dynamic arrivals

Aperiodic tasks

 $au_{\rm A}$

EDA421/DIT171 - Parallel and Distributed Real-Time Systems, Chalmers/GU, 2011/2012



 μ_1

Handling on-line changes

Static (periodic) tasks

Handling on-line changes

Consequences of on-line changes:

Overload situations:

mode charges

Target

Updated November 8, 2011

CHALMERS

- 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 non-intuitive negative effects of system schedulability.
 - Question: How do we avoid certain changes or use feasibility tests to guarantee that anomalies do not occur?



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 a priori) arrive
 - New tasks (with characteristics not known a priori) enter the system at run-time.
- Changing hardware configuration:
 - Transient/intermittent/permanent hardware faults
 - Controlled hardware re-configuration (mode change)

CHALMERS

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.

Lecture #8





Lecture #8











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 no period.
 - 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.

EDA421/DIT171 - Parallel and Distributed Real-Time Systems, Chalmers/GU, 2011/2012 Updated November 8, 2011

Handling aperiodic tasks

Approaches for handling aperiodic tasks:

• Server-based approach:

CHALMERS

CHALMERS

- Reserve capacity to a "server task" that is dedicated to handling aperiodic tasks.
- All aperiodic tasks are accepted, but can only be handled in a best-effort fashion => no guarantee on schedulability
- Server-less approach:
 - A schedulability test is made on-line for each arriving aperiodic task => guaranteed schedulability for accepted task.
 - Rejected aperiodic tasks could either be dropped or forwarded to another processor (in case of multiprocessor systems)

Aperiodic servers

Handling (soft) aperiodic tasks on uniprocessors:

- Static-priority servers:
 - Handles aperiodic/sporadic tasks in a system where periodic tasks are scheduled based on a static-priority scheme (RM).
- Dynamic-priority servers:
 - Handles aperiodic/sporadic tasks in a system where periodic tasks are scheduled based on a dynamic-priority scheme (EDF).
- Slot-shifting server:
 - Handles aperiodic/sporadic tasks in a system where periodic tasks are scheduled based on a time-driven scheme.

Primary goal: to minimize the response times of aperiodic tasks in order to increase the likelihood of meeting their deadlines.

CHALMERS

Handling aperiodic tasks

Challenges in handling aperiodic tasks:

• Server-based approach:

Lecture #8

- 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: Advantage: If no aperiodic request occurs at beginning of server period, the entire server capacity for that period is lost.

Lecture #8

CHALMERS

Static-priority servers

Deferrable Server (DS): (Lehoczky, Sha & Strosnider, 1987) Service aperiodic tasks using a dedicated task with a period Ts and a capacity Cs. Server maintains its capacity until end of period so that requests can be serviced as capacity is not exhausted.

Advantage:

- Even better average response time because capacity is not lost







Lecture #8

CHALMERS

(Other) Static-priority servers

Priority Exchange Server: (Lehoczky, Sha & Strosnider, 1987)
Preserves its capacity by temporarily exchanging it for the execution time of a lower-priority periodic task.
Sporadic Server: (Sprunt, Sha & Lehoczky, 1989)
Replenishes its capacity only after it has been consumed by aperiodic task execution.
Slack Stealing: (Lehoczky & Ramos-Thuel, 1992)
Does not use a periodic server task. Instead, it creates a passive task which attempts to make time for servicing aperiodic tasks by "stealing" processing time from periodic tasks without causing their deadlines to be missed.



CHALMERS

Slot-shifting server

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.

CHALMERS 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.