

Spin Locks

- Busy waiting (*spinning*) is injurious on mono-processors!
- On *multi-processors*, spinning may be faster than context switch
- If process holding lock is descheduled, spinning wastes CPU time
- Locks with *bounded spin* may be provided by OS.

Caveats of Low-level Sharing

- Consider `var x, y : int; x := 1; y := 2`
- *x* and *y* may not be assigned as expected:
 - The *compiler* may reorder the assignments
 - The *processor* may reorder the store instructions
 - The *memory system* may reorder the write operations
 - The *compiler/processor* may keep values in registers
- Proper ordering must be ensured by *synchronization instructions*

Java low-level memory model

- Access to variables (except long and double) is *atomic*
- Monitor entry/exit synchronize memory operations
- Access to volatile variables is always synchronized

Critical Regions with Shared Variables

Test-and-Set Solution

- Let `TS(var x) = {var r; (r,x) := (x,1); return r}`
- `var lock : int := 0;`

```
process P1
loop
  while TS(lock) = 1 do skip;
  critical section1;
  lock := 0;
  noncritical section1
end loop
```

```
process P2
loop
  while TS(lock) = 1 do skip;
  critical section2;
  lock := 0;
  noncritical section2
end loop
```

Critical Regions with Shared Variables

Non-spinning Test-and-Set Solution

- Let `TS(var x) = {var r; (r,x) := (x,1); return r}`
- `var lock : int := 0; s : semaphore := 0;`

```
process P1
loop
  while TS(lock) = 1 do P(s);
  critical section1;
  lock := 0;
  V(s);
  noncritical section1
end loop
```

```
process P2
loop
  while TS(lock) = 1 do P(s);
  critical section2;
  lock := 0;
  V(s);
  noncritical section2
end loop
```

- Uses at least one semaphore operation for each entry/exit
- Semaphore should be binary

Critical Regions with Shared Variables

Non-spinning Test-and-Set Attempt

- Let $\text{TS}(\text{var } x) = \langle \text{var } r; (r, x) := (x, 1); \text{return } r \rangle$
- $\text{var lock : int} := 0; s : \text{semaphore} := 0; \text{waiting : bool} := \text{false};$
 process P_1
loop
 while $\text{TS}(\text{lock}) = 1$ do
 {waiting := true; $P(s)$ };
 critical section₁;
 lock := 0;
 if waiting then
 {waiting := false; $V(s)$ };
 noncritical section₁
 end loop
 process P_2
loop
 while $\text{TS}(\text{lock}) = 1$ do
 {waiting := true; $P(s)$ };
 critical section₂;
 lock := 0;
 if waiting then
 {waiting := false; $V(s)$ };
 noncritical section₂
 end loop

- Suffers from *race conditions*

Futex (Fast user-space mutual exclusion)

Idea

- Combine user-space atomic operations with OS-based blocking
- Atomic test and wait using *read-check-modify* principle

Data Type

- $\text{type futex} = \text{struct} \{ \underbrace{\text{val : int}}_{\text{userspace}}; \underbrace{\text{queue : condition_queue}}_{\text{kernelspace}} \}$
- Fields related by hash-table based *mapping*

Kernel operations

- $\text{wait_if_eq}(\text{var } f : \text{futex}, \text{old} : \text{int}) \triangleq \langle \text{if } f.\text{val} = \text{old} \text{ then } \text{wait}(f.\text{queue}) \rangle$
- $\text{wake}(\text{var } f : \text{futex}, n : \text{int}) \triangleq \langle \text{for } 1..n \text{ do } \text{signal}(f.\text{queue}) \rangle$

Futex Example: Critical Region

- Let `TS(var x) = <var r; (r,X) := (x,1); return r>`
`INC(var x : int) = <x := x + 1; return x>`
`DEC(var x : int) = <x := x - 1; return x>`
- `var f : futex; waiting : int := 0;`
`f.val := 0;`
`process P1`
`loop`
`while TS(f.val) = 1 do`
`{INC(waiting);`
`wait_if_eq(f,1);`
`DEC(waiting); }`
`critical section1;`
`f.val := 0;`
`if waiting > 0 then wake(f,1);`
`noncritical section1`
`end loop`
- `process P2`
`loop`
`while TS(f.val) = 1 do`
`{INC(waiting);`
`wait_if_eq(f,1);`
`DEC(waiting); }`
`critical section2;`
`f.val := 0;`
`if waiting > 0 then wake(f,1);`
`noncritical section2`
`end loop`