

Specifying concurrent components

Ian J. Hayes

July 19, 2017 15:57

Specification of talk

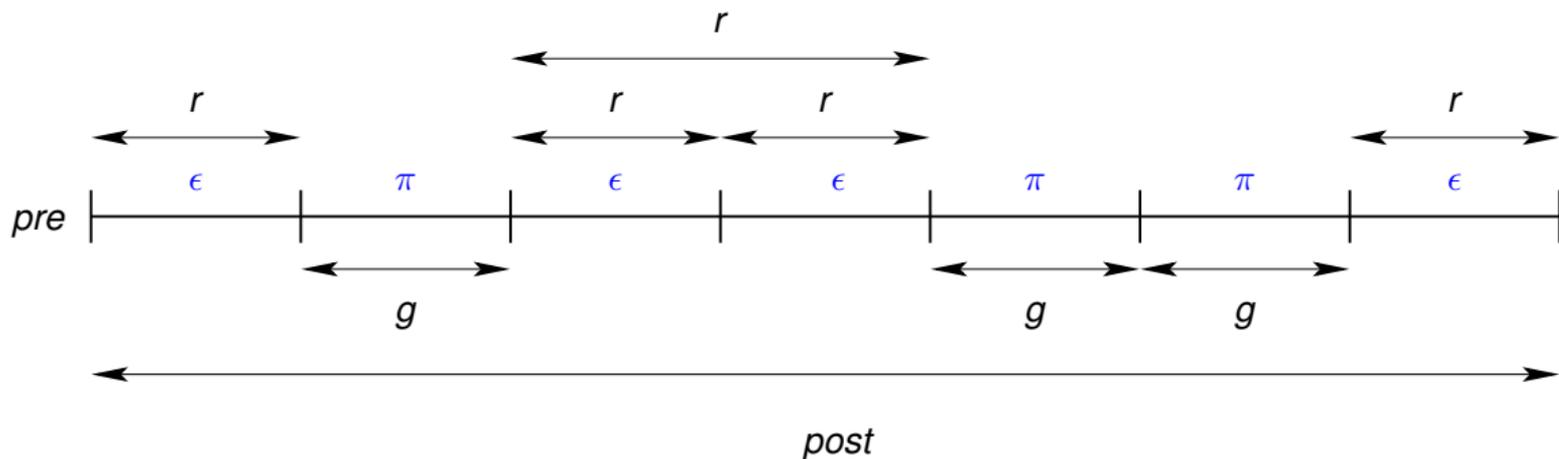
Me		You
	pre start on time	pre true
	rely you ask the right questions	rely I'll understand Ian's talk
guarantee you'll understand my talk	\parallel^1	guarantee I'll ask the right questions
assume questions will eventually stop		ensure I'll stop asking questions eventually
post finish talk		post finish listening

¹Assuming Michael implements real-time scheduling for the parallel

- ▶ Specification for the Linux ticket lock (and Peterson's) algorithm
- ▶ The specification emphasises decoupling
 - ▶ the specification of the module implementing the locking algorithm
 - ▶ from the context in which it is used,
 - ▶ allowing its development to take place independent of a particular context
- ▶ Need to clearly define the assumptions it makes about its context

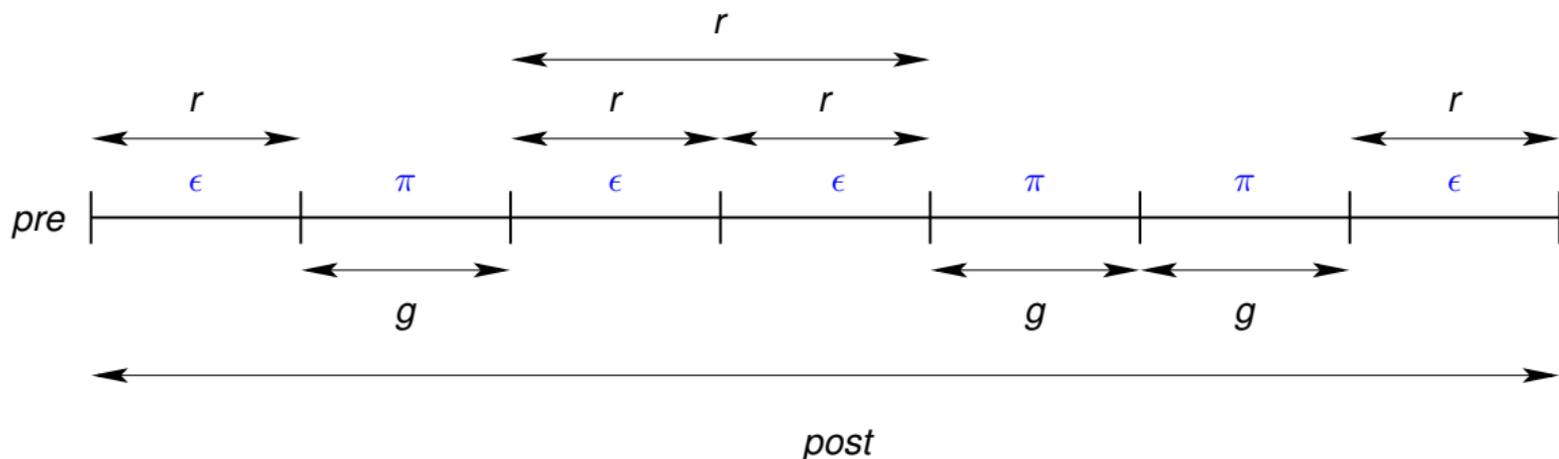
Concurrency and interference (Jones [1])

- ▶ pre condition
- ▶ post condition
- ▶ rely condition r
- ▶ guarantee condition g
- ▶ π atomic program step
- ▶ ϵ atomic environment step



Rely conditions

- ▶ If there is one environment step it must satisfy r
- ▶ If there are no environment steps, the interference is the identity relation, id , on states and hence $id \Rightarrow r$, i.e. r is **reflexive**
- ▶ If there are two environment steps, overall they satisfy $r \circ r$ and hence $r \circ r \Rightarrow r$, i.e. r is **transitive**
- ▶ Hence $r^* = r$



- ▶ To show termination of operations that may block (e.g. **await** b)
 - ▶ one needs to be able to assume that b will eventually be made (stably) true by the environment
- ▶ Because r is reflexive
 - ▶ it cannot be a well-founded relation and hence
 - ▶ it cannot be used to show that the environment eventually achieves some desired state
- ▶ The extension used here is to
 - ▶ add an additional assumption about the environment
 - ▶ expressed using temporal logic operators, like leads-to (e.g. $p_1 \rightsquigarrow p_2$).
 - ▶ these progress assumptions are in addition to the rely relation of Jones.

- ▶ It is an interesting concurrent algorithm that can be used in a multiprocessor context where multiple processors are vying for access to a resource
- ▶ Processes requiring the resource first take a “ticket” (as in a bakery or deli) and are given access to the resource in ticket order
- ▶ Provided each process that gains access to the resource releases it, all processes with a ticket will eventually receive control of the resource

- ▶ A module with two operations *acquire* and *release*
- ▶ The process at the head of the queue, if there is one, has the resource
- ▶ Processes are allocated the resource in queue order

```
module Lock(PID : finite type with equality)  
var q : seq PID init q = [] inv no_duplicates(q)
```

Acquiring the lock

The acquire operation gains access to the resource.

$acquire(k : PID) \hat{=}$

wr	q	// can only modify q
pre	$k \notin elems(q)$	// must not already be waiting in q
rely	$(hd(q) = k \Rightarrow hd(q') = k) \wedge$ $(k \in elems(q) \Rightarrow k \in elems(q'))$	// can't be preempted if k has resource // k stays in the queue
guar	$q' = q \vee$ $(k \notin elems(q) \wedge q' = q \hat{\ } [k])$	// may only append k (atomically)
assume	$k \in elems(q) \rightsquigarrow hd(q) = k$	// k gets to the head of queue
post	$hd(q') = k$	// k has the resource on termination

Releasing the lock

The *release* operation gives up control of the resource.

$release(k : PID) \hat{=}$

wr	q	// can only modify q
pre	$hd(q) = k$	// k must hold the resource
rely	$hd(q) = k \Rightarrow hd(q') = k$	// can't be preempted
guar	$q' = q \vee (hd(q) = k \wedge q' = tl(q))$	// can only remove k (atomically)
post	$k \notin elems(q)$	// it must release the resource

The set of process identifiers

- ▶ Any (large enough) set can be used for *PID*
- ▶ It is the responsibility of the calling context to ensure that
 - ▶ two different processes do not call *acquire* using the same process identifier
 - ▶ a process only calls *release* if it has the lock
- ▶ If the set of process identifiers is of size two, the specification can be implemented by Peterson's algorithm

$PID_Peterson ::= A \mid B$

$PetersonSpecification \hat{=} Lock(PID_Peterson)$

Example code for the use

```
use_lock(k : PID){  
  {k ∉ elems(q)}  
  non_critical_section; // guarantees  $R \wedge q' = q$   
  {k ∉ elems(q)}  
  acquire(k);  
  {hd(q) = k}  
  critical_section; // relies on  $R$ , guarantees  $q' = q$ , and terminates  
  {hd(q) = k}  
  release(k)  
  {k ∉ elems(q)}  
}
```

use_lock($k : PID$)

wr q, \dots

pre $k \notin \text{elems}(q) \wedge \dots$

rely $(\text{hd}(q) = k \Rightarrow (\text{hd}(q') = k \wedge R)) \wedge$
 $(k \in \text{elems}(q) \Rightarrow k \in \text{elems}(q'))$

assume $k \in \text{elems}(q) \rightsquigarrow \text{hd}(q) = k$

guar $((q = [] \vee \text{hd}(q) \neq k) \Rightarrow R) \wedge$
 $(\text{acquire.guarantee} \vee \text{release.guarantee})$

ensures $\forall j \in PID, i \in \mathbb{N} \bullet$

$j \in \text{elems}(q) - \{k\} \wedge \text{hd}(q) = k \wedge \text{posn}(q, j) = i \rightsquigarrow \text{posn}(q, j) < i$

post $k \notin \text{elems}(q) \wedge \dots$

rely R

$\{q = []\} \parallel_{k \in 1..N} \dots; use_lock(k); \dots$

Correctness of progress property

One process is enabled

$$q = [] \vee (\exists k \in PID \bullet hd(q) = k)$$

For each process the ensures of all other processes imply its assumption

$$\forall k \in PID \bullet (\forall n \in PID \bullet n \neq k \Rightarrow ensure_n) \Rightarrow assume_k$$

$$\forall k \in PID \bullet$$

$$(\forall n \in PID \bullet n \neq k \Rightarrow$$

$$\forall j \in PID, i \in \mathbb{N} \bullet$$

$$j \in elems(q) - \{n\} \wedge hd(q) = n \wedge posn(q, j) = i \rightsquigarrow posn(q, j) < i)$$

\Rightarrow

$$\forall i \in \mathbb{N} \bullet k \in elems(q) \wedge hd(q) \neq k \wedge posn(q, k) = i \rightsquigarrow posn(q, k) < i$$

$$(\forall i \in \mathbb{N} \bullet k \in elems(q) \wedge hd(q) \neq k \wedge posn(q, k) = i \rightsquigarrow posn(q, k) < i)$$

\Rightarrow

$$k \in elems(q) \rightsquigarrow hd(q) = k$$

```
module AbstractLock(PID : finite type with equality)  
var q : seq PID init q = [] inv no_duplicates(q)
```

Acquiring the lock

An abstract implementation of the *acquire* operation atomically adds k to the queue and then waits until k is at the head of the queue.

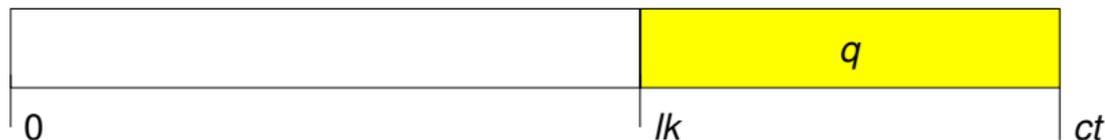
```
acquire( $k : PID$ ) {  
  { $k \notin elems(q)$ }  
  with  $q$  do  $q := q \hat{\ } [k]$ ;  
  await  $k = \langle hd(q) \rangle$   
}
```

Releasing the lock

The *release* operation removes k from the head of the queue.

```
release( $k : PID$ ){  
  { $k = hd(q)$ }  
  with  $q$  do  $q := tl(q)$   
}
```

Ticket lock implementation



The queue is empty if $lk = ct$. If the queue contains n elements then $lk + n = ct$.

```
module TicketLock(PID : finite type with equality)
var ct, lk :  $\mathbb{N}$  init  $ct = 0 \wedge lk = 0$  inv  $lk \leq ct$ 
abs-inv  $\#q = ct - lk$ 
```

Acquiring the lock

The abstract *acquire* operation adds k to the end of the queue.

```
acquire( $k : PID$ ){  
  var  $tk : \mathbb{N}$ ;  
  with  $ct$  do( $tk := ct$ ;  $ct := ct + 1$ ); // take a ticket  
                                           //  $q(tk - lk) = k$   
  await  $tk = \langle lk \rangle$  // wait until  $lk$  reaches the ticket's value  
}
```

Releasing the lock

The abstract release operation removes the head of the queue.

```
release( $k : PID$ ){  
   $\langle Ik \rangle := Ik + 1$   
}
```

The **with** statement within *acquire* can be implemented using a compare-and-swap (CAS) instruction.

$$\text{CAS}(x, \textit{old}, \textit{new}, \textit{done}) \hat{=} \text{with } x \text{ do} (\text{if } x = \textit{old} \text{ then } x := \textit{new}; \textit{done} := \textit{true} \\ \text{else } \textit{done} := \textit{false})$$

Taking ticket using CAS

Because an execution of a compare-and-swap can fail, it is placed in a loop that repeats until it succeeds.

```
var done :  $\mathbb{B}$ ;  
repeat  
   $tk := \langle ct \rangle$ ;  
  CAS( $ct, tk, tk + 1, done$ )  
until done
```

Unfortunately this loop may never terminate.

```
module PetersonLock implements PetersonSpecification  
var req : PID  $\rightarrow$   $\mathbb{B}$    init req(A) = req(B) = false  
    tk : PID           init tk  $\in$  PID
```

Acquiring the lock

The complement of a process identifier, \bar{k} , gives the identifier of the other process, i.e. $\bar{\bar{A}} = A$ and $\bar{\bar{B}} = B$.

```
acquire( $k : PID$ ) {  
  req( $k$ ) := true; tk :=  $\bar{k}$ ;  
  await req( $\bar{k}$ )  $\Rightarrow$  tk =  $k$   
}
```

Releasing the lock

```
release(k : PID) {  
  req(k) := false  
}
```

For an application using calls to *acquire/release* with a constant process identifiers, e.g. *acquire*(*A*), one can replace *req* by two variables *a* and *b*, corresponding to *req*(*A*) and *req*(*B*), respectively. That is, one can just inline the call and optimise it for the particular, constant parameter.

- ▶ compositionality
- ▶ modularisation
- ▶ abstract specification of operations
- ▶ data refinement

-  C.B. Jones.
Tentative steps toward a development method for interfering programs.
ACM ToPLaS, 5(4):596–619, October 1983.