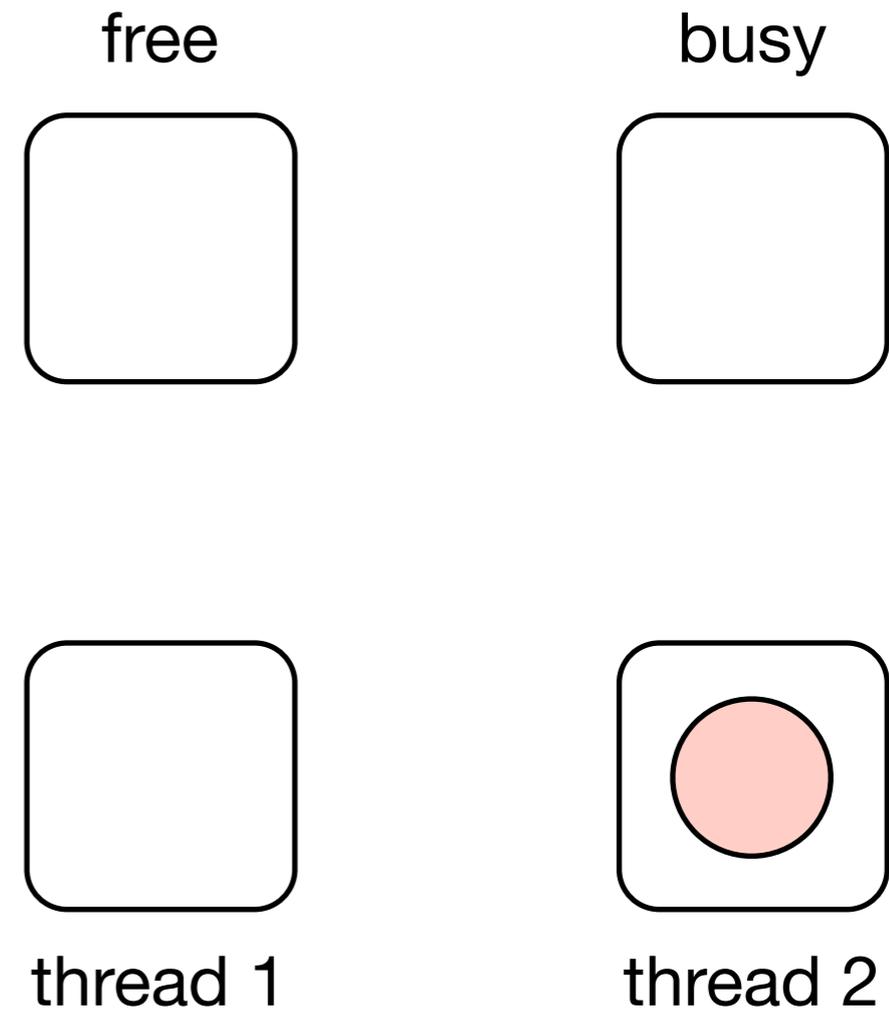
$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

Producer/consumer via semaphores

Initially $free = 1$ and $busy = 0$.

```

    { empty }
{ empty }   { empty }
P(free)     P(busy)
{ 10 ↦ _ } { 10 ↦ _ }
[10] := 43  || x := [10]
{ 10 ↦ _ } { 10 ↦ _ }
V(busy)     V(free)
{ empty }   { empty }
    { empty }
  
```

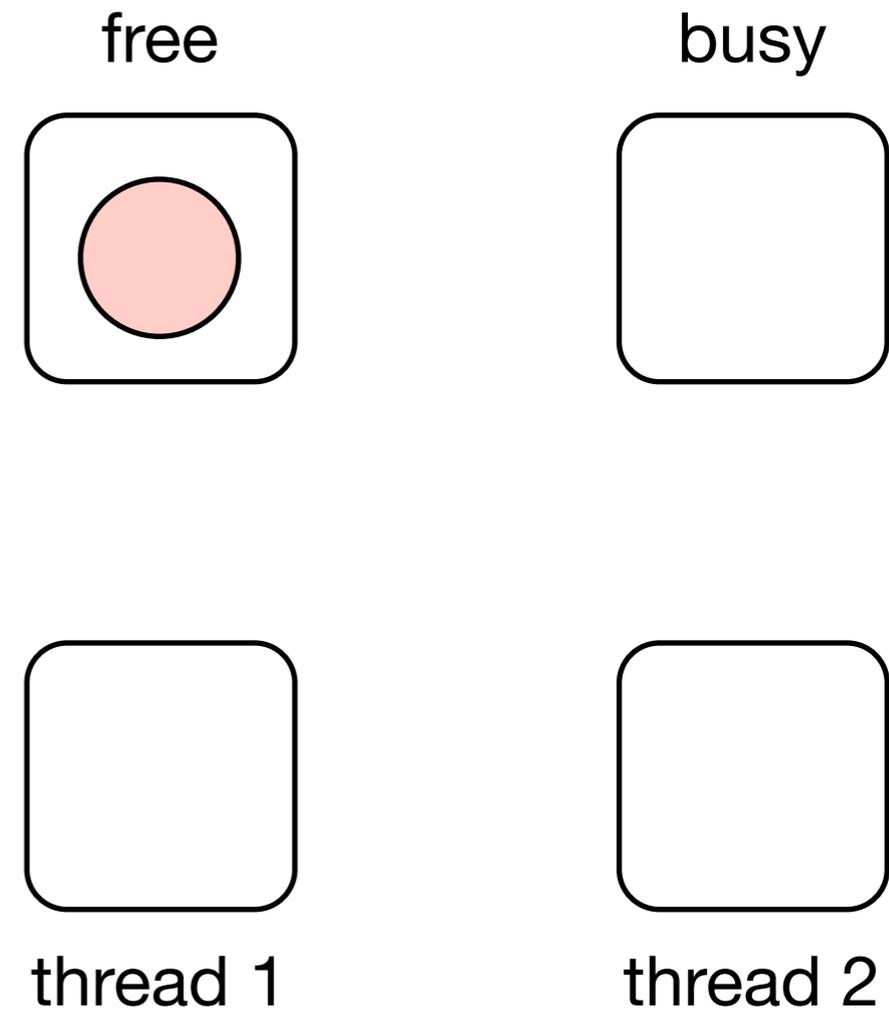


$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

Producer/consumer via semaphores

Initially $free = 1$ and $busy = 0$.

{ empty }		{ empty }
{ empty }		{ empty }
P($free$)		P($busy$)
{ 10 \mapsto _ }		{ 10 \mapsto _ }
[10] := 43		x := [10]
{ 10 \mapsto _ }		{ 10 \mapsto _ }
V($busy$)		V($free$)
{ empty }		{ empty }
		{ empty }



$$RI_s = (s = 0 \wedge \text{empty}) \vee (s = 1 \wedge 10 \mapsto _)$$

Remarks

- Each semaphore invariant talks only about itself, not about other semaphores or processes.
- Each assertion within a process talks about only its own state, not the state of the other process or even the semaphores.
- We do not maintain $0 \leq \text{free} + \text{busy} \leq 1$ as a global invariant.
- Semaphores are “logically attached” to resources. P and V are *ownership transformers*.

Example: a single-place buffer

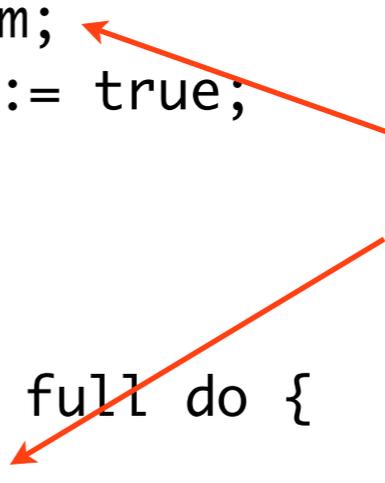
Initially `full := false`.

resource `buf(c, full)`

Filling the buffer:

```
put (m) = with buf when ¬full do {  
    c := m;  
    full := true;  
}
```

passing a pointer,
not the value.



Emptying the buffer:

```
get (n) = with buf when full do {  
    n := c;  
    full := false;  
}
```

Invariant:

$RI = (\text{empty} \wedge \neg \text{full}) \vee (c \mapsto _ \wedge \text{full})$

Example: a single-place buffer

RI = (empty \wedge \neg full) \vee (c \mapsto _ \wedge full)

{ empty }

full := false;

{ empty \wedge \neg full }

resource buf (c, full)

{ empty * empty * RI }

{ empty }

x := cons(3);

{ x \mapsto 3 }

with buf when \neg full do

{ (x \mapsto 3)*(empty \wedge \neg full) }

c := x; full := true;

{ c \mapsto 3 * (empty \wedge full) }

{ empty * RI }

put(x)

{ empty }

with buf when full do

{ empty * ((c \mapsto _) \wedge full) }

||

y := c; full := false;

{ (y \mapsto _) * (empty \wedge \neg full) }

dispose (y);

{ empty }

get(y)

{ empty * empty * RI }

{ RI }

Ownership *is in the eye of theasserter*

Transfer of ownership is not determined operationally:

whatever we transfer depends on what we want to prove.

In the last example ownership of the location allocated by thread 1 had to be transferred to thread 2, so that thread 2 could safely dispose it:

```
x:=cons(3);      get(y);
put(x);          || use(y);
                 dispose(y)
```

Reminder: $RI = (\text{emp} \wedge \neg \text{full}) \vee (c \mapsto _ \wedge \text{full})$

Ownership *is in the eye of the asserter*

Transfer of ownership is not determined operationally:

whatever we transfer depends on what we want to prove.

The code below is silly, but should be provable:

```
x:=cons(3);      get(y);  
put(x);         ||  
dispose(x);
```

Exercise: prove the code above, using the invariant

$$RI = (\text{emp} \wedge \neg \text{full}) \vee (\text{emp} \wedge \text{full})$$

Ownership *is in the eye of the asserter*

Transfer of ownership is not determined operationally:

whatever we transfer depends on what we want to prove.

However you won't be able to prove:

```
x:=cons(3);      get(y);  
put(x);          ||  dispose(y);  
dispose(x);
```

because ownership cannot flow both to thread 1 and thread 2.

This is fortunate: this program attempts to dispose the same pointer twice.

Simple exercises

- Is the triple below derivable?

```
{ empty }  
  x := cons(3);  
  z := cons(3);  
  [x] := 4 || [z] := 5;  
{ x ↦ 4 * z ↦ 5 }
```

- And this?

```
{ empty }  
  x := cons(3);  
  [x] := 4 || [x] := 5;  
{ x ↦ - }
```

- And this?

```
{ empty }  
  x := 4 || x := 5;  
{ empty }
```

- And this?

```
{ y = x+1 }  
  x := 4 || y := y+1;  
{ y = x+2 }
```

Simple exercises

- Is the triple below derivable?

```
{ empty }  
  x := cons(3);  
  z := cons(3);  
  [x] := 4 || [z]  
{ x ↦ 4 * z ↦ 5 }
```

This is a stack race!

- And this?

```
{ empty }  
  x := 4 || x := 5;  
{ empty }
```

- And

Here the race is between $x := 4$
and the proof of

```
{ y = x+1 } y := y+1 { y = x+2 }
```

```
{  
  x := cons(3);  
  [x] := 4 || [x] := 5;  
  x ↦ - }
```

- And this?

```
{ y = x+1 }  
  x := 4 || y := y+1;  
{ y = x+2 }
```

Simple exercises

- Is the triple below derivable?

```
{ empty }
  x := cons(3);
  z := cons(3);
  [x] := 4 || [z]
{ x ↦ 4 * z ↦ 5 }
```

This is a stack race!

- And this?

```
{ empty }
  x := 4 || x := 5;
{ empty }
```

- And this?

- And

Here the race is between $x := 4$ and the proof of

```
{ y = x+1 } y := y+1 { y = x+2 }
```

```
{ empty }
  x := cons(3);
  [x] := 4 || [x] := 5;
{ x ↦ - }
```

```
{ y = x+1 }
  x := 4 || y := y+1;
{ y = x+2 }
```

The logic forbids all kinds of races.

Exercise: parallel mergesort

Let $ls(p) = (\text{empty} \wedge p = \text{nil}) \vee \exists j. p \mapsto _ * (p+1 \mapsto j) * ls(j)$.

Suppose that the functions `split` and `merge` obey to the specifications

$$\{ ls(p) \} \text{split}(r,p) \{ ls(r) \}$$
$$\{ ls(p) * ls(q) \} \text{merge}(r,p,q) \{ ls(r) \}$$

Prove that:

$$\{ ls(p) \}$$

```
mergesort(r,p) {  
  if p = Nil then r := p;  
  else {  
    split(q,p);  
    mergesort(q1,q) || mergesort(p1,p);  
    merge(r,p1,q1)
```

$$\{ ls(r) \}$$

Exercise: parallel mergesort

Let $ls(p) = (\text{empty} \wedge p = \text{nil}) \vee \exists j. p \mapsto _ * (p+1 \mapsto j) * ls(j)$.

Suppose that the functions `split` and `merge` obey to the specifications

$\{ ls(p) \}$ `split`

$\{ ls(p) * ls(q) \}$

This is another example of a **shape predicate**: it only describes the memory layout of the data structure, not the actual content.

Prove that:

$\{ ls(p) \}$

`mergesort(r, p)` {

Concurrent separation logic is *decidable* for shape predicates.

(well, you still have to supply the loop invariants).

p);



Check out the SmallFoot tool.

Exercise: parallel mergesort

Let $ls(p) = (\text{empty} \wedge p = \text{nil}) \vee \exists j. p \mapsto _ * (p+1 \mapsto j) * ls(j)$.

Suppose that the functions `split` and `merge` obey to the specifications

$$\{ ls(p) \} \text{split}(r,p) \{ ls(r) \}$$
$$\{ ls(p) * ls(q) \} \text{merge}(r,p,q) \{ ls(r) \}$$

Prove that:

$$\begin{array}{l} \{ ls(p) \} \\ \text{mergesort}(r,p) \{ \\ \quad \text{if } p = \text{Nil} \text{ then } r := p; \\ \quad \text{else } \{ \\ \quad \quad \text{split}(q,p); \\ \quad \quad \text{mergesort}(q1,q) \parallel \text{mergesort}(p1,p); \\ \quad \quad \text{merge}(r,p1,q1) \\ \quad \} \\ \{ ls(r) \} \end{array}$$

Dynamic partitioning idioms

- Memory Managers, Thread Pools, Connection Pools;
- efficient Message Passing (copy avoiding);
- double-buffered I/O;
- many semaphore programs.

These idioms underlie much fundamental code: Microkernel OS designs, web servers, network packet processing, etc...

Old program design ideas, reflected in concurrent separation logic.

Question: are we done?

No: Reynolds counterexample

This logic is *inconsistent!* We can derive:

$$\{ x \mapsto _ \} \text{ with } r \text{ when true do skip } \{ F \}$$

where the resource $r()$ has invariant $RI = T$.



The triple states that the program diverges, while obviously it does not.

Exercise: can you find such derivation?

No: Reynolds counterexample

This logic is *inconsistent!* We can derive:

$$\{ x \mapsto _ \} \text{ with } r \text{ when true do skip } \{ F \}$$

where the resource r () has invariant $RI = T$.

Let one be a shorthand for $x \mapsto _$. From:

$$\frac{\frac{\{ T \} \text{ skip } \{ T \}}{\{ (\text{emp} \vee \text{one}) * \text{True} \} \text{ skip } \{ \text{emp} * \text{True} \}}}{\{ \text{emp} \vee \text{one} \} \text{ with } r \text{ when true do skip } \{ \text{emp} \}}$$

we can derive:

$$\frac{\frac{\frac{\{ \text{emp} \vee \text{one} \} \text{ with } \dots \{ \text{emp} \}}{\{ \text{emp} \} \text{ with } \dots \{ \text{emp} \}}}{\{ \text{emp} * \text{one} \} \text{ with } \dots \{ \text{emp} * \text{one} \}} \quad \frac{\{ \text{emp} \vee \text{one} \} \text{ with } \dots \{ \text{emp} \}}{\{ \text{one} \} \text{ with } \dots \{ \text{emp} \}}}{\frac{\{ \text{one} \} \text{ with } \dots \{ \text{one} \}}{\{ \text{one} \wedge \text{one} \} \text{ with } r \text{ when true do skip } \{ \text{emp} \wedge \text{one} \}} \quad \frac{\{ \text{one} \} \text{ with } \dots \{ \text{emp} \}}{\{ \text{one} \} \text{ with } r \text{ when true do skip } \{ \text{false} \}}}$$

What the Reynolds counterexample implies

Trouble if you have all of:

$$\frac{\{P_1\} C \{Q_1\} \quad \{P_2\} C \{Q_2\}}{\{P_1 \wedge P_2\} C \{Q_1 \wedge Q_2\}} \quad \frac{\{(P * RI_r) \wedge S\} C \{Q * RI_r\}}{\{P\} \text{ with } r \text{ when } S \text{ do } C \{Q\}}$$
$$\frac{\{P\} C \{Q\}}{\{P * R\} C \{Q * R\}}$$

The semantics of $P * Q$ is nondeterministic:

$$\exists h_1, h_2. \text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset \wedge h_1 \oplus h_2 = h \wedge (s, h_1) \models P \wedge (s, h_2) \models Q$$

The resource invariant T does not precisely nail down the storage owned by the resource; it is ambiguous. And the connective $*$ can be satisfied with different splittings.

What the Reynolds counterexample implies

Trouble if you have all of:

$$\frac{\{P_1\} \subset \{Q_1\} \quad \{P_2\} \subset \{Q_2\}}{\{P_1 \wedge P_2\} \subset \{Q_1 \wedge Q_2\}}$$

$$\frac{\{(P * RI_r) \wedge S\} \subset \{Q * RI_r\}}{\{P\} \text{ with } r \text{ when } S \text{ do } \subset \{Q\}}$$

If we can nail down the storage owned more precisely, perhaps we can get around this problem...

The semantics of $P * Q$ is nondeterministic:

$$\exists h_1, h_2. \text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset \wedge h_1 \oplus h_2 = h \wedge (s, h_1) \models P \wedge (s, h_2) \models Q$$

The resource invariant T does not precisely nail down the storage owned by the resource; it is ambiguous. And the connective $*$ can be satisfied with different splittings.

Precise predicates

A predicate P is *precise* if for every state, there is at most one substate satisfying it. Formally:

P is *precise* if for all s, h , there exists at most one $h' \sqsubseteq h$ where $s, h' \models P$.

Examples of *imprecise predicates*:

$\top \quad 10 \mapsto _ \vee 11 \mapsto _$

$$ls(x,y) = (x = y \wedge \text{empty}) \vee (\exists x' . x \mapsto x' * ls(x',y))$$

Examples of *precise predicates*:

$\text{empty} \quad 10 \mapsto _ \quad (\text{empty} \wedge \neg \text{full}) \vee (c \mapsto _ \wedge \text{full})$

$$ls(x,y) = \text{if } (x = y) \text{ then empty else } \exists x' . x \mapsto x' * ls(x',y)$$

A sound separation logic

Consider concurrent separation logic with the restriction that

all the resource invariants are precise.

Notation: let Γ define all the resources, and let $\text{inv}(\Gamma)$ *-conjunction of all the resource invariants.

Theorem:

If $\{ P \} \text{c} \{ Q \}$ is provable, every finite computation from a state satisfying $P * \text{inv} \Gamma$,

- is *error free*; and
- ends in a state satisfying $Q * \text{inv} \Gamma$.

A sound separation logic

Consider concurrent separation logic with the restriction that

all the resource invariants are precise.

Notation: let Γ define all the resources, and let $\text{inv}(\Gamma)$ *-conjunction of all the resource invariants.

Theorem:

If $\{ P \} c \{ Q \}$ is provable, every finite computation from a state satisfying $P * \text{inv } \Gamma$,

- is *error free*; and
- ends in a state satisfying $Q * \text{inv } \Gamma$.

We are not done: we must define what is a *computation* of c .



Soundness of concurrent separation logic

Brookes proof

Alternative proofs:

- Vafeiadis, *Concurrent separation logic and operational semantics*
- Hayman, Winskel, *Independence and concurrent separation logic* (LICS 06)



?

Disclaimer

The purpose of the following section is only to give an overview of the proof of soundness of concurrent separation logic, to characterise what *error-free* means operationally, and discover which is the role of precise predicates.

A simpler and more elegant proof (which unfortunately does not prove that programs verified using CSL do not have data-races) can be found here:

Vafeiadis, *Concurrent separation logic and operational semantics*.

<http://www.mpi-sws.org/~viktor/cslsound/>

Brookes's semantic analysis: the big picture

1. The denotation of a command is a set of traces:
 - traces captures all the interactions of the command with an arbitrary state;
 - the trace *abort* captures a *race*.
2. An LTS defines the action of a trace on a state:
 - the LTS goes to the state *abort* if the trace performs an access outside of the domain of the state, or if the trace is *abort*.
3. Command C is race-free from state s , if for all traces $\alpha \in [[C]]$, $\neg s \xrightarrow{\alpha} \text{abort}$.
4. Intuition (but we'll need one more idea):

if $\{P\} \subset \{Q\}$, then every *finite computation* of C from a state satisfying $P * \text{inv } \Gamma$, is race-free, and ends in a state satisfying $Q * \text{inv } \Gamma$.

1. denotation of commands

1. The denotation of a command is a set of traces

- traces captures all the interactions of a command with an arbitrary state:

$$\text{e.g. } \llbracket x := i+1 \rrbracket = \{ i=v . x:=v+1 \mid v \in \text{Value} \} .$$

- the trace *abort* captures a race:

$$\text{e.g. } \llbracket x := 1 \parallel x := 2 \rrbracket = \{ x:=1 . x:=2 , x:=2 . x:=1 , \text{abort} \}$$

Special care required to define the denotation of \parallel .

1. actions and traces

A command denotes a set of *traces*. A trace is a *sequence of actions*:

δ	idle
$i=v, i:=v$	read, write
$[v]=v', [v]:=v'$	lookup, update
$\text{alloc}(v, L), \text{disp}(v)$	allocate, dispose
$\text{try}(r), \text{acq}(r), \text{rel}(r)$	try, acquire, release
abort	race detected

λ ranges over actions. A trace can be finite or infinite. α, β range over traces.

Concatenation of traces is defined modulo:

$$\alpha.\delta.\beta = \alpha.\beta$$

$$\alpha.\text{abort}.\beta = \alpha.\text{abort}$$

1. clauses (1)

$$\llbracket \mathbf{skip} \rrbracket = \{\delta\}$$

$$\llbracket i := e \rrbracket = \{\rho \ i := v \mid (\rho, v) \in \llbracket e \rrbracket\}$$

$$\llbracket c_1; c_2 \rrbracket = \llbracket c_1 \rrbracket \llbracket c_2 \rrbracket$$

$$\llbracket \mathbf{if} \ b \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \rrbracket = \llbracket b \rrbracket_{true} \llbracket c_1 \rrbracket \cup \llbracket b \rrbracket_{false} \llbracket c_2 \rrbracket$$

$$\llbracket \mathbf{while} \ b \ \mathbf{do} \ c \rrbracket = (\llbracket b \rrbracket_{true} \llbracket c \rrbracket)^* \llbracket b \rrbracket_{false} \cup (\llbracket b \rrbracket_{true} \llbracket c \rrbracket)^\omega$$

$$\llbracket i := [e] \rrbracket = \{\rho \ [v] = v' \ i := v' \mid (\rho, v) \in \llbracket e \rrbracket\}$$

$$\llbracket i := \mathbf{cons}(E) \rrbracket = \{\rho \ \mathit{alloc}(l, L) \ i := l \mid (\rho, L) \in \llbracket E \rrbracket\}$$

$$\llbracket [e] := e' \rrbracket = \{\rho \ \rho' \ [v] := v' \mid (\rho, v) \in \llbracket e \rrbracket \ \& \ (\rho', v') \in \llbracket e' \rrbracket\}$$

$$\llbracket \mathbf{dispose} \ e \rrbracket = \{\rho \ \mathit{disp} \ l \mid (\rho, l) \in \llbracket e \rrbracket\}$$



sequential constructs



pointer operations

1. clauses (2)

$$\llbracket \mathbf{with} \ r \ \mathbf{when} \ b \ \mathbf{do} \ c \rrbracket = \mathit{wait}^* \ \mathit{enter} \cup \ \mathit{wait}^\omega$$

synchronisation

where

$$\mathit{wait} = \{ \mathit{try} \ r \} \cup \ \mathit{acq} \ r \ \llbracket b \rrbracket_{\mathit{false}} \ \mathit{rel} \ r$$

$$\mathit{enter} = \mathit{acq} \ r \ \llbracket b \rrbracket_{\mathit{true}} \ \llbracket c \rrbracket \ \mathit{rel} \ r$$

$$\llbracket c_1 \parallel c_2 \rrbracket = \llbracket c_1 \rrbracket_{\{\}} \parallel_{\{\}} \llbracket c_2 \rrbracket$$

parallel composition

Key ideas:

- 1) processes start with no resources;
- 2) resources are mutually exclusive;
- 3) races produce abort.

1. clauses (3)

$$\llbracket c_1 \parallel c_2 \rrbracket = \llbracket c_1 \rrbracket_{\{\}} \parallel_{\{\}} \llbracket c_2 \rrbracket$$

We rely on an auxiliary operator on traces:

- it builds all the traces obtained by interleaving the actions of the two threads
- in doing so, it keeps track of the resources allocated by each thread, and looks for data races.

Key ideas:

- 1) processes start with no resources;
- 2) resources are mutually exclusive;
- 3) races produce abort.

1. (3) resource enabling

Let A_1 and A_2 be sets of resources.

What a process can do depends on its resources and those of its environment.
For that, we define the *resource enabling relation*:

$$(A_1, A_2) \xrightarrow{\lambda} (A'_1, A_2)$$

Intuition: a process holding the resources A_1 can do λ in an environment that holds the resources A_2 .

$$(A_1, A_2) \xrightarrow{acq\ r} (A_1 \cup \{r\}, A_2) \quad \text{if } r \notin A_1 \cup A_2$$

$$(A_1, A_2) \xrightarrow{rel\ r} (A_1 - \{r\}, A_2) \quad \text{if } r \in A_1$$

$$(A_1, A_2) \xrightarrow{\lambda} (A_1, A_2) \quad \text{if } r \neq acq\ r, rel\ r$$

1. (3) interference and interleaving

Two actions *interfere* if one write to a variable or a cell used by the other:

$$\lambda_1 \bowtie \lambda_2 \text{ iff } \text{free}(\lambda_1) \cap \text{writes}(\lambda_2) \neq \{\} \text{ or } \text{free}(\lambda_2) \cap \text{writes}(\lambda_1) \neq \{\}$$

We can then define (fair, resource sensitive, race-detecting) interleaving:

$$\alpha_{A_1} \parallel_{A_2} \epsilon = \{\alpha \mid (A_1, A_2) \xrightarrow{\alpha} \cdot\}$$

$$\epsilon_{A_1} \parallel_{A_2} \alpha = \{\alpha \mid (A_2, A_1) \xrightarrow{\alpha} \cdot\}$$

$$(\lambda_1 \alpha_1)_{A_1} \parallel_{A_2} (\lambda_2 \alpha_2) =$$

$$\{\lambda_1 \beta \mid (A_1, A_2) \xrightarrow{\lambda_1} (A'_1, A_2) \ \& \ \beta \in \alpha_{1A'_1} \parallel_{A_2} (\lambda_2 \alpha_2)\}$$

$$\cup \{\lambda_2 \beta \mid (A_2, A_1) \xrightarrow{\lambda_2} (A'_2, A_1) \ \& \ \beta \in (\lambda_1 \alpha_1)_{A_1} \parallel_{A'_2} \alpha_2\}$$

$$\cup \{\text{abort} \mid \lambda_1 \bowtie \lambda_2\}$$

1. Examples

$$[[x := 1 \parallel y := 1]] = \{ x:=1 y:=1, y:=1 x:=1 \}$$

$$[[x := 1 \parallel x := 1]] = \{ x:=1 x:=1, \text{abort} \}$$

$$[[\text{with } r \text{ do } x := 1]] = (\text{try } r)^* \text{acq } r \text{ } x:=1 \text{ rel } r \cup (\text{try } r)^\omega$$

$$[[\text{with } r \text{ do } x := 1 \parallel \text{with } r \text{ do } x := 2]] =$$

$$(\text{try } r)^* \text{acq } r \text{ } x:=1 \text{ rel } r \cup (\text{try } r)^\omega \quad \{\} \parallel \{\} \quad (\text{try } r)^* \text{acq } r \text{ } x:=2 \text{ rel } r \cup (\text{try } r)^\omega =$$

$$(\text{try } r)^* \text{acq } r (\text{try } r)^* x:=1 (\text{try } r)^* \text{rel } r (\text{try } r)^* \text{acq } r x:=2 \text{ rel } r$$

$$\cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=2 (\text{try } r)^* \text{rel } r (\text{try } r)^* \text{acq } r x:=1 \text{ rel } r$$

$$\cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=1 (\text{try } r)^* \text{rel } r (\text{try } r)^\omega$$

$$\cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=2 (\text{try } r)^* \text{rel } r (\text{try } r)^\omega$$

$$\cup (\text{try } r)^\omega$$

1. Examples

$$[[x := 1 \parallel y := 1]] = \{x:=1 y:=1, y:=1 x:=1\}$$

$$[[x := 1 \parallel x := 1]] = \{x:=1 x:=1, \text{abort}\}$$

$$[[\text{with } r \text{ do } x := 1]] = (\text{try } r)^* \text{acq } r \text{ } x:=1 \text{ rel } r \cup (\text{try } r)^\omega$$

$$[[\text{with } r \text{ do } x := 1 \parallel \text{with } r \text{ do } x := 2]] =$$

here acq r by
thread 2 is not
enabled

$$(\text{try } r)^* \text{acq } r \text{ } x:=1 \text{ rel } r \cup (\text{try } r)^* \text{acq } r \text{ } x:=2 \text{ rel } r \cup (\text{try } r)^\omega =$$

$$\begin{aligned} & (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=1 (\text{try } r)^* \text{rel } r (\text{try } r)^* \text{acq } r x:=2 \text{ rel } r \\ & \cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=2 (\text{try } r)^* \text{rel } r (\text{try } r)^* \text{acq } r x:=1 \text{ rel } r \\ & \cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=1 (\text{try } r)^* \text{rel } r (\text{try } r)^\omega \\ & \cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=2 (\text{try } r)^* \text{rel } r (\text{try } r)^\omega \\ & \cup (\text{try } r)^\omega \end{aligned}$$

1. Examples

$$[[x := 1 \parallel y := 1]] = \{ x:=1 y:=1, y:=1 x:=1 \}$$

$$[[x := 1 \parallel x := 1]] = \{ x:=1 x:=1, \text{abort} \}$$

$$[[\text{with } r \text{ do } x := 1]] = (\text{try } r)^* \text{acq } r \text{ } x:=1 \text{ rel } r \cup (\text{try } r)^\omega$$

$$[[\text{with } r \text{ do } x := 1 \parallel \text{with } r \text{ do } x := 2]]$$

here acq r by thread 2 is not enabled

now, acq r by thread 2 is enabled

$$(\text{try } r)^* \text{acq } r \text{ } x:=1 \text{ rel } r \cup (\text{try } r)^\omega \cup (\text{try } r)^* \text{acq } r \text{ } x:=2 \text{ rel } r \cup (\text{try } r)^\omega =$$

$$\begin{aligned} & (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=1 (\text{try } r)^* \text{rel } r (\text{try } r)^* \text{acq } r x:=2 \text{ rel } r \\ & \cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=2 (\text{try } r)^* \text{rel } r (\text{try } r)^* \text{acq } r x:=1 \text{ rel } r \\ & \cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=1 (\text{try } r)^* \text{rel } r (\text{try } r)^\omega \\ & \cup (\text{try } r)^* \text{acq } r (\text{try } r)^* x:=2 (\text{try } r)^* \text{rel } r (\text{try } r)^\omega \\ & \cup (\text{try } r)^\omega \end{aligned}$$

2. The action of a trace on a state

The state is store + heap + resource:

$$(s, h, A)$$

- global store: $s : \text{var} \rightarrow \text{value}$;
- global heap: $h : \text{loc} \rightarrow \text{value}$;
- resources A held by the process.

Actions cause *state change*, and either end in a new state, or abort.

$$(s, h, A) \xRightarrow{\alpha} (s', h', A')$$

$$(s, h, A) \xRightarrow{\alpha} \mathbf{abort}$$

2. the LTS that relates states and actions (1)

$$(s, h, A) \xRightarrow{\delta} (s, h, A)$$

$$(s, h, A) \xRightarrow{i=v} (s, h, A) \quad \text{if } (i, v) \in s$$

$$(s, h, A) \xRightarrow{i:=v} ([s \mid i : v], h, A) \quad \text{if } i \in \text{dom } s$$

$$(s, h, A) \xRightarrow{[v]=v'} (s, h, A) \quad \text{if } (v, v') \in h$$

$$(s, h, A) \xRightarrow{[v]:=v'} (s, [h \mid v : v'], A) \quad \text{if } v \in \text{dom } h$$

$$(s, h, A) \xRightarrow{\text{alloc}(v, [v_0, \dots, v_n])} (s, [h \mid v : v_0, \dots, v + n : v_n], A) \\ \text{if } v, v + 1, \dots, v + n \notin \text{dom } h$$

$$(s, h, A) \xRightarrow{\text{disp } v} (s, h \setminus v, A) \quad \text{if } v \in \text{dom } h$$

$$(s, h, A) \xRightarrow{\text{acq } r} (s, h, A \cup \{r\}) \quad \text{if } r \notin A$$

$$(s, h, A) \xRightarrow{\text{rel } r} (s, h, A - \{r\}) \quad \text{if } r \in A$$

$$(s, h, A) \xRightarrow{\text{try } r} (s, h, A)$$

2. the LTS that relates states and actions (2)

if $i \notin \text{dom } s$

$$(s, h, A) \xRightarrow{i=v} \mathbf{abort}$$

$$(s, h, A) \xRightarrow{i:=v} \mathbf{abort}$$

$$(s, h, A) \xRightarrow{\mathbf{abort}} \mathbf{abort}$$

$$\mathbf{abort} \xRightarrow{\lambda} \mathbf{abort}$$

if $v \notin \text{dom } h$

$$(s, h, A) \xRightarrow{[v]=v'} \mathbf{abort}$$

$$(s, h, A) \xRightarrow{[v]:=v'} \mathbf{abort}$$

$$(s, h, A) \xRightarrow{\mathbf{disp } v} \mathbf{abort}$$

A global computation is an *executable sequence of actions*:

$$(s, h, A) \xRightarrow{\alpha} (s', h', A')$$

$$(s, h, A) \xRightarrow{\alpha} \mathbf{abort}$$

3. Error freedom

Definition: a command c is *error-free* if from (s,h) iff

forall $\alpha \in [[c]]$. $\neg ((s,h,\{\}) \xrightarrow{\alpha} \text{abort})$.

Example:

`dispose x || dispose y`

is error-free from all the states s such that $\neg(s(x) = s(y)) \wedge s(x), s(y) \in \text{dom}(h)$.

3.

Defin

Exam

dis

is er



4. A theorem?

It would be natural to define validity of $\{ P \} c \{ Q \}$ as:

Theorem:

$\{ P \} c \{ Q \}$ if every *finite computation* of c from a state satisfying $P * \text{inv } \Gamma$,

1) is error free,

2) ends in a state satisfying $Q * \text{inv } \Gamma$.

Proof: It is natural to proceed by induction on the derivation of $\{P\} c \{Q\}$. But... can you prove the case where c is $c_1 \parallel c_2$ and the last rule is the rule for parallel composition?

4. A theorem? Not yet!

It would be natural to define validity of $\{ P \} c \{ Q \}$ as:

Theorem:

$\{ P \} c \{ Q \}$ if every *finite computation* of c from a state satisfying $P * \text{inv } \Gamma$,

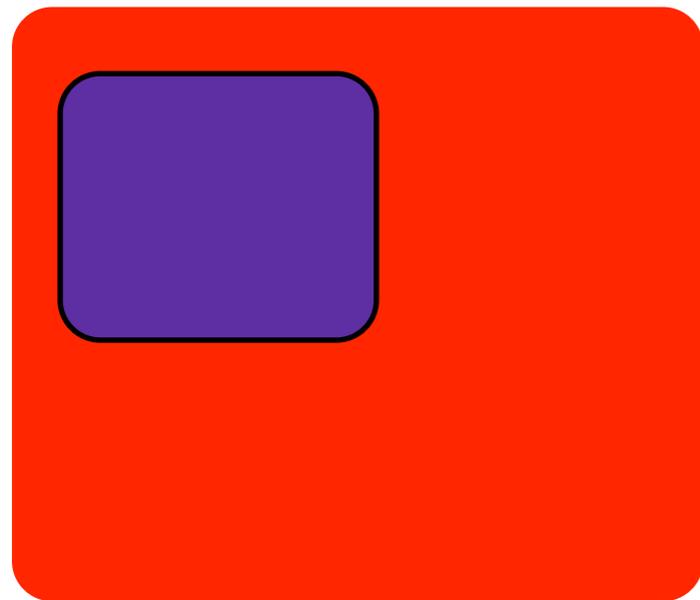
1) is error free,

2) ends in a state satisfying $Q * \text{inv } \Gamma$.

Proof: It is natural to proceed by induction on the derivation of $\{P\} c \{Q\}$. But... can you prove the case where c is $c_1 \parallel c_2$ and the last rule is the rule for parallel composition?

NO! This definition is not compositional: finite computations of c look at c in its entirety, and do not give enough informations about the computations performed by c_1 and c_2 .

Compositionality



heap

$C1 \parallel C2$

imagine C1 owns the **blue** part of the heap...

- C1 behaviour defines the evolution of the **blue** part of the heap;
- but all the red part is not constrained at all by C1, and might change under the influence of C2;
- if the semantics of C1 is defined in terms of the whole heap, it is tricky to derive it from the behaviour of $C1 \parallel C2$...

Idea: define the semantics of the thread only in terms of the heap it owns!

Local computations

Idea: keep track of the *local state* of each thread. The local state is defined by:

$$(s, h, A)$$

subject the condition

$$\text{dom}(s) \cap \text{owned}(\Gamma) = \text{owned}(\Gamma|A).$$

Local store only contains protected variables for which the process has resources.

A process starts with only non-critical data in its local state:

- local state grows when resource is acquired;
- local state shrinks when resource is released;
- error if program breaks design rules.

We can define another LTS, that captures the *local effects*.

Local effects: acquire and release

The rule for $\text{acq } r$ imports into the local state the part of the stack and of the heap protected by r .

The heap h' is uniquely determined because the invariant R is precise.

$$(s, h, A) \xrightarrow[\Gamma]{\text{acq } r} (s \cdot s', h \cdot h', A \cup \{r\})$$

$$\text{if } r(X):R \in \Gamma$$

$$\text{and } s \perp s', h \perp h', \text{dom } s' = X,$$

$$(s \cdot s', h') \models R$$

$$(s, h, A) \xrightarrow[\Gamma]{\text{rel } r} (s \setminus X, h - h', A - \{r\})$$

$$\text{if } r(X):R \in \Gamma$$

$$h' \subseteq h, (s, h') \models R$$

Similarly for $\text{rel } r$.

Local effects: other transitions

All the other transitions are inherited from the global semantics. E.g. (excerpt):

$$(s, h, A) \xrightarrow[\Gamma]{\delta} (s, h, A)$$

$$(s, h, A) \xrightarrow[\Gamma]{i=v} (s, h, A) \quad \text{if } i \in \text{dom } s$$

$$(s, h, A) \xrightarrow[\Gamma]{i:=v} ([s \mid i : v], h, A) \\ \text{if } i \in \text{dom } s - \text{free}(\Gamma \setminus A)$$

(in the last rule, the extra condition ensures that the variable being updated is does not belong to a resource).

Local effects: abort transitions

Two new abort rules (plus all the cases as in the global semantics):

- cannot update a variable protected by a resource not owned.

$$(s, h, A) \xrightarrow[\Gamma]{i:=v} \mathbf{abort}$$

if $i \in \text{free}(\Gamma \setminus A)$ or $i \notin \text{dom } s$

- when releasing a resource, the associated invariant must hold.

$$(s, h, A) \xrightarrow[\Gamma]{rel r} \mathbf{abort}$$

*if $r(X):R \in \Gamma$
and $\forall h' \subseteq h. (s, h') \models \neg R$*

Local computations

$$(s, h, A) \xrightarrow[\Gamma]{\alpha} (s', h', A')$$

$$(s, h, A) \xrightarrow[\Gamma]{\alpha} \mathbf{abort}$$

Local computation captures what a *thread* sees of a computation.

Assumes that the environment:

- respects the resource rules;
- interferes only on synchronisation.

A local computation of put || (get ; dispose y)

$$\Gamma = \text{buf}(c, \text{full}) : (\text{full} \wedge c \mapsto _) \vee (\neg \text{full} \wedge \mathbf{emp})$$

$$([x : v, y : _], [v : _], \{\})$$

$$\frac{\text{acq buf}}{\Gamma} \rightarrow ([x : v, y : _, \text{full} : \text{false}, c : _], [v : _], \{\text{buf}\})$$

$$\frac{\text{full}=\text{false put } v}{\Gamma} \rightarrow ([x : v, y : _, \text{full} : \text{true}, c : v], [v : _], \{\text{buf}\})$$

$$\frac{\text{rel buf}}{\Gamma} \rightarrow ([x : v, y : _], [], \{\})$$

$$\frac{\text{acq buf}}{\Gamma} \rightarrow ([x : v, y : _, \text{full} : \text{true}, c : v], [v : _], \{\text{buf}\})$$

$$\frac{\text{full}=\text{true get } v}{\Gamma} \rightarrow ([x : v, y : v, \text{full} : \text{false}, c : v], [v : _], \{\text{buf}\})$$

$$\frac{\text{rel buf}}{\Gamma} \rightarrow ([x : v, y : v], [v : _], \{\})$$

$$\frac{y=v \text{ disp } v}{\Gamma} \rightarrow ([x : v, y : v], [], \{\})$$

This can be decomposed into local computations of put and of get; dispose y...

The local computations of put and get || dispose y

put

$$\begin{aligned} & ([x : v], [v : -], \{\}) \\ \frac{acq\ buf}{\Gamma} & \rightarrow ([x : v, full : false, c : -], [v : -], \{buf\}) \\ \frac{full=false\ put\ v}{\Gamma} & \rightarrow ([x : v, full : true, c : v], [v : -], \{buf\}) \\ \frac{rel\ buf}{\Gamma} & \rightarrow ([x : v], [], \{\}) \end{aligned}$$

get || dispose y

$$\begin{aligned} & ([y : -], [], \{\}) \\ \frac{acq\ buf}{\Gamma} & \rightarrow ([y : -, full : true, c : v], [v : -], \{buf\}) \\ \frac{full=true\ get\ v}{\Gamma} & \rightarrow ([y : v, full : false, c : v], [v : -], \{buf\}) \\ \frac{rel\ buf}{\Gamma} & \rightarrow ([y : v], [v : -], \{\}) \\ \frac{y=v\ disp\ v}{\Gamma} & \rightarrow ([y : v], [], \{\}) \end{aligned}$$

Validity

$\{ P \} C \{ Q \}$ is *valid* if every *finite local computation* of C from a state satisfying $P * \text{inv } \Gamma$, is 1) error free and 2) ends in a state satisfying $Q * \text{inv } \Gamma$.

Theorem: all provable formulas are valid.

Proof: uses local states and local effects, shows that each rule preserves validity, for parallel uses the parallel lemma:

- a local computation of $C_1 \parallel C_2$ decomposes into local computations of C_1 and C_2 ;
- A local error of $C_1 \parallel C_2$ is caused by a local error of C_1 or C_2 (not by interference);
- A successful local computation of $C_1 \parallel C_2$ is consistent with all successful local computations of C_1 and C_2 .

Local vs. global

1. Soundness shows that *provable formulas are valid*;
2. *Validity* refers to *local* computations.

Need to connect local computations with conventional notions: global state, traditional partial correctness.

Theorem: Suppose $\alpha \in \llbracket c \rrbracket$, $h = h_1 \cdot h_2$, $(s, h_2) \models \text{inv}(\Gamma)$.

1. If $(s, h) \xRightarrow{\alpha} \mathbf{abort}$ then $(s \setminus \text{owned } \Gamma, h_1) \xrightarrow[\Gamma]{\alpha} \mathbf{abort}$

2. If $(s, h) \xRightarrow{\alpha} (s', h')$ then $(s \setminus \text{owned } \Gamma, h_1) \xrightarrow[\Gamma]{\alpha} (s'_1, h'_1)$ where

$$s'_1 = s' \setminus \text{owned } \Gamma \quad \text{and} \quad \exists h'_2. h' = h'_1 \cdot h'_2 \ \& \ (s', h'_2) \models \text{inv}(\Gamma)$$

Corollary: validity implies error freedom

$\{ P \} c \{ Q \}$ if every *finite computation* of c from a state satisfying $P * \text{inv } \Gamma$,

1) is error free,

2) ends in a state satisfying $Q * \text{inv } \Gamma$.

Many concurrent separation logics?

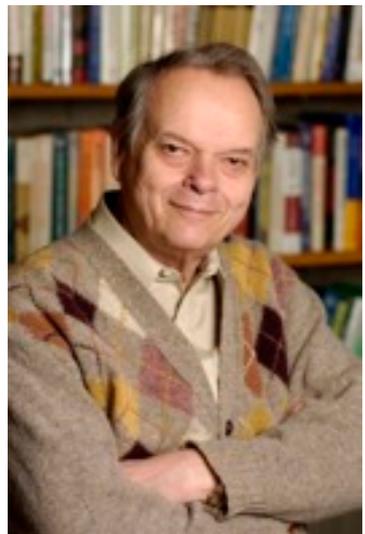
The logic presented here is not entirely realistic, and is not as expressive as one might hope/desire/expect.



Several variants have been proposed, including:

- Gotsman, Berdine, Cook, Rinetzky and Sagiv (APLAS 07)
- ... *plenty of other logics...*
- Hobor, Appel and me (ESOP 08)





Thanks to:

Stephen Brookes

John Reynolds

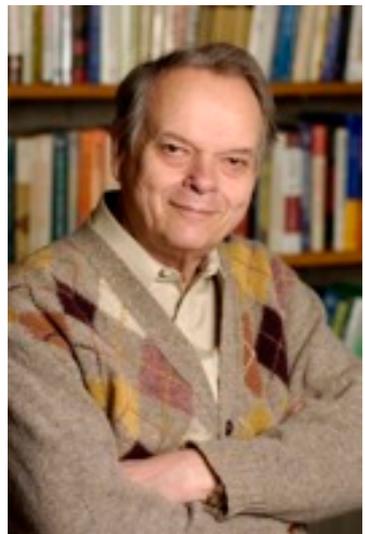
Tony Hoare

Edgser Dijkstra

Peter O'Hearn

Per Brinch Hansen

Exercise: associate each picture with its owner...



Thanks to:

Stephen Brookes

John Reynolds

Tony Hoare

Edgser Dijkstra

Peter O'Hearn

Per Brinch Hansen

Exercise: associate each picture with its owner...

References:

Peter O'Hearn, *Resources, concurrency and local reasoning*;

Stephen Brookes, *A semantics for concurrent separation logic*;

Viktor Vafeiadis, *Concurrent separation logic and operational semantics*

all available from <http://moscova.inria.fr/~zappa/teaching/mpri/2010/> .

Next lecture:

can we reason about racy programs?