



A theorem for the RT scheduling latency (and a measurement tool too)

Daniel Bristot de Oliveira, Daniel Casini, Rômulo Silva de Oliveira and Tommaso Cucinotta
Principal Software Engineer



Episode III – Showing the math

Daniel Bristot de Oliveira, Daniel Casini, Rômulo Silva de Oliveira and Tommaso Cucinotta
Principal Software Engineer



Real-Time Linux



"Real-Time" Linux


Real-Time Linux vs Real-Time theory

Experimental vs Analytical



```
root@realtime-01 ~]# cat -clictest.txt
[root@realtime-01 ~]# cyclicttest --smp -p 95 -m
/dev/cpu_dma_latency set to 0us
policy: fifo: loadavg: 14.90 6.21 3.98 2/387 2735923 1

T: 0 (2735898) P:95 I:1000 C: 66520 Min: 4 Act: 5 Avg: 5 Max: 15
T: 1 (2735899) P:95 I:1500 C: 44341 Min: 4 Act: 6 Avg: 5 Max: 20
T: 2 (2735900) P:95 I:2000 C: 33251 Min: 4 Act: 6 Avg: 5 Max: 15
T: 3 (2735901) P:95 I:2500 C: 26598 Min: 4 Act: 5 Avg: 5 Max: 15
T: 4 (2735902) P:95 I:3000 C: 22162 Min: 4 Act: 5 Avg: 5 Max: 16
T: 5 (2735903) P:95 I:3500 C: 18993 Min: 4 Act: 6 Avg: 5 Max: 13
T: 6 (2735904) P:95 I:4000 C: 16617 Min: 4 Act: 5 Avg: 5 Max: 14
T: 7 (2735905) P:95 I:4500 C: 14769 Min: 4 Act: 5 Avg: 5 Max: 12
T: 8 (2735906) P:95 I:5000 C: 13290 Min: 4 Act: 6 Avg: 5 Max: 14
T: 9 (2735907) P:95 I:5500 C: 12030 Min: 8 Act: 12 Avg: 13 Max: 24
T:10 (2735908) P:95 I:6000 C: 11072 Min: 4 Act: 5 Avg: 5 Max: 13
T:11 (2735909) P:95 I:6500 C: 10219 Min: 5 Act: 6 Avg: 5 Max: 20
T:12 (2735910) P:95 I:7000 C: 9488 Min: 5 Act: 6 Avg: 5 Max: 13
T:13 (2735911) P:95 I:7500 C: 8854 Min: 5 Act: 5 Avg: 5 Max: 14
T:14 (2735912) P:95 I:8000 C: 8200 Min: 4 Act: 6 Avg: 5 Max: 14
T:15 (2735913) P:95 I:8500 C: 7801 Min: 5 Act: 9 Avg: 5 Max: 17
T:16 (2735914) P:95 I:9000 C: 7370 Min: 4 Act: 6 Avg: 5 Max: 13
T:17 (2735915) P:95 I:9500 C: 6987 Min: 5 Act: 6 Avg: 6 Max: 20
T:18 (2735916) P:95 I:10000 C: 6638 Min: 5 Act: 9 Avg: 6 Max: 25
```



Real-Time Linux vs Real-Time theory

Linux approach



```
root@realtime-01 ~]# cat -cyclictest.txt
root@realtime-01 ~]# cyclictest --smp -p 95 -m
/dev/cpu_dma_latency set to 0us
policy: fifo loadavg: 14.90 6.21 3.98 2/387 2735923 1


T: 0 (2735898) P:95 I:1000 C: 66520 Min: 4 Act: 5 Avg: 5 Max: 15
T: 1 (2735899) P:95 I:1500 C: 44341 Min: 4 Act: 6 Avg: 5 Max: 20
T: 2 (2735900) P:95 I:2000 C: 33251 Min: 4 Act: 6 Avg: 5 Max: 15
T: 3 (2735901) P:95 I:2500 C: 26598 Min: 4 Act: 5 Avg: 5 Max: 15
T: 4 (2735902) P:95 I:3000 C: 22162 Min: 4 Act: 5 Avg: 5 Max: 16
T: 5 (2735903) P:95 I:3500 C: 18993 Min: 4 Act: 6 Avg: 5 Max: 13
T: 6 (2735904) P:95 I:4000 C: 16617 Min: 4 Act: 5 Avg: 5 Max: 14
T: 7 (2735905) P:95 I:4500 C: 14769 Min: 4 Act: 5 Avg: 5 Max: 12
T: 8 (2735906) P:95 I:5000 C: 13290 Min: 4 Act: 6 Avg: 5 Max: 14
T: 9 (2735907) P:95 I:5500 C: 12030 Min: 8 Act: 12 Avg: 13 Max: 24
T:10 (2735908) P:95 I:6000 C: 11072 Min: 4 Act: 5 Avg: 5 Max: 14
T:11 (2735909) P:95 I:6500 C: 10219 Min: 5 Act: 6 Avg: 5 Max: 13
T:12 (2735910) P:95 I:7000 C: 9488 Min: 5 Act: 6 Avg: 5 Max: 20
T:13 (2735911) P:95 I:7500 C: 8854 Min: 5 Act: 5 Avg: 5 Max: 13
T:14 (2735912) P:95 I:8000 C: 8200 Min: 4 Act: 6 Avg: 5 Max: 14
T:15 (2735913) P:95 I:8500 C: 7801 Min: 5 Act: 9 Avg: 5 Max: 17
T:16 (2735914) P:95 I:9000 C: 7370 Min: 4 Act: 6 Avg: 5 Max: 13
T:17 (2735915) P:95 I:9500 C: 6987 Min: 5 Act: 6 Avg: 6 Max: 20
T:18 (2735916) P:95 I:10000 C: 6638 Min: 5 Act: 9 Avg: 6 Max: 20
```



- Linux was adapted to become a RTOS
- PREEMPT_RT: *De facto* standard
- Evaluated (mainly) with cyclictest
- Cyclictest:
 - Practical: lightweight and out-of-the-box
 - It is a “closed-box” test
 - No demonstration
 - Does not provide evidence of “root-cause”

Real-Time Linux vs Real-Time theory

Real-time analysis





- Based on the timing description of the system
- Capture all behaviors
- Precisely define the worst cases
- But depends on a precise definition of the system
- Often overly-simplified

```
yc: cat cyclictest.txt
# cyclictest --smp -p 95 -m
cy set to 0us
vg: 14.90 6.21 3.98 2/387 2735923

I:1000 C: 66520 Min: 4 Act: 5
I:1500 C: 44341 Min: 4 Act: 6
I:2000 C: 33251 Min: 4 Act: 6
I:2500 C: 26598 Min: 4 Act: 5
I:3000 C: 22162 Min: 4 Act: 5
I:3500 C: 18993 Min: 4 Act: 6
I:4000 C: 16617 Min: 4 Act: 5
I:4500 C: 14769 Min: 4 Act: 6
I:5000 C: 13290 Min: 4 Act: 5
I:5500 C: 12080 Min: 4 Act: 6
I:6000 C: 11072 Min: 8 Act: 12
I:6500 C: 10219 Min: 4 Act: 5
I:7000 C: 9488 Min: 5 Act: 6
I:7500 C: 8854 Min: 5 Act: 6
I:8000 C: 8290 Min: 5 Act: 5
I:8500 C: 7801 Min: 4 Act: 6
I:9000 C: 7370 Min: 5 Act: 9
I:9500 C: 6987 Min: 4 Act: 6
I:10000 C: 6638 Min: 5 Act: 9
```

But, I like both.





Demystifying the Real-Time Linux Scheduling Latency

Daniel Bristot de Oliveira 
Red Hat, Italy
bristot@redhat.com

Daniel Casini 
Scuola Superiore Sant'Anna, Italy
daniel.casini@santannapisa.it

Rômulo Silva de Oliveira 
Universidade Federal de Santa Catarina, Brazil
romulo.deoliveira@ufsc.br

Tommaso Cucinotta 
Scuola Superiore Sant'Anna, Italy
tommaso.cucinotta@santannapisa.it

Abstract

Linux has become a viable operating system for many real-time workloads. However, the black-box approach adopted by `cyclicttest`, the tool used to evaluate the main real-time metric of the kernel, the scheduling latency, along with the absence of a theoretically-sound description of the in-kernel behavior, sheds some doubts about Linux meriting the real-time adjective. Aiming at clarifying the PREEMPT_RT Linux scheduling latency, this paper leverages the *Thread Synchronization Model* of Linux to derive a set of properties and rules defining the Linux kernel behavior from a scheduling perspective. These rules are then leveraged to derive a sound bound to the scheduling latency, considering all the sources of delays occurring in all possible sequences of synchronization events in the kernel. This paper also presents a tracing method, efficient in time and memory overheads, to observe the kernel events needed to define the variables used in the analysis. This results in an easy-to-use tool for deriving reliable scheduling latency bounds that can be used in practice. Finally, an experimental analysis compares the `cyclicttest` and the proposed tool, showing that the proposed method can find sound bounds faster with acceptable overheads.

2012 ACM Subject Classification Computer systems organization → Real-time operating systems

Keywords and phrases Real-time operating systems, Linux kernel, PREEMPT_RT, Scheduling latency


Digital Object Identifier 10.4230/LIPICs.ECRTS.2020.9

Supplementary Material ECRTS 2020 Artifact Evaluation approved artifact available at <https://doi.org/10.4230/DARTS.6.1.3>.
Supplement material and the code of the proposed tool is available at: <https://bristot.me/demystifying-the-real-time-linux-latency/>

Funding This work has been partially supported by CAPES, The Brazilian Agency for Higher Education, project PrInt CAPES-UFSC "Automation 4.0."

Acknowledgements The authors would like to thank Thomas Gleixner, Peter Zijlstra, Steven Rostedt, Arnaldo Carvalho De Melo and Clark Williams for the fruitful discussions about the model, analysis, and tool.

© Daniel Bristot de Oliveira, Daniel Casini, Rômulo Silva de Oliveira, and Tommaso Cucinotta;
licensed under Creative Commons License CC-BY
32nd Euromicro Conference on Real-Time Systems (ECRTS 2020).
Editor: Marcus Völz; Article No. 9; pp. 9:1–9:23
Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany




Episode I: getting formal

Math side: Talk is cheap...




Show me the math!

Generators of events



21

Specifications: Sufficiency conditions



redhat

Demystifying the Real-Time Linux Scheduling Latency

Approach

Formal specification



Scheduling latency bound

provides an overall bound that is valid for all the possible events sequences.

► **Lemma 7.**

$$L^{IF} \leq \max(D_{ST}, D_{FOID}) + D_{FAIR} + D_{FID}$$

Proof. The lemma follows by noting that cases (3-a), (3-b), (3-c), (3-d), (3-e) are mutually exclusive and cover all the possible sequences of events from the occurrence of `set_need_resched`, to the time instant in which `rt_mutex_unlock` is allowed to execute (as required by Definition 1), and the right-hand side of Equation 4 simultaneously upper bounds the right-hand sides of Equations 2, 3, 4, and 5.

Theorem 8 summarizes the results derived in this section.


► **Theorem 8.** The scheduling latency experienced by an arbitrary thread τ_i^{RT} is bounded by the least positive value that fulfills the following recursion equation:

$$L = \max(D_{ST}, D_{FOID}) + D_{FAIR} + D_{FID} + L^{IF}(\delta) + L^{IF}(\delta)$$

Proof. The theorem follows directly from Lemma 7 and Equation 1.


² Note that, internally to the IRQ handler, the preemption state may be changed (e.g., by calling `set_need_resched`).

Measurement and analysis



From formal specification to synchronization rules

Formally backed natural language arguments



- Generators
 - Basic/Independent behavior
 - e.g., irq_disable/enable, scheduler call
- Translated into a set of operations
- Specifications
 - Relations among generators
 - e.g., necessary conditions to call the scheduler
- Translated into a set of synchronization rules

Scheduling latency definition

The **scheduling latency** experienced by an arbitrary thread τ is:

- the **longest time** elapsed **between** the *time A* in which any job of τ becomes **ready and with the highest priority**,
- and the *time F* in which the scheduler returns and allows τ to **execute its code**.

**ready AND with the
highest priority:**

It covers the case in which
these **two actions are
not a single event.**

It is **scheduler
independent.**

**There is only one
highest priority thread**
on a CPU: it is **the one
selected to run** by the
scheduler.

Interference and blocking

```
void __sched notrace __schedule(bool  
struct task_struct *prev, *next;  
unsigned long *switch_count;  
struct rq_flags rf;  
struct rq *rq;  
int cpu;  
  
cpu = smp_processor_id();  
rq = cpu_rq(cpu);  
prev = rq->curr;  
  
schedule_debug(prev, preempt);  
  
if (sched_feat(HRTICK))  
    hrtick_clear(rq);  
  
local_irq_disable();  
rcu_note_context_switch(preempt);
```

The **scheduling latency** is caused by:


- **Blocking** from the current (and so lower) priority thread;
 - Including scheduling.
- **Interference** from IRQs and NMI.

This are well established terms in the real-time scheduling literature:

Interference from higher priority, blocking from lower priority.

Blocking bound

From the specification that bounds the block to a timeline



Scheduling latency: start

- The **longest time** elapsed **between** the *time A* in which any job of τ becomes **ready and with the highest priority**:
 - Generalized to the **need_resched** event
 - Works for all schedulers
 - `cyclictest` does not work for DL with `NR_TASKS > CPUs`.
 - Works for all conditions
 - E.g., a throttled DL task after a replenishment will cause a need resched without a wakeup.
 - Has **preempt and IRQ disabled** as **necessary conditions**
 - So we use the occurrence of the first **necessary condition** as the starting point of the critical window.
 - E.g., when preemption was disabled for the first time.

The wakeup is the only event that causes a need resched, and that is why it was not used here.

But ready means that the task was awakened.

Scheduling latency: end

- And the *time* F in which the scheduler returns and allows τ to **execute its code**.
 - Generalized to the **preempt_enable** after **__schedule()**
 - Implies that the system **crossed the context switch** code path.
 - Context switch implies **__schedule()**
 - **Context switch needs:**
 - **Preempt disable to schedule as necessary condition**
 - **irqs disabled by thread as necessary condition**


We are looking for a safe-bound, and so we have to put pessimism values.

We can latter reduce the pessimism, but with safe arguments.

How do we bound that?


Need resched -> ctxsw

All possible cases



Timeline and cases


Variables in the the timeline



Blocking variables


- **DPOID**: preemption or interrupts disabled to postpone the scheduler;
- **DPAIE**: preemption and interrupts enabled, as a transient state from **poid** to **psd**; when scheduling a new highest priority thread.
- **DPSD**: preemption disable to schedule;
- **DST**: delay caused by the scheduling tail; the “non return” point in which a new arrived task will have to wait for the current scheduling operation to finish before scheduling.

In the model, the preemption control is specialized into two different operations: to *postpone the scheduler* (the most known behavior) or to *protect the execution of the `__schedule()` function* from recursion.




Timeline and cases

Variables in the the timeline



And the *scheduling latency* bounds to:




$$\mathbf{L = \max(D_{ST}, D_{POID}) + D_{PAIE} + D_{PSD} + I^{NMI}(L) + I^{IRQ}(L)}$$

The bound considers all possible cases. Note that the Latency L is present in both sides of the equation.

So, L is bounded by the least positive value fulfilling the equation (like on RTA).

And the *scheduling latency* bounds to:




$$L = \max(D_{ST}, D_{POID}) + D_{PAIE} + D_{PSD} + I^{NMI}(L) + I^{IRQ}(L)$$

The bound considers all possible cases. Note that the Latency L is present in both sides of the equation.

So, L is bounded by the least positive value fulfilling the equation (like on RTA).

Interrupts are workload dependent

- Instead of proposing “the best” interrupt characterization, the rtsl reports the scheduling latency based on some well-known characterizations:
 - No interrupt
 - Worst single interrupt
 - Single occurrence of all interrupts
 - Sporadic
 - Sliding window (Author’s preferred)
 - Sliding window with oWCET



This topic was heavily discussed at the Real-time Micro Conference (inside Linux Plumbers) in 2019, more info here:



Episode II: getting practical (and efficient)

A practical scheduling latency estimation tool

Method and challenges



```


root@realtime-01 ~]# cat -clictest.txt
root@realtime-01 ~]# cyclicttest --smp -p 95 -m
# /dev/cpu_dma_latency set to 0us
policy: fifo: loadavg: 14.90 6.21 3.98 2/387 2735923      1

T: 0 (2735898) P:95 I:1000 C: 66520 Min: 4 Act: 5 Avg: 5 Max: 15
T: 1 (2735899) P:95 I:1500 C: 44341 Min: 4 Act: 6 Avg: 5 Max: 20
T: 2 (2735900) P:95 I:2000 C: 33251 Min: 4 Act: 6 Avg: 5 Max: 15
T: 3 (2735901) P:95 I:2500 C: 26598 Min: 4 Act: 5 Avg: 5 Max: 15
T: 4 (2735902) P:95 I:3000 C: 22162 Min: 4 Act: 5 Avg: 5 Max: 16
T: 5 (2735903) P:95 I:3500 C: 18993 Min: 4 Act: 6 Avg: 5 Max: 13
T: 6 (2735904) P:95 I:4000 C: 16617 Min: 4 Act: 5 Avg: 5 Max: 14
T: 7 (2735905) P:95 I:4500 C: 14769 Min: 4 Act: 5 Avg: 5 Max: 12
T: 8 (2735906) P:95 I:5000 C: 13290 Min: 4 Act: 6 Avg: 5 Max: 14
T: 9 (2735907) P:95 I:5500 C: 12030 Min: 8 Act: 12 Avg: 13 Max: 24
T:10 (2735908) P:95 I:6000 C: 11072 Min: 4 Act: 5 Avg: 5 Max: 14
T:11 (2735909) P:95 I:6500 C: 10219 Min: 5 Act: 6 Avg: 5 Max: 13
T:12 (2735910) P:95 I:7000 C: 9488 Min: 5 Act: 6 Avg: 5 Max: 20
T:13 (2735911) P:95 I:7500 C: 8854 Min: 5 Act: 5 Avg: 5 Max: 13
T:14 (2735912) P:95 I:8000 C: 8200 Min: 4 Act: 6 Avg: 5 Max: 14
T:15 (2735913) P:95 I:8500 C: 7801 Min: 5 Act: 9 Avg: 5 Max: 17
T:16 (2735914) P:95 I:9000 C: 7370 Min: 4 Act: 6 Avg: 5 Max: 13
T:17 (2735915) P:95 I:9500 C: 6987 Min: 5 Act: 6 Avg: 6 Max: 20
T:18 (2735916) P:95 I:10000 C: 6638 Min: 5 Act: 9 Avg: 6 Max: 25
  
```



- Based on the latency bound
- The latency bound is based on the model
- The *model* is based on tracing of events
 - but high frequency events
 - hundreds MB/sec/CPU
- Challenges:
 - To minimize the (runtime) overhead
 - Work out-of-the-box

Rtsl: a measurement tool



Kernel space:

- Rtsl events

User space:


- Rtsl command
 - Python

Has three commands:

- The **record** command saves the trace data;
- The **report** command process the trace and does the analysis.
- The **stats** command produces a histogram of the thread variables

rtsl events


Low overhead tracer



- Hooks to events
 - Filters the high frequency trace
 - Doing in-kernel processing
 - Use a knob on debugfs to enable the tracing
- For blocking variables:
 - Reports all values or only the discover of new max values
- For IRQ and NMI:
 - Reports one event for each occurrence
- Discounts the interference:
 - e.g., IRQ interference on a **poid**

Kernel changes

- The rtsl events depends on:
- **preemptirq tracepoints**
 - So it needs a "debug/trace" kernel (yeah...)
 - But life finds a way
- Annotations on the preempt_disable to sched
 - No functional changes
- NMI tracepoints
 - Or change in the current one to the extreme points of the handler




The parser was developed as a kernel module. In this way I can leave it off tree... but it would be better to have it in.

If we get it in, we can change the debugfs for the tracefs.

rtsl record


Trace recording



- Captures the values for the variables
 - Only new max values for thread variables
 - Saving them into a trace file
- Calls real tracers to do the tracing:
 - Perf
 - Ftrace
- Controls the trace section
- Saves the trace in the rtsl_data/ dir

rtsl report


Trace processing



- Analyzes the trace!
 - All in user-space
- Most of the tool is done in python
 - Easy to extend the analysis (researchers like)
- Parses the trace file in parallel
 - Per cpu trace parsing (e.g., `perf script-c $...`)
 - Generates per-cpu database with the data
 - In the `rtsl_data/` dir
 - Uses a C trace-plugin create the database
- Database in a sqlite3 file

rtsl report

Data processing



- The analysis is done on the database
 - IRQ analysis needs to read data back and forth
 - Trace can reach tens of GB/per-cpu
- The analysis is done in parallel
- Two outputs:
 - Textual output
 - Charts
 - Using matplotlib

rtsl report output


Textual output

```
Interference Free Latency:
  paie is lower than 1 us -> neglectable
  latency = max(poid,    dst) + paie +   psd
  42212 = max(22510, 19312) +    0 + 19702
Cyclicttest:
  Latency =      27000 with Cyclicttest
No Interrupts:
  Latency =      42212 with No Interrupts
Sporadic:
  INT:    oWCET          oMIAT
  NMI:      0              0
  33:    16914          257130
  35:    12913          1843 <- oWCET > oMIAT
  236:   20728          1558 <- oWCET > oMIAT
  246:    3299          1910321
Did not converge.
```

```
continuing....
Sliding window:
  Window: 42212
          NMI:           0
          33:    16914
          35:    14588
          236:   20728
          246:    3299
  Window: 97741
          236:   21029 <- new!
  Window: 98042
Converged!
Latency =      98042 with Sliding Window
```


rtsl report output

Chart output



rtsl stats


Online view



- Monitor the thread variables
 - poid/psd/dst/paie...
- Uses BCC
 - Saves histograms in kernel
 - Display in user-space
 - Can plot data

Experiments


- Scheduling latency measurements on two systems:
 - workstation: eighth CPUs
 - server: twelve CPUs server
- Experiments:
 - Single-core
 - Different duration
 - Different workload
 - Multi-core
- Running in parallel with cyclicttest
- Note: The goal of the experiments is to demonstrate the tool, not to define worst values.




The experiments passed
by the artifact evaluation!



Single-core experiments




Multicore experiments



Remarks

- The PREEMPT_RT preemption model is deterministic, and the scheduling latency is bounded.
- The approach presented in the paper opens the door for a new set of real-time analysis for Linux;
 - The analytical interpretation of Linux thread model developed in this paper untight the Linux complexity, **enabling the reasoning at a more sophisticated level.**
- Even though rtsl finds higher scheduling latency values, they are still low enough to justify Linux as RTOS on the current scenarios.
- rtsl is practical, and resolves many problems of cyclicttest.
 - E.g., it can be used to point to the root causes of the latency;
 - But still can, and should, be improved:
 - Both with code, and other analysis.




For more information about this paper, like source code, other comments, Q&A, check its companion page!



rtsl vs cyclictest? nah

- They help the same people
 - But they do different things
- rtsl is a more specific tool
 - Covers a single aspect: sched latency
 - Covers all cases at synchronization level
 - In the worst condition, even those that happened at different points in time.
 - With strong arguments
 - Depends on kernel features (PREEMPT_RT/preemptirq...)
- cyclictest is a more generic tool
 - Covers many aspects: external activation of the timer
 - Hardware delays? Hardware bugs?
 - Without analysis - a closed-box test
 - Run on the potato that runs Linux
- rtsl adds only 4-ish us of overhead on cyclictest



For more information about this paper, like source code, other comments, Q&A, check its companion page!



Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 linkedin.com/company/red-hat

 youtube.com/user/RedHatVideos

 facebook.com/redhatinc

 twitter.com/RedHat

In the next episode....