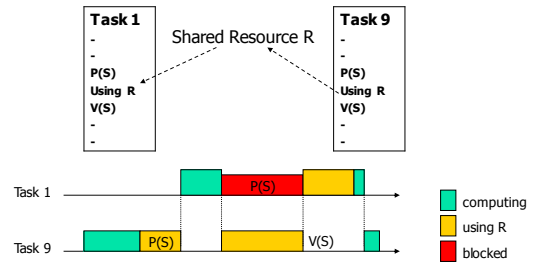


Today's topic:
Resource/Data Sharing and Synchronization

Priority Ceiling Protocols

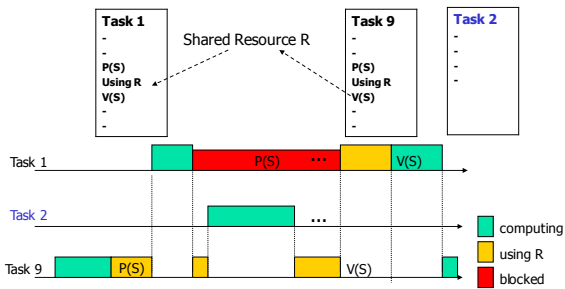
The simplest form of priority inversion



1

2

Un-bounded priority inversion



3

4

Solutions

- Tasks are 'forced' to follow certain rules when locking and unlocking semaphores (requesting and releasing resources)
- The rules are called 'Resource access protocols'
 - NPP, BIP, HLP, PCP


Non Preemption Protocol (NPP)

- Modify $P(S)$ so that the caller is assigned the highest priority if it succeeds in locking S
 - Highest priority=non preemption!
- Modify $V(S)$ so that the caller is assigned its own priority back when it releases S

This is the simplest method to avoid Priority Inversion!

NPP: + and -

- Simple and easy to implement (+), how?
- Deadlock free (++) , why?
- Number of blocking = 1 (+), Why?
- Allow low-priority tasks to block high-priority tasks including those that have no sharing resources (-)



5


6

Basic Priority Inheritance Protocol (BIP)

- supported in RT POSIX
- Idea:**
 - A gets semaphore S
 - B with higher priority tries to lock S, and blocked by S
 - B transfers its priority to A (so A is resumed and run with B's priority)
- Run time behaviour:** whenever a lower-priority task blocks a higher priority task, it inherits the priority of the blocked task


7

Example



8


Problem 1: potential deadlock



Task 2: ... P(S2) ... P(S1) ...
Task 1: ... P(S1) ... P(S2) ...

9

Problem 2: chained blocking – many preemptions



Task 1 needs M resources may be blocked M times:
→ many preemptions/much run-time overheads
→ maximal blocking = sum of all CS sections for lower-priority tasks

10

BIP: Blocking time calculation

- Let
 - $CS(k,S)$ denote the computing time for the critical section that task k uses semaphore S.
 - $Use(S)$ is the set of tasks using S
- The maximal blocking time for task i:
 - $B = \sum \{CS(k,S) \mid i, k \text{ in } Use(S), pr(k) < pr(i) \leq C(S)\}$

11

Properties of BIP: + and -

- Bounded Priority inversion (+)
- Reasonable Run-time performance (+)
- Require no info on resource usage of tasks (+)
- Potential deadlocks (-)
- Chain-blocking – many preemptions (-)

12

Implementation of Ceiling Protocols

- Main ideas:
 - Priority-based scheduling
 - Implement P/V operations on Semaphores to assign task priorities dynamically

13

Semaphore Control Block for BIP

counter
queue
Pointer to next SCB
Holder

14

Standard P-operation (without BIP)

- P(scb):


```

Disable-interrupt;
If scb.counter>0 then {scb.counter - -1;
else
    {save-context( );
    current-task.state := blocked;
    insert(current-task, scb.queue);
    dispatch();
    load-context( ) }
Enable-interrupt
            
```

15

P-operation with BIP

- P(scb):


```

Disable-interrupt;
If scb.counter>0 then {scb.counter - -1;
                        scb.holder:= current-task
                        add(current-task.sem-list,scb)}
else
    {save-context( );
    current-task.state := blocked;
    insert(current-task, scb.queue);
    /* queue sorted according to task priority */
    save(scb.holder.priority);
    scb.holder.priority := current-task.priority;
    dispatch();
    load-context( ) }
Enable-interrupt
            
```

16

Standard V-operation (without BIP)

- V(scb):


```

Disable-interrupt;
If not-empty(scb.queue) then
    { next-to-run := get-first(scb.queue);
    next-to-run.state := ready;
    insert(next-to-run, ready-queue);
    save-context();
    schedule(); /* dispatch invoked*/
    load-context( ) }
else scb.counter ++1;
Enable-interrupt
            
```

17

V-operation with BIP

- V(scb):


```

Disable-interrupt;
current-task.priority := "original/previous priority"
/* highest-priority of tasks blocked by smaphors ownd by current-task*/
/* check all blocked tasks waiting for sem in current-tcb.sem-list*/
If not-empty(scb.queue) then
    { next-to-run := get-first(scb.queue);
    /*queue sorted according to task priority */
    next-to-run.state := ready;
    scb.holder := next-to-run;
    add(next-to-run.sem-list, scb);
    insert(next-to-run, ready-queue);
    save-context();
    schedule(); /* dispatch invoked*/
    load-context( ) }
else scb.counter ++1;
Enable-interrupt
            
```

18

Immediate Priority Inheritance:

=Highest Locker's Priority Protocol (HLP)
 =Priority Protect Protocol (PPP)

- Adopted in Ada95 (protected object), POSIX mutexes
- Idea:** define the ceiling $C(S)$ of a semaphore S to be the highest priority of all tasks that use S during execution. Note that $C(S)$ can be calculated statically (off-line).

19

Run-time behaviour of HLP

- Whenever a task succeeds in holding a semaphore S , its priority is changed dynamically to the maximum of its current priority and $C(S)$.
- When it finishes with S , it sets its priority back to what it was before


20

Example

	priority	use
Task 1	H	S3
Task 2	M	S1, S
Task 3	L	S1, S2
Task 4	Lower	S2, S

$C(S1)=M$
 $C(S2)=L$
 $C(S3)=H$
 $C(S)=M$


Example: Highest Locker's Priority Protocol



21

22


Property 1: Deadlock free (HLP)



Once task 2 gets S1, it runs with pri H, task 1 will be blocked (no chance to get S2 before task 2)

23

Property 2: Tasks will be blocked at most once



24

HLP: Blocking time calculation

- Let
 - $CS(k,S)$ denote the computing time for the critical section that task k uses semaphore S .
 - $Use(S)$ is the set of tasks using S
- Then the maximal blocking time B for task i is as follows:
 - $B = \max\{CS(k,S) \mid i, k \in Use(S), pr(k) < pr(i) \leq C(S)\}$

25

Implementation of HLP

- Calculate the ceiling for all semaphores
- Modify SCB
- Modify P and V-operations

26

Semaphore Control Block for HLP

counter
queue
Pointer to next SCB
Ceiling

27

P-operation with HLP

- P(scb):


```

Disable-interrupt;
If scb.counter > 0 then
  { scb.counter - -1;
    save(current-task.priority);
    current-task.priority := Ceiling(scb) }
else
  {save-context();
   current-task.state := blocked
   insert(current-task, scb.queue);
   dispatch();
   load-context() }
Enable-interrupt
            
```

28

V-operation with HLP

- V(scb):


```

Disable-interrupt;
current-task.priority := get(previous-priority)
If not-empty(scb.queue) then
  next-to-run := get-first(scb.queue);
  next-to-run.state := ready;
  next-to-run.priority := Ceiling(scb);
  insert(next-to-run, ready-queue);
  save-context();
  schedule(); /* dispatch invoked */
  load-context();
end then
else scb.counter ++1;
end else
Enable-interrupt
            
```

29

Properties of HLP: + and -

- Bounded priority inversion
- Deadlock free (+), **Why?**
- Number of blocking = 1 (+), **Why?**
- HLP is a simplified version of PCP (+)
- The extreme case of HLP=NPP (-)
 - E.g when the highest priority task uses all semaphores, the lower priority tasks will inherit the highest priority

30

Summary

	NPP	BIP	HLP
Bounded Priority Inversion	yes	yes	yes
Avoid deadlock	yes	no	yes
Avoid Un-necessary blocking	no	yes	yes/no
Blocking time calculation	Easy	hard	easy

31


Priority Ceiling Protocol (combining HLP and BIP)

- Each semaphore S has a Ceiling $C(S)$
- Run-time behaviour:
 - Assume that S is the semaphore with highest ceiling locked by other tasks currently: $C(S)$ is "the current system priority"
 - If A wants to lock a semaphore (not necessarily S), it must have a strictly higher priority than $C(S)$ i.e. $P(A) > C(S)$. Otherwise A is blocked, and it transmits its priority(+ ϵ) to the task currently holding S

32

Example: PCP

A: ...P(S1)...V(S1)... Prio(A)=H
 B: ...P(S2)...P(S3)...V(S3)...V(S2)... Prio(B)=M C(S1)=H
 C: ...P(S3)...P(S2)...V(S2)...V(S3) Prio(C)=L C(S2)=C(S3)=M



33

PCP: Blocking time calculation

- Let
 - $CS(k,S)$ denote the computing time for the critical section that task k uses semaphore S .
 - $Use(S)$ is the set of tasks using S
- The maximal blocking time for task i :
 - $B = \max\{CS(k,S) \mid i,k \in Use(S), pr(k) < pr(i) \leq C(S)\}$

34

Exercise: implementation of PCP

- Implement P,V-operations that follow PCP
- (this is not so easy)

35

Properties of PCP: + and -

- Bounded priority inversion (+)
- Deadlock free (+)
- Number of blocking = 1 (+)
- Better response times for high priority tasks (+)
 - Avoid un-necessary blocking
- Not easy to implement (-)

36

Summary

	NPP	BIP	HLP	PCP
Bounded Priority Inversion	yes	yes	yes	yes
Avoid deadlock	yes	no	yes	yes
Avoid Un-necessary blocking	no	yes	yes/no	yes
Blocking time calculation	easy	hard	easy	easy
Number of blocking	1	>1	1	1
Implementation	easy	easy	easy	hard

37