

PyTimechart practical

Pierre Tardy
Software Engineer - UMG
October 2011

Intel Employee since 2009

Working on Intel's phone platforms

- Meego
- Android
- Power Management
- Tools (pytimechart, buildbot)
- Open-Source expertise

Formerly Freescale Employee

DVBH, LTE (gpe, oe, poky, cairo)

I will not rephrase the documentation, you can access it following this link:

<http://packages.python.org/pytimechart/index.html>

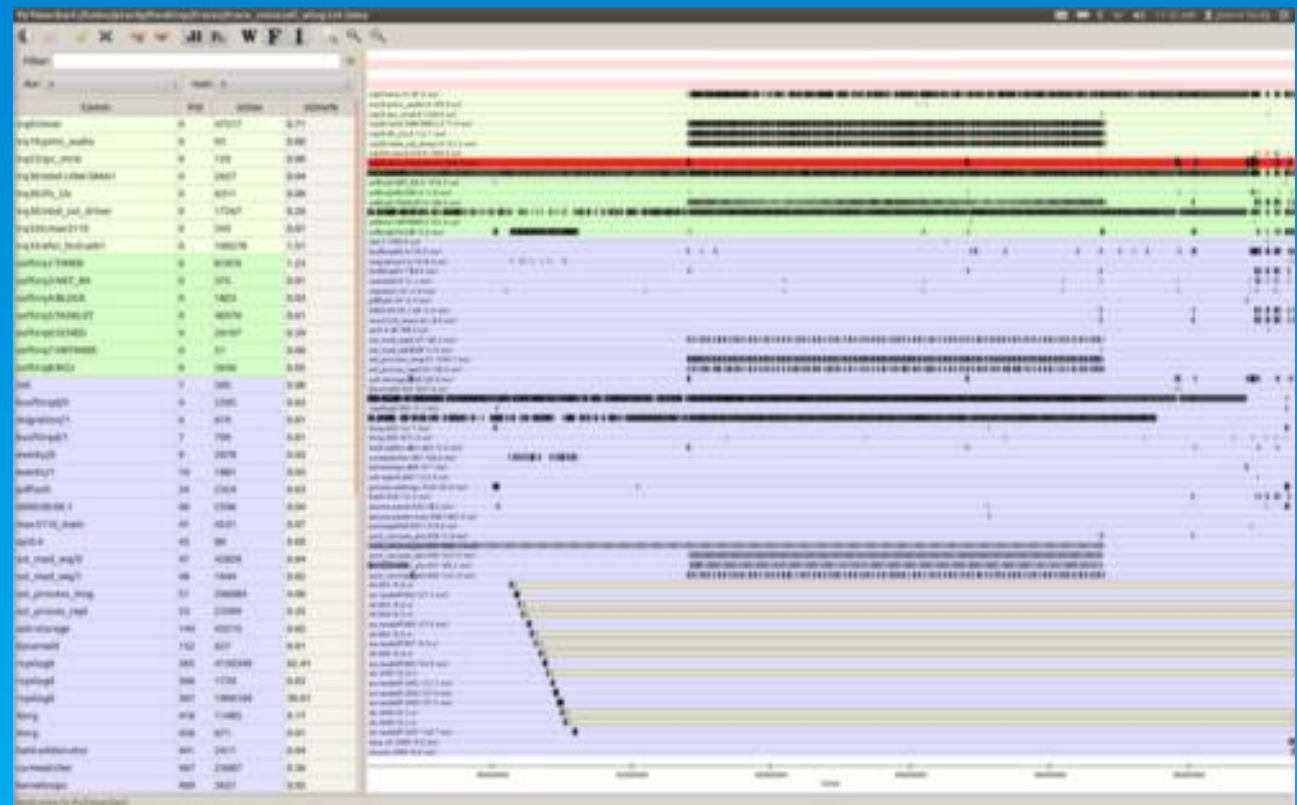
Goal is to present the tool using real-world traces. See what you can get, what you won't get.

1. PyTimechart Overview
2. Audio Player usecase – Hunting the wakes
3. Bootcharting Ubuntu – Filtering with pytimechart
4. Modem driver traces – Hacking PyTimechart to decode your own traces



Questions?

PyTimechart Overview



Ftrace/perf are linux kernel tracing system.

- Very low overhead
- Already a lot of tracepoints available
- Can trace all function calls without kernel compile
- Hard to dive into trace
- Most people need visual representation or get lost

PyTimechart aims to let you quickly find what you are looking for

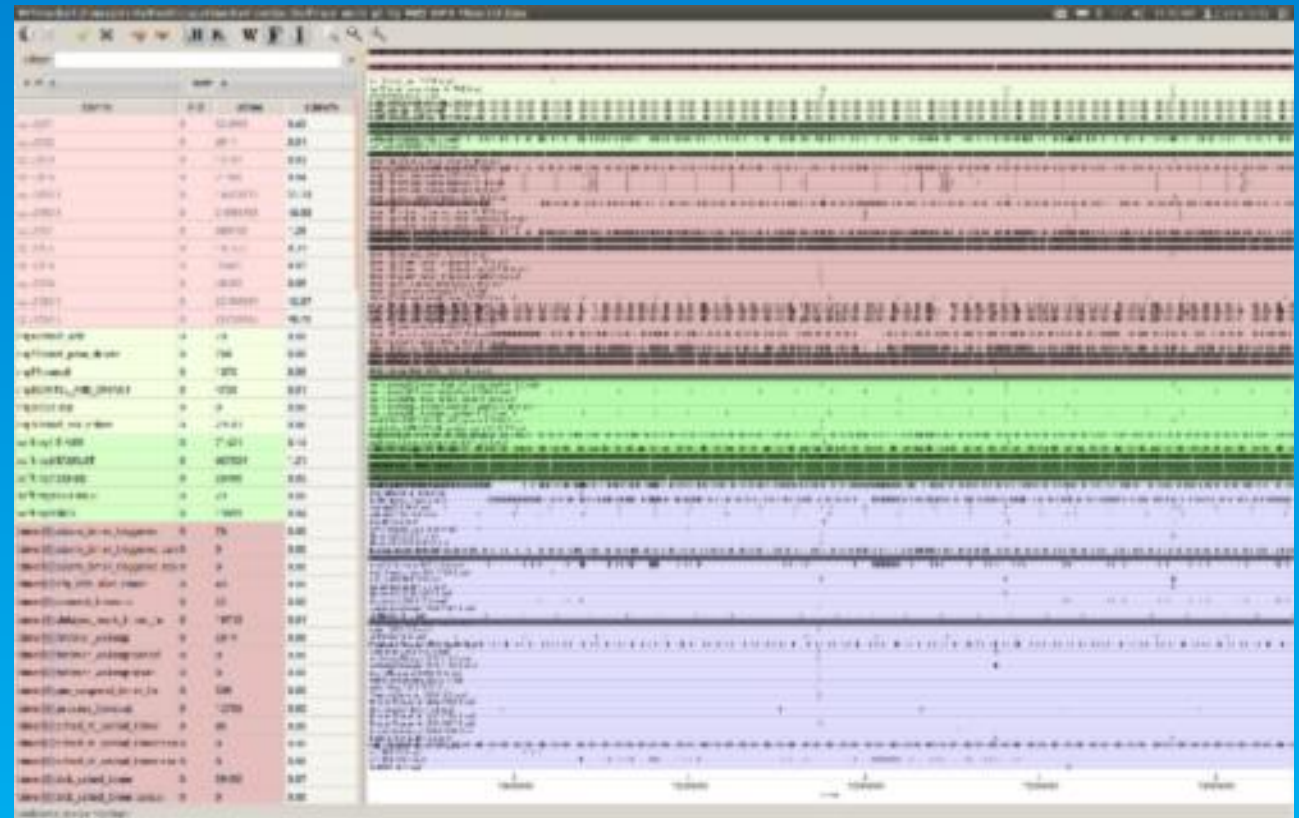
- Based on chaco, the browsing is very fast
- Based on python, parsing is very simple

One line per Process. Process in pytimechart can be:

- An IRQ callback
- A workqueue
- A tasklet
- A linux process, identified by pid, and comm
- A power state (cpuidle, runtime_pm)

Simple browsing

- Zoom with mouse wheel
- Select a part of the chart
 - Show time range
 - Processor share
 - Original trace as text



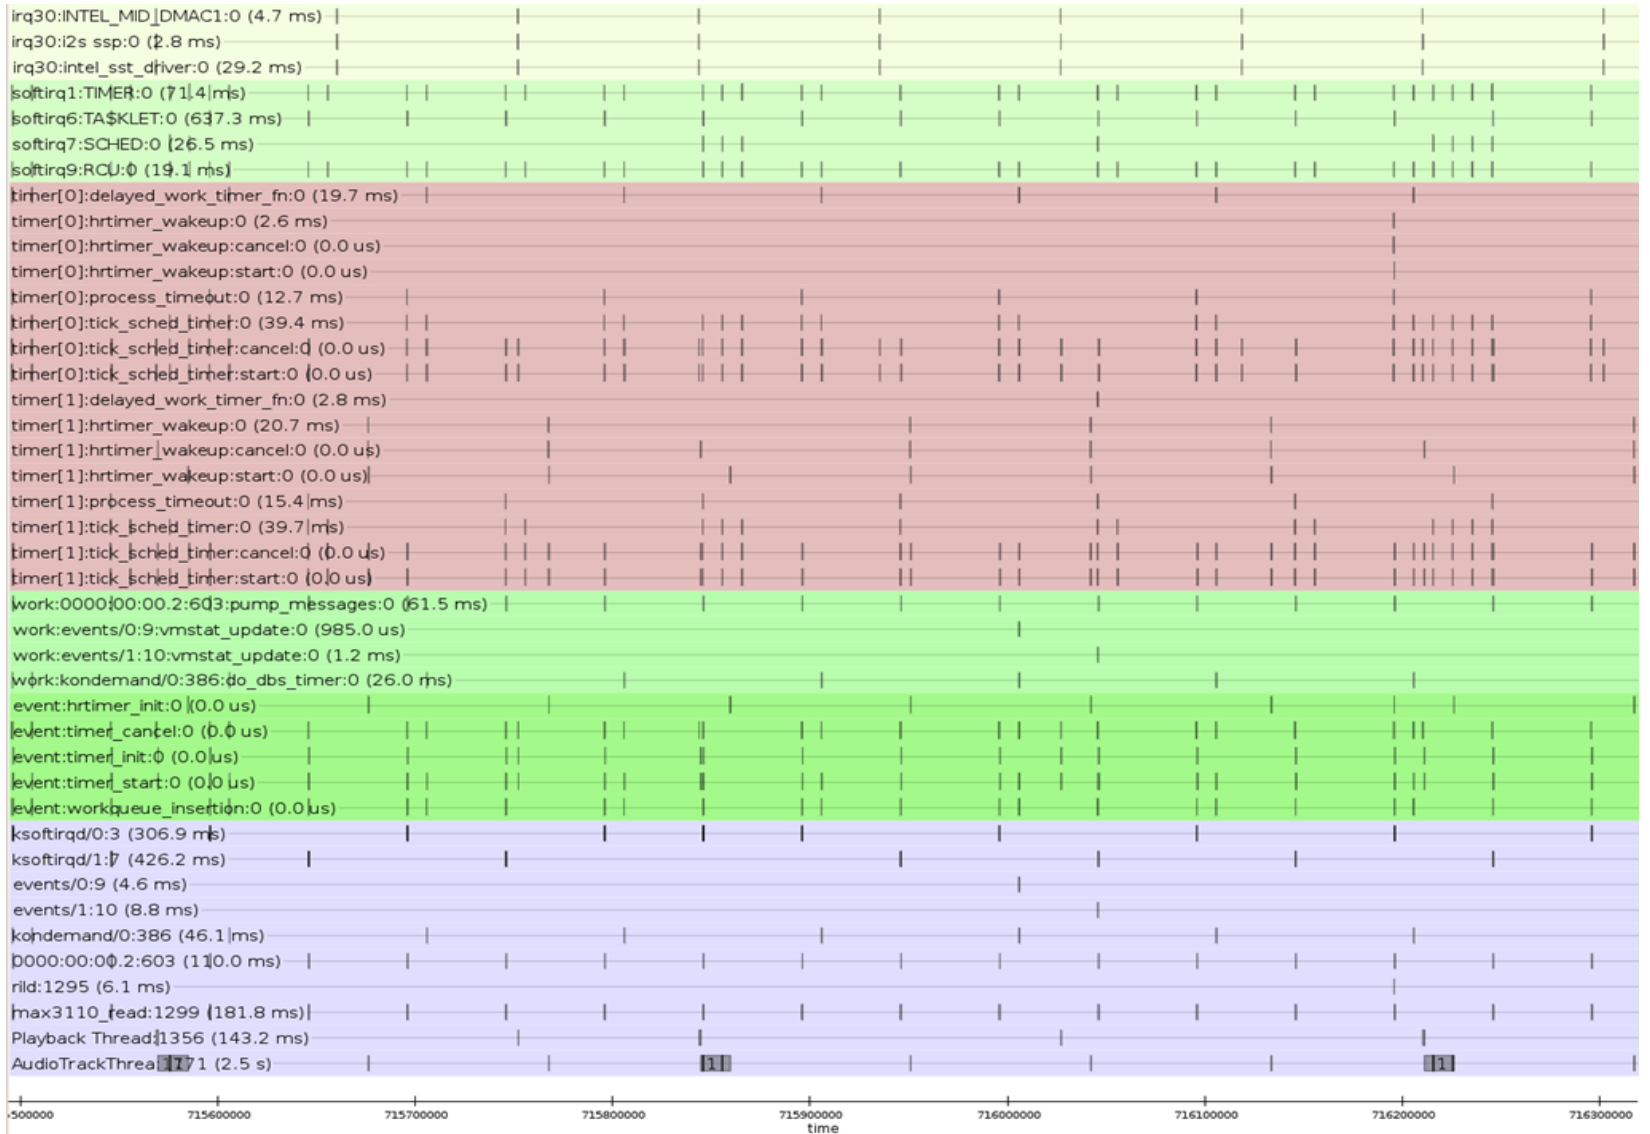
Playing mp3 is easy. Playing your whole music library in one battery charge is much harder.

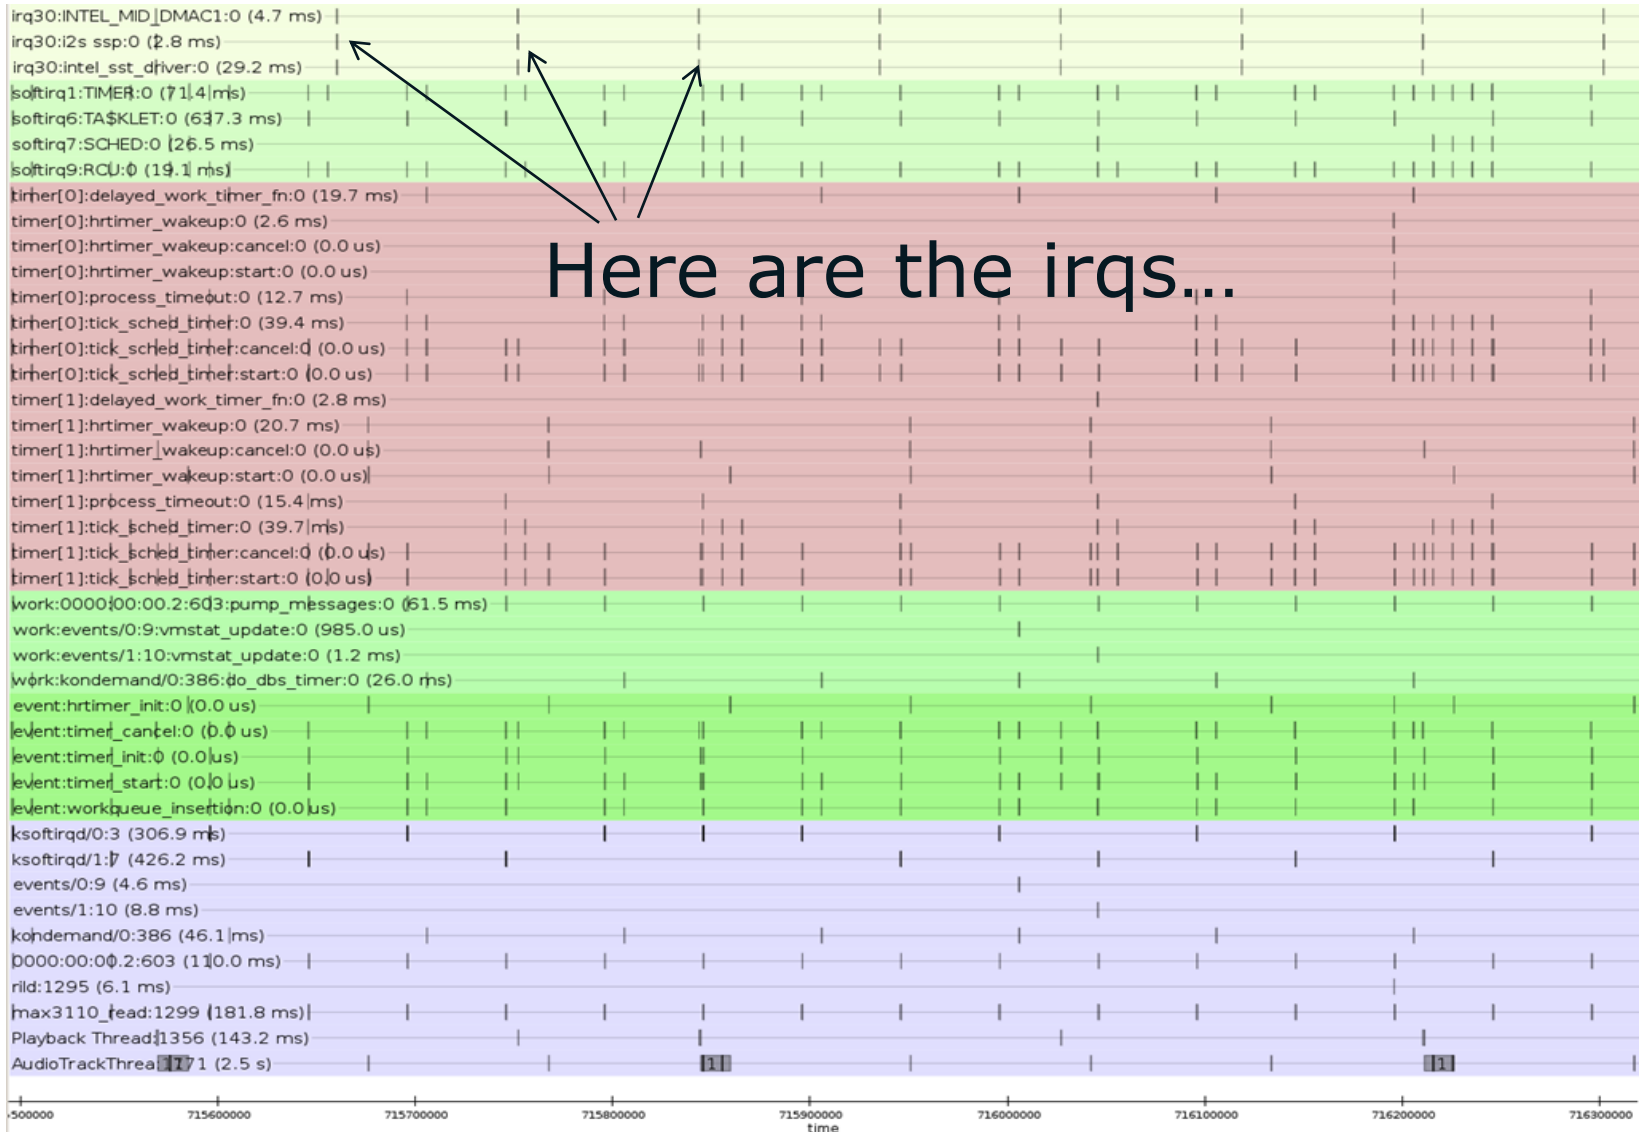
Need to improve wake up count. Each time the cpu wakes up, energy is spent (cache, PLL, etc)

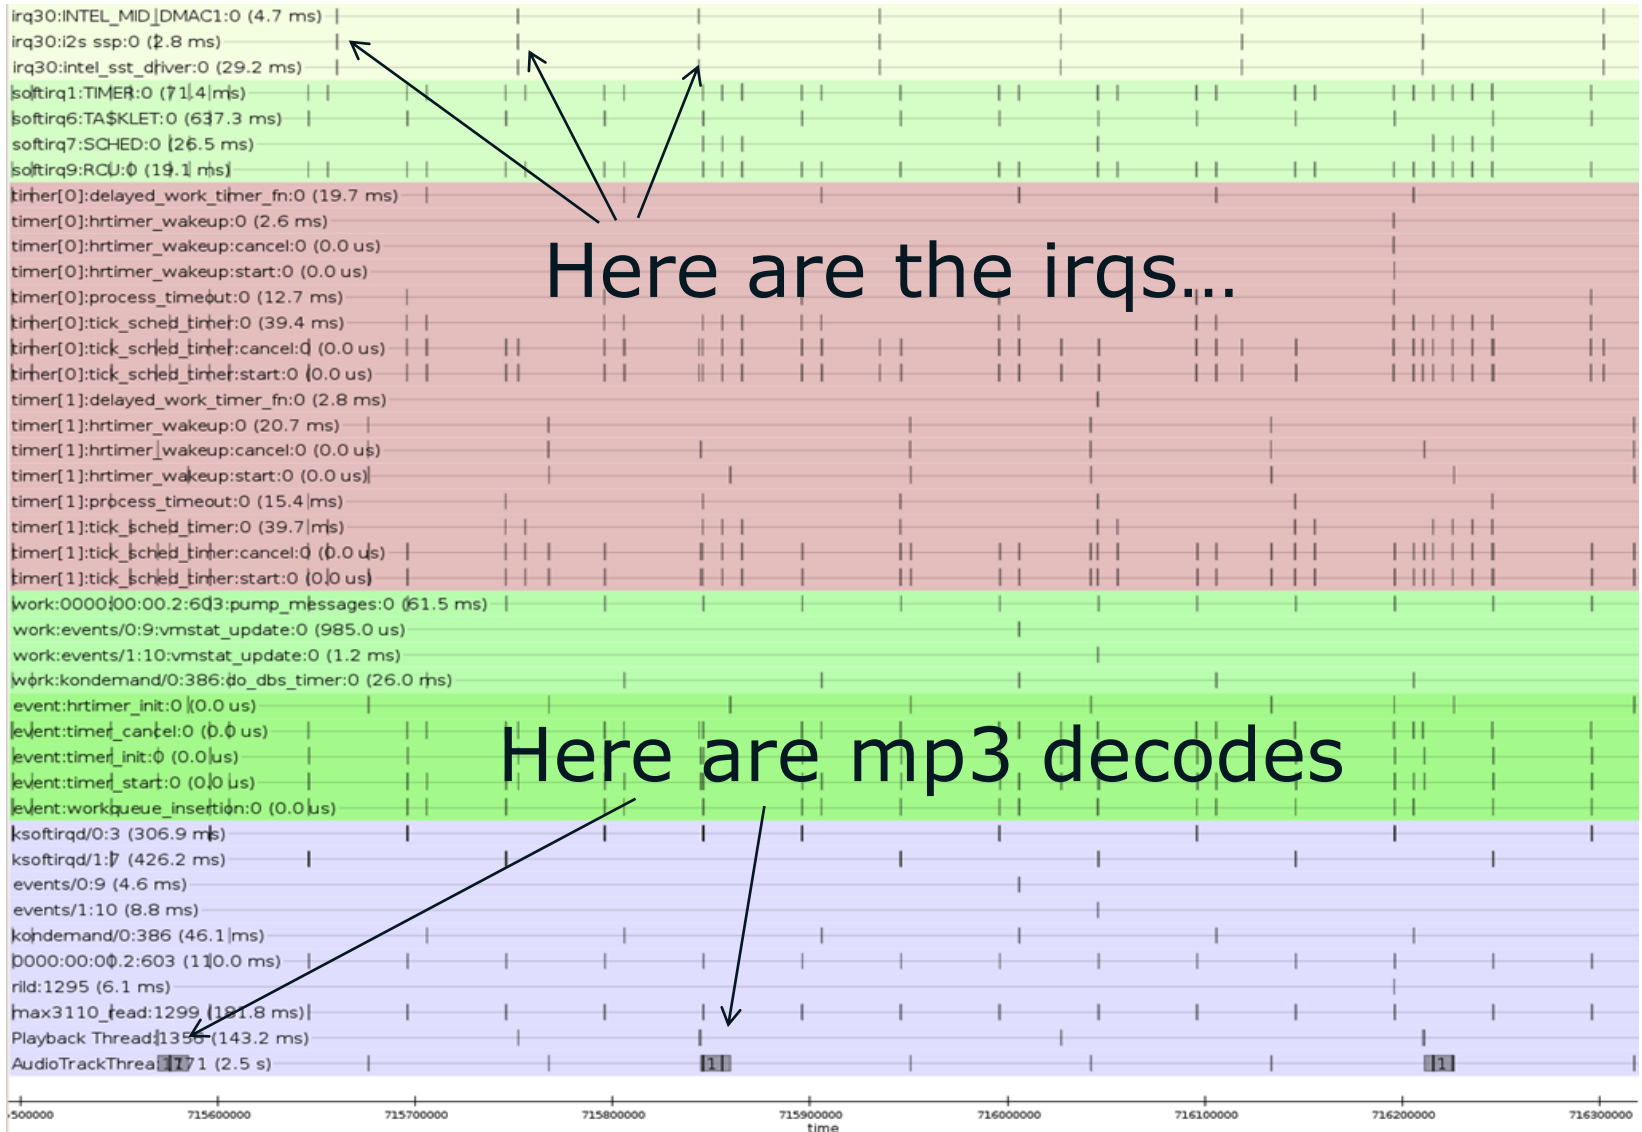
Audio HW is generating IRQ when a new buffer is needed. This goes all the way through user space, where the mp3 decode is done.

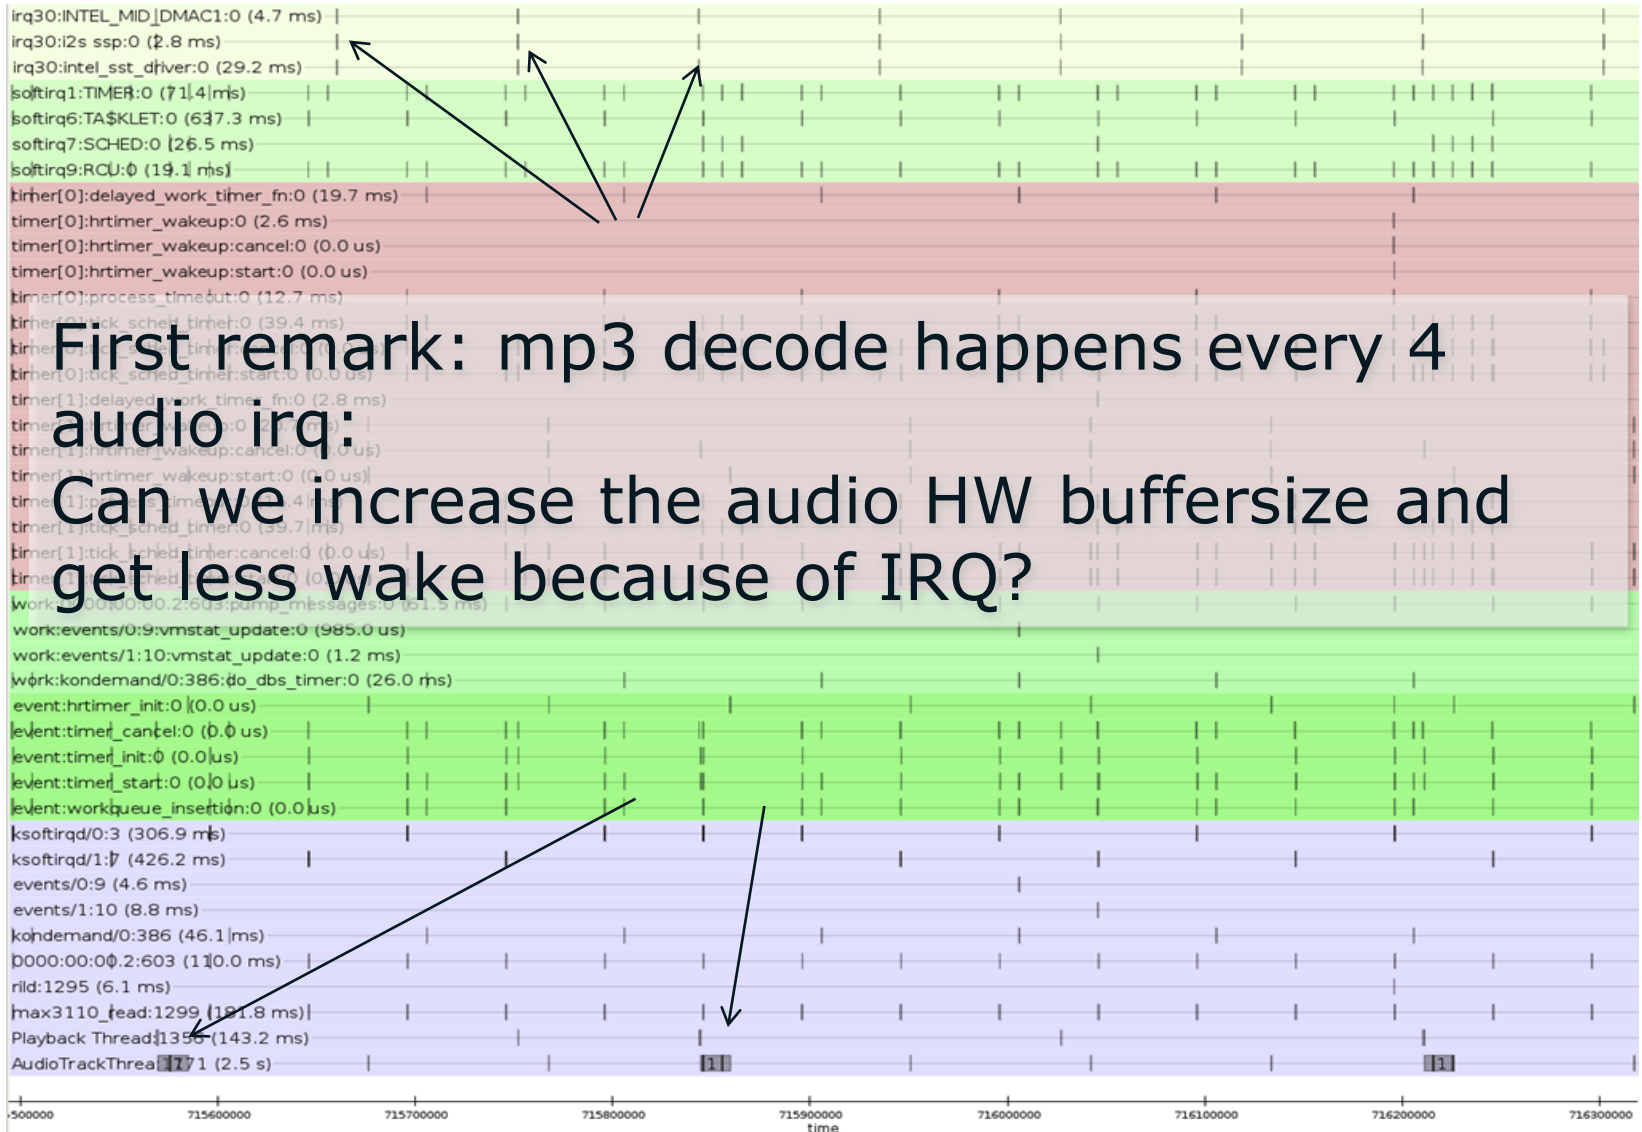
We are waiting for:

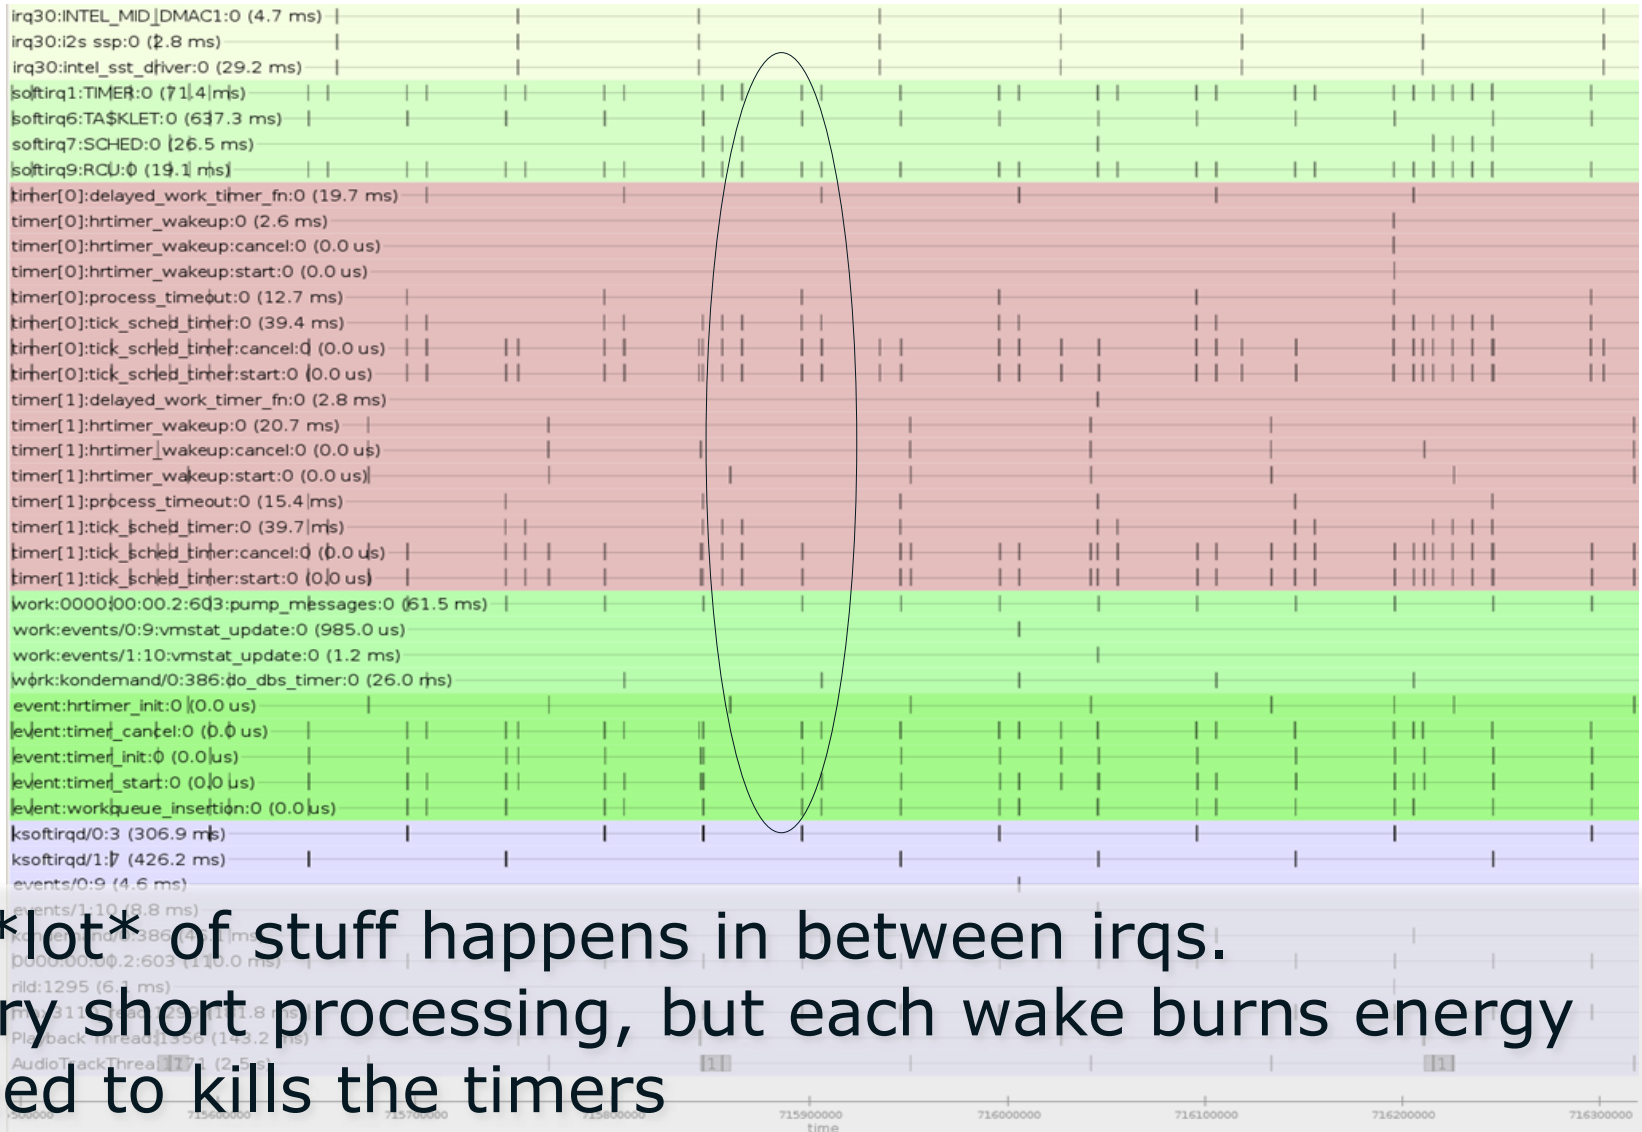
- Audio IRQ
- MP3Decode
- Sleep
- Audio IRQ
- MP3Decode
- Sleep
- Etc.



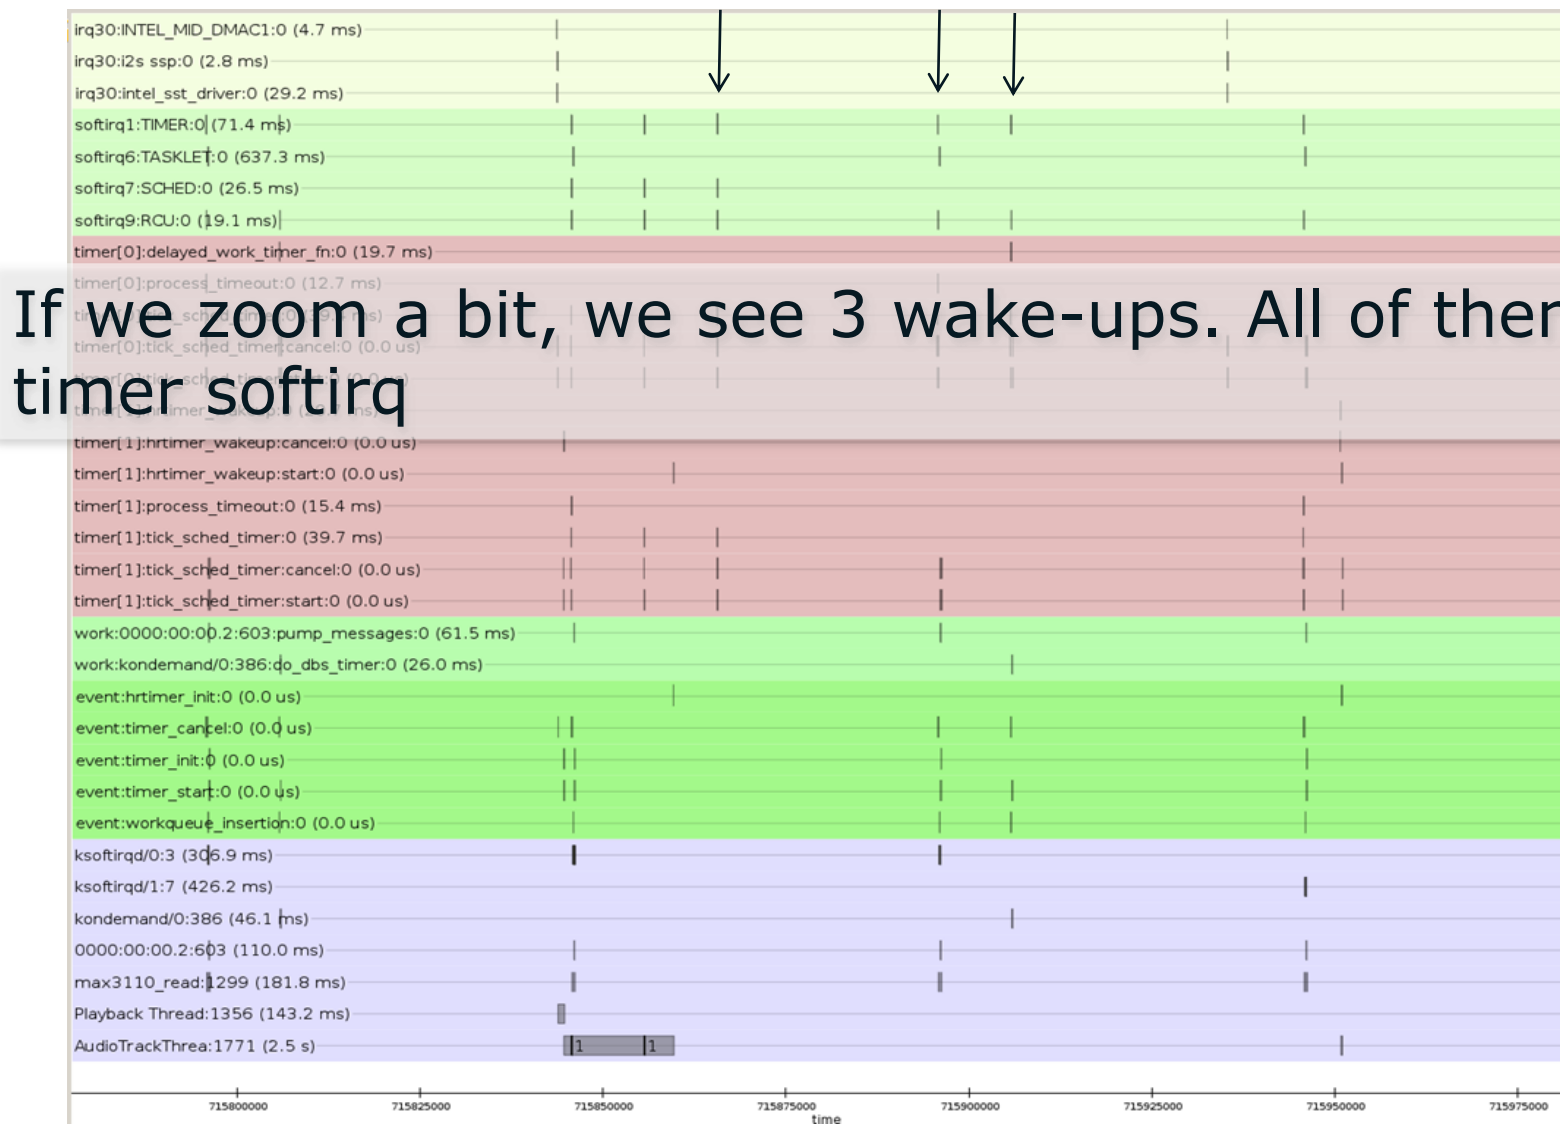




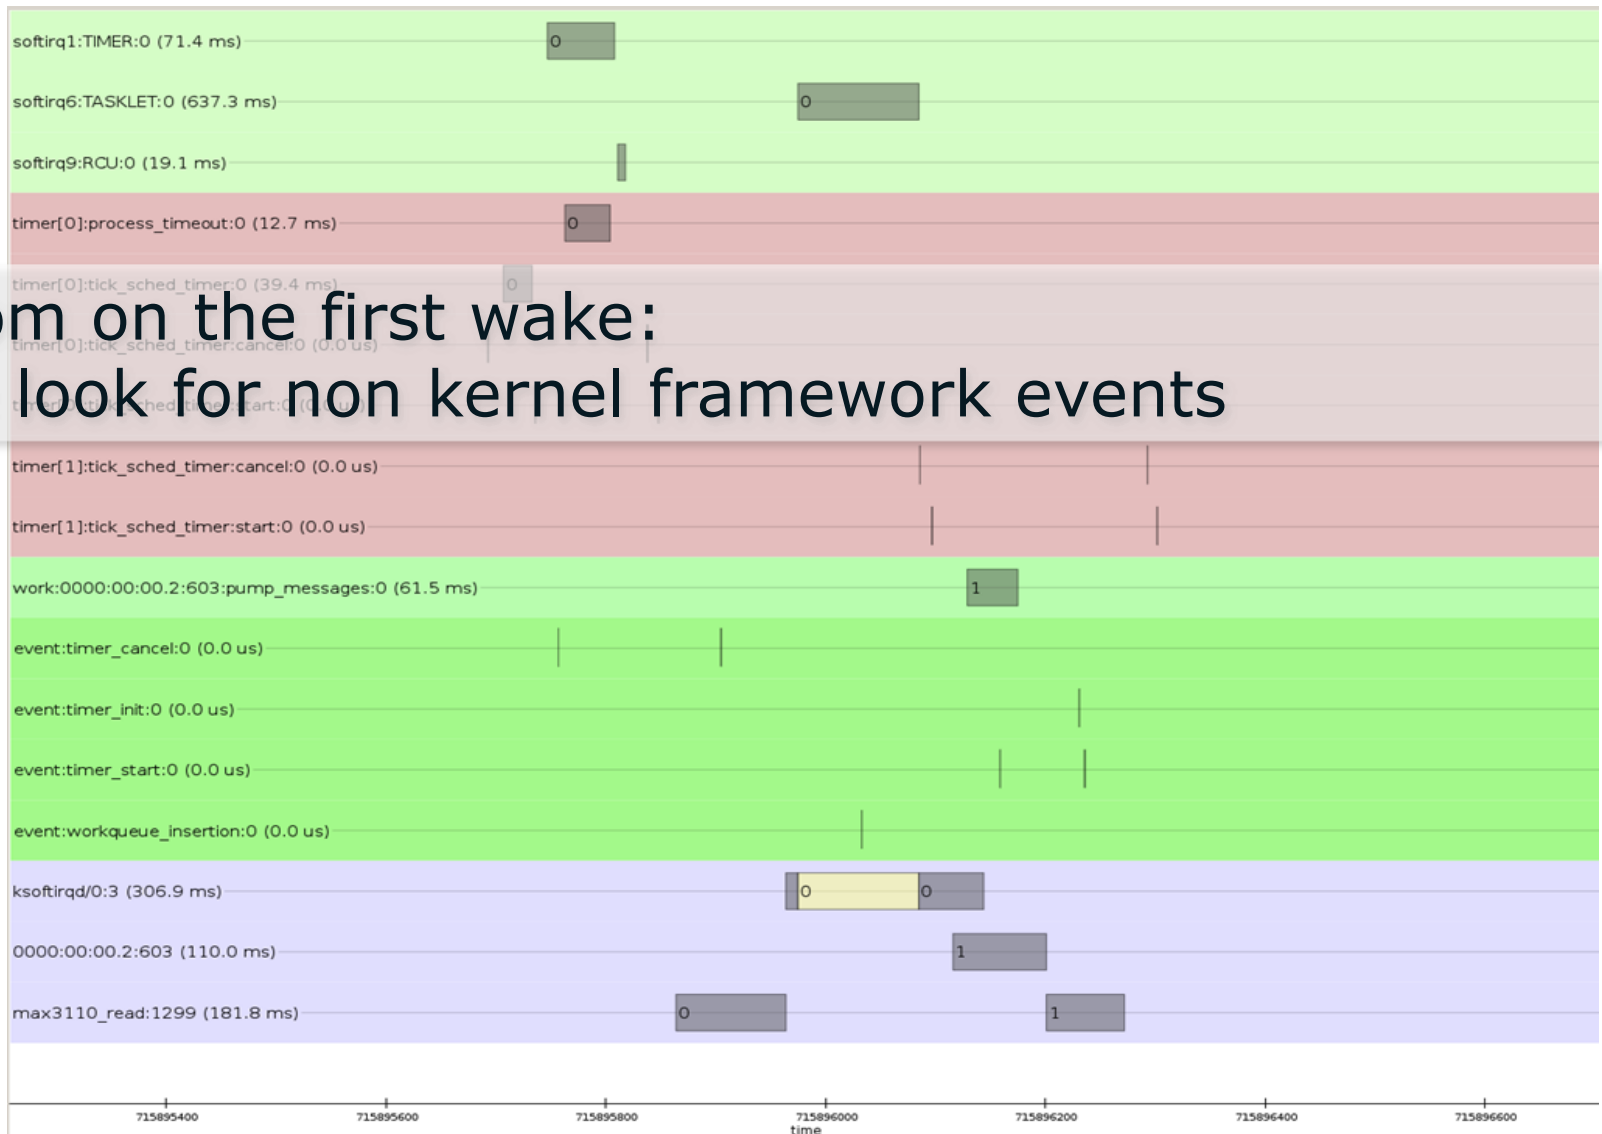




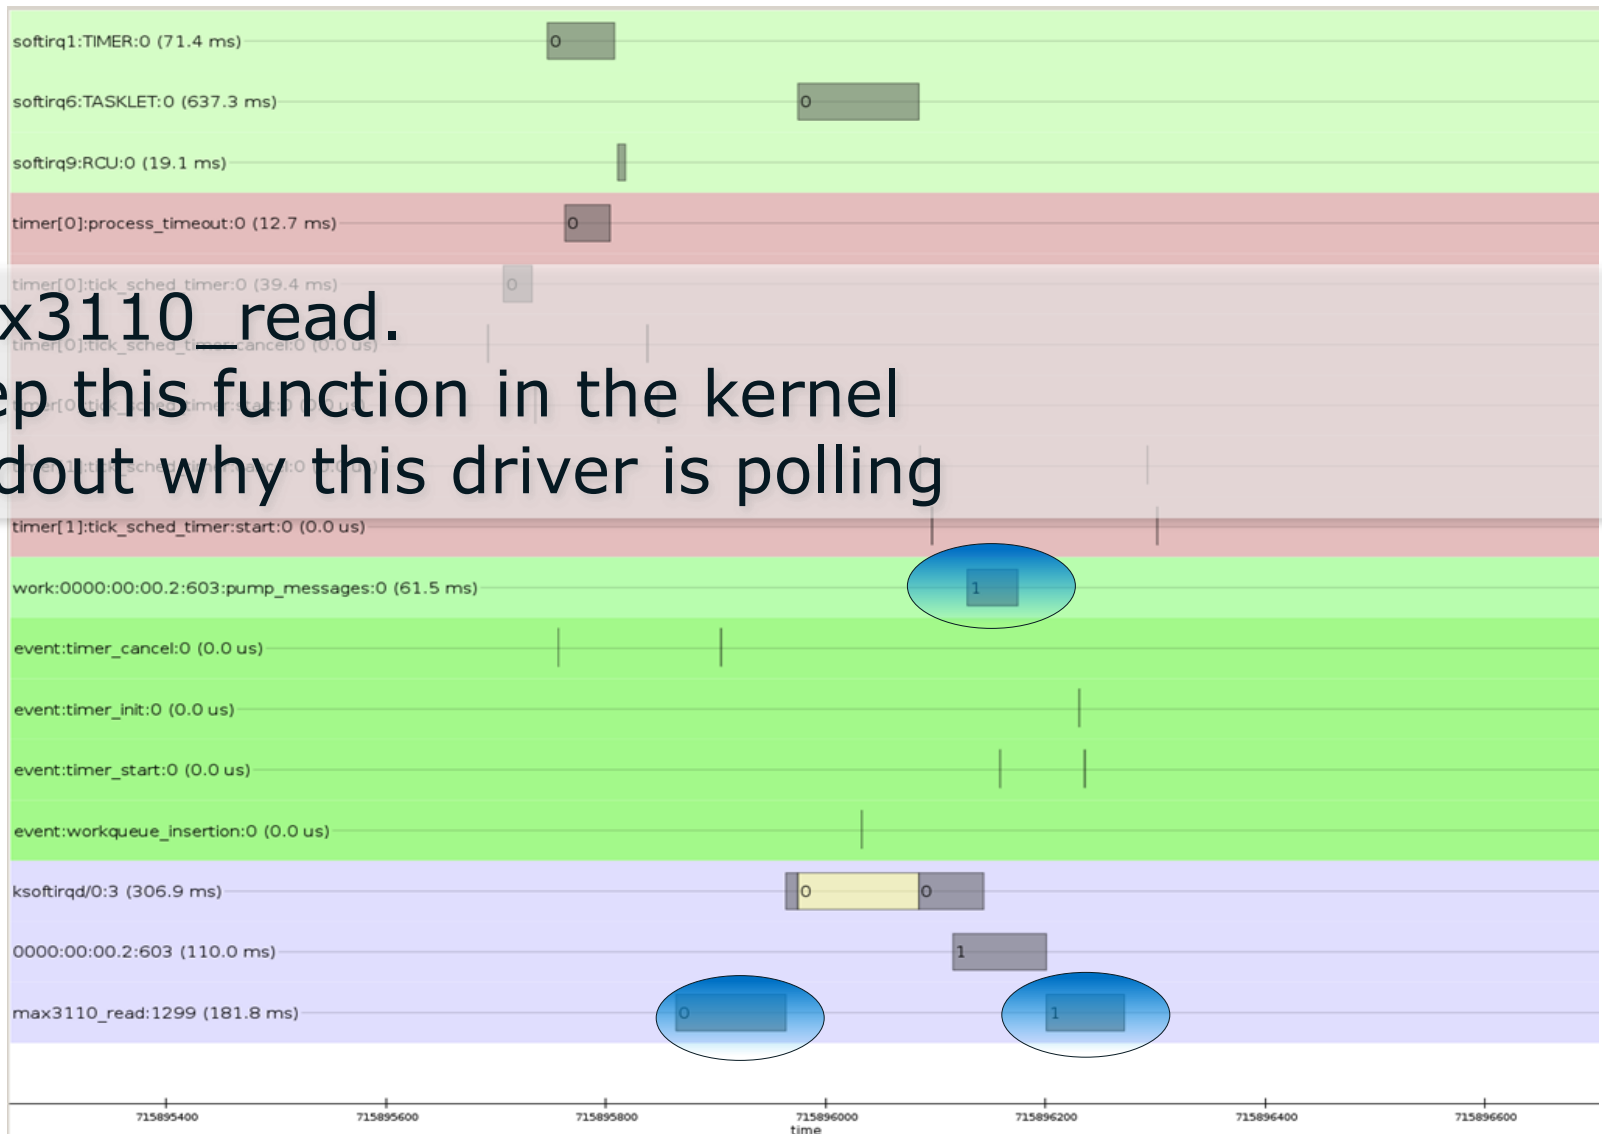
A *lot* of stuff happens in between irqs.
Very short processing, but each wake burns energy
Need to kill the timers



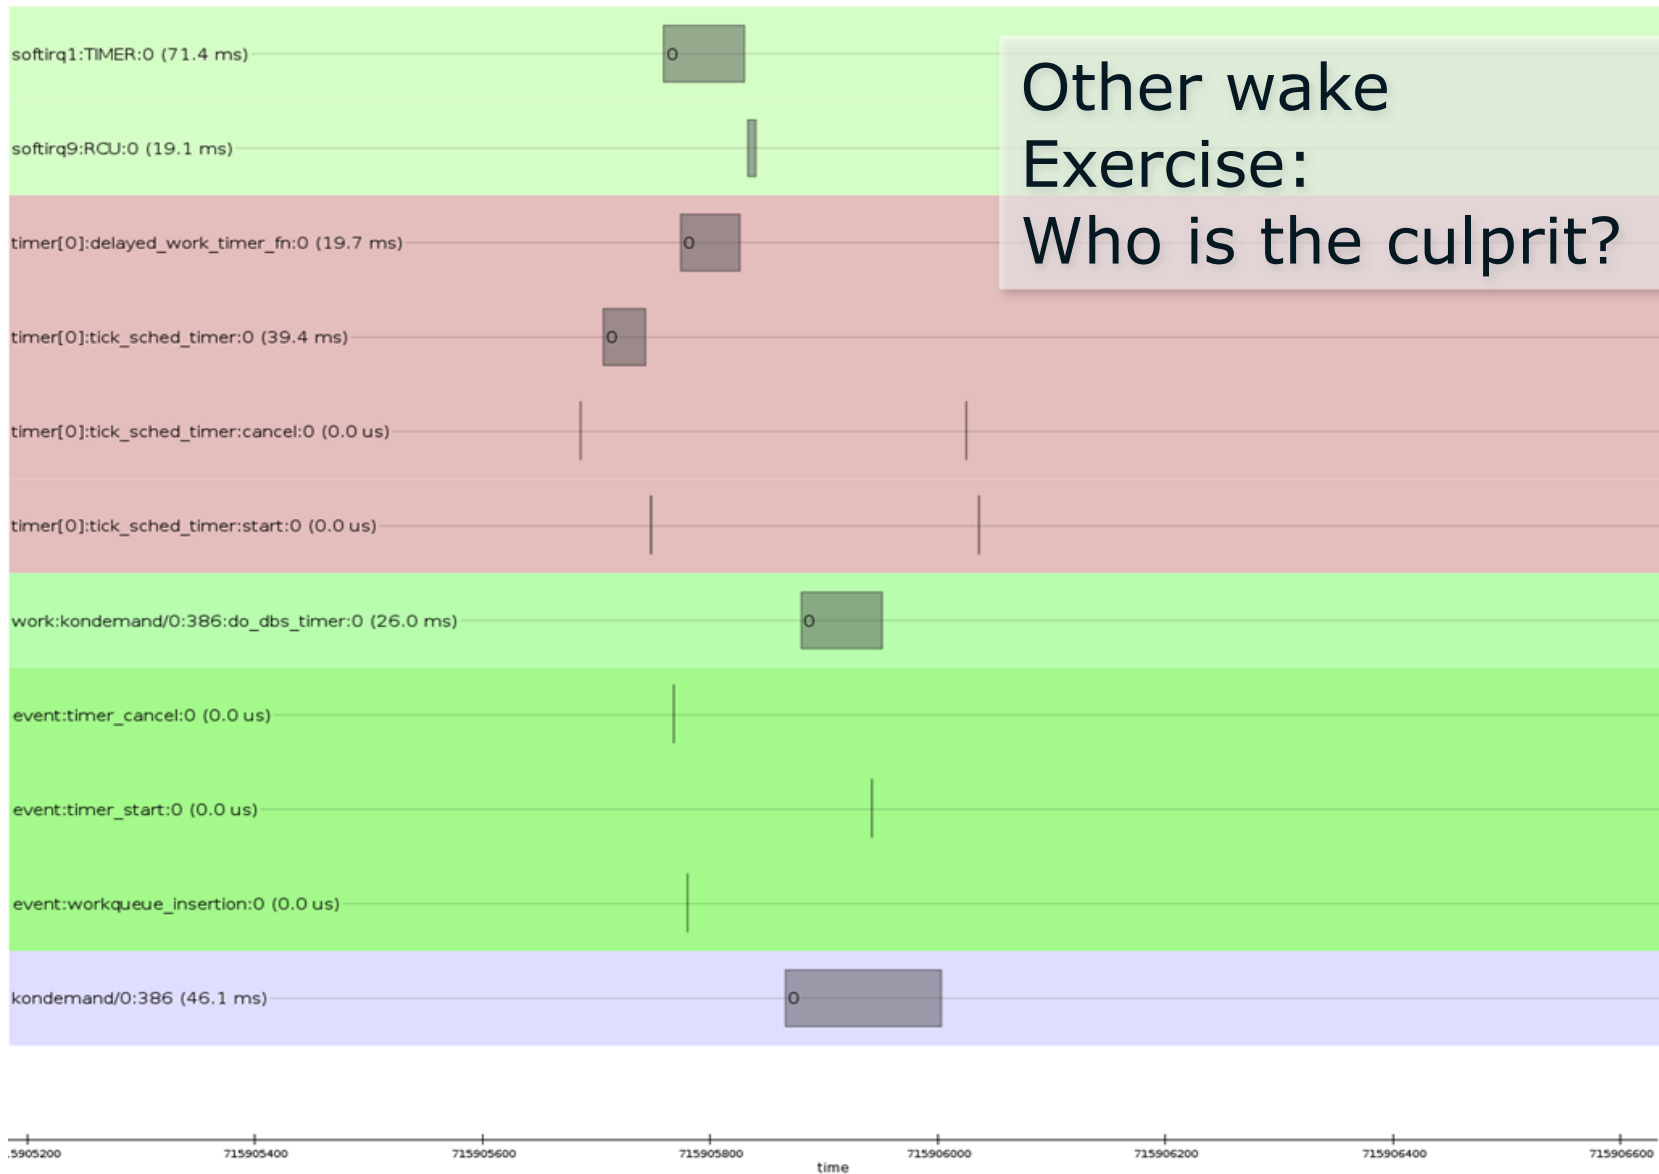
zoom on the first wake:
We look for non kernel framework events

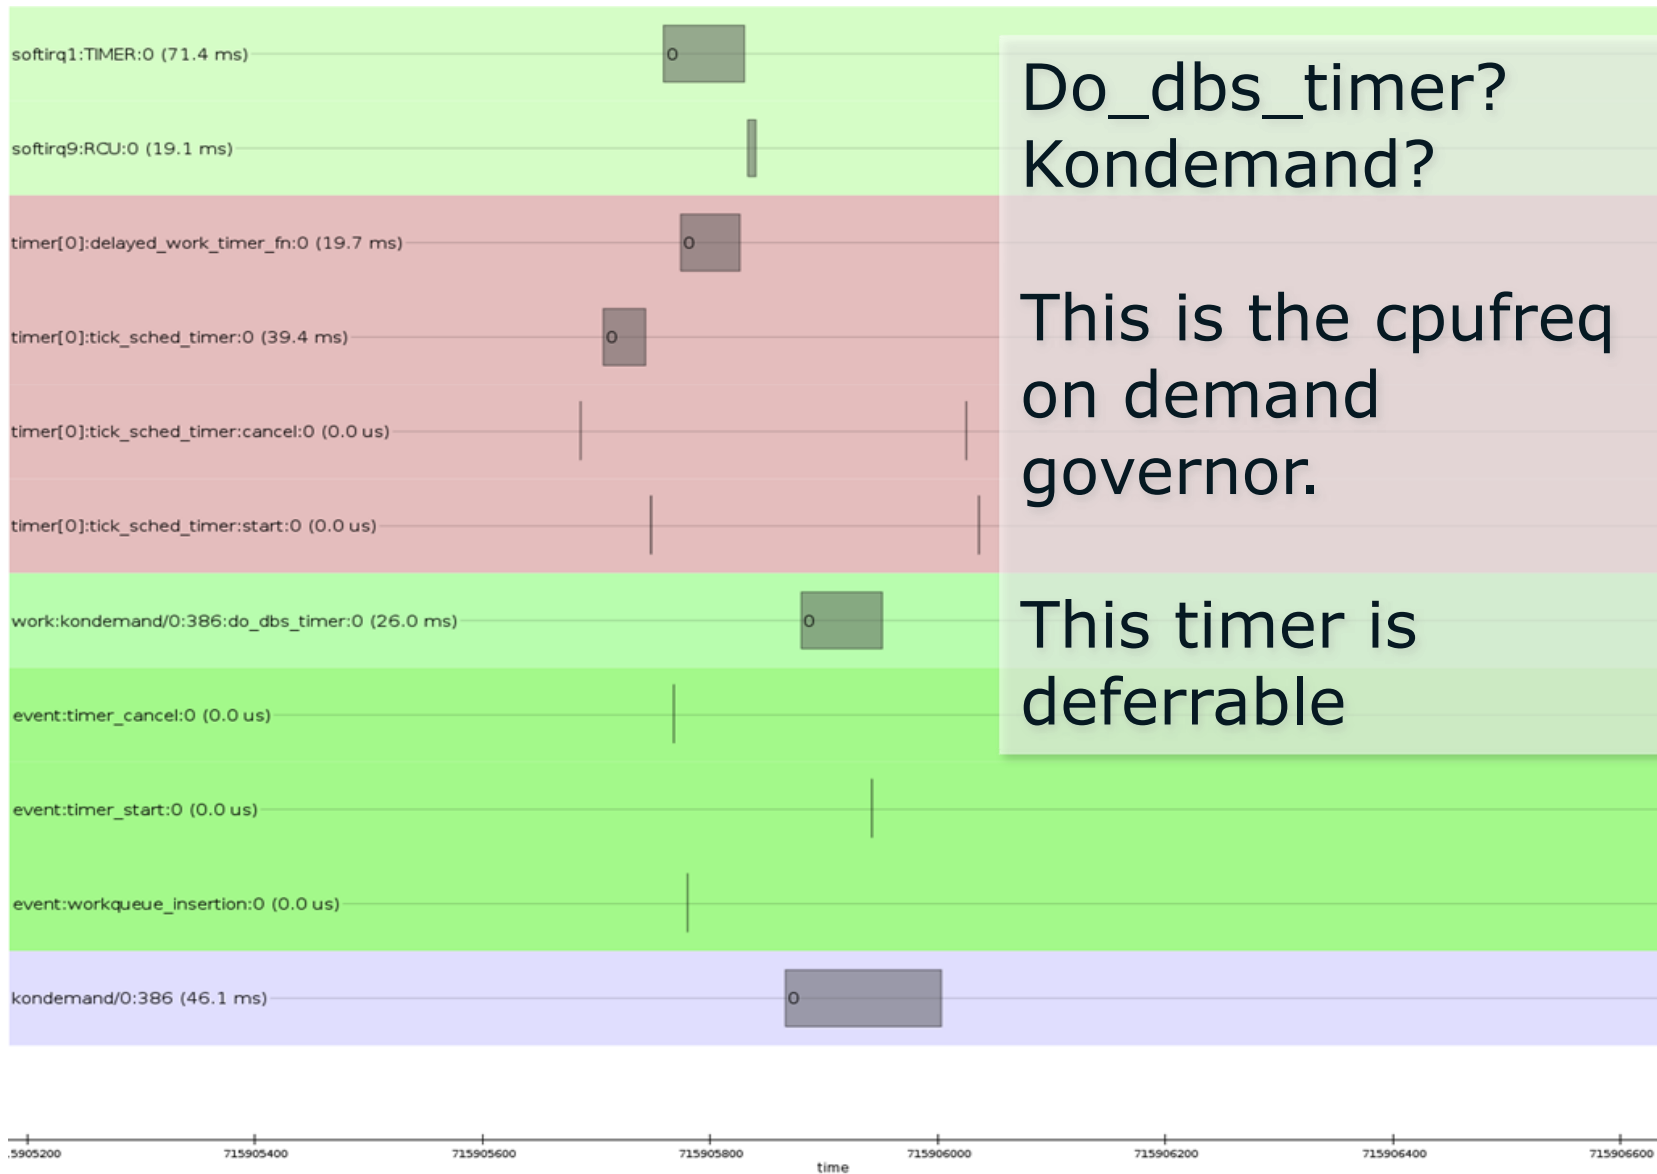


max3110_read.
Grep this function in the kernel
Findout why this driver is polling



Other wake Exercise: Who is the culprit?





Do_dbs_timer?
Kondemand?

This is the cpufreq
on demand
governor.

This timer is
deferrable



Index

PowerTOP

Tickless Idle

Introduction

Deferrable timers

round_jiffies

FAQ

Downloads

Applications Power
Management

Processor Power
Management

Power and Performance
Measurement

Linux ACPI

ACPICA

BLTK

Power QoS

Display and Graphics

Power Saving

Device and Bus Power
Management

Tickless Idle

Deferrable timers

In the Linux kernel, timers provide a way for any driver or kernel routine to register an event that needs to be handled at some point in the future.

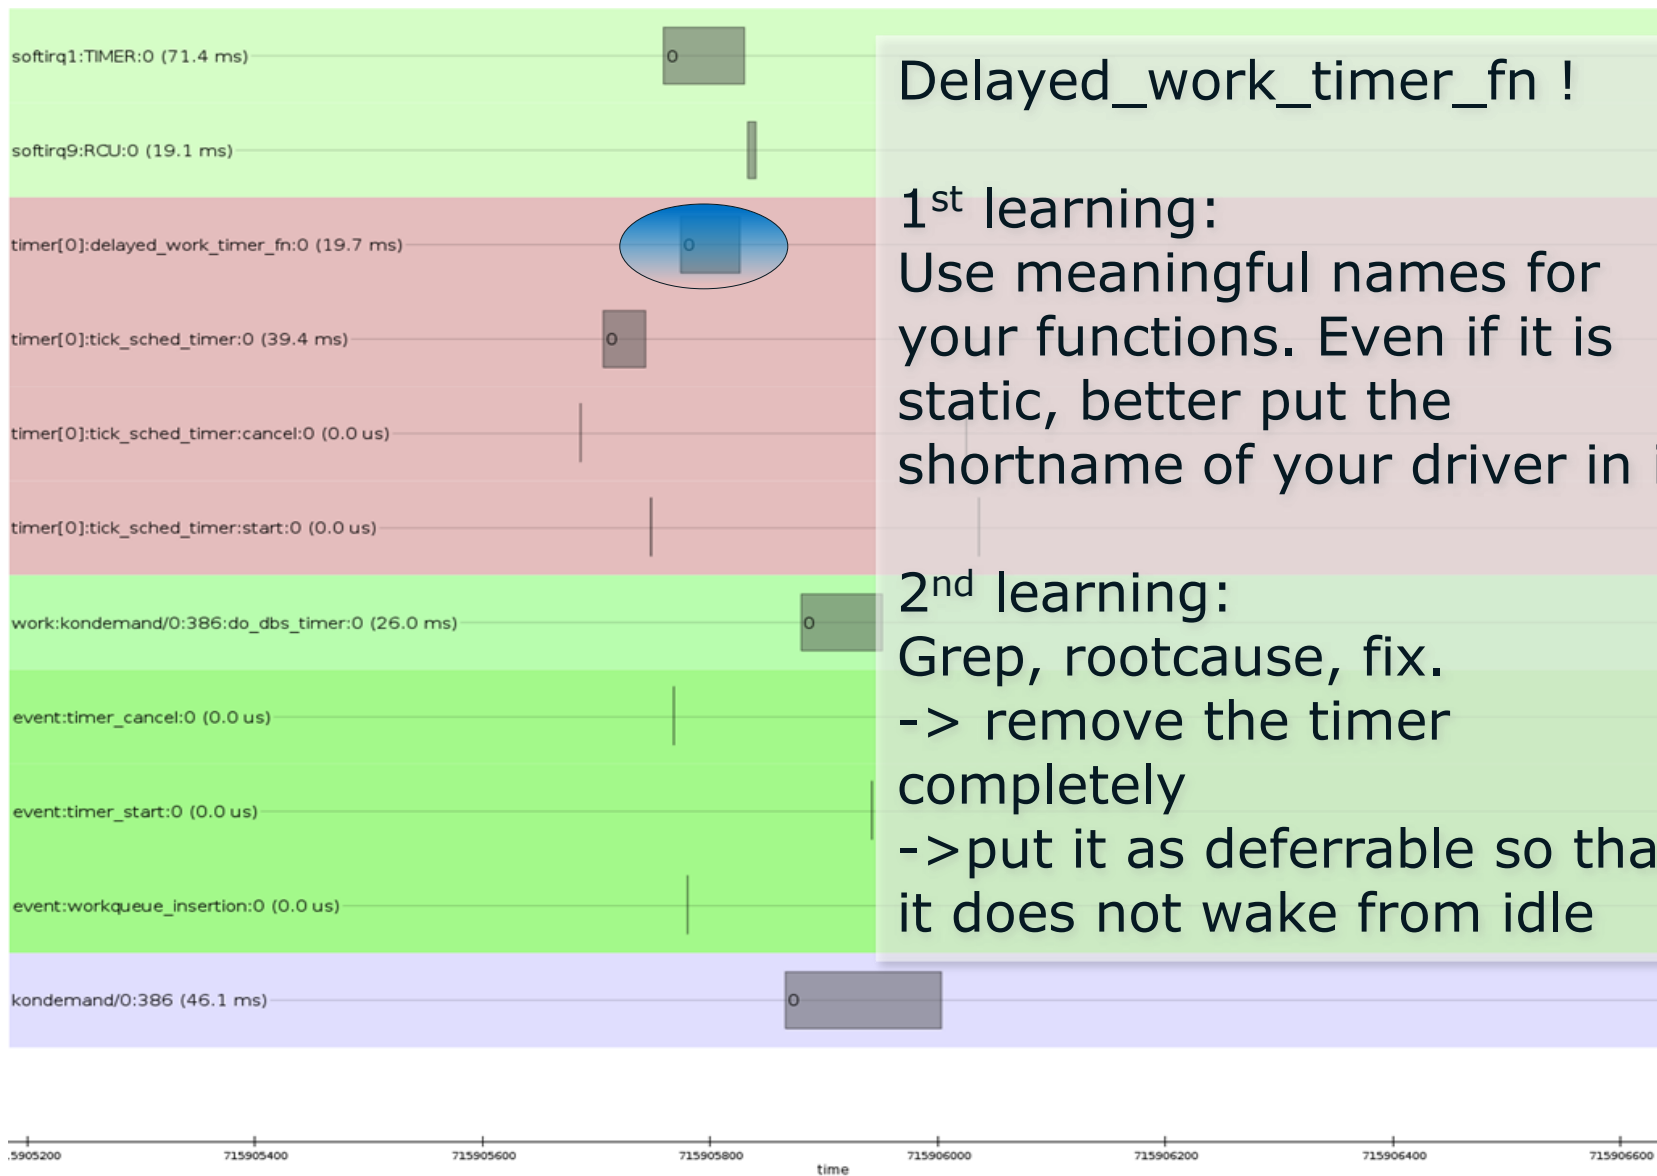
Some timers, used inside the kernel, are important to service in a sub-second interval when the system is busy, but are not critical when the system is idle. A prime example, is the timer usage in ondemand driver, which does a periodic sampling of CPU use.

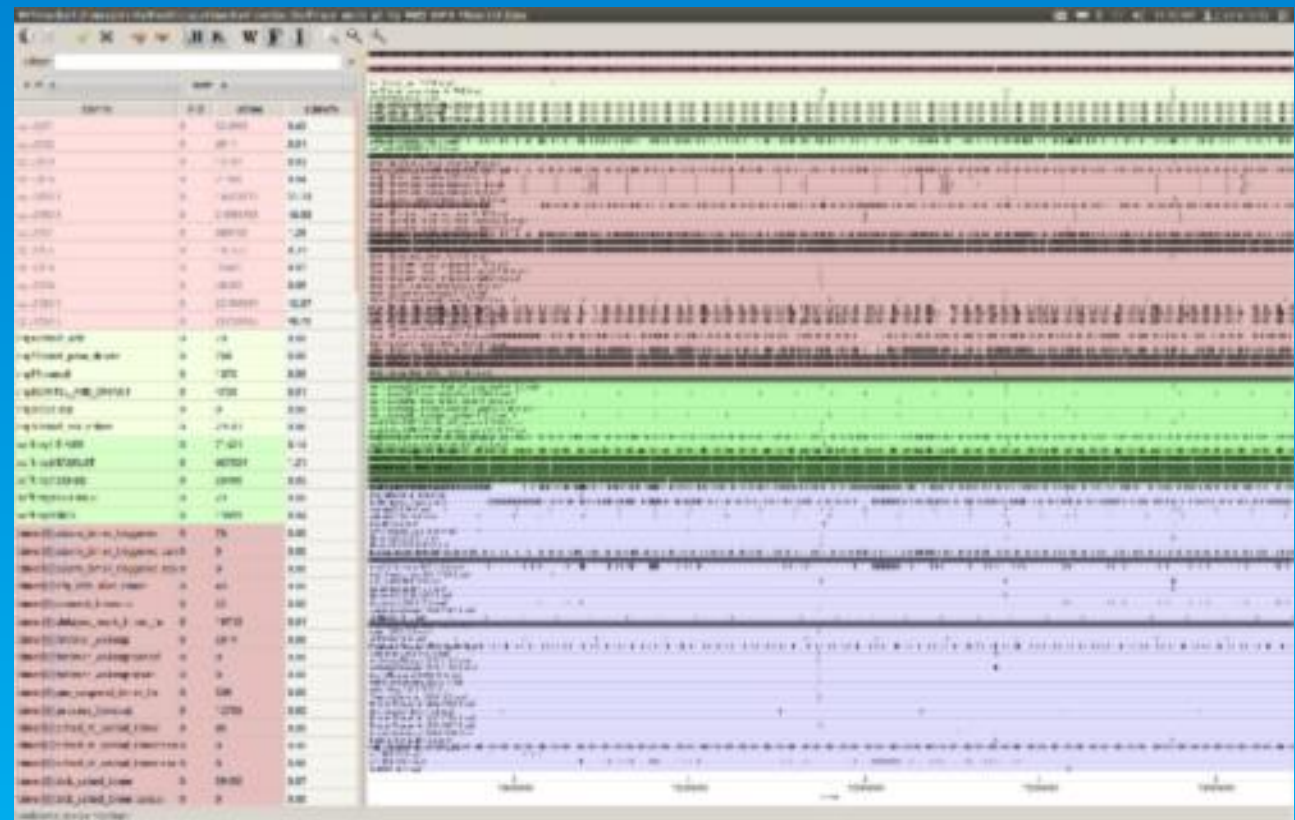
Accuracy is important for these timers when the CPU is busy because they are used to increase the frequency of a CPU to meet increased demand. But, while system is idle, ondemand governor can tolerate some delay and let the CPU remain idle for a longer time.

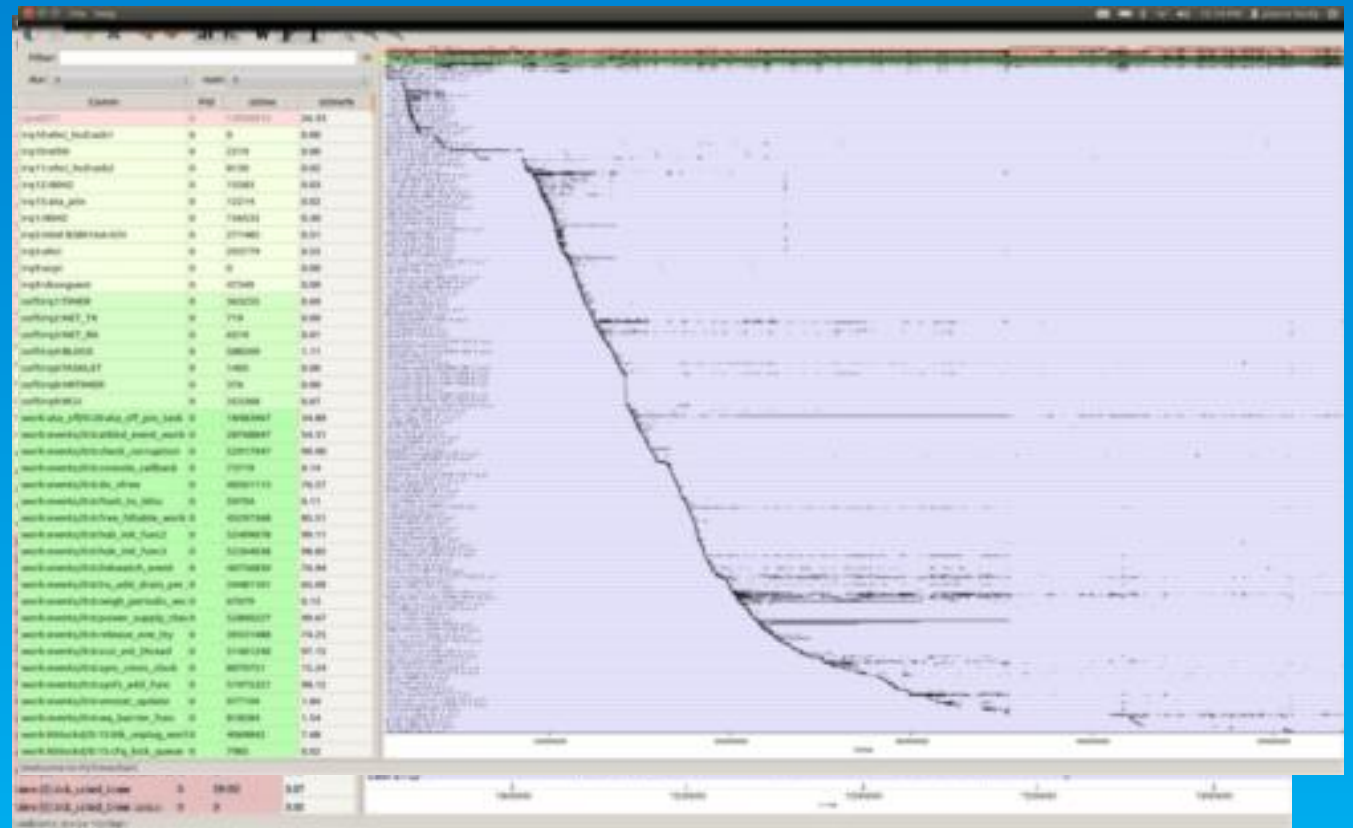
To effectively handle such timers, the "deferrable timer" infrastructure was introduced in the Linux kernel. These deferrable timers are handled normally when the CPU is busy. However, when the CPU is idle, deferrable timers are queued until there is a non-deferrable timer or an interrupt that wakes the CPU from idle.

Using deferrable timers significantly reduces the number of times a CPU is woken up from idle. The data in the table below is based on an Intel® Core™ 2 Duo-based test platform. The table shows the number of interrupts per second and the average time the CPU spends in idle state (in uS). Note that the kernel used during the measurement had tickless idle enabled with HZ being 1000.

	# interrupts	#events	Avg CPU idle residency (uS)
Ondemand	118	60.60	10161
Ondemand + deferrable timer	89	17.17	20312

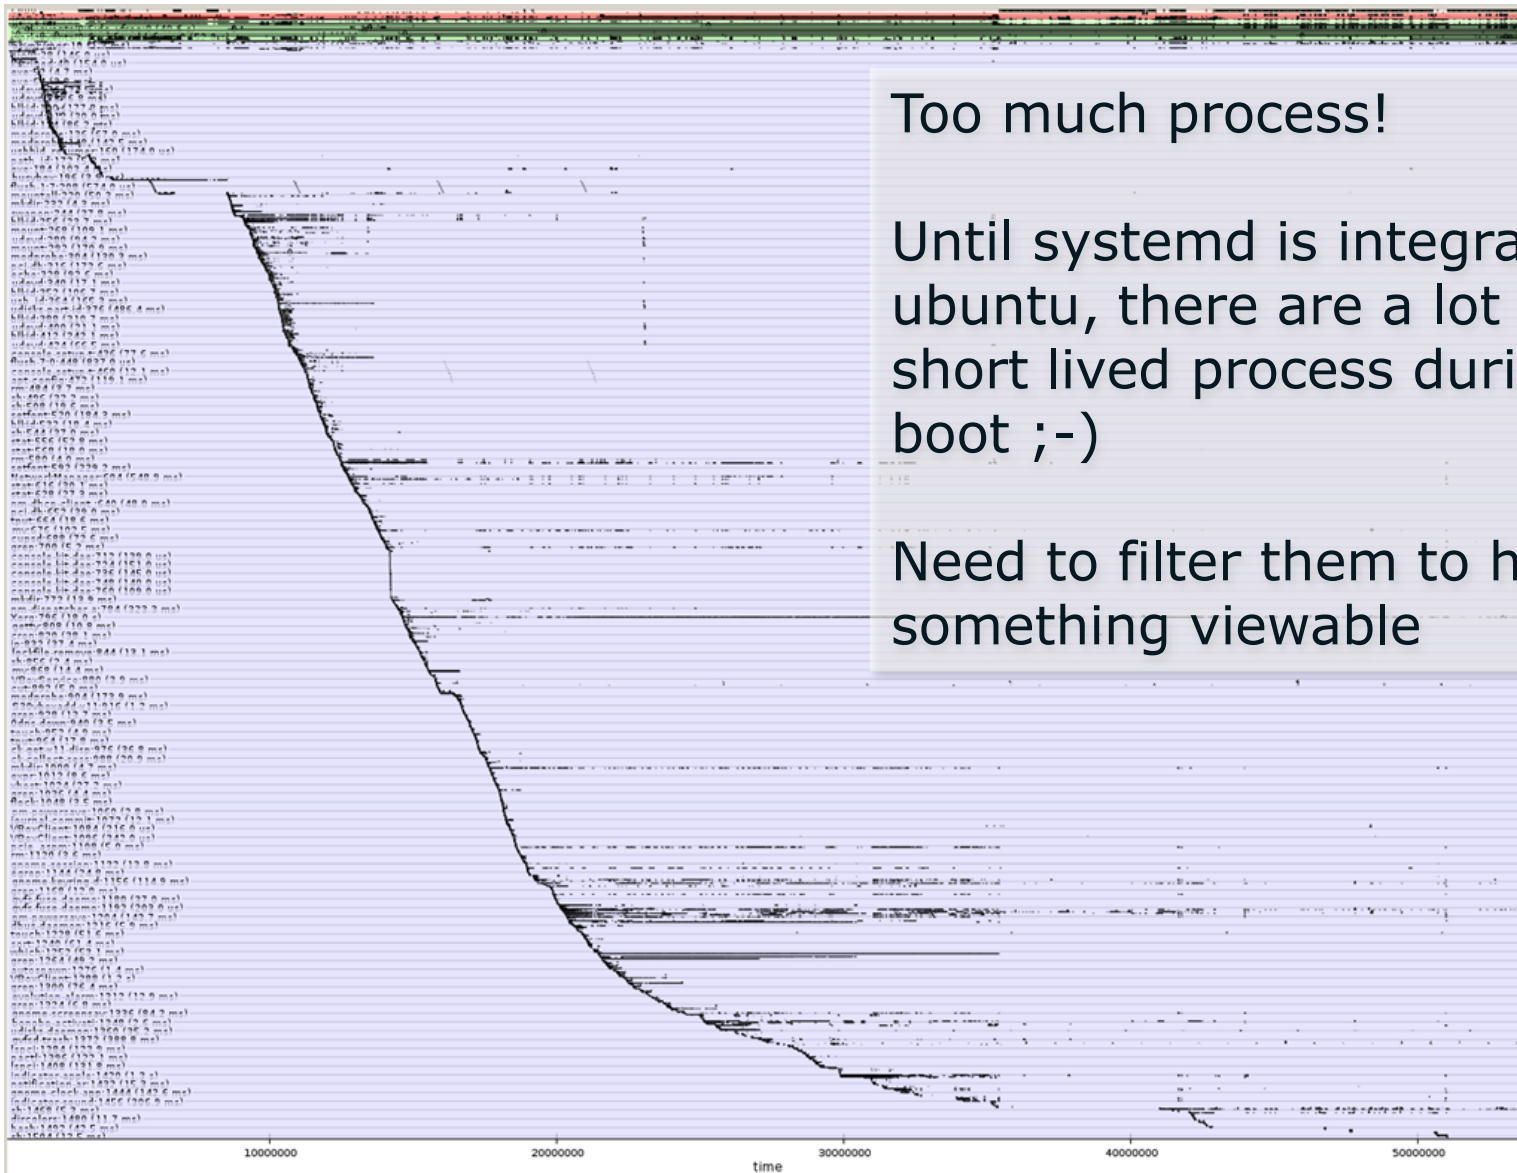


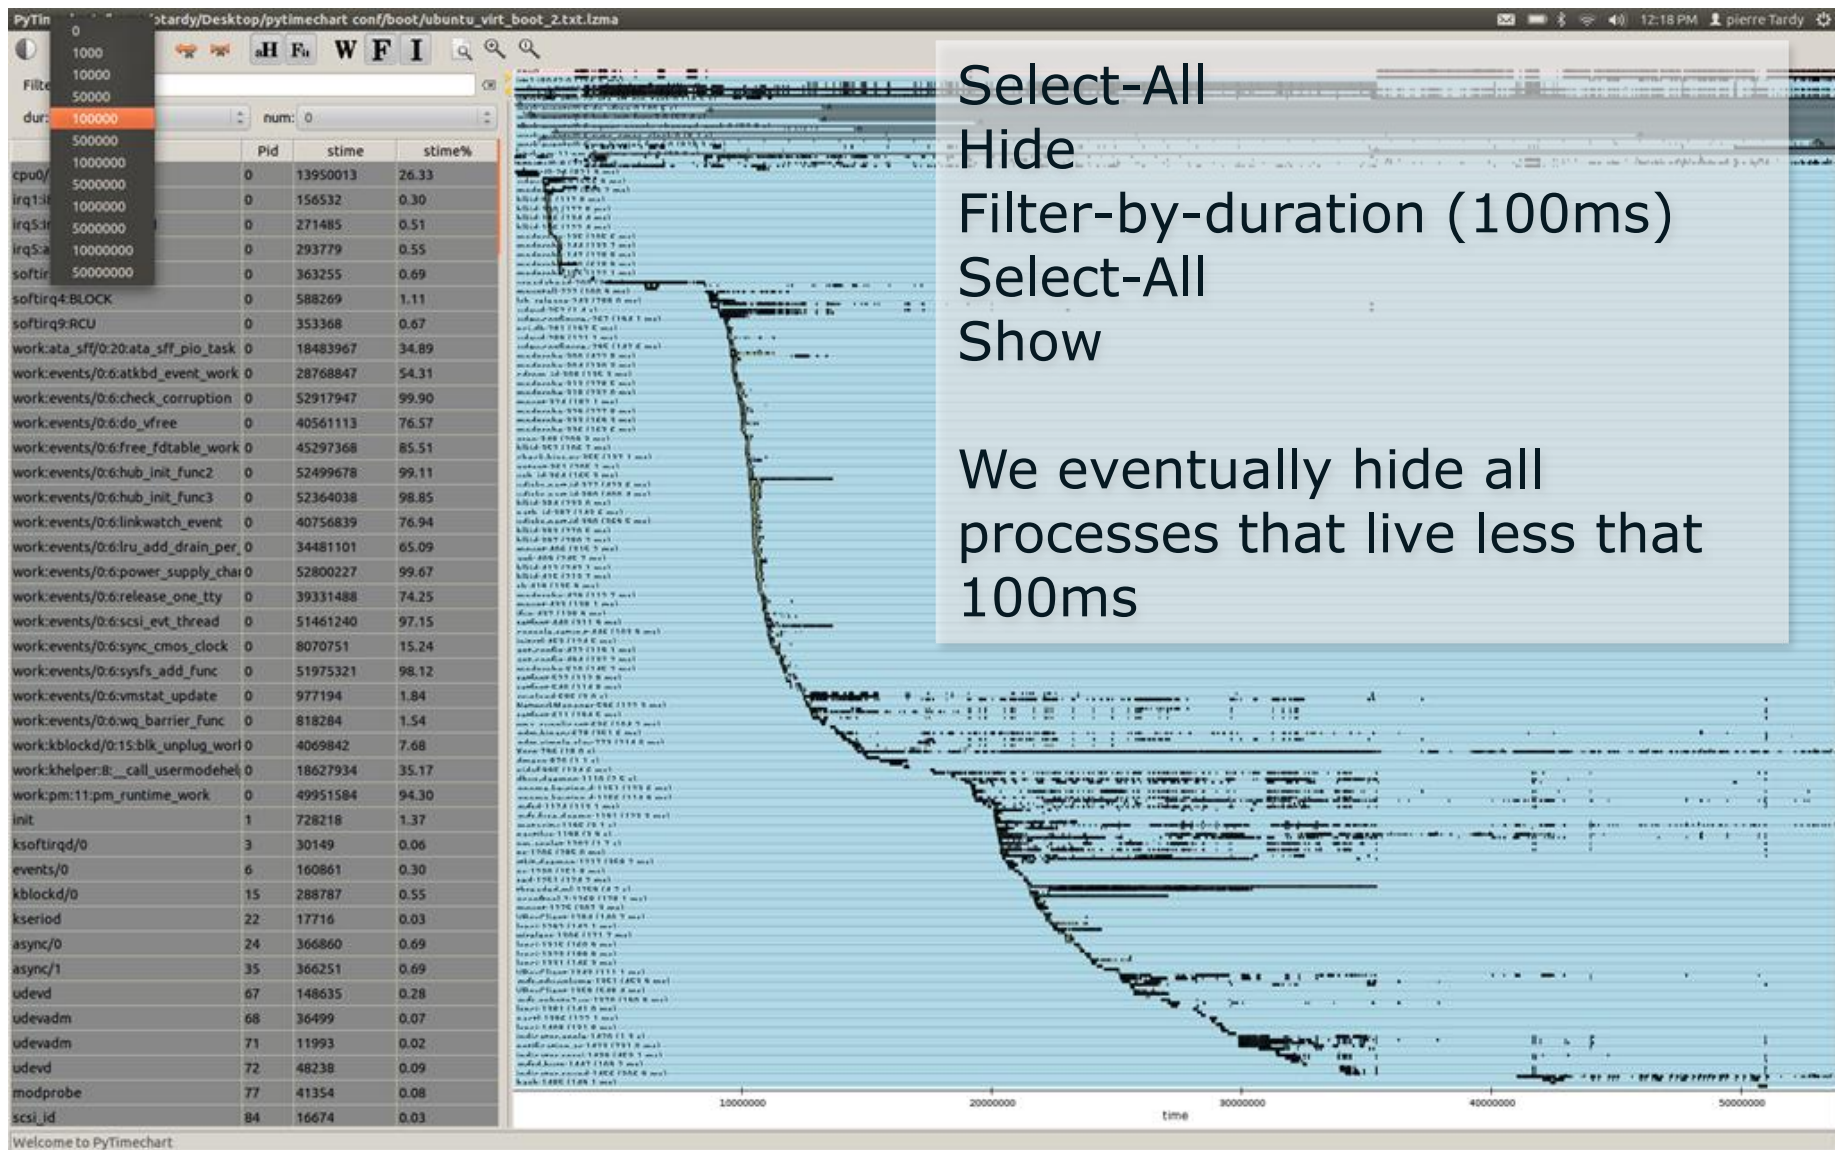




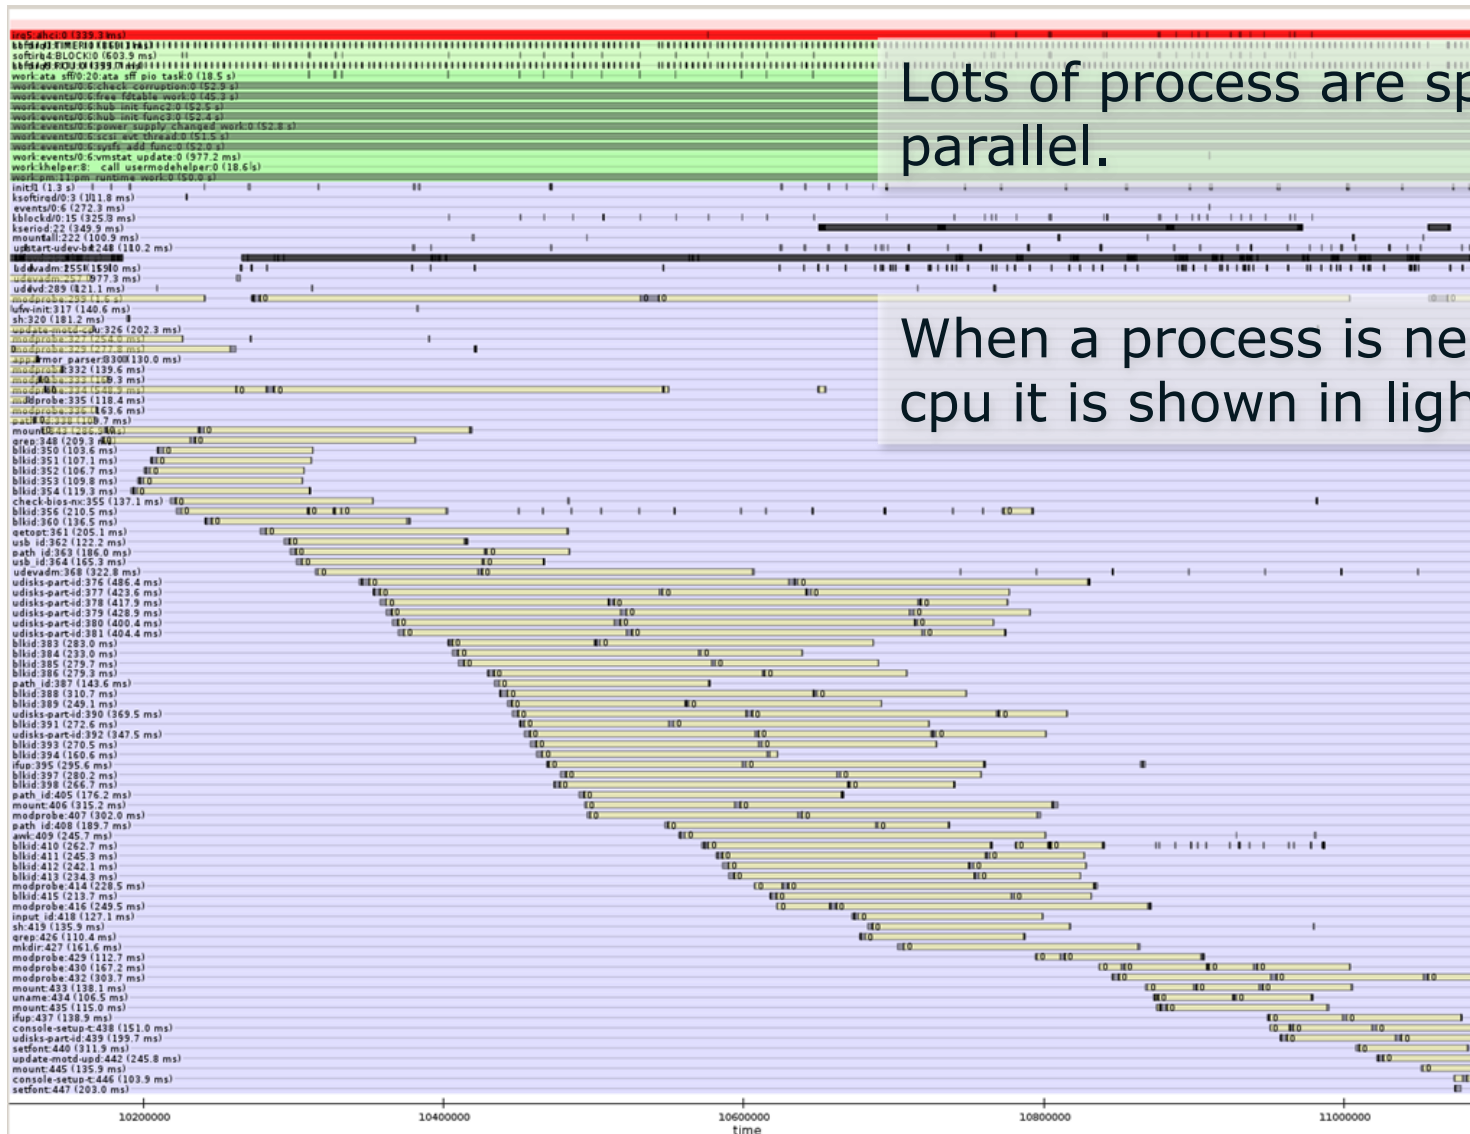
Whats wrong here?



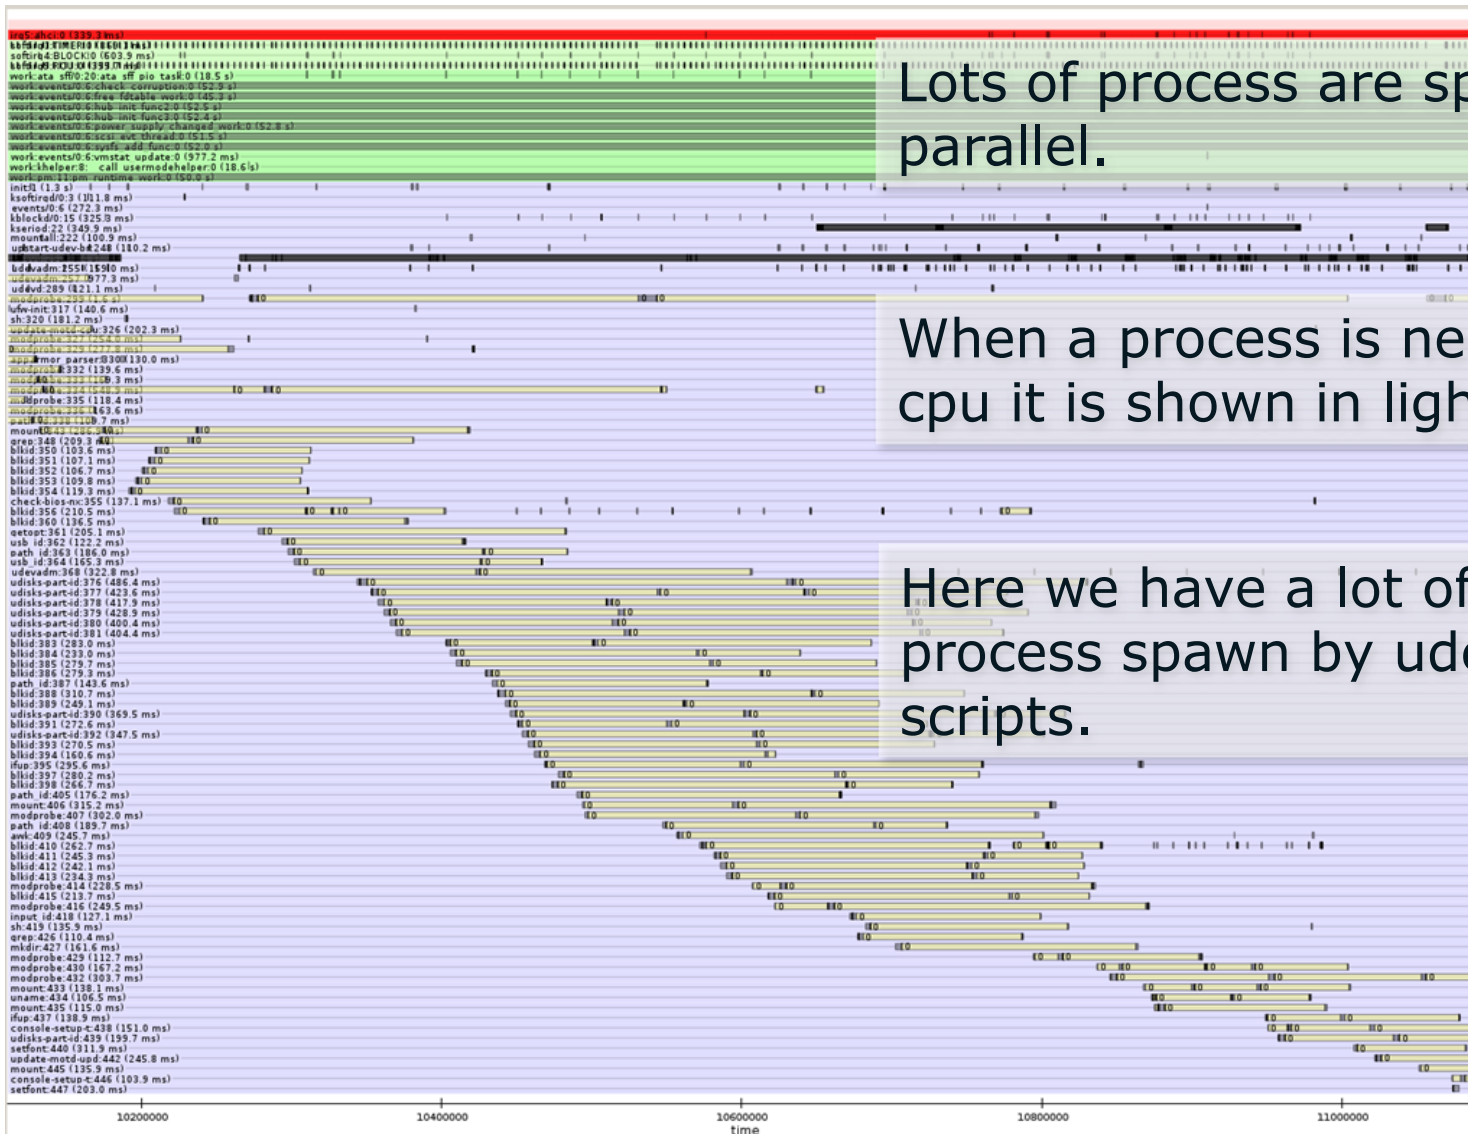




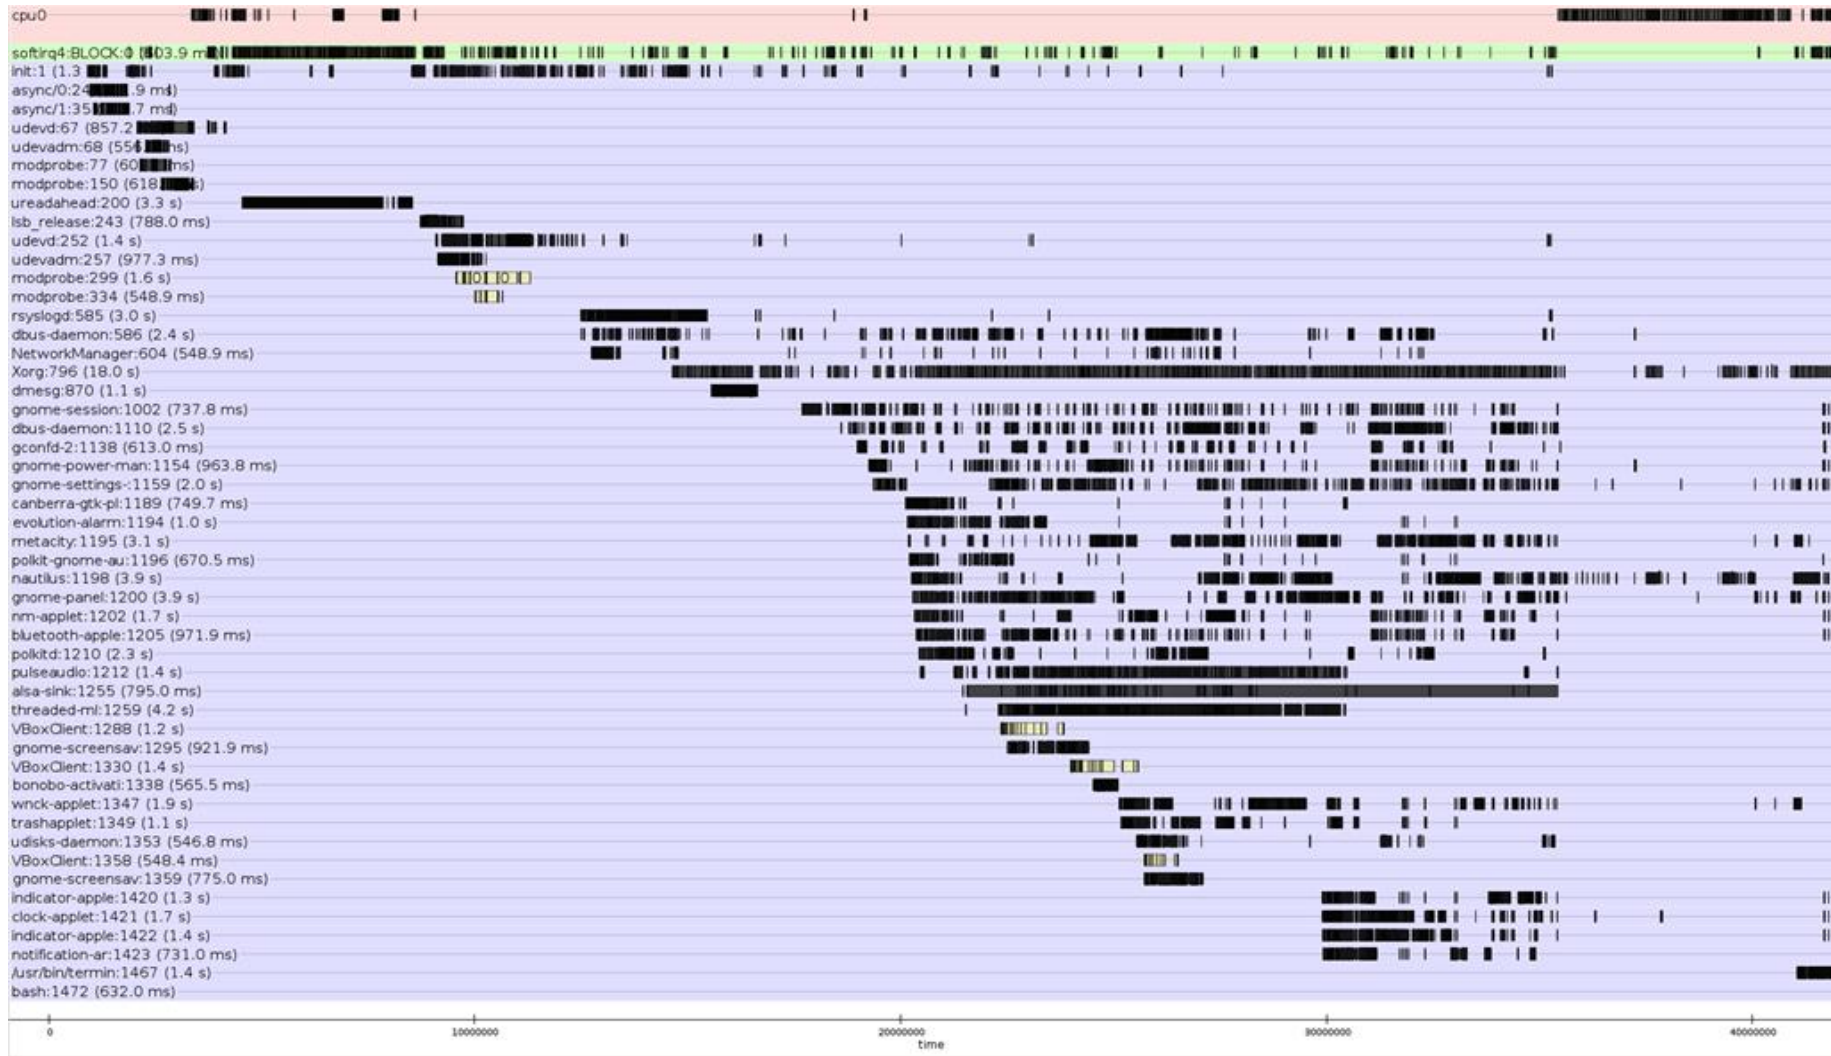
Lots of process are spawned in parallel.



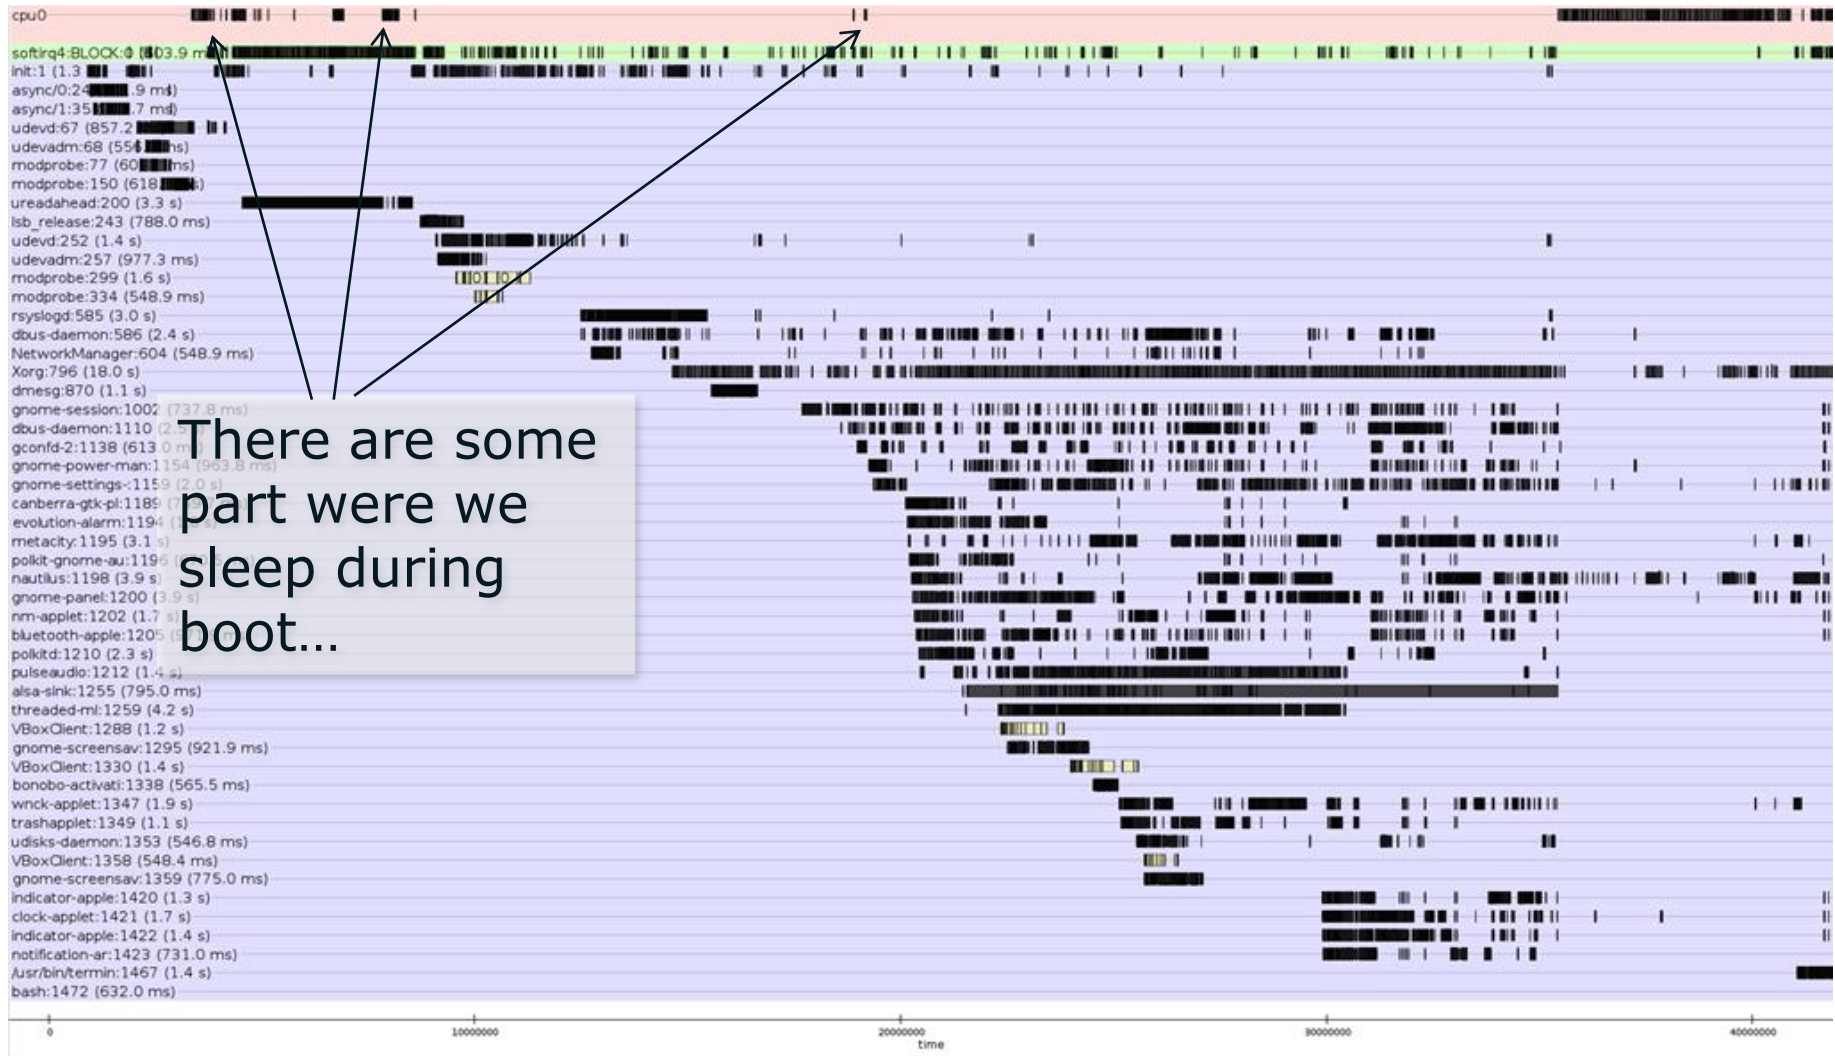
When a process is needing
cpu it is shown in light yellow.



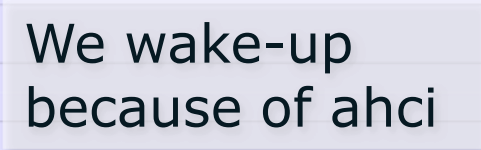
Filtering more process...



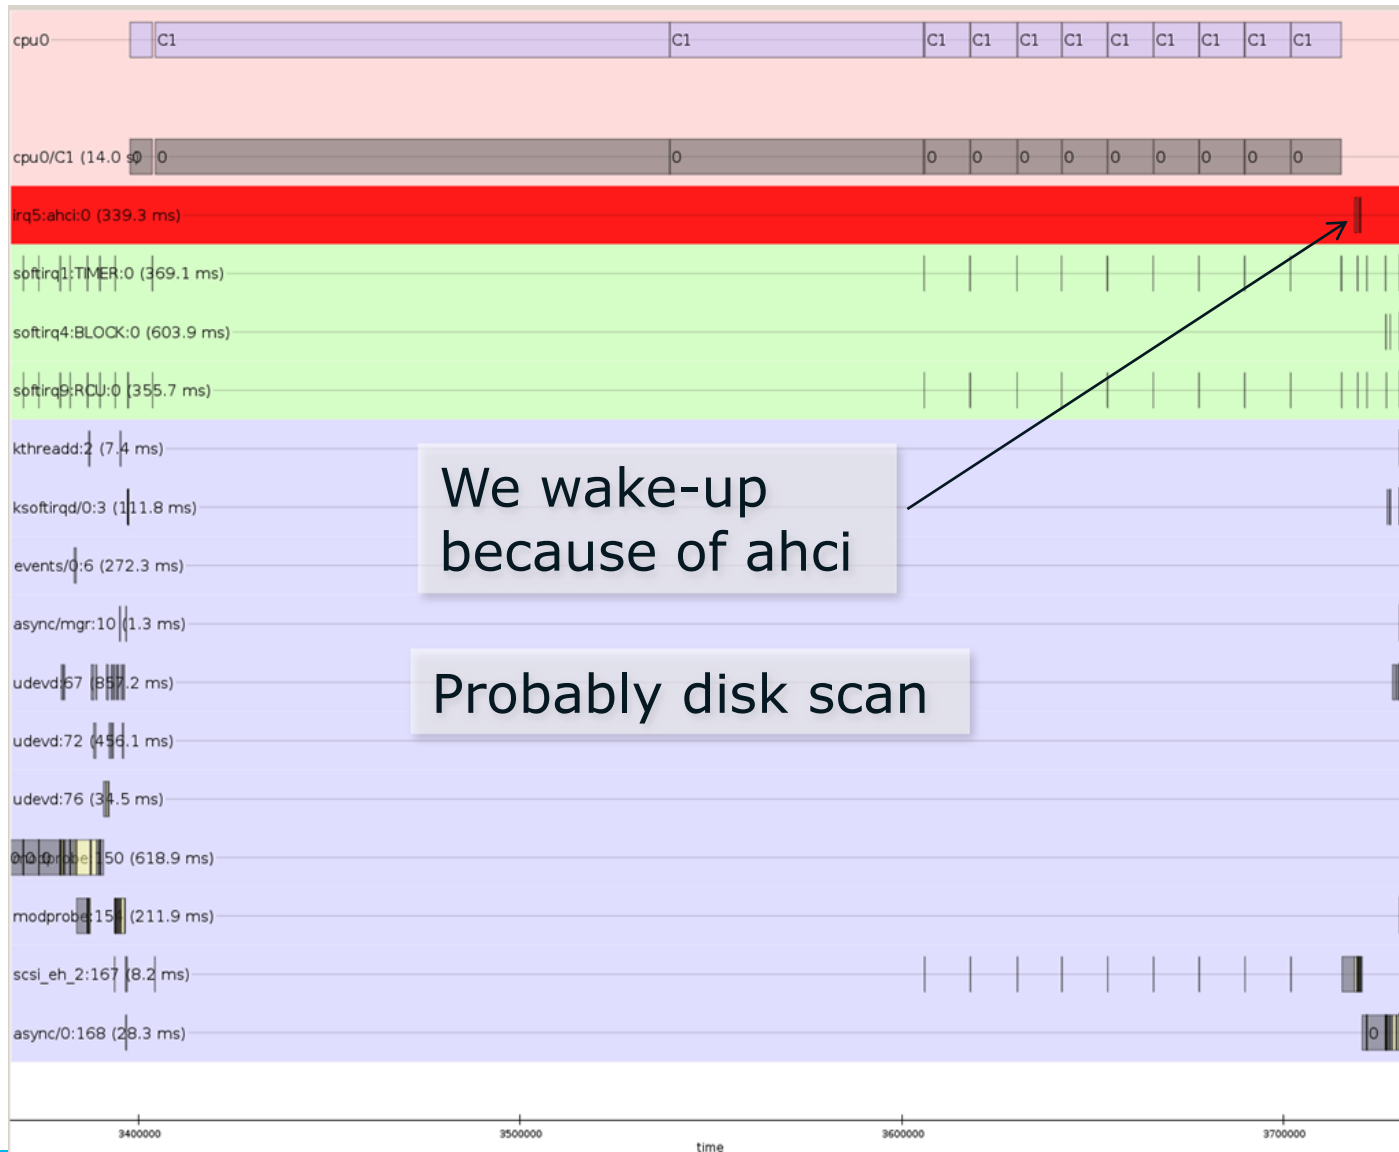
Filtering more process...



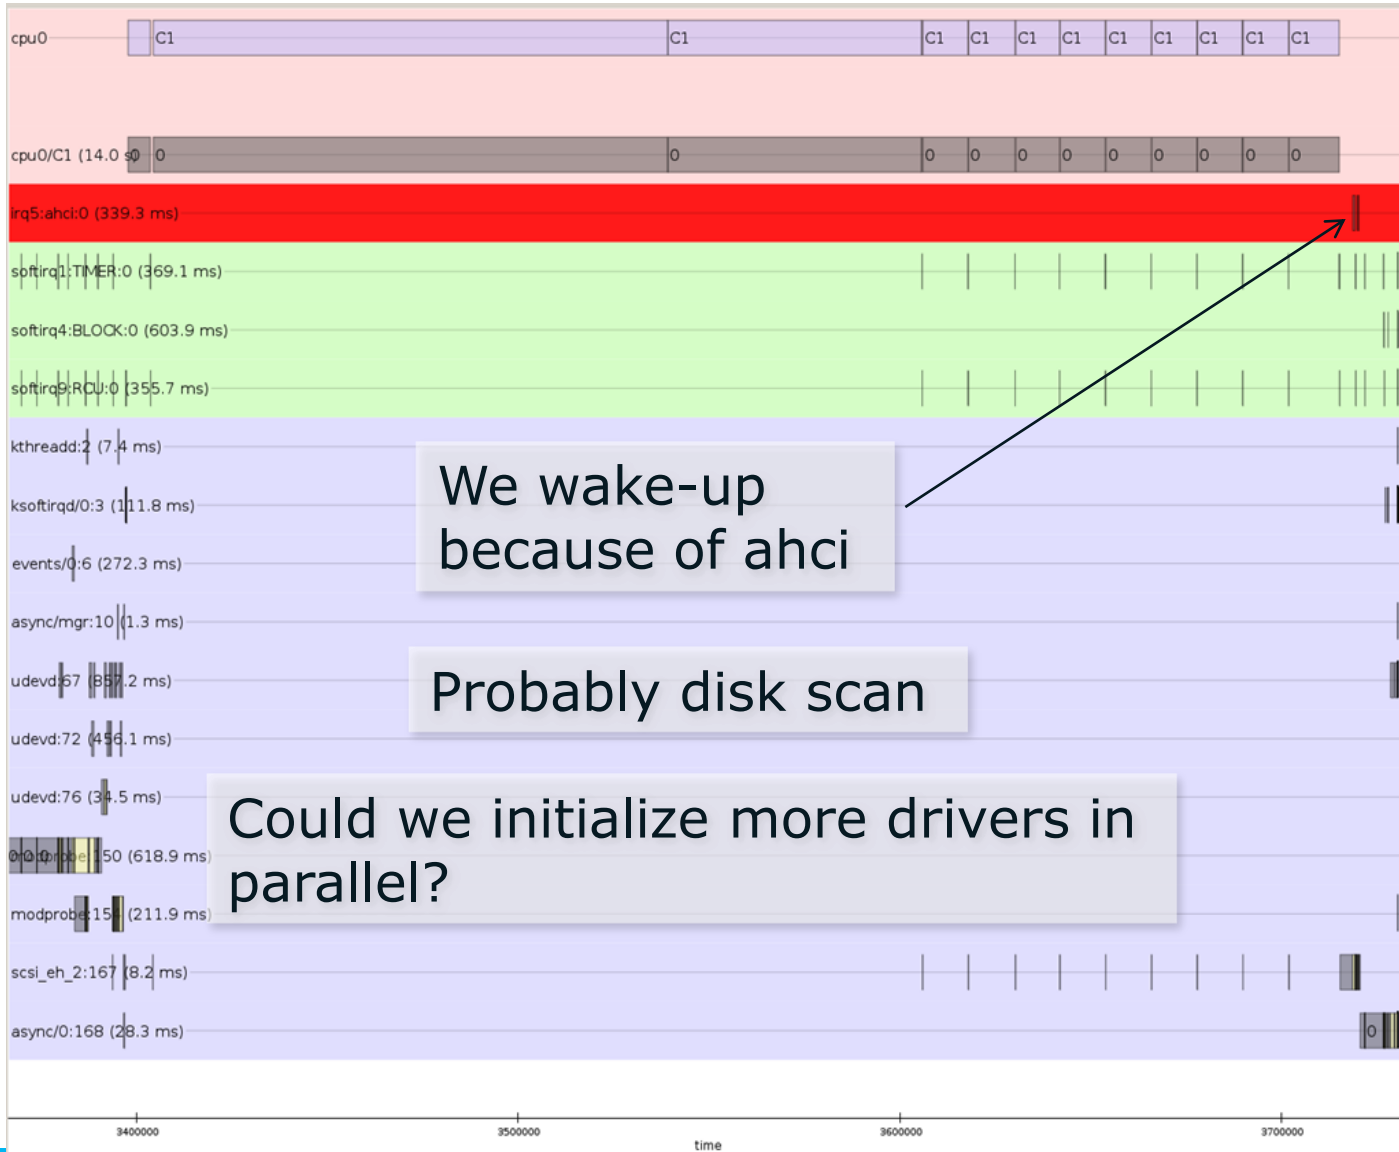
Zoom to first sleep



Zoom to first sleep



Zoom to first sleep



An asynchronous function call infrastructure

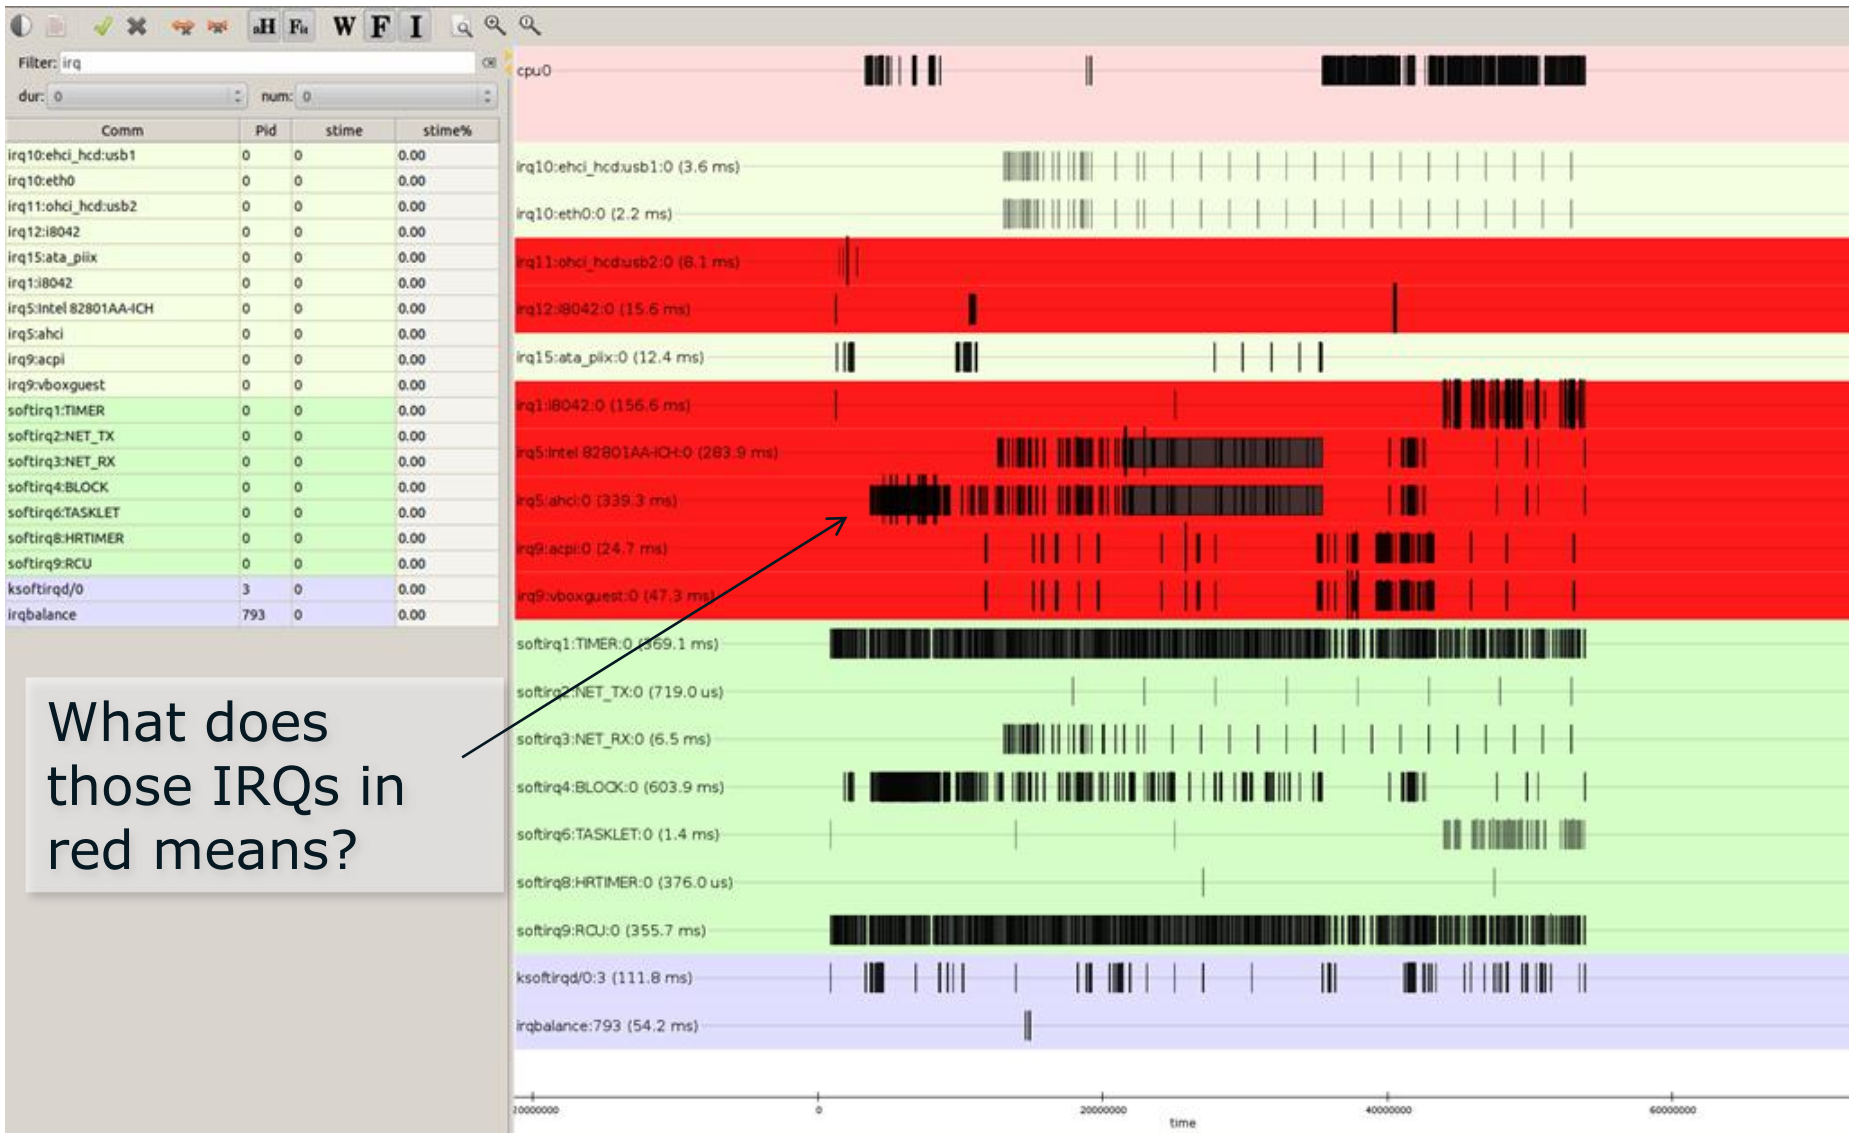
By **Jonathan Corbet**
January 13, 2009

Arjan van de Ven's [fast boot project](#) will be familiar to most LWN readers by now. Most of Arjan's work has not yet found its way into the mainline, though, so most of us still have to wait for our systems to boot the slow way. That said, the 2.6.29 kernel will contain one piece of the fast boot work, in the form of the asynchronous function call infrastructure. Users will need to know where to find it, though, before making use of it.

There are many aspects to the job of making a system boot quickly. Some of the lowest-hanging fruit can be found in the area of device probing. Figuring out what hardware exists on the system tends to be a slow task at best; if it involves physical actions (such as spinning up a disk) it gets even worse. Kernel developers have long understood that they could gain a lot of time if this device probing could, at least, be done in a parallel manner: while the kernel is waiting for one device to respond, it can be talking to another. Attempts at parallelizing this work over the years have foundered, though. Problems with device ordering, concurrent access, and more have adversely affected system stability, with the inevitable result that the parallel code is taken back out. So early system initialization remains almost entirely sequential.

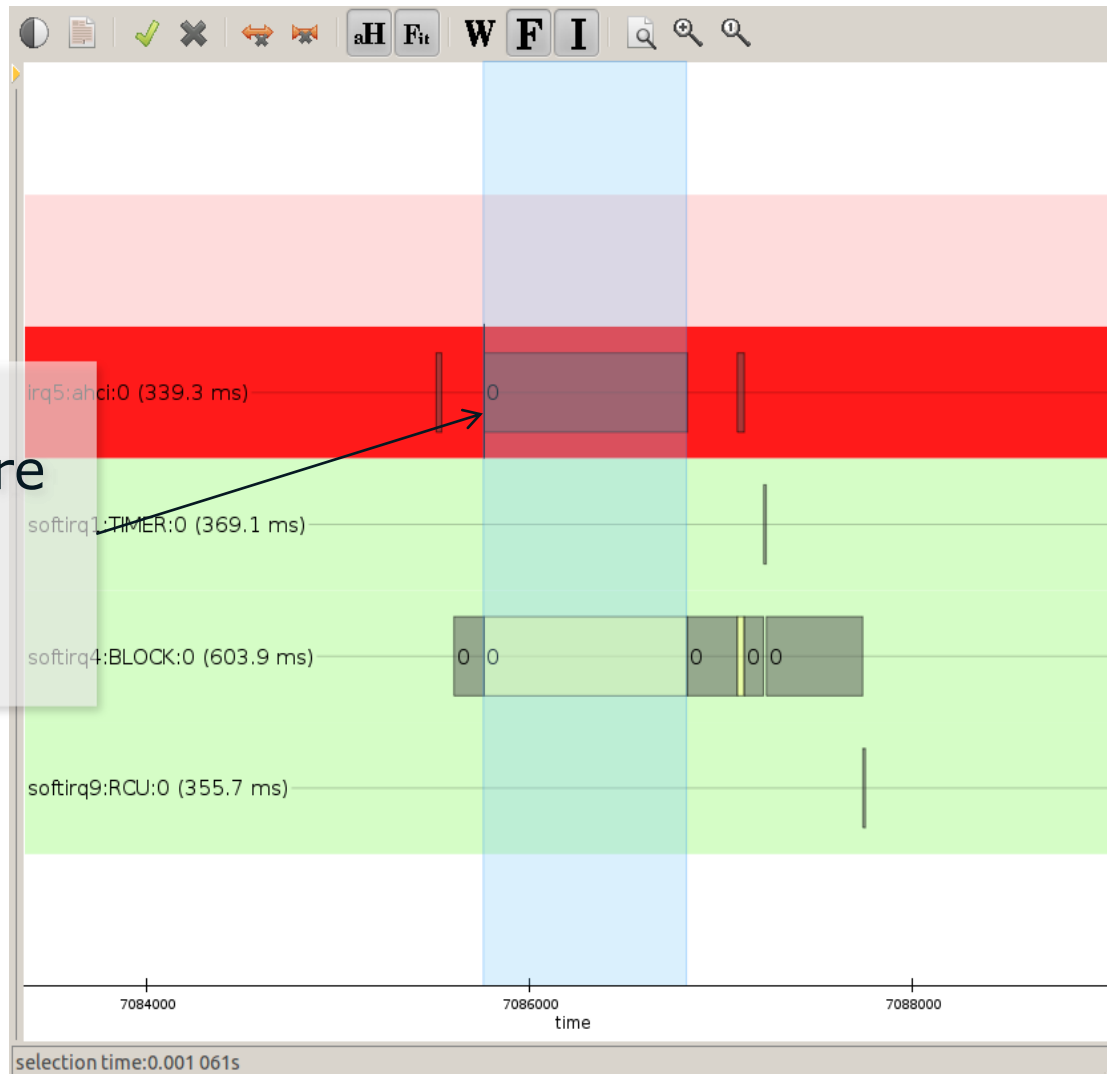
Although this is far to be new stuff, we don't see a lot of drivers using this, AFAIK.

Disgression on red irq



Disgression on red irq

IRQ handler
that lasts more
than 1ms
Bad for RT
latency!



Modem driver traces

Hacking PyTimechart to decode your own traces

```
from timechart.plugin import *
from timechart import colors
from timechart.model import tcProcess

class hsi(plugin):
    additional_colors = """
    hsi_bg                #80ff80
    """
    additional_ftrace_parsers = [
    ]
    additional_process_types = {
        "hsi":(tcProcess, MISC_TRACES_CLASS),
    }
    @staticmethod
    def do_function_ffl_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_ffl_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

    @staticmethod
    def do_function_hsi_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        pidcaller = proj.generic_find_process(event.common_pid,event.common_comm,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_hsi_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

plugin_register(hsi)
```

Hacking PyTimechart to decode your own traces

Pytimechart is written in python

- Easy to decode its own tracepoints
- Even if you don't add tracepoints, you can take advantage of function tracing to make more sense of your traces

The trace I took for this is a simple "receive SMS" trace

- I want to see when the hsi driver is in receive mode
- First, I trace hsi* and ffl* which are the low level, and protocol driver of our modem interface
- I can see 4 interesting functions in the trace:
 - hsi_start_rx()
 - hsi_stop_rx()
 - ffl_start_rx()
 - ffl_stop_rx()

timechart/plugins/template.py

Provided for convenience as a good starting point.

See doc for more detailed info

```
from timechart.plugin import *
from timechart import colors
from timechart.model import tcProcess

# to use with start_spi.sh
class template(plugin):
    additional_colors = ""
    template_bg = #80ff80
    """
    additional_fttrace_parsers = [
    ]
    additional_process_types = {
        "template":(tcProcess, MISC_TRACES_CLASS),
    }
    @staticmethod
    def do_function_my_start_function(proj,event):
        """This method will be called when the function "my_start_function" appears in the trace
        in this example, we start a process, and mark its caller as waked it
        """
        process = proj.generic_find_process(0,"template","template")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)

        # the function caller
        caller = proj.generic_find_process(0,event.caller,"template")
        proj.generic_add_wake(caller, process,event)

        # the calling process
        pidcaller = proj.generic_find_process(event.common_pid,event.common_comm,"hsi")
        proj.generic_add_wake(pidcaller, process,event)

    @staticmethod
    def do_function_my_stop_function(proj,event):
        """This method will be called when the function "my_stop_function" appears in the trace
        in this example, we stop the "template" process
        """
        process = proj.generic_find_process(0,"template","template")
        proj.generic_process_end(prev,event,build_p_stack=False)

# this plugin is disabled... uncomment to enable it.
#plugin_register(template)
```

timechart/plugins/hsi.py

```
sed s/template/hsi/g template.py > hsi.py
```

```
from timechart.plugin import *
from timechart import colors
from timechart.model import tcProcess

class hsi(plugin):
    additional_colors = """
hsi_bg                #80ff80
"""

    additional_fttrace_parsers = [
    ]
    additional_process_types = {
        "hsi":(tcProcess, MISC_TRACES_CLASS),
    }
    @staticmethod
    def do_function_ffl_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_ffl_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)
    @staticmethod
    def do_function_hsi_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        pidcaller = proj.generic_find_process(event.common_pid,event.common_comm,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_hsi_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

plugin_register(hsi)
```

timechart/plugins/hsi.py

sed s/template/hsi/g template.py > hsi.py

I want that the start functions begins a process event

```
from timechart.plugin import *
from timechart import colors
from timechart.model import tcProcess

class hsi(plugin):
    additional_colors = """
    hsi_bg                #80ff80
    """
    additional_fttrace_parsers = [
    ]
    additional_process_types = {
        "hsi":(tcProcess, MISC_TRACES_CLASS),
    }
    @staticmethod
    def do_function_ffl_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_ffl_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

    @staticmethod
    def do_function_hsi_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        pidcaller = proj.generic_find_process(event.common_pid,event.common_comm,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_hsi_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

plugin_register(hsi)
```

timechart/plugins/hsi.py

sed s/template/hsi/g template.py > hsi.py

I want that the start functions begins a process event

and the stop function ends it

```
from timechart.plugin import *
from timechart import colors
from timechart.model import tcProcess

class hsi(plugin):
    additional_colors = """
    hsi_bg                #80ff80
    """
    additional_fttrace_parsers = [
    ]
    additional_process_types = {
        "hsi":(tcProcess, MISC_TRACES_CLASS),
    }
    @staticmethod
    def do_function_ffl_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_ffl_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

    @staticmethod
    def do_function_hsi_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        pidcaller = proj.generic_find_process(event.common_pid,event.common_comm,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_hsi_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

plugin_register(hsi)
```

timechart/plugins/hsi.py

sed s/template/hsi/g template.py > hsi.py

I want that the start functions begins a process event

and the stop function ends it

I also add a wake event to see where I come from

```
from timechart.plugin import *
from timechart import colors
from timechart.model import tcProcess

class hsi(plugin):
    additional_colors = """
    hsi_bg                #80ff80
    """
    additional_fttrace_parsers = [
    ]
    additional_process_types = {
        "hsi":(tcProcess, MISC_TRACES_CLASS),
    }
    @staticmethod
    def do_function_ffl_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

    @staticmethod
    def do_function_ffl_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsiffl","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

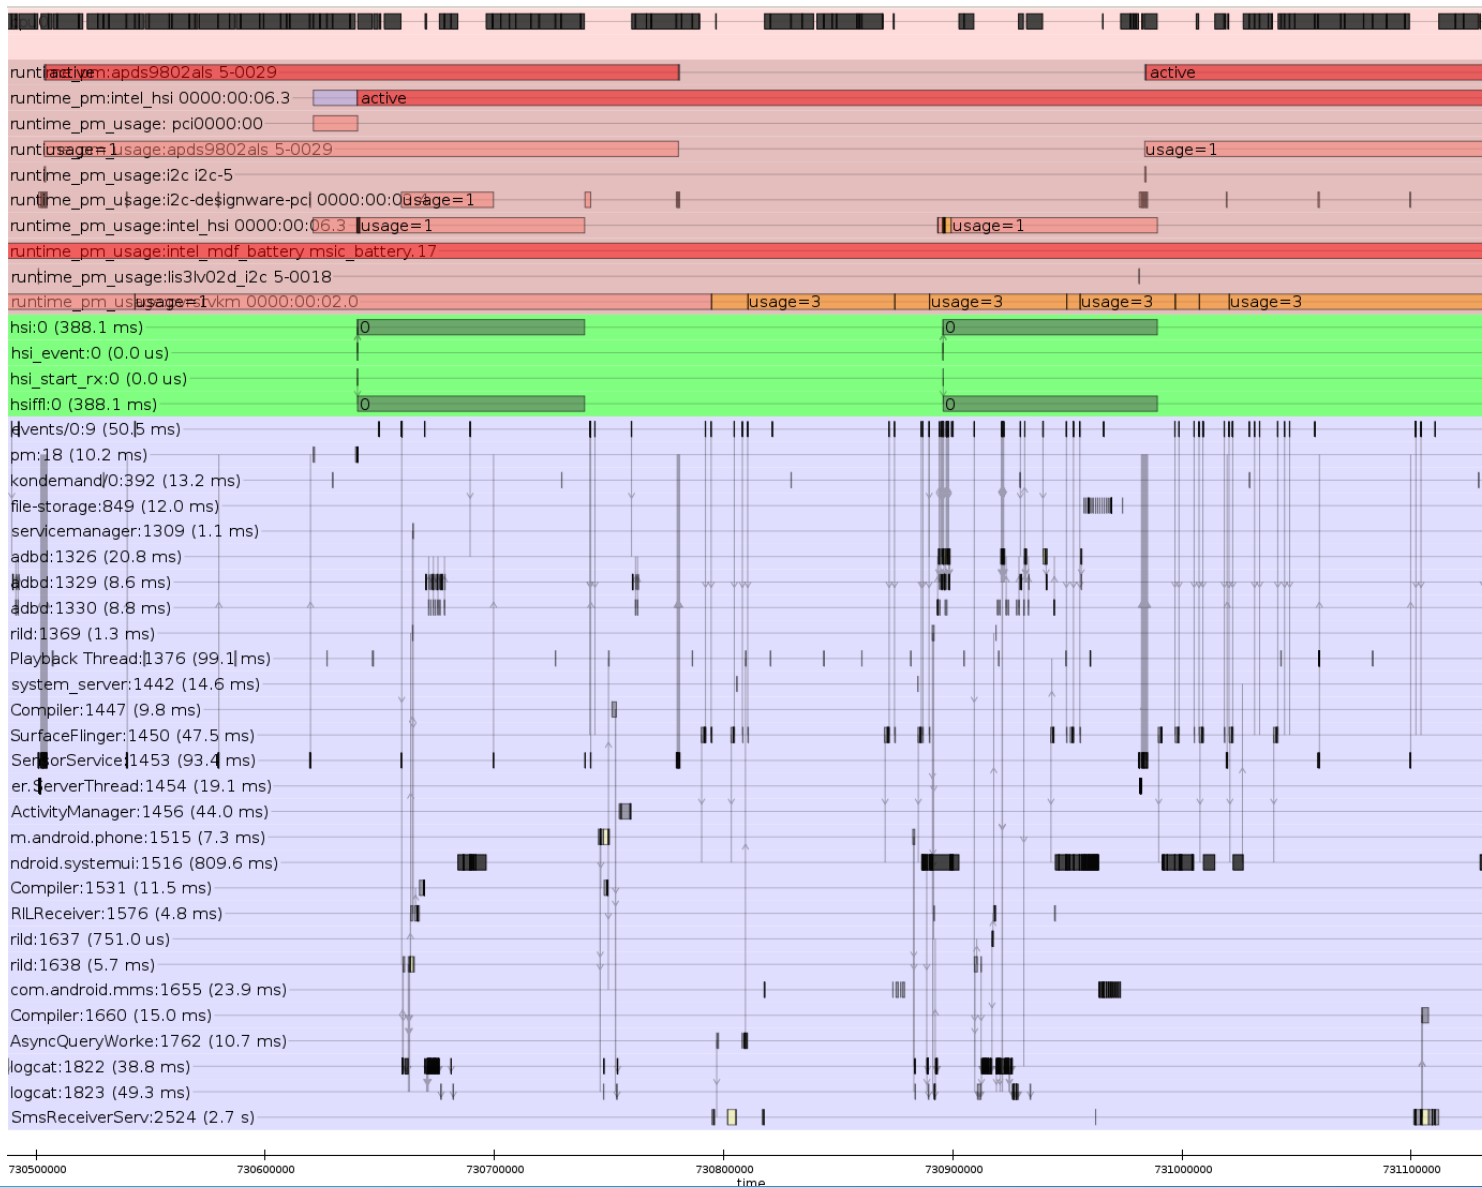
    @staticmethod
    def do_function_hsi_start_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        caller = proj.generic_find_process(0,event.caller,"hsi")
        pidcaller = proj.generic_find_process(event.common_pid,event.common_comm,"hsi")
        proj.generic_process_start(process,event,build_p_stack=False)

        proj.generic_process_single_event(caller,event)
        proj.generic_add_wake(caller, process,event)

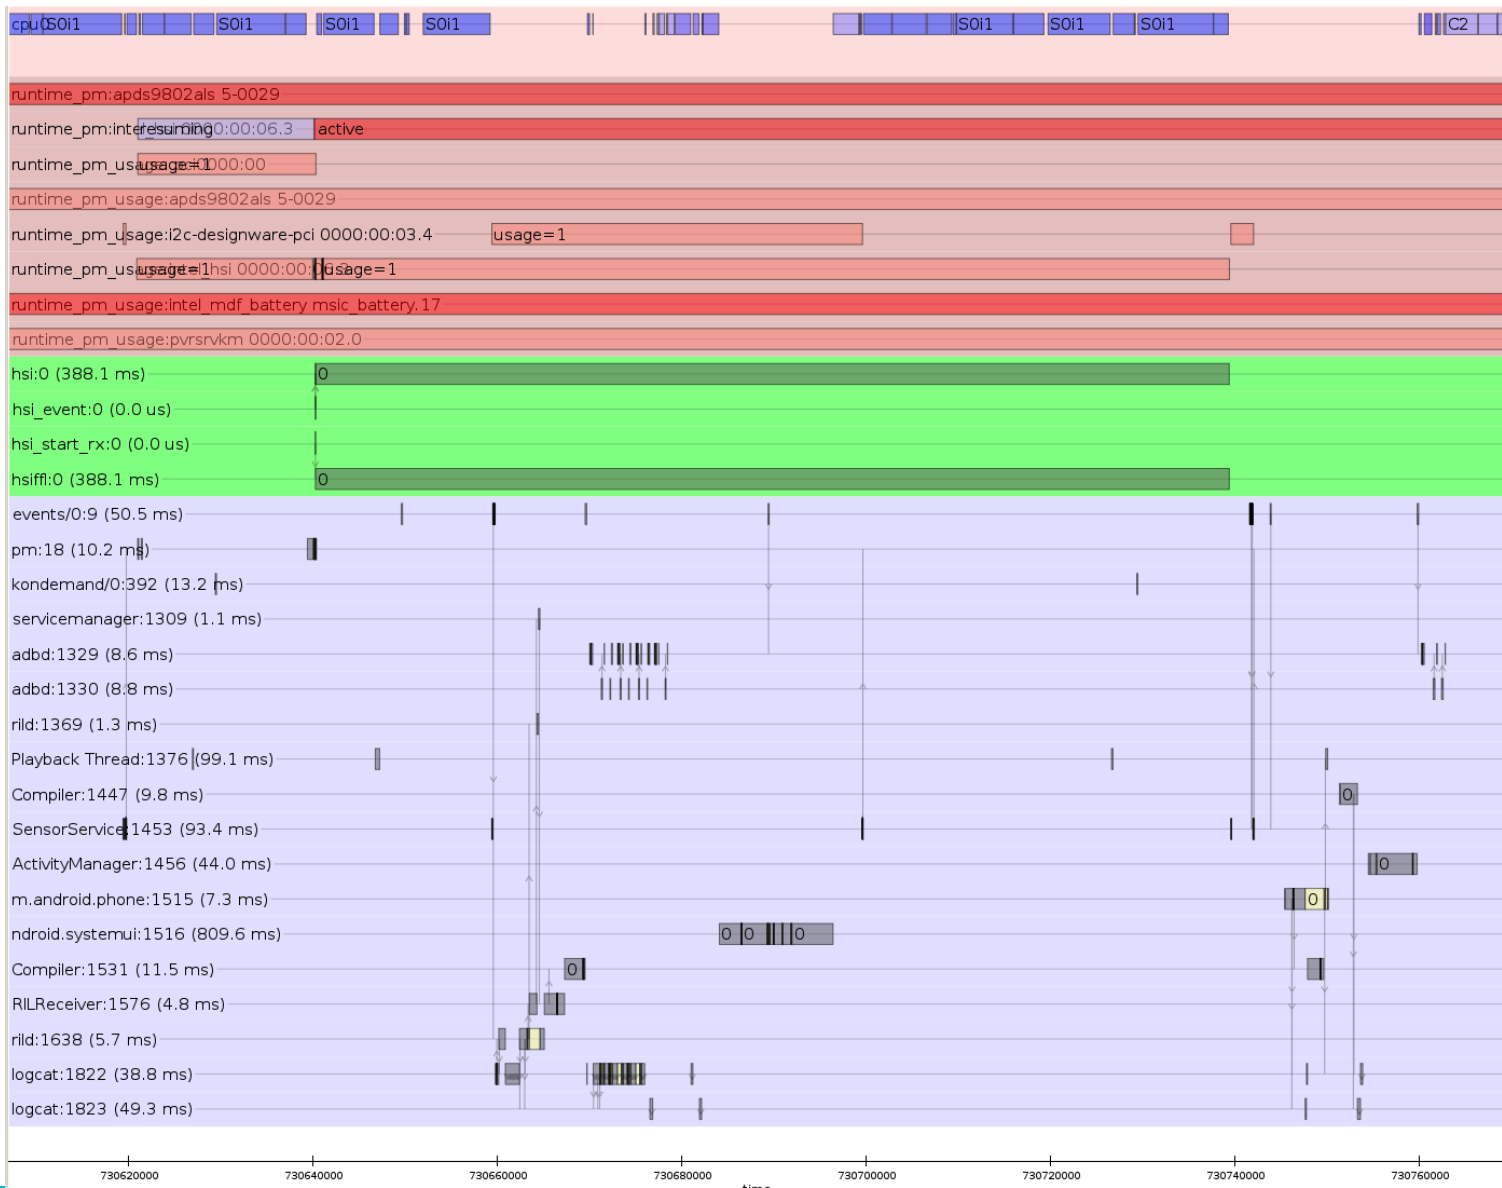
    @staticmethod
    def do_function_hsi_stop_rx(proj,event):
        process = proj.generic_find_process(0,"hsi","hsi")
        proj.generic_process_end(process,event,build_p_stack=False)

plugin_register(hsi)
```

Results



Results Zoomed





Questions?



Thank You

