The Earliest Deadline First scheduler

- In the *Earliest Deadline First* (EDF) scheduler, service class i is characterized by its service deadline d_i [seconds]
- At the end of each packet transmission, the scheduler fetches, among the packets waiting for transmission, the packet with the smallest residual time
- This is done by marking each packet on its arrival with the time stamp indicating its arrival time t_k
- At time t, the residual time of this packet, belonging to service class i, is calculated as $t_k + d_i - t$, representing the amount of time left before the packet service deadline expires
- Note that the packet residual time can be negative, indicating that the service deadline has already expired
- Smaller residual time means more urgent need of service: this is the reason why the packet with the smallest residual time is selected for service

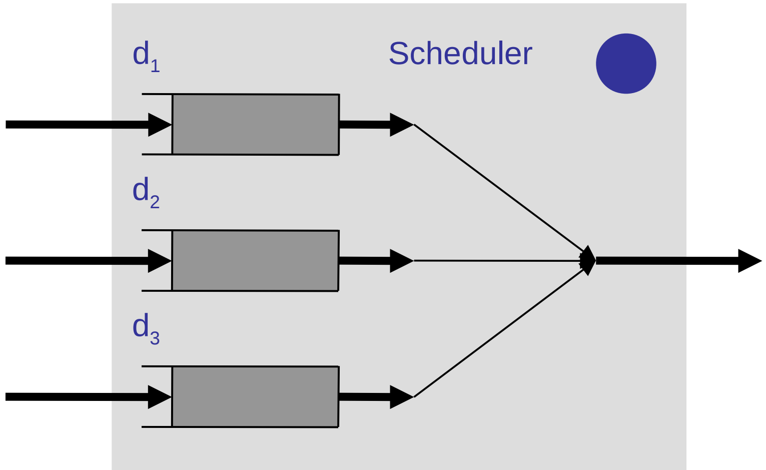

Back to admission control

- A simple example of admission control in a scenario with a single link served with a FIFO scheduler is shown in the figure
- *N* statistically identical flows share the link
- The statistical SLA is (*d*, *p*)
- Delay increases as the number of flows grows
- There exists a maximum number of admissible flows, N_{max}, such that the SLA is fulfilled
- If $N>N_{max}$, the SLA is violated

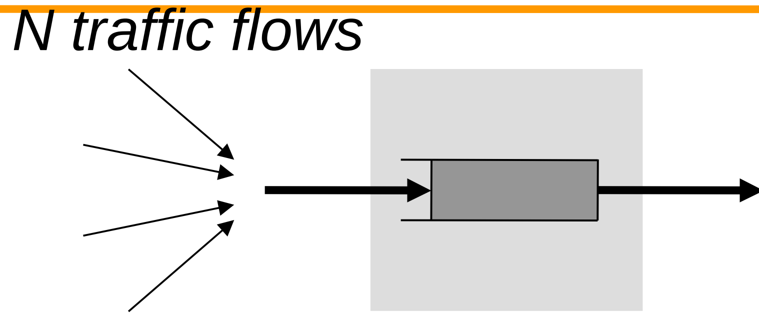

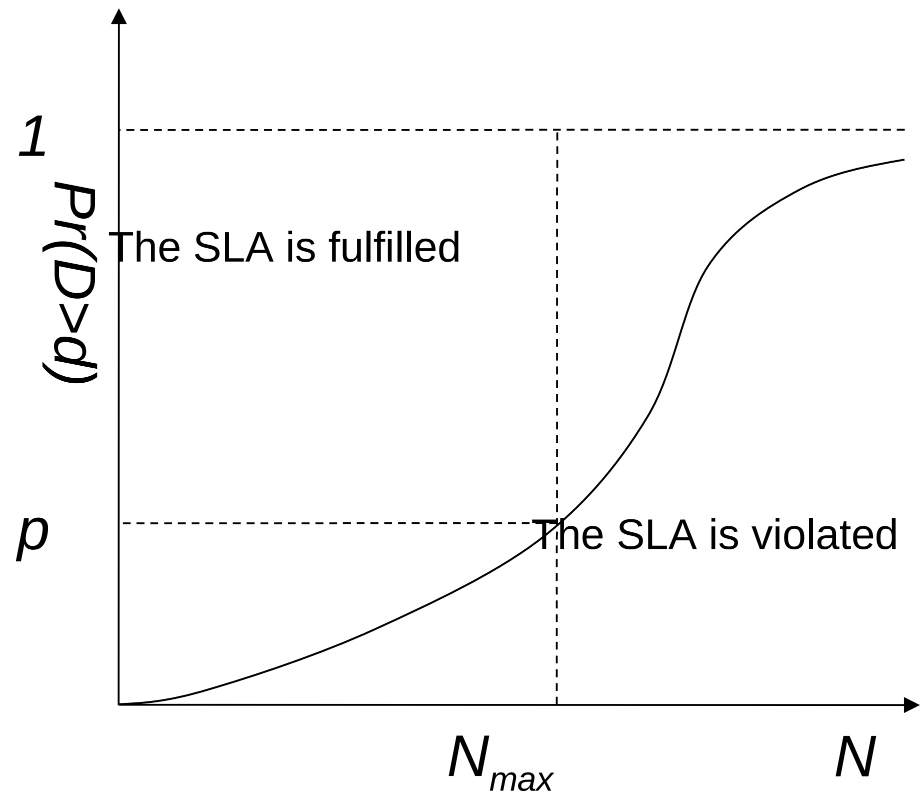

Back to admission control

- For a single-class scheduler, the value of N_{max} defines the ***admission region*** of the scheduler
- For a single-class scheduler, the admission region is uni-dimensional and defines the set of admissible values of *N*, in [0, N_{max}]
- In the figure, the admission region is highlited in greed and, in order to meet the SLA, the red region should be avoided, that is, requests that would lead the system into the red region should be blocked
- The amplitude of the admission region depends on the link's capacity and on the statistical properties of traffic

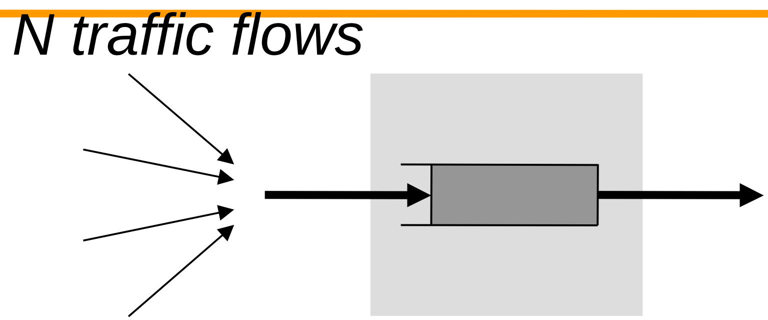

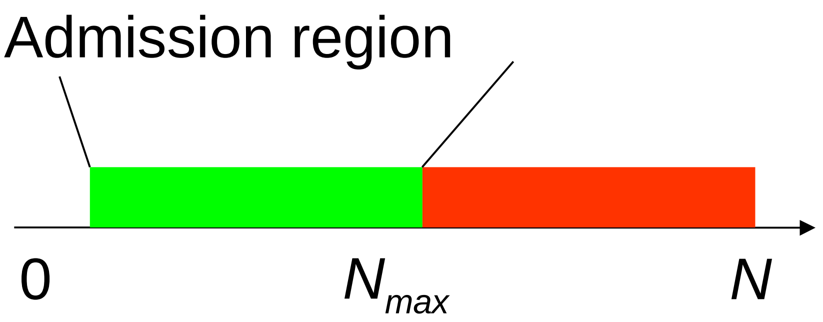

Back to admission control

- For a two-class scheduler, the admission region is clearly bi-dimensional
- The admission region is the set of all pairs (N_1, N_2) such that the SLA (d_1, p_1) of class-1 flows and the SLA (d_2, p_2) of class-2 flows is fulfilled
- The admission region (in green in the figure) is delimited by the admission line $N_2(N_1)$
- The admission line is the curve representing the maximum number of class-2 flows that can be admitted, given N_1 admitted flows in class 1

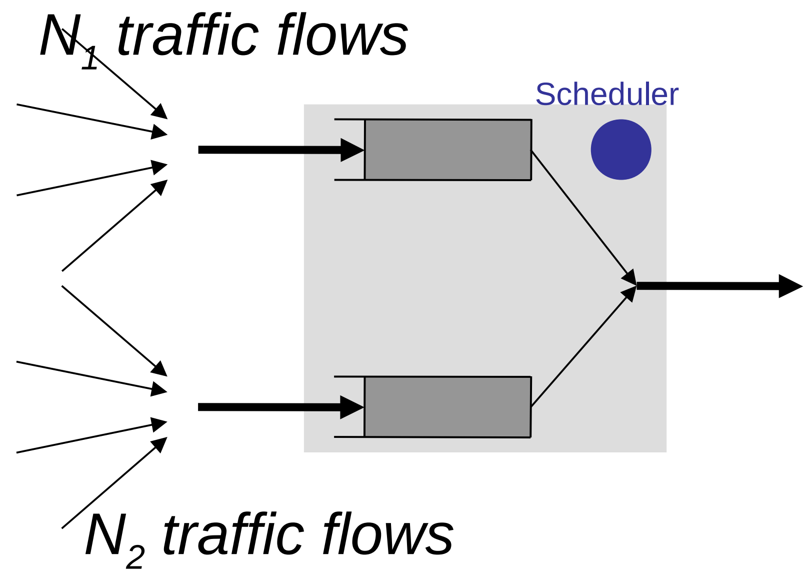

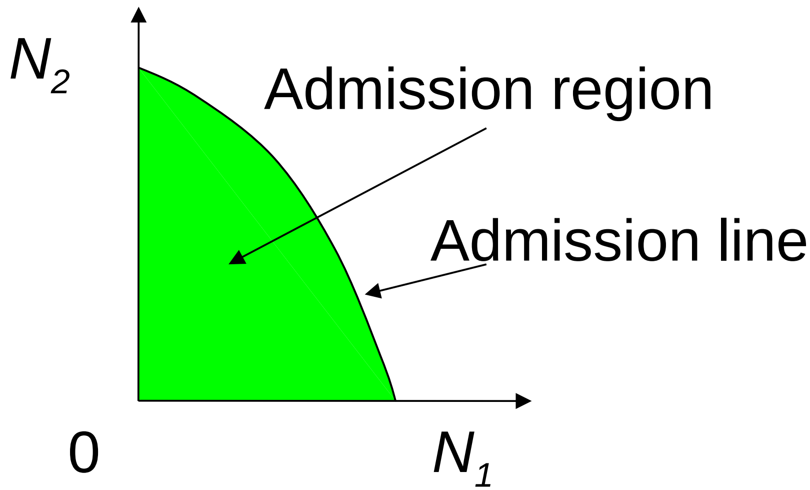

Back to admission control

- In a three-class scheduler, the admission region is three-dimensional
- There are N_1, N_2 and N_3 flows in class 1, 2 and 3, respectively
- The SLAs of the three classes are (d_1, p_1), (d_2, p_2) and (d_3, p_3)
- The admission region is the set of triplets (N_1, N_2, N_3) such that the three SLAs are simultaneously met
- The extension to more than three classes is intuitive, even if calculations to determined the shape of the admission region are increasingly complex