

# Real-time Embedded Systems

**Colin Perkins**

<http://csperkins.org/teaching/2003-2004/rtes4/lecture18.pdf>

**Reminder: Programming assignment due at 5pm tomorrow**

UNIVERSITY  
*of*  
GLASGOW



# Lecture Outline

- Considerations of real-time on embedded systems
- Open system architecture
  - Discussion of concepts
  - Advantages and disadvantages
  - Implementation using a two level scheduler
- Case study of an open system: RTLinux

Reading for this lecture: Chapters 7.9 and 12.5–12.7

# Real-time Embedded Systems

What is an embedded system?

“An embedded system is a special-purpose computer system built into a larger device.”

“An embedded system is some combination of computer hardware and software, either fixed in capability or programmable, that is specifically designed for a particular kind of application device. Industrial machines, automobiles, medical equipment, cameras, household appliances, airplanes, vending machines, and toys (as well as the more obvious cellular phone and PDA) are among the myriad possible hosts of an embedded system.”

⇒ Special purpose, (often) limited resources, not generally upgradeable

# Real-time Embedded Systems

- Embedded systems are **closed**:
  - Run a fixed suite of applications, known *a-priori*
  - Tasks are scheduled according to some well-known algorithm
  - Generally static and require predictability
  - Prove, or exhaustively demonstrate, correctness
  - Limited resources, tailored to the task at hand
  - Dedicated operating systems, scheduler support, etc.
- Embedded systems form part of the wider world:
  - Interact with the world through sensors and actuators
  - Often part of a wider system, comprising other embedded and general purpose systems
- Separation of concerns:
  - Embedded controllers engineered separately to other systems, including other embedded systems

⇒ Strongly engineered, according to the design principles taught in this module

# An Open Architecture for Real-Time Systems

- The advantage of a traditional embedded system is that resources are dedicated and predictability is guaranteed
- The disadvantage is that resources are dedicated, and typically underused
- This ensures predictability, but is wasteful since many applications have both general purpose and embedded components as part of the complete system
- Desirable to have a single device that can run both general purpose and embedded applications simultaneously
  - An **open system architecture** that can support many different classes of application, removing the distinction between embedded and general purpose systems
  - Not always suitable, but can give large savings for some classes of application

# Objectives of an Open System Architecture

- Independent design choice
  - The developer of an application can use a scheduling discipline best suited to that application to control thread execution and resource access, independent of other applications on system
- Independent validation
  - If the system validates assuming it runs alone on a processor with normalised speed  $S$ , it will run on a virtual share of a real processor with equivalent performance
- Admission and timing guarantee
  - New applications are subjected to an admission test before they are scheduled. If the admission test is passed, and the application accepted as a real-time task, the open system will guarantee its schedulability, regardless of other applications in the system

⇒ Independently developed and validated hard real-time applications can share the system with other real time and non-real time applications

# Implementing an Open System Architecture

- Should be clear that the open system architecture can only be implemented on a strictly partitioned virtual machine
  - Partition processor time
  - Control access to global resources
- Each application to run submits its requirements (e.g. task characteristics, type of scheduler needed, etc.) to a virtual machine monitor that performs an acceptance test
- The monitor partitions the physical resources into distinct virtual machines and runs schedulers for each application
- The partitioning can be implemented using a **two-level scheduler**

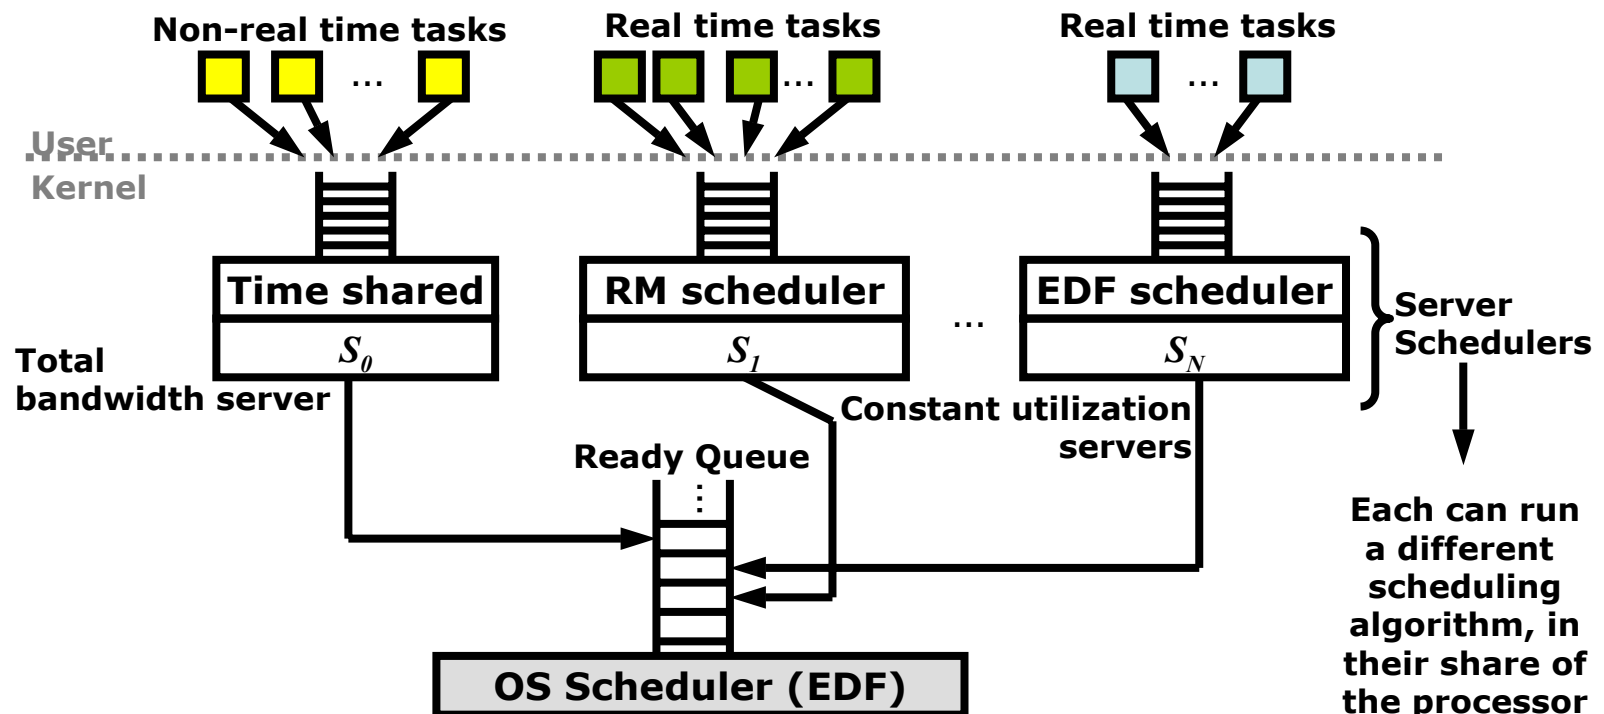
# Two-Level Scheduler

- Recall that a **constant utilization server** can be defined to use a fraction  $\tilde{u}_i$  of the processor
- A **total bandwidth server** will use at least the fraction  $\tilde{u}_i$  and can also claim idle time
- Both run under an EDF scheduling algorithm and are defined by certain consumption and replenishment rules See lecture 9
- Consider a system comprising:
  - A set of constant utilization servers  $S_i$  for  $i=1, 2, \dots, n$  each using fraction  $\tilde{u}_i$  of the processor
  - A total bandwidth server,  $S_0$ , using a fraction  $\tilde{u}_0$
  - An EDF scheduler, running these servers
- If the total size of the servers is less than some threshold, the tasks will fairly share a processor with  $S_0$  using the idle time
  - The threshold, the maximum schedulable utilization, depends on the properties of the server tasks




# Operation of a Two-Level Scheduler

- The servers,  $S_i$ , are scheduled according to an EDF algorithm by the **OS scheduler**
- Each server then runs an internal **server scheduler** which schedules jobs within the server, to subdivide the fraction of time allocated to that server



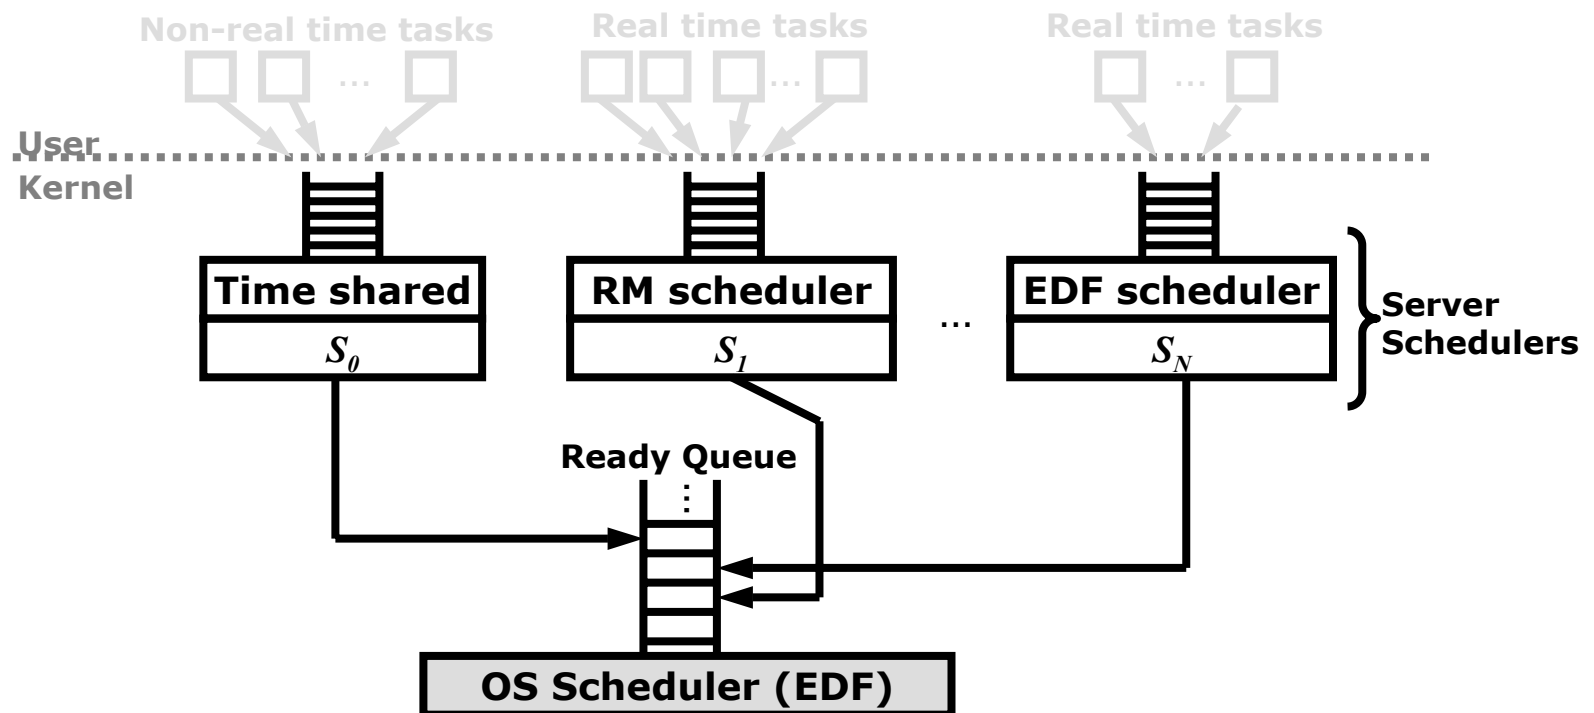
# Operation of a Two-Level Scheduler

- The OS scheduler maintains an EDF ready queue, used to select which of the servers to execute
  - Each server is eligible to run if it has work to do, and budget remaining, in the usual manner
  - The OS scheduler schedules the server with earliest deadline among the ready servers



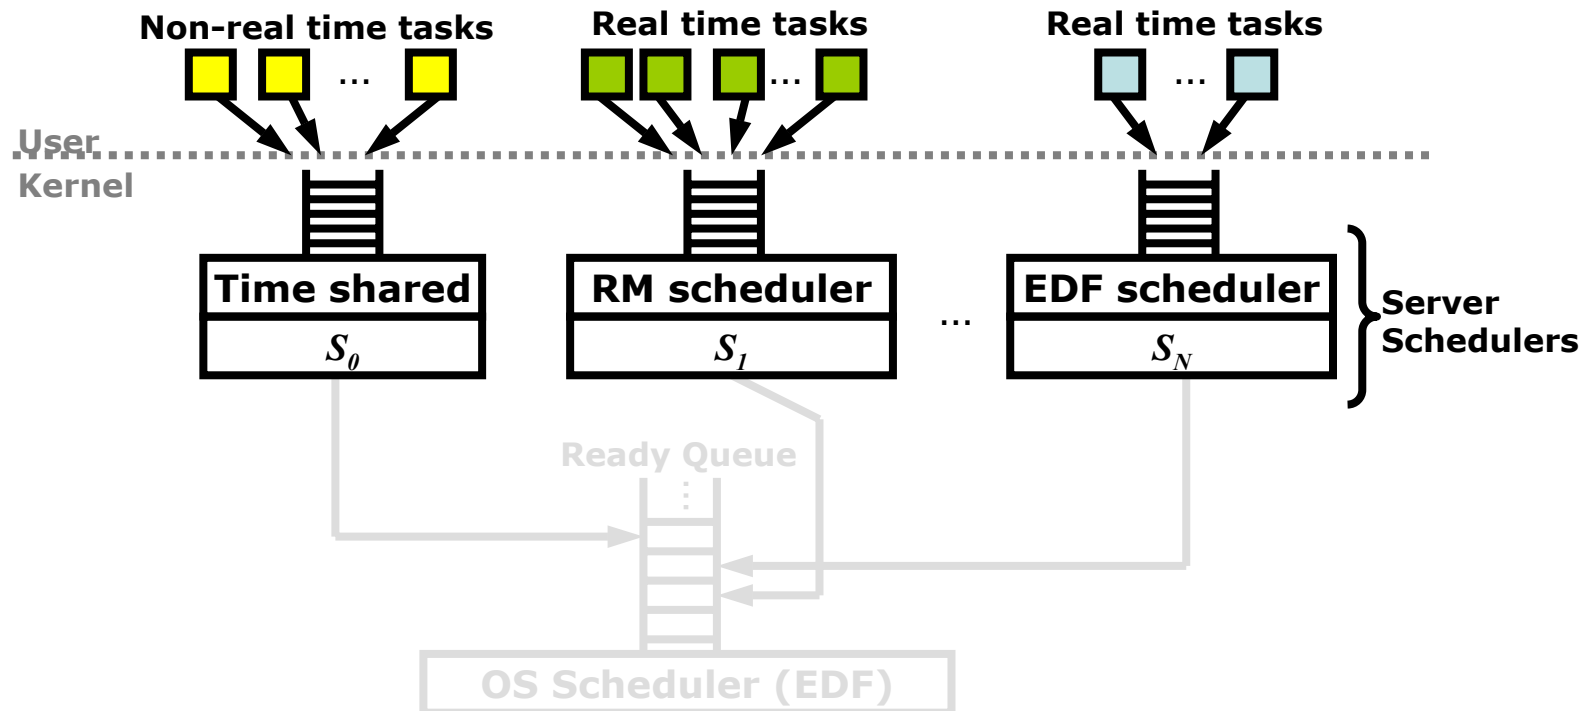
# Operation of a Two-Level Scheduler

- Both levels of scheduler are implemented in the kernel, to avoid doubling the context switch overhead
  - Applications see only a virtual machine, comprising their server and it's scheduler
  - The underlying OS scheduler is invisible to applications



# Operation of a Two-Level Scheduler

- Each server maintains a list of the threads comprising the applications running on that virtual machine
  - When executed by the OS scheduler, the server scheduler picks one of these threads to execute, according to its local policy
  - Each server scheduler can have different policy, and schedule its threads according to a different algorithm



# Types of Application

- Each application running on the system is allocated a server
- The jobs comprising the application can be:
  - Non-real time
  - Real-time and predictable
  - Real-time and unpredictable
- To perform EDF scheduling, the OS scheduler needs to know when the next scheduling event (=deadline) will occur on each server
  - Non-real time applications have no deadlines
  - Predictable real-time applications are those where the next event time is known
  - For unpredictable applications the next event time is not known and must be estimated
- Based on the event times, the server execution budgets are replenished and deadlines set

# Non-real Time Applications

- The server scheduler for  $S_0$  schedules all non-real time tasks according to a time sharing algorithm with time slice  $x$
- The server for  $S_0$  is a total bandwidth server
- Replenishment rules for the server:
  - When starting:
    - the server budget is set to  $x$
    - the deadline is set to  $x/\tilde{u}_0$  ( $\tilde{u}_0 < 1$  so deadline scaled out)
  - When budget exhausted at time  $t$ , if there are ready jobs:
    - the server budget is reset to  $x$
    - the deadline is set to either  $t+x/\tilde{u}_0$  or current deadline plus  $x/\tilde{u}_0$ , whichever is greater
  - When a busy interval ends:
    - the server budget is reset to  $x$
    - the deadline is set to the current time plus  $x/\tilde{u}_0$
- These rules ensure non-real time applications are scheduled as if on a slower processor with speed  $\tilde{u}_0$

# Real-Time Applications

- To schedule real-time applications, the OS scheduler needs to know the time when each event that will trigger a context switch will occur
- An application scheduled according to a pre-emptive priority algorithm is **unpredictable** if it contains aperiodic or sporadic tasks, or periodic tasks with significant release time jitter
  - Occurrence of scheduling events only known at run-time
- Other applications are **predictable**, since they contain only periodic tasks with fixed release times and known resource access patterns
  - The scheduler can compute event times for predictable applications before the system begins execution

# Predictable Applications

- Predictable applications contain periodic tasks with fixed release times and known resource access patterns
- These include:
  - Cyclic executives
  - Priority scheduled but not preemptable
  - Priority scheduled and preemptable, with known event times
- Each can run on a constant utilization server, and meet deadlines, but each type of predictable task has slightly different constraints
  - Required capacity of the task
  - Replenishment rules for the server



# Required Capacity

- A real-time application has to meet timing constraints
- To do this, it has been verified to need a certain amount of execution time,  $e$
- The job is to run on a slower processor, speed  $u < 1$ 
  - $u$  denotes the fraction of the original processor speed
- Can multiply the execution time  $e$  of all jobs by  $1/u$  to check if the system is still schedulable on the slow processor
- The minimum fraction of speed at which the application is schedulable is its **required capacity**, and is a critical parameter when scheduling jobs on the open system

# Cyclic Executives

- A cyclic executive is characterised by a cyclic frame of size  $f$  and the workload appears to the server as a single thread
- Scheduled on a constant utilization server of size  $\tilde{u}_i$  equal to required capacity
  - Ready for execution at the start of each cyclic frame
  - Execution time of  $f \cdot \tilde{u}_i$  each cycle
  - Budget replenished each cycle, deadline set to beginning of next cycle
- Loops using a fixed fraction of the processor time
- The server executes the application according to its pre-computed cycle

# Priority Scheduled, Non-preemptable Applications

- A priority scheduled, but not pre-emptive, application is scheduled by a constant utilization server
  - Server size  $\tilde{u}_i$  equal to the required capacity
  - Usual consumption and replenishment rules
- The server scheduler orders jobs in the application according to the scheduling algorithm requested by the application
- Since jobs are non-preemptable, the server must not be pre-empted while a job is running
  - Otherwise a job could be pre-empted by another server running on the OS scheduler
- This limits schedulable utilization of the **complete** system (not just this server):
  - Let  $B$  denote maximum execution time of all jobs
  - Let  $D_{min}$  denote the minimum relative deadline of all jobs
  - All servers are schedulable provided  $\Sigma \tilde{u}_i < 1 - B/D_{min}$
  - Implications on the acceptance test for the open system, due to presence of non-preemptable applications...

# Priority Scheduled, Preemptable Applications

- Preemptable priority scheduled applications are executed on a server that is similar to a constant utilization server
  - Slightly different replenishment rules
- Why different rules?
- Consider two jobs  $J_1$  and  $J_2$  running on a slow processor with speed 0.25
  - Each job has execution time 0.25
  - Job  $J_1$  is released at 0.5 and must complete by 1.5
  - Job  $J_2$  is released at 0.0 and must complete by 2.0
- Execution:
  - Job  $J_2$  starts at time 0.0, by time 0.5 has executed for 0.125
  - Job  $J_1$  starts at time 0.5, pre-empts  $J_2$ , and executes to completion by 1.5
  - Job  $J_2$  resumes at 1.5, executes to completion at 2.0
  - All deadlines are met

# Priority Scheduled, Preemptable Applications

- Consider the same jobs, on a constant utilization server of size  $\frac{1}{4}$  on a normal speed processor
- A possible scenario is:
  - Job  $J_2$  starts at time 0.0 with server budget 0.25 and deadline 1
  - Job  $J_2$  uses the budget, and completes before time 0.5
  - Job  $J_1$  is released at time 0.5, but the server budget is gone
  - At time 1.0 the server budget is replenished to 0.25 and its deadline set to 2.0; the server is eligible to run, and will execute job  $J_1$  when it runs
  - Because the server deadline is 2.0, the server may not execute until after  $J_1$  has missed its deadline at 1.5
- Problem: because running on a faster processor  $J_2$  was allowed to consume more execution time than it would running on the slower processor, preventing  $J_1$  from running
  - Used 0.25 compared to 0.125 on slow processor
  - Can be considered a form of priority inversion

# Priority Scheduled, Preemptable Applications

- When running preemptable priority scheduled applications, the server uses slightly modified replenishment rules to avoid this problem
  - Let  $t$  be latest of the current deadline, or current time
  - Let  $t'$  be the release time of the next job in the task
  - Budget =  $\min(e_i, (t' - t) \tilde{u}_i)$
  - Deadline =  $\min(t + e_i / \tilde{u}_i, t')$
- These rules prevent jobs consuming more time than they would on a slower processor, if they would be limited by pre-emption, by explicitly taking into account pre-emption time
- With this slight modification, preemptable priority scheduled applications can be supported

# Unpredictable Applications

- If jobs are unpredictable, the server cannot derive accurate replenishment rules
  - Run the server under the constant utilization rules, giving  $q\tilde{u}_i$  units of budget every  $q$  time units
  - The scheduling quantum,  $q$ , is a key parameter
- The server can over-budget the application up to the size of the scheduling quantum
  - Can result in priority inversion, as before...
- If a bound,  $t'$ , on time to the next event can be given, the server can be scheduled with
  - Budget =  $(t' + q - t).s_i$
  - Deadline =  $t' + q$
- Bounds the size of the scheduling quantum, and hence the duration of any priority inversion

# Summary of Scheduling Behaviour

- Given certain constraints, can schedule non-real time and predictable real-time applications with correct behaviour
- Unpredictable real-time applications may see occasional priority inversion of their jobs
- Applications are isolated from each other, provided an acceptance test is enforced to ensure constraints are met



# Scheduling Overhead

- Might think that the two-level scheduler is very inefficient
- Not so if all applications are predictable:
  - Same number of context switches
  - More work to determine what to run, but insignificant compared to context switch overhead
- In unpredictable applications running, overhead depends on the scheduling quantum
  - Small quantum gives better real-time performance, at the expense of more overheads
  - 30% overhead not unusual, but may still be better than using dedicated hardware
- If unpredictable jobs are rare, the two level scheduler works well and allows real-time and non-real-time jobs to share a processor

# Admission Control

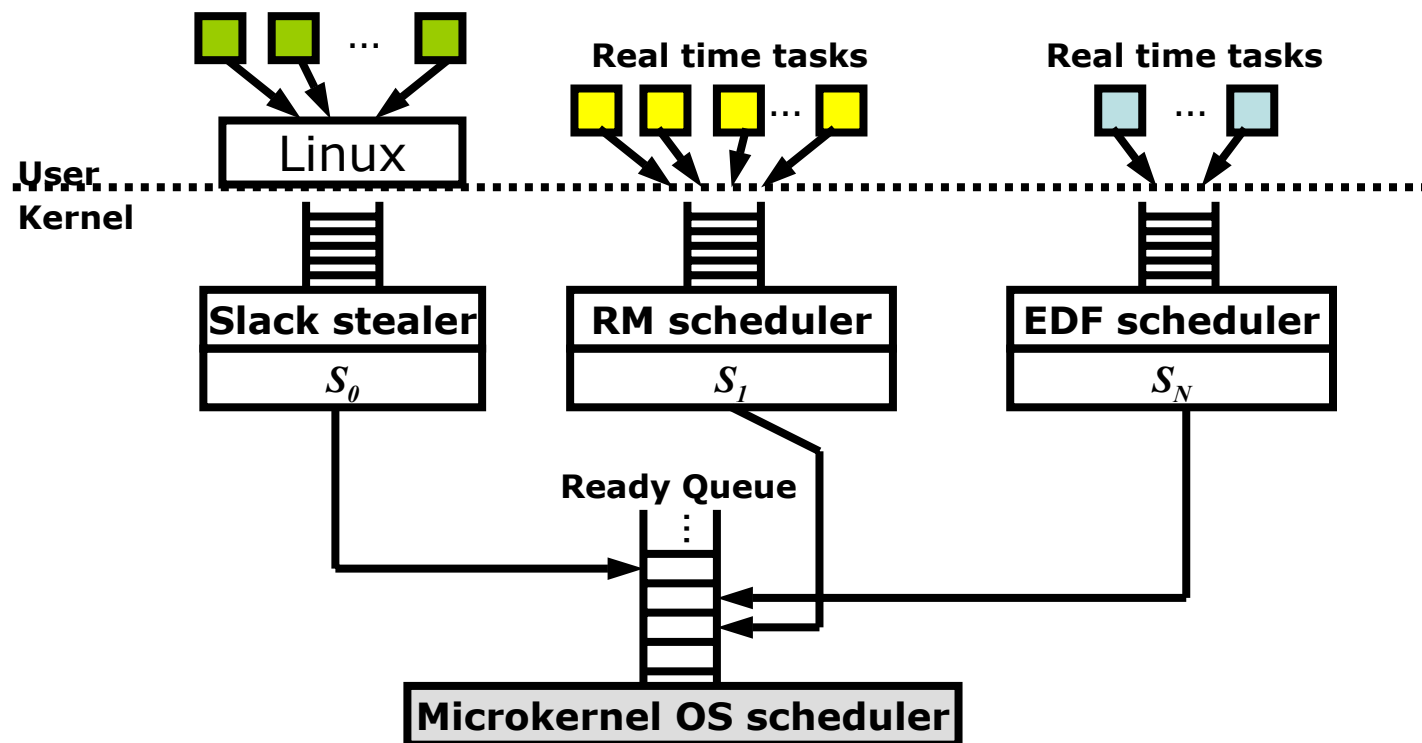
- All jobs start execution in non-real time mode, so they don't disrupt already running real-time jobs
- A job may switch to real-time scheduling on its own server, subject to an acceptance test
- Jobs must provide:
  - Required capacity  $\tilde{u}_i$  and scheduling algorithm
  - Maximum execution time  $B$  of all non-preemptable sections
  - Existence of aperiodic/sporadic tasks, if any
  - Shortest relative deadline  $D_{min}$
- Job is accepted if  $\Sigma \tilde{u}_i < 1 - \max(B/D_{min})$

# Properties of the Open System Architecture

- Conceptually clean way to share resources between tasks with different requirements
- Disadvantage:
  - Several applications share a hardware resource, so failure of the hardware or OS scheduler can take down an entire set of applications
  - Trade-off cost saving for potential reduction in reliability
- Not widely implemented:
  - Prototype demonstrated within Windows NT
  - Similar, but less powerful, systems are used commercially
    - Symbian mobile phones running a real-time microkernel to handle voice processing, with Symbian OS running as a background task to support the UI and user applications
    - RTLinux

# Case Study: RTLinux

- A simple example of a two-level scheduler
  - The OS scheduler is a microkernel real-time operating system
  - Real-time tasks run directly on the microkernel
    - RM and EDF schedulers provided
  - Linux runs as the idle task



# Case Study: RTLinux

- A modified Linux kernel runs above the microkernel
  - All hardware access is arbitrated by the microkernel
  - Interrupts emulated in software on the microkernel
  - Linux can **always** be pre-empted if a real-time task needs to run
- Communication between real-time and non-real-time tasks done by FIFO buffers, locked into memory
  - Appear as normal devices (`/dev/rft1`) under Linux
  - Non-blocking and atomic access from the real-time kernel
- Conceptually, RTLinux maps closely onto the open system architecture
  - Differs in the details

# Summary

By now, you should know...

- Concepts of real-time on embedded systems
- The idea of an open system architecture, to support a range of application types on a single system
- Strategies for implementing the open system architecture, using a two-level scheduler
- Overview of RTLinux, as a simple system using a two-level scheduler

Research seminar at 3pm tomorrow:

“Xen and the art of virtualisation”

Rolf Neugebauer, Intel Research