


CPU Idle Loop Rework

Rafael J. Wysocki

Intel Open Source Technology Center

September 26, 2018

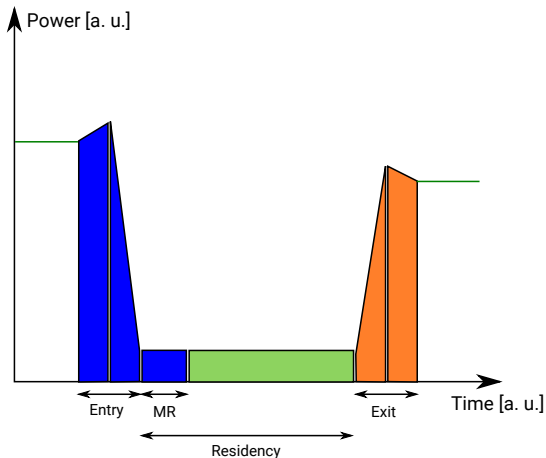
Terminology: CPUs = Logical CPUs



CPUs: Busy Vs Idle

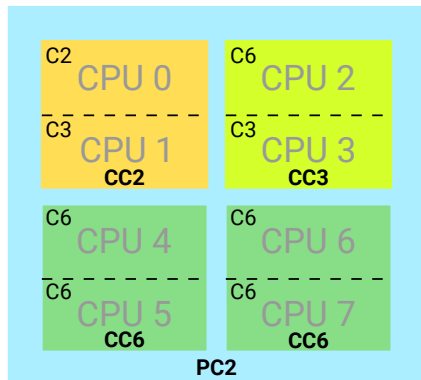


CPU Idle States Properties

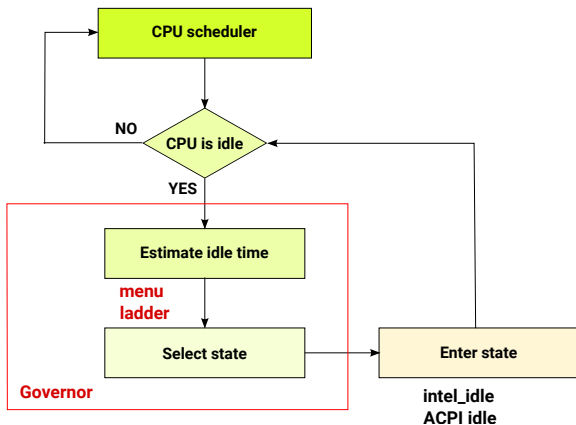


CPU Idle States Example: C-states On Intel Hardware

MWAIT <HINT>



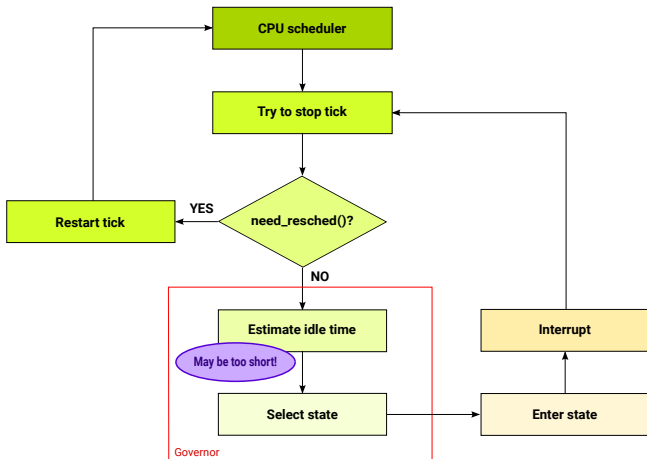
High-level CPU Idle Time Management Control Flow



Complication: CPU Scheduler Tick Timer



Original Idle Loop Design Issue



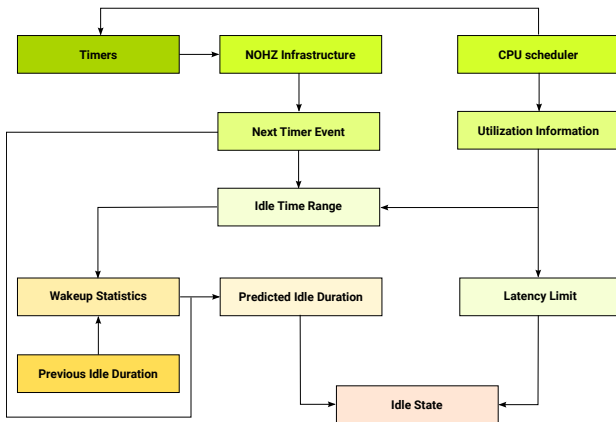
Short Idle Duration Prediction Problem

Actual Predicted	Long Idle	Short Idle
Long Idle	WIN	LOSS
Short Idle	LOSS	LOSS

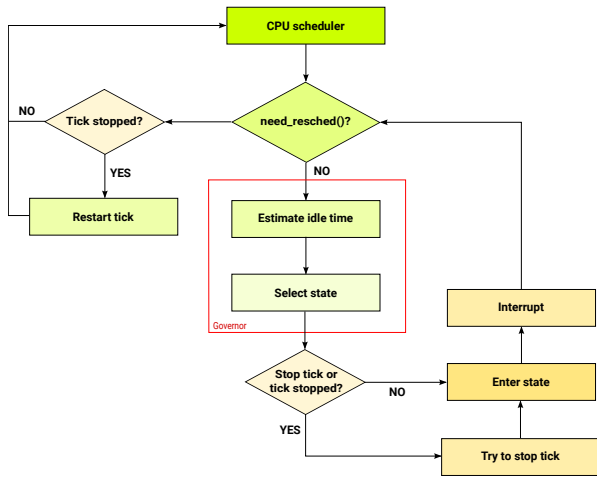
Idle Governor Problem: Two Categories Of Information



The menu Idle Governor



Redesigned Idle Loop (Linux* 4.17 And Later)







* Other names and brands may be claimed as the property of others





Short Idle Duration Prediction Problem Mitigated

Actual Predicted	Long Idle	Short Idle
Long Idle	WIN	LOSS
Short Idle	Neutral	WIN


Idle Power (Intel OTC Server Power Lab)



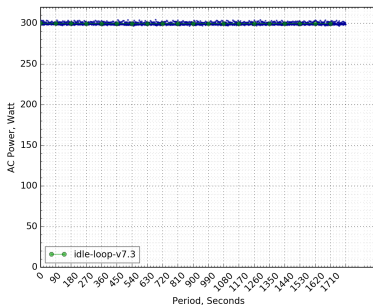
Idle Power (Intel OTC Server Power Lab) Continued



Idle Power (Intel OTC Server Power Lab) Continued




Before



After (different scale)

Questions?



References



Linux Kernel Bug 199227, *Idle loop ordering issue* (https://bugzilla.kernel.org/show_bug.cgi?id=199227).



T. Ilsche, M. Hähnel, R. Schöhne, M. Bielert, D. Hackenberg,
Powernightmares: The Challenge of Efficiently Using Sleep States on Multi-Core Systems
(<https://tu-dresden.de/zih/forschung/ressourcen/dateien/projekte/haec/powernightmares.pdf?lang=en>).



R. J. Wysocki, *Power Management Challenges in Linux* (https://www.linuxplumbersconf.org/2017/ocw//system/presentations/4652/original/linux_pm_challenges.pdf).



R. J. Wysocki, *PM Infrastructure in the Linux Kernel – Current Status and Future*
(https://events.linuxfoundation.org/sites/events/files/slides/kernel_PM_infra_0.pdf).



J. Corbet, *The cpuidle subsystem* (<http://lwn.net/Articles/384146/>).

Disclaimer

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© Intel Corporation