Future Research Challenges in Temperature-aware Systems

An enterprise perspective

Partha Ranganathan
Principal Research Scientist
Hewlett Packard Labs

Enterprise challenges:
Compute equipment consume power…

- Electricity costs
  - For large data center, recurring costs: $4-$8 million/yr

  "… energy costs for [data center] building $1.7 million last year...", Cincinnati Bell, 2003

  "… electricity costs large fraction of data center operations…," Google 2003

- Environmental friendliness
  - Compute equipment energy use: 22M GJ + 3.9M tons CO2
  - EnergyStar (US), TopRunner (Japan), FOE (Switzerland),…

  "…goal to increase computer energy efficiency by 85% by 2005." Japan's "TopRunner" energy program, 2002

Enterprise challenges:

Power delivery and heat extraction costs

- Power delivery & cooling costs large fraction of TCO
  - Capital costs:
    - For 10MW data center, $2-$4 million for cooling equipment
  - Recurring costs:
    - At data center, 1W of cooling for 1W of power
    - For 10MW data center, $4-$8 million for cooling power

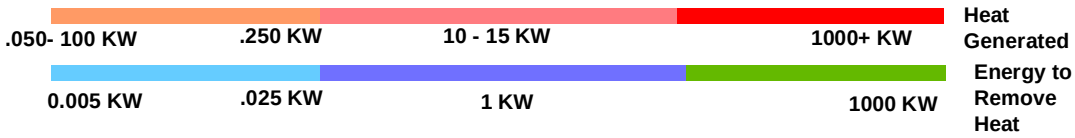

[Source: SmartCooling, HP Labs]

"The issue at those levels isn't cooling per se, but affordable cooling."
Ron Hughes – President California Data Center Design Group, Computerworld 11/04

"IBM said … satisfied customer will require more dual SCSI drives on blade centers without sacrificing blade server density because of potential overheating." CRN 8/6/04

"Power supplies were a large fraction of building the data center," NACIO 2004

Enterprise challenges:
Implications on density, compaction, reliability

- Compaction / density not possible without solving heat
  - Blades increasing density; Data center pushback on cooling

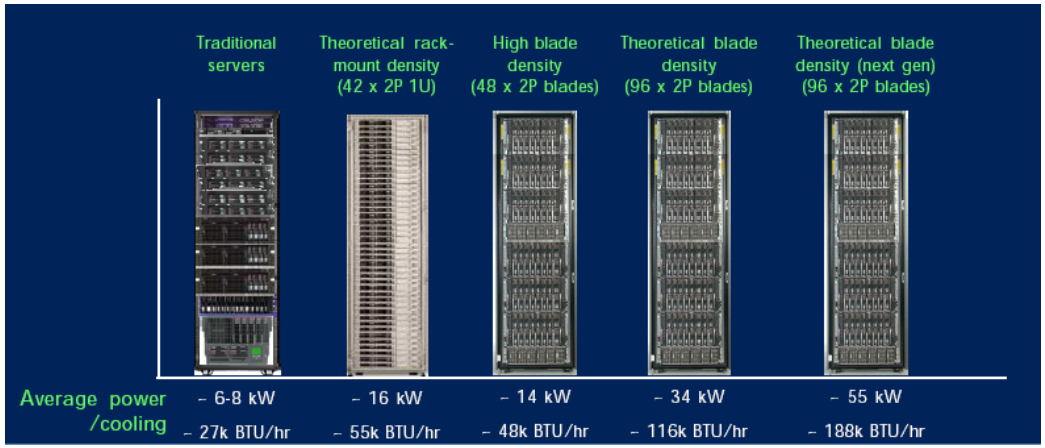

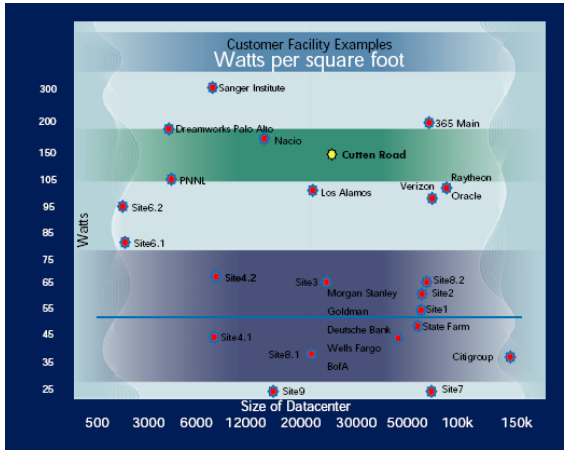

- Availability / reliability issues with higher temperature
  - 50% server reliability degradation for 10°C over 20°C [Uptime]
  - 50% decrease in hard disk lifetime for 15°C increase [Colt]

Two hard challenges

Challenge 1:
Cooling is not the problem!

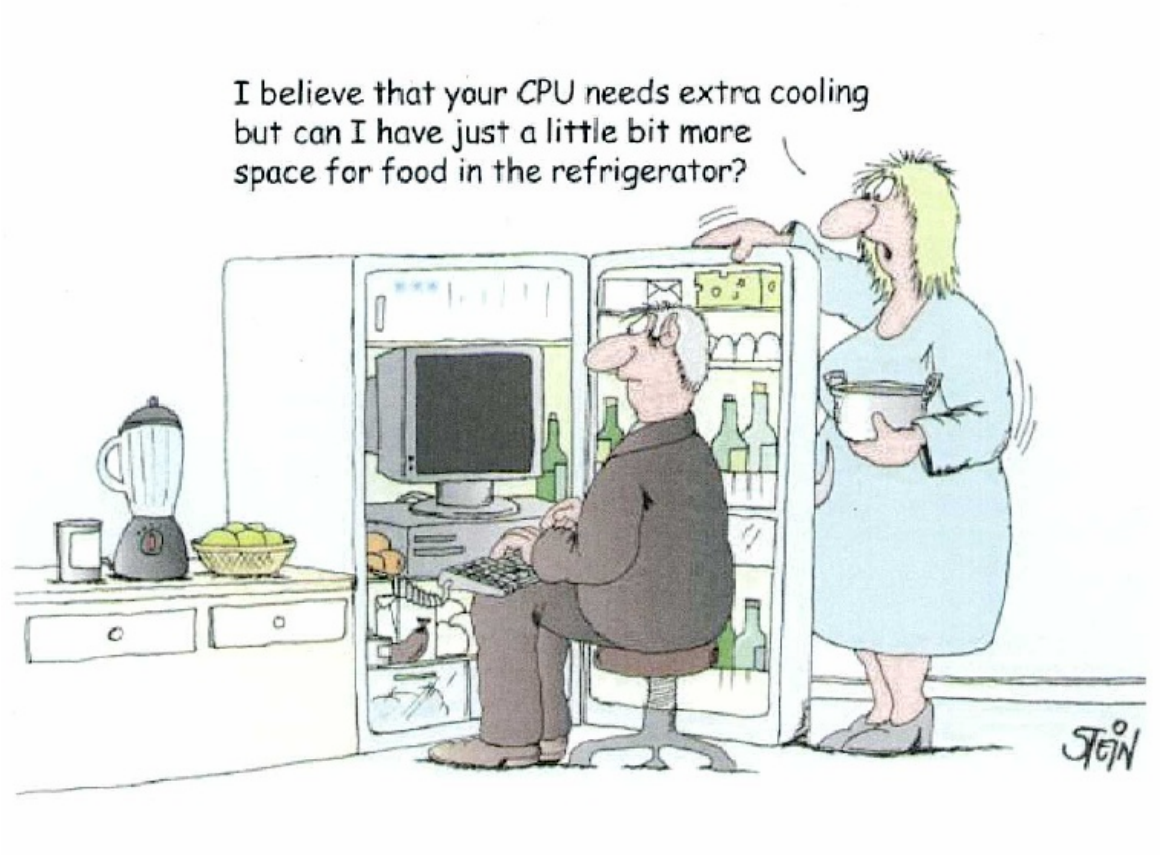

It is all about <u>end-user experience!</u>

The Intersecting Control Loops Problem

- Control loops for different objective functions
  - Example: Shrek2 rendering on HPL utility environment
  - Cooling control
  - Power control
  - Performance control
  - Other business SLO (availability, costs, …)

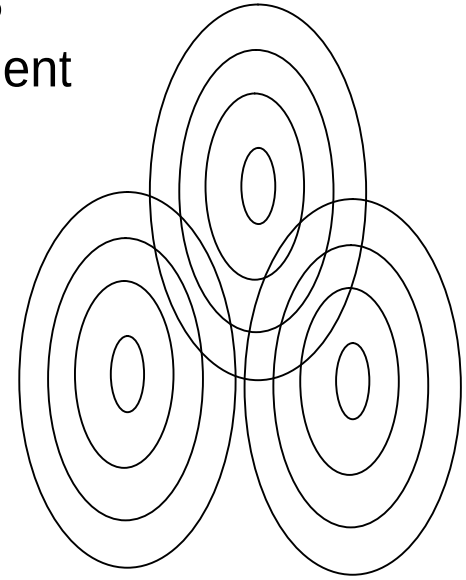

- Control loops at different layers
  - at chip, server, enclosure and rack, data center, grid
  - E.g.: LongRun + PowerRegulator + DenseAndSmart + Weatherman

- No co-ordination between levels
  - Same knobs, conflicting control
  - Interaction of one on another

What is needed?

Holistic cross-layer solutions for TCO

Going beyond buzzwords

- Integrated IT/Facilities management
  - Multiple service providers, multiple resource providers
  - How do you specify and negotiate cost and price?
- Multi-level management
  - Multiple levels of co-ordinated control and management
  - What is the information to expose?
- Holistic system architecture design

Challenge 2

It is all about predicting heat problems!

The Heat Understanding Problem

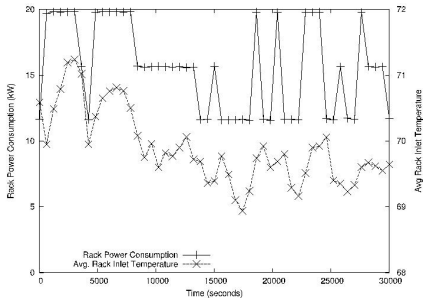

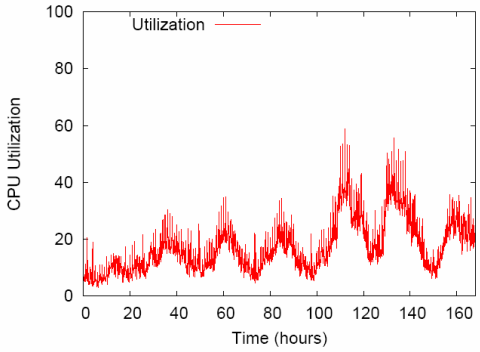

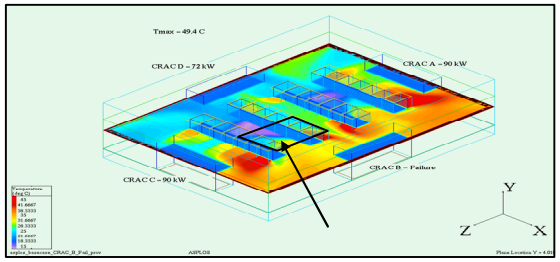

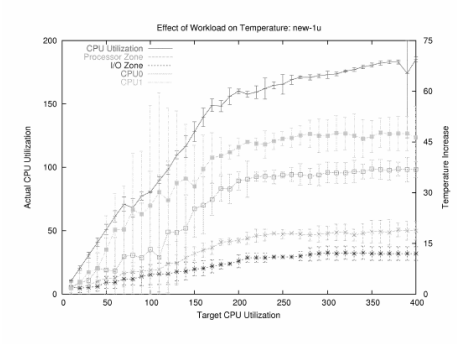

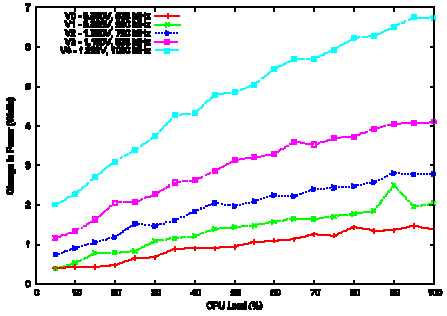

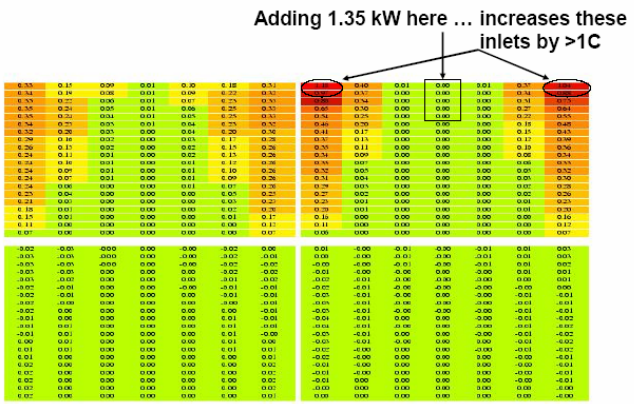

- Correlating heat to power to compute to SLO (currency)
- Measuring and modeling (low-cost)

What is needed?

Benchmarks, metrics, tools, models

Heat-aware approach to modeling systems

- What are the right benchmarks?
  - Currently: SPEC, SPLASH, TPC
  - Needed: ERP, CRM, file server, …
- What is the right objective function?
  - Currently: MIPS, energy-delay
  - Needed: SLA, risk, exergy, SHI, utility, … (EnergyStar ratings?)
- How do we measure/simulate/model?
  - Current: vtune/SimpleScalar/Wattch/HotSpot
  - Needed: integrated IT/facilities environments

Summary: the next big challenges

- Understanding heat inefficiencies
  - Benchmarks, tools, models, measurements

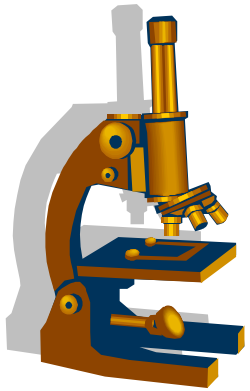

- Holistic cross-layer heat management
  - Optimize for total cost of ownership across different layers

Questions?

Speaker contact:
Partha.Ranganathan@hp.com