

Simple Lock Program and Service

Shankar

September 18, 2014

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

Producer-Consumer using Lock Service

- Lock for threads $0, \dots, N-1$ $// N \geq 1$
 - input functions: `acq()`, `rel()`, `end()`
 - non-input function: `serve()`
- Main
 - `bool xreq[N]`: `xreq[i]` true iff user `i` has ongoing request
 - `bool xacq`: true iff a user holds the lock
 - start thread executing `serve()`
- Function `serve()`
 - cycle through entries of `xreq`
 - if `xreq[j]` true: set `xacq`, unset `xreq[j]`, wait for `xacq` false
- Input functions
 - `acq()`: set `xreq[mytid]`, wait for it to be false; return
 - `rel()`: unset `xacq`; return
 - `end()`: execute `endSystem()`; return

```
program SimpleLock(int N) {  
    ia {N ≥ 1}  
    boolean[N] xreq ← false;  
    boolean xacq ← false;  
    int xp ← 0;  
    Tid t ← startThread(serve());  
    return mysid;  
  
    function void serve() {  
        while (true)  
        a0: if ● (xreq[xp])  
        a1:   ● xacq ← true;  
        a2:   ● xreq[xp] ← false;  
        a3:   while ● (xacq) skip;  
        a4: xp ← mod(xp+1, N);  
    }  
}
```

Note the ●'s

- ignore them for now
- later we refer to them as “atomicity breakpoints”

```
input void mysid.acq()
    ia {mytid in 0..N-1}
a5: xreq[mytid] ← true;
a6: while • (xreq[mytid]) skip;
    return;
}

input void mysid.rel() {
    ia {mytid in 0..N-1}
a7: xacq ← false;
    return;
}

input void mysid.end() {
    ia {true}
    endSystem();
}
```

atomicity assumption:
reads and writes of
xacq,
xreq[0], ..., xreq[N-1]

progress assumption:
weak fairness
for threads

}

- Input assumptions of `acq()` and `rel()` are “weak”
 - only require caller tid to be in $0..N-1$
 - allow `acq()` caller to hold lock
 - allow `rel()` caller to not hold lock
- Hence the program has some odd **allowed** evolutions
 - e.g., two users hold lock simultaneously
[but it does implement SimpleLockService]
- Input assumptions are sufficient to ensure following
 - SimpleLock(N) is **fault-free**
// no allowed evolution is faulty
 - the **●**'s are a valid set of **atomicity breakpoints**
// code between two successive **●**'s is **effectively atomic**

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

Producer-Consumer using Lock Service

- Lock for threads $0, \dots, N-1$
- Main
 - vars indicating: whether ending; which user (if any) has lock
- Input functions `acq()`, `rel()`, `end()`
- No output function
- Defines **all** acceptable io sequences
- Constrains both environment and lock, e.g.,
 - `acq.ia`: not ending, caller in $0..N-1$, does not hold lock
- Atomicity assumptions: input parts and output parts
- Progress assumptions:
 - `acq()` returns eventually if lock becomes repeatedly free
 - `rel()` and `end()` each returns eventually


```
service SimpleLockService(int N) {  
  ic {N ≥ 1}  
  boolean[N] acqd ← false;    // acqd[i] true iff i has lock  
  ending ← false;             // termination initiated  
  return mysid;  
  
  input void mysid.acq() {  
    ic {not ending and (mytid in 0..N-1) and not acqd[mytid]}  
    oc {forall(j in 0..N-1: not acqd[j])}  
    acqd[mytid] ← true;  
    return;  
  }  
}
```

```
input void mysid.rel() {  
    ic {not ending and (mytid in 0..N-1) and acqd[mytid]}  
    acqd[mytid] ← false;  
    oc {true}  
    return;  
}
```

```
input void mysid.end() {  
    ic {not ending}  
    ending ← true;  
    oc {true}  
    return;  
}
```

atomicity assumption {input parts and output parts}

progress assumption {

 // rel returns

 forall(i: (i in mysid.rel) *leads-to* (not i in mysid.rel));

 // if no one holds the lock forever then acq returns

 forall(i: acqd[i] *leads-to* not acqd[i]) \Rightarrow

 forall(i: (i in mysid.acq) *leads-to* (not i in mysid.acq));

 // end returns

 forall(i: (i in mysid.end) *leads-to* (not i in mysid.end));

} }

■ Convention: i, j range over $0..N-1$

- Program is fault-free
 - otherwise it's useless as a service
- Atomicity breakpoints at (and only at) output conditions
 - natural consequence of atomicity assumptions
- Progress stated by leads-to (and not fairness) assertions
- Comparing against SimpleLock
 - input conditions stronger than SimpleLock's input assumptions
 - so precludes some ("odd") evolutions of SimpleLock
 - has io sequences not achievable by SimpleLock(N)

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

Producer-Consumer using Lock Service

- Define lock-service **inverse** program
 - most general environment for a lock implementation
- Define program Z:
 - concurrently executes implementation and service inverse
- Define the assertions that Z must satisfy
 - safety: Z satisfies inverse's input conditions
 - progress: Z inverse's progress assertions
- Prove that Z satisfies above assertions

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

- Simple Lock Service Inverse

- Implements conditions

- Proving the Implements Conditions

Producer-Consumer using Lock Service

```
service SimpleLockService(int N) {  
program SimpleLockServiceInverse(int N, Sid lck) {  
  // lck: lock system being tested  
  ic {N ≥ 1}  
  boolean[N] acqd ← false;  
  ending ← false;  
  return mysid;  
  
  input void mysid.acq() {  
  output doAcq() {  
    ie oc {not ending and (mytid in 0..N-1) and not acqd[mytid]}  
    lck.acq();  
    oe ic {forall(j in 0..N-1: not acqd[j])}  
    acqd[mytid] ← true;  
    return;  
  }  
}
```



```
output doRel() input void mysid.rel() {  
    ie oc {not ending and (mytid in 0..N-1) and acqd[mytid]}  
    acqd[mytid] ← false;  
    lck.rel();  
    oe ic {true}  
    return;  
}  
  
input void mysid.end() {  
output doEnd() {  
    ie oc {not ending}  
    ending ← true;  
    lck.end();  
    oe ic {true}  
    return;  
}
```

atomicity assumption {input parts and output parts}

progress assumption **condition** {

forall(i: (i in **mysid lck.rel**)
 leads-to (not i in **mysid lck.rel**));

forall(i: **acqd[i]** *leads-to* not **acqd[i]**) \Rightarrow
 forall(i: (i in **mysid lck.acq**)
 leads-to (not i in **mysid lck.acq**));

forall(i: (i in **mysid lck.end**)
 leads-to (not i in **mysid lck.end**));

}

}

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

Simple Lock Service Inverse

Implements conditions

Proving the Implements Conditions

Producer-Consumer using Lock Service

```
program Z(int N) {  
  ic { $N \geq 1$ }  
  inputs(); outputs();           // aggregate sys-  
tem Z is closed  
  Sid lck  $\leftarrow$  startSystem(SimpleLock(N));  
  Sid lsi  $\leftarrow$  startSystem(SimpleLockServiceInverse(N, lck));  
  return mysid;  
  atomicity assumption {}  
  progress assumption {weak fairness}  
}
```

$B_0 : \text{Inv } [(i \text{ at } \text{lsi.doAcq.ic}) \Rightarrow \text{forall}(j: \text{not } \text{acqd}[j])]$

$B_1 : (i \text{ in } \text{lck.rel}) \text{ leads-to } (\text{not } i \text{ in } \text{lck.rel})$

$B_2 : \text{forall}(i: \text{acqd}[i] \text{ leads-to } \text{not } \text{acqd}[i]) \Rightarrow$
 $\text{forall}(i: (i \text{ in } \text{lck.acq}) \text{ leads-to } (\text{not } i \text{ in } \text{lck.acq}))$

$B_3 : (i \text{ in } \text{lck.end}) \text{ leads-to } (\text{not } i \text{ in } \text{lck.end})$

■ Recall conventions

- i, j range over $0..N-1$
- free variables are universally quantified
 e.g., B_3 equivalent to $\text{forall}(i: B_3)$

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

Simple Lock Service Inverse

Implements conditions

Proving the Implements Conditions

Producer-Consumer using Lock Service

system $lck(N)$

<main>

fn serve(){...●...}

input acq(){...●...}

input rel(){...}

input end(){...}

system $lsi(N, lck)$

<main>

output doAcq(){●oc ...}

output doRel(){●oc ...}

output doEnd(){●oc ...}

- step $Z.init$: $Z.main, lck.init, lsi.main$
- step $doAcq.call$: $lsi.doAcq.oc● \rightarrow lck.acq●$
- step $acq.ret$: $lck.acq● \rightarrow lsi.doAcq.end$
- step $doRel$: $lsi.doRel.oc● \rightarrow lck.rel \rightarrow lsi.doRel.end$
- step $doEnd$: $lsi.doEnd.oc● \rightarrow lck.end \rightarrow lsi.doEnd.end$
- steps in $lck.serve()$ defined by its ●'s
 - valid in Z because lck gets only allowed inputs (from lsi)

- Recall B_0 : if thread at `doAcq.ic` then every `acqd[j]` is false
- Given Z 's effective atomicity, B_0 is equivalent to $Inv\ C_0$

$$C_0 : ((i \text{ on } lck.acq\bullet) \text{ and not } lck.xreq[i]) \Rightarrow \\ \text{forall}(j: \text{not } lsi.acqd[j])$$

- $Inv\ C_1$ and $Inv\ C_2$ hold // operational reasoning

$$C_1 : (lck.alive \text{ and } (\text{not } t \text{ on } a3)) \Rightarrow \text{forall}(j: \text{not } acqd[j])$$

$$C_2 : (t \text{ on } a3) \Rightarrow \\ ((acqd[xp] \text{ or } \\ (\text{not } acqd[xp] \text{ and } (xp \text{ on } a6) \text{ and not } xreq[xp])) \\ \text{and } \text{forall}(j, j \neq xp: \text{not } acqd[j]))$$

- $Inv\ C_0$ holds from $Inv\ C_1$ and $Inv\ C_2$ // operational reasoning

- Recall B_1 : thread in `lck.rel` eventually leaves `lck.rel`
- B_1 holds
 - `lck.rel.body` has no loops and no blocking
 - thread has weak fairness (from `lck` progress assumption)
- Recall B_3 : thread in `lck.end` eventually leaves `lck.end`
- B_3 holds just like B_1

- Recall B_2 : $D_0 \Rightarrow D_1$, where

D_0 : $\text{acqd}[i]$ *leads-to* $\text{not acqd}[i]$

D_1 : $(k \text{ in } \text{lck.acq})$ *leads-to* $(\text{not } k \text{ in } \text{lck.acq})$

- We will establish the following

D_2 : $[t \text{ at } a_0, xp = j, j \text{ in } \text{lck.acq}]$ *leads to*
 $[xp \text{ not in } \text{lck.acq}]$

D_4 : $[t \text{ at } a_0, xp = j]$ *leads to*
 $[t \text{ at } a_0, xp = \text{mod}(j+1, N)]$

- D_2 and D_4 imply D_1

- We establish

$D_2 : ((t \text{ on } a0) \text{ and } xp = j \text{ and } xreq[j]) \text{ leads-to}$
 $((t \text{ on } a3) \text{ and } xp = j \text{ and } acqd[j])$

- Proof

- “j in lck.acq” equivalent to “j at a6” // Z's atomicity
- [j at a6, t at a0, xp=j] leads to // via wfair t
[j at a6, t at a3, xreq[j] is false] leads to // via wfair j
[j not at a6, t at a3, xreq[j] is false]

- We establish

$D_3 : ((t \text{ on } a3) \text{ and } xp = j \text{ and } \text{acqd}[j]) \text{ leads-to } ((t \text{ on } a0) \text{ and } xp = \text{mod}(j+1, N))$

- Proof

- $[D_2 \text{'s rhs}] \text{ leads to } // \text{ via } D_0, j.\text{doRel}$
 $[xacq \text{ false}, t \text{ on } a3] \text{ leads to } // \text{ via wfair } t$
 $[t \text{ at } a0, xp \text{ is } \text{mod}(j+1, N)]$

- D_2 and D_3 imply

$D_4 : ((t \text{ on } a0) \text{ and } xp = j) \text{ leads-to } ((t \text{ on } a0) \text{ and } xp = \text{mod}(j+1, N))$

- See text

Simple Lock Program

Simple Lock Service

Proving Lock Implements Service

Producer-Consumer using Lock Service

- Program ProdCons1
 - start systems: producer, consumer, lock service
 - producer and consumer use lock service
- Show that ProdCons1 is fault-free
 - show that it satisfies input conditions of lock service system
- Obtain atomicity breakpoints // effective atomicity
- Establish desired properties
 - still hold when lock service is replaced by a lock implementation

```
program ProdConsLck(...) {  
  ia {...}  
  <hide lck inputs>;  
  lck ← startSystem(SimpleLockService());  
  cons ← startSystem(Consumer(lck));  
  prod ← startSystem(Producer(lck, cons));  
  return [0, mysid];  
  
  atomicity assumption {} // none  
  progress assumption {weak fairness}  
}
```


SimpleLockService(N): ... input mysid.acq(): ic {...} ● oc {...} input mysid.rel(): ... input mysid.end(): ...	Consumer(lck): start- Thd(consum()); fn consum(): while (...) lck.acq(); ... lck.rel(); lck.end(); endSystem(); input mysid.put(): ...	Producer(lck,cons): start- Thd(prod()); fn produce(): while (...) lck.acq(); cons.put(); lck.rel(); endSystem();
--	---	--

- Single atomicity breakpoint in entire program text
 - ProdCons.init: start → only 2 threads at lck.acq
 - cons.step: lck.acq → lck.acq or exit
 - prod.step: lck.acq → lck.acq or exit