Smart Chip, System and Data Center
Dynamic Provisioning of Power and Cooling from Chip Core to the Cooling Tower

Chandrakant D. Patel

Distinguished Technologist
Hewlett Packard Laboratories
chandrakant.patel@hp.com
www.hpl.hp.com/personal/Chandrakant_Patel

Economics and sustainability in the global compute utility necessitates management of energy as a key resource. The overwhelming "burdened" cost of delivery of power to a dense aggregation of industry standard compute, networking and storage hardware in a data center requires a close examination of all the energy flow paths from *chip core to the cooling tower*. The energy savings and associated environmental impact justifies the need for devising a control system that can precisely provision the power required by the cooling resources and dynamically adapt the heat dissipation profile at chip, system and data center level.

This talk examines the thermal challenges arising from highly integrated high power density microprocessors and servers. The miniaturization of semiconductor technologies has pushed the current power density of the microprocessor core over 200 W/cm^2 - an order of magnitude increase in 5 years - resulting in the use of active micro-scale heat removal techniques. State of the art application of heat removal technologies, applied based on maximum heat load and managed with a lack of knowledge of the overall system requirements, do not suffice from an energy efficiency point of view. Currently chips have the flexibility to scale power. This variability in heat generation must be utilized to enable balanced chip performance based on the most efficient provisioning of the cooling resources. To enable "right" provisioning of cooling resources, flexibility must be devised at all levels of the heat removal stack – chip, system and data center. The ability to change the temperature and mass flow of the coolant are the key flexibilities with respect to the cooling system. With these underlying flexibilities in heat generation and heat removal, one can overlay a low-cost sensing network and create a control system that can modulate the cooling resources and work "hand in hand" with a power scheduling mechanism to create an energy aware global computing utility. However, to address this multi-scale problem, a metric that scales from microns to meters is required. An analysis and evaluation engine based on the second law of thermodynamics quantifying the destruction of available energy (exergy) from chip core to the cooling tower is proposed.

An example of a "smart" data center with a distributed sensing layer, flexibilities in compute, power and cooling resources, combined with data center management based on high level thermo-fluids policies is described. Preliminary results that substantiate the energy savings advantage are presented.

As a future research direction, a global control system that provisions power and cooling resources based on the need, and uses a metric such as MIPS per unit of exergy destroyed as a performance criteria, is proposed. The salient message is that such a vision can be realized through multi-disciplinary collaboration between computer science, mechanical engineering and electrical engineering.