Power-centric Cloud Management

Ada Gavrilovska, Karsten Schwan, Chad Huneycutt, Matthew Wolf, …
Hrishikesh Amur, Jhenkar Vidyashankar…
Center for Experimental Systems in Computer Science
Pramod Kumar, Yogendra Joshi
CEETHERM/Mechanical Engineering
Georgia Tech

Green Computing Research Initiative at GT

http://img.all2all.net/main.php?g2_itemId=157

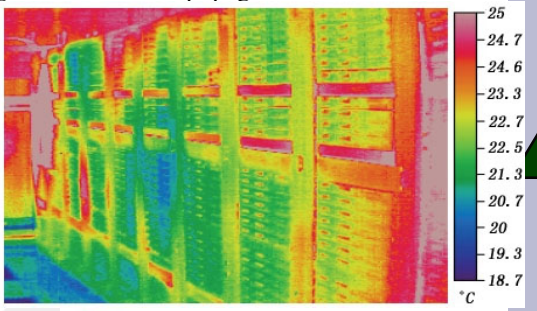

focus of our work:

Datacenter and beyond: design, IT management, HVAC control… (ME, SCS, OIT…)

Rack: mechanical design, thermal and airflow analysis, VPTokens, OS and management (ME, SCS)

Board: VirtualPower, scheduling/ scaling/operating system… (SCS, ME, ECE)

Chip and Package: power multiplexing, spatiotemporal migration (SCS, ECE)

Circuit level: DVFS, power states, clock gating (ECE)

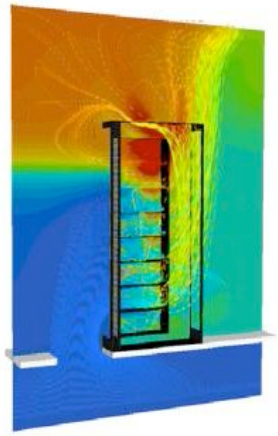

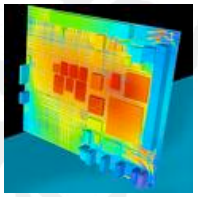

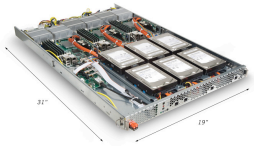

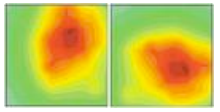

Power distribution and delivery (ECE)

Instrumented Facility

- ~3000 cores
  - VMware VCloud
  - physical partitioning
- Fully instrumented
  - RPDUs; sensors for temperature, air velocity…; fan speeds…;
  - Software tools for per-VM resource usage
- Dynamic controls on cooling infrastructure

Questions pursued

- Placement decisions
  - What's the performance cost of placing a workload on hot vs. cold isles?
  - What's the impact of operating at higher temperatures?
- Management costs
  - When is it more energy-efficient (perf/J) to consolidate VMs on fewer cores vs. run longer in lower power state vs. 'run-to-completion' and idle sooner?
- Usage monitoring
  - How much energy is consumed by this client's VMs?
  - How much energy is consumed by all Hadoop VMs?
- Insights into (hw/sw) component designs
  - How should fan controllers be designed?

Cloud "*Power Map*"

- Physical and software sensors
  - temperature, PDUs, CPU and memory usage, VM-core mappings, application-specific…
- Distributed HBase database
  - Configured in VMs running on small set of nodes, scalable
  - Inputs via SNMP, EVPath, VirtualCenter… ; OSISoft Pi server
  - Schema easily adapted to monitor various entities

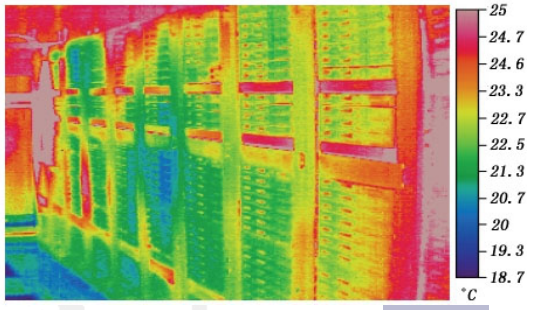

TABLE: Power usage

Row Key	Column Family "VM:"		Column Family "HOST"		Column Family "CLUSTER"	
DATACENTER-ID	Column "VM Name"	Value Watts	Column "Host Name"	Value Watts	Column "Cluster Name"	Value Watts
	VM-0	72	Host 1	40	Cluster 1	40
		71		42		42
		73		46		39
		72		46		39
		76		48		40
	VM1	65	Host 2	67	Cluster 2	55
		66		66		60
		66		66		62
		65		65		66
		65		65		54
	VM2	20	Host 3	64	Cluster 3	82
		22		62		80
		24		60		79
		23		64		70
		23		68		76
	VM3	82	Host 4	71	Cluster 3	36
		82		72		37
		82		72		36
		80		72		36
		76		78		36
	VM4	75	Host 5	82	Cluster 4	44
		75		81		45
		77		81		44
		76		81		44
		76		83		44

TABLE: Cpu Counter

Row Key	Column Fam "VM:"								
Datacenter ID	Column "VM Name"	V							
	Cpu-burn								
	VM1								
	VM2								
		23		68		76		75	58
	VM3	82	Host 4	71	Cluster 3	36	RP-4	56	VC-4	62
		82		72		37		56		61
		82		72		36		56		62
		80		72		36		56		64
		76		78		36		56		65
	VM4	75	Host 5	82	Cluster 4	44	RP-5	49	VC-5	75
		75		81		45		48		74
		77		81		44		48		74
		76		81		44		47		74
		76		83		44		46		76

Dynamic Resource Management

- PowerMap to close-the-loop distributed management tasks
  - algorithms like VMPToken (Ripal Nathuji)
  - integrate Monalytics monitoring infrastructure (Chengwei Wang)
- Extend with VM-level power-metering (Bhavani Krishnan)
- Understand power-related properties of important classes of applications
  - Big Data Hadoop jobs (Hrishikesh Amur); enterprise services like RUBiS (Calton Pu, …)

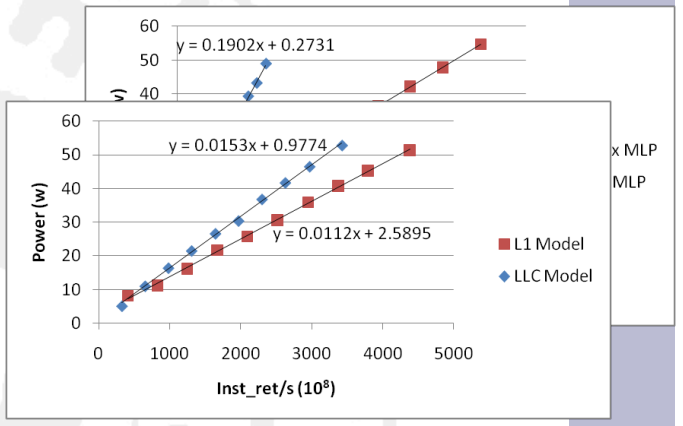

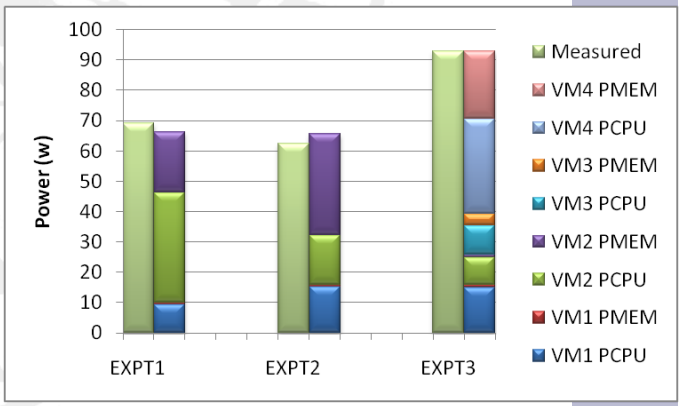

Realizing CoolIT

- Thermal- and energy-aware load management mechanisms
  - Power caps derived from first-order models on datacenter thermal behavior
- Understanding impact of load on behavior of components of cooling infrastructure
  - e.g., fans

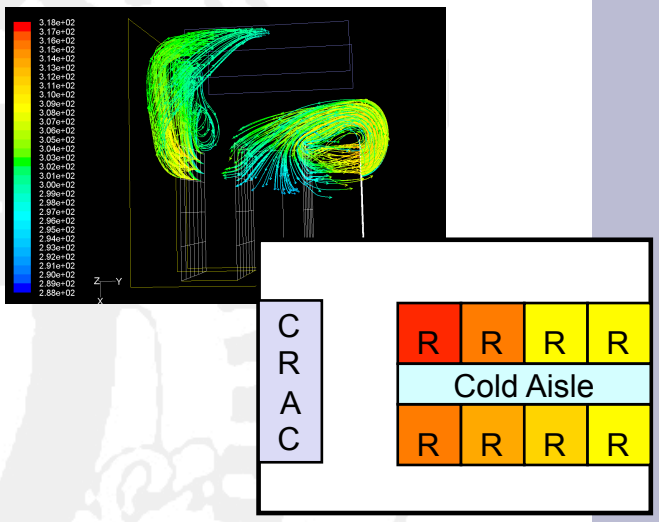

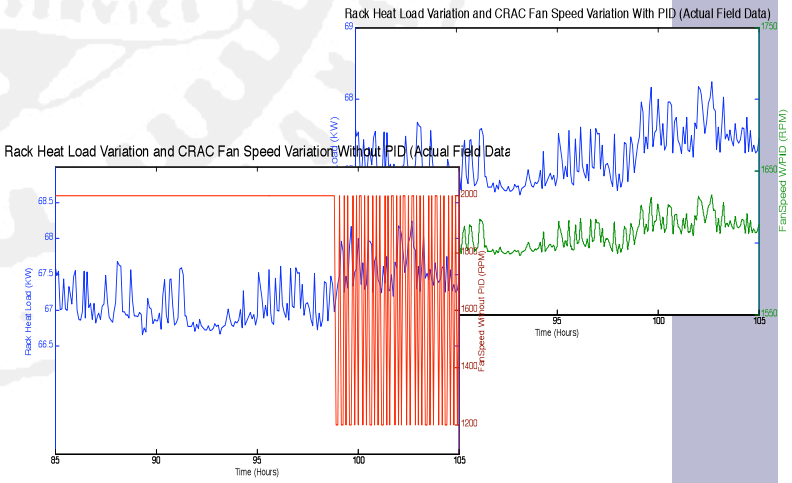


cercs