

Concurrent Separation Logic

Mike Dodds

slides: Matthew J. Parkinson, Alexey Gotsman

Concurrency

Concurrent:

“Running together in space, as parallel lines; going on side by side, as proceedings; occurring together, as events or circumstances; existing or arising together; conjoint”

- Oxford English Dictionary

Programming language

$C ::= \dots \mid C \parallel C \mid \dots$

Motivation

- **Concurrency is hard:**

“If you can get away with it, avoid using threads. Threads can be difficult to use, and they make programs harder to debug.”

Java Sun Tutorial “Threads and Swing”

- **Multi-core means concurrency everywhere!**

Testing is hard

“Testing concurrent software is hard. Even simple tests require invoking methods from multiple threads and worrying about issues such as timeouts and deadlock. Unlike in sequential programs, many failures are rare, probabilistic events and numerous factors can mask potential errors.”

JavaOne Technical session

Testing is hard

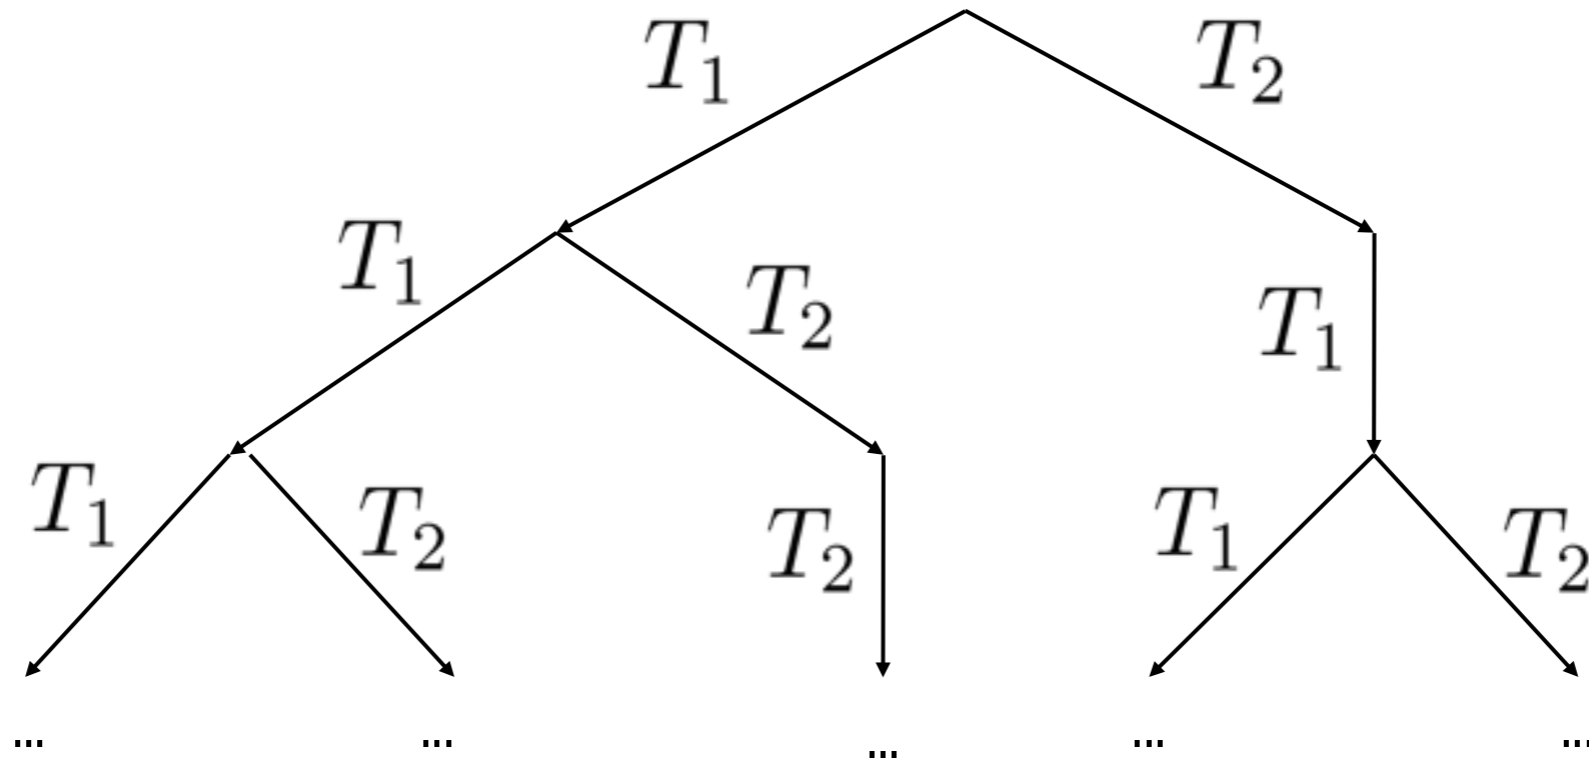
“Testing concurrent software is hard. Even simple tests require invoking methods from multiple threads and worrying about issues such as timeouts and deadlock. Unlike in sequential programs, many failures are rare, probabilistic events and numerous factors can mask potential errors.”

JavaOne Technical session

Verification to the rescue?

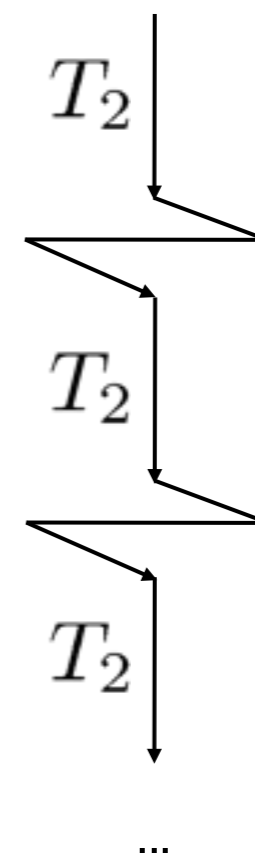
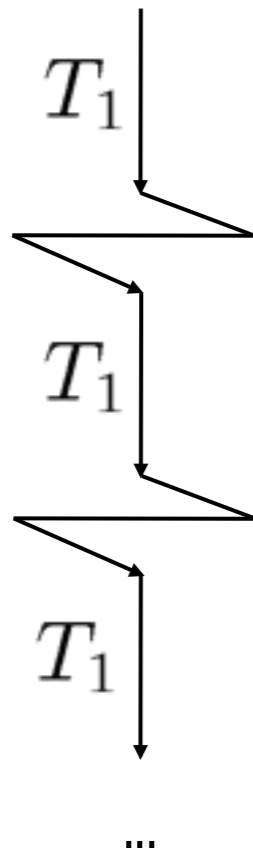
Verifying concurrent programs is hard

Have to consider all possible interleavings:



Thread-modular reasoning

- Considers every thread in isolation under some assumption on its environment:



Thread-modular reasoning

- Considers every thread in isolation under some assumption on its environment:



Thread-modular reasoning

- Considers every thread in isolation under some assumption on its environment:

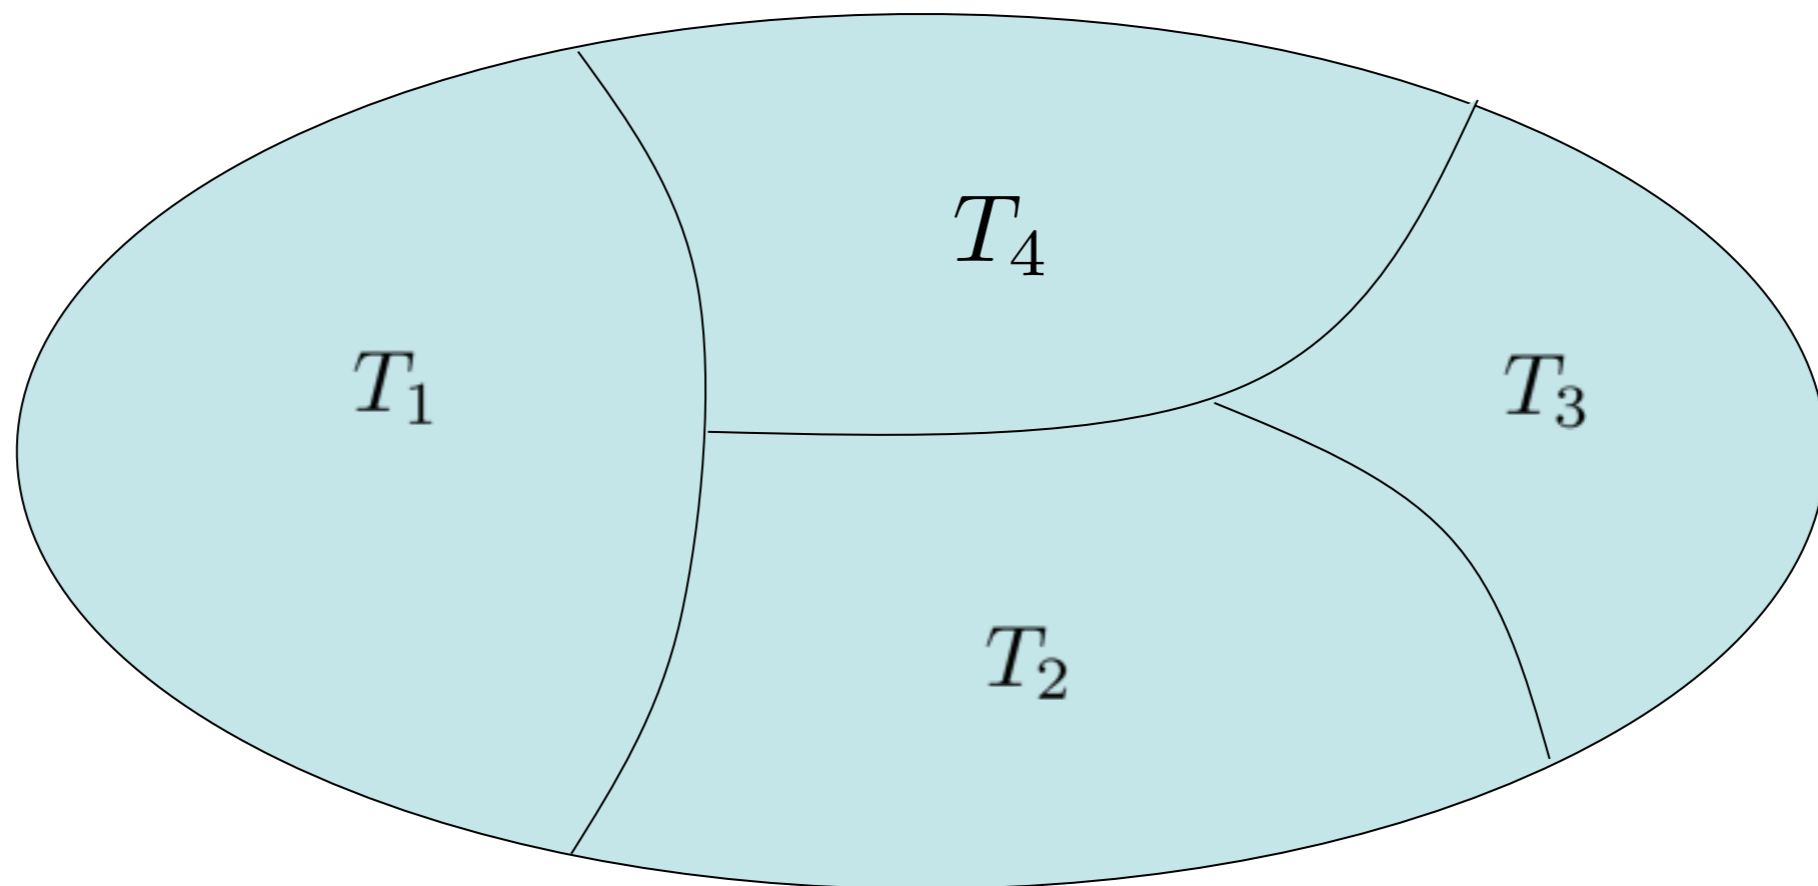


- Avoids direct reasoning about all interleavings

Disjoint Concurrency

Disjoint concurrency

- Language with parallel composition: $C_1 \parallel C_2$
- Every thread operates on its own part of the heap:



Parallel proof rule

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

variables used in C_1 , P_1 and Q_1 not modified by C_2 ;
variables used in C_2 , P_2 and Q_2 not modified by C_1

Parallel proof rule

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

variables used in C_1 , P_1 and Q_1 not modified by C_2 ;
variables used in C_2 , P_2 and Q_2 not modified by C_1

Parallel proof rule

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

variables used in C_1 , P_1 and Q_1 not modified by C_2 ;
variables used in C_2 , P_2 and Q_2 not modified by C_1

Parallel proof rule

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

variables used in C_1 , P_1 and Q_1 not modified by C_2 ;
variables used in C_2 , P_2 and Q_2 not modified by C_1

- Remember semantics of triples: C_1 accesses only the memory in P_1 and the one it allocates itself
- No way to mess up the heap owned by C_2 !

Example

$$\begin{array}{ccc} & \{ x \mapsto _ * y \mapsto _ \} & \\ \{ x \mapsto _ \} & & \{ y \mapsto _ \} \\ [x] := 3 & \parallel & [y] := 4 \\ \{ x \mapsto 3 \} & & \{ y \mapsto 4 \} \\ & \{ x \mapsto 3 * y \mapsto 4 \} & \end{array}$$

Parallel Dispose tree

```
struct Tree {
```

```
    Tree *Left;
```

```
    Tree *Right; }
```

```
disposetree(Tree *x) {
```

```
    if (x != NULL) {
```

```
        i = x->Left;
```

```
        j = x->Right;
```

```
        (disposetree(i) || disposetree(j) || free(x));
```

```
    }
```

```
}
```

Parallel Dispose tree

```
struct Tree {  
    Tree *Left;  
    Tree *Right; }  
}
```

$$\text{Tree}(x) = (x = \text{NULL} \wedge \text{emp}) \vee (\exists i, j. x \mapsto i, j * \text{Tree}(i) * \text{Tree}(j))$$

```
{ tree(x) }
```

```
disposetree(Tree *x) {
```

```
    if (x != NULL) {
```

```
        i = x->Left;
```

```
        j = x->Right;
```

```
        (disposetree(i) || disposetree(j) || free(x));
```

```
    }
```

```
} { emp }
```

Parallel Dispose tree

$$\text{Tree}(x) = (x = \text{NULL} \wedge \text{emp}) \vee \\ (\exists i, j. x \mapsto i, j * \text{Tree}(i) * \text{Tree}(j))$$

{ tree(x) ∧ x ≠ NULL }

{ ∃i,j. tree(i) * tree(j) * x ↦ i,j }

i = x->Left;

{ ∃j. tree(i) * tree(j) * x ↦ i,j }

j = x->Right;

{ tree(i) * tree(j) * x ↦ i,j }

(disposetree(i) || disposetree(j) || free(x));

{ emp }

Example

$$\{ \text{tree}(i) * \text{tree}(j) * x \mapsto i,j \}$$
$$\{ \text{tree}(i) \} \quad \{ \text{tree}(j) \} \quad \{ x \mapsto i,j \}$$
$$\text{disposetree}(i) \parallel \text{disposetree}(j) \parallel \text{dispose } x \quad .$$
$$\{ \text{emp} \} \quad \{ \text{emp} \} \quad \{ \text{emp} \}$$
$$\{ \text{emp} * \text{emp} * \text{emp} \}$$
$$\{ \text{emp} \}$$

Can we verify these?

$\{ \text{emp} \}$
 $x := \text{new};$
 $z := \text{new};$
 $[x] := 4 \parallel [z] := 5;$
 $\{ x \mapsto 4 * z \mapsto 5 \}$

$\{ \text{emp} \}$
 $x := \text{new};$
 $[x] := 4 \parallel [x] := 5;$
 $\{ x \mapsto _ \}$

$\{ \text{emp} \}$
 $x := 4 \parallel x := 5;$
 $\{ \text{emp} \}$

$\{ y = x + 1 \}$
 $x := 4 \parallel y := y + 1;$
 $\{ y = x + 2 \}$

Merge sort

```
mergesort(x, n)
  if n > 1 then
    local m in
      m := n/2;
      mergesort(x,m) || mergesort(x+m,n-m);
      merge(x,m,n-m)
```

Merge sort

{ array(x,n) }

mergesort(x, n)

{ sorted_array(x,n) }

{ sorted_array(x,m) * sorted_array(x+m,n) }

merge(x,m,n)

{ sorted_array(x,m+n) }

Merge sort

```
{ array(x,n) }  
mergesort(x, n)  
  if n > 1 then  
    local m in  
    m := n/2;  
    mergesort(x,m) || mergesort(x+m,n-m);  
    merge(x,m,n-m)  
{ sorted_array(x,n) }
```

Concurrent Separation Logic

Multiple access

How do we verify a program where several threads want access to the same memory? e.g.

`[x] := 43` || `[x] := 47`

We protect shared values with **locks**

Multiple access

How do we verify a program where several threads want access to the same memory? e.g.

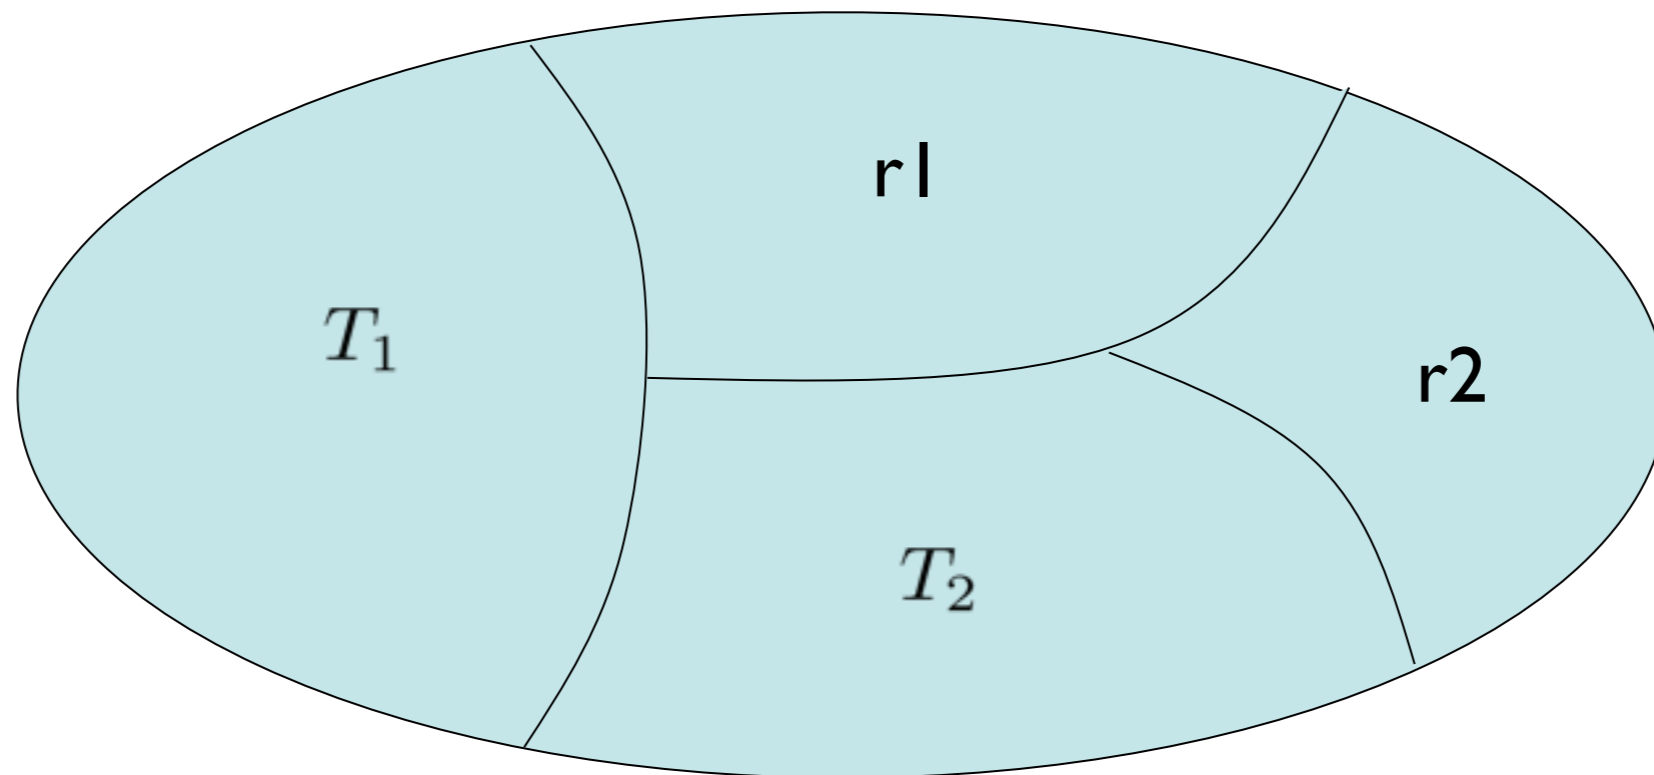
`[x] := 43` || `[x] := 47`

We protect shared values with **locks**



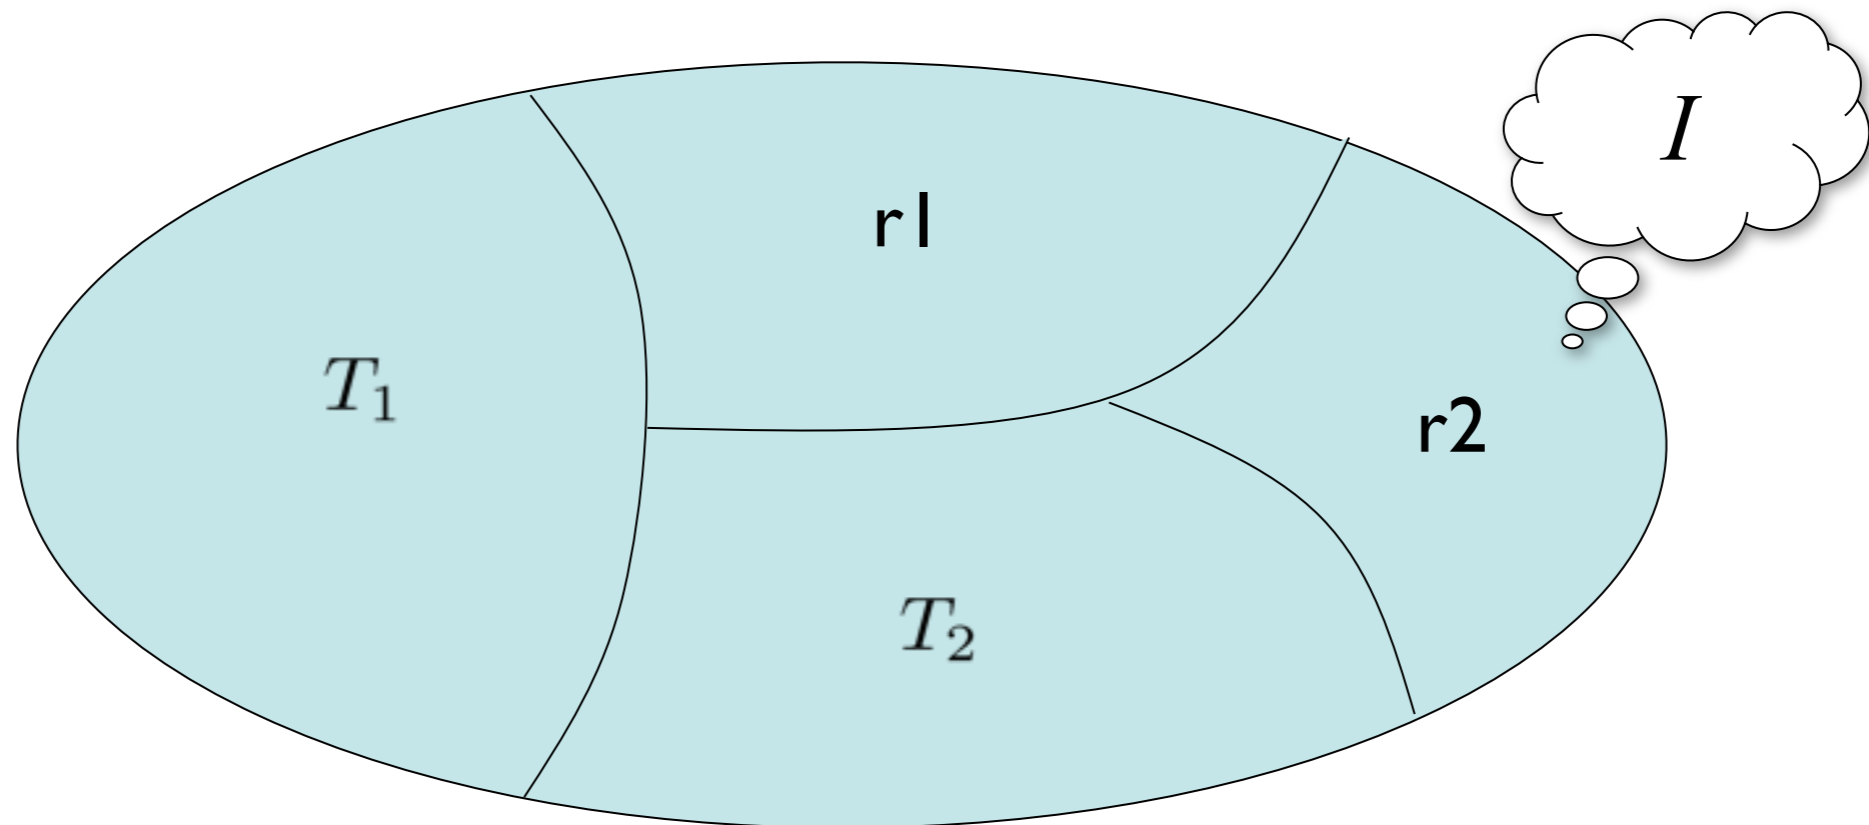
Reasoning principle

Separation property: at any time, the state of the program can be partitioned into that owned by each thread and each free lock



Reasoning principle

Separation property: at any time, the state of the program can be partitioned into that owned by each thread and each free lock



Assign a **resource invariant** I to every lock r : describes the part of the heap protected by the lock

Programming language

$C ::= \dots \mid \text{resource } r \text{ in } C \mid \text{with } r \text{ when } B \text{ in } C \mid \dots$

Resource Rule

$$\frac{\Delta, r : I \vdash \{P\} C \{Q\}}{\Delta \vdash \{P * I\} \text{resource } r \text{ in } C \{Q * I\}}.$$

Lock Rule

$$\frac{\Delta \vdash \{ (P * I) \wedge B \} C \{ Q * I \}}{\Delta, r : I \vdash \{ P \} \text{ with } r \text{ when } B \text{ in } C \{ Q \}}.$$

Caveat: side-conditions

There are subtle variable side-conditions used to allow locks to refer to global variables.

Each variable is either associate to

- a single thread; or
- a single lock.

It can then only be modified and used in assertions by the thread, or while the thread holds the associate lock.

Binary Semaphore

We can encode a semaphore as a critical region

$P(s) = \text{with } r_s \text{ when } s=1 \text{ do } s := 0$

$V(s) = \text{with } r_s \text{ when } s=0 \text{ do } s := 1$

Resource invariant

$(s=0 \wedge \text{emp}) \vee (s=1 \wedge Q)$

Initially,

$s=0$

Example

{ emp }
P(s)
[x] := 43
V(s)
{ emp }



{ emp }
P(s)
[x] := 47
V(s)
{ emp }

Example

{ emp }

P(s)

{ x ↦ _ }

[x] := 43

{ x ↦ _ }

V(s)

{ emp }

Example

{ emp }

P(s)

{ x ↦ _ }

[x] := 43

{ x ↦ _ }

V(s)

{ emp }

$$\frac{\{ \text{emp} * (I_s \wedge s=1) \} s := 0 \{ x \mapsto _ * I_s \}}{\{ \text{emp} \} \text{ with } r_s \text{ when } s=1 \text{ do } s:=0 \{ x \mapsto _ \}}$$

Example

{ emp }

P(s)

{ x ↦ _ }

[x] := 43

{ x ↦ _ }

V(s)

{ emp }

$$\frac{\{ \text{emp} * (I_s \wedge s=1) \} s := 0 \{ x \mapsto _ * I_s \}}{\{ \text{emp} \} \text{ with } r_s \text{ when } s=1 \text{ do } s:=0 \{ x \mapsto _ \}}$$
$$\frac{\{ x \mapsto _ * (I_s \wedge s=0) \} s := 1 \{ \text{emp} * I_s \}}{\{ x \mapsto _ \} \text{ with } r_s \text{ when } s=0 \text{ do } s:=1 \{ \text{emp} \}}$$

One place buffer

full := false

with buff when full do
 full := false
 y := c
dispose y

x := new
with buff when \neg full do
 full := true;
 c := x;

One place buffer

full := false

$\{ (\text{full} \wedge c \mapsto _) \vee (\neg \text{full} \wedge \text{emp}) \}$

Resource
Invariant

with buff when full do
 full := false
 y := c
dispose y

x := new
with buff when \neg full do
 full := true;
 c := x;

One place buffer

$\{ (\text{full} \wedge c \mapsto _) \vee (\neg\text{full} \wedge \text{emp}) \}$

with buff when full do

$\{ \text{full} \wedge c \mapsto _ \}$

full := false

y := c

$\{ (\neg\text{full} \wedge \text{emp})$

$* y \mapsto _ \}$

$\{ y \mapsto _ \}$

dispose y

x := new

$\{ x \mapsto _ \}$

with buff when $\neg\text{full}$ do

$\{ (\neg\text{full} \wedge \text{emp})$

$* x \mapsto _ \}$

full := true;

c := x;

$\{ \text{full} \wedge c \mapsto _ \}$

Ownership is in the eye of the assessor

full := false

with buff when full do
 full := false
 y := c

x := new
with buff when \neg full do
 full := true;
 c := x;
dispose x

Can we verify the following?

Ownership is in the eye of the assertor

full := false

{ emp }

Resource
Invariant

with buff when full do
 full := false
 y := c

x := new
with buff when ¬full do
 full := true;
 c := x;
dispose x

Can we verify the following?

next time:
Semantics