

Interplay of Power Management at Core and Server Level

Authors : Jörg Lenhardt, Wolfram Schiffmann, Jörg Keller

Abstract : While the feature sizes of recent Complementary Metal Oxid Semiconductor (CMOS) devices decrease the influence of static power prevails their energy consumption. Thus, power savings that benefit from Dynamic Frequency and Voltage Scaling (DVFS) are diminishing and temporal shutdown of cores or other microchip components become more worthwhile. A consequence of powering off unused parts of a chip is that the relative difference between idle and fully loaded power consumption is increased. That means, future chips and whole server systems gain more power saving potential through power-aware load balancing, whereas in former times this power saving approach had only limited effect, and thus, was not widely adopted. While powering off complete servers was used to save energy, it will be superfluous in many cases when cores can be powered down. An important advantage that comes with that is a largely reduced time to respond to increased computational demand. We include the above developments in a server power model and quantify the advantage. Our conclusion is that strategies from datacenters when to power off server systems might be used in the future on core level, while load balancing mechanisms previously used at core level might be used in the future at server level.

Keywords : power efficiency, static power consumption, dynamic power consumption, CMOS

Conference Title : ICSRD 2020 : International Conference on Scientific Research and Development

Conference Location : Chicago, United States

Conference Dates : December 12-13, 2020