

Problem detection in real-time systems by trace analysis

Mathieu Côté

Laboratoire DORSAL

mathieu.cote@polymtl.ca

13 mai 2015



POLYTECHNIQUE
MONTRÉAL



Outline

- Introduction
- Literature review
- Modeling
- Views
- Results
- Conclusion



Introduction : problematic

Music player trace in Trace Compass

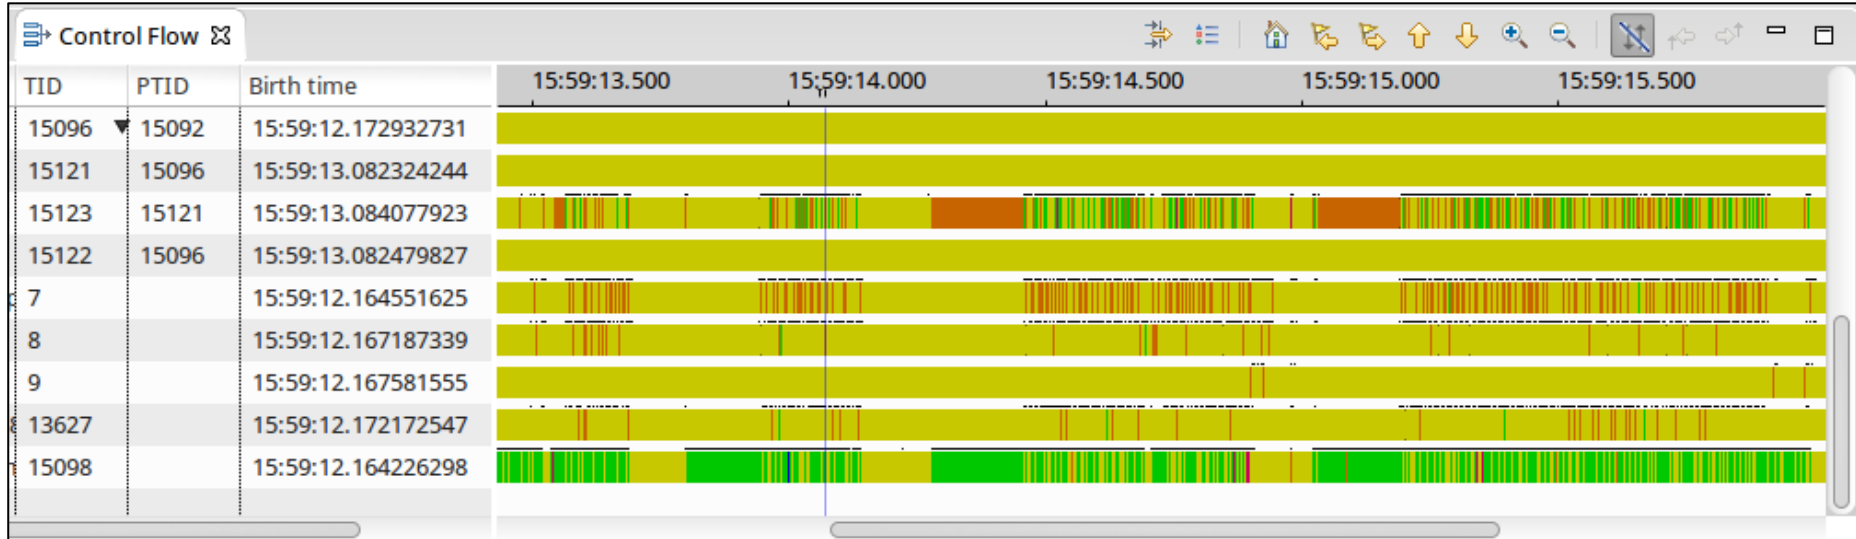


Figure 1 : Multiple executions of an audio player

Introduction : problematic

Introduction

Literature

Modeling

Views

Results

Conclusion

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

Views

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**

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

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

Scheduling policies

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

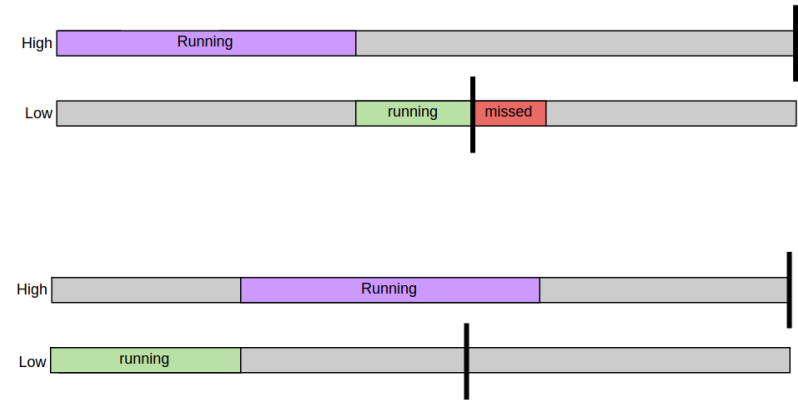


Figure 2 : Deadline missed

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)

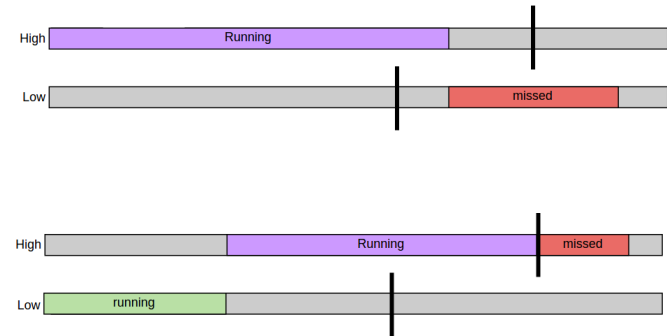


Figure 3 : Highest priority

Priority inversion

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

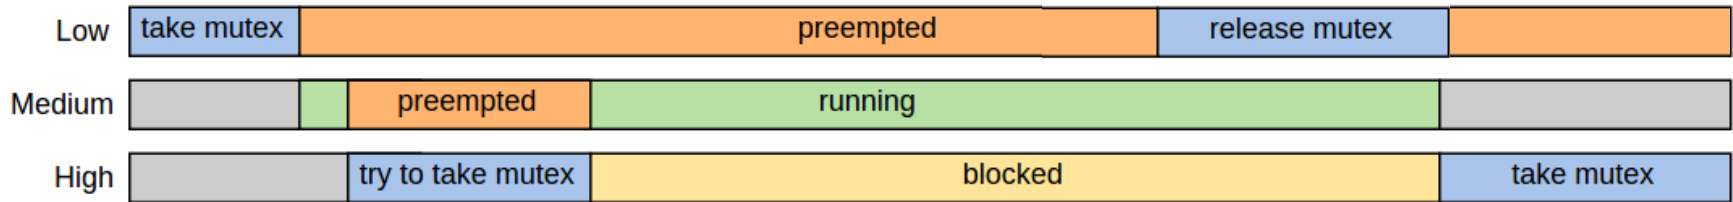


Figure 4 : Priority inversion

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 5 : Priority ceiling protocol

Priority inversion

Priority inheritance

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

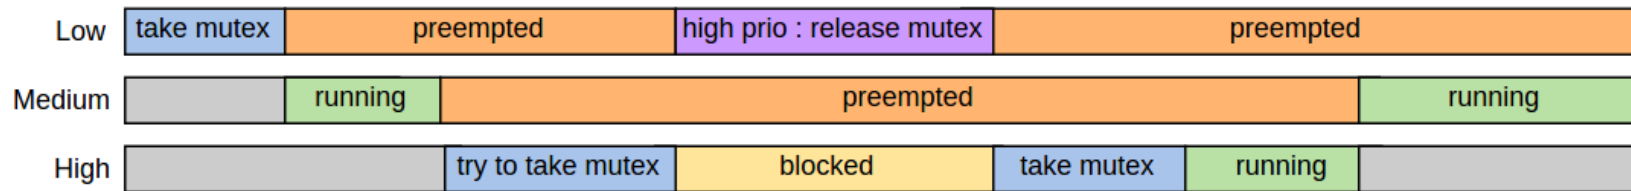


Figure 6 : Priority inheritance

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
- Sort the executions by running time
- Limitations
 - Threads need to have different priorities
 - Model is fixed
 - Not working with SCHED_DEADLINE
 - Manual analysis to find problems
 - Problems when more than one processor

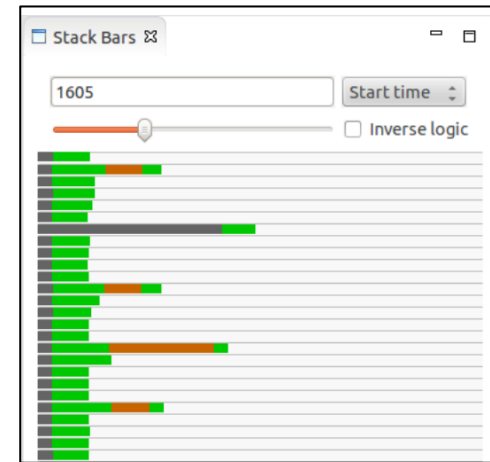


Figure 7 : Original stackbars view

Modeling

Advantage of using only kernel events

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

Modeling

- Identify executions automatically and then let the users choose between some valid models
 - Define a support ratio
 - Find all event types that are more frequent than the ratio
 - Increase the episode sizes using the fact that the sub-episodes must also be supported
 - Difficulties :
 - Using only event types
 - Execution time and memory usage
 - Many possible resulting models

Modeling : method

State machine

- User identifies :
 - an execution or
 - events that define the start and the end (name, parameters with operations, etc.)
 - TIDs for start and end
 - Presets for common cases

Events Selection

Define executions

Enter the starting tid(s) or process name (blank for current only, separate by coma, * for all)

3152,test_priority,3153

Enter the ending tid(s) or process name (blank for same, separate by coma, * for all)

*

Check to use the selected range as time limit

Load this preset

Enter start event name

sched_wakeup | sched_wakeup_new

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

tid=\$tid

Enter end event name or blank to use only the start event

sched_switch

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

prev_tid=\$tid,prev_state!=0,

Enter the deadline for this execution (blank for none)

50000

Depth options

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

Change current depth selection

OK Cancel

Figure 8 : 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

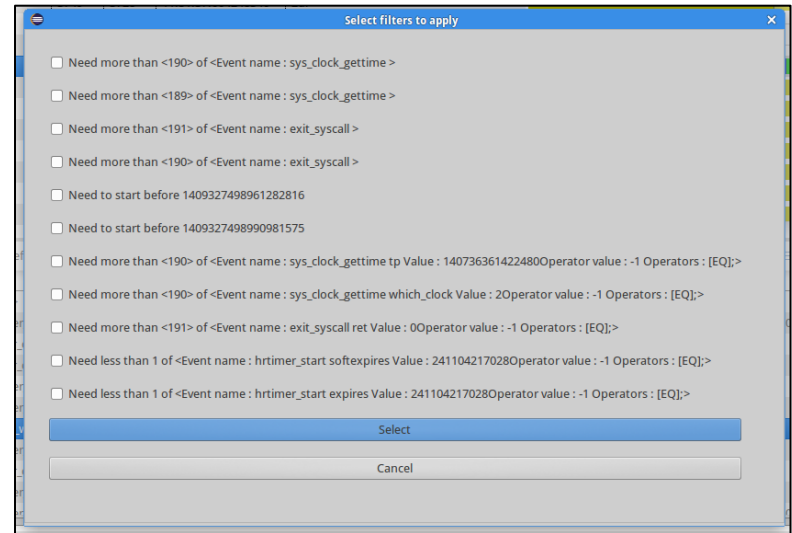




Figure 9 : Dialog to select modifications to apply

Overview

1) Control Flow View

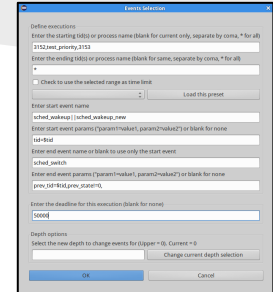


3) Stackbars View




5) Other views

2) Define executions



4) Critical Flow View with CP Complement view



Views

Stackbars view

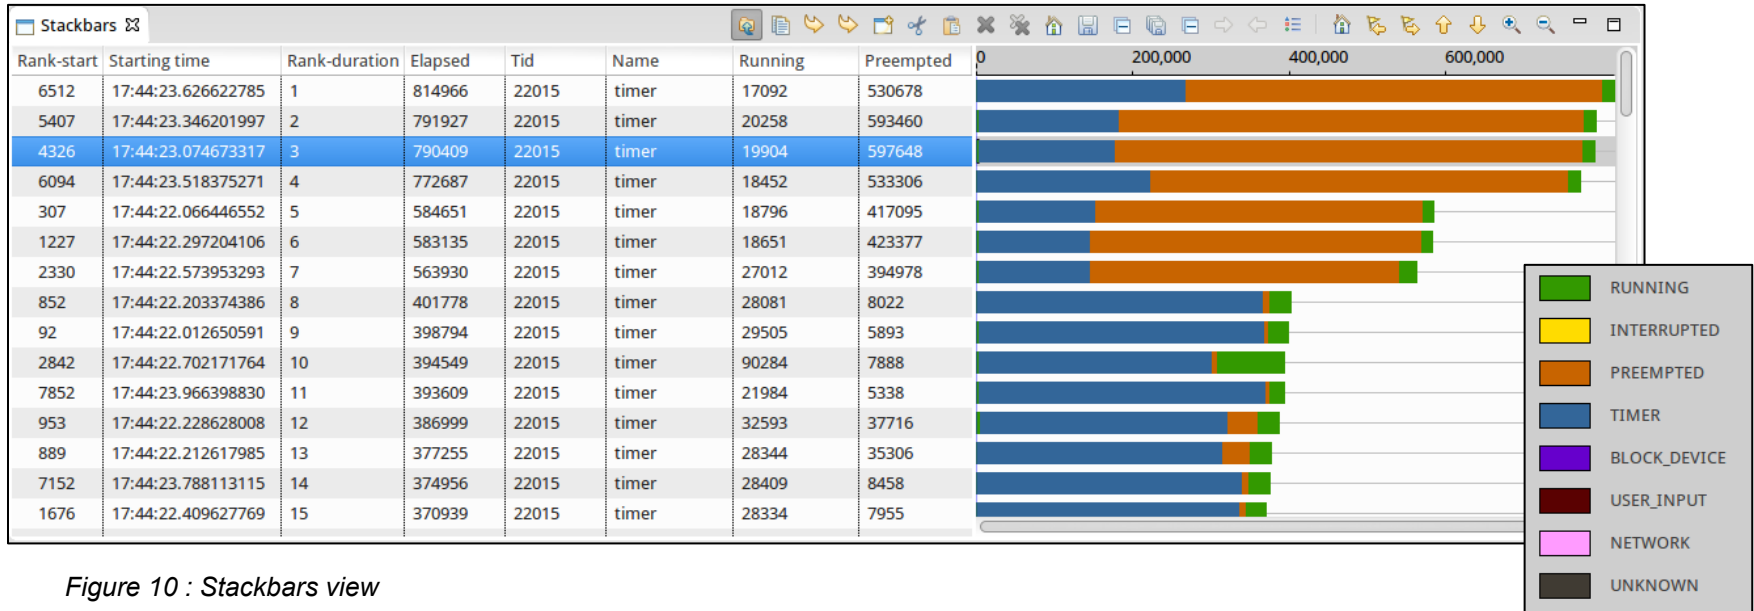


Figure 10 : Stackbars view

Views

- Supports
 - Thread pool
 - Nested executions

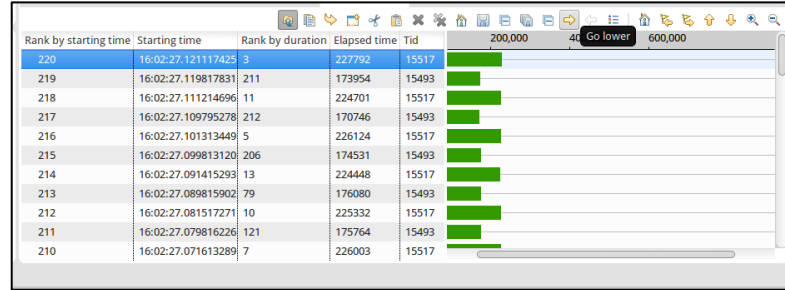


Figure 11 : Task on multiple threads




Figure 12 : Nested executions

Views

Time View

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

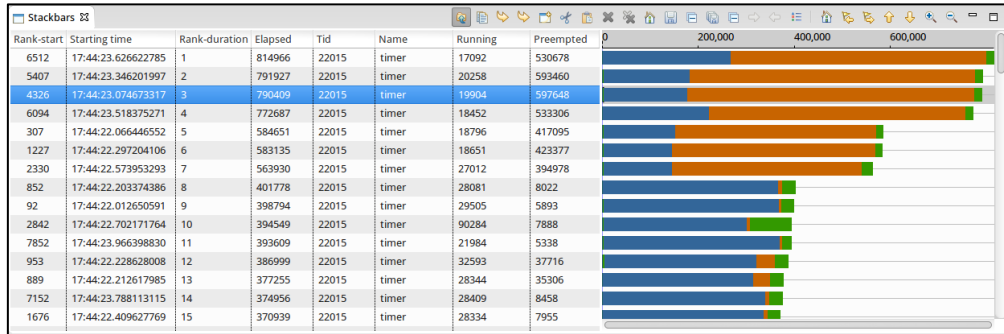


Figure 14 : Stackbars view

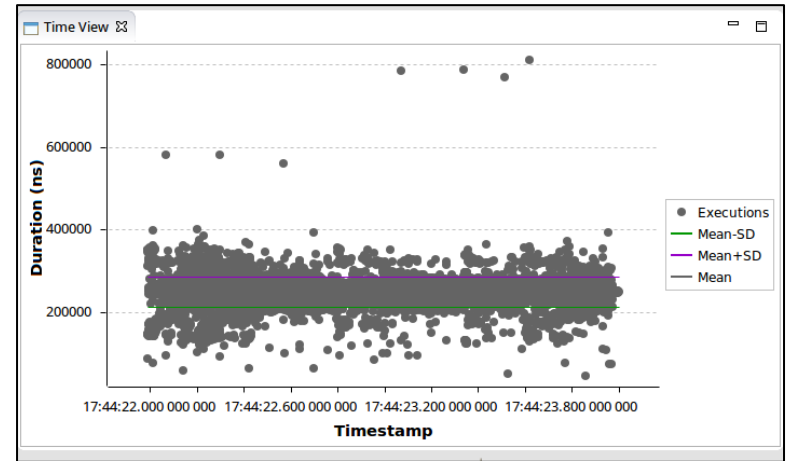


Figure 13 : Time view

Views

CP Complement View

- Show the priority of all running threads during preemption period of any thread in the critical path

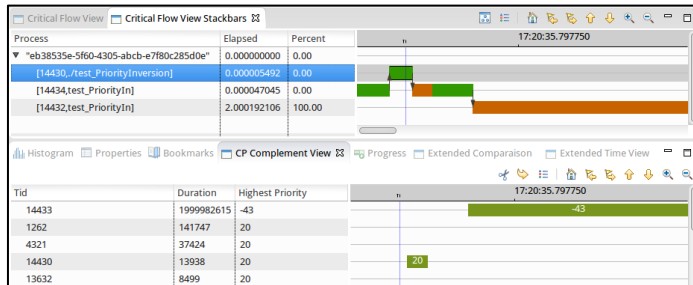


Figure 15 : CP Complement view

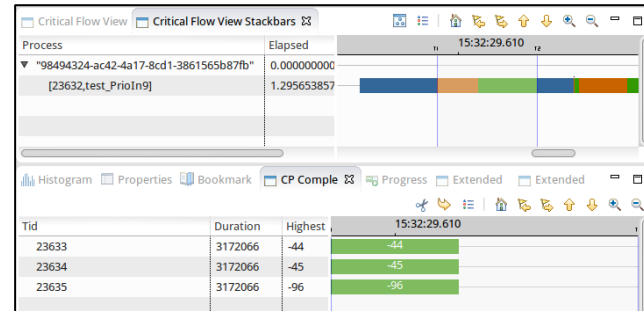


Figure 16 : CP Complement view

Views

CP Complement View

- Detect priority inversion

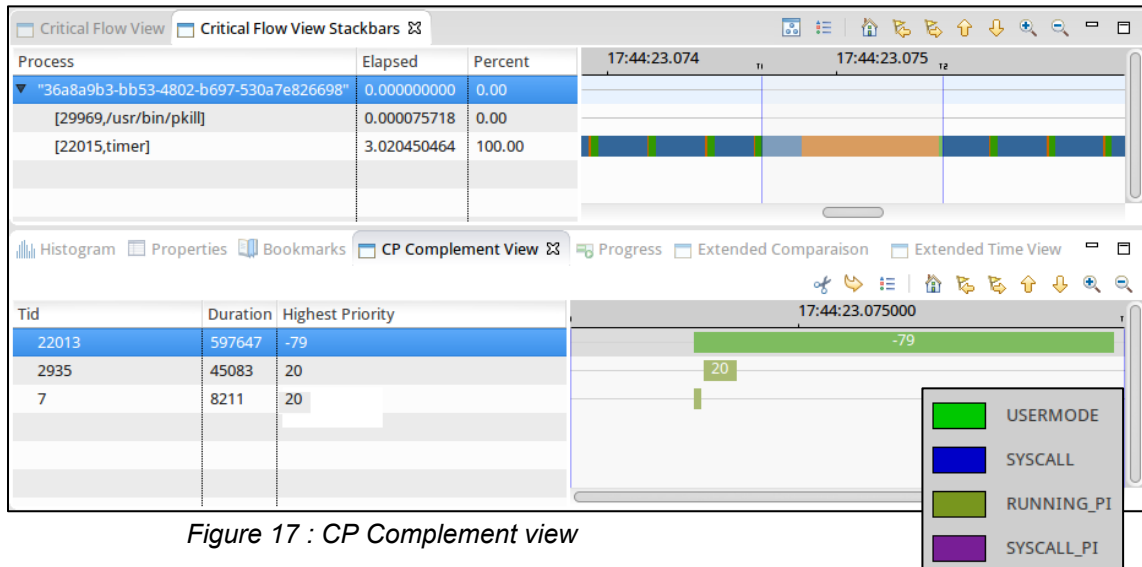


Figure 17 : CP Complement view

Both

Related only

Select tid (blank for current execution)

Select start time (blank for current execution)

Select end time (blank for current execution)

Select running cpus (separated by a comma, empty for all)

Example

Find out why some executions take more time




Figure 18 : Problematic executions




Figure 19 : Normal executions

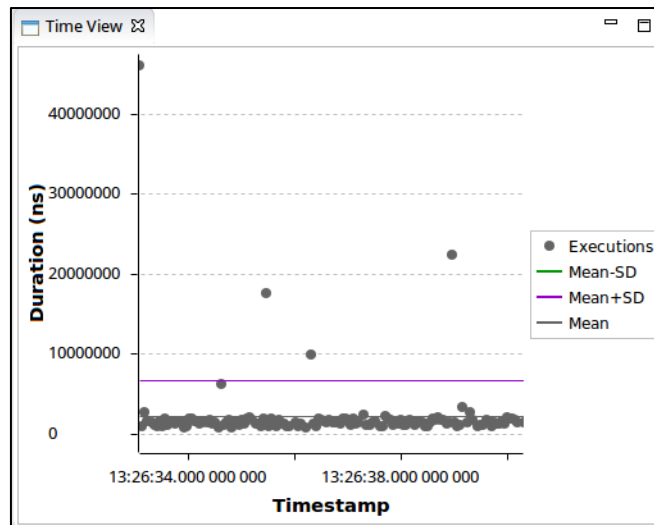


Figure 20 : Time View

Example

Normal execution


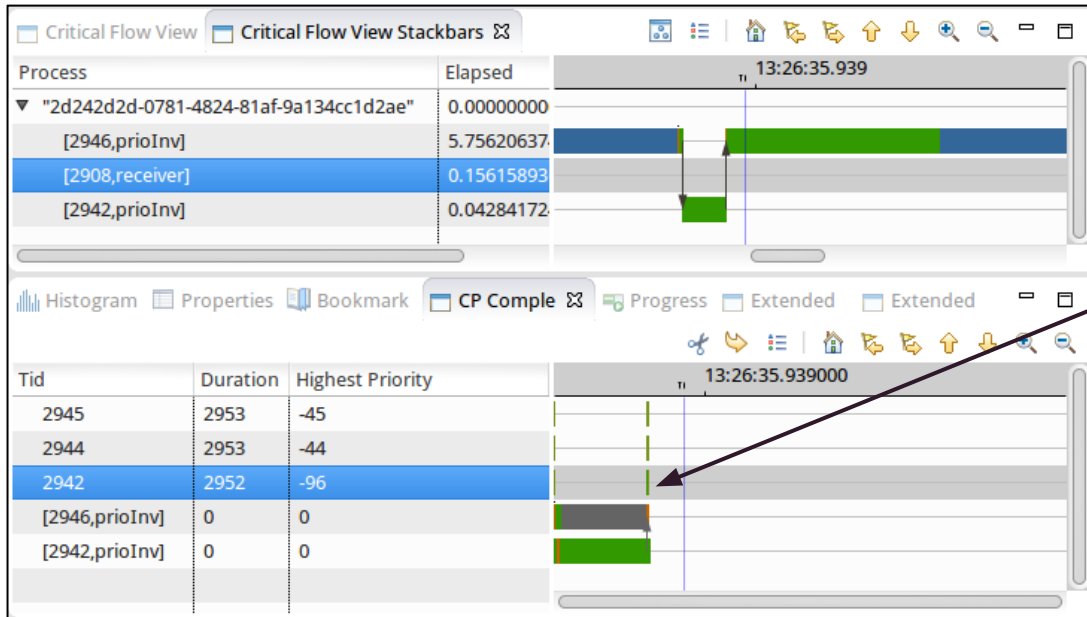


Figure 22 : There was priority inheritance

Figure 21 : CP Complement of a normal execution

Example

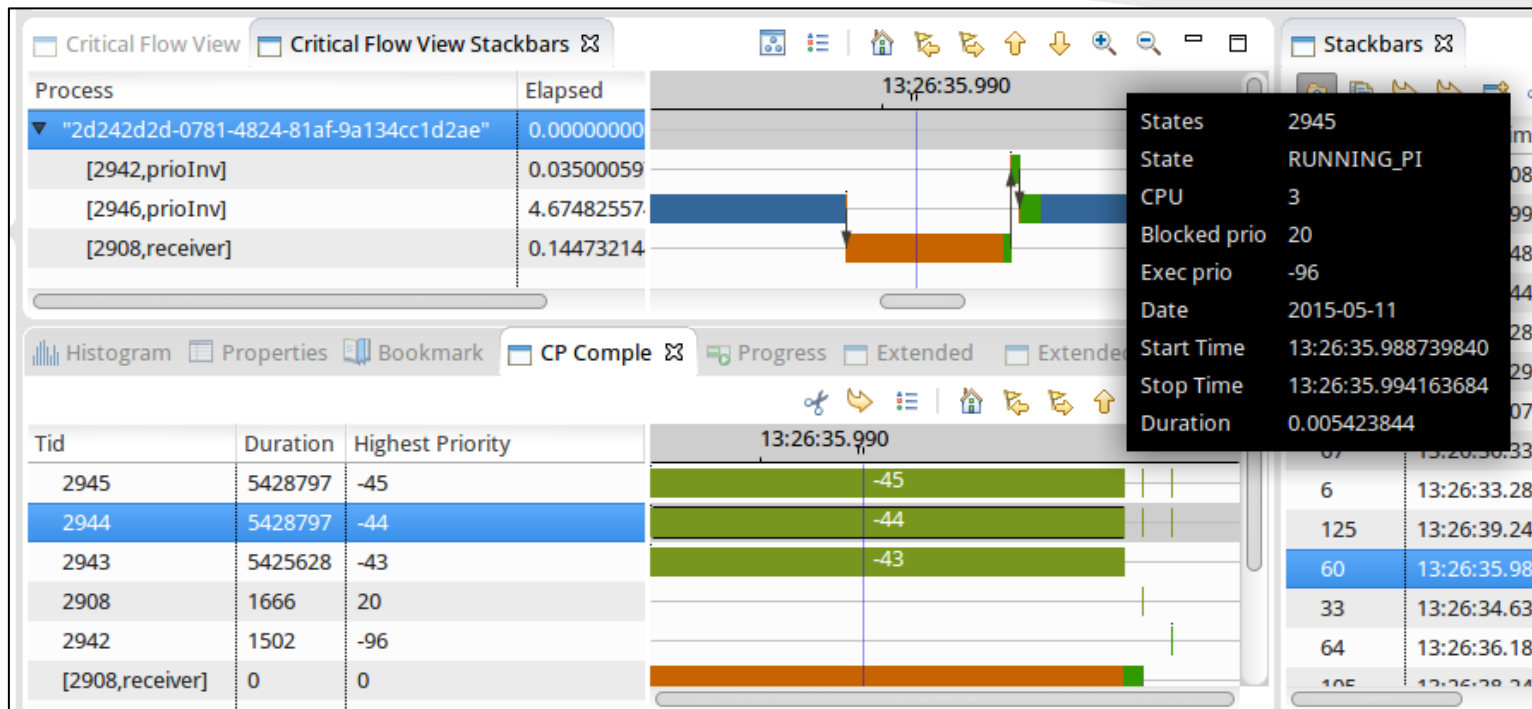


Figure 23 : CP Complement of a problematic execution

Other results

Output the dependencies during an execution

- priority
- directly related option

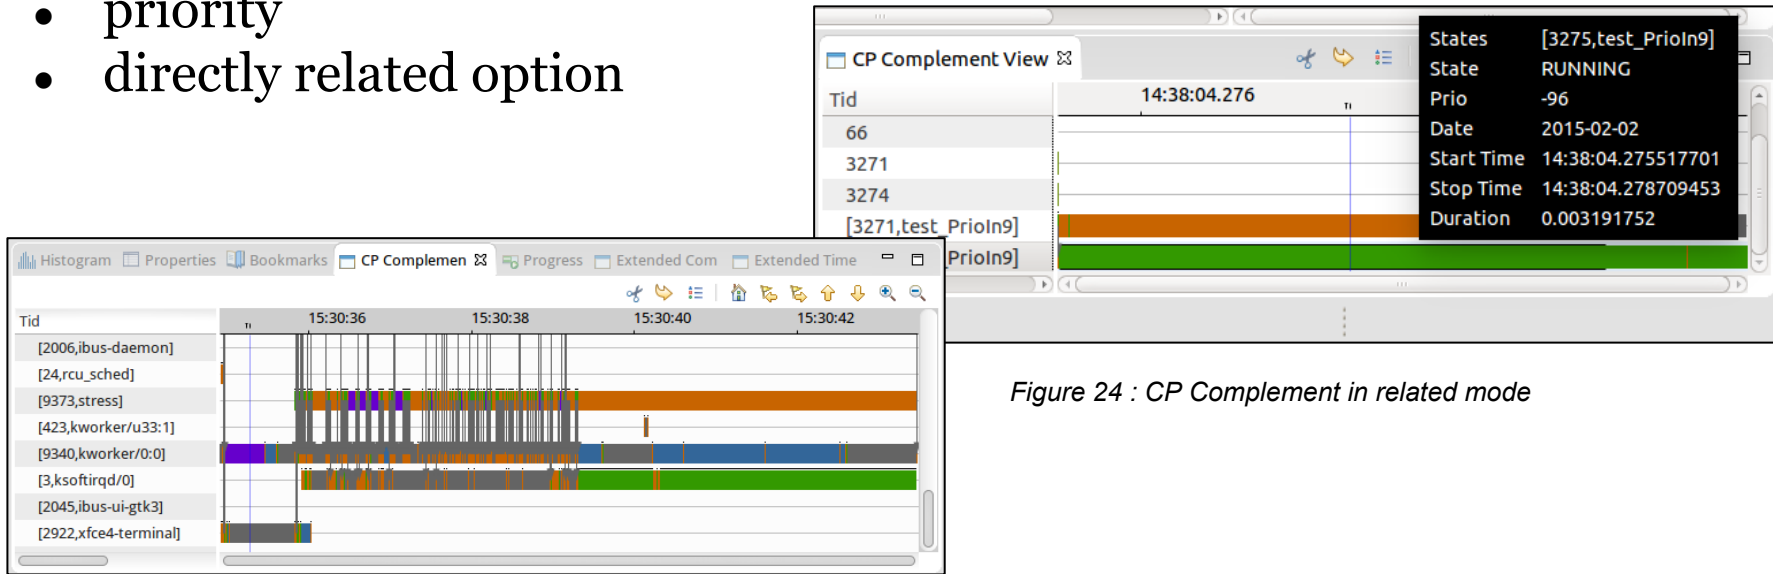


Figure 24 : CP Complement in related mode

Other results

Extended comparison view

Select types

- Futex
- Mq_send
- Mq_receive
- Wakeup
- Irq
- Preempted in userspace
- Preempted in other syscall
- Softirq
- HRTimer

Select parameters

Tid(s) (separated by a comma)

Start time

End time

Load the current time range

Cancel OK

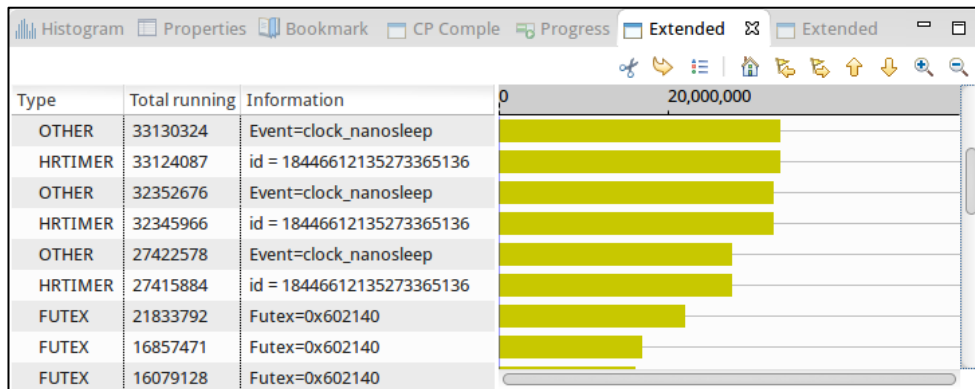

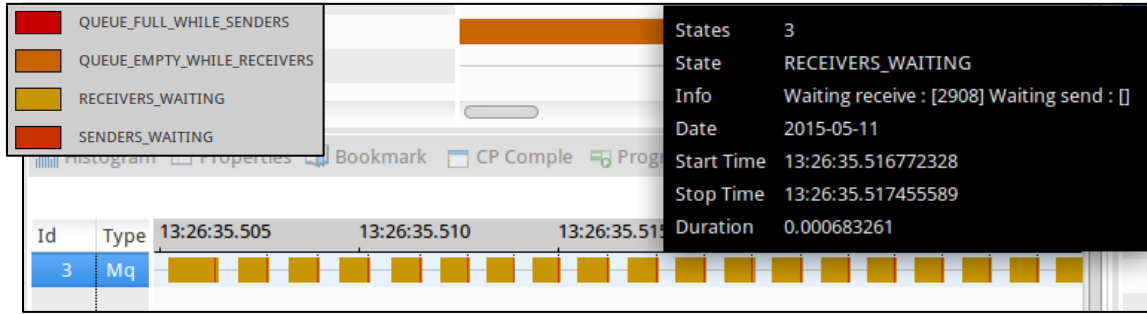


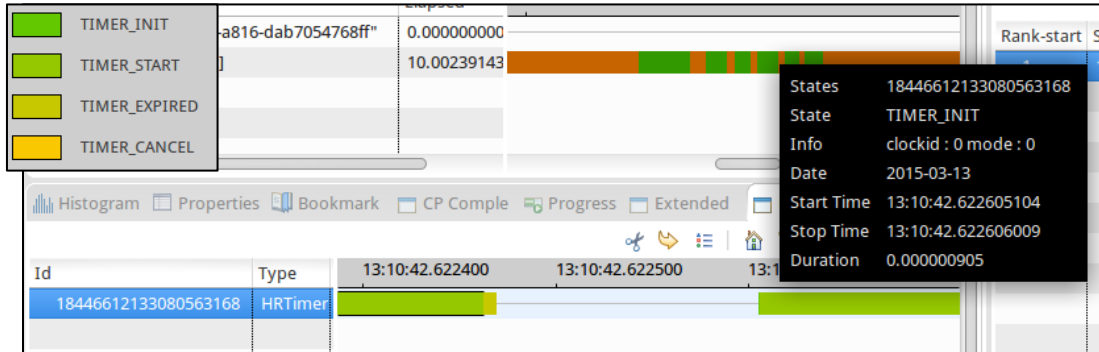
Figure 25 : Extended comparison view

Other results

Extended time view : queue



HRTimer



Select parameters

Queue(s) (separated by a comma)

Futex (separated by a comma)

Timers (separated by a comma)

Select time options

Start time (Blank for selected extended entry)

End time (Blank for selected extended entry)

Nb events back (Blank for default)

Figure 26 : Extended time view

Other results

- Deadline analysis
 - Tell which executions missed their deadlines
 - User input

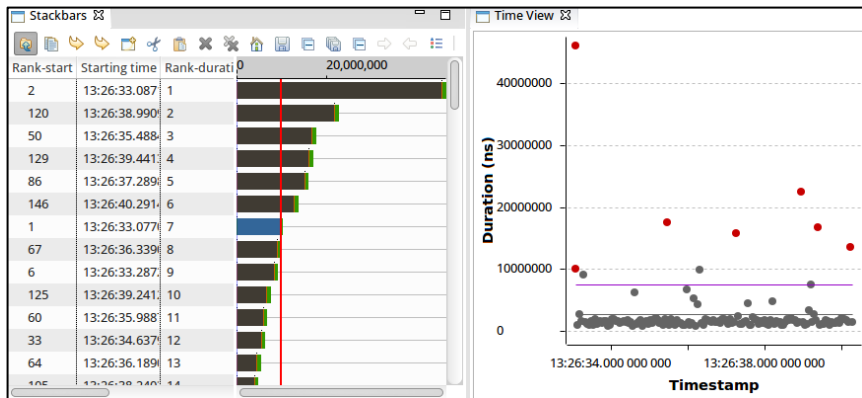


Figure 27 : Deadline

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?