

Simple locking and priority inversion /1

 To illustrate an initial example of priority inversion, consider the execution of the task set shown below, under *simple locking* (i.e., by use of binary semaphores)

Task	Priority	Execution sequence	Release time
a	1 (low)	eQQQQe	0
Ь	2	ee	2
с	3	eVVe	2
d	4 (high)	eeQVe	4

Legend: e: one unit of execution; Q (or V): one unit of use of resource R_q (or R_v)

Real-Time Systems

```
2015/16 UniPD / T. Vardanega
```

237 of 446









Bounding direct blocking under BPIP If the system has {r_{j=1,..,K}} critical sections that can lead to a task τ_i being blocked under BPIP then the maximum number of times that τ_i can be blocked is K The upper bound on the blocking time B_i(rc) for τ_i with K critical sections in the system is given by B_i(rc) = Σ_{j=1}^K use(r_j, i) × C_{max}(r_j) use(r_j, i) = 1 if r_j is used by at least one task τ_i: π_i < π_i and one task τ_i: π_h ≥ π_i | 0 otherwise C_{max}(r_j) being the duration of use of r_j by any such task τ_l. The worst case for task τ_i with BPIP is to block for the longest duration of contending use on access to <u>all</u> the resources it needs

Real-Time Systems

240 of 446

2015/16 UniPD / T. Vardanega























Arbitrary deadlines /1

 The RTA equation must be modified to cater for situations where D > T in which multiple jobs of the same task compete for execution

- The number *q* of additional releases to consider is bounded by the lowest value of *q*:*R_i(q)* ≤ *T_i*
 - $\omega_i(q)$ represents the level-i busy period, which extends as long as qT_i falls within it
- The worst-case response time is then $R_i = max_q R_i(q)$

```
2015/16 UniPD / T. Vardanega
```

Real-Time Systems

255 of 446















2015/16 UniPD / T. Vardanega







