



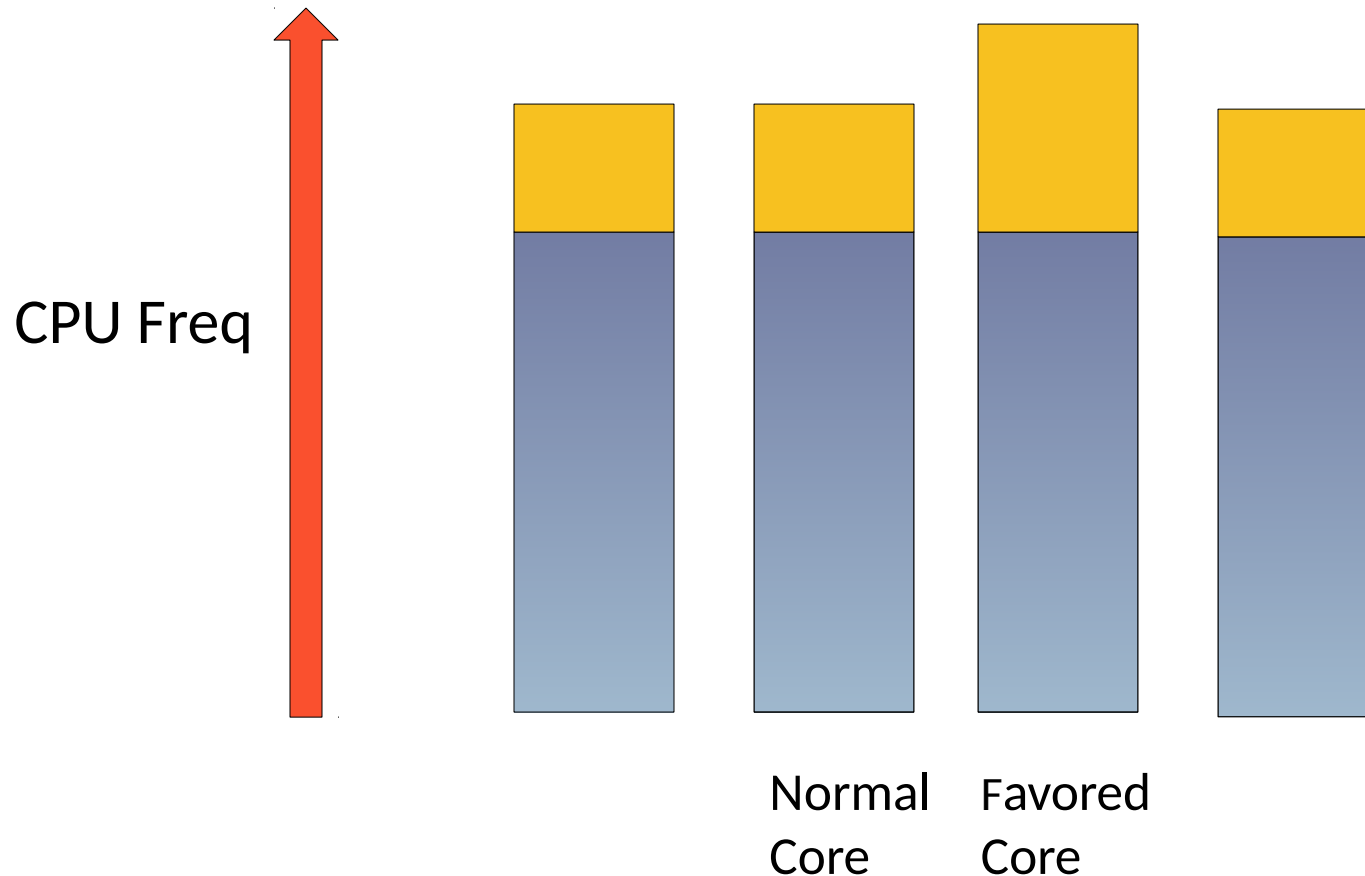
Favored Cores a.k.a. Intel Turbo Boost Max Technology 3.0

Srinivas Pandruvada / Tim Chen



What is Favored Core?

- Favored core (core with extra turbo boost frequency headroom)



Taking Advantage of Favored Core

- Characterize the capabilities of each core and program the results into fuses.
- A means exist for software to enumerate the per core capabilities
- Move work onto the favored core
 - By scheduler
 - By pinning prioritized task manually

Capability Enumeration

- Require HWP (Hardware P-State enabled)
- ACPI_CPC.Highest Performance
 - Enable the CPPC support for x86
 - The OSC_SB_CPC_DIVERSE_HIGH_SUPPORT bit needs to be set to reveal the diverse core capabilities.
 - Program the max turbo frequency of each core as a priority metric, that the scheduler can use

Linux Scheduler Support

Extend and Leverage existing ASYM_PACKING scheduler feature first developed for POWER7

- Pack the processes in the order specified by the priority metric of each CPU during load balancing. Pull processes to idle CPU that has higher priority.
- Each CPU's priority metric is assigned during boot time when OS enumerates the top turbo capability of each CPU.
- Currently we base our priority metric on the max turbo boost frequency. The smt sibling's priority is discounted by half to make sure we don't overpack a core.