

Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

Mac Mollison

Björn Brandenburg

James H. Anderson

Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

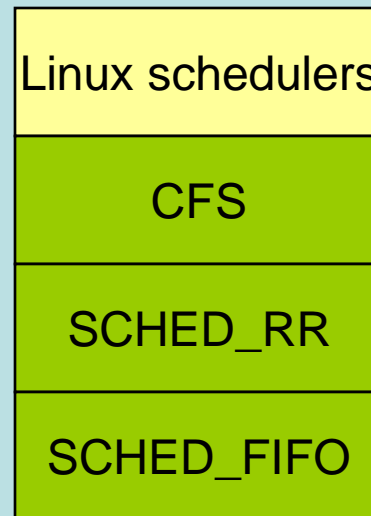
- What is LITMUS^{RT}?
- What is this talk about?
- What is a typical scheduling policy?
- Why do we need a test tool?
- How do we test? (Answer: Unit Testing)
- What are the specific tests?

LITMUS^{RT}

Linux **T**estbed for **M**ultiprocessor **S**cheduling in **R**eal-**T**ime Systems

Patch to Linux 2.6.24 Kernel

Before the patch:

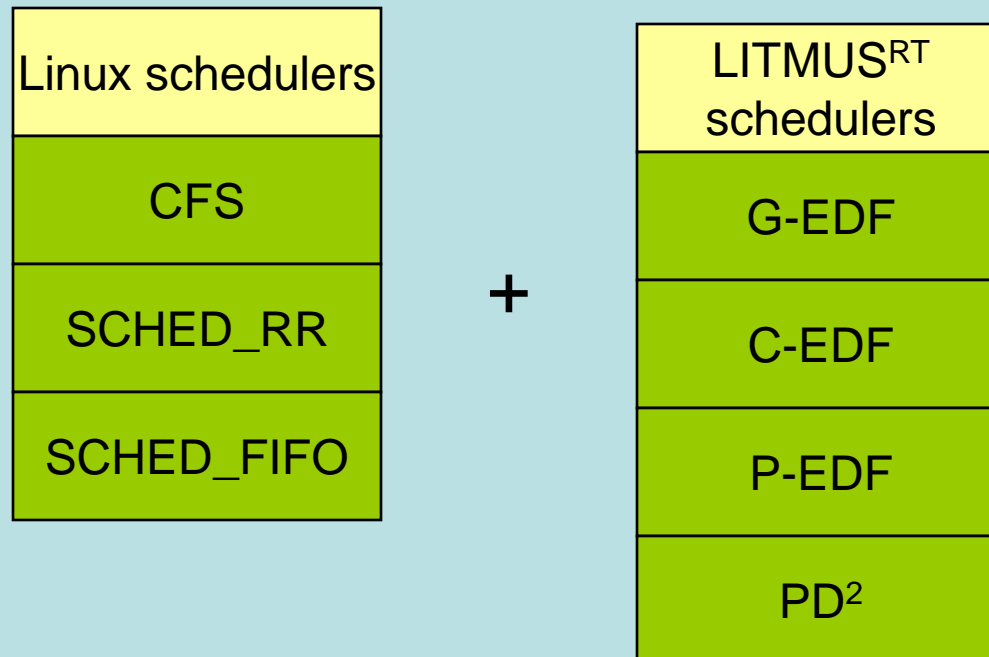


LITMUS^{RT}

Linux **T**estbed for **M**ultiprocessor **S**cheduling in **R**eal-**T**ime **S**ystems

Patch to Linux 2.6.24 Kernel

After the patch:




Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

- What is LITMUS^{RT}?
- **What is this talk about?**
- What is a typical scheduling policy?
- Why do we need a test tool?
- How do we test? (Answer: Unit Testing)
- What are the specific tests?

Overview

- What LITMUS^{RT} is
- Why we want to test LITMUS^{RT} schedulers
 - *Implementing real-time schedulers is nontrivial – bugs can be subtle*
- How to test LITMUS^{RT} schedulers
 - *Unit Testing - testing small pieces of code programmatically – **with a twist***



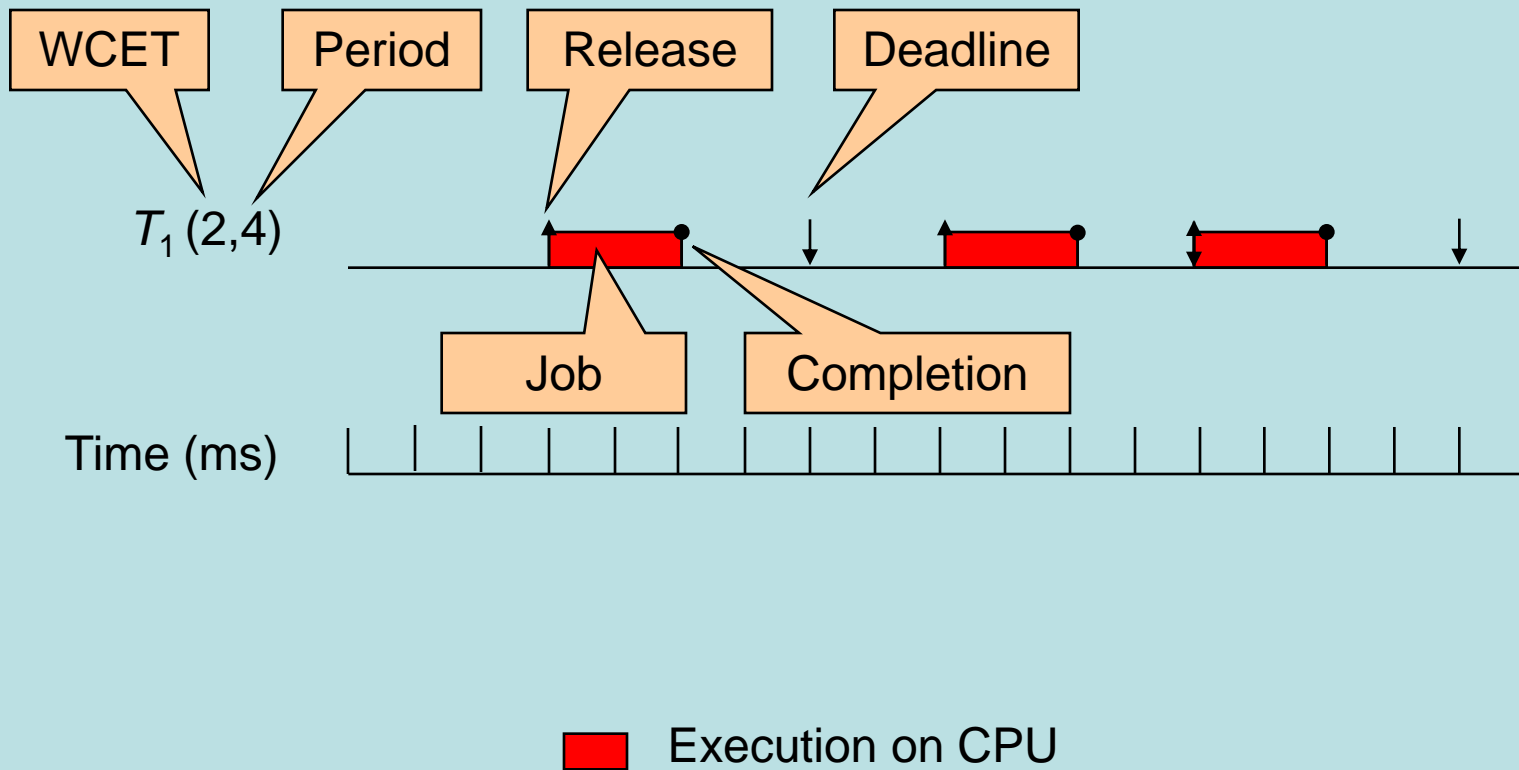
My Work

- Developed specification for a tool to test schedulers
- Implemented prototype of the tool for the G-EDF scheduling policy


Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

- What is LITMUS^{RT}?
- What is this talk about?
- **What is a typical scheduling policy?**
- Why do we need a test tool?
- How do we test? (Answer: Unit Testing)
- What are the specific tests?

Sporadic Task Model



Global Scheduling Policies



M CPUs

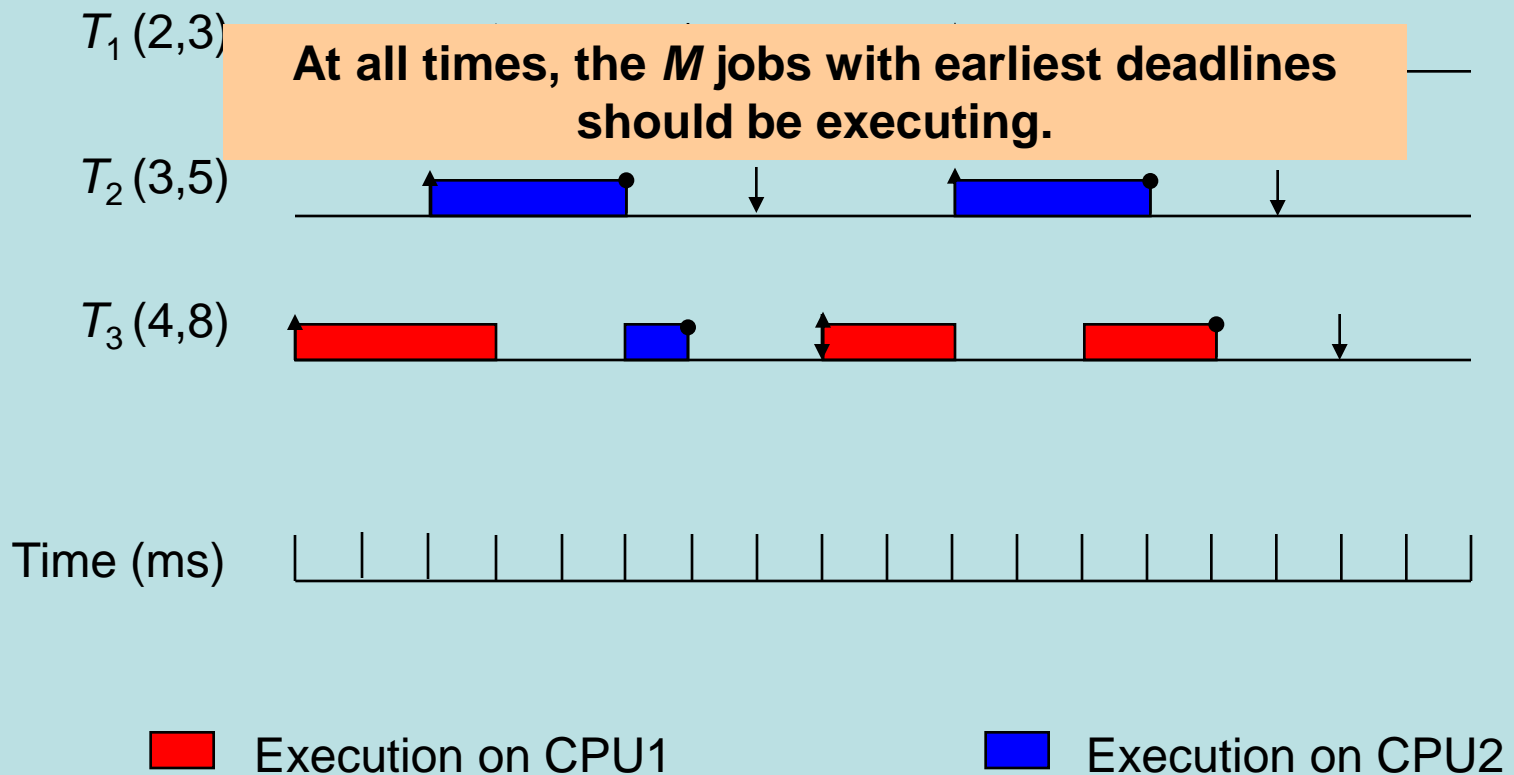
CPU 1

CPU 2


CPU 3

CPU 4

Global Earliest Deadline First (G-EDF)



Global Earliest Deadline First (G-EDF)



Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

- What is LITMUS^{RT}?
- What is this talk about?
- What is a typical scheduling policy?
- **Why do we need a test tool?**
- How do we test? (Answer: Unit Testing)
- What are the specific tests?

Why Develop a Test Tool?

1) Code is very complex

- We cannot easily tell if the code is correct
- We need help debugging

```
1 sched litmu... sched gsn ...
sched_litmus.c (~/kernels/litmus2008-mod/litmus) - GVIM
File Edit Tools Syntax Buffers Window Help
litmus->tick(p);
}
#define NO_CPU -1
static void litmus_schedule(struct rq *rq, struct task_struct *prev)
{
    struct rq* other_rq;
    long prev_state;
    lt_t_maybe_deadlock = 0;
    /* WARNING: rq is not locked! */
    if (is_realtime(prev))
        update_time_litmus(rq, prev);

    /* let the plugin schedule */
    rq->litmus_next = litmus->schedule(prev);

    /* check if a global plugin pulled a task from a different RQ */
    if (rq->litmus_next && task_rq(rq->litmus_next) != rq) {
        /* we need to migrate the task */
        other_rq = task_rq(rq->litmus_next);
        TRACE_TASK(rq->litmus_next, "migrate from %d\n", other_rq->cpu)

        /* while we drop the lock, the prev task could change its
         * state
         */
        prev_state = prev->state;
        mb();
        spin_unlock(&rq->lock);

        /* Don't race with a concurrent switch.
         * This could deadlock in the case of cross or circular migrat
         * It's the job of the plugin to make sure that doesn't happen
         */
        TRACE_TASK(rq->litmus_next, "stack in use=%d\n",
            rq->litmus_next->rt_param.stack_in_use);
        if (rq->litmus_next->rt_param.stack_in_use != NO_CPU) {
            TRACE_TASK(rq->litmus_next, "waiting to deschedule\n");
            _maybe_deadlock = litmus_clock();
        }
        while (rq->litmus_next->rt_param.stack_in_use != NO_CPU) {
            cpu_relax();
        }
    }
}

sched_gsn_edf.c (~/kernels/litmus2008-mod/litmus) - GVIM1
File Edit Tools Syntax Buffers Window Help
for(last = lowest_prio_cpu();
    edf_preemption_needed(&gsnedf, last->linked);
    last = lowest_prio_cpu()) {
    /* preemption necessary */
    task = __take_ready(&gsnedf);
    TRACE("check_for_preemptions: attempting to link task %d to %d\n",
        task->pid, last->cpu);
    if (last->linked)
        requeue(last->linked);
    link_task_to_cpu(task, last);
    preempt(last);
}

/* gsnedf_job_arrival: task is either resumed or released */
static ninline void gsnedf_job_arrival(struct task_struct* task)
{
    BUG_ON(!task);

    requeue(task);
    check_for_preemptions();
}

static void gsnedf_release_jobs(rt_domain_t* rt, struct heap* tasks)
{
    unsigned long flags;

    spin_lock_irqsave(&gsnedf_lock, flags);

    __merge_ready(rt, tasks);
    check_for_preemptions();

    spin_unlock_irqrestore(&gsnedf_lock, flags);
}

/* caller holds gsnedf_lock */
static ninline void job_completion(struct task_struct *t, int forced)
{
    BUG_ON(!t);

    sched_trace_task_completion(t, forced);
}

36,42-45
286,2-5
```

Why Develop a Test Tool?


1) Code is very complex

- We cannot easily tell if the code is correct
- We need help debugging

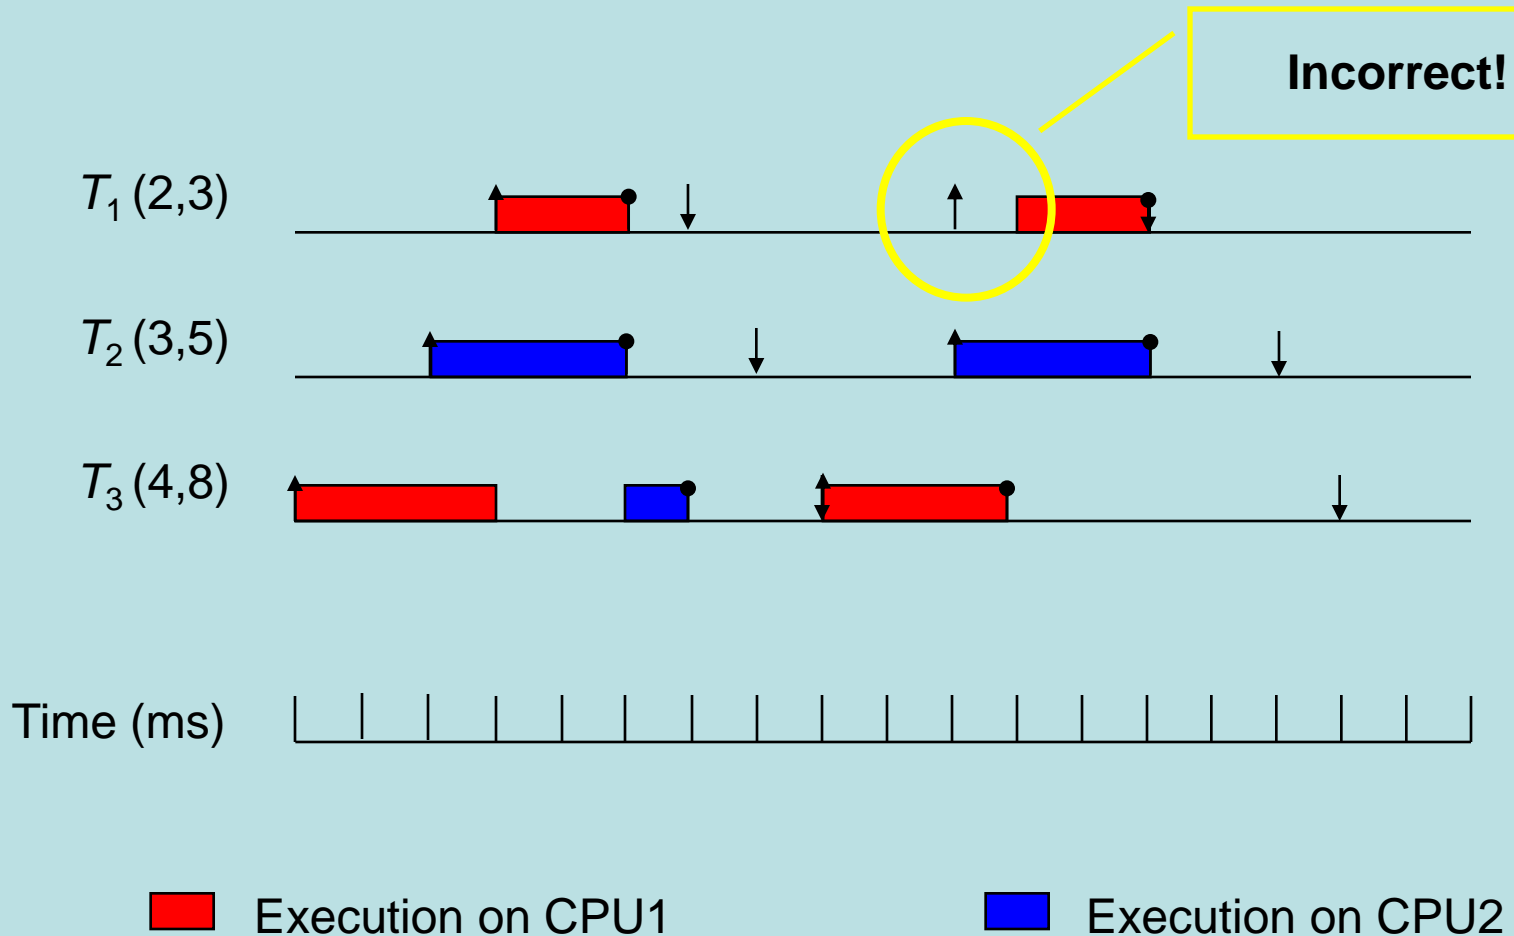
2) Resulting schedules are very complex

- We cannot easily tell if correct schedules are produced


Global Earliest Deadline First (G-EDF) ???




Global Earliest Deadline First (G-EDF)



Global Earliest Deadline First (G-EDF)




Global Earliest Deadline First (G-EDF)



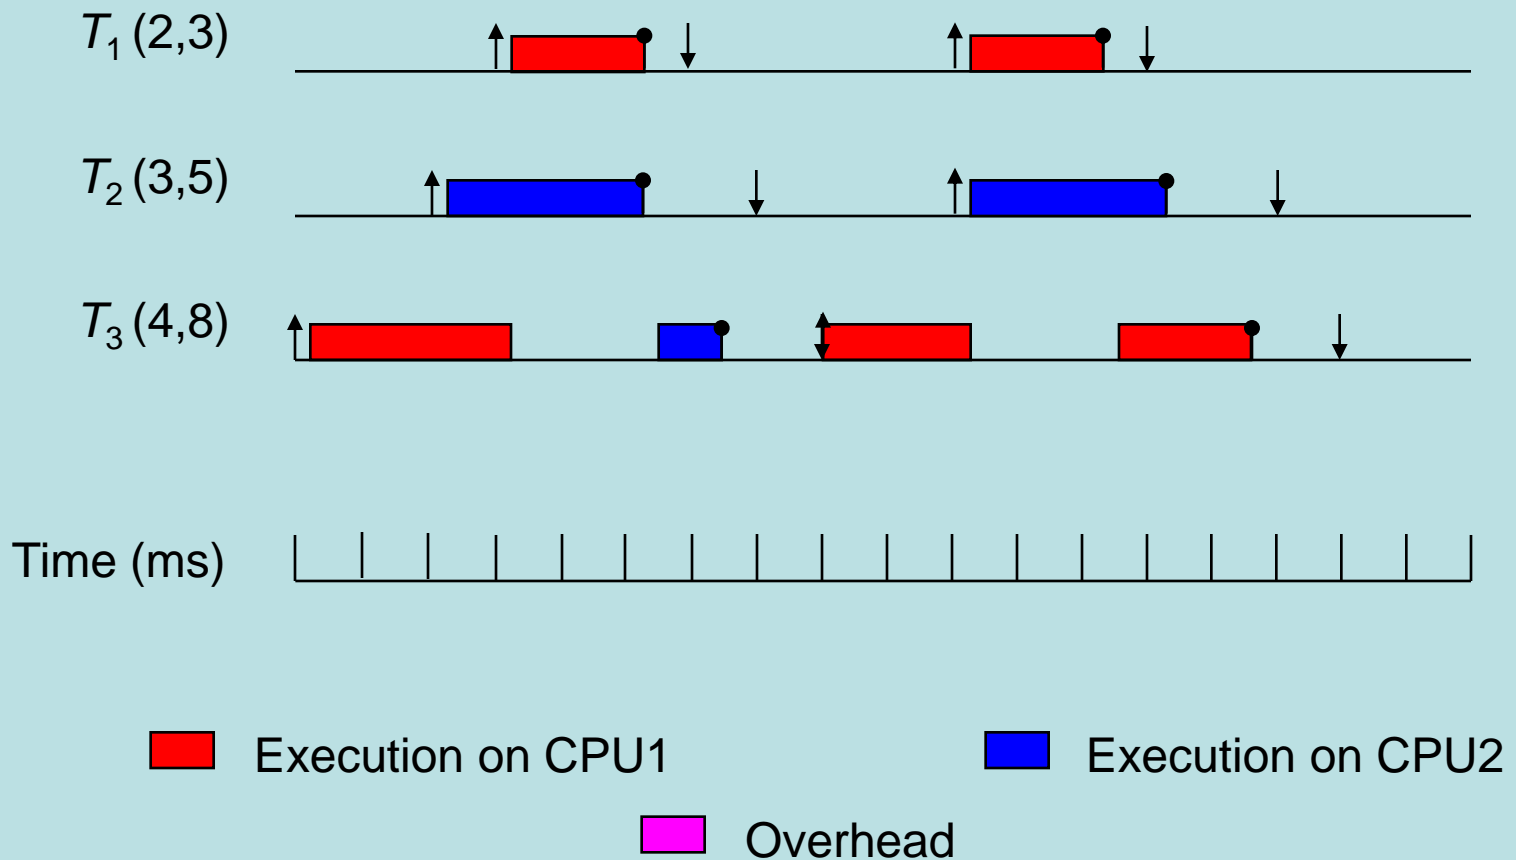
Why Develop a Test Tool?


- 1) Code is very complex
 - We cannot easily tell if the code is correct
 - We need help debugging
- 2) Resulting schedules are very complex
 - We cannot easily tell if correct schedules are produced
- 3) We need to minimize overhead
 - Detailed regression testing is necessary

Overhead



Overhead





Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

- What is LITMUS^{RT}?
- What is this talk about?
- What is a typical scheduling policy?
- Why do we need a test tool?
- How do we test? (Answer: Unit Testing)
- What are the specific tests?

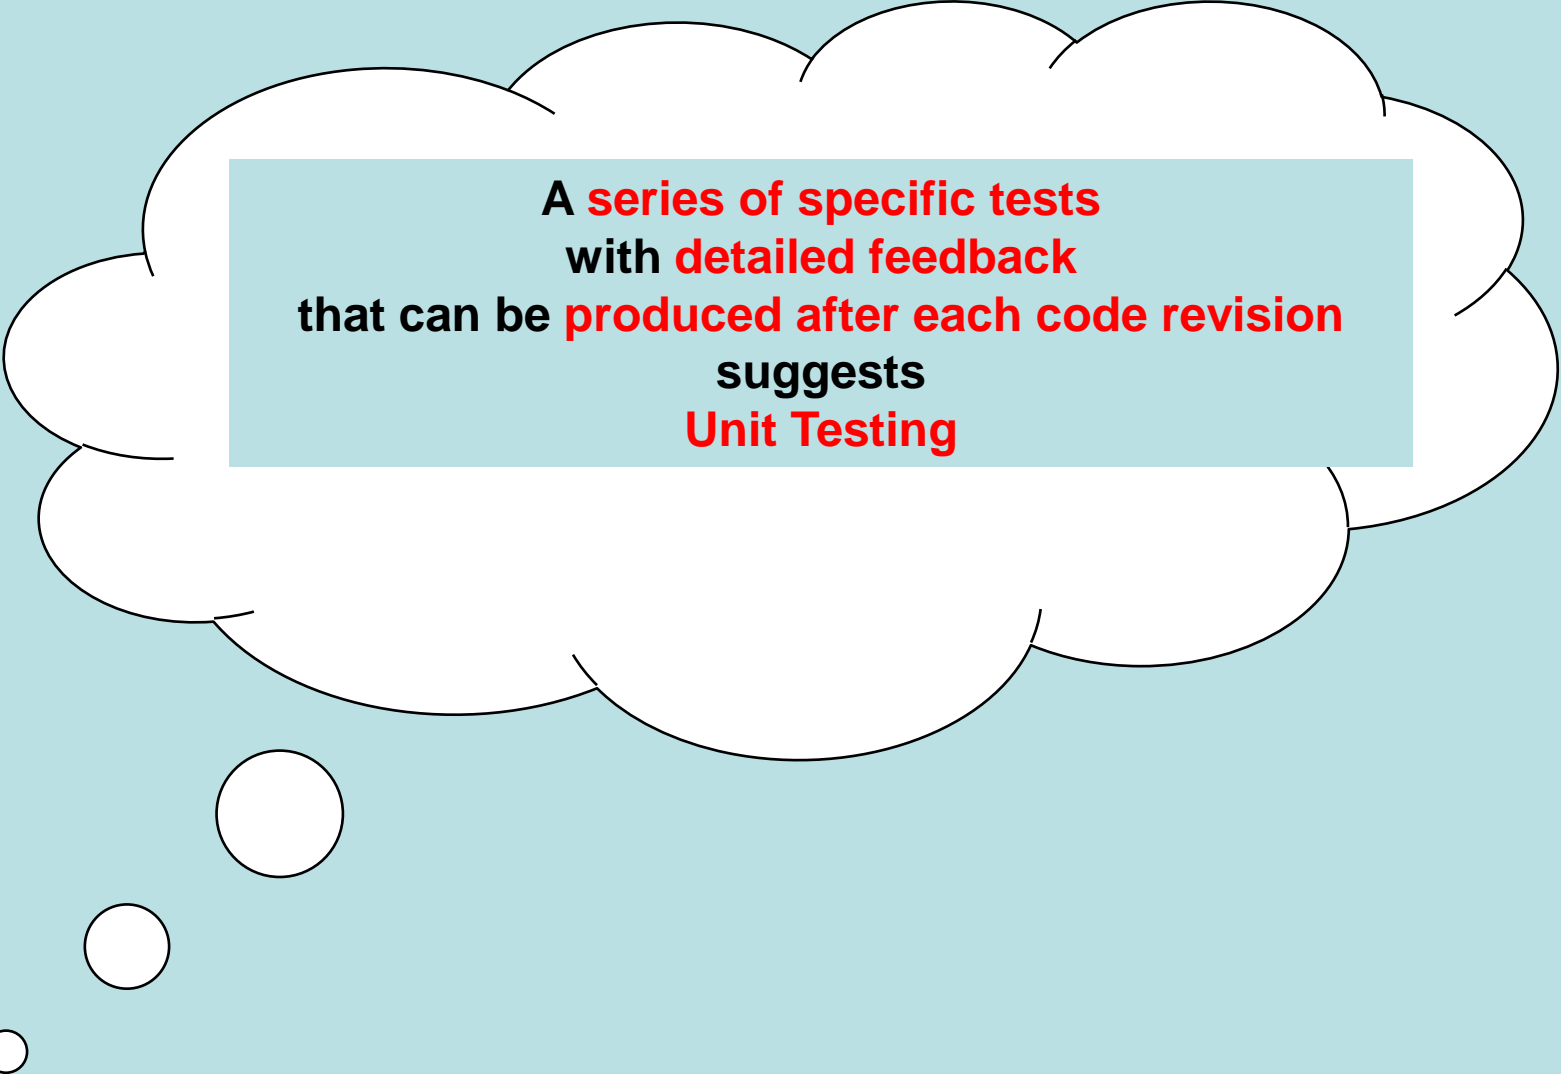
The Challenge

**Without overhead, and if we did not need very detailed feedback, we could check the invariant:
At all times, the M jobs with earliest deadlines
should be executing.**

The Challenge

Instead, we have to use a series of specific tests to detect anomalies in scheduling and to measure overhead.

Unit Testing



**A series of specific tests
with detailed feedback
that can be produced after each code revision
suggests
Unit Testing**

Unit Testing

A **series of specific tests**
with **detailed feedback**
that can be **produced after each code revision**
suggests
Unit Testing

Unit Testing: programmatically
testing small
modules of *code*
after each revision


Unit Testing

A **series of specific tests**
with **detailed feedback**
that can be **produced after each code revision**
suggests
Unit Testing

Unit Testing: programmatically
testing small
modules of *code*
after each revision

We test
recorded
schedule
“traces” instead

Feather-Trace



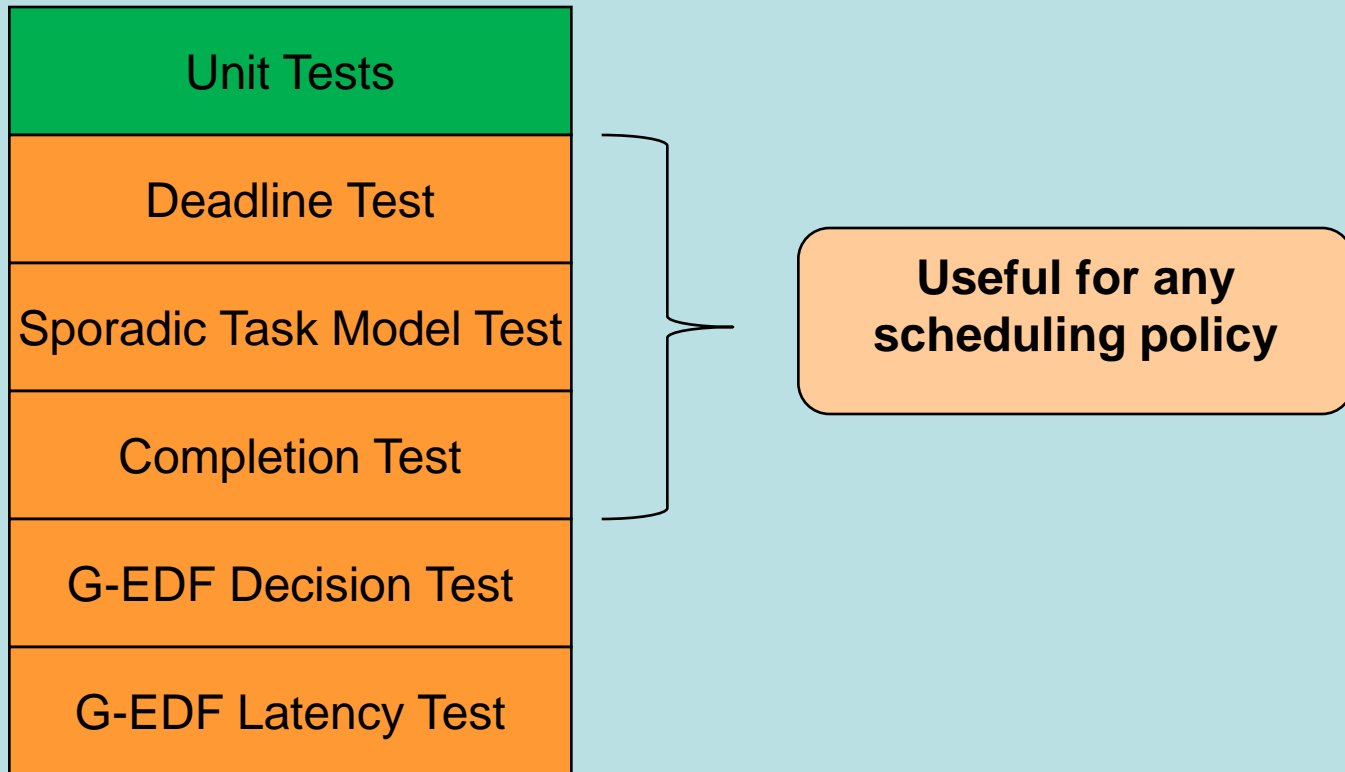
Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

- What is LITMUS^{RT}?
- What is this talk about?
- What is a typical scheduling policy?
- Why do we need a test tool?
- How do we test? (Answer: Unit Testing)
- What are the specific tests?

Unit Tests

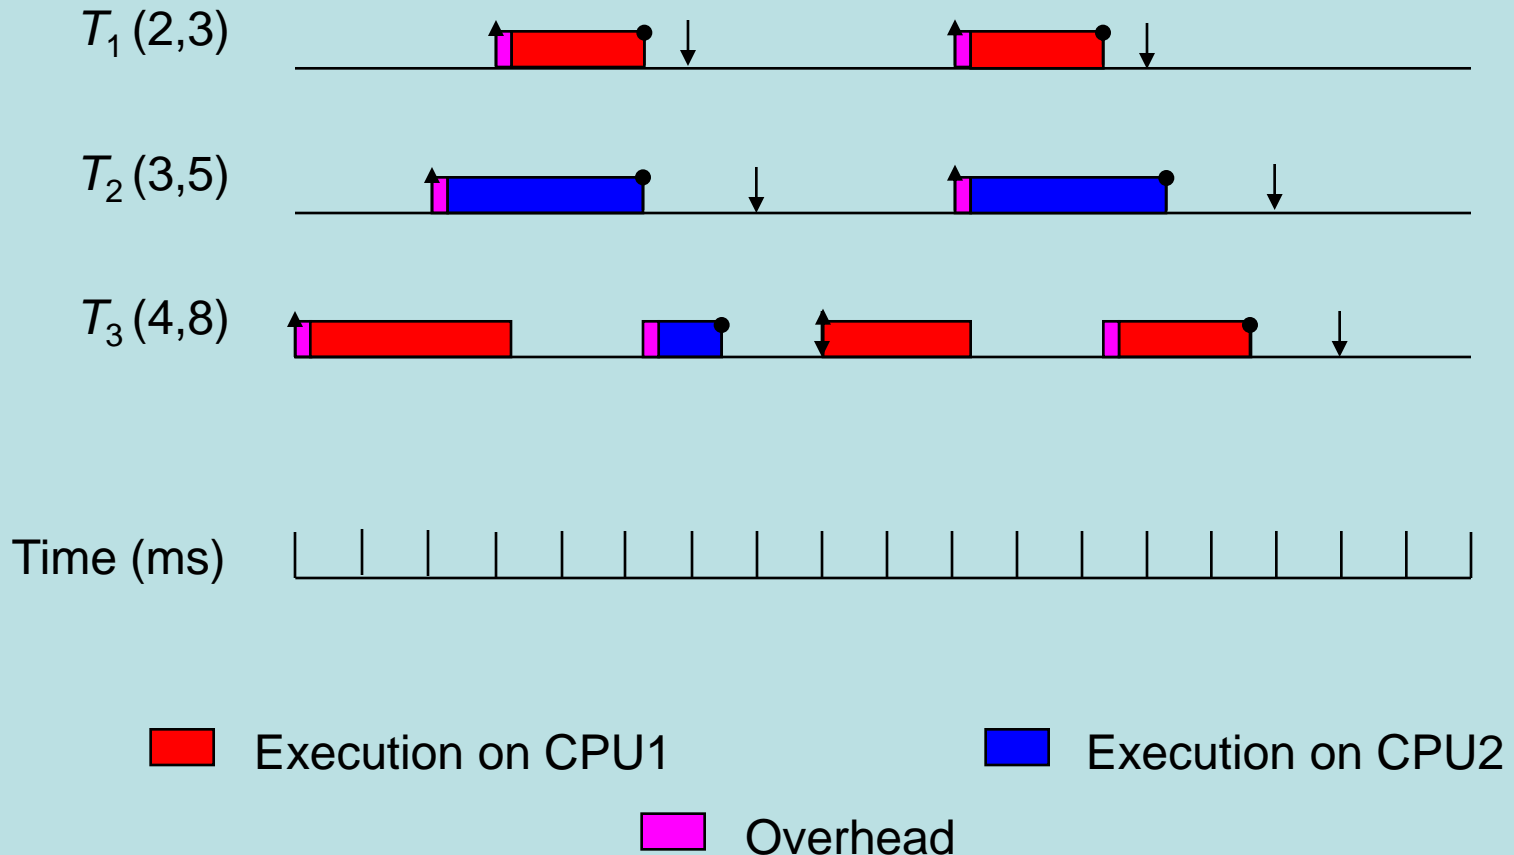
Unit Tests
Deadline Test
Sporadic Task Model Test
Completion Test
G-EDF Decision Test
G-EDF Latency Test

Unit Tests




Deadline Test

Did all jobs complete by their deadlines?



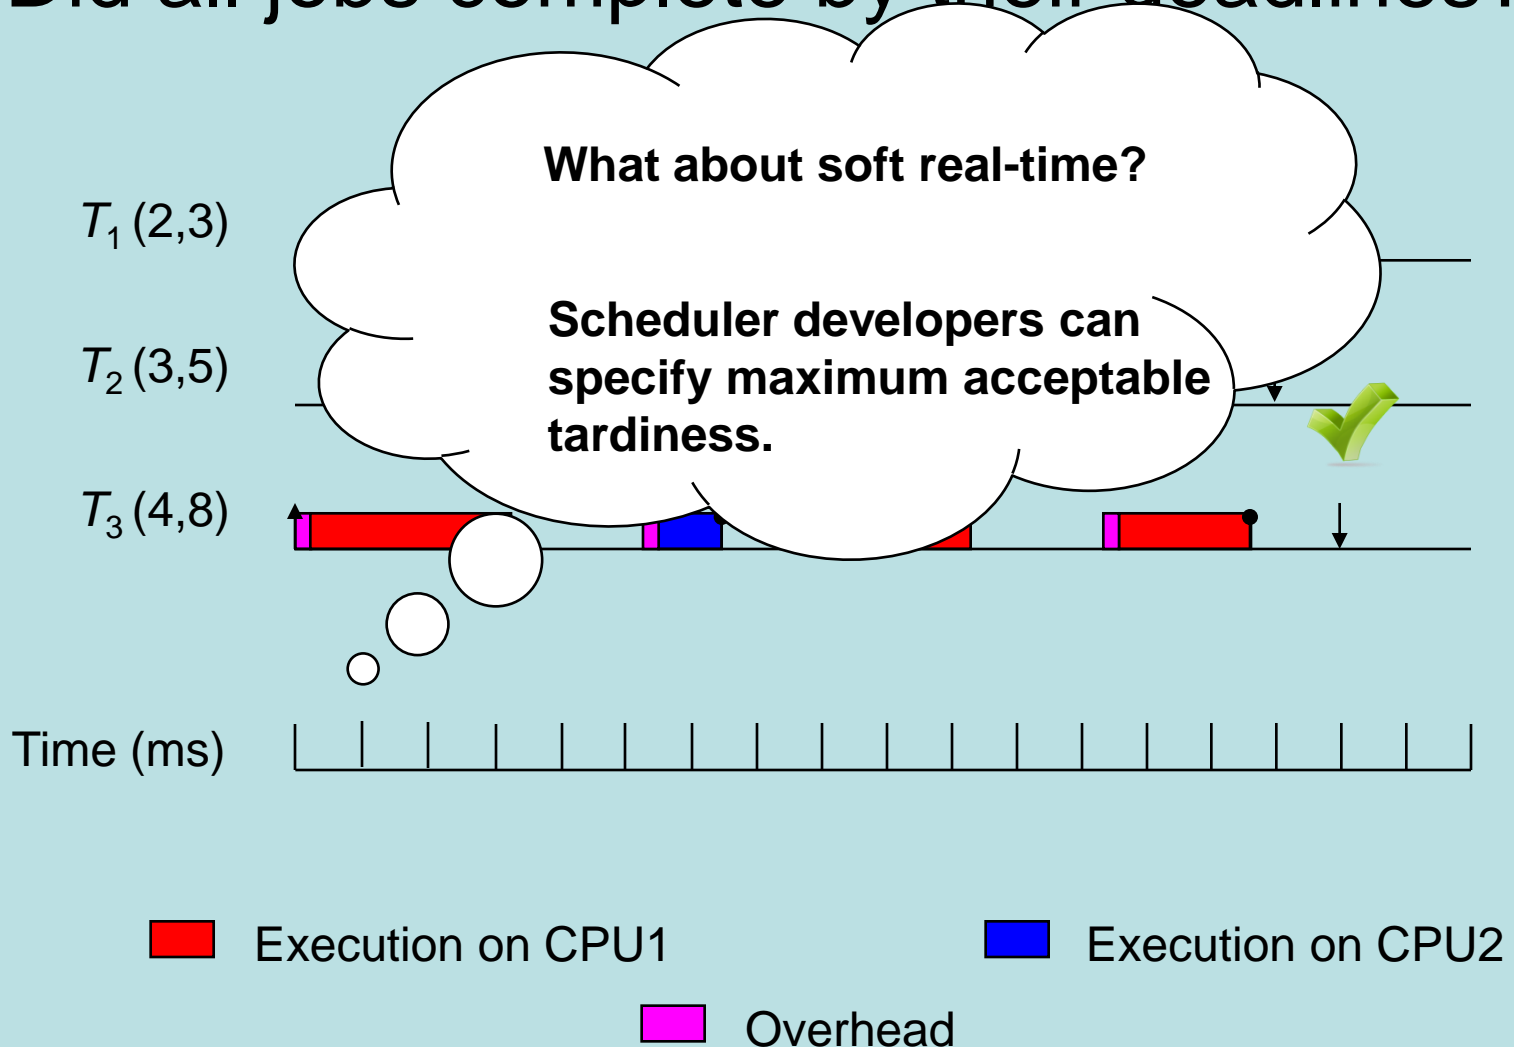
Deadline Test

Did all jobs complete by their deadlines?




Deadline Test

Did all jobs complete by their deadlines?

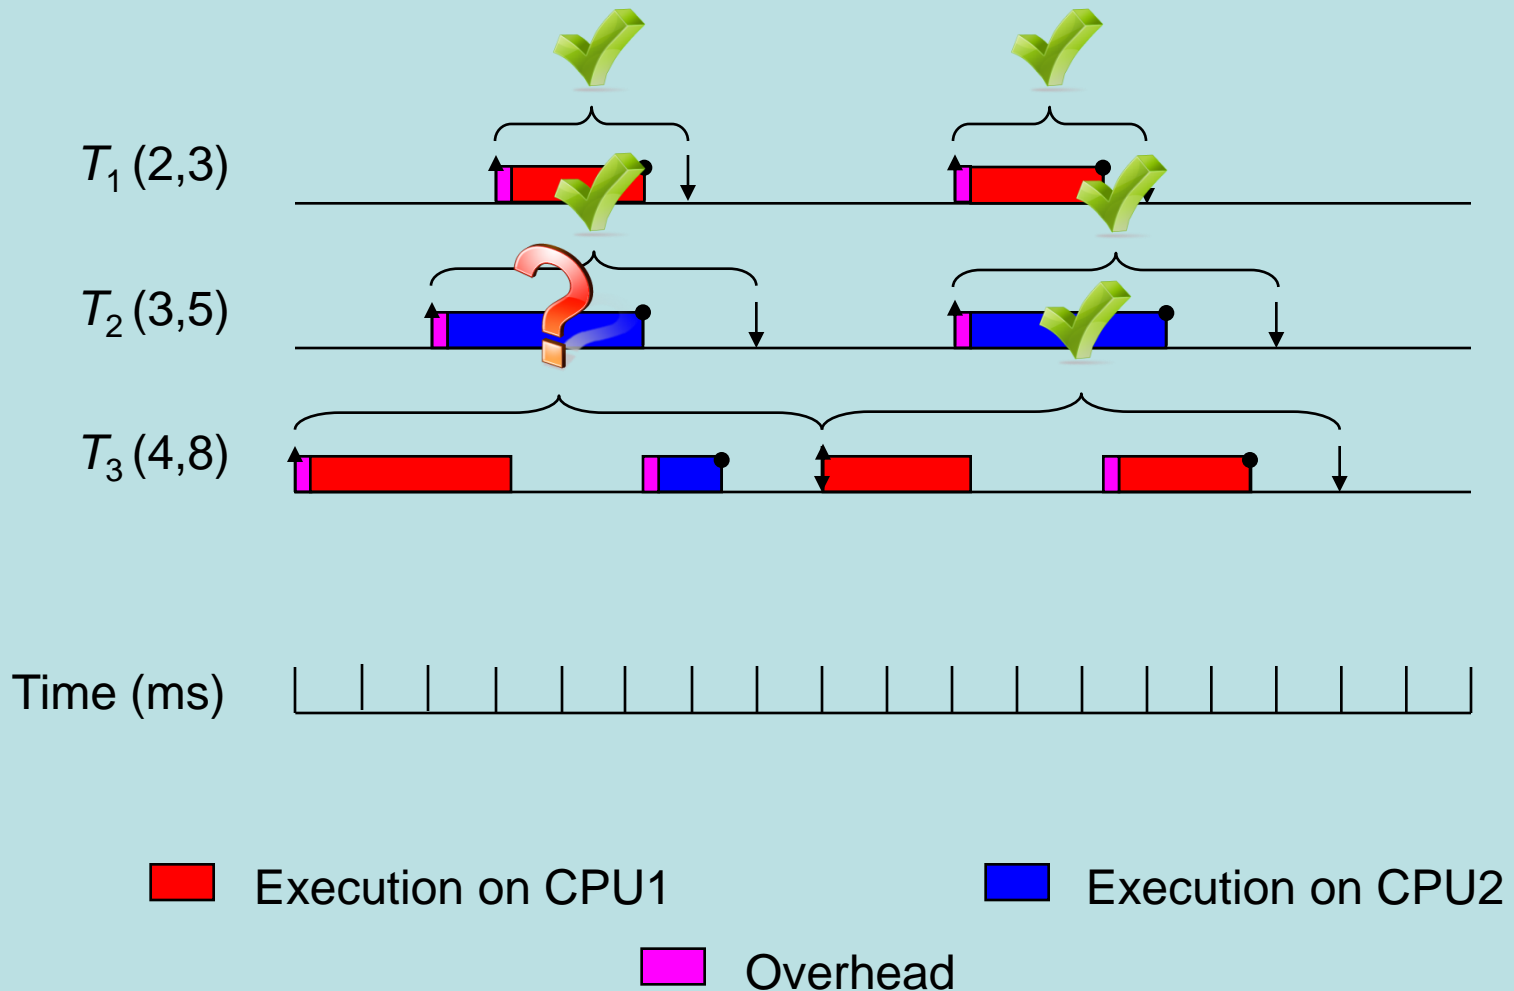


Sporadic Task Model Test

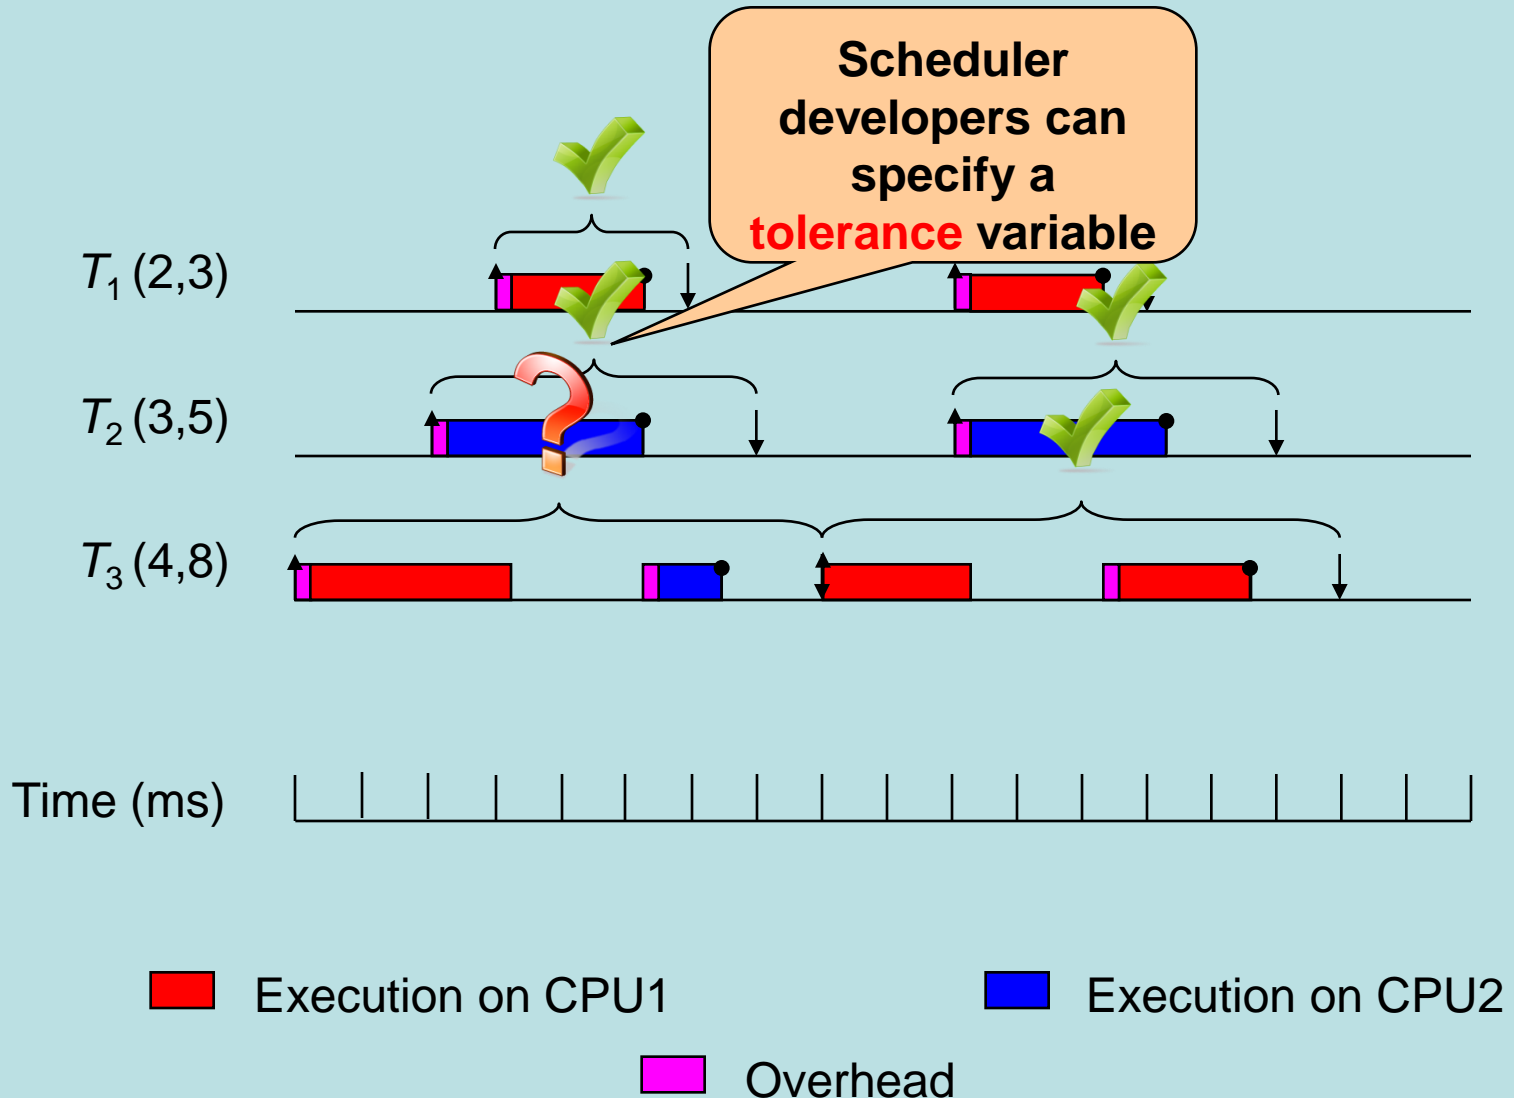
Were job releases separated by at least the period of the task?



Sporadic Task Model Test




Sporadic Task Model Test



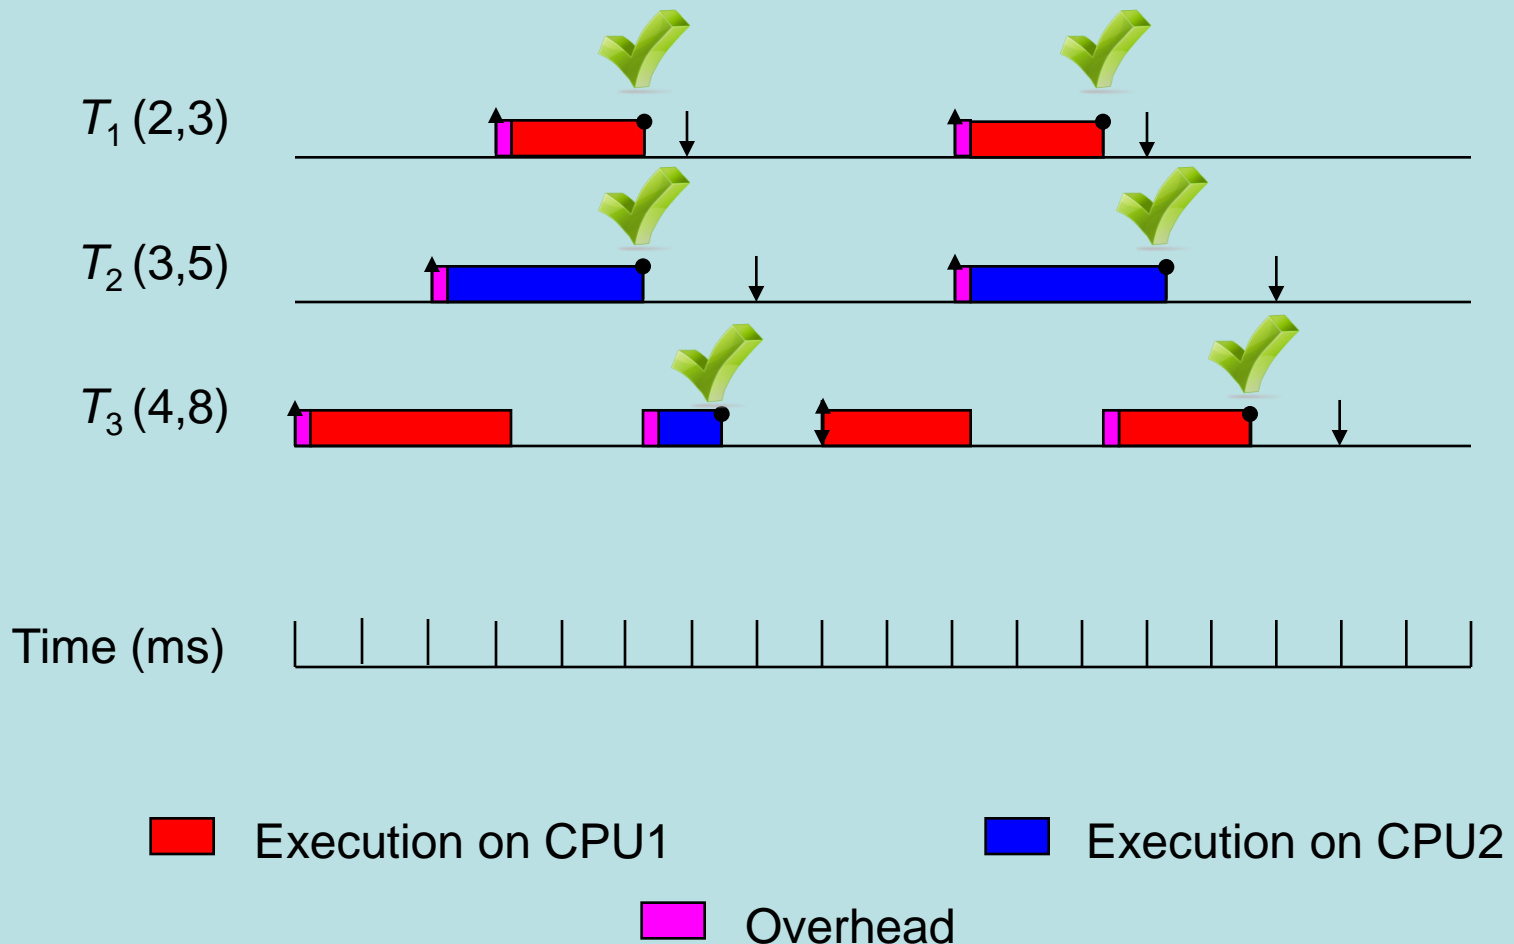
Completion Test

Did all released jobs actually complete?

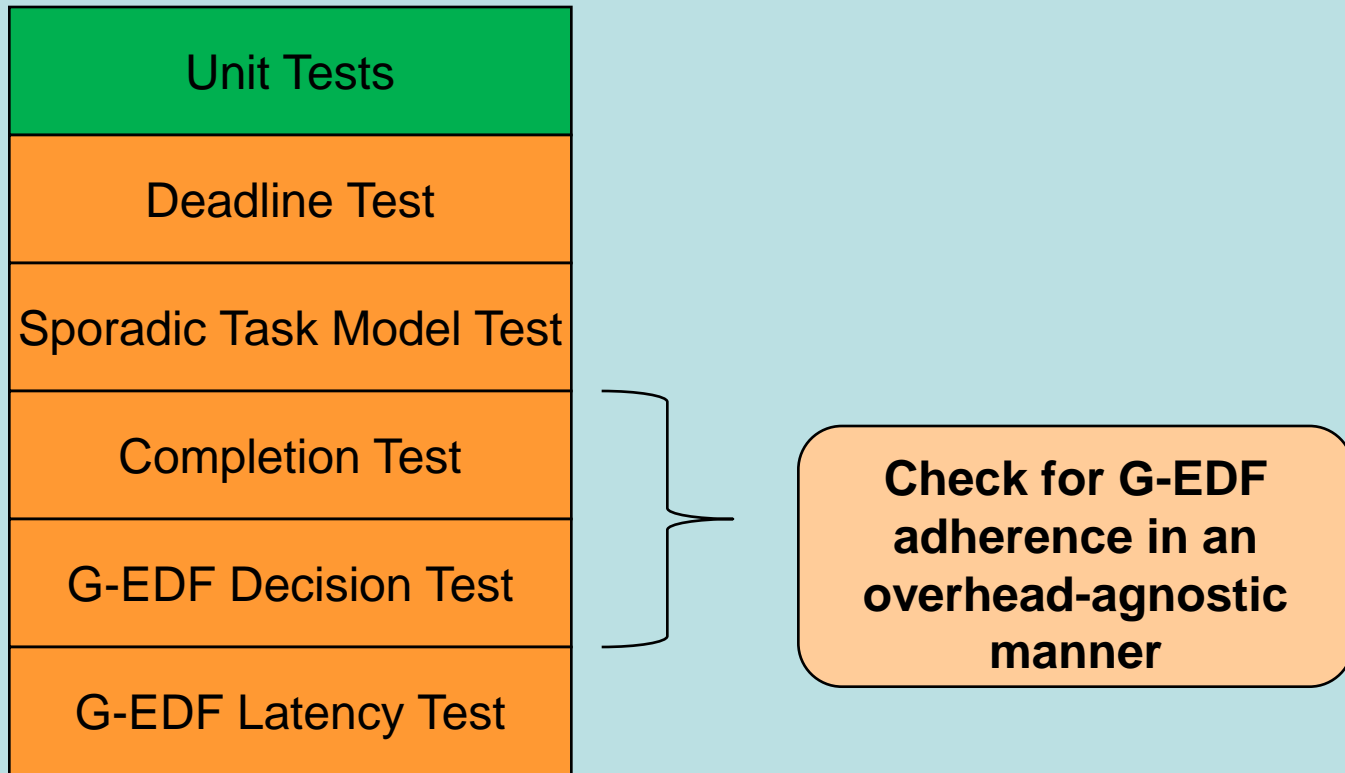


Completion Test

Did all released jobs actually complete?




Unit Tests




G-EDF Decision Test

Are jobs switched to execution in EDF order?




G-EDF Decision Test

Are jobs **Test algorithm models execution state to check for correct decisions.** execution



G-EDF Decision Test

Are jobs switched to execution
in EDF order?




G-EDF Decision Test

Are jobs switched to execution

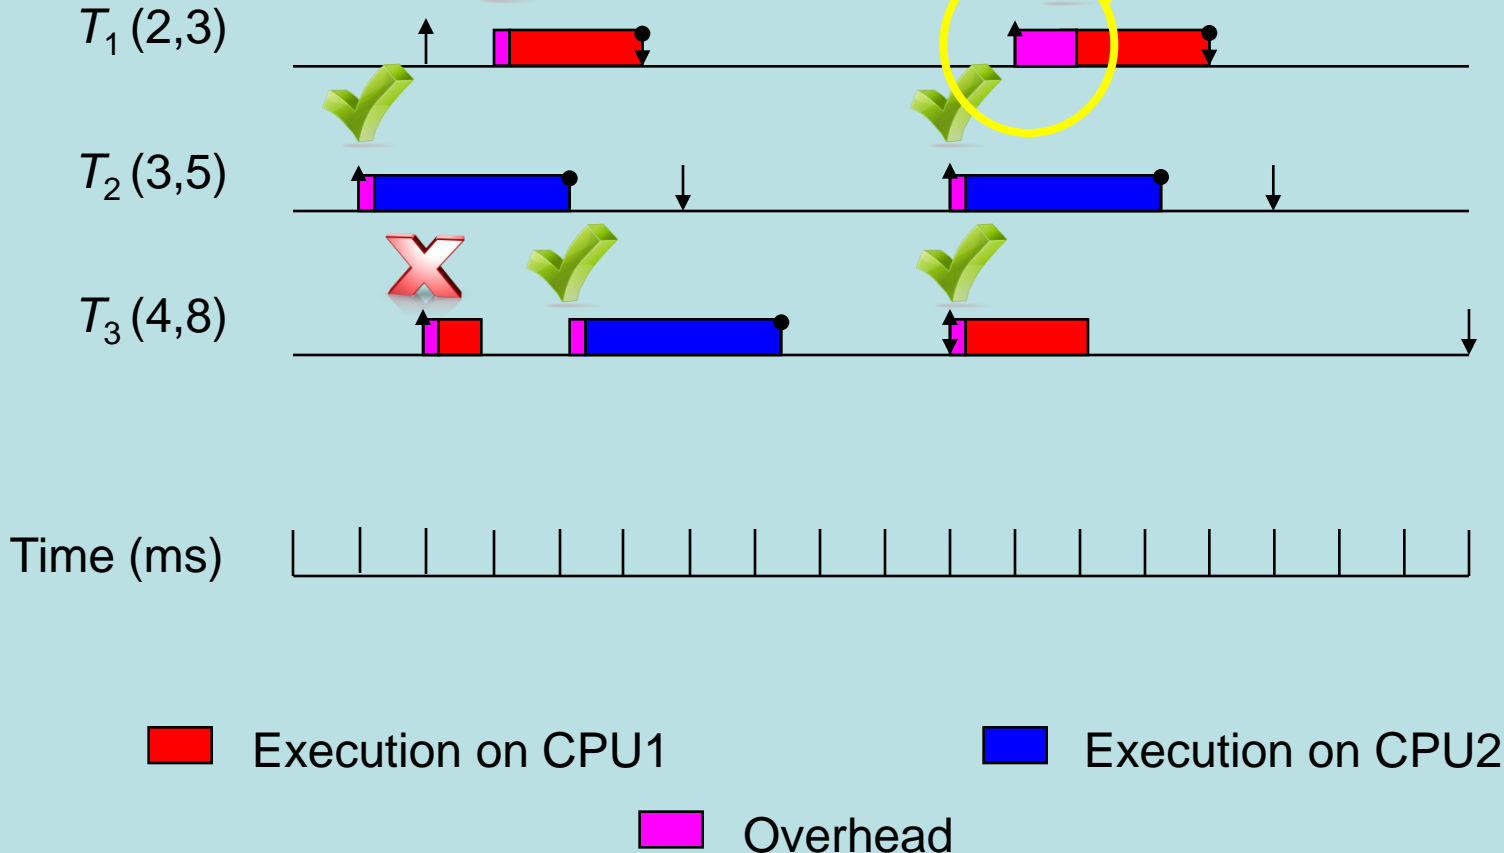
in EDF order?

Incorrect?

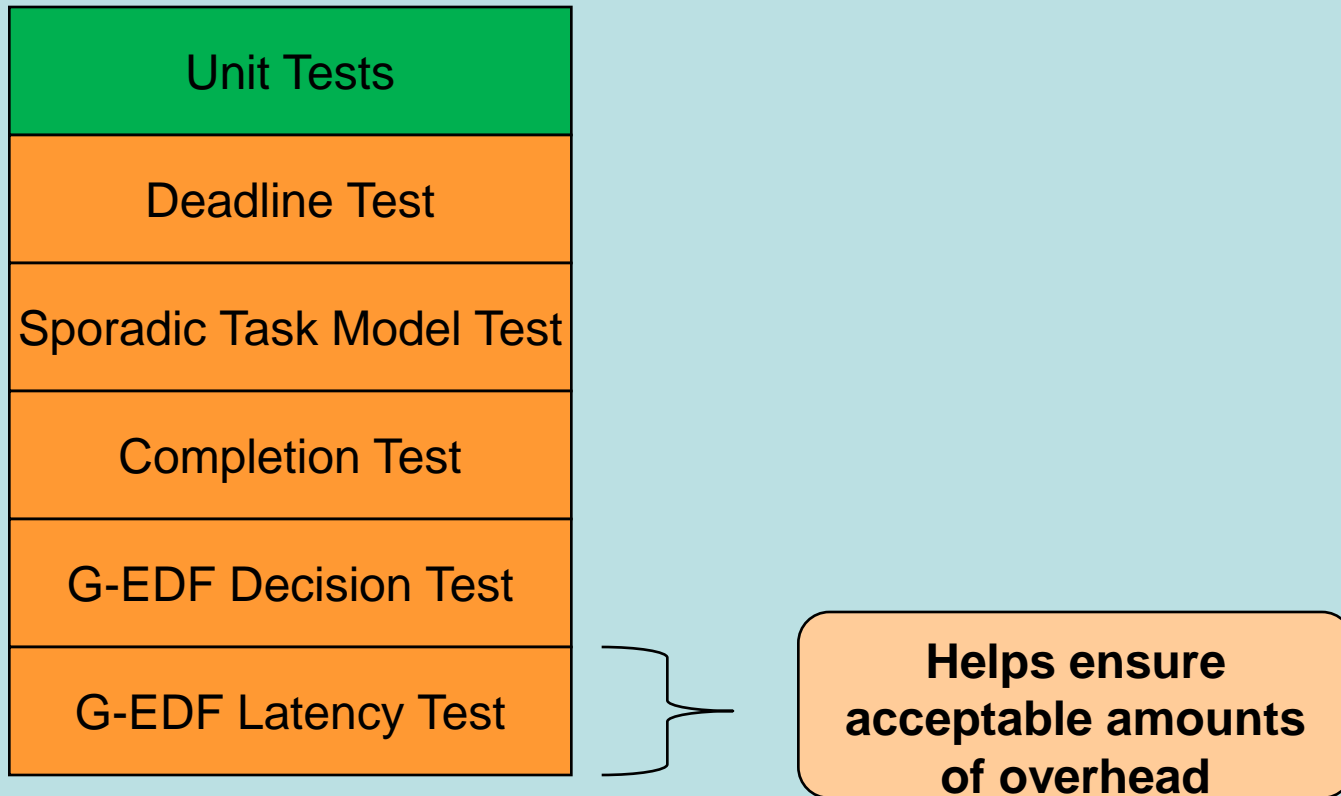


G-EDF Decision Test

Are jobs T_1 and T_2 executed in EDF order?
This cannot be distinguished from overhead.
Incorrect?




Unit Tests




G-EDF Latency Test

Measures latency



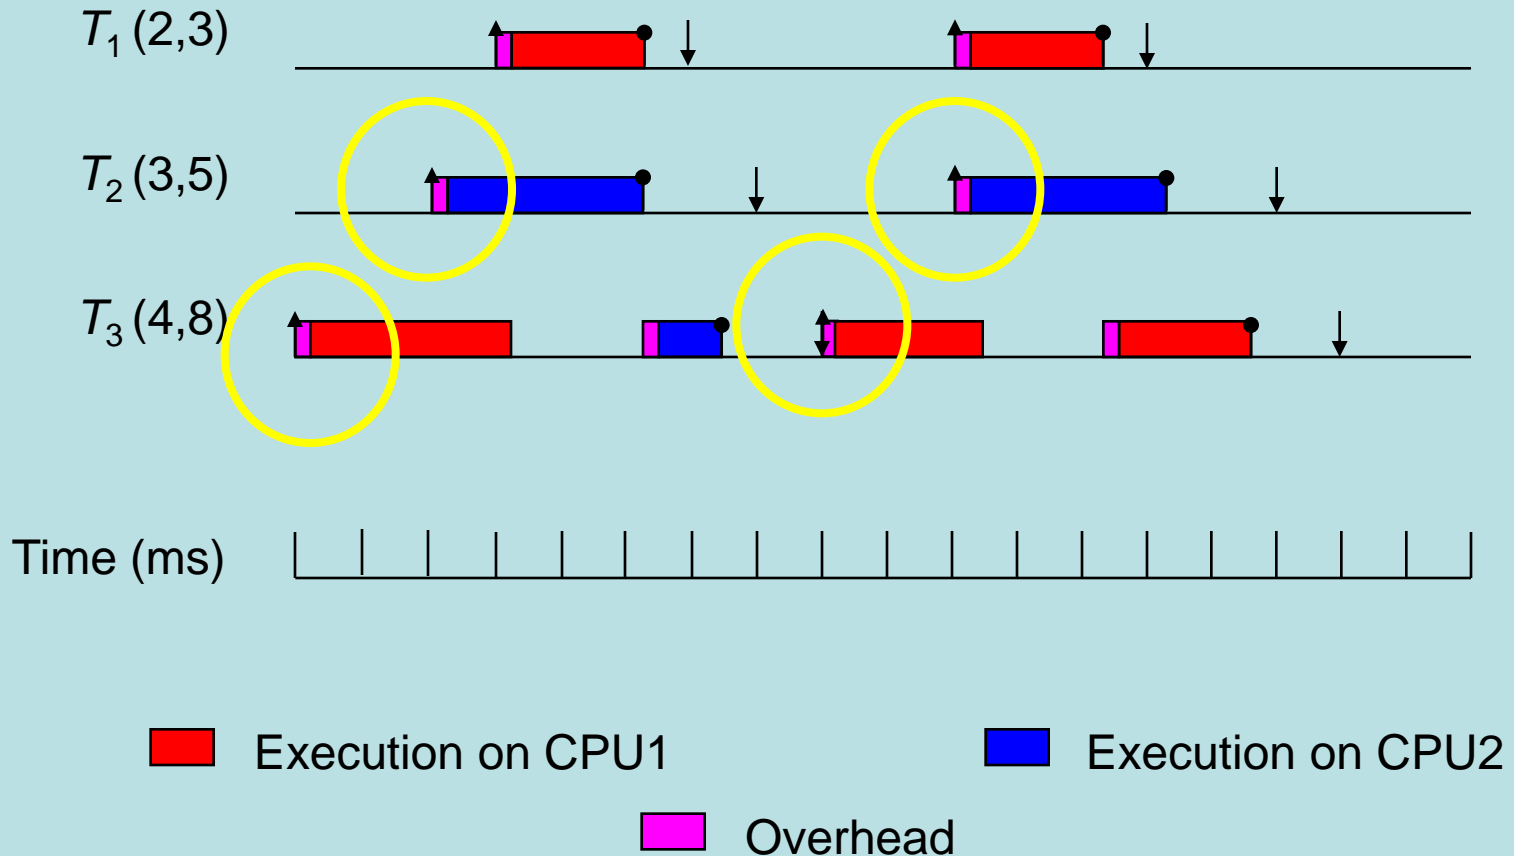
G-EDF Latency Test

Test algorithm can use sequence of events on each processor to determine type of latency




G-EDF Latency Test

Type 1




G-EDF Latency Test

Type 2



G-EDF Latency Test

Type 3



Summary

- What LITMUS^{RT} is
- Why we want to test LITMUS^{RT} schedulers
 - *Implementing real-time schedulers is nontrivial – bugs can be subtle*
- How to test LITMUS^{RT} schedulers
 - *Unit Testing - testing small pieces of code programmatically – **with a twist***

Towards Unit Testing Real-Time Schedulers in LITMUS^{RT}

Questions?