



Contribution ID: 90

Type: **not specified**

Frequency-invariance gaps in current kernel

Wednesday, September 14, 2022 3:00 PM (25 minutes)

The kernel's load tracking scales the observed load by the frequency the CPU is running at, this scaled value is used to determine how loaded a CPU truly is and how its frequency should change.

Currently, on X86, four-core turbo level is used as the maximum ratio for every CPU. However, Intel client Hybrid platforms have Pcores and Ecores, and Intel server platforms with Intel-Speed-Select-Technology enabled have high-priority cores and low-priority cores.

The Pcore/High-Priority-Core can run at higher maximum frequency, while the remaining cores can only run at lower maximum frequency.

In these cases, unified maximum ratio for every CPU doesn't reflect the truth and brings unfairness to the load balance.

Also, the current code doesn't handle special cases where the frequencies for one or more CPUs are clamped via sysfs.

We would like to demonstrate the impacts brought by those issues for further discussion.

I agree to abide by the anti-harassment policy

Yes

Primary authors: ZHANG, Rui; CHEN, Yu

Presenter: ZHANG, Rui

Session Classification: Power Management and Thermal Control MC

Track Classification: LPC Microconference: Power Management and Thermal Control MC