

3.b Task interactions and blocking

Where we allow tasks to contend for shared resources, we see what access control protocols can do to prevent disaster, and discuss their pros and cons

Inhibiting preemption /1

- In the real world, certain application procedures do *not* allow preemption
 - The execution of *non-reentrant* code shared by multiple jobs, whether directly (by application-level calls) or indirectly (within system calls), cannot tolerate preemption
 - A reentrant procedure allows calls to itself to overlap – owing to preemption – while *safeguarding* the local context of every such call
 - Reentrancy needs the stack, which may need a lot of memory ...
- For reasons of integrity or efficiency, some system-level activities should *not* be preempted
 - Operations on hardware devices do *not* allow preemption
- At its simplest, preemption is inhibited by just disabling dispatching
 - Uh: how does that happen?



How does preemption happen?

- All the CPU does is to repeat a simple cycle of basic micro-operations forever (until stopped)
 - Fetch, Decode, Read, Execute, Write
 - One such cycle for each instruction of the application program
 - The throughout of this cycle (# of instructions executed / unit of time) is increased by *pipelining* those micro-operations
 - This is why branching and jumping are bad news!
- Those micro-operations *cannot* be interrupted
 - Electrons save no context: if you stop, you lose the whole pipeline!
- The *only* way to preempt program execution is to *prefix* a “check-for-interrupt” clause to the Fetch stage
 - The source of an interrupt is an event that needs attention (which *cannot* be given by the current program): the execution must move elsewhere, which *is* preemption
 - This is why, if an interrupt request is found asserted, the CPU is *hijacked*
- Omitting that check or not allowing interrupt requests to register, effectively *disables* preemption



Inhibiting preemption /2

- A higher-priority job J_h that, at its release time, finds a lower-priority job J_l executing with disabled preemption, gets **blocked** for a time duration that depends on J_l
 - Under FPS, this is a flagrant case of **priority inversion**
- The feasibility of J_h now depends on J_l !
 - Under FPS, this form of blocking for J_i is upper-bounded by $B_i(np) = \max_{k=i+1, \dots, n}(\theta_k)$ where $\theta_k \leq e_k$ is the longest span of J_k 's non-preemptible execution
 - This cost is paid by of J_i only *once* per release because lower-priority jobs *cannot* preempt J_i

The drag of self suspension /1

- A job J_i that invokes suspending operations or self suspends (`sleep()`), suffers a time penalty that worsens its response time
- J_i incurs a degenerate form of blocking that can be bounded as $B_i(ss) = \max(\delta_i) + \sum_{k=1, \dots, i-1} \min(e_k, \max(\delta_k))$
 - $\max(\delta_i)$ is the longest duration of self suspension by job J_i
 - The \sum term is the cumulative interference caused by self-suspending high-priority jobs that may become ready during the (shifted) busy period of J_i
 - Every J_k might resume from self-suspension exactly when J_i does, and therefore interfere up to $\max(\delta_k)$ but never more than e_k
- In general, a job J_i that self suspends K times during execution incurs total blocking $B_i = B_i(ss) + (K + 1)B_i(np)$
 - As $B_i(np)$ is potentially incurred at at *every* resumption

2019/2020 UniPD - T. Vardanega Real-Time Systems 156 of 538

The drag of self suspension /2

- Self suspension with independent tasks on single-core processors causes *scheduling anomalies*
 - Deadlines can be missed when task utilization or suspension delays are *decreased*
- **Example:** consider a feasible task set under EDF
 - $\tau_1 = \{0, 10, (2, 2, 2), 6\}$
 - $\tau_2 = \{5, 10, (1, 1, 1), 4\}$
 - $\tau_3 = \{7, 10, (1, 1, 1), 3\}$
 - τ_3 would miss its deadline if τ_1 's execution or suspension was 1 time unit shorter

2019/2020 UniPD - T. Vardanega Real-Time Systems 157 of 538

The drag of self suspension /3

Under Rate Monotonic Scheduling

Selfish self-suspension

2019/2020 UniPD - T. Vardanega Real-Time Systems 158 of 538

The drag of self suspension /4

(ϕ_i, p_i, e_i, D_i) $\tau_1 = \{0, 4, 2.5, 4\}, \tau_2 = \{3, 10, 2, 10\} U = 0.875$

τ_1 self-suspends for 1.5 τ_2 misses its deadline

$B_2(ss) = 0 + \min(2.5, 1.5) = 1.5$ is a pessimistic upperbound!
 With $\phi_2 = 3$, the actual blocking for τ_2 in $[3, 10]$ reduces to 1
 But still $B_2(ss) = 1 > \sigma_{2,1}(0) = 0.5$

2019/2020 UniPD - T. Vardanega Real-Time Systems 159 of 538

Access contention


- Concurrent access to shared resources causes potential for contention that needs specialized control
 - A **resource access control protocol**
- Such a protocol specifies (1) when, (2) on what conditions, (3) in which order, a resource access request may be granted
 - Access contention situations may cause **priority inversion** to arise (see following examples)

Example /1

$(\varphi_i, p_i, e_i, D_i)$ Max use of shared resource per execution
 $\tau_1 = \{-, -, 2, 20, \mathbf{R(4)}\}$, $\tau_2 = \{2, -, 3, 17, \mathbf{R(4)}\}$, $\tau_3 = \{6, -, 3, 14, \mathbf{R(2)}\}$

under EDF (periods *not* specified, as they do not matter here)

$\tau_1 :: e; \mathbf{R(4)}; e.$ $\tau_2 :: e; \mathbf{R(4)}; e.$ $\tau_3 :: e; \mathbf{R(2)}; e.$




Example /2

$(\varphi_i, p_i, e_i, D_i)$
 $\tau_1 = \{-, -, 2, 20, \mathbf{R(2.5)}\}$, $\tau_2 = \{2, -, 3, 17, \mathbf{R(4)}\}$, $\tau_3 = \{6, -, 3, 14, \mathbf{R(2)}\}$

under EDF

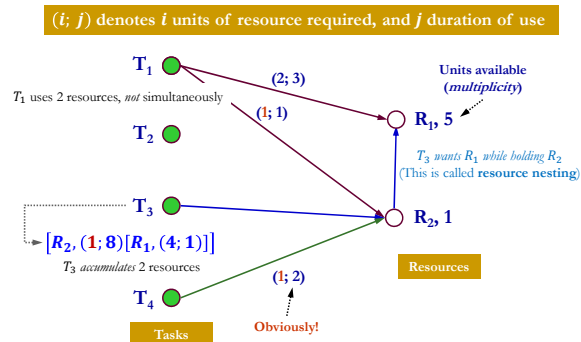
Same as before, except with *shorter* use of R by τ_1



Assumptions and notations

- In order to minimize the extent of interference
 - Jobs should *not* self suspend (directly or indirectly)
 - Jobs can be *preempted*
- We say that job J_h is **directly blocked** by a lower-priority job J_l when
 - J_l is granted exclusive access to a shared resource R
 - J_h has requested R and its request has *not* been granted
- To study the problem we may want to use a **wait-for graph**

Example (wait-for-graph)



2019/2020 UniPD - T. Vardanega

Real-Time Systems

164 of 538

Resource access control [option a]

- **Inhibiting preemption** in critical sections
 - A job that requires access to a resource is *always* granted it
 - A job that has been assigned a resource runs at a priority higher than any other job
 - These two clauses imply each other (why?)
 - They jointly prevent deadlock situations from occurring (why?)
- This protocol causes **bounded** priority inversion
 - At most *once* per job (we already understood why)
 - For a maximum duration of $B_i(rc) = \max_{k=i+1, \dots, n}(C_k)$
 - For job indices in monotonically non-increasing order and C_k denoting the worst-case duration of critical section for job J_k

2019/2020 UniPD - T. Vardanega

Real-Time Systems

165 of 538

Critique of [option a]

- This strategy causes **distributed overhead**
 - *All* jobs – including those that do *not* compete for resource access – incur some time penalty
 - Very unfair: undesirable
- It should be preferable that time overhead was *solely* (or at least mostly) incurred by the jobs that *do* compete for resource access
 - The priority of the job that is granted the resource should be *no less* than that of its *competitor* jobs (but of no other)
 - This principle has two possible realizations
 - One is called **priority inheritance**, the other is called **priority ceiling**
 - We shall now examine how each of them operates

2019/2020 UniPD - T. Vardanega

Real-Time Systems

166 of 538

Resource access control [option b]

- **Basic priority inheritance protocol** (BPIP)
 - The job's priority may vary over time
 - The variation follows inheritance principles
- **Protocol rules**
 - **Scheduling**: jobs are dispatched by preemptive priority-driven scheduling; at release time, they assume their *assigned priority*
 - **Allocation**: when job J requires access to resource R at time t
 - If R is free, R is assigned to J until release
 - If R is busy, the request is denied and J becomes *blocked*
 - **Priority inheritance**: when job J becomes blocked, job J_1 that blocks it takes on J 's current priority as its *inherited priority* and retains it until R is released; at that point J_1 reverts to its previous priority

2019/2020 UniPD - T. Vardanega

Real-Time Systems

167 of 538

Critique of [option b]

- BPIP suffers two forms of blocking
 - **Direct blocking**, owing to resource contention
 - **Inheritance blocking**, owing to priority raising
- Priority inheritance is *transitive*
 - Direct blocking *is* transitive as jobs may need to accumulate resources
- BPIP does *not* prevent deadlock
 - Cyclic blocking proceeds from transitive direct blocking
- BPIP incurs *reducible* distributed overhead
 - Under BPIP, a job may become blocked every time it competes for a shared resource, hence multiple times in the same run
- BPIP needs *no* prior knowledge on which resources are shared
 - It is inherently dynamic, hence usable for open (non real-time) systems

2019/2020 UniPD - T. Vardanega Real-Time Systems 168 of 538

Resource access control [option c]

- **Basic priority ceiling protocol** (BPCP)
 - Similar to BPIP, except that it needs *all* resource requirements to be *statically known*
 - Every resource R is assigned a **priority ceiling attribute** set statically to the highest priority of the jobs that require R
 - At time t , the system has a ceiling $\pi_s(t)$ attribute set to the highest priority ceiling of all resources currently in use
 - If no resource is currently in use at t , $\pi_s(t)$ defaults to $\Omega <$ the lowest priority of all jobs

2019/2020 UniPD - T. Vardanega Real-Time Systems 169 of 538

BPCP protocol rules

- **Scheduling**: jobs are dispatched by preemptive priority-driven scheduling; at release time they assume their assigned priority
- **Allocation**: when job J requests access to resource R at time t
 - If R is already assigned, the request is denied and J becomes blocked
 - If R is free and J 's priority $\pi_J(t) > \pi_s(t)$, the request is granted
 - If J currently owns the resource whose priority ceiling = $\pi_s(t)$, the request is granted
 - Otherwise the request is denied and J becomes blocked Avoidance blocking
- **Priority inheritance**: when job J becomes blocked by job J_l , J_l takes on J 's current priority $\pi_J(t)$ until J_l releases all resources with priority ceiling $> \pi_J(t)$; at that point J_l 's priority reverts to the level that preceded access to those resources

2019/2020 UniPD - T. Vardanega Real-Time Systems 170 of 538

Critique of [option c] /1

- BPCP is *not* greedy (BPIP is!)
 - Under BPCP, a request for a free resource may be denied
- Hence, BPCP causes each job J to incur **three** distinct forms of blocking caused by lower-priority job J_l

2019/2020 UniPD - T. Vardanega Real-Time Systems 171 of 538

Critique of [option c] /2

- **Avoidance blocking** is what makes BPCP not greedy and also prevents deadlock from occurring
 - If, at time t , job J has $\pi_J(t) > \pi_s(t)$ then it must be so that
 - J will never use any of the resources in use at time t
 - So won't all jobs with higher priority than J
- The system ceiling $\pi_s(t)$ determines which jobs can be assigned a resource free at time t without risking deadlock
 - All jobs with priority higher than the system ceiling $\pi_s(t)$
- **Caveat**
 - To stop job J from blocking itself when attempting to accumulate resources, BPCP must grant its request in case $\pi_J(t) \leq \pi_s(t)$, but J at t holds the resources $\{X\}$ whose priority ceiling is $= \pi_s(t)$

2019/2020 UniPD - T. Vardanega Real-Time Systems 172 of 538

Critique of [option c] /3

- BPCP does *not* incur reducible distributed overhead as it does *not* permit transitive blocking
- **Theorem** [Sha & Rajkumar & Lehoczky, 1990]
Under BPCP a job may become blocked for *at most* the duration of *one* critical section
 - Under BPCP, when a job becomes blocked, its blocking can *only* be caused by a single ready job
 - The job that causes others to block cannot itself be blocked
 - Hence BPCP does not permit transitive blocking
 - Demonstration: **By exercise**
- The maximum possible value of that duration for job J_i is termed the **blocking time $B_i(rc)$** due to resource contention
 - $B_i(rc)$ must be accounted for in the schedulability test for J_i

2019/2020 UniPD - T. Vardanega Real-Time Systems 173 of 538

Computing the BPCP blocking time /1

High
Priority
Low

Directly blocked by

	J2	J3	J4	J5	J6
J1		6			2
J2			5		
J3					4
J4					
J5					

Priority-inheritance blocked by

	J2	J3	J4	J5	J6
J1					
J2		6			2
J3			5		2
J4					4
J5					4

Avoidance blocked by

	J2	J3	J4	J5	J6
J1					
J2		6			2
J3			5		2
J4					4
J5					

$B_i(rc) = \max \text{ value in row } J_i \text{ across all tables}$


2019/2020 UniPD - T. Vardanega Real-Time Systems 174 of 538

Computing the BPCP blocking time /2

- Table rows are sorted by priority
 - Jobs are assigned distinct priorities (i.e., no overlap)
- Table “*directly blocked by*” is easy to understand ...
- Table “*priority-inheritance blocked by*”
 - Job J_{i+1} causes direct blocking inherits the blocked job's priority: all jobs with priority lower than the inherited one but higher than J_{i+1} 's suffer blocking
 - The value in cell $[i, k]$ is max across (rows $1, \dots, i - 1$; column k) in Table “*directly blocked by*”
- Table “*avoidance blocked by*”
 - The resource is free but another resource with priority ceiling higher than your current priority is being used by a job with assigned priority lower than yours
 - The cells here are as in Table “*priority-inheritance blocked by*” except for the jobs that do *not* request resources (e.g., J_5), which are exempt from this blocking

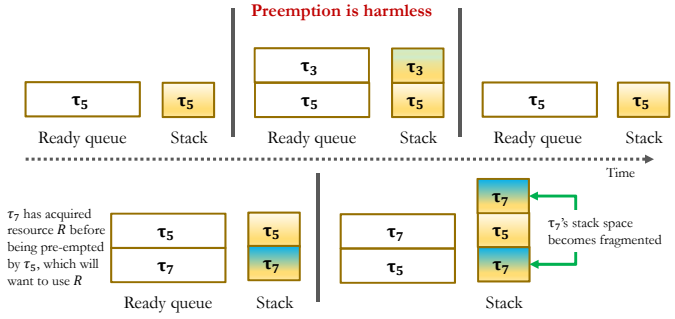
2019/2020 UniPD - T. Vardanega Real-Time Systems 175 of 538

Resource access control [option d]

- **Stack-based ceiling priority protocol**
 - SB-CPP uses the system ceiling same as BPCP, but it allows jobs to *share* stack space, saving tons of memory
 - Try and see why!
 - The other protocols seen so far don't
 - SB-CPP does it by ensuring that *no* request for resources will *ever* be denied to a running job 
 - This prevents jobs' stack space from fragmenting
 - Blocking causes stack fragmentation
 - Preemption does not!
 - One more reason to discourage self-suspension ...

2019/2020 UniPD - T. Vardanega Real-Time Systems 176 of 538

What blocking and preemption do to the stack




Preemption is harmless

τ₇ has acquired resource *R* before being pre-empted by τ₅, which will want to use *R*

[Inheritance-]Blocking is disastrous

2019/2020 UniPD - T. Vardanega Real-Time Systems 177 of 538

SB-CPP protocol rules [Baker, 1991]

- Computation of and updates to ceiling $\pi_s(t)$:
 - When all resources are free, $\pi_s(t) = \Omega$
 - $\pi_s(t)$ is updated any time t a resource is assigned or released
- Scheduling: on release at time t , job J stays *blocked* until its *assigned priority* $\pi_J(t) > \pi_s(t)$
 - Jobs that are not blocked are dispatched to execution by preemptive priority-driven scheduling
- Allocation: whenever a job issues a request for a resource, the request is granted 


2019/2020 UniPD - T. Vardanega Real-Time Systems 178 of 538

Critique of [option d]

- Under SB-CPP, a job J can only begin execution when the resources it may need are free
 - Otherwise $\pi_J(t) > \pi_s(t)$ cannot hold
- Under SB-CPP, a job J that may get preempted does *not* become blocked on resumption
 - The preempting job *cannot* contend resources with J
- SB-CPP prevents deadlock from occurring
- Under SB-CPP, $B_i(rc)$ for any job J_i is the same as BPCP's
- SB-CPP has lower algorithmic complexity in time and space than BPCP, as it needs *less* checks against $\pi_s(t)$

2019/2020 UniPD - T. Vardanega Real-Time Systems 179 of 538

Resource access control [option e]

- **Ceiling priority protocol** (base version)
 - CPP does *not* use the system ceiling $\pi_s(t)$
 - Resources continue to have a ceiling priority attribute
- **Scheduling**: jobs are scheduled with FPS with “FIFO *within priorities*” ruling
 - A job that does not hold any resource, runs with its *assigned priority*
 - A job that acquires a resource has its *current priority* set to the highest value among the ceiling priority of the resources that it holds
- **Allocation**: whenever a job issues a request for a resource, the request is granted 

2019/2020 UniPD - T. Vardanega

Real-Time Systems

180 of 538

Summary

- Issues arising from task contention of shared resources under preemptive priority-based scheduling
- Survey of resource access control protocols
- Critique of the surveyed protocols

2019/2020 UniPD - T. Vardanega

Real-Time Systems

181 of 538

Selected readings

- L. Sha, R. Rajkumar, J.P. Lehoczky (1990)
Priority inheritance protocols: an approach to real-time synchronization
DOI: 10.1109/12.57058
- T. Baker (1990)
A Stack-Based Resource Allocation Policy for Real-time Processes
DOI: 10.1109/REAL.1990.128747

2019/2020 UniPD - T. Vardanega

Real-Time Systems

182 of 538