A Science of Power Management
a systems perspective

Parthasarathy Ranganathan
Distinguished Technologist, Hewlett Packard Labs

What does a science of power management mean to me?

Approach

Prior body of work
- chips to datacenters
- mobile to enterprise
- metrology to solutions

Energy-aware user interfaces
Heterogeneity-based architectures
Power-aware blade servers
Integrated IT/facilities resource mgmt
Profit-aware scheduling
Unified power management arch
Microblades and megaservers
Power-aware networking
Tools, metrics, benchmarks, …
Joulesort, Zesti, Weatherman
Consil, Splice, eprof, …

hpl.hp.com/personal/partha_ranganathan

Interesting topics for us to ponder…

Why do we need a science of power management?

Of course, power is important

Battery, electricity, environment, heat, compaction, reliability, …

But, havent we done a lot?

[Ranganathan, EE282 class notes]

$P = C*Vdd^2*F_{0\to1} + Tsc*Vdd*Ipeak*F_{0\to1} + Vdd*I_{leakage}$
Average power, peak power, power density, energy-delay, …

CIRCUITS	ARCHITECTURE	COMPILER, OS, APP
• Voltage scaling/islands • Clock gating/routing • Redesigned latches/flip-flops • Redesigned memory cells • Other optimizations	• Voltage/freq scaling • Gating • Split instrucn windows • SMT thread throttling • Bank partitioning • Cache redesign • Lowpower states • DRAM refresh-control • Switching control • Code compression • Data packing/buffering	• Switching control • Memory access reduce • Powermode-control • CPU/resource schedule • Memory/disk control • Networking • Distributed computing • Fidelity control • Dynamic data types • Power API

hp invent

Yes!

...but

Need much lower power

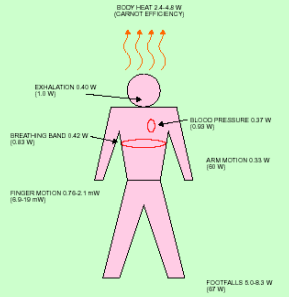

[Starner, IBM Systems Journal,1996]

Solar energy?

164W/sqm * 9% * 3cmx2cm * 50 deration for non-sunny periods = 10mW

Increased compaction…

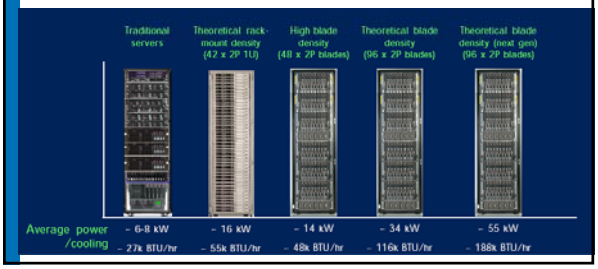

not enough cooling?

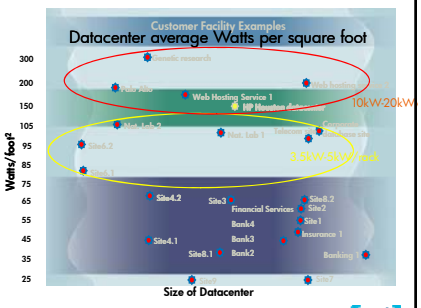

The cost of power and cooling is likely to exceed that of hardware…

Luiz Barroso, Google

More work needed on power mgmt…

So, how can a formal science help?

Formulating the problem

$$P = k_B T \frac{dn^2}{\theta}$$

$$n = \sqrt{\frac{\theta P}{k_B T d}}$$

One billion pentiums in one handheld!

Another quick experiment

- Laptop: 165X
- Handheld: 15X
- Cell phone: 6X
- RIM pager: 92 mW

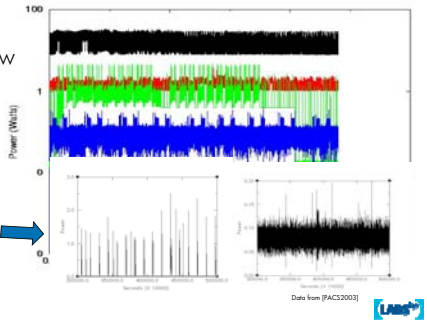

Radio wakeup:
- 100ms (iPAQ);
- 1.2 sec (cell)
- 5 sec (RIM)

❶Can we identify minimum energy needed for a task?

(optimality?)

minimum energy needed for a task…

How do we define energy & work done?

MIPS per Watt?

Energy-Aware User Interfaces

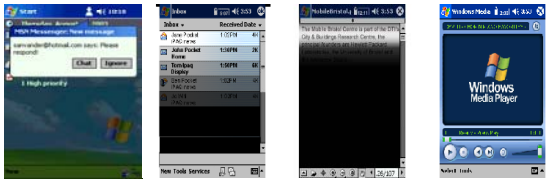

~~MIPS~~ per Watt?

End-user experience!

Burdened costs of power

$$(1 + K_1 + L_1 + K_2 L_1)\, U_{\$,grid} P_{\substack{consumed \\ hardware}}$$

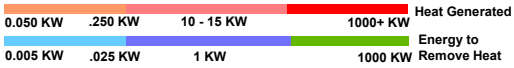

Supply vs demand side power efficiencies

Supply vs demand side power efficiencies

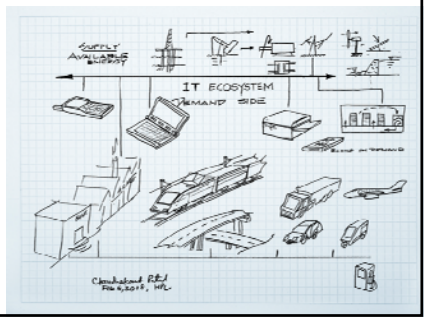

Energy in extraction, reclamation, transport, manufacturing, …

Energy in operation

~~MIPS~~ per ~~Watt~~?

end-user experience per exergy watts?

❷How do we measure and reason about energy efficiency?

Thinking about the solution

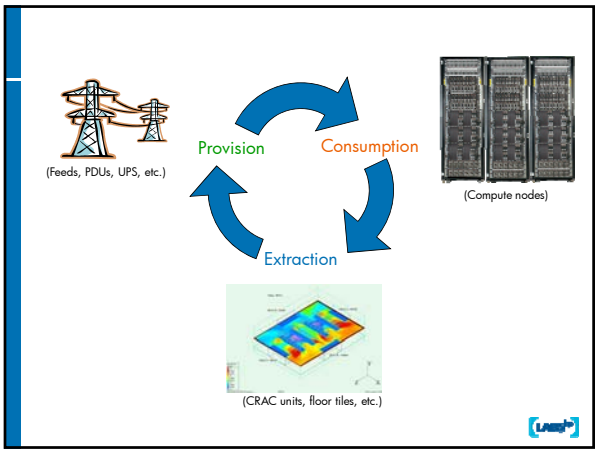

Landscape of Power Optimizations

$P = C*Vdd^2*F_{0\to1} + Tsc*Vdd*Ipeak*F_{0\to1} + Vdd*I_{leakage}$

Average power, peak power, power density, energy-delay, …

CIRCUITS	ARCHITECTURE	COMPILER, OS, APP
• Voltage scaling/islands • Clock gating/routing Clock-tree distribution, half-swing clocks	• Voltage/freq scaling • Gating Pipeline, clock, functional units, branch prediction, data path	• Switching control Register relabeling, operand swapping, instruction scheduling • Memory access reduce
	• Low-power states • DRAM refresh-control divided module, scratch	Power-aware routing, proximity-based routing, balancing hop count, … • Distributed computing Mobile agents placement, network-driven computation
	• Switching control Gray, bus-invert, address-increment • Code compression • Data packing/buffering	• Fidelity control • Dynamic data types • Power API

Avoid waste!

How do we avoid waste?

A framework to optimize power

[CACM09]

A taxonomy of inefficiencies

(how do we identify waste?)

A taxonomy of techniques

(how do we reduce waste?)

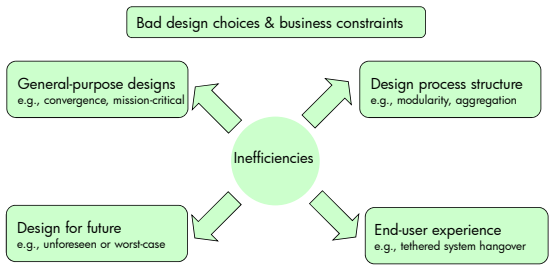

avoiding waste & improving power efficiency

❶Energy-efficient technologies
e.g., replace disk with flash, replace copper with optics, …

❷Match power to work
e.g., energy proportionality; turn-off/dial-down things, …

❸Match work to power
e.g., asymmetric multicores, temp-aware scheduling, …

❹Piggy back energy events
e.g., shared caches, coalesced request streams, …

❺Special purpose solutions
e.g., GPUs, ASICs, …

38 10 April 2009

❻Cross layers for efficiency
e.g., ensemble power management, …

❼Tradeoff some other metric
e.g., fidelity-aware energy management

❽Tradeoff the uncommon-case
e.g., power-supply efficiency, …

❾Spend somebody's power
e.g., remote server offload for mobile power

❿Spend power to save power
e.g., periodic cleanup to save energy, …

39 10 April 2009

❸How do we formalize these principles?

Building the solution

Measure & model

Predict & analyze

Actuate & control

Measure & model
Predict & analyze
Actuate & control

"Science" challenges in each category

Measure & model

e.g., neural-network models for thermal maps (Consil/Weatherman)

Predict & analyze

e.g., models to predict power-performance-cooling (Zephyr, Zesti)

Actuate & control

e.g., federated control theory, scheduling algorithms

One example

[ASPLOS08]

Power Struggles!

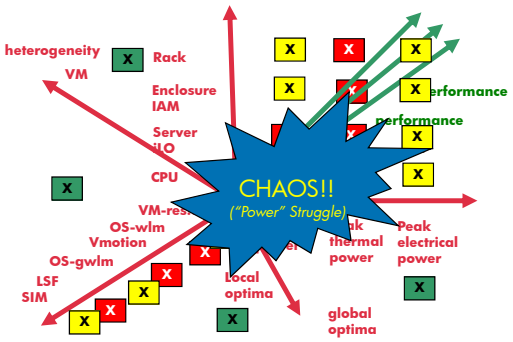

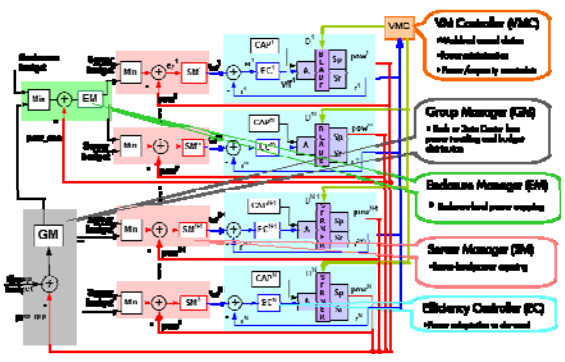

Feedback Controller @ Core

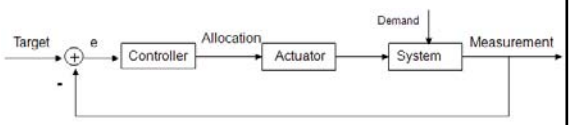

Standard feedback control loop

- Formal theoretical guarantees of stability and performance
- Can account for inaccuracies in model, workload variations

Efficiency + capping at single server

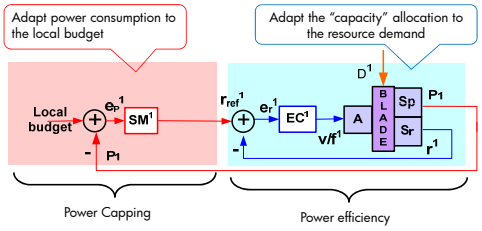

Multi-level Power Capping

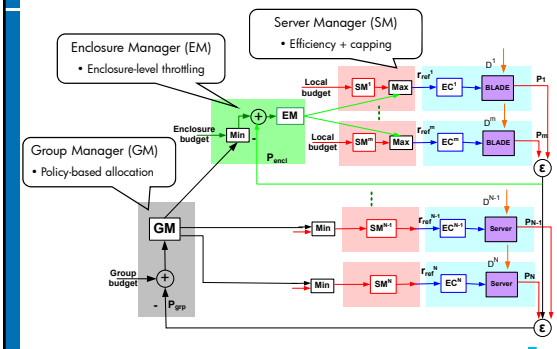

Virtual Machine Workload Distributor

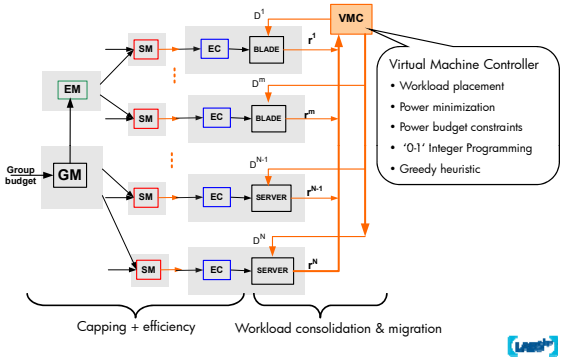

Unified and Extensible Architecture

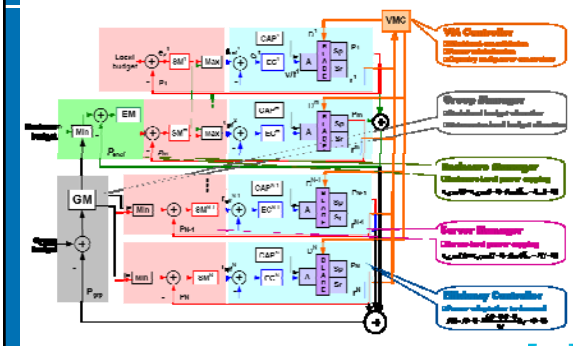

It works!

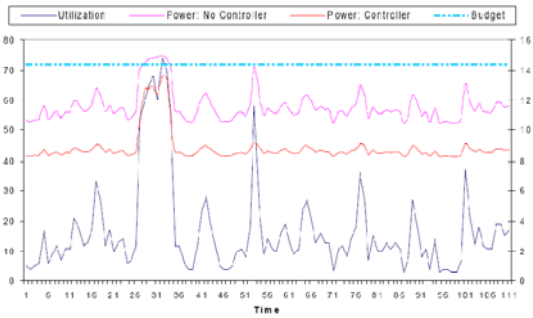

It works well!

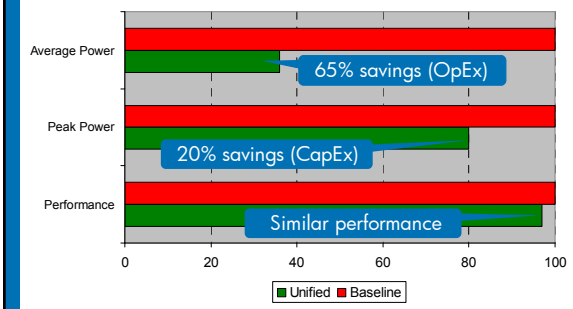

Other interesting insights…

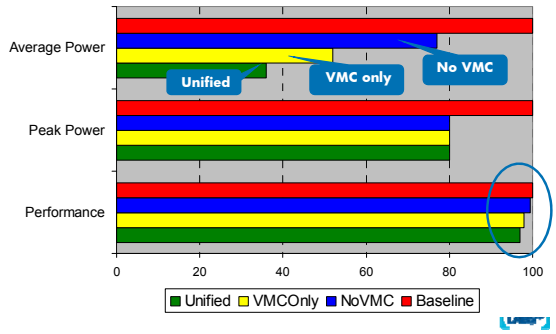

Interfaces

Formal rigor

Flexible

Extensible

Federated

Closing

Power mgmt ready for a science

important future challenges & prior body of work

Several interesting challenges

- formulating the problem
  - e.g., optimality, metrics
- thinking about the solution
  - e.g., principles of power management
- designing the solution
  - e.g., models, analysis, control, algorithms, …

Thanks!!