

Problem detection in real-time systems by trace analysis

Mathieu Côté

Laboratoire DORSAL

mathieu.cote@polymtl.ca



**POLYTECHNIQUE
MONTRÉAL**



Outline

- Introduction
- Literature review
- Approach
 - Modeling
 - Problems
 - Analysis
- Results
- Conclusion



Introduction : definition

- Real-time task : execution time, deadline, period (optional)
- Execution : periodic, sporadic
- Hard/soft real-time

PREEMPT_RT

- Priority inheritance for mutex in kernel
- Reduce non-preemptive sections in kernel

Introduction : problematic

Introduction

Literature

Modeling

Problems

Analysis

Results

Conclusion

Music player trace in Trace Compass




Figure 1 : Multiple executions of an audio player

Introduction : problematic

Advantages of tracing real-time systems

- Low **overhead**
- Low **jitter**
- Access to **specific** information (priority, scheduling policy, etc.)

What is missing?

- Real-time **specific** user tools
- Show **useful** data

Introduction : goals

Introduction

Literature

Modeling

Problems

Analysis

Results

Conclusion

1. Develop a **model** to define real-time task **executions** in a trace
2. Identify common **problems** in real-time systems and useful **information** to analyze them
3. Develop a method to analyze the **trace segment** corresponding to an execution to identify if the execution presents a **problem**

Literature review

Linux low-latency tracing for multicore hard real-time systems
(Beamonte, 2013)

- LTTng-UST **modification** to reduce the added latency
- Demonstrated **low latency** tracing with LTTng

Literature review

Real-time Linux analysis using low-impact tracer (Rajotte, 2014)

- Recreate the task states using kernel events
- Compare executions of a task
- Limitations
 - Model
 - Threads need to have different priorities
 - Fixed
 - Analysis
 - Manual
 - Some statistics

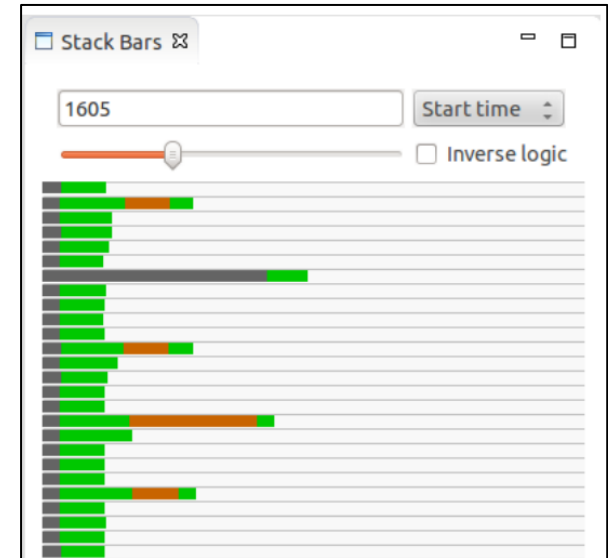


Figure 2 : Original stackbars view

Modeling

Advantage of using only kernel events

- No need to modify the application source code to add tracepoints manually

Modeling : view

Stackbars view in Trace Compass




Figure 3 : Stackbars view

Modeling : view

States in *Stackbars* view

- Running : in userspace or in system calls
- Ready : between sched_wakeup and sched_switch
- Blocked or preempted : when you are still in a task execution but are scheduled out




Figure 4 : *Stackbars* view legend

Modeling

- Identify executions automatically and then let the users choose between some valid models
 - Estimate the number of executions
 - Find the longest subsequence repeated at least n times
 - Difficulties :
 - Execution time
 - Too many possible resulting models

Modeling : method

State machine

- User identifies :
 - an execution or
 - events that defined the start and the end (name, parameters with operations, etc.)

Events Selection

Enter the deadline for this execution (-1 for none)

-1

Enter start event name or blank for default

sched_wakeup | | sched_wakeup_new

Enter start event params ("param1=value1, param2=value2") or blank for none

tid=\$tid

Enter the tid(s) for the start event (blank for current only, separate by coma)

Enter end event name or blank for default

sched_switch

Enter end event params ("param1=value1, param2=value2") or blank for none

prev_state!=0,prev_tid=\$tid

Enter the tid(s) for end event (blank for end event to be on the same thread than the corresponding start event)

Select the new depth to change events for (Upper = 0). Current = 0

Change current depth selection

OK Cancel

Figure 5 : Dialog to define model

Modeling : method

State machine

- Remove execution
- Add execution
- Define an execution as invalid and recalculate
 - Will suggest some modifications to the model based on differences between valid and invalid executions
 - The user can select the ones he wants to apply

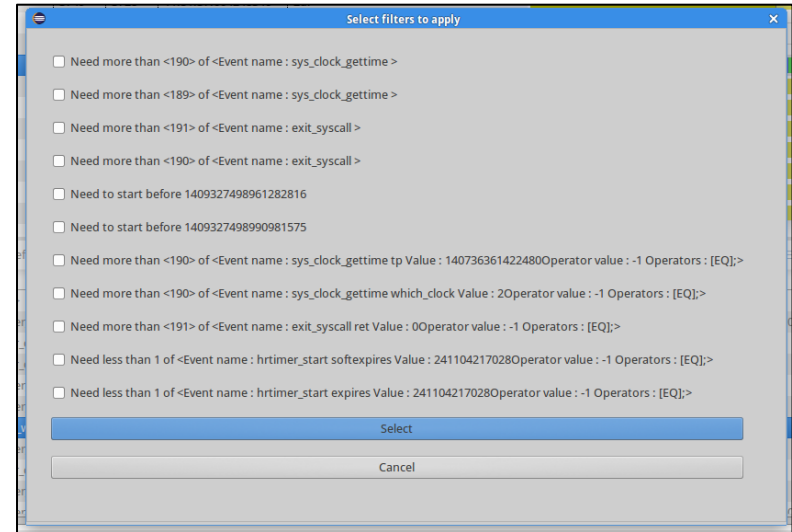


Figure 6 : Dialog to select modifications to apply

Modeling : method

State machine

- Supports
 - Thread pool
 - Nested executions

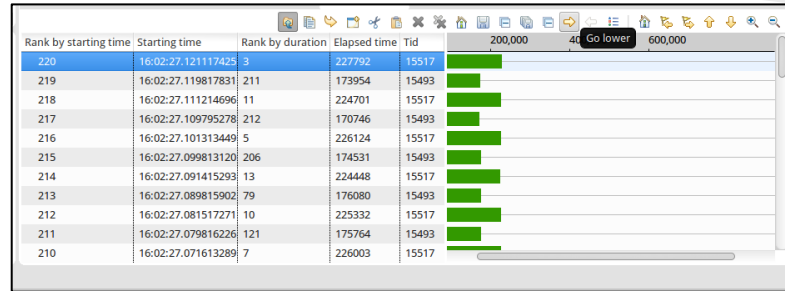


Figure 7 : Task on multiple threads




Figure 8 : Nested executions

Specific informations

Scheduling policies

- Normal
 - SCHED_OTHER : standard
 - SCHED_BATCH
 - SCHED_IDLE
- Real-time
 - SCHED_FIFO
 - SCHED_RR : with time quantum
 - SCHED_DEADLINE : Global Earliest Deadline First, highest user controllable priority

Specific informations

Scheduling policies

- SCHED_FIFO and SCHED_RR
 - A deadline can be missed even if there was a valid scheduling to respect all deadlines to respect all deadlines
- SCHED_DEADLINE
 - No deadline will be missed if there is a valid scheduling

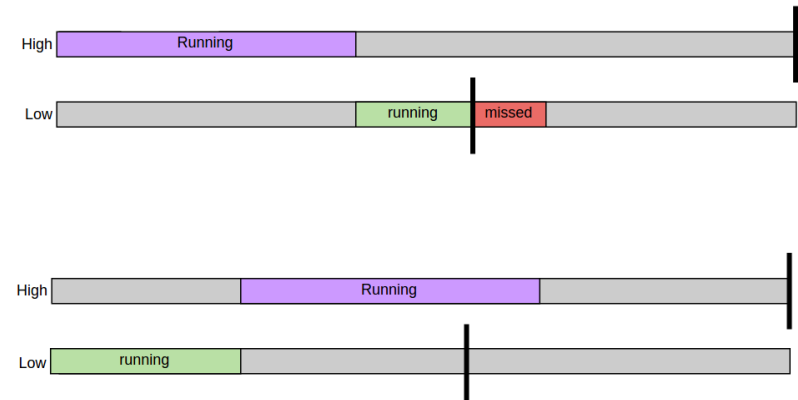


Figure 9 : Deadline missed

Specific informations

Scheduling policies

- SCHED_FIFO and SCHED_RR
 - The highest priority task will always execute if it is able to
- SCHED_DEADLINE
 - If there is a missed deadline, it can be on a highest priority task (for the user, because there is no priority set)

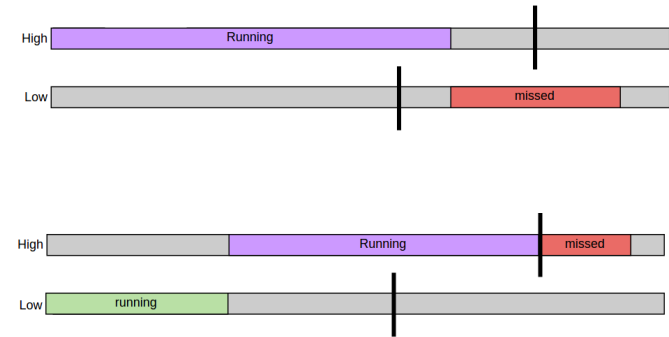


Figure 10 : Highest priority

Specific informations

Events to track to get *policy* :

`sched_setscheduler`, `sched_setparam`, `sched_setattr`

Additional events to track to get *priority* :

`setpriority`, `sched_pi_setprio`, `sched_switch`

Events to track to get *cpus_allowed*:

`sched_setaffinity`, need to add some

Results : views

- View of duration by starting timestamp
- Synced with other views

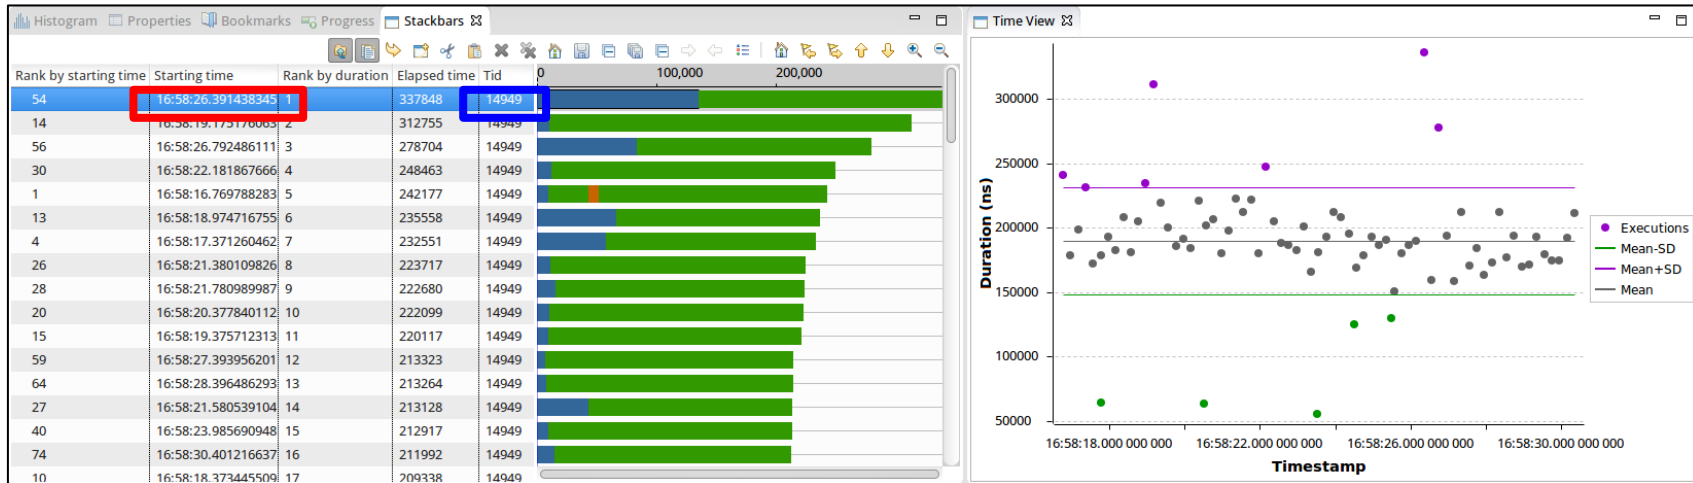



Figure 11 : Stackbars view and stackbars time view

Results : periodic conflict

Analysis for the thread : [8837, ./test_sched]

Priority : -49 from time : 14:08:26.155926228

Policy : SCHED_RR

The analysed thread was preempted from time : 14:08:26.155935758 for : 160916

This thread was running when [8837,./test_sched] was preempted.

First time : 14:08:26.155935758

Thread ID : 8812

Duration : 160917

Priority : -50

Policy : SCHED_RR

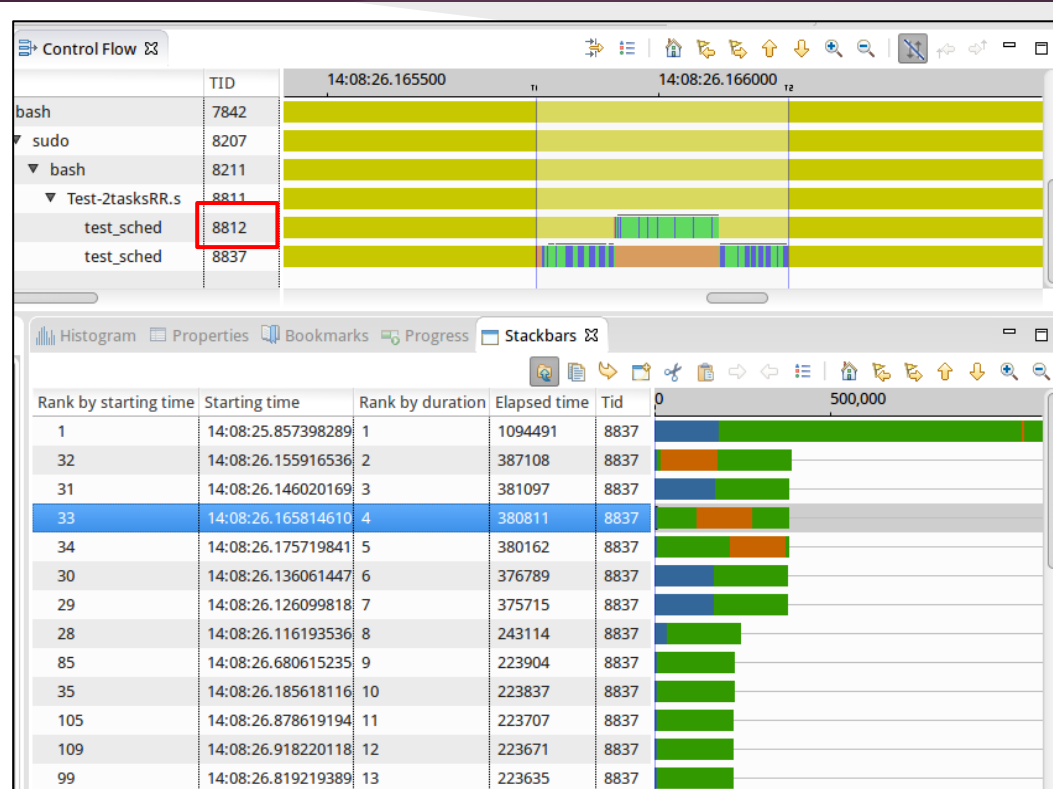


Figure 12 : Periodic conflict

Results : priority inversion

The high priority task is blocked by the low priority task that is preempted because the medium priority task is running

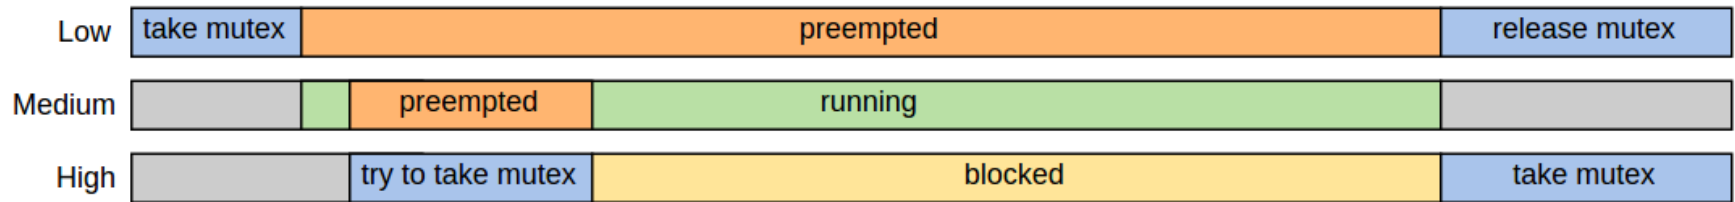


Figure 13 : Priority inversion

Results : priority inversion

Priority ceiling protocol

- Better if the high priority task accesses the resource more often than the low priority task, because it is faster and has fewer context switches, but it can give an unnecessary high priority to the lower task



Figure 14 : Priority ceiling protocol

Results : priority inversion

Priority inheritance

- Better if the low priority task accesses the resource more often

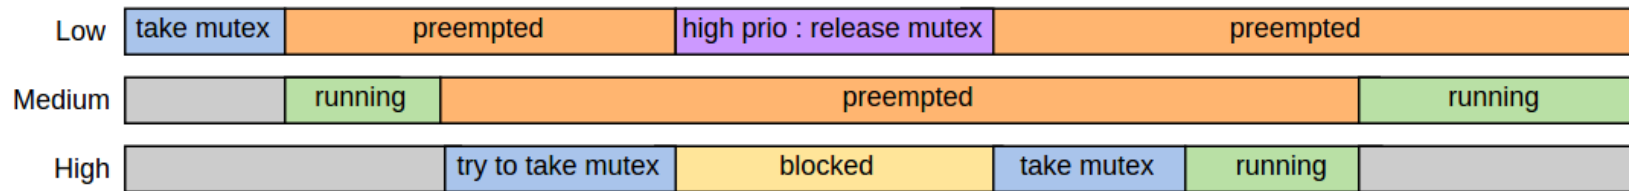


Figure 15 : Priority inheritance

Results : priority inversion

- Introduction
- Literature
- Modeling
- Problems
- Analysis
- Results**
- Conclusion

Tid:15684 -> low priority

Tid:15685 -> medium
priority

Tid:15686 -> high priority

Analysis for the thread : [15686,test_PriorityIn]
 Priority : -96 from time : 16:03:54.507283434
 Policy : SCHED_FIFO

 The thread : [15684,test_PriorityIn] was preempted
 when in the critical path of the analysed thread
 from time : 16:03:54.507316303 for : 10077919 ns
 Priority : 20 (L)

This thread was running when [15684,
 test_PriorityIn] was preempted. (L)
 First time : 16:03:54.507316303
 Thread ID : 15685 (M)
 Duration : 10027986 ns
 Priority : -43
 Policy : SCHED_FIFO

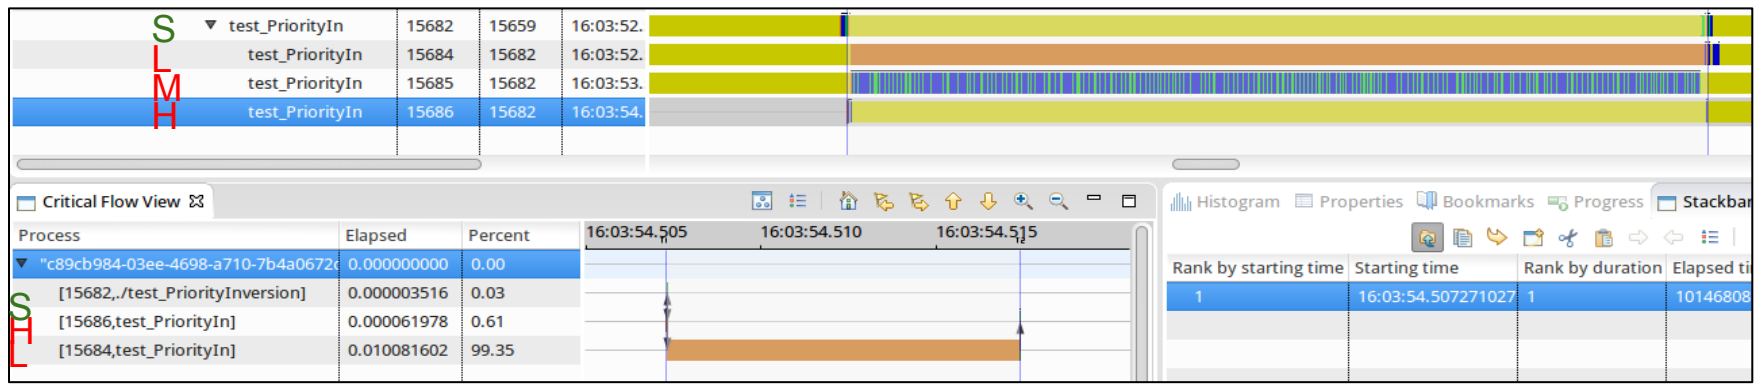


Figure 16 : Priority inversion

Results : priority

Priority inheritance (PTHREAD_PRIO_INHERIT)




Figure 17 : Priority inheritance protocol

Low priority temporarily set to the same priority as the high priority thread (-96) when high is blocked

Results : priority

Priority ceiling (PTHREAD_PRIO_PROTECT)




Figure 18 : Priority ceiling protocol

Low priority set to -96

Other results

- Deadline analysis
 - Tell which executions missed their deadlines
 - User input
 - Get it from events for SCHED_DEADLINE policy

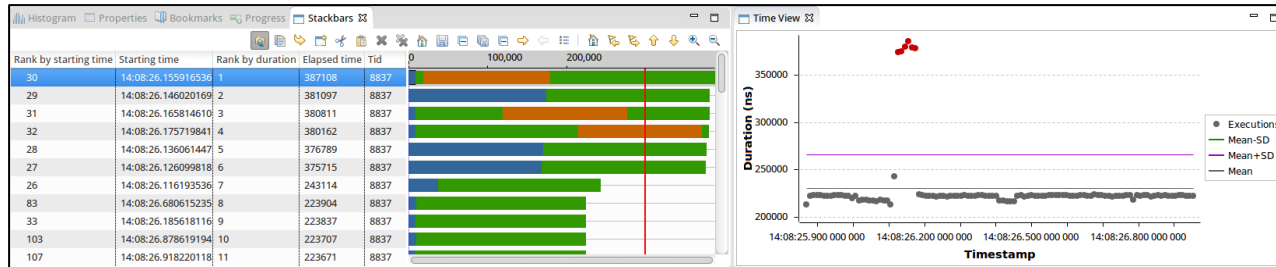


Figure 19 : Deadline

- Device blocked analysis

Conclusion

- Future work
 - Modeling
 - Instrument complex real-time application in user-space and for each task, validate if it is possible to model only with kernel events
 - Analysis
 - Validate with real bugs
 - Add new analysis
- Questions?