


## Mars Pathfinder Incident

- Landing on July 4, 1997
- "...experiences software glitches..."
- Pathfinder experiences repeated RESETs after starting gathering of meteorological data.



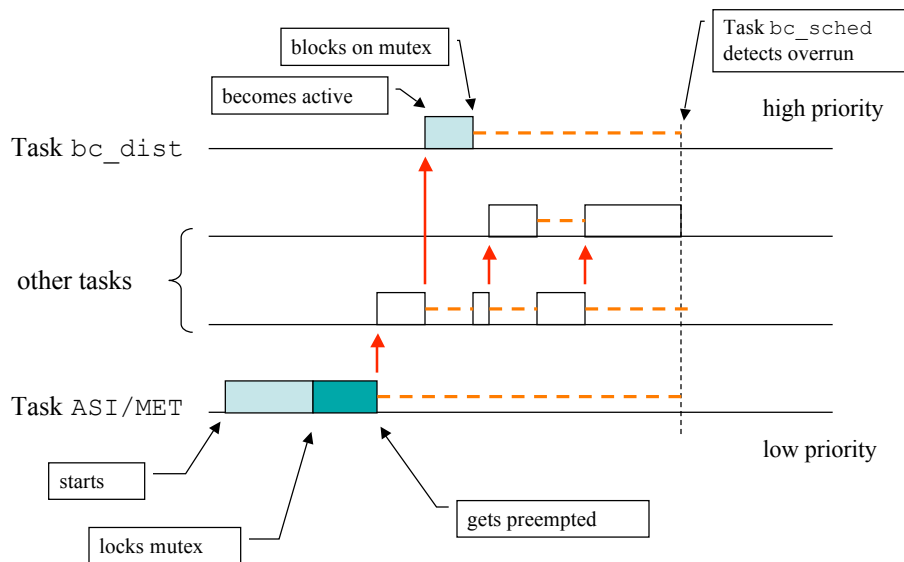
- RESETs generated by watchdog process.
- Timing overruns caused by priority inversion.



- Resources:

`research.microsoft.com/~mbj/Mars_Pathfinder/Mars_Pathfinder.html`

## Priority Inversion on Mars Pathfinder



## Mars Pathfinder: Resolution


- “Faster, better, cheaper” had NASA and JPL using “shrink-wrap” hardware (IBM RS6000) and software (Wind River VxWorks RTOS).
- Logging designed into VxWorks enabled NASA and Wind River to reproduce the failure on Earth. This reproduction made the priority inversion obvious.
- NASA patched the lander’s software to enable priority inheritance.

## Resource Access

- Processor(s)
  - $m$  types of serially reusable resources  $R_1, \dots, R_m$
  - An execution of a job  $J_i$  requires:
    - a processor for  $e_i$  units of time
    - some resources for exclusive use
- Resources
  - serially Reusable: Allocated to one job at a time. Once allocated, held by the job until no longer needed.
  - examples: semaphores, locks, servers, ...
  - operations:  
`lock(Ri) -----<critical section>----- unlock(Ri)`
  - resources allocated non-preemptively
  - critical sections properly nested

## Preemption During Critical Sections


Example:



- Negative effect on schedulability and predictability.


## Unpredictability: Scheduling Anomalies

- Example:  $T_1 = (c_1=2, e_1 = 5, p_1 = 8)$   $T_2 = (4, 7, 22)$   $T_3 = (4, 6, 26)$




- Shorten critical section of  $T_3$ :

$$T_1 = (c_1=2, e_1 = 5, p_1 = 8) \quad T_2 = (4, 7, 22) \quad T_3 = (2.5, 6, 26)$$



## Disallow Process Preemption in CS





- Analysis identical to analysis with non-preemptable portions
- Define:  $\beta$  = maximum duration of all critical sections
- Task  $T_i$  is schedulable if

$$\sum_{k=1} e_k + \frac{\beta}{p_i} = U_X(i)$$

$X$ : scheduling algorithm

- Problem: critical sections can be rather long.

## Priority Inheritance



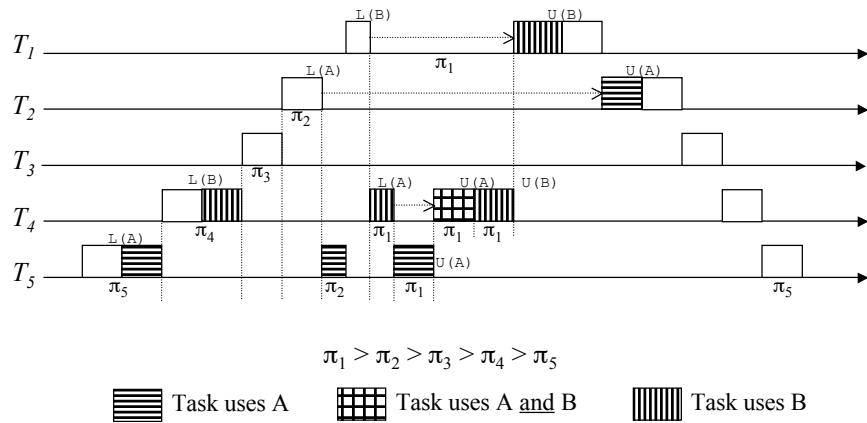
## Terminology

- A job is directly blocked when it requests a resource  $R_i$ , i.e. executes a  $\text{lock}(R_i)$ , but no resource of type  $R_i$  is available.
- The scheduler grants the lock request, i.e. allocates the requested resource to the job, according to the resource allocation rules, as soon as the resources become available.
- $J'$  directly blocks  $J$  if  $J'$  holds some resources that  $J$  has requested.
- Priority Inheritance:
  - Basic strategy for controlling priority inversion:  
Let  $\pi$  be the priority of  $J$   
and  $\pi'$  be the priority of  $J'$   
and  $\pi' < \pi$   
then the priority of  $J'$  is set to  $\pi$  whenever  $J'$  directly blocks  $J$ .
- New forms of blocking may be introduced by the resource management policy to control priority inversion and/or prevent deadlocks.


## Basic Priority-Inheritance Protocol

- Jobs that are not blocked are scheduled according to a priority-driven algorithm preemptively on a processor.
- Priorities of tasks are fixed, except for the conditions described below:
  - A job  $J$  requests a resource  $R$  by executing  $\text{lock}(R)$
  - If  $R$  is available, it is allocated to  $J$ .  $J$  then continues to execute and releases  $R$  by executing  $\text{unlock}(R)$
  - If  $R$  is allocated to  $J'$ ,  $J'$  directly blocks  $J$ . The request for  $R$  is denied.
  - However: Let  $\pi =$  priority of  $J$  when executing  $\text{lock}(R)$   
 $\pi' =$  priority of  $J'$  at the same time
  - For as long as  $J'$  holds  $R$ , its priority is  $\max(\pi, \pi')$  and returns to  $\pi'$  when it releases  $R$ .
  - That is:  $J'$  inherits the priority of  $J$  when  $J'$  directly blocks  $J$  and  $J$  has a higher priority.
- Priority Inheritance is transitive.

## Example: Priority Inheritance Protocol




## Example: Priority Inheritance Protocol



**Problem:** If  $T_5$  tries to lock (B) while it has priority  $\pi_5$ , we have a deadlock!

## Properties of PIP

- It does not prevent deadlock.
- Task can be blocked directly by a task with a lower priority at most once, for the duration of the (outmost) critical section.
- Consider a task whose priority is higher than  $n$  other tasks:



- Each of the lower-priority tasks can directly block the task at most once.
- A task outside the critical section cannot directly block a higher-priority task.


## Priority Ceiling Protocol

- Assumptions:
  - Priorities of tasks are fixed
  - Resources required by tasks are known
- Definition (Priority Ceiling of  $R$ )  
Priority Ceiling  $\Pi_R$  of  $R$  = highest priority of all tasks that will request  $R$ .
- Any task holding  $R$  may have priority  $\Pi_R$  at some point; either its own priority is  $\Pi_R$ , or it inherits  $\Pi_R$ .
- Motivation:
  - Suppose there are resource  $A$  and  $B$ .
  - Both  $A$  and  $B$  are available.  $T_1$  requests  $A$ .
  - $T_2$  requests  $B$  after  $A$  is allocated.
    - If  $\pi_2 > \Pi_A$ :  $T_1$  can never preempt  $T_2 \Rightarrow B$  should be allocated to  $T_2$ .
    - If  $\pi_2 \leq \Pi_A$ :  $T_1$  can preempt  $T_2$  (and also request  $B$ ) at some later time.  $B$  should not be allocated to  $T_2$ , to avoid deadlock.

## Priority Ceiling Protocol

- Same as the basic Priority Inheritance Protocol, except for the following:
- When a task  $T$  requests for allocation of a resource  $R$  by executing `lock(R)` :
  - The request is denied if
    1.  $R$  is already allocated to  $T'$ . ( $T'$  directly blocks  $T$ .)
    2. The priority of  $T$  is not higher than all priority ceilings for resources allocated to tasks other than  $T$  at the time. (These tasks block  $T$ .)
  - Otherwise,  $R$  is allocated to  $T$ .
- When a task blocks other tasks, it inherits the highest of their priorities.

## Example: Priority Ceiling Protocol






## Example: Priority Ceiling Protocol

$$\pi_1 > \pi_2 > \pi_3 > \pi_4 > \pi_5$$


$$\Pi_A = \pi_2, \quad \Pi_B = \pi_1$$




## Types of Blocking

- **Blocking:** A higher-priority task waits for a lower-priority task.
- A task  $T_H$  can be blocked by a lower-priority task  $T_L$  in three ways:


- directly, i.e.




- when  $T_L$  inherits a priority higher than the priority  $\pi_H$  of  $T_H$ .



- When  $T_H$  requests a resource the priority ceiling of resources held by  $T_L$  is equal to or higher than  $\pi_H$ :




## Closer Look At Avoidance Blocking



## Stack Sharing

- Sharing of the stack among tasks eliminates stack space fragmentation and so allows for memory savings:



- However:
  - Once job is preempted, it can only resume when it returns to be on top of stack.
  - Otherwise, it may cause a deadlock.
  - Stack becomes a resource that allows for “one-way preemption”.


## Stack-Sharing Priority-Ceiling Protocol

- To avoid deadlocks: Once execution begins, make sure that job is not blocked due to resource access.
- Otherwise: Low-priority, preempted, jobs may re-acquire access to CPU, but can not continue due to unavailability of stack space.
- Define:  $\hat{\Pi}(t)$ : highest priority ceiling of all resources currently allocated,  
If no resource allocated,  $\hat{\Pi}(t) = \Omega$ .

Protocol:

1. *Update Priority Ceiling*: Whenever all resources are free,  $\Pi(t) = \Omega$ . The value of  $\Pi(t)$  is updated whenever resource is allocated or freed.
2. *Scheduling Rule*: After a job is released, it is blocked from starting execution until its assigned priority is higher than  $\hat{\Pi}(t)$ . At all times, jobs that are not blocked are scheduled on the processor in a priority-driven, preemptive fashion according to their assigned priorities.
3. *Allocation Rule*: Whenever a job requests a resource, it is allocated the resource.

## SSPCP



## Stack-Sharing Priority Ceiling Protocol

- The Stack-Based Priority-Ceiling Protocol is **deadlock-free**:
  - When a job begins to execute, all the resources it will ever need are free.
  - Otherwise,  $\Pi(t)$  would be higher or equal to the priority of the job.
  - Whenever a job is preempted, all the resources needed by the preempting job are free.
  - The preempting job can complete, and the preempted job can resume.
- Worst-case blocking time of Stack-Based Protocol is the same as for Basic Priority Ceiling Protocol.
- Stack-Based Protocol smaller context-switch overhead (2 CS) than Priority Ceiling Protocol (4 CS)
  - Once execution starts, job cannot be blocked.

## Ceiling-Priority Protocol

- Stack-Based Protocol does not allow for self-suspension
  - Stack is shared resource
- Re-formulation for multiple stacks (no stack-sharing) straightforward:

### Ceiling-Priority Protocol

#### *Scheduling Rules:*

1. Every job executes at its assigned priority when it does not hold resources.
2. Jobs of the same priority are scheduled on FIFO basis.
3. Priority of jobs holding resources is the highest of the priority ceilings of all resources held by the job.

#### *Allocation Rule:*

- Whenever a job requests a resource, it is allocated the resource.

## Priority-Ceiling Locking in Ada 9X

- Task definitions allow for a pragma Priority as follows:  
`pragma Priority(expression)`
- Task priorities:
  - *base priority*: priority defined at task creation, or dynamically set with `Dynamic_Priority.Set_Priority()` method.
  - *active priority*: base priority or priority inherited from other sources (activation, rendez-vous, protected objects).
- Priority-Ceiling Locking:
  - Every protected object has a *ceiling priority*: Upper bound on active priority a task can have when it calls a protected operation on objects.
  - While task executes a protected action, it inherits the ceiling priority of the corresponding protected object.
  - When a task calls a protected operation, a check is made that its active priority is not higher than the ceiling of the corresponding protected object. A Program Error is raised if this check fails.

## Priority-Ceiling Locking in Ada 9X: Implementation

- Efficient implementation possible that does not rely on explicit locking.
- Mutual exclusion is enforced by priorities and priority ceiling protocol only.
- We show that Resource  $R$  can never be requested by Task  $T2$  while it is held by Task  $T1$ .
- Simplified argument:
  - $AP(T2)$  can never be higher than  $C(R)$ . Otherwise, run-time error would occur.  $\Rightarrow AP(T2) \leq C(R)$
  - As long as  $T1$  holds  $R$ , it cannot be blocked.
    - Therefore, for  $T2$  to request  $R$  after  $T1$  seized it,  $T1$  must have been preempted (priority of  $T1$  does not change while  $T1$  is in ready queue).
  - For  $T2$  to request  $R$  while  $T1$  is in ready queue,  $T2$  must have higher active priority than  $T1$ .  $\Rightarrow AP(T2) \leq C(R)$
  - $T1$  is holding  $R$   $\Rightarrow C(R) \leq AP(T1) < AP(T2)$ 
    - Before  $T2$  requests  $R$ ,  $T2$ 's priority must drop to  $\leq C(R)$
    - Case 1:  $AP(T2)$  drops to below  $AP(T1)$   $\Rightarrow T2$  preempted
    - Case 2:  $AP(T2)$  drops to  $AP(T1)$   $\Rightarrow T2$  must yield to  $T1$  (by rule)