


# Real Time Systems1

## lecture 9

### Priority Inheritance Protocol



# Priority Inheritance Protocol

- ▶ The *Priority Inheritance Protocol* (PIP), proposed by Sha, Rajkumar and Lehoczky [SRL90], avoids unbounded priority inversion by modifying the priority of those tasks that cause blocking.
- ▶ In particular, when a task  $\tau_i$  blocks one or more higher-priority tasks, it temporarily assumes (*inherits*) the highest priority of the blocked tasks.

# Priority Inheritance Protocol

- ▶ This prevents medium-priority tasks from preempting  $\tau_i$  and prolonging the blocking duration experienced by the higher-priority tasks.
- ▶ The Priority Inheritance Protocol can be defined as follows:
  - ▶ 1. Tasks are scheduled based on their active priorities. Tasks with the same priority are executed in a First Come First Served discipline.

# Priority Inheritance Protocol

- ▶ 2. When task  $\tau_i$  tries to enter a critical section  $z_{i,k}$  and resource  $R_k$  is already held by a lower-priority task  $\tau_j$ , then  $\tau_i$  is blocked.  $\tau_i$  is said to be blocked by the task  $\tau_j$  that holds the resource. Otherwise,  $\tau_i$  enters the critical section  $z_{i,k}$ .
- ▶ 3. When a task  $\tau_i$  is blocked on a semaphore, it transmits its active priority to the task, say  $\tau_j$ , that holds that semaphore.

## Priority Inheritance Protocol

- ▶ Hence,  $\tau_j$  resumes and executes the rest of its critical section with a priority  $p_j = p_i$ . Task  $\tau_j$  is said to *inherit* the priority of  $\tau_i$ . In general, a task inherits the highest priority of the tasks it blocks. That is, at every instant,  $p_j(R_k) = \max\{P_j, \max_h \{P_h / \tau_h \text{ is blocked on } R_k\}\}$ .
- ▶ 4. When  $\tau_j$  exits a critical section, it unlocks the semaphore, and the highest-priority task blocked on that semaphore, if any, is awakened.

# Priority Inheritance Protocol

- ▶ The active priority of  $\tau_j$  is updated as follows: if no other tasks are blocked by  $\tau_j$ ,  $p_j$  is set to its nominal priority  $P_j$ ; otherwise it is set to the highest priority of the tasks blocked by  $\tau_j$ , according to Equation (7.8).
- ▶ 5. Priority inheritance is transitive; that is, if a task  $\tau_3$  blocks a task  $\tau_2$ , and  $\tau_2$  blocks a task  $\tau_1$ , then  $\tau_3$  inherits the priority of  $\tau_1$  via  $\tau_2$ .


# Priority Inheritance Protocol Example

- ▶ There are five jobs and two resources *Black* and *Shaded*. The parameters of the jobs and their critical sections are listed in part (a). As usual, jobs are indexed in decreasing order of their priorities:
- ▶ The priority  $\pi_i$  of  $J_i$  is  $i$ , and the smaller the integer, the higher the priority. In the schedule in part (b) of this figure, black boxes show the critical sections when the jobs are holding *Black*.
- ▶ Shaded boxes show the critical sections when the jobs are holding *Shaded*

# Priority Inheritance Protocol Example

Job	$r_i$	$e_i$	$\pi_i$	Critical Sections
$J_1$	7	3	1	[ <i>Shaded</i> ; 1]
$J_2$	5	3	2	[ <i>Black</i> ; 1]
$J_3$	4	2	3	
$J_4$	2	6	4	[ <i>Shaded</i> ; 4] [ <i>Black</i> ; 1.5]]
$J_5$	0	6	5	[ <i>Black</i> ; 4]

(a)



(b)

FIGURE 8-8 Example illustrating transitive inheritance of priority inheritance. (a) Parameters of jobs. (b) Schedule under priority inheritance.



# Priority Inheritance Protocol

- ▶ 1. At time 0, job  $J_5$  becomes ready and executes at its assigned priority 5. At time 1, it is granted the resource *Black*.
- ▶ 2. At time 2,  $J_4$  is released. It preempts  $J_5$  and starts to execute.
- ▶ 3. At time 3,  $J_4$  requests *Shaded*. *Shaded*, being free, is granted to the job. The job continues to execute.
- ▶ 4. At time 4,  $J_3$  is released and preempts  $J_4$ . At time 5,  $J_2$  is released and preempts  $J_3$ .

# Priority Inheritance Protocol

- ▶ 5. At time 6,  $J_2$  executes  $L(Black)$  to request *Black*;  $L(Black)$  fails because *black* is in use by  $J_5$ .  $J_2$  is now directly blocked by  $J_5$ .
- ▶ According to rule 3,  $J_5$  inherits the priority 2 of  $J_2$ .
- ▶ Because  $J_5$ 's priority is now the highest among all ready jobs,  $J_5$  starts to execute
- ▶ 6.  $J_1$  is released at time 7. Having the highest priority 1, it preempts  $J_5$  and starts to execute.

# Priority Inheritance Protocol

- ▶ 7. At time 8,  $J_1$  executes  $L(Shaded)$ , which fails, and becomes blocked. Since  $J_4$  has  $Shaded$  at the time, it directly blocks  $J_1$  and, consequently, inherits  $J_1$ 's priority 1.
- ▶  $J_4$  now has the highest priority among the ready jobs  $J_3$ ,  $J_4$ , and  $J_5$ . Therefore, it starts to execute.
- ▶ 8. At time 9,  $J_4$  requests the resource *Black* and becomes directly blocked by  $J_5$ .
- ▶ At this time the current priority of  $J_4$  is 1, the priority it has inherited from  $J_1$  since time 8.

# Priority Inheritance Protocol

- ▶ Therefore,  $J_5$  inherits priority 1 and begins to execute.
- ▶ 9. At time 11,  $J_5$  releases the resource *Black*. Its priority returns to 5, which was its priority when it acquired *Black*.
- ▶ The job with the highest priority among all unblocked jobs is  $J_4$ .
- ▶ Consequently,  $J_4$  enters its inner critical section and proceeds to complete this and the outer critical section.

# Priority Inheritance Protocol

- ▶ 10. At time 13,  $J_4$  releases *Shaded*. The job no longer holds any resource; its priority returns to 4, its assigned priority.
- ▶  $J_1$  becomes unblocked, acquires *Shaded*, and begins to execute.
- ▶ 11. At time 15,  $J_1$  completes.  $J_2$  is granted the resource *Black* and is now the job with the highest priority. Consequently, it begins to execute.
- ▶ 12. At time 17,  $J_2$  completes. Afterwards, jobs  $J_3$ ,  $J_4$ , and  $J_5$  execute in turn to completion

# Priority Inheritance Protocol

- ▶ From this example, we notice that a high-priority task can experience two kinds of blocking:
- ▶ **Direct blocking.** It occurs when a higher-priority task tries to acquire a resource already held by a lower-priority task. Direct blocking is necessary to ensure the consistency of the shared resources.
- ▶ **Push-through blocking.** It occurs when a medium-priority task is blocked by a low-priority task that has inherited a higher priority from a task it directly blocks. Push-through blocking is necessary to avoid unbounded priority inversion.

# Priority Inheritance Protocol

- ▶ Note that in most situations when a task exits a critical section, it resumes the priority it had when it entered. This, however, is not always true.
- ▶ Consider the example illustrated in Figure 7.9. Here, task  $\tau_1$  uses a resource  $R_a$  guarded by a semaphore  $S_a$ , task  $\tau_2$  uses a resource  $R_b$  guarded by a semaphore  $S_b$ , and task  $\tau_3$  uses both resources in a nested fashion ( $S_a$  is locked first).
- ▶ At time  $t_1$ ,  $\tau_2$  preempts  $\tau_3$  within its nested critical section; hence, at time  $t_2$ , when  $\tau_2$  attempts to lock  $S_b$ ,  $\tau_3$  inherits its priority,  $P_2$ .

# Priority Inheritance Protocol

- ▶ Similarly, at time  $t_3$ ,  $\tau_1$  preempts  $\tau_3$  within the same critical section, and at time  $t_4$ , when  $\tau_1$  attempts to lock  $S_a$ ,  $\tau_3$  inherits the priority  $P_1$ .
- ▶ At time  $t_5$ , when  $\tau_3$  unlocks semaphore  $S_b$ , task  $\tau_2$  is awakened but  $\tau_1$  is still blocked; hence,  $\tau_3$  continues its execution at the priority of  $\tau_1$ .
- ▶ At time  $t_6$ ,  $\tau_3$  unlocks  $S_a$  and, since no other tasks are blocked,  $\tau_3$  resumes its original priority  $P_3$ .



# Priority Inheritance Protocol

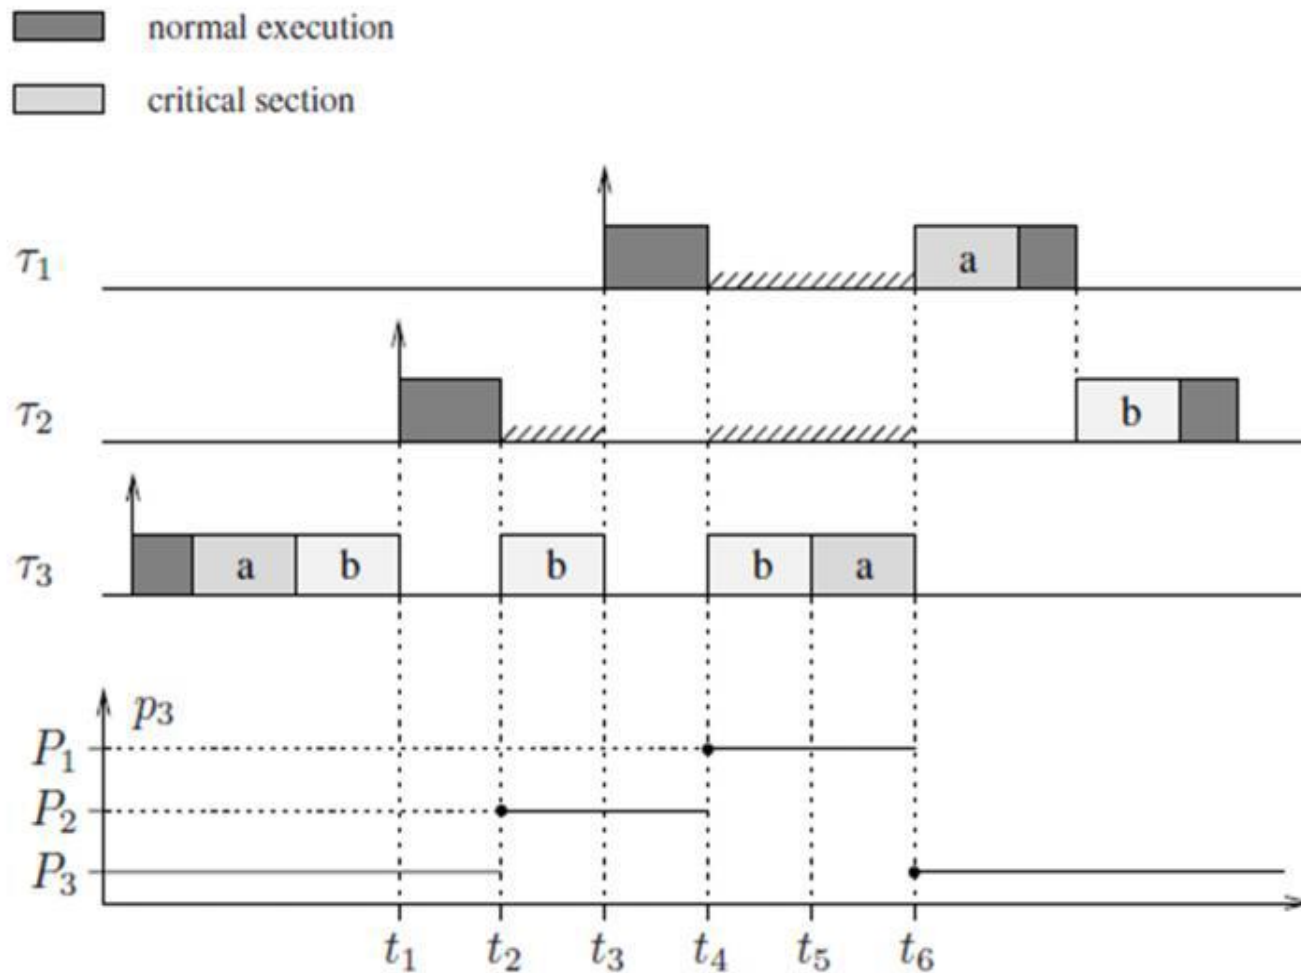


Figure 7.9 Priority inheritance with nested critical sections.

# Priority Inheritance Protocol

- ▶ An example of transitive priority inheritance is shown in Figure 7.10.
- ▶ Here, task  $\tau_1$  uses a resource  $R_a$  guarded by a semaphore  $S_a$ , task  $\tau_3$  uses a resource  $R_b$  guarded by a semaphore  $S_b$ , and task  $\tau_2$  uses both resources in a nested fashion ( $S_a$  protects the external critical section and  $S_b$  the internal one).
- ▶ At time  $t_1$ ,  $\tau_3$  is preempted within its critical section by  $\tau_2$ , which in turn enters its first critical section (the one guarded by  $S_a$ ), and at time  $t_2$  it is blocked on semaphore  $S_b$ . As a consequence,  $\tau_3$  resumes and inherits the priority  $P_2$ .

# Priority Inheritance Protocol

- ▶ At time  $t_3$ ,  $\tau_3$  is preempted by  $\tau_1$ , which at time  $t_4$  tries to acquire  $R_a$ . Since  $S_a$  is locked by  $\tau_2$ ,  $\tau_2$  inherits  $P_1$ . However,  $\tau_2$  is blocked by  $\tau_3$ ; hence, for transitivity,  $\tau_3$  inherits the priority  $P_1$  via  $\tau_2$ .
- ▶ When  $\tau_3$  exits its critical section, no other tasks are blocked by it; thus it resumes its nominal priority  $P_3$ . Priority  $P_1$  is now inherited by  $\tau_2$ , which still blocks  $\tau_1$  until time  $t_6$ .

# Priority Inheritance Protocol




Figure 7.10 Example of transitive priority inheritance.