Embedded Information Systems: Addendum (October 1, 2019)

This document introduces the Total Bandwidth Server method to schedule periodic and aperiodic tasks using the EDF algorithm.

Total Bandwidth Server (TBS) [Spuri & Butazzo]

This approach assigns a possible earlier deadline to each aperiodic request. This assignment should be done in such a way that the overall utilization of the aperiodic load never exceeds a specified maximum value μ_S. The name of the server comes from the fact that, each time an aperiodic request enters the system; the total bandwidth of the server is immediately assigned to it, whenever possible.

When the k-th aperiodic request arrives at time $t = r_k$, it receives a deadline:

$$d_k = max(r_k, d_{k-1}) + \frac{C_{ak}}{\mu_S},$$

where C_{ak} is the execution time of the request and μ_S is the server utilization factor (that is, its bandwidth). By definition $d_0 = 0$. In the deadline assignment rule the bandwidth allocated to previous aperiodic requests is considered through the deadline d_{k-1}.

Once a deadline is assigned, the request is inserted into the ready queue of the system and scheduled by EDF as any other periodic instance. As a consequence, the implementation overhead of this algorithm is practically negligible.

The Figure below illustrates this method. We have two periodic tasks: $T_1 = 6ms, C_1 = 3ms$, and $T_2 = 8ms, C_2 = 2ms$. Consequently $\mu_P = 0.75$ and thus $\mu_S = 0.25$.

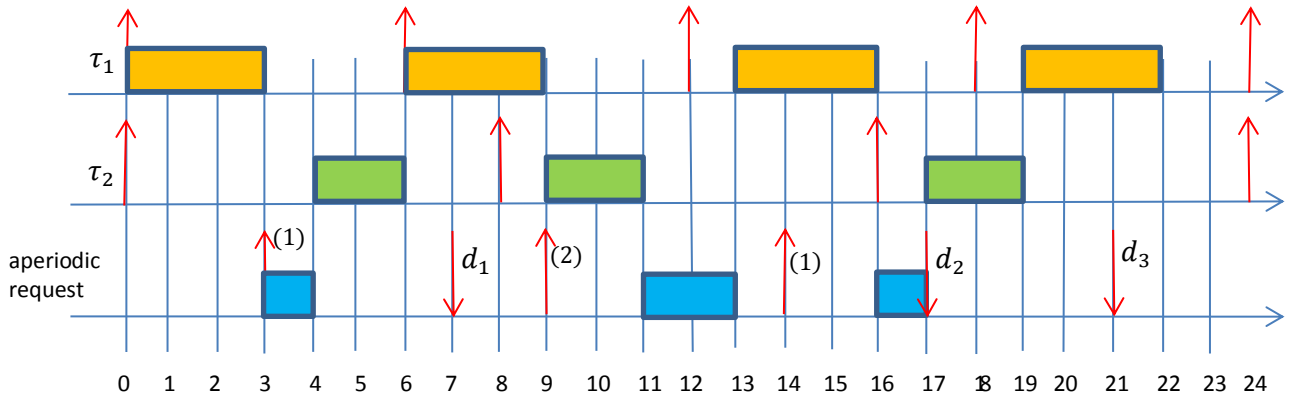

The first aperiodic request arrives at time $t = 3ms$, and is serviced with deadline $d_1 = r_1 + C_{a1}/\mu_S = (3 + 1/0.25)ms = 7ms$. Being this value the earliest deadline in the system, the aperiodic request is executed immediately. The second request, which arrives at time $t = 9ms$, receives a deadline $d_2 = r_2 + C_{a2}/\mu_S = (9 + 2/0.25)ms = 17ms$, however this is not serviced immediately, because at time $t = 9ms$ there is an active periodic task, τ_2 with a shorter deadline: $16\,ms$. Finally, the third aperiodic request arrives at time $t = 14ms$ and gets a deadline $d_3 = max(r_3, d_2) + C_{a3}/\mu_S = (17 + 1/0.25)ms = 21ms$. It does not receive immediate service, since at time $t = 14ms$ task τ_1 is active and has an earlier deadline: $18\,ms$.

It can be proved that if the processor utilization factor of the periodic tasks is μ_P, and that of the Total Bandwidth Server is μ_S, then this task set can be scheduled using EDF if and only if

$$\mu_P + \mu_S \leq 1.$$

Proof: If in every $[t_1, t_2]$ interval C_a is the total computation time of those aperiodic requests, which arrived at t_1 or later, and served with deadlines less than or equal to t_2, then

$$C_a \leq (t_2 - t_1)\mu_S,$$

because

$$C_a = \sum_{k=k_1}^{k_2} C_{ak} = \mu_S \sum_{k=k_1}^{k_2} \left(d_k - max(r_k, d_{k-1})\right) \leq \mu_S \left(d_{k_2} - max\left(r_{k_1}, d_{k_1-1}\right)\right) \leq \mu_S(t_2 - t_1).$$

After this, the proof of the schedulability test follows closely that of the periodic case.