

Real-Time Systems

Fall 2007

Notes for HW2

Notes by: Rohit Jain, David Ott, & Kevin Jeffay

Processor Demand

Consider a real-time system that has executed successfully up to some time a under some priority assignment scheme (*i.e.*, no deadline has been missed). The *processor demand* in the interval $[a, b]$, denoted $D[a, b]$, is the minimum amount of time that must be spent executing tasks in the interval $[a, b]$ to ensure that no job misses a deadline at or before time b . Processor demand is determined by the priority assignment scheme in use and the number of jobs released in and around the interval $[a, b]$. Thus for an interval $[a, b]$, $D[a, b]$ under an *RM* priority assignment may be different than $D[a, b]$ under an *EDF* priority assignment. For example, consider the pair of tasks shown executing under an *EDF* priority assignment in Figure 2. $D[0, 5] = 3$ because 3 units of work are required to satisfy jobs of T_1 and T_2 released in $[0, 5]$ with deadlines at or before time 5. However, $D[0, 3] = 0$ because there are no jobs with deadlines in the interval $[0, 3]$, whereas for an interval of identical length such as $[1, 4]$, $D[1, 4] = 1$. Moreover, if a different priority assignment scheme is used, such as a *latest deadline first* priority assignment, then $D[1, 4] = 2$.




Figure 2

In general, there is no simple expression for processor demand. However, at a minimum, $D[a, b]$ is at least as big as the amount of processor time required by jobs of tasks released in the interval $[a, b]$ with deadlines at or before time b , and no bigger than this amount plus the amount of processor time required by jobs of tasks released before time a with deadlines in the interval $[a+1, b]$ and by jobs of tasks released in $[a, b-1]$ with deadlines after time b . These bounds on $D[a, b]$ are illustrated in Figure 3. Although no general expression for processor demand exists, for certain priority assignment schemes and certain types of intervals, we will shortly develop useful analytic bounds on processor demand.




Figure 3: Lower (left) and upper (right) bounds on processor demand.

Processor demand gives us a concise way of expressing schedulability of a task set under a given priority assignment scheme. If a task set is schedulable under a particular priority assignment, then for all t and l , the processor demand in $[t, t+l]$ must be less than or equal to l .

Theorem 0: For any priority assignment scheme, a set of periodic tasks will be schedulable if and only if for all non-negative t and L ,

$$D[t, t+L] \leq L. \quad (1)$$

Proof: The proof follows immediately from the definition of processor demand. \square

Thus schedulability can be expressed as a function of processor demand. While (1) holds for any priority assignment, for specific priority assignments we can develop (numerous) refinements to (1) that will be more useful.

Theorem 0': For the *EDF* and *RM* priority assignment schemes, a set of periodic tasks will be schedulable if and only if for all $L \geq 0$, $D[0, L] \leq L$.

Proof: Left as an exercise.