Concepts of Concurrent Computation Spring 2014 Lecture 4: Semaphores

Bertrand Meyer Sebastian Nanz Chris Poskitt

Chair of Software Engineering



Last week: synchronisation, but lacking the simplicity

- we looked at various solutions to the mutual exclusion problem
- algorithms were limited to the simplest tools atomic read and write to shared memory
 - => difficult to implement; complex
 - => reliance on busy waiting
 - => no encapsulation of synchronisation variables

















































































































Today's lecture: semaphores

- we will discuss <u>semaphores</u>, an important synchronisation primitive
- conceptually simple, although their implementations require stronger atomic operations
- widespread use in operating systems
- invented by Dijkstra in 1965



Next on the agenda

- I. general and binary semaphores
- 2. implementing semaphores
- 3. beyond the mutual exclusion problem
- 4. simulating general semaphores

General semaphores (aka "counting semaphores")

• a general semaphore is an object consisting of:

(1) an integer variable count such that count ≥ 0

(2) two <u>atomic</u> operations: <u>down</u> and <u>up</u>

if a process calls *down* when *count* > 0, then *count* is decremented by 1 (otherwise it first waits)

if a process calls *up*, then *count* is incremented by I

General semaphores (in Eiffel-like pseudocode)

class SEMAPHORE feature

count : INTEGER

down

do-atomic
 await count > 0
 count := count - |
end

up do-atomic count := count +/ end end

General semaphores (in Eiffel-like pseudocode)

class SEMAPHORE feature

count : INTEGER



will discuss how to ensure <u>atomicity</u> and how to avoid <u>busy wait</u> later!

up do-atomic count := count + / end end

• create a semaphore s and initialise s.count to 1; then:

s.down critical section s.up

create a semaphore s and initialise s.count to 1; then:

s.down critical section s.up

one process at a time; or one hot desk!



• or in the style of last week's mutual exclusion problems:

| count := 1 | | | | |
|------------------|---|------------------|---|--|
| P1 | | P2 | | |
| 1 2 3 4 | <pre>while true loop await count > 0 count := count - 1 critical section count := count + 1 non-critical section end</pre> | 1 2 3 4 | <pre>while true loop await count > 0 count := count - 1 critical section count := count + 1 non-critical section end</pre> | |

 mutual exclusion and deadlock freedom can be proven

=> remember the atomicity of *down* and *up*!

• solution does <u>not</u> satisfy starvation freedom

=> a different implementation later will fix this

The general semaphore invariant

- general semaphores are characterised by the following invariant -- important for proofs!
- given some semaphore, let:
 - => k denote its <u>initial</u> value with $k \ge 0$
 - => count denote its <u>current</u> value
 - => #down denote the number of completed down operations
 - => #up denote the number of completed up operations
- then the following equations are invariant:

(1)
$$count \ge 0$$

(2) $count = k + #up - #down$

Binary semaphores

- in the previous example, s.count is always either 0 or 1
- such a semaphore is called a binary semaphore and can be implemented using a Boolean variable

b : BOOLEAN

```
down
do-atomic
await b
b := false
end
up
do-atomic
b := true
end
```

Binary semaphores

- in the previous example, s.count is always either 0 or 1
- such a semaphore is called a binary semaphore and can be implemented using a Boolean variable

b : BOOLEAN

```
down
do-atomic
await b
b := false
end
up
do-atomic
b := true
end
```

This is <u>deceptively</u> similar to the previous weeks' early, and wrong attempts at providing mutual exclusion. What's different?

Next on the agenda

- I. general and binary semaphores
- 2. implementing semaphores
- 3. beyond the mutual exclusion problem
- 4. simulating general semaphores

Avoiding busy waiting

• busy-wait semaphores are not ideal

=> they are not starvation free
=> inefficient in the context of multitasking

 more preferable would be for processes to block themselves when having to wait

=> thus freeing processing resources as early as possible

 idea: keep track of blocked processes, "waking them" upon up calls on the semaphore

Avoiding busy waiting



Avoiding busy waiting



Implementing the scheme

- to avoid starvation, we will track blocked processes in a collection *blocked*
- we equip *blocked* with the following operations, which will be integrated into *down* and *up*

=> add(P) -- insert process P into collection
=> remove -- select, remove, and return an item from the collection
=> is_empty -- true if collection empty; false otherwise

• if *blocked* is implemented as a set, we call the semaphore weak; if as a FIFO queue, then strong

Weak semaphore

• a weak semaphore is a blocking semaphore in which the collection *blocked* is implemented as a set

=> *remove* will pick and remove a <u>random</u> element

down do-atomic if count > 0 then count := count - 1 else blocked.add(P) P.state := blocked end end up do-atomic if blocked.is_empty then count := count + 1 else Q := blocked.remove Q.state := ready end end

Weak semaphore

• a weak semaphore is a blocking semaphore in which the collection *blocked* is implemented as a set

=> *remove* will pick and remove a <u>random</u> element



Weak semaphore

• a weak semaphore is a blocking semaphore in which the collection *blocked* is implemented as a set

=> *remove* will pick and remove a <u>random</u> element



 weak semaphores provide starvation-freedom in the two process scenario

=> why?

• what about mutual exclusion for *n* processes?

• create a semaphore s and initialise s.count to 1; then:



- starvation is possible for n > 2 with weak semaphores because we select a process from blocked at random
- solution is to use a strong semaphore, in which blocked is implemented as a FIFO queue

Strong semaphores provide a solution to the mutual exclusion problem with *n* processes (how to prove)

• mutual exclusion -- show that

f#cs + count = I

where **#cs** is the number of processes in critical sections

starvation freedom -- apply proof by contradiction

=> begin by assuming that a process in *blocked* is starved

A note on atomicity

• operations *down* and *up* are typically built in software using lower-level primitives (e.g. synchronisation algorithms)

• alternatively:

=> using test-and-set instructions (atomic read and write)
=> disabling interrupts (only realistic on a single processing unit)

A note on Java

• java.util.concurrent.Semaphore

http://docs.oracle.com/javase/7/docs/api/java/util/concurrent/Semaphore.html

constructors

- => Semaphore(int k)
- -- a weak semaphore
- => Semaphore(int k, boolean b) -- a strong semaphore if b true

operations

- => acquire() -- corresponds to down
- => release() -- corresponds to up

Next on the agenda

- I. general and binary semaphores
- 2. implementing semaphores
- 3. beyond the mutual exclusion problem
- 4. simulating general semaphores

The k-exclusion problem

• in the *k*-exclusion problem, we allow up to *k* processes to simultaneously be in their critical sections

=> mutual exclusion is the k = 1 instance

 use a general semaphore corresponding to the number of processes allowed to be in their critical sections

| s.count := k | | |
|------------------|---|--|
| Pi | | |
| 1 2 3 4 | while true loop s.down critical section s.up non-critical section | |
| | end | |

The k-exclusion problem

 in the k-exclusion problem, we allow up to k processes to simultaneously be in their critical sections

=> mutual exclusion is the k = 1 instance

 use a general semaphore corresponding to the number of processes allowed to be in their critical sections



Barriers

- semaphores can be used to control the ordering of events in a system
- a barrier is a form of synchronisation that determines a point in a program's execution that all processes in a group have to reach before any of them may move on

=> important for concurrent iterative algorithms

Barriers

- semaphores can be used to control the ordering of events in a system
- a barrier is a form of synchronisation that determines a point in a program's execution that all processes in a group have to reach before any of them may move on

=> important for concurrent iterative algorithms

| s1.count := 0 s2.count := 0 | | | | |
|--------------------------------|---|------------------|---|--|
| P1 | | P2 | | |
| 1 2 3 4 | code before the barrier s1.up s2.down code after the barrier | 1 2 3 4 | code before the barrier s2.up s1.down code after the barrier | |

Barriers

- semaphores can be used to control the ordering of events in a system
- a barrier is a form of synchronisation that determines a point in a program's execution that all processes in a group have to reach before any of them may move on

=> important for concurrent iterative algorithms

| s1.count := 0 s2.count := 0 s2.count := 0 | | | | |
|---|---|------------------|---|--|
| P1 | | P2 | | |
| 1 2 3 4 | code before the barrier s1.up s2.down code after the barrier | 1 2 3 4 | code before the barrier s2.up s1.down code after the barrier | |

The producer-consumer problem



The producer-consumer problem



The producer-consumer problem

• a good solution would:

=> ensure that every data item produced is eventually consumed
=> be deadlock-free
=> be starvation-free

- need a semaphore for mutual exclusion (the buffer)
- but additional semaphore(s) for condition synchronisation

=> e.g. consumer should block until the buffer is non-empty

Solution for an <u>un</u>bounded buffer

| mu [.] not | mutex.count := 1 not_empty.count := 0 | | | | |
|---|--|-----------------------|--|--|--|
| Producer _i Consumer _i | | | nsumer _i | | |
| 1 2 3 4 5 | <pre>while true loop d := produce mutex.down b.append(d) mutex.up not_empty.up end</pre> | 1 2 3 4 5 | <pre>while true loop not_empty.down mutex.down d := b.remove mutex.up consume(d) end</pre> | | |

Solution for an <u>un</u>bounded buffer

| mutex.count := 1 | | | | |
|-----------------------|--|-----------------------|--|--|
| Pro | Producer _i Consumer _i | | | |
| 1 2 3 4 5 | <pre>while true loop d := produce mutex.down b.append(d) mutex.up not_empty.up end</pre> | 1 2 3 4 5 | <pre>while true loop not_empty.down mutex.down d := b.remove mutex.up consume(d) end</pre> | |

Solution for an <u>un</u>bounded buffer

| mutex.count := 1 | | | | |
|------------------|---|-----|-----------------|--|
| Pro | not_empty.count := 0 | | | |
| | uucei i | C01 | isumer i | |
| 1 | while true loop d := produce | 1 | while true loop | |
| 2 | blocks until not_empty.count > 0 mutex.down | | | |
| 3 | p.uppena(a) | | a := b.remove | |
| 4 | mutex.up | 4 | mutex.up | |
| 5 | not_empty.up | 5 | consume(d) | |
| | end | | end | |

Solution for a bounded buffer

| <pre>mutex.count := 1 not_empty.count := 0 not_full.count := k</pre> | | | where k is the size of the buffer | |
|--|--|-----------------------|-----------------------------------|--|
| Producer _i | | Consumer _i | | |
| 1 2 3 4 5 | <pre>while true loop d := produce not_full.down mutex.down b.append(d) mutex.up not_empty.up end</pre> | | 1 2 3 4 5 | <pre>while true loop not_empty.down mutex.down d := b.remove mutex.up not_full.up consume(d) end</pre> |

Dining philosophers problem (a solution that can deadlock)

 multiple semaphores must be used with care -- they are prone to deadlock!

s[1].count := 1, ..., s[n].count := 1
Philosopher,
while true loop

| | while true loop |
|---|-----------------------|
| 1 | think |
| 2 | s[i].down |
| 3 | s[(i mod n) + 1].down |
| 4 | eat |
| 5 | s[(i mod n) + 1].up |
| 6 | s[i].up |
| | end |



()



 multiple semaphores must be used with care -- they are prone to deadlock! ()



Dining philosophers problem (a fix!)

- assume that philosopher n picks up the left fork before the right fork
- this breaks the circle of resource requests; there will always be one philosopher who can acquire both forks and release them again

| Phi | Philosophern | | |
|-----|-----------------|--|--|
| | while true loop | | |
| 1 | think . | | |
| 2 | s[1].down | | |
| 3 | s[n].down | | |
| 4 | eat | | |
| 5 | s[n].up | | |
| 6 | s[1].up | | |
| | end | | |

Next on the agenda

- I. general and binary semaphores
- 2. implementing semaphores
- 3. beyond the mutual exclusion problem
- 4. simulating general semaphores

General semaphores are superfluous

 while conceptually useful, general semaphores (theoretically) are not necessary -- they can be implemented through binary semaphores alone

General semaphores are superfluous

```
mutex.count := 1 -- binary semaphore
delay.count := 1 -- binary semaphore
count := k
```

general_down do delay.down mutex.down count := count - 1if count > 0 then delay.up end mutex.up end

general_up do mutex.down count := count + 1 if count = 1 then delay.up end mutex.up end

General semaphores are superfluous



Next on the agenda

- I. general and binary semaphores
- 2. implementing semaphores
- 3. beyond the mutual exclusion problem
- 4. simulating general semaphores

Summary

- semaphores are conceptually simple but powerful tools for solving synchronisation problems
- choice of implementation can affect starvation-freedom
- applications beyond mutual exclusion: k-exclusion, barriers, condition synchronisation

but: correct usage is still far from trivial

• essential reading: Chapter 4 of the CCC textbook