



Which timer infra for runtime pm ?

Vincent Guittot

LEADING
COLLABORATION
IN THE ARM
ECOSYSTEM



Agenda

- Runtime PM autosuspend
- Move on hrtimer
- Impacts and constraints
- Statistics

Runtime PM autosuspend

- Based on timer and jiffies
- Min granularity is between 1 tick and 2 ticks
 - Which means [4-8 ms[on arm64
 - And [10-20 ms[on arm32
- 8 ms is quite long
 - At least on embedded systems

Runtime PM autosuspend

- Example : autosuspend sets to 1ms for GPU on an arm64 platform
 - Timeout moved from the lower bound to the higher bound (4 to 8ms)
 - Power impact to wait 7 more ms before power gating

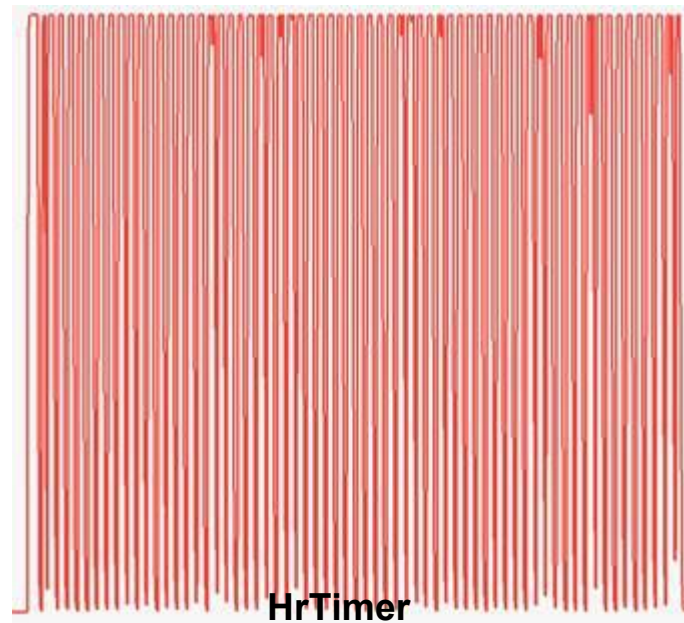
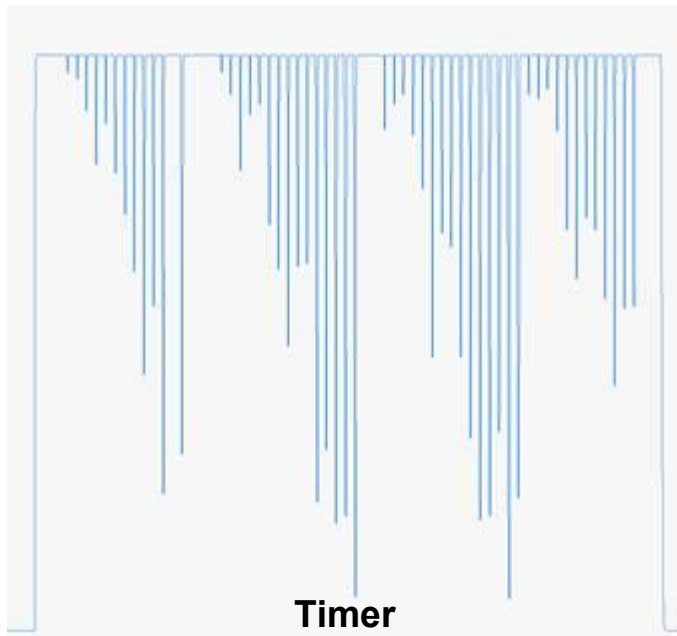
- What if we can be closer to 1ms ?
 - Finer grain is only necessary for short timeout

Move on hrtimer

- Get better granularity
 - Up to HW capabilities
- Take advantage of slack to gather updates
 - Similar to what is done above 1 seconds with timer
 - 12% of slack
 - Keep short autosuspend accurate
 - Gather long autosuspend in one wake up

Power impact

- Back on the 1ms autosuspend of the CPU
 - 10% of power decrease for GPU for idle case
- GPU voltage domain



Perf impact and constraint

- Move on ktime
- Use 64bits variable
 - Impact on 32bits arch
- Performance impact ?
 - mark_last_busy : +5% on hikey arm64 octo cores
 - 1.05 usec vs 1.11 usec
 - rpm_suspend: -5% on hikey arm64 octo cores
 - 14.92 usec vs 14.21 usec

PM runtime statistics

- Based on jiffies
- Anything lower than 1 tick might be not seen
 - 4ms on arm64
 - 10ms on arm
- Move on ktime
 - Get better statistics
 - Aligned with genpd

What else to take care ?

- Early boot availability
 - Availability at early boot ?
 - Use `ktime_get_mono_fast_ns` instead ?



Thank You

For further information: www.linaro.org