

Phantom Monitors: A Simple Foundation for Modular Proofs of Fine-Grained Concurrent Programs

Christian J. Bell, Mohsen Lesani, Adam Chlipala,
Stephan Boyer, Gregory Malecha, Peng Wang

MIT CSAIL

Goal: verification of concurrent client programs...

```
global stack jobs
fun client(input)
  foreach v in input
    push jobs (In v)
  results = []
  while length(results) < length(input)
    x = pop jobs
    case x = Some (Out v): results := v :: results
    case x = Some (In v): push jobs (In v)
  return results
fun worker(compute)
  while true
    x = pop jobs
    case x = Some (In v): push jobs (Out (compute v))
    case x = Some (Out v): push jobs (Out v)
fun testCase()
  for 1..4 do
    fork worker(factorial)
  client_pid = fork client([1,2,3,4])
  results = join client_pid
  assert (sum results = 33)
```

Goal: ... and verification of fine-grained concurrent datastructures

```
fun push(head, v)
  node = alloc 2
  write node.item v
  while true
    oldHead = read head
    write node.next oldHead
    if cas head oldHead node = 1 then
      return

fun pop(head)
  while true
    oldHead = read head
    if oldHead = 0 then
      return None
    else
      newHead = read oldHead.next
      if cas head oldHead newHead = 1 then
        v = read oldHead.item
      return (Some v)
```

General Challenges

1. What does the concurrent program logic look like?
 - **abstraction**: high level
 - **local reasoning**: modular/manageable proofs
 - **generality**: can we prove real & interesting programs?
2. End-to-end verification
 - what does the **machine code** actually do?
 - can we **trust** our program logic?
3. Verification framework development
 - how do we **quickly test** new ideas?

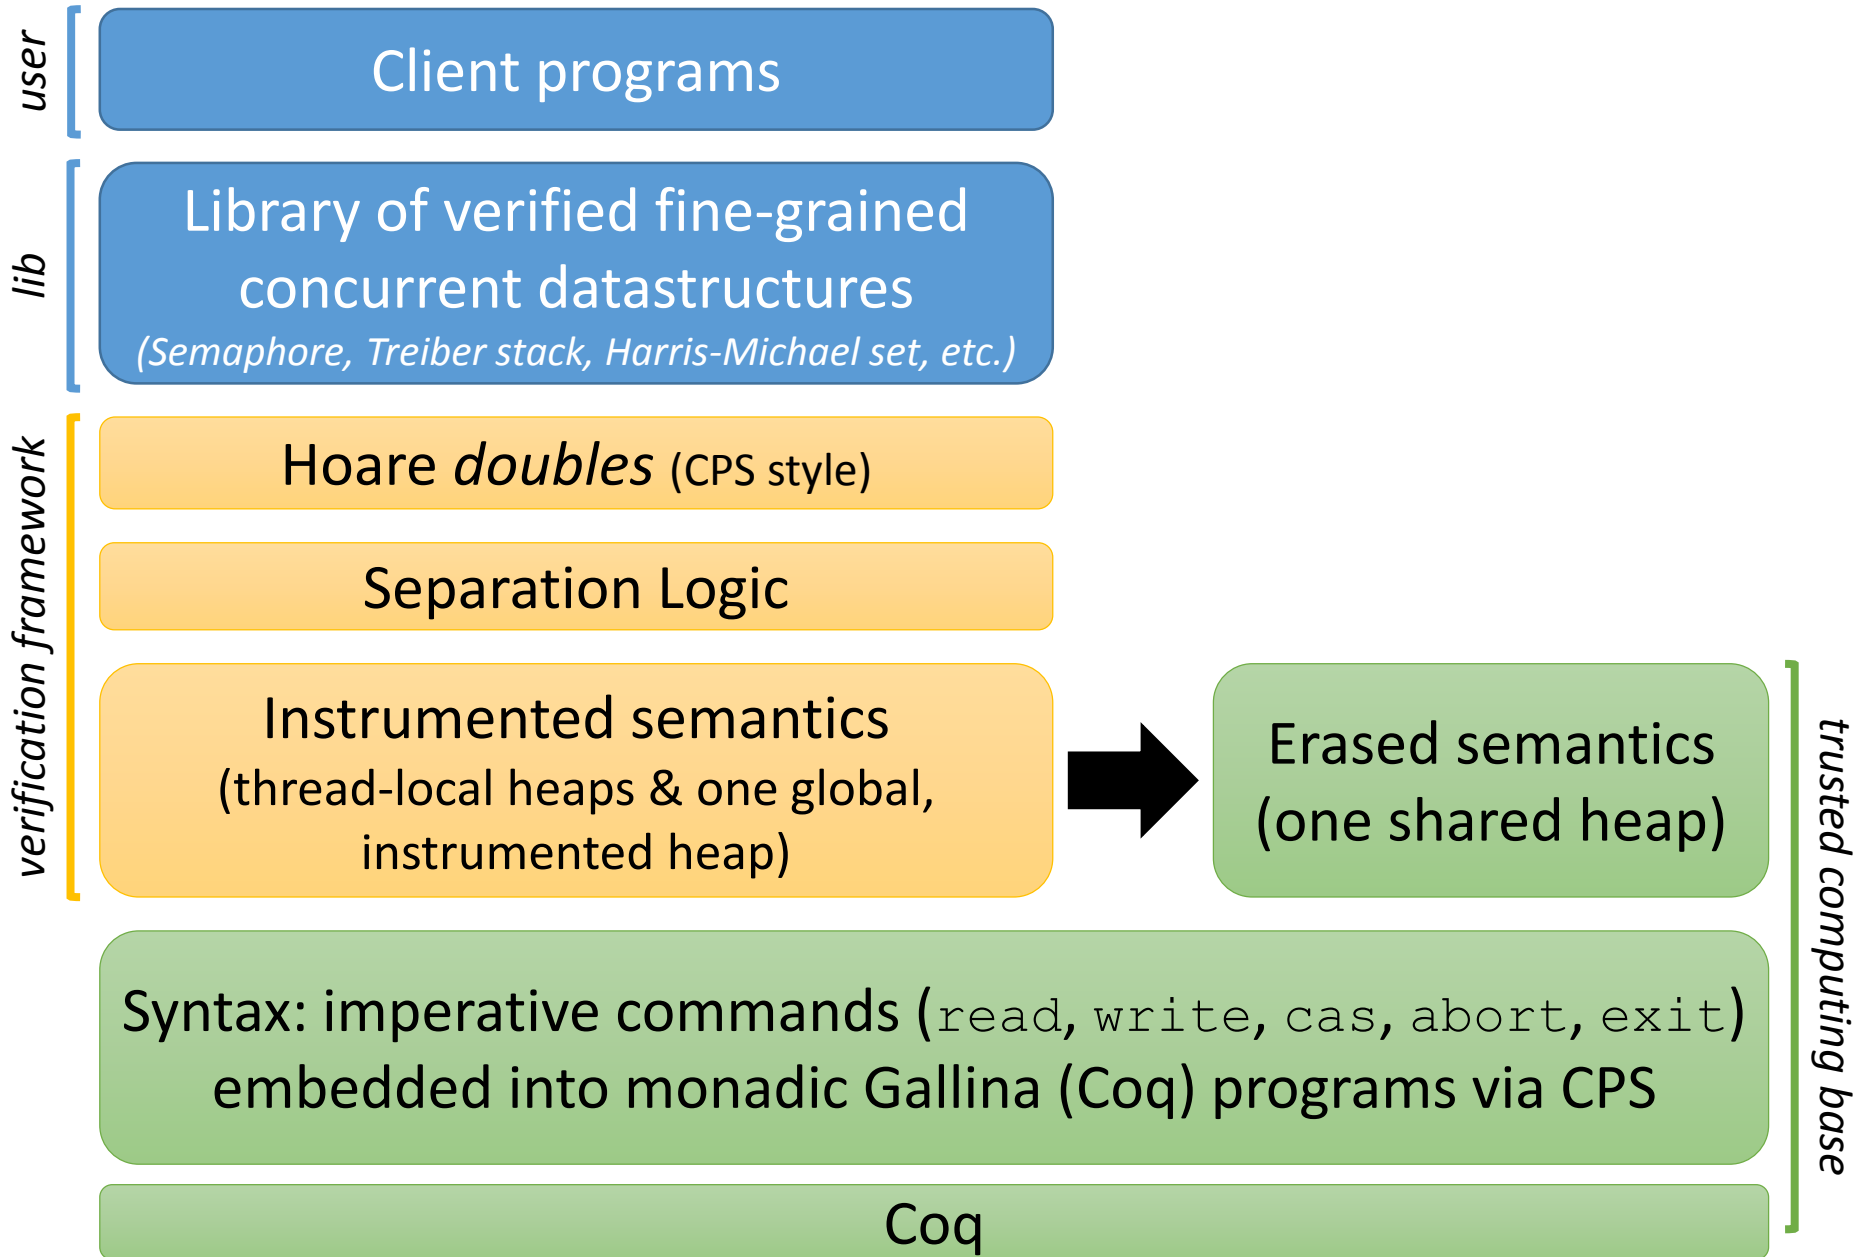
Our Focus

1. What does the concurrent program logic look like?
 - **abstraction**: high level
 - **local reasoning**: modular/manageable proofs
 - **generality**: can we prove real programs?
2. End-to-end verification
 - what does the **machine code** actually do?
 - can we **trust** our program logic?
3. Verification framework development
 - how do we **quickly test** new ideas?

Method

- Minimal operational semantics
 - Syntax: Imperative commands + Gallina programs
 - Erased & Instrumented semantics
- Minimal instrumentation for global state
 - “Phantom monitors” vs ghost state
- Verification framework is built on top
- Machine-checked proofs in Coq

Method



Trusted Computing Base

$v \in V$ Value
 $a \in A$ Address
 $\alpha \in O$::= `read a | write a v | cas a v0 v1`
 $s \in S$::= `x ← α; s | nil | abort`

$$\frac{h(a) = v}{(h, \text{read } a) \rightarrow_h (h, v)} \quad \frac{h(a) = v_1}{(h, \text{cas } a v_1 v_2) \rightarrow_h (h[a \mapsto v_2], 1)}$$

$$\frac{a \in \text{dom}(h)}{(h, \text{write } a v) \rightarrow_h (h[a \mapsto v], 1)} \quad \frac{a \in \text{dom}(h) \quad h(a) \neq v_1}{(h, \text{cas } a v_1 v_2) \rightarrow_h (h, 0)}$$

(a) Syntax

$$\frac{(h, \alpha) \rightarrow_h (h', v)}{(h, P \uplus [i \mapsto x \leftarrow \alpha; s]) \xrightarrow{i'} (h', P \uplus [i \mapsto s[v/x]])}$$

$i \in I$ Thread ID
 $h \in H = A \rightarrow V$ Heap
 $P = I \rightarrow S$ Processes

$$\frac{(h, P) \xrightarrow{i'} (h', P')}{(h, P) \rightarrow (h', P')}$$

(a) Semantic domains

(b) Erased operational semantics

$$\frac{\forall i. P(i) \neq \text{nil} \Rightarrow \exists h', P'. (h, P) \xrightarrow{i'} (h', P') \quad \forall h', P'. (h, P) \rightarrow (h', P') \Rightarrow \text{safe-program } h' P'}{\text{safe-program } h P} \text{ SAFE}$$

(c) Safety

TCB: Syntax

Inductive action: Set :=

```
| read: action
| write: address -> value -> action
| cas: address -> value -> value -> action.
```

CoInductive proc : Set :=

```
(* safely terminated thread *)
| p_nil: proc
(* crashed thread *)
| p_abort: proc
(* perform action, then call the continuation with its result *)
| p_act: action -> (value -> proc) -> proc.
```


Instrumented Global State

“Phantom Monitors”

- General idea: access to shared data structures is coordinated by a *global policy*:
 - what can the current thread do?
 - what can interfering threads do?
- We write a policy for a shared datastructure as a monadic [corecursive] Coq function that *monitors* every operation acting on the structure, rejecting any operation that violates the protocol, and evolving over time.


Instrumented Global State

“Phantom Monitors”



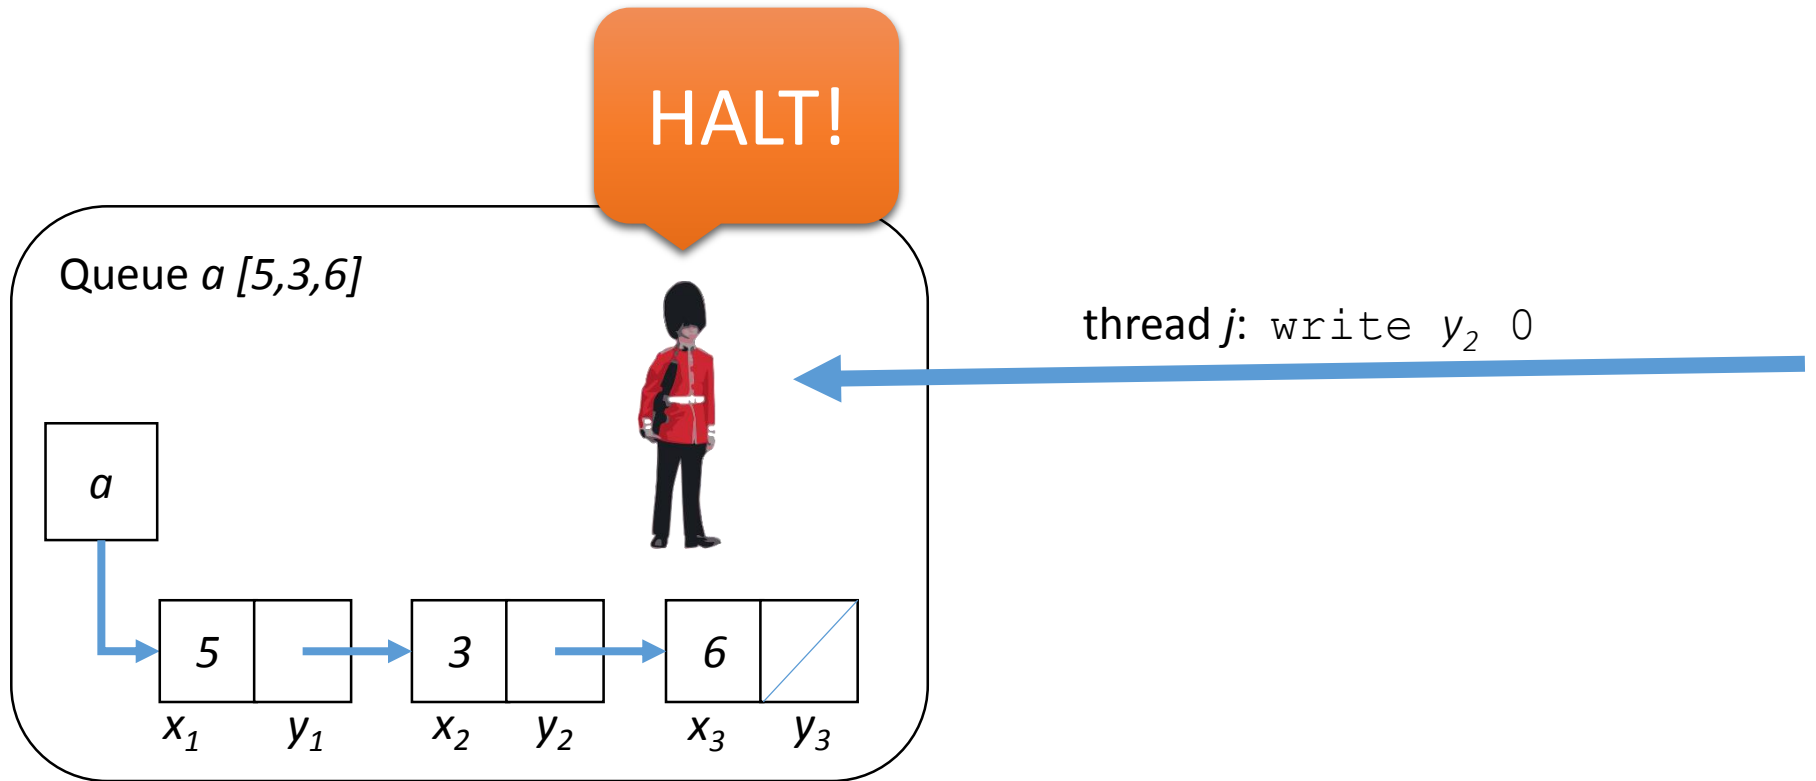
Instrumented Global State

“Phantom Monitors”



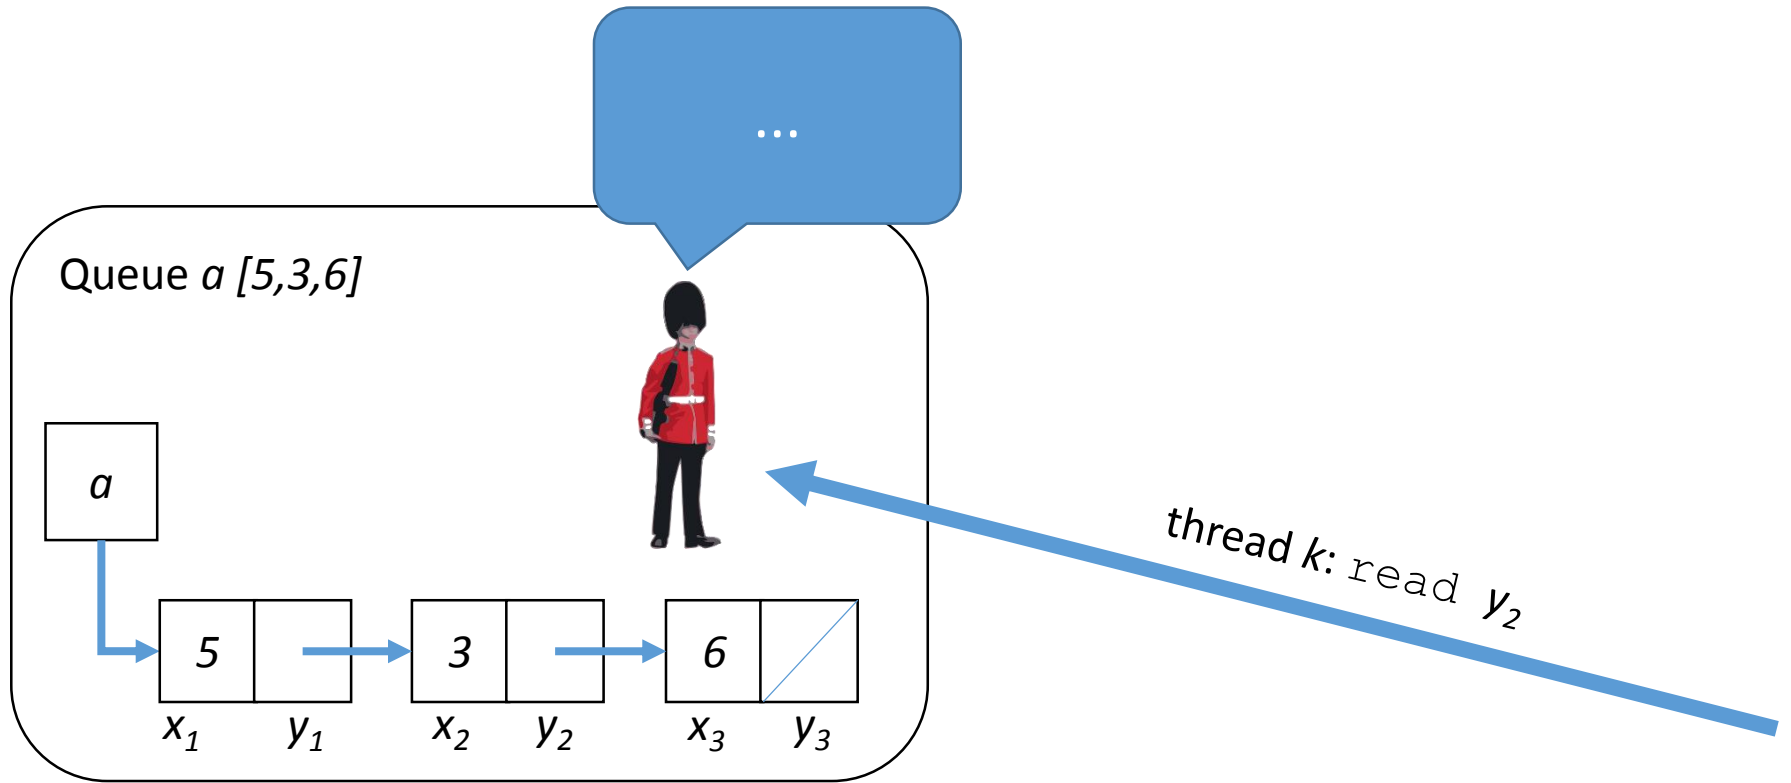
Instrumented Global State

“Phantom Monitors”



Instrumented Global State

“Phantom Monitors”



Phantom Monitors


Is a Coq function that:

1. Observes all operations on a data structure
2. Accepts or rejects each operation
3. May generate an abstract operation (“dequeue”) or silently accept it
4. Can change state
5. *Can be composed together*

Client Program Policy

(single-value case)

- client thread: c
 - pushes unfinished value into a shared stack ($\varsigma_0 \rightarrow \varsigma_1$)
 - collects the finished value ($\varsigma_1 \rightarrow \varsigma_3$)
- worker thread: w
 - checks the stack for (unfinished) values ($\varsigma_1 \rightarrow \varsigma_2$)
 - pushes the computed value of each pop ($\varsigma_2 \rightarrow \varsigma_1$)



Client Program Policy

(general case)

```
protocol JobsProto(input, client_pid, compute) implements StackProtocol
  list loading = input      // unfinished values to be pushed
  map processing = empty   // values held by worker threads

  onPush(i, x)
    case x = (Out v): //  $\zeta_2 \rightarrow \zeta_1$ 
      assert processing[i] = (Out v)
       $\forall \exists v'. \text{processing}[i] = (\text{In } v') \wedge v = \text{compute}(v')$ 
      processing.remove(i)
    case x = (In v): //  $\zeta_0 \rightarrow \zeta_1$ 
      assert  $\exists l'. \text{loading} = v :: l'$ 
      loading := tail(loading)

  onPop(i, x)
    if i = client_pid then
      case x = (Out v): assert True //  $\zeta_1 \rightarrow \zeta_3$ 
      case x = (In v): loading := v :: loading //  $\zeta_1 \rightarrow \zeta_0$ 
    else //  $\zeta_1 \rightarrow \zeta_2$ 
      assert  $i \notin \text{dom}(\text{processing})$ 
      processing.add(i, x)
```

Stack Specification

$\alpha ::= \text{push } v \mid \text{pop } v$ and:

$$\frac{}{\llbracket \sigma_0 \rrbracket = \epsilon} \quad \frac{\sigma \xrightarrow[i]{\text{pop } v} \sigma'}{\llbracket \sigma \rrbracket = v :: \llbracket \sigma' \rrbracket} \quad \frac{\sigma \xrightarrow[i]{\text{push } v} \sigma'}{\llbracket \sigma' \rrbracket = v :: \llbracket \sigma \rrbracket}$$

$$\frac{}{\text{stable}(\text{Stack}_\Sigma a \sigma)} \quad \frac{\sigma \rightsquigarrow_{\neq i}^* \sigma'}{\text{Stack}_\Sigma a \sigma' * \text{pid } i \vdash \text{Stack}_\Sigma a \sigma * \text{pid } i}$$

$$\frac{\text{stable } \mathcal{W} \quad \forall \sigma'. \sigma \rightsquigarrow_{\neq i}^* \sigma' \Rightarrow \sigma' \xrightarrow[i]{\text{push } v} \bullet \quad \forall \sigma'. \sigma \rightsquigarrow_{\neq i}^* \xrightarrow[i]{\text{push } v} \sigma' \Rightarrow \mathcal{V} \Vdash_i \{\text{Stack}_\Sigma a \sigma' * \mathcal{W}\} s}{\mathcal{V} \Vdash_i \{\text{Stack}_\Sigma a \sigma * \mathcal{W}\} \text{push } a v; s}$$

$$\frac{\text{stable } \mathcal{W} \quad \forall \sigma'. \sigma \rightsquigarrow_{\neq i}^* \sigma' \Rightarrow (\llbracket \sigma' \rrbracket = \epsilon \vee \exists v. \sigma' \xrightarrow[i]{\text{pop } v} \bullet) \quad \forall \sigma'. \sigma \rightsquigarrow_{\neq i}^* \sigma' \Rightarrow \llbracket \sigma' \rrbracket = \epsilon \Rightarrow \mathcal{V} \Vdash_i \{\text{Stack}_\Sigma a \sigma' * \mathcal{W}\} s \text{ None} \quad \forall \sigma', v. \sigma \rightsquigarrow_{\neq i}^* \xrightarrow[i]{\text{pop } v} \sigma' \Rightarrow \mathcal{V} \Vdash_i \{\text{Stack}_\Sigma a \sigma' * \mathcal{W}\} s \text{ (Some } v)}{\mathcal{V} \Vdash_i \{\text{Stack}_\Sigma a \sigma * \mathcal{W}\} x \leftarrow \text{pop } a; s x}$$

Hypotheses

Our minimal TCB, semantically derived framework, and Coq proofs enable:

- quick(-ish*) development cycle of our logical framework
- automated proofs (via Ltac)
- exploring *general* logics
 - ex: do not bake in:
 - composable global reasoning
 - permission accounting
 - derive restricted principles as needed
- verifying challenging concurrent programs

* Coq proofs take time, but concurrency is tricky enough that it is easy to make mistakes with pen & paper proofs of logics

Thanks!

Client Program Policy

(general case)

```
protocol StackMonitor  $\Sigma$  (address head,  $\sigma_0$ ) implements Monitor
   $\Sigma$   $\sigma = \sigma_0$  (* abstract client protocol *)

onRead(i, a, h, hAcq, hRel)
  assert hAcq = hRel = empty

onWrite(i, a, v, h, hAcq, hRel)
  assert False

onCAS(i, a, oldHead, newHead, h, hAcq, hRel)
  assert a = head  $\wedge$  hRel = empty
  if h(head) = oldHead then
    if h(oldHead.next) = newHead  $\wedge$  oldHead  $\neq$  0 then
       $\sigma$ .onPop(i, h(oldHead.item))
      assert hAcq = empty
    else if hAcq(newHead.next) = oldHead then
       $\sigma$ .onPush(i, hAcq(newHead.item))
      assert newHead  $\neq$  0  $\wedge$  dom(hAcq) = {newHead.item, newHead.next}
    else
      assert False
  else
    assert hAcq = empty
```