



Introduction to Embedded Systems

Sanjit A. Seshia

UC Berkeley
EECS 149/249A
Fall 2015

© 2008-2015: E. A. Lee, A. L. Sangiovanni-Vincentelli, S. A. Seshia. All rights reserved.

Chapter 12: Operating Systems, Microkernels, &
Scheduling

Source

This lecture draws heavily from:

Giorgio C. Buttazzo, *Hard Real-Time Computing Systems*, Springer, 2004.

Responsibilities of a Microkernel (a small, custom OS)

- Scheduling of threads or processes
 - Creation and termination of threads
 - Timing of thread activations
- Synchronization
 - Semaphores and locks
- Input and output
 - Interrupt handling

Scheduling, UC Berkeley: 3

A Few More Advanced Functions of an Operating System – Not discussed here...

- Memory management
 - Separate stacks
 - Segmentation
 - Allocation and deallocation
- File system
 - Persistent storage
- Networking
 - TCP/IP stack
- Security
 - User vs. kernel space
 - Identity management

Scheduling, UC Berkeley: 4

Outline of a Microkernel Scheduler

- Main Scheduler Thread (Task):
 - set up periodic timer interrupts;
 - create default thread data structures;
 - dispatch a thread (procedure call);
 - execute main thread (idle or power save, for example).
- Thread data structure:
 - copy of all state (machine registers)
 - address at which to resume executing the thread
 - status of the thread (e.g. blocked on mutex)
 - priority, WCET (worst case execution time), and other info to assist the scheduler

Scheduling, UC Berkeley: 5

Outline of a Microkernel Scheduler

- Timer interrupt service routine:
 - dispatch a thread.
- Dispatching a thread:
 - *disable interrupts;*
 - determine which thread should execute (scheduling);
 - if the same one, enable interrupts and return;
 - save state (registers) into current thread data structure;
 - save return address from the stack for current thread;
 - copy new thread state into machine registers;
 - replace program counter on the stack for the new thread;
 - *enable interrupts;*
 - return.

Scheduling, UC Berkeley: 6

The Focus in this Lecture:
How to decide which thread to schedule?

Considerations:

- Preemptive vs. non-preemptive scheduling
- Periodic vs. aperiodic tasks
- Fixed priority vs. dynamic priority
- Priority inversion anomalies
- Other scheduling anomalies

Scheduling, UC Berkeley: 7

When can a new thread be dispatched?

- *Under Non-Preemptive scheduling:*
 - When the current thread completes.
- *Under Preemptive scheduling:*
 - Upon a timer interrupt
 - Upon an I/O interrupt (possibly)
 - When a new thread is created, or one completes.
 - When the current thread blocks on or releases a mutex
 - When the current thread blocks on a semaphore
 - When a semaphore state is changed
 - When the current thread makes any OS call
 - file system access
 - network access
 - ...

Scheduling, UC Berkeley: 8

Preemptive Scheduling

Assumptions:

1. All threads have priorities
 - either statically assigned (constant for the duration of the thread)
 - or dynamically assigned (can vary).
2. Kernel keeps track of which threads are “enabled” (able to execute, e.g. not blocked waiting for a semaphore or a mutex or for a time to expire).

Preemptive scheduling:

- At any instant, the **enabled thread with the highest priority** is executing.
- Whenever any thread changes priority or enabled status, the kernel can dispatch a new thread.

Scheduling, UC Berkeley: 9

Rate Monotonic Scheduling

Assume n tasks invoked periodically with:

- periods T_1, \dots, T_n (impose real-time constraints)
- worst-case execution times (WCET) C_1, \dots, C_n
 - assumes no mutexes, semaphores, or blocking I/O
- no precedence constraints
- fixed priorities
- preemptive scheduling

Rate Monotonic Scheduling: Priorities ordered by period (smallest period has the highest priority)

Scheduling, UC Berkeley: 10

Rate Monotonic Scheduling

Assume n tasks invoked periodically with:

- periods T_1, \dots, T_n (impose real-time constraints)
- worst-case execution times (WCET) C_1, \dots, C_n
 - assumes no mutexes, semaphores, or blocking I/O
- no precedence constraints
- fixed priorities
- preemptive scheduling

Theorem: If any priority assignment yields a feasible schedule, then RMS (smallest period has the highest priority) also yields a feasible schedule.

RMS is optimal in the sense of feasibility.

Liu and Leland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," J. ACM, 20(1), 1973.


Scheduling, UC Berkeley: 11

Feasibility for RMS

Feasibility is defined for RMS to mean that every task executes to completion once within its designated period.

Scheduling, UC Berkeley: 12


Showing Optimality of RMS:
Consider two tasks with different periods



Is a non-preemptive schedule feasible?

Scheduling, UC Berkeley: 13


Showing Optimality of RMS:
Consider two tasks with different periods



Non-preemptive schedule is not feasible. Some instance of the Red Task (2) will not finish within its period if we do non-preemptive scheduling.

Scheduling, UC Berkeley: 14

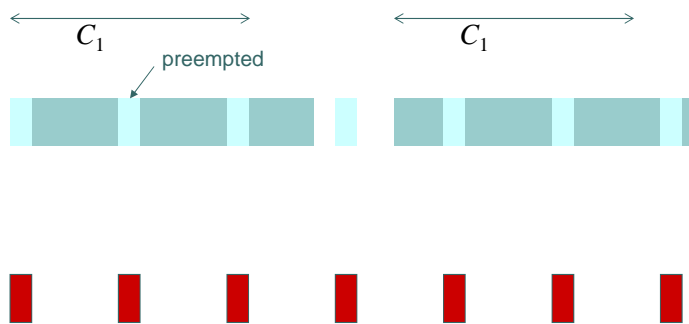
Showing Optimality of RMS:
Consider two tasks with different periods



What if we had a preemptive scheduling with higher priority for red task?

Scheduling, UC Berkeley: 15

Showing Optimality of RMS:
Consider two tasks with different periods



Preemptive schedule with the red task having higher priority is feasible. Note that preemption of the blue task extends its completion time.

Scheduling, UC Berkeley: 16

Any Implicit Assumptions in Our Reasoning?

Assume zero time needed to switch between tasks

Scheduling, UC Berkeley: 17

Showing Optimality of RMS for 2 Tasks: Key Proof Ideas


1. Worst case: when both tasks arrive together.
2. For this case, show if non-RMS is feasible, then RMS is feasible.

Scheduling, UC Berkeley: 18

Showing Optimality of RMS:
Alignment of tasks

Completion time of the lower priority task is worst (latest) when its *starting phase* matches that of higher priority tasks.

Thus, when checking schedule feasibility, it is sufficient to consider only the worst case: All tasks start their cycles at the same time.




Scheduling, UC Berkeley: 19

Showing Optimality of RMS:
(for two tasks)

It is sufficient to show that if a non-RMS schedule is feasible, then the RMS schedule is feasible.

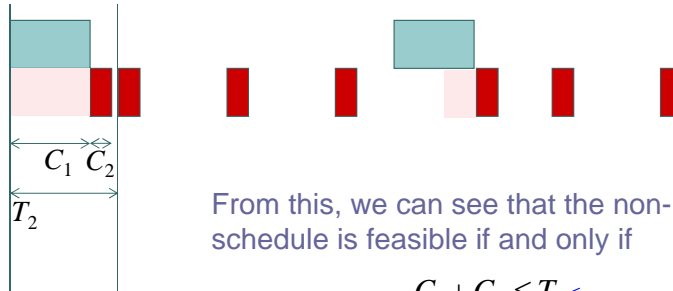
Consider two tasks as follows:



Scheduling, UC Berkeley: 20

Showing Optimality of RMS: (for two tasks)

The non-RMS, fixed priority schedule looks like this:



From this, we can see that the non-RMS schedule is feasible if and only if


$$C_1 + C_2 \leq T_2 \leftarrow \text{Actually: } \min(T_1, T_2)$$

We can then show that this condition implies that the RMS schedule is feasible.

Scheduling, UC Berkeley: 21

Showing Optimality of RMS: (for two tasks)

The RMS schedule looks like this: (task with smaller period moves earlier)



The condition for the non-RMS schedule feasibility:

$$C_1 + C_2 \leq T_2$$

is clearly sufficient (though not necessary) for feasibility of the RMS schedule. QED.

Scheduling, UC Berkeley: 22

Comments

- This proof can be extended to an arbitrary number of tasks (though it gets much more tedious).
- This proof gives optimality only w.r.t. feasibility. It says nothing about other optimality criteria.
- Practical implementation:
 - Timer interrupt at greatest common divisor of the periods.
 - Multiple timers

Scheduling, UC Berkeley: 23

Deadline Driven Scheduling:

1. Jackson's Algorithm: EDD (1955)

Given n independent *one-time* tasks with deadlines d_1, \dots, d_n , schedule them to minimize the maximum *lateness*, defined as

$$L_{\max} = \max_{1 \leq i \leq n} \{f_i - d_i\}$$

where f_i is the finishing time of task i . Note that this is negative iff all deadlines are met.

Earliest Due Date (EDD) algorithm: Execute them in order of non-decreasing deadlines.

Note that this does not require preemption.

Scheduling, UC Berkeley: 24

Theorem: EDD is Optimal in the Sense of Minimizing Maximum Lateness


To prove, use an interchange argument. Given a schedule S that is not EDD, there must be tasks a and b where a immediately precedes b in the schedule but $d_a > d_b$. Why?

We can prove that this schedule can be improved by interchanging a and b . Thus, no non-EDD schedule achieves smaller max lateness than EDD, so the EDD schedule must be optimal.


Scheduling, UC Berkeley: 25

Consider a non-EDD Schedule S

There must be tasks a and b where a immediately precedes b in the schedule but $d_a > d_b$



$$L_{\max} = \max\{f_a - d_a, f_b - d_b\} = f_b - d_b$$



$$L'_{\max} = \max\{f'_a - d_a, f'_b - d_b\}$$

Case 1: $f'_a - d_a > f'_b - d_b$.

Then: $L'_{\max} \leq f'_a - d_a = f_b - d_a \leq L_{\max}$
(because: $d_a > d_b$).

Case 2: $f'_a - d_a \leq f'_b - d_b$.

Then: $L'_{\max} \leq f'_b - d_b \leq L_{\max}$
(because: $f'_b < f_b$).

Theorem: $L'_{\max} \leq L_{\max}$.

Hence, S' is no worse than S .

Deadline Driven Scheduling:

1. Horn's algorithm: EDF (1974)

Extend EDD by allowing tasks to “arrive” (become ready) at any time.

Earliest deadline first (EDF): Given a set of n independent tasks with *arbitrary arrival times*, any algorithm that at any instant executes the task with the earliest absolute deadline among all arrived tasks is optimal w.r.t. minimizing the maximum lateness.

Proof uses a similar interchange argument.

Scheduling, UC Berkeley: 27

Using EDF for Periodic Tasks

- The EDF algorithm can be applied to periodic tasks as well as aperiodic tasks.
 - Simplest use: Deadline is the end of the period.
 - Alternative use: Separately specify deadline (relative to the period start time) and period.

Scheduling, UC Berkeley: 28

RMS vs. EDF? Which one is better?

What are the pros and cons of each?


Scheduling, UC Berkeley: 29

Comparison of EDF and RMS

- Favoring RMS
 - Scheduling decisions are simpler (fixed priorities vs. the dynamic priorities required by EDF. EDF scheduler must maintain a list of ready tasks that is sorted by priority.)
- Favoring EDF
 - Since EDF is optimal w.r.t. maximum lateness, it is also optimal w.r.t. feasibility. RMS is only optimal w.r.t. feasibility. For infeasible schedules, RMS completely blocks lower priority tasks, resulting in unbounded maximum lateness.
 - EDF can achieve full utilization where RMS fails to do that
 - EDF results in fewer preemptions in practice, and hence less overhead for context switching.
 - Deadlines can be different from the period.

Scheduling, UC Berkeley: 30

Precedence Constraints





DAG, showing that task 1 must complete before tasks 2 and 3 can be started, etc.

A directed acyclic graph (DAG) shows precedences, which indicate which tasks must complete before other tasks start.

Scheduling, UC Berkeley: 31


Example: EDF Schedule



Is this feasible? Is it optimal?

Scheduling, UC Berkeley: 32

EDF is not optimal under precedence constraints




The EDF schedule chooses task 3 at time 1 because it has an earlier deadline. This choice results in task 4 missing its deadline.

Is there a feasible schedule?

Scheduling, UC Berkeley: 33



LDF is optimal under precedence constraints



The LDF schedule shown at the bottom respects all precedences and meets all deadlines.

Scheduling, UC Berkeley: 34



Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 35



Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 36



Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 37



Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 38



Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 39



Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 40

Latest Deadline First (LDF) (Lawler, 1973)



The LDF scheduling strategy builds a schedule backwards. Given a DAG, choose the leaf node with the latest deadline to be scheduled last, and work backwards.

Scheduling, UC Berkeley: 41

Latest Deadline First (LDF) (Lawler, 1973)

LDF is optimal in the sense that it minimizes the maximum lateness.



It does not require preemption. (We'll see that EDF does.)

However, it requires that all tasks be available and their precedences known before any task is executed.

Scheduling, UC Berkeley: 42

EDF with Precedences

With a preemptive scheduler, EDF can be modified to account for precedences and to allow tasks to arrive at arbitrary times. Simply adjust the deadlines and arrival times according to the precedences.



Recall that for the tasks at the left, EDF yields the schedule above, where task 4 misses its deadline.



Scheduling, UC Berkeley: 43

EDF with Precedences Modifying release times

Given n tasks with precedences and release times r_i , if task i immediately precedes task j , then modify the release times as follows:

$$r'_j = \max(r_j, r_i + C_i)$$

assume:
 $r_i = 0$


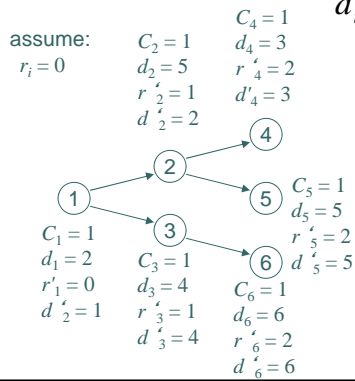


Scheduling, UC Berkeley: 44

EDF with Precedences Modifying deadlines

Given n tasks with precedences and deadlines d_i , if task i immediately precedes task j , then modify the deadlines as follows:

$$d'_i = \min(d_i, d'_j - C_j)$$



Using the revised release times and deadlines, the above EDF schedule is optimal and meets all deadlines.

Scheduling, UC Berkeley: 45

Optimality

EDF with precedences is optimal in the sense of minimizing the maximum lateness.

Scheduling, UC Berkeley: 46