





CHALMERS	
Scheduling constraint	S
Examples of scheduling constraints:	
 Non-preemptive scheduling: Once started, a task cannot be preempted by a Greedy scheduling: Once started, a task cannot be preempted by a No processor sharing: A processor sharing: A processor can only execute one task at a time No dynamic task parallelism: A task can only execute on one processor at a No task migration: A task can only execute on one given processor change processor during its execution 	nother task lower-priority task e time or, or cannot

CHALMERS	
Scheduling constraints	
Non-preemptive scheduling:	
 Advantages: Mutual exclusion is automatically guaranteed Existing methods for WCET analysis works well 	
 Disadvantages: Negative effect on schedulability Scheduling decision takes effect after a task has executed Once a task starts executing, all other tasks on the same processor will be blocked until execution is complete 	

CHALMERS
Scheduling constraints
Preemptive scheduling:
 Advantages: Schedulability is not negatively affected Scheduling decisions can take effect as soon as the system state changes (even in the middle of task execution) The capacities of task priorities can be used in full Disadvantages: Mutual exclusion has to be guaranteed by e.g. semaphores (or similar constructs) WCET analysis is more complicated since cache and pipeline contents will be affected by a task switch Program security may be compromised (through so-called <i>covert channels</i>) if full preemption is allowed
 Program security may be compromised (through so-called covert channels) if full preemption is allowed



CHALMERS
Scheduling constraints
Fair scheduling:
 Example: p-fair scheduling (Baruah et al. 1995) Although a task has started executing, lower-priority tasks receive a guaranteed time quantum per time unit for execution All tasks hence make some kind of progress per time unit
 Advantages: – Schedulability maximized when task switch cost is negligible
 Disadvantages: Scheduler is relatively complicated to implement Poor schedulability when task switch cost is non-negligible Fairness implies significantly more task switches than greediness
L



CHALMERS
Scheduling algorithm
How much an oracle is the scheduling algorithm?
 Myopic scheduler: Scheduling algorithm only knows about currently ready tasks. Scheduling decisions are only taken whenever a new task instance arrives or a running task instance terminates.
 Clairvoyant scheduler: Scheduling algorithm "knows the future"; that is, it knows in advance the arrival times of the tasks. On-line clairvoyant scheduling is difficult to realize in practice.
"Predictions are always hard to make. In particular about the future." (Yogi Berra)

CHALMERS
Static scheduling
General properties:
 Off-line schedule generation: Explicit start and finishing times for each task is derived Cyclic schedule with a meta period equal to the least common multiple (LCM) of the task periods
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$

CHALMERS	
Static scheduling	
General properties:	
 Automatic techniques for schedule generation <u>Simulate</u> a run-time system with dynamic scheduling and record the executions (start and finish times), or Search for a feasible schedule using an intelligent heuristic, such as branch-and-bound (A*) or simulated annealing 	
 Schedulability test obtained "for free" Generated schedule can easily be checked for feasibility 	
 Mutual exclusion and precedence is handled explicitly Heuristic algorithm can be constrained to never perform a task switch in a critical region, and to obey execution order requirements 	







Dynamic scheduling
 Advantages: High flexibility Schedule can easily adapt to changes in the system Effective for different types of tasks Sporadic tasks easily supported (via suitable priority assignment) Implementation is not affected by task characteristics
 Disadvantages: Less predictable execution Temporary variations (jitter) in periodicity can occur Complicated inter-task communication Task must synchronize to exchange data Difficult to adapt to TDMA networks (but simple for e.g. CAN)



CHALMERS
Dynamic scheduling
Deadline-monotonic scheduling (DM):
 Uses <u>static</u> priorities Priority is determined by task deadline Tasks with shorter (relative) deadlines are assigned higher priorities Note: RM is a special case of DM, with D_i = T_i
 Theoretically well-established (for the uniprocessor) Exact schedulability test is an NP-complete problem DM is optimal among all scheduling algorithms that uses static priorities under the assumption that Di ≤ Ti for all tasks (shown by J. YT. Leung & J. Whitehead in 1982)











CHALMERS	
Handling shared resources	
Avoiding priority inversion:	
 Non-preemptive critical sections: Creates unnecessary blocking Only recommended for short critical sections 	
 Access-control protocols for critical sections: Priority Inheritance Protocol (PIP) [static priority] Priority Ceiling Protocol (PCP) [static priority] Immediate Ceiling Priority Protocol (ICPP) [static priority] Stack Resource Policy (SRP) [static and dynamic priority] PIP/PCP for dynamic priorities (EDF) Distributed PCP 	









CHALMERS
Handling shared resources
Alternative approach:
Lock-free and wait-free object sharing
If several tasks attempt to access a lock-free (wait-free) object concurrently, and if some proper subset of these tasks stop taking steps, then one (each) of the remaining tasks completes its access in a finite number of its own steps.

CHALMERS	
Handling shared resources	
Lock-Free Object Sharing: (Anderson et al., 1996)	
 Basic idea: The lock-free object sharing scheme is implemented using "retry loops". Object accesses are implemented using compare-and-swap instructions typically found in modern RISC processors. 	
 Advantage: Resource accesses are non-blocking Deadlock-free Avoids priority inversion Requires no kernel-level support 	
 Disadvantage: – Potentially unbounded retry loops 	

