



Scheduler Load tracking updates and improvements

Vincent Guittot
OSPM 3rd April 2017

LEADING
COLLABORATION
IN THE ARM
ECOSYSTEM



Agenda

- Status
- Remaining problems
- What's next ?
- Discussion

Status



Load tracking in scheduler

- CFS class
 - PELT geometric serie
- RT class
 - `rt_avg`
- Deadline class
 - Per rq active utilization

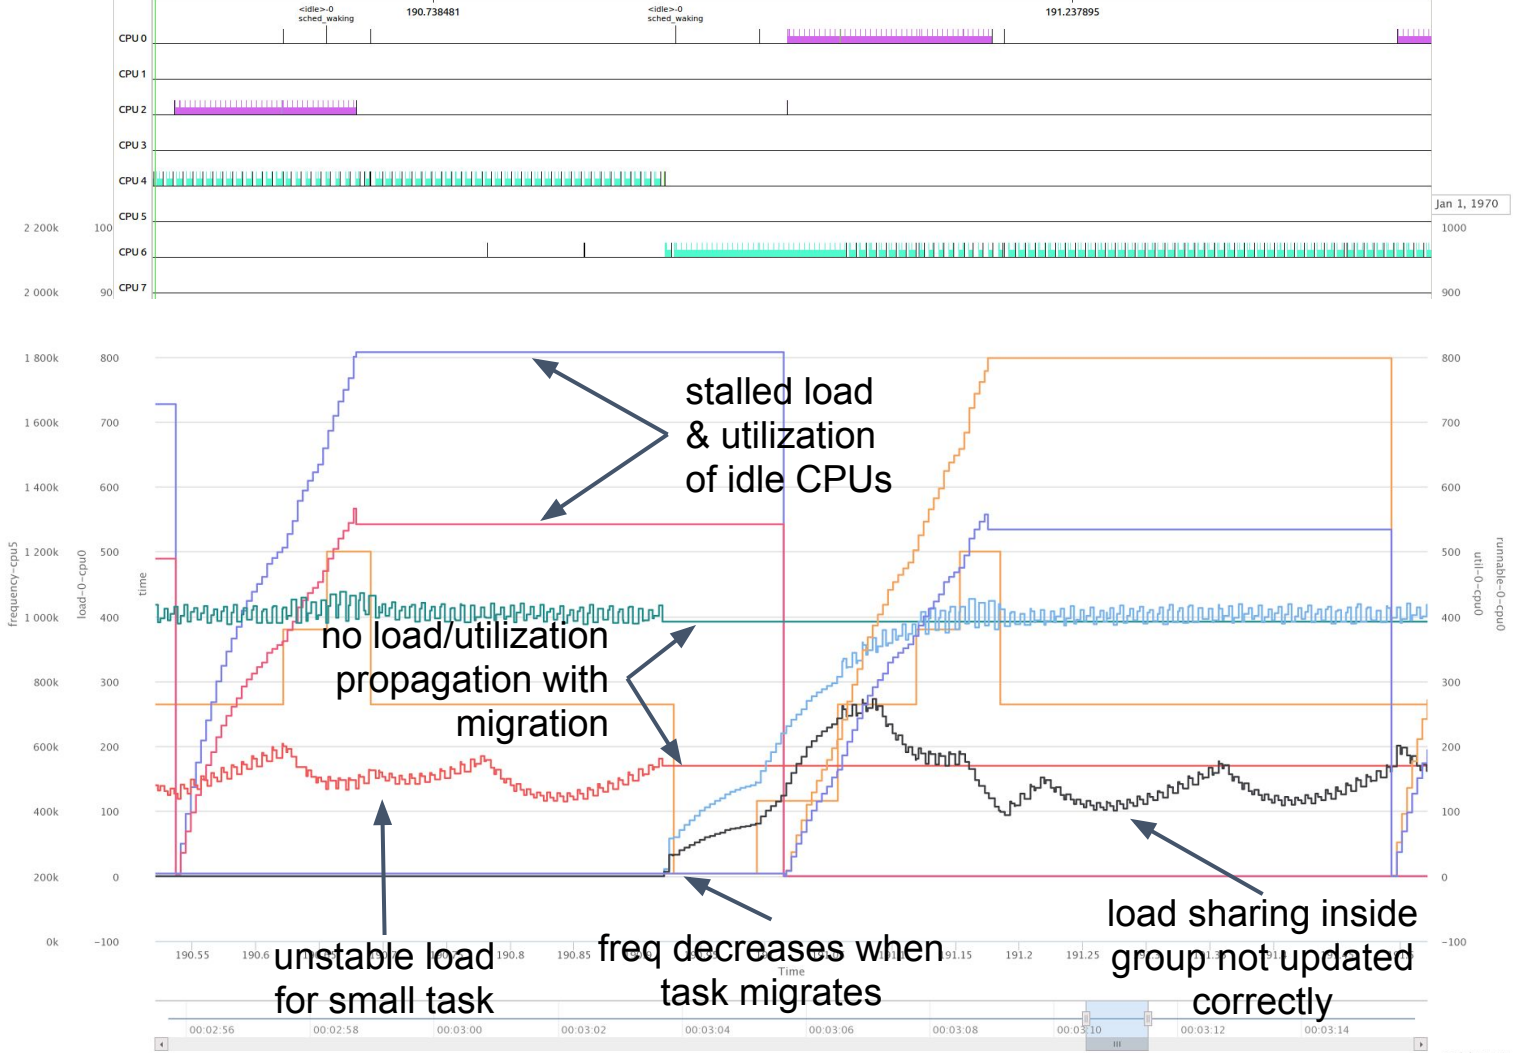


Load tracking in scheduler

- Most of the focus has been on CFS and PELT
- RT load tracking
 - `rt_avg` is not a good metric
- Deadline utilization
 - See Juri's patchset

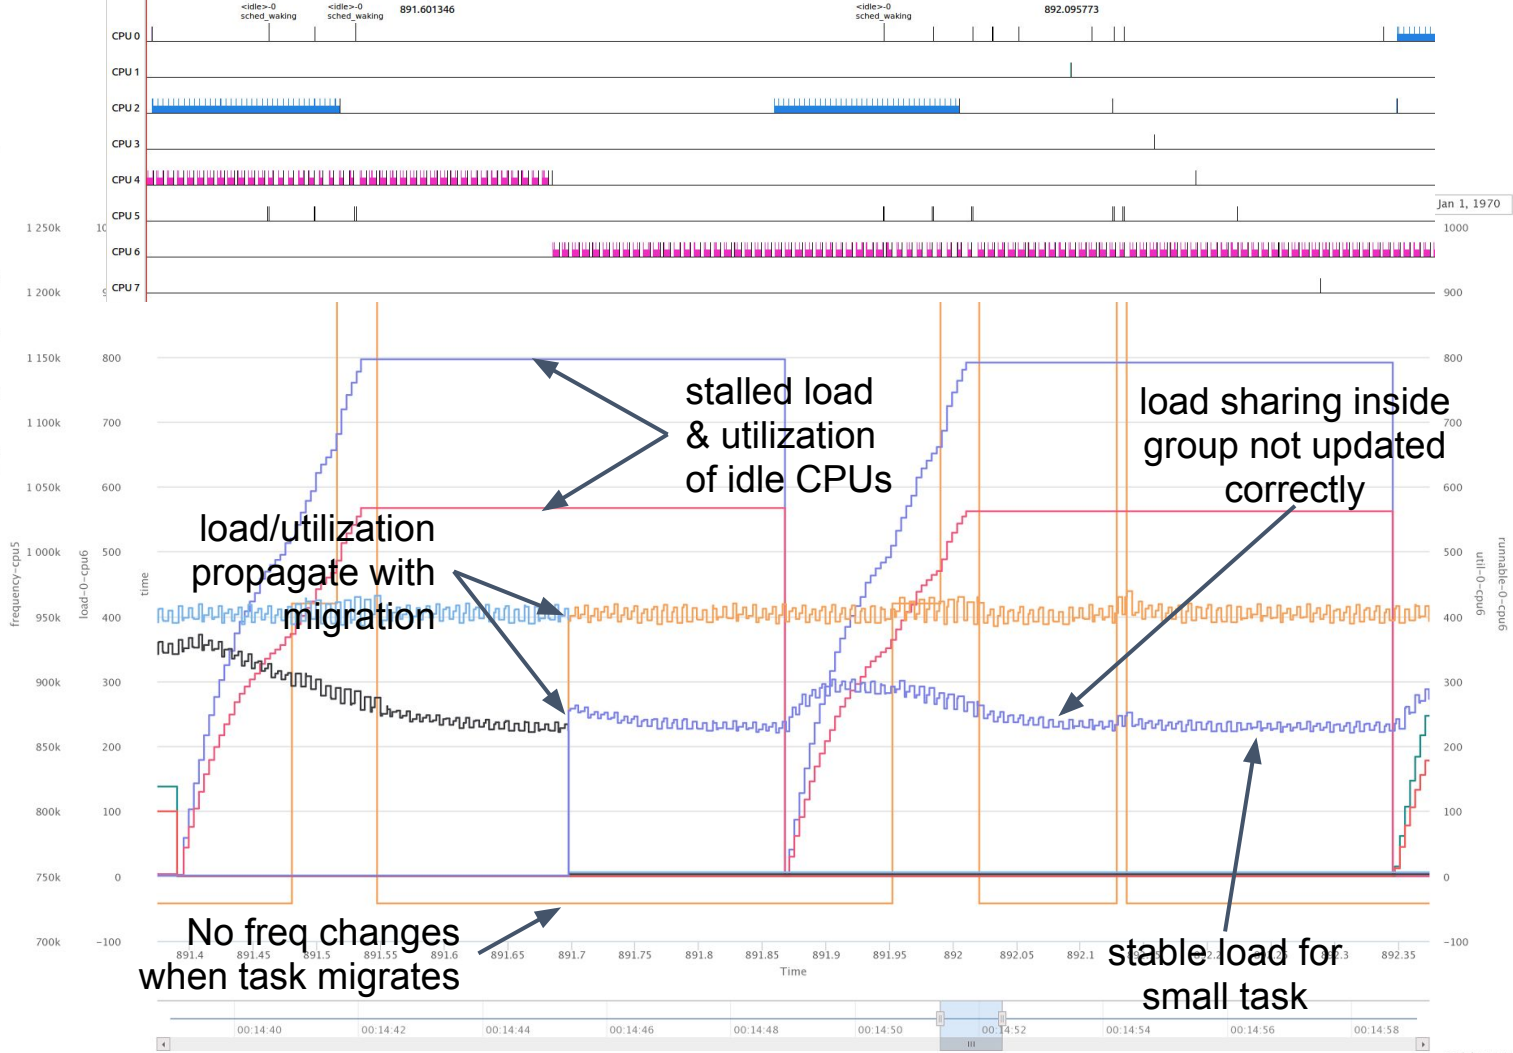
v4.9

- 0
- 0
- 0
- frequency-cpu5
- util-0-cpu2
- load-0-cpu2
- runnable-0-cpu2
- util-0-cpu4
- load-0-cpu4
- runnable-0-cpu4
- util-0-cpu6
- load-0-cpu6
- runnable-0-cpu6
- util-0-cpu0
- load-0-cpu0
- runnable-0-cpu0



tip

- 0
- 0
- 0
- frequency-cpu5
- util-0-cpu2
- load-0-cpu2
- runnable-0-cpu2
- util-0-cpu0
- load-0-cpu0
- runnable-0-cpu0
- util-0-cpu4
- load-0-cpu4
- runnable-0-cpu4
- util-0-cpu6
- load-0-cpu6
- runnable-0-cpu6



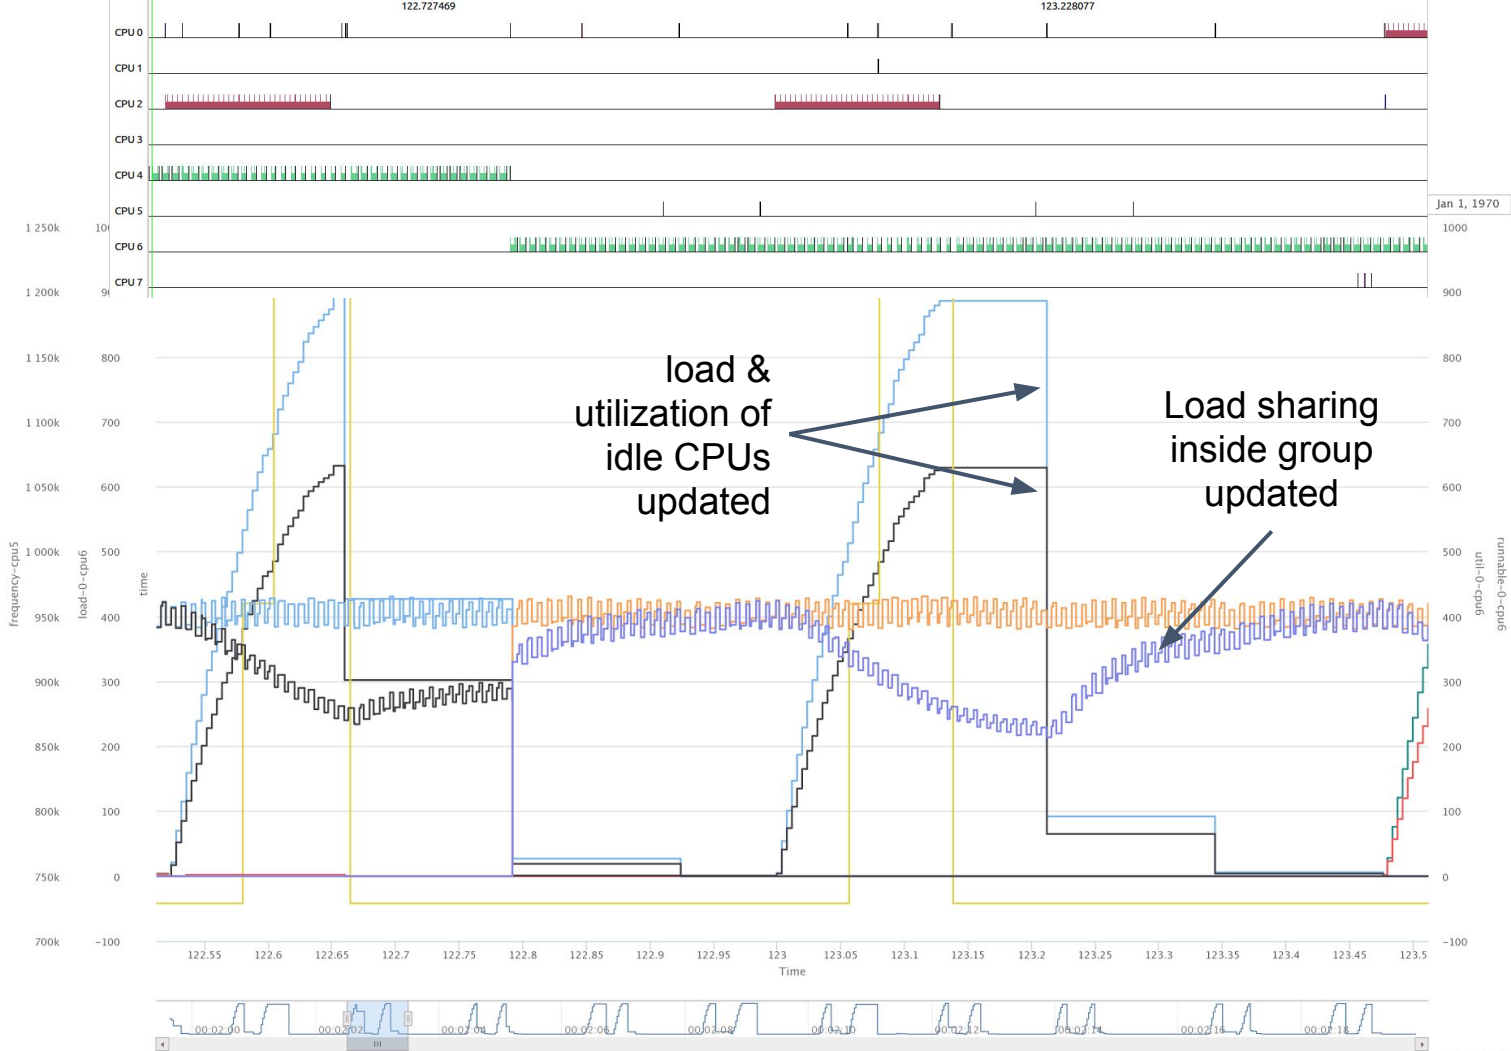

No freq changes when task migrates

stalled load & utilization of idle CPUs

load sharing inside group not updated correctly

stable load for small task

Target

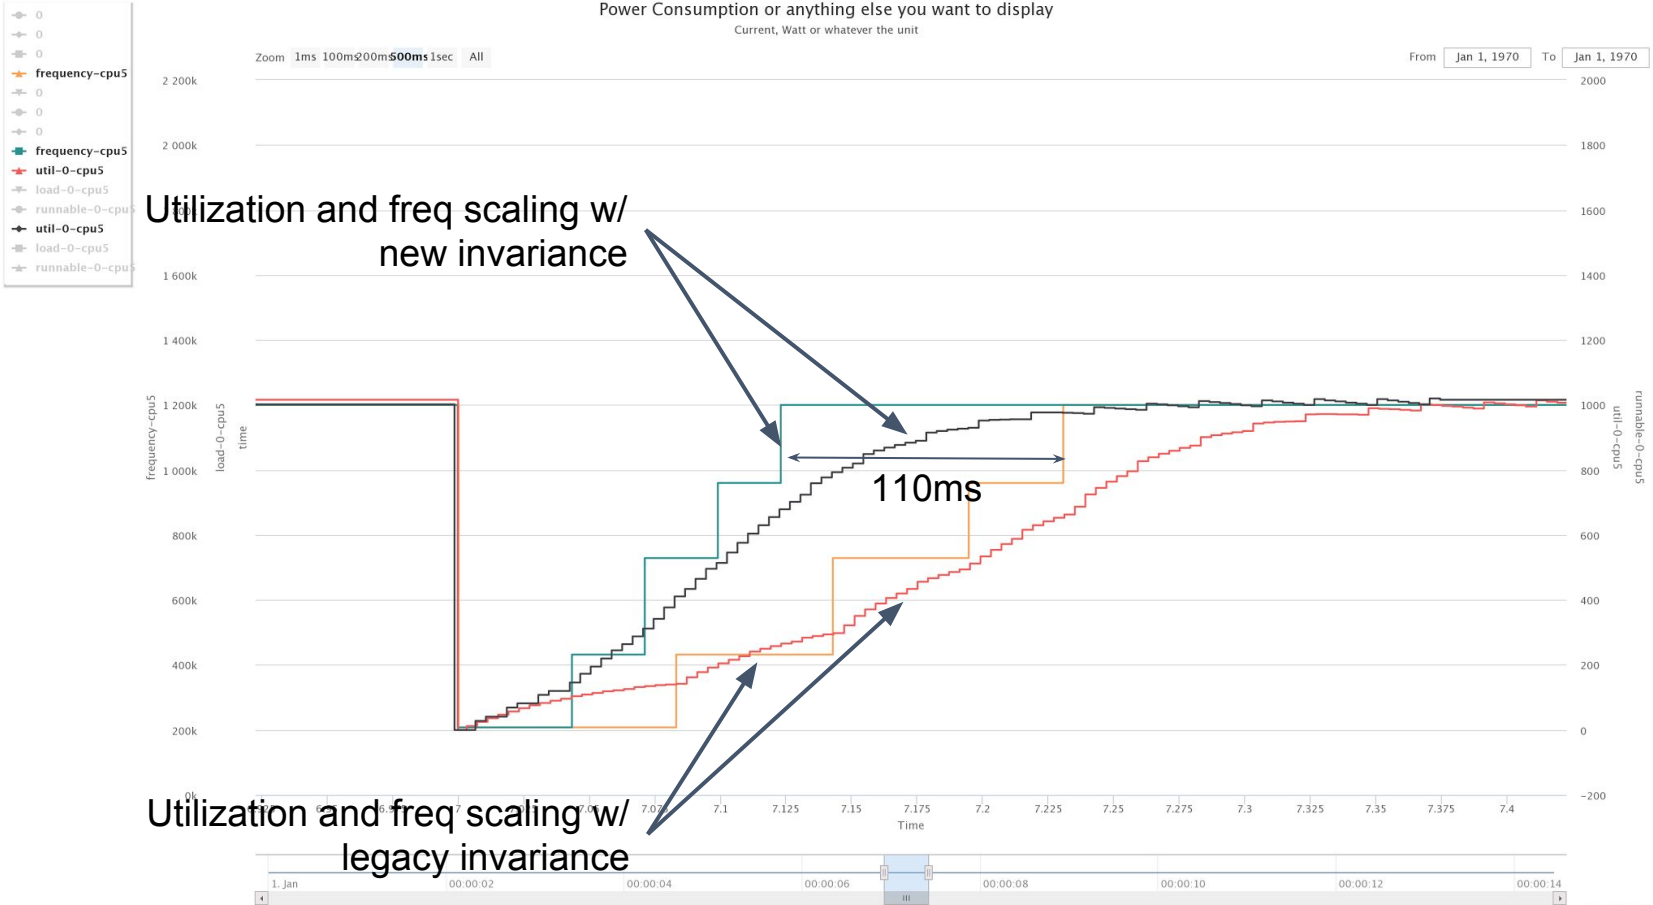


A decorative background on the left side of the slide. It features a glowing green globe with a white grid overlay, showing the Earth's continents. Above the globe, there are several floating icons: a smartphone, a circular target-like icon, and a glowing green arrow pointing right. The overall color scheme is dark green and black with bright green highlights.

Remaining points

- Frequency invariance
- Update of blocked idle load
- Utilization drop on RT/DL preemption
- Utilization drop on migration

Frequency invariance



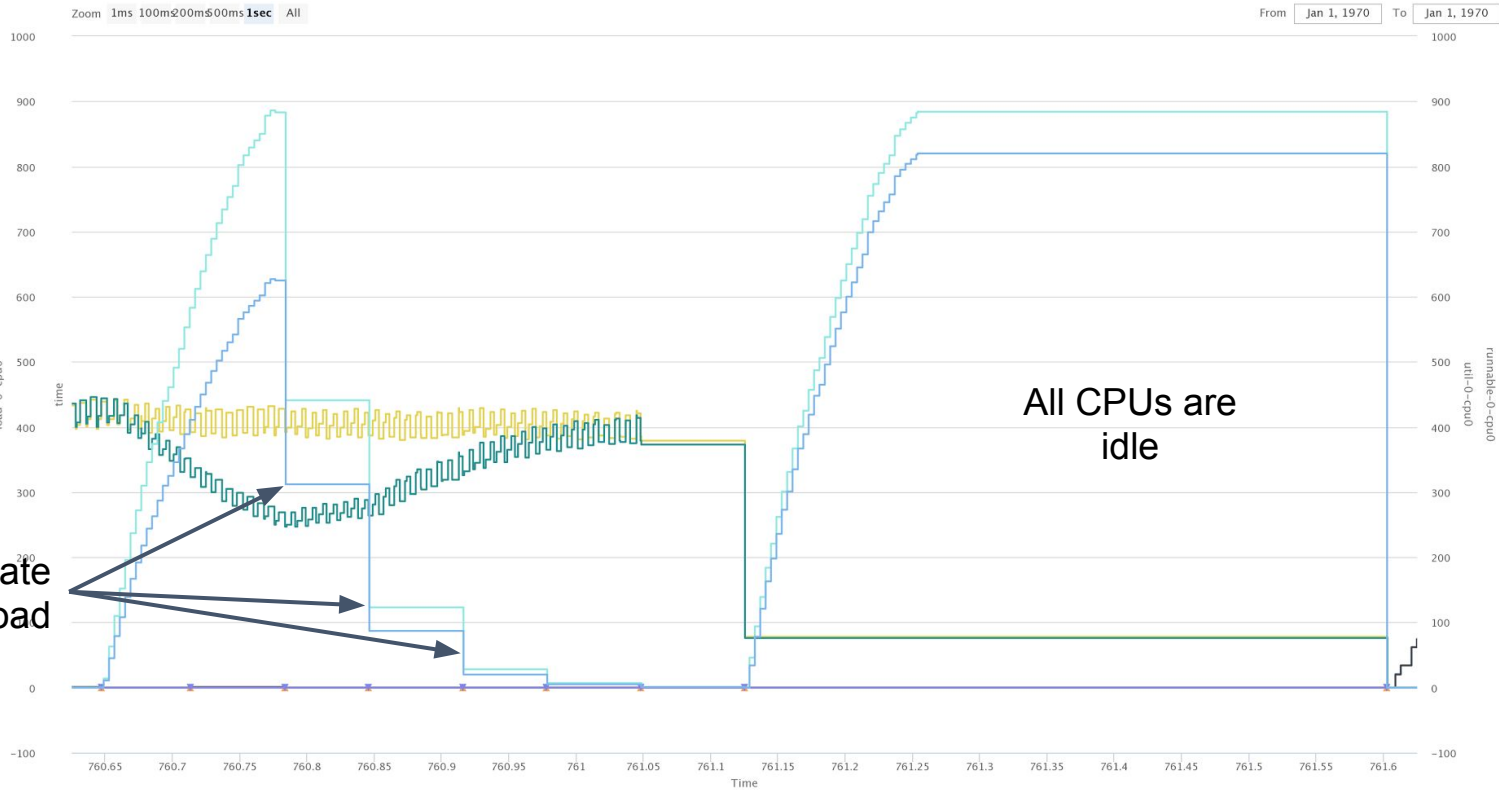
Frequency invariance

- Faster increase
- Same range of variation across freq and micro arch
- Same behavior for Load and utilization

Updating blocked idle load

- util-0-cpu2
- load-0-cpu2
- runnable-0-cpu2
- util-0-cpu4
- load-0-cpu4
- runnable-0-cpu4
- util-0-cpu6
- load-0-cpu6
- runnable-0-cpu6
- util-0-cpu0
- load-0-cpu0
- runnable-0-cpu0

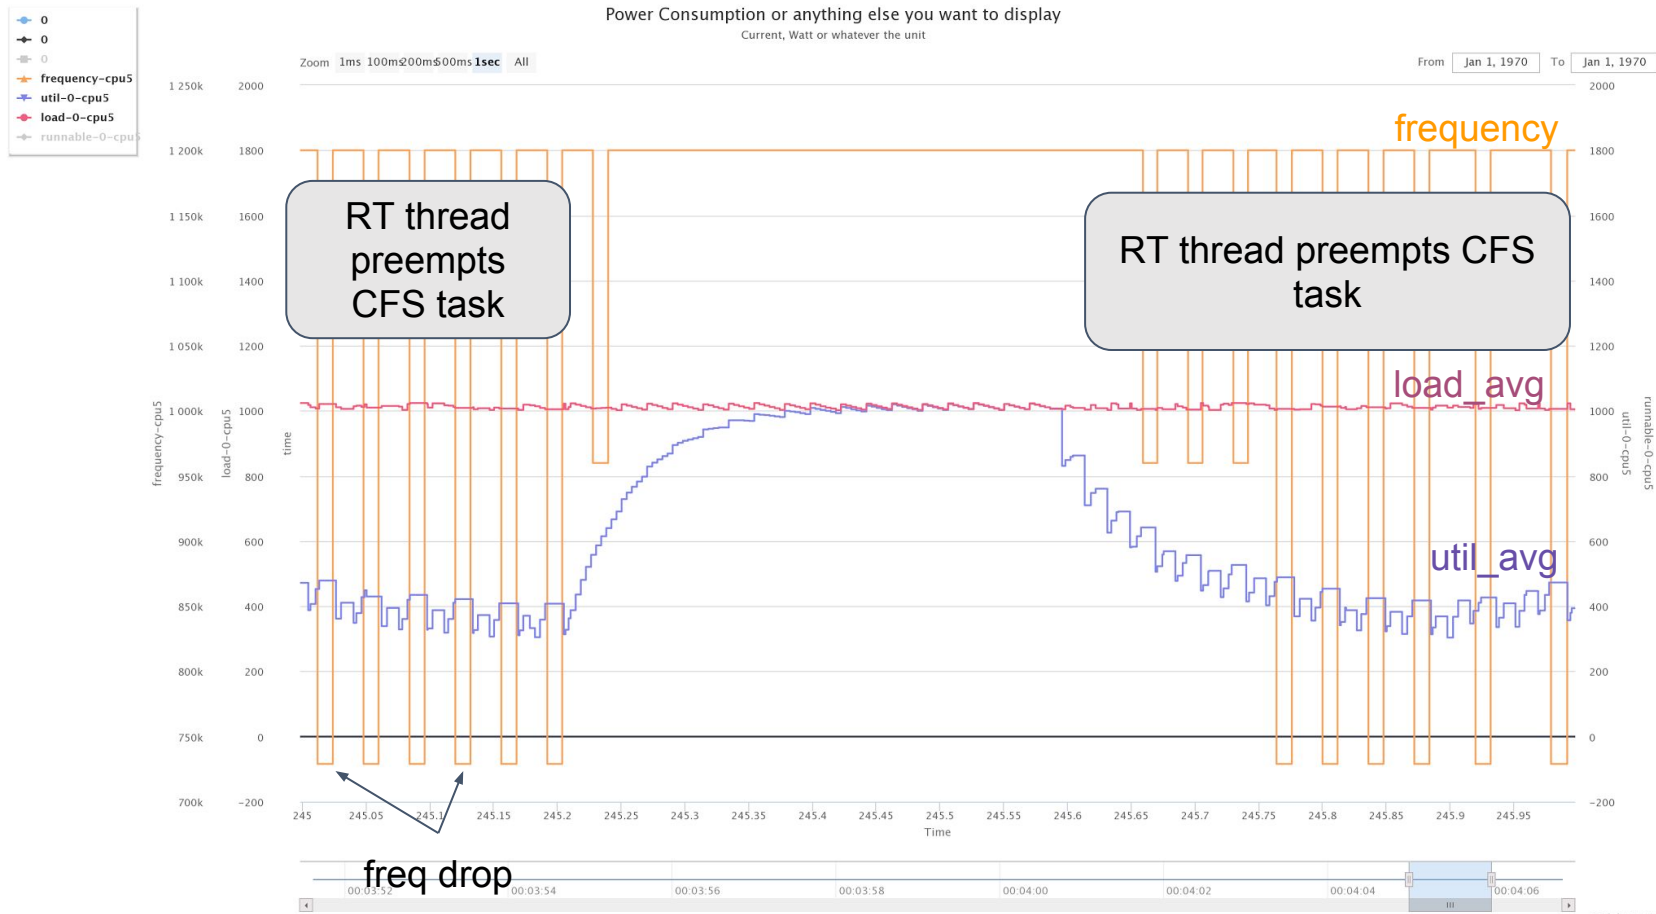
Power Consumption or anything else you want to display
Current, Watt or whatever the unit



Updating blocked idle load

- Periodic load balance updates blocked load periodically
 - Only if load balance is needed
- Need to update block load more frequently
 - Load is used to set share in task group
 - Utilization is used to set OPP (w/ schedutil)
- ILB already update blocked load periodically
 - Kick it more frequently but only to update blocked load
- Kick triggered in `select_task_rq_fair`
 - If task pinned on 1 CPU, don't kick. Should we ?
 - If RT/DL task don't kick. Should we ?
- If system fully idle, don't wake up
- But not enough
 - When a task wake up, it's already too late


utilization drop on RT/DL preemption



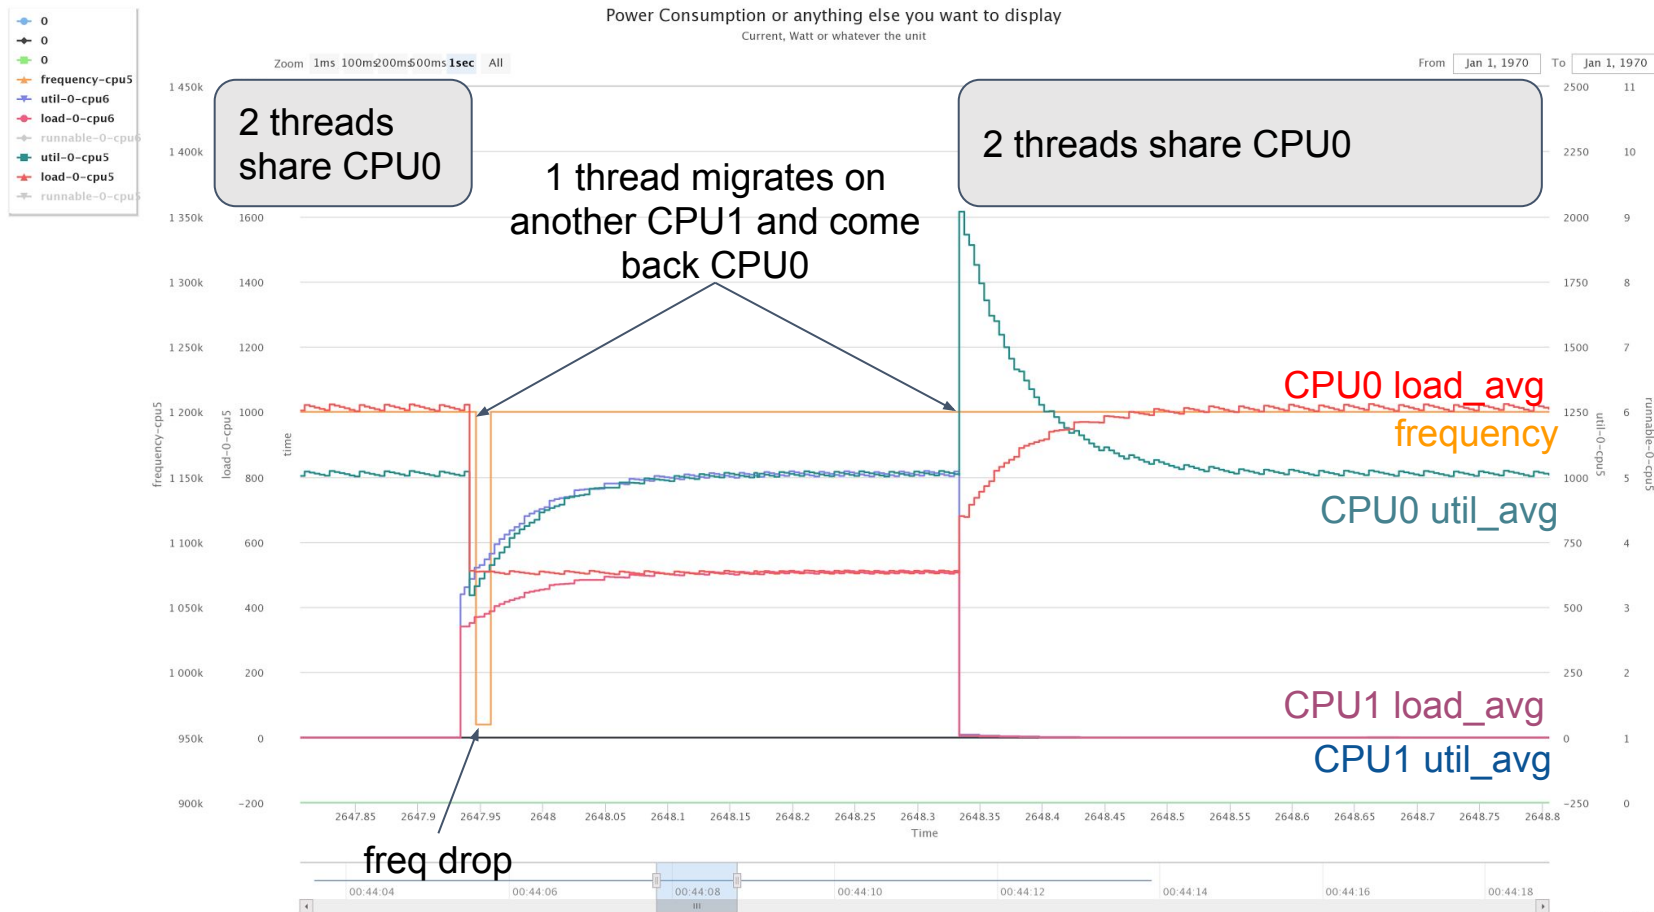
Prototype RT utilization tracking

- Add a PELT like utilization tracking for `rt_rq`
 - Only for root `rt_rq`
- Update in 3 places in RT class
 - `pick_next_task_rt`
 - `put_prev_task_rt`
 - `task_tick_rt`
- And 1 place in CFS
 - `update_blocked_averages`
- Add CFS+ RT (+ DL) utilization to estimate current utilization
- Must ensure a periodic update of blocked load

utilization drop on RT/DL preemption



utilization drop on migration



Evaluate waiting time

- How to evaluate that a thread doesn't have all the running time it wants ?
- Not overloaded CPU
 - Can use `util_avg` without issue
- Overloaded CPU
 - `Util_avg` is not enough
 - `Util_avg` doesn't reflect usage of entity
- Use runnable time ?

Others ?



Thank You

For further information: www.linaro.org

This is a Standard Content Slide - Enter Title Here

- Enter bullet points here
- Add images on the right or below

Alternative Content Slide

- If you have a more succinct message, use this slide






This Slide is for when Two Columns are Needed

- You can use this for two columns of bullets
- Or you can replace one column with an image or diagram
- This could be a second column of bullets
- Or it could be a table, image or graphic

Use this Slide for a Large Image

Useful Logos



Download Hi Res logos from [here](http://link.linaro.org/logos)* to use on your slides

*<http://link.linaro.org/logos>



LEADING COLLABORATION
IN THE ARM ECOSYSTEM

Some Images

- Download Hi Res photos from [here](http://link.linaro.org/groupphotos)* to use on your slides

* <http://link.linaro.org/groupphotos>

