

CHALMERS
Handling on-line changes
Origins of on-line changes:
 Changing task characteristics: Tasks execute shorter than their worst-case execution time. Tasks increase/decrease the values of their static parameters as a result of, for example, mode changes.
 Dynamically arriving tasks: Aperiodic tasks (with characteristics known <i>a priori</i>) arrive New tasks (with characteristics not known <i>a priori</i>) enter the system at run-time.
 Changing hardware configuration: Transient/intermittent/permanent hardware faults Controlled hardware re-configuration (mode change)

CHALMERS
Handling on-line changes
Consequences of on-line changes:
 Overload situations: Changes in workload/architecture characteristics causes the accumulated processing demands from all tasks to exceed the capacities of the available processors. Question: How do we reject certain tasks in a way such that the inflicted damage is minimized?
 Scheduling anomalies: Changes in workload/architecture causes <u>non-intuitive</u> negative effects of system schedulability. Question: How do we avoid certain changes or use feasibility tests to guarantee that anomalies do not occur?

CHALMERS
Handling overload conditions
How do we handle a situation where the system becomes temporarily overloaded?
 Best-effort schemes: – No prediction for overload conditions.
 Guarantee schemes: Processor load is controlled by continuous acceptance tests.
 Robust schemes: Different policies for task acceptance and task rejection.
 Negotiation schemes: Modifies workload characteristics within agreed-upon bounds.















CHALMERS
Handling aperiodic tasks
Aperiodic task model:
 Spatial: The aperiodic task arrival is handled <u>centralized</u>; this is the case for multiprocessor servers with a common run-time system. The aperiodic task arrival is handled <u>distributed</u>; this is the case for distributed systems with separate run-time systems.
 Temporal: The aperiodic task is assumed to only arrive <u>once</u>; thus, it has <u>no period</u>. The actual arrival time of an aperiodic task is not known in advance (unless the system is clairvoyant). The actual parameters (e.g., WCET, relative deadline) of an aperiodic task may not be known in advance.



CHALMERS
Handling aperiodic tasks
Challenges in handling aperiodic tasks:
 Server-based approach: How do we reserve enough capacity to the server task without compromising schedulability of hard real-time tasks, while yet offering good service for future aperiodic task arrivals? Server-less approach:
 How do we design a schedulability test that accounts for arrived aperiodic tasks (remember: they do not have periods)? To what other processor do we off-load a rejected aperiodic task (in case of multiprocessor systems)?



CHALMERS
Static-priority servers
Background scheduling: Schedule aperiodic activities in the background; that is, when there are no periodic task instances to execute.
Advantage: – Very simple implementation
Disadvantage: – Response time can be too long



CHALMERS
Static-priority servers
 Polling Server (PS): (Lehoczky, Sha & Strosnider, 1987) Service aperiodic tasks using a dedicated task with a period Ts and a capacity Cs. If no aperiodic tasks need service in the beginning of PS:s period, PS suspends itself until beginning of next period. Unused server capacity is used by periodic tasks.
Advantage: – Much better average response time
 Disadvantage: If no aperiodic request occurs at beginning of server period, the entire server capacity for that period is lost.







CHALMERS
Static-priority servers
Feasibility test for RM + DS: A set of <i>n</i> periodic tasks and one aperiodic server are schedulable using RM if the processor utilization does not exceed:
$U_{RM+DS} = U_S + n \left(\left(\frac{U_S + 2}{2U_S + 1} \right)^{1/n} - 1 \right)$





CHALMERS
Static-priority servers
Non-existence of optimal servers: (Tia, Liu & Shankar, 1995)
For any set of periodic tasks ordered on a given static-priority scheme and aperiodic requests ordered according to a given aperiodic queuing discipline, there does not exist any valid algorithm that minimizes the response time of every soft aperiodic request.
For any set of periodic tasks ordered on a given static-priority scheme and aperiodic requests ordered according to a given aperiodic queuing discipline, <u>there does not exist any on-line algorithm</u> that minimizes the <u>average</u> response time of the soft aperiodic requests.

CHALMERS
Dynamic-priority servers
 Dynamic Priority Exchange Server: (Spuri & Buttazzo, 1994) Preserves its capacity by temporarily exchanging it for the execution time of a lower-priority (longer deadline) task. Dynamic Sporadic Server: (Spuri & Buttazzo, 1994) Replenishes its capacity only after it has been consumed by aperiodic task execution.
Total Bandwidth Server: (Spuri & Buttazzo, 1994) Assign a (possibly earlier) deadline to each aperiodic task and schedule it as a normal task. Deadlines are assigned such that the overall processor utilization of the aperiodic load never exceeds a specified maximum value Us.

CHALI	Slot-shifting server
S	Slot-Shifting Server: (Fohler, 1995) Schedules aperiodic tasks in the unused time slots in a schedule generated for time-driven dispatching
	Associated with each point in time is a <u>spare capacity</u> that indicates by how much the execution of the next periodic task can be shifted in time without missing any deadline.
	Whenever an aperiodic task arrives, task instances in the static workload may be shifted in time – by as much as the spare capacity indicates – in order to accommodate the new task