

Modular Verification of a Non-Blocking Stack

Matthew Parkinson Richard Bornat Peter O'Hearn

Lecture for 15-818A4 Separation Logic

Henry DeYoung

April 25, 2011

Introduction

Concurrent Separation Logic

Guiding Principles:

- ▶ *Dijkstra*: Limit interference to rare moments of synchronization.
- ▶ *Separation Logic*: “Separate parts of a program that depend on separated resources can be dealt with independently.”

An Approach: Use units of mutual exclusion to synchronize.

- ▶ Inefficient — Threads often must wait for exclusive access.
- ▶ Is there an effective treatment of *non-blocking concurrency*?

3

Non-Blocking Concurrency

Threads attempt to make concurrent changes to shared data structures, looping and trying again if an attempt fails.

- ▶ Algorithms rely on atomic hardware operations
 - ▶ Atomic reads and writes to single-word store (heap?) locations
 - ▶ CAS (compare-and-set) instruction

$\text{CAS}(\ell, v_1, v_2)$: If $\ell \mapsto v_1$, then atomically mutate so $\ell \mapsto v_2$.
Otherwise, skip.

Isn't this just more finely-grained mutual exclusion?

- ▶ “An atomic access to the [shared] data structure is treated as a unit of mutual exclusion.”

4

Rely-Guarantee for Non-Blocking Concurrency

Each thread must specify

- ▶ properties which it *guarantees* to preserve, and
- ▶ properties which it *relies* on other threads to preserve.

Thread interference floods the proofs of clients!

- ▶ Every instruction in every thread must ensure the guarantees of its thread, because every specification must include those guarantees.

Can separation logic be used to contain thread interference?

5

Separation Logic and Non-Blocking Concurrency

- ▶ Shared data structure, equipped with an invariant
- ▶ Operations on data structure as procedures:
 - ▶ Non-blocking interaction through concurrent execution of procedures
 - ▶ Pre- and post-conditions of procedures do not involve invariant
 - ▶ Invariant is exposed within atomic operations in procedure bodies

“Even though the proofs [of the procedure bodies] may be horrible, their horrors are confined, and we can consider threads to be independent when outside those procedures.”

6

Concurrent Separation Logic, “Simplified”

Concurrent Separation Logic, Simplified

$$\Gamma; I \vdash \{Q\} C \{R\}$$

- ▶ Procedure specifications $\{Q'\} f(\vec{x}) \{R'\}$
- ▶ *Single* resource invariant for the shared data structure

$$\frac{\Gamma; I \vdash \{Q_1\} C_1 \{R_1\} \quad \Gamma; I \vdash \{Q_2\} C_2 \{R_2\}}{\Gamma; I \vdash \{Q_1 \star Q_2\} C_1 \parallel C_2 \{R_1 \star R_2\}}$$

Standard rule for parallel composition in concurrent separation logic.

The atomic Rule

$$\frac{\Gamma; \mathbf{emp} \vdash \{Q \star I\} C \{R \star I\}}{\Gamma; I \vdash \{Q\} \mathbf{atomic}\{C\} \{R\}}$$

C must execute in mutual exclusion with $\mathbf{atomic}\{\dots\}$ commands!

- ▶ Rely on serialization properties of the hardware.
- ▶ C may invoke at most one single-word read or write in the shared data structure, or must be a CAS instruction.

Special case of the CCR rule when there is only one resource.

- ▶ It follows that $\mathbf{atomic}\{\mathbf{atomic}\{\dots\}\}$ should not terminate.

Can one enforce the side condition more syntactically? E.g.,

$$\frac{\vdash \{Q \star I\} A \{R \star I\}}{\Gamma; I \vdash \{Q\} \mathbf{atomic}\{A\} \{R\}}$$

where A is an atomic hardware operation.

9

The module Rule

$$\frac{\Gamma; I \vdash \{Q_f\} C_f \{R_f\} \quad \Gamma, \{Q_f\} f(\vec{x}_f) \{R_f\}; \mathbf{emp} \vdash \{Q\} C \{R\}}{\Gamma; \mathbf{emp} \vdash \{Q \star I\} \mathbf{module} f(\vec{x}_f) = C_f \mathbf{in} C \{R \star I\}}$$

- ▶ Atomic instructions inside C_f can use the invariant — the rest of the program cannot.
- ▶ The procedure body itself is not executed atomically.
- ▶ Procedures cannot be recursive, apparently.
- ▶ “Extension to several procedures is obvious.”

10

The module Rule

$$\frac{\begin{array}{l} \Gamma; I \vdash \{Q_f\} C_f \{R_f\} \\ \Gamma; I \vdash \{Q_g\} C_g \{R_g\} \\ \Gamma, \{Q_f\} f(\vec{x}_f) \{R_f\}, \{Q_g\} g(\vec{x}_g) \{R_g\}; \mathbf{emp} \vdash \{Q\} C \{R\} \end{array}}{\Gamma; \mathbf{emp} \vdash \{Q \star I\} \text{ module } f(\vec{x}_f) = C_f \text{ and } g(\vec{x}_g) = C_g \text{ in } C \{R \star I\}}$$

- ▶ Atomic instructions inside C_f can use the invariant — the rest of the program cannot.
- ▶ The procedure body itself is not executed atomically.
- ▶ Procedures cannot be recursive, apparently.
- ▶ “Extension to several procedures is obvious.”

11

The module Rule, cont.

“A more general approach: the module rule can be derived from

$$\frac{\Gamma; I_1 \star I_2 \vdash \{Q\} C \{R\}}{\Gamma; I_1 \vdash \{Q \star I_2\} C \{R \star I_2\}} \quad \text{and} \quad \frac{\Gamma; I_1 \vdash \{Q\} C \{R\}}{\Gamma; I_1 \star I_2 \vdash \{Q\} C \{R\}} .”$$

Define $\text{module } f(\vec{x}_f) = C_f \text{ in } C$ as $\text{resource } r \text{ in } \text{let } f(\vec{x}_f) = C_f \text{ in } C$
with $r \notin \text{res}(C)$.

$$\frac{\Gamma; \mathbf{emp} \star I \vdash \{Q_f\} C_f \{R_f\} \quad \frac{\Gamma, \{Q_f\} C_f \{R_f\}; \mathbf{emp} \vdash \{Q\} C \{R\}}{\Gamma, \{Q_f\} C_f \{R_f\}; \mathbf{emp} \star I \vdash \{Q\} C \{R\}}}{\frac{\Gamma; \mathbf{emp} \star I \vdash \{Q\} \text{ let } f(\vec{x}_f) = C_f \text{ in } C \{R\}}{\Gamma; \mathbf{emp} \vdash \{Q \star I\} \text{ resource } r \text{ in let } f(\vec{x}_f) = C_f \text{ in } C \{R \star I\}}}$$

12

The module Rule, cont.

“A more general approach: the module rule can be derived from

$$\frac{\Gamma; l_1 \star l_2 \vdash \{Q\} C \{R\}}{\Gamma; l_1 \vdash \{Q \star l_2\} C \{R \star l_2\}} \quad \text{and} \quad \frac{\Gamma; l_1 \vdash \{Q\} C \{R\}}{\Gamma; l_1 \star l_2 \vdash \{Q\} C \{R\}} .”$$

“Note that these two rules make the standard frame rule derivable.”

$$\frac{\frac{\Gamma; l_1 \vdash \{Q\} C \{R\}}{\Gamma; l_1 \star l_2 \vdash \{Q\} C \{R\}}}{\Gamma; l_1 \vdash \{Q \star l_2\} C \{R \star l_2\}}$$

This seems odd because the bottom rule does not declare a resource.

13

The CAS Rule

$$\frac{\Gamma; l \vdash \{(Q \star e \mapsto e_2) \wedge e' = e_1\} C_1 \{R\} \quad \Gamma; l \vdash \{(Q \star e \mapsto e') \wedge e' \neq e_1\} C_2 \{R\}}{\Gamma; l \vdash \{Q \star e \mapsto e'\} \text{if } (\text{CAS}(e, e_1, e_2)) \text{ then } C_1 \text{ else } C_2 \{R\}}$$

$\text{CAS}(b; v, v_1, v_2)\{v'\} =$

$\{v \mapsto v'\}$

newvar v'' **in**

$v'' := [v];$

$b := (v'' = v_1);$

if b **then**

$[v] := v_2$

else

skip

$\{(b \wedge v \mapsto v_2 \wedge v' = v_1) \vee$
 $(\neg b \wedge v \mapsto v' \wedge v' \neq v_1)\}$

$\{Q \star e \mapsto e'\}$

newvar b **in**

$\text{CAS}(b; e, e_1, e_2)\{e'\}; \} \star Q$

$\{(b \wedge (Q \star e \mapsto e_2) \wedge e' = e_1) \vee$
 $(\neg b \wedge (Q \star e \mapsto e') \wedge e' \neq e_1)\}$

if b **then**

$\{(Q \star e \mapsto e_2) \wedge e' = e_1\}$

C_1

else

$\{(Q \star e \mapsto e') \wedge e' \neq e_1\}$

C_2

$\{R\}$

14

Heap Cell Permissions

Three types of permissions:

1. Read-only permissions
2. Write permissions
3. Existence permissions

This paper requires only that the invariant shares permissions with each thread — not that threads share permissions with each other.

Read-Only Permissions:

$$E \mapsto F \iff E \overset{r}{\mapsto} F \star E \overset{r}{\mapsto} F$$

A read permission allows a heap cell to be read.

$$E \overset{r}{\mapsto} F \star E \overset{r}{\mapsto} F' \Rightarrow F = F'$$

15

Heap Cell Permissions

Three types of permissions:

1. Read-only permissions
2. Write permissions
3. Existence permissions

This paper requires only that the invariant shares permissions with each thread — not that threads share permissions with each other.

Write and Existence Permissions:

$$E \mapsto F \iff E \overset{w}{\mapsto} F \star E \overset{e}{\mapsto} -$$

A write permission allows mutation of heap cells. An existence permission allows nothing, but ensures existence.

$$E \overset{e}{\mapsto} - \star E' \overset{e}{\mapsto} - \Rightarrow E \neq E'$$

16

Correctness of Michael's Algorithm

Case Study: Michael's Algorithm

Michael's algorithm implements a shared stack.

- ▶ Think storage allocator for single-cell heap records.
 - ▶ $\text{Stack} \approx \text{freelist}$, $\text{pop}() \approx \text{malloc}()$, and $\text{push}(x) \approx \text{free}(x)$
- ▶ `push` can fail because of interference from other threads.

Specifications of `pop()` and `push(x)`:

$$\{\mathbf{emp}\} \text{pop}() \{(ret \mapsto -) \vee (ret = \mathbf{nil} \wedge \mathbf{emp})\}$$
$$\{x \mapsto -\} \text{push}(x) \{(ret \wedge \mathbf{emp}) \vee (\neg ret \wedge x \mapsto -)\}$$

A Client of the Shared Stack

Interference between threads is confined to pop and push.

alloc() of memory manager:

- ▶ Uses (slower) system allocator if the stack is empty.
- ▶ Proof does not need to (directly) account for thread interference.

```
{emp}
alloc() {
  local y;
  {emp}
  y=pop();
  {(y = nil ∧ emp) ∨ y ↦ -}
  if (y==nil) {
    {emp}
    y=new();
    {y ↦ -}
  }
  {y ↦ -}
  return y;
}
{ret ↦ -}
```

19

push for a Shared Stack?

```
{b ↦ -}
1  push(b) {
2    local t,n;
3    while(true) {
4      t = TOP;
5      b->t1 = t;
6      if (CAS(&TOP, t, b)) // If TOP still matches b->t1,
7        break;           // then atomically set TOP to b.
8    }                   // Else, try pushing again.
9    return true;
10 }
    {ret ∧ emp}
```

20

pop for a Shared Stack?


```
1 pop() {  
2   local t,n;  
3   while(true) {  
4     t = TOP;  
5     if (t == nil)  
6       break;  
7     n = t->tl;  
8     if (CAS(&TOP, t, n)) // If TOP still matches t, then  
9       break;           // atomically set TOP to t->tl.  
10  }                    // Else, try popping again.  
11  return t;  
12 }
```



pop() is broken by the ABA problem!

21

The ABA Problem



22

Fixing the ABA Problem

Solution: Add a global array H of 'hazard pointers'.

```
1 pop() {
2   local t,n;
3   while(true) {
4     t = TOP;
5     if (t == nil)
6       break;
7     H[tid] = t;
8     if (t != TOP) continue;
9     n = t->t1;
10    if (CAS(&TOP, t, n))
11      break;
12  }
13  H[tid] = nil;
14  return t;
15 }
```

```
1 push(b) {
2   local t,n;
3   for (n=0; n<=THREADS; n++)
4     if (H[n] == b)
5       return false;
6   while(true) {
7     t = TOP;
8     b->t1 = t;
9     if (CAS(&TOP, t, b))
10      break;
11  }
12  return true;
13 }
```

(†) Between execution of `H[tid] = t;` and assignment of another value to `H[tid]`, the cell pointed to by `t` will not be removed from the stack and subsequently re-inserted.

23

Fixing the ABA Problem

Solution: Add a global array H of 'hazard pointers'.

```
1 pop() {
2   local t,n;
3   while(true) {
4     t = TOP;
5     if (t == nil)
6       break;
7     H[tid] = t;
8     if (t != TOP) continue;
9     n = t->t1;
10    if (CAS(&TOP, t, n))
11      break;
12  }
13  H[tid] = nil;
14  return t;
15 }
```

“Don’t try to re-insert t;
I’m trying to pop it!”

Check that t is still on top.

“I’m done popping t.”

(†) Between execution of `H[tid] = t;` and assignment of another value to `H[tid]`, the cell pointed to by `t` will not be removed from the stack and subsequently re-inserted.

24

Fixing the ABA Problem

Solution: Add a global array H of 'hazard pointers'.

Check that no other thread is trying to pop record b.

```
1  push(b) {
2    local t,n;
3    for (n=0; n<=THREADS; n++)
4      if (H[n] == b)
5        return false;
6    while(true) {
7      t = TOP;
8      b->t1 = t;
9      if (CAS(&TOP, t, b))
10       break;
11   }
12   return true;
13 }
```

- (†) Between execution of $H[tid] = t;$ and assignment of another value to $H[tid]$, the cell pointed to by t will not be removed from the stack and subsequently re-inserted.

25

Fixing the ABA Problem

Solution: Add a global array H of 'hazard pointers'.

```
1  pop() {
2    local t,n;
3    while(true) {
4      t = TOP;
5      if (t == nil)
6        break;
7      H[tid] = t;
8      if (t != TOP) continue;
9      n = t->t1;
10     if (CAS(&TOP, t, n))
11       break;
12   }
13   H[tid] = nil;
14   return t;
15 }
```

```
1  push(b) {
2    local t,n;
3    for (n=0; n<=THREADS; n++)
4      if (H[n] == b)
5        return false;
6    while(true) {
7      t = TOP;
8      b->t1 = t;
9      if (CAS(&TOP, t, b))
10       break;
11   }
12   return true;
13 }
```

- (†) Once a cell is inserted into the stack, its $t1$ value will not be altered until it is removed.

26

Auxiliary State to Express Key Properties

Q: How can we express the (\dagger) and (\ddagger) properties?

A: Use auxiliary state to track the status of hazard pointers.

- ▶ A hazard pointer can be in one of four states:

Unset $H[tid] = \text{nil}$.

Req $H[tid] = b$ and b has not been observed in the stack.

Tail(k) $H[tid] = b$ and b has been observed in the stack with tail k .

Left $H[tid] = b$ and b has been observed in the stack, but is now known not to be in the stack.

- ▶ If a hazard pointer is *Tail(k)* or *Left*, it shouldn't be pushed.

27

Adding Auxiliary Assignments to pop

```
1 pop() {
2   local t,n;
3   while(true) {
4     atomic{ t = TOP; }
5     if (t == nil) break;
6     atomic{ H[tid] = t; H'[tid] = Req; }   t not yet observed in stack.
7     atomic{ if (TOP != t) continue;
8             else H'[tid] = Tail(t->t1); }   H[tid] is observed in stack.
9     atomic{ n = t->t1;
10            if (H'[tid] != Tail(n))
11                H'[tid] = Left; }           If t->t1 changed, then t was
12            atomic{ if (CAS(&TOP, t, n)) break; }   popped ( $\ddagger$ ). So, t can't
13        }                                       have been re-inserted ( $\dagger$ ).
14    atomic{ H[tid] = nil; H'[tid] = Unset; }   H[tid] is no longer set.
15    return t;
16 }
```

Over-cautious: Test $TOP == t$ as surrogate for 't is somewhere in the stack'.

28

Adding Auxiliary Assignments to push

```
1  push(b) {
2    local t,n;
3    for (n=0; n<=THREADS; n++) {
4      atomic{
5        if (H[n] == b) { G[tid] = Unset; return false; }
6        G[tid] = NotHaz(b, {0,...,n});
7      }
8    }
9    while(true) {
10     atomic{ t = TOP; }
11     b->t1 = t;
12     atomic{
13       if (CAS(&TOP, t, b)) { G[tid] = Unset; break; }
14     }
15   }
16   return true;
17 }
```

“I give up on pushing b.”

“I’m trying to push b, and it’s safe for 0, ..., n.”

“I’m done pushing b.”

What if hazard pointers change after we’ve checked them?

- ▶ Can’t happen — A cell being pushed is not in the stack. So, during the push, no other thread can set its hazard pointer to that cell’s address.

29

Building Up the Invariant

Q: How does the stack constrain the hazard-pointer array h and hazard-status array h' ?

A:

- ▶ If y is in the stack and is hazard-pointered for thread i , then its status is either $Tail(z)$ or Req — it cannot be $Left$ or $Unset$.

$$\text{RNode}(y, z, h, h') \stackrel{\text{def}}{=} y \mapsto z \wedge \forall i \in T. ((h(i) = y) \Rightarrow (h'(i) = \text{Tail}(z) \vee h'(i) = \text{Req}))$$

- ▶ Apply this restriction to every element of the stack:

$$\text{RList}(y, h, h') \stackrel{\text{def}}{=} (y = \mathbf{nil} \wedge \mathbf{emp}) \vee \exists z. (\text{RNode}(y, z, h, h') \star \text{RList}(z, h, h'))$$

30

Building Up the Invariant, cont.

Q: How does G restrict the hazard-pointer and -status arrays?

A:

$$\text{HCons}(h, h') \stackrel{\text{def}}{=} \bigodot_{i \in T} \left(\begin{array}{l} G[i] \stackrel{r}{\mapsto} \text{Unset} \vee \\ \exists b, S. \left((G[i] \stackrel{r}{\mapsto} \text{NotHaz}(b, S) \star b \stackrel{e}{\mapsto} -) \right. \\ \left. \wedge \forall j \in S. (h(j) = b \Rightarrow h'(j) = \text{Req}) \right) \end{array} \right)$$

For each thread i , either:

- ▶ thread i is not pushing anything; or
- ▶ thread i is pushing some b and:
 - ▶ b is not a confirmed hazard for any threads checked so far in the for-loop, and
 - ▶ b is not in the stack and not being pushed by another thread.

31

Building Up the Invariant, cont.

Q: How does G restrict the hazard-pointer and -status arrays?

A:

$$\text{HCons}(h, h') \stackrel{\text{def}}{=} \bigodot_{i \in T} \left(\begin{array}{l} G[i] \stackrel{r}{\mapsto} \text{Unset} \vee \\ \exists b, S. \left((G[i] \stackrel{r}{\mapsto} \text{NotHaz}(b, S) \star b \stackrel{e}{\mapsto} -) \right. \\ \left. \wedge \forall j \in S. (h(j) = b \Rightarrow h'(j) = \text{Req}) \right) \end{array} \right)$$

- ▶ If thread i is pushing some b , the thread gives $b \stackrel{e}{\mapsto} -$ to the invariant and only keeps $b \stackrel{w}{\mapsto} -$ for itself. Thus, the thread loses total permission and it cannot begin to push a cell and then pass the cell to another thread that finishes pushing b .
- ▶ The invariant provides only read permissions for G . Thread i will have the other read permission for $G[i]$.

32

Finally, the Invariant

$$\text{Inv} \stackrel{\text{def}}{=} \exists h, h'. \left(\begin{array}{l} (\odot_{j \in T}. \mathbb{H}[j] \xrightarrow{r} h(j) \star \mathbb{H}'[j] \xrightarrow{r} h'(j)) \star \\ \text{HCons}(h, h') \star \text{RList}(\text{TOP}, h, h') \star \\ \forall i \in T. (h(i) = \mathbf{nil} \Leftrightarrow h'(i) = \text{Unset}) \end{array} \right)$$

- ▶ The arrays \mathbb{H} and \mathbb{H}' represent h and h' , respectively.
- ▶ The invariant provides only read permissions for \mathbb{H} and \mathbb{H}' . Thread j will have the other read permission for $\mathbb{H}[j]$ and $\mathbb{H}'[j]$.
- ▶ h and h' are well-behaved with respect to the array G .
- ▶ h and h' are well-behaved with respect to the stack.
- ▶ h and h' are consistent for unset hazard pointers.

Should this instead be \wedge ?

Some Details of the Proof

Some Details of pop

```

{ (H[tid] ↦ t * H'[tid] ↦ Req) ∧ t ≠ nil }
atomic{
  { (H[tid] ↦ t * H'[tid] ↦ Req * Inv) ∧ t ≠ nil }
  if (TOP != t)
    { H[tid] ↦ - * H'[tid] ↦ - * Inv }
    continue;
  else
    { (H[tid] ↦ TOP * H'[tid] ↦ Req * Inv) ∧ t ≠ nil ∧ TOP = t }
    H'[tid] = Tail(t->t1);
    { H[tid] ↦ t * H'[tid] ↦ Tail(-) * Inv } ∧ t ≠ nil }
}
{ (H[tid] ↦ t * H'[tid] ↦ Tail(-)) ∧ t ≠ nil }

```

Recall

$$\text{Inv} \stackrel{\text{def}}{=} \exists h, h'. \left(\left(\bigodot_{j \in T}. H[j] \overset{r}{\mapsto} h(j) * H'[j] \overset{r}{\mapsto} h'(j) \right) * \right. \\ \left. \text{HCons}(h, h') * \text{RList}(\text{TOP}, h, h') * \right. \\ \left. \forall i \in T. (h(i) = \text{nil} \Leftrightarrow h'(i) = \text{Unset}) \right)$$

It suffices to show:

$$h(\text{tid}) = t \wedge \text{RNode}(t, n, h, h') \Rightarrow \text{RNode}(t, n, h, h'[\text{tid} \mapsto \text{Tail}(n)])$$

35

Some Details of pop

```

{ (H[tid] ↦ t * H'[tid] ↦ Req) ∧ t ≠ nil }
atomic{
  { (H[tid] ↦ t * H'[tid] ↦ Req * Inv) ∧ t ≠ nil }
  if (TOP != t)
    { H[tid] ↦ - * H'[tid] ↦ - * Inv }
    continue;
  else
    { (H[tid] ↦ TOP * H'[tid] ↦ Req * Inv) ∧ t ≠ nil ∧ TOP = t }
    H'[tid] = Tail(t->t1);
    { H[tid] ↦ t * H'[tid] ↦ Tail(-) * Inv } ∧ t ≠ nil }
}
{ (H[tid] ↦ t * H'[tid] ↦ Tail(-)) ∧ t ≠ nil }

```

Recall

$$\text{Inv} \stackrel{\text{def}}{=} \exists h, h'. \left(\left(\bigodot_{j \in T}. H[j] \overset{r}{\mapsto} h(j) * H'[j] \overset{r}{\mapsto} h'(j) \right) * \right. \\ \left. \text{HCons}(h, h') * \text{RList}(\text{TOP}, h, h') * \right. \\ \left. \forall i \in T. (h(i) = \text{nil} \Leftrightarrow h'(i) = \text{Unset}) \right)$$

It suffices to show:

$$h(\text{tid}) = t \wedge \text{HCons}(h, h') * t \mapsto n \\ \Rightarrow \text{HCons}(h, h'[\text{tid} \mapsto \text{Tail}(n)]) * t \mapsto n$$

36

Some Details of pop, cont.

We need to show:

$$\begin{aligned} h(\text{tid}) = t \wedge \text{HCons}(h, h') \star t \mapsto n \\ \Rightarrow \text{HCons}(h, h'[\text{tid} \mapsto \text{Tail}(n)]) \star t \mapsto n \end{aligned}$$

Recall

$$\text{HCons}(h, h') \stackrel{\text{def}}{=} \bigodot_{i \in T} \left(\begin{array}{l} G[i] \xrightarrow{r} \text{Unset} \vee \\ \exists b, S. \left((G[i] \xrightarrow{r} \text{NotHaz}(b, S) \star b \xrightarrow{e} -) \right. \\ \left. \wedge \forall j \in S. (h(j) = b \Rightarrow h'(j) = \text{Req}) \right) \end{array} \right)$$

For each thread i , there are two cases:

1. If $G[i] = \text{Unset}$, then setting $h'(\text{tid}) = \text{Tail}(n)$ is safe for i .
2. Otherwise, we have $t \xrightarrow{w} n \star t \xrightarrow{e} - \star b \xrightarrow{e} - \Rightarrow t \neq b$,
and setting $h'(\text{tid}) = \text{Tail}(n)$ is safe for i .

37

More Details of pop

```

{ (H[tid]  $\xrightarrow{r}$  t  $\star$  H'[tid]  $\xrightarrow{r}$  Tail(-))  $\wedge$  t  $\neq$  nil }
atomic{
  { (H[tid]  $\xrightarrow{r}$  t  $\star$  H'[tid]  $\xrightarrow{r}$  Tail(-)  $\star$  Inv)  $\wedge$  t  $\neq$  nil }
  n = t->t1;
  if (H'[tid]  $\neq$  Tail(n))
    H'[tid] = Left;
}
{ (H[tid]  $\xrightarrow{r}$  t  $\star$  (H'[tid]  $\xrightarrow{r}$  Tail(n)  $\vee$  H'[tid]  $\xrightarrow{r}$  Left))  $\wedge$  t  $\neq$  nil }

```

Case: If t is in the stack, i.e., $\text{RList}(\text{TOP}, h, h') \Rightarrow t \mapsto - \star \text{true}$,
then Inv gives us total permission for t .

Moreover, the test $H'[\text{tid}] \neq \text{Tail}(n)$ will fail because

$$\begin{aligned} h(\text{tid}) = t \wedge h'(\text{tid}) = \text{Tail}(a) \\ \wedge \text{RList}(\text{TOP}, h, h') \wedge (t \mapsto x \star \text{true}) \\ \Rightarrow x = a \end{aligned}$$

because $(p \star q) \wedge e \hookrightarrow e' \Rightarrow ((p \wedge e \hookrightarrow e') \star q) \vee (p \star (q \wedge e \hookrightarrow e'))$.

38

More Details of pop

```

{(H[tid] ↦ t * H'[tid] ↦ Tail(-)) ∧ t ≠ nil}
atomic{
  {(H[tid] ↦ t * H'[tid] ↦ Tail(-) * Inv) ∧ t ≠ nil}
  n = t->t1;
  if (H'[tid] != Tail(n))
    H'[tid] = Left;
}
{(H[tid] ↦ t * (H'[tid] ↦ Tail(n) ∨ H'[tid] ↦ Left)) ∧ t ≠ nil}

```

Case: Otherwise, $RList(TOP, h, h') \Rightarrow \neg(t \mapsto - * \mathbf{true})$, and we do not have any permissions for t .

- ▶ Therefore, we must use a racy read:

$$\overline{\{\mathbf{emp}\} \ n = [t] \ \{\mathbf{emp}\}}$$

This reflects the algorithm's optimism in reading $t \rightarrow t1$.

- ▶ But, how do we justify setting $h'(tid) = Left$?

39

More Details of pop, cont.

We may set the status to *Left* because

- ▶ it follows from distributivity of $\wedge \neg(t \mapsto - * \mathbf{true})$ over $*$ that

$$\begin{aligned}
 h(tid) &= t \wedge RList(TOP, h, h') \wedge \neg(t \mapsto - * \mathbf{true}) \\
 &\Rightarrow RList(TOP, h, h'[tid \mapsto Left])
 \end{aligned}$$

- ▶ and because

$$\begin{aligned}
 h(tid) &= t \wedge h'(tid) = Tail(-) \wedge HCons(h, h') \\
 &\Rightarrow HCons(h, h'[tid \mapsto Left])
 \end{aligned}$$

Recall

$$HCons(h, h') \stackrel{\text{def}}{=} \bigodot_{i \in T} \left(G[i] \overset{r}{\mapsto} Unset \vee \left(\exists b, S. \left((G[i] \overset{r}{\mapsto} NotHaz(b, S) * b \overset{e}{\mapsto} -) \wedge \forall j \in S. (h(j) = b \Rightarrow h'(j) = Req) \right) \right) \right)$$

40

Still More Details of pop

```

{H[tid] ↦ t * (H'[tid] ↦ Tail(n) ∨ H'[tid] ↦ Left) ∧ t ≠ nil}
atomic{
  {(H[tid] ↦ t * H'[tid] ↦ Tail(n) * Inv ∧ t ≠ nil) ∨
   (H[tid] ↦ t * H'[tid] ↦ Left * Inv ∧ t ≠ nil)}
  if (CAS(&TOP, t, n))
    {H[tid] ↦ - * H'[tid] ↦ - * (t ≐ nil ∨ t ↦ -) * Inv}
    break;
  {H[tid] ↦ - * H'[tid] ↦ - * Inv}
}
{H[tid] ↦ - * H'[tid] ↦ -}

```

Case: Status is *Left*.

- The CAS will fail because:

$$H[tid] \mapsto t * H'[tid] \mapsto \text{Left} * \text{Inv} \Rightarrow \text{TOP} \neq t$$

The postcondition follows immediately.

41

Still More Details of pop

```

{H[tid] ↦ t * (H'[tid] ↦ Tail(n) ∨ H'[tid] ↦ Left) ∧ t ≠ nil}
atomic{
  {(H[tid] ↦ t * H'[tid] ↦ Tail(n) * Inv ∧ t ≠ nil) ∨
   (H[tid] ↦ t * H'[tid] ↦ Left * Inv ∧ t ≠ nil)}
  if (CAS(&TOP, t, n))
    {H[tid] ↦ - * H'[tid] ↦ - * (t ≐ nil ∨ t ↦ -) * Inv}
    break;
  {H[tid] ↦ - * H'[tid] ↦ - * Inv}
}
{H[tid] ↦ - * H'[tid] ↦ -}

```

Case: Status is *Tail(n)* and $\text{TOP} \neq t$.

- The CAS will fail, and the postcondition follows immediately.

42

Still More Details of pop

```

{H[tid] ↦ t * (H'[tid] ↦ Tail(n) ∨ H'[tid] ↦ Left) ∧ t ≠ nil}
atomic{
  {(H[tid] ↦ t * H'[tid] ↦ Tail(n) * Inv ∧ t ≠ nil) ∨
   (H[tid] ↦ t * H'[tid] ↦ Left * Inv ∧ t ≠ nil)}
  if (CAS(&TOP, t, n))
    {H[tid] ↦ - * H'[tid] ↦ - * (t ≐ nil ∨ t ↦ -) * Inv}
    break;
  {H[tid] ↦ - * H'[tid] ↦ - * Inv}
}
{H[tid] ↦ - * H'[tid] ↦ -}

```

Case: Status is $Tail(n)$ and $TOP = t$.

- The postcondition follows from

$$\begin{aligned}
 h(tid) = t \wedge h'(tid) = Tail(n) \wedge RList(t, h, h') \wedge t \neq \mathbf{nil} \\
 \Rightarrow t \mapsto n * RList(n, h, h')
 \end{aligned}$$

43

Some Details of push

```

{G[tid] ↦ NotHaz(b, T) * b ↦ t}
atomic{
  {G[tid] ↦ NotHaz(b, T) * b ↦ t * Inv}
  if (CAS(&TOP, t, b)) {
    G[tid] = Unset;
    {G[tid] ↦ Unset * Inv}
    break;
  }
  {G[tid] ↦ NotHaz(b, T) * b ↦ - * Inv}
}
{G[tid] ↦ NotHaz(b, T) * b ↦ -}

```

- If the CAS succeeds, then b is not a confirmed hazard because

$$\begin{aligned}
 HCons(h, h') * G[tid] \mapsto NotHaz(b, T) \\
 \Rightarrow \forall i \in T. (h(i) = b \Rightarrow h'(i) = Req).
 \end{aligned}$$

Recall that

$$HCons(h, h') \stackrel{\text{def}}{=} \bigodot_{i \in T} \left(G[i] \mapsto Unset \vee \left(\exists b, S. \left((G[i] \mapsto NotHaz(b, S) * b \mapsto -) \wedge \forall j \in S. (h(j) = b \Rightarrow h'(j) = Req) \right) \right) \right)$$

44

Some Details of push

```
{G[tid]  $\overset{r}{\mapsto}$  NotHaz(b, T) * b  $\overset{w}{\mapsto}$  t}
atomic{
  {G[tid]  $\overset{r}{\mapsto}$  NotHaz(b, T) * b  $\overset{w}{\mapsto}$  t * Inv}
  if (CAS(&TOP, t, b)) {
    G[tid] = Unset;
    {G[tid]  $\overset{r}{\mapsto}$  Unset * Inv}
    break;
  }
  {G[tid]  $\overset{r}{\mapsto}$  NotHaz(b, T) * b  $\overset{w}{\mapsto}$  - * Inv}
}
```

- Therefore, it is safe to add b to the stack:

$$b \mapsto t * \text{RList}(t, h, h') * (\forall i \in T. (h(i) = b \Rightarrow h'(i) = \text{Req})) \\ \Rightarrow \text{RList}(b, h, h')$$

Recall that

$$\text{RNode}(b, t, h, h') \stackrel{\text{def}}{=} b \mapsto t \wedge \\ \forall i \in T. ((h(i) = b) \Rightarrow \\ (h'(i) = \text{Tail}(t) \vee h'(i) = \text{Req}))$$

Conclusion

Conclusion

Key Points:

- ▶ Formally contained inter-thread interference to operations on shared, non-blocking data structure.
- ▶ “Even though the proofs [of the procedure bodies] may be horrible, their horrors are confined.”

Future Work:

- ▶ Liveness properties — lock/wait-freedom?
- ▶ Marriage of separation logic and rely/guarantee or temporal logic?