

CS 839: Systems Verification
Lecture 18: Lock invariants
Learning outcomes

1. Understand how concurrent separation logic extends sequential separation logic
2. Recall the rules for using lock invariants

Concurrent separation logic

1. Semantics is different, so need new soundness theorem
2. Need a rule for spawn
3. What else do we need to verify real programs?

Definition (CSL soundness): For some pure $\Phi(v)$ (a Prop), if $\{P\} e \{\lambda v \Phi(v)\}$ and $([e], h) \rightsquigarrow_{tp} ([e'] \mathbin{++} T, h')$, then if e' is an expression then $([e'] \mathbin{++} T, h')$ is not stuck, or $e' = v'$ for some value v' and $\Phi(v')$ holds. Furthermore, no thread in T is stuck in h' .

Exercise: soundness for spawned threads

Suppose we said (T, h) is stuck only if *no* threads could take a step.

What does soundness say now?

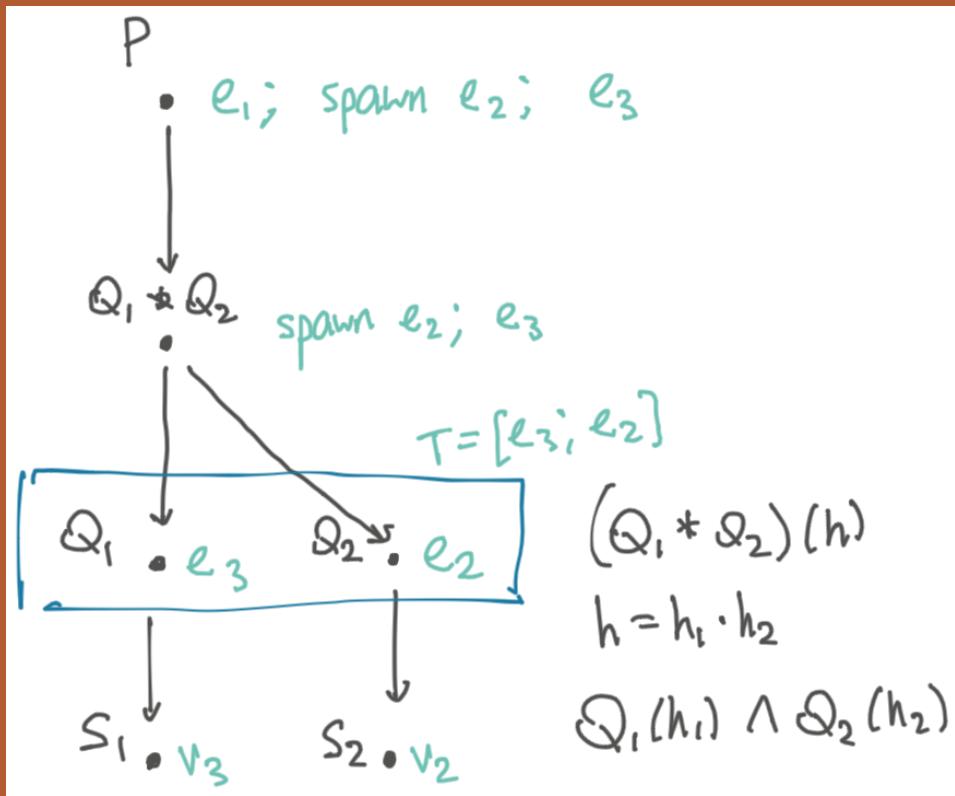
What program and specification would be true now that wasn't before?

Reasoning about spawn

$$\frac{\{P\} \ e' \ \{Q\}}{\{\text{wp}(e, \text{True}) * P\} \ \text{spawn } e; \ e' \ \{Q\}}$$

This rule is actually just framing P around the call to spawn, but it's easier to follow how this would work.

Note that the `wp(e, _)` here can only be used once - it's not a Hoare triple.



Intuition behind CSL: resources are split and transferred to other threads. Similar to how we sent resources to function calls, except this time both the spawned thread and the following code are running; with a function we waited for it to finish before proceeding.

Code demo

Locks (`sync.Mutex`)

```
// a zero sync.Mutex is an unlocked mutex
func (m *sync.Mutex) Lock()
func (m *sync.Mutex) Unlock()
```

Lock invariants

$\{R\} \ m := \text{NewMutex}() \ \{ \text{isLock}(m, R) \}$

$\{\text{isLock}(\ell_m, R)\} \ \text{Lock } \ell_m \ \{R\}$
 $\{\text{isLock}(\ell_m, R) * R\} \ \text{Unlock } \ell_m \ \{True\}$

Exercise: mutex invariant

Suppose we could get $isLock(\ell, R_1) * isLock(\ell, R_2)$ (note the same mutex). What would go wrong?

More code demo