Green High Performance Computing Initiative at Rutgers University (GreenHPC@RU)

Center for Autonomic Computing

Manish Parashar, Ivan Rodero

Contact: {parashar, irodero}@cac.rutgers.edu

Green HPC

http://nsfcac.rutgers.edu/GreenHPC

Motivations and Challenges

Motivations for HPC Power Management

- Growing scales of High-end High Performance Computing systems
  - Power consumption has become critical in terms of operational costs (dominant part of IT budgets)
- Existing power/energy management focused on single layers
  - May result in suboptimal solutions

Challenges in HPC Power Management

- GreenHPC
  - $/W/MFLOP, defining energy efficiency
  - Infrastructure: thermal sensors, instrumentation, monitoring, cooling, etc.
- Architectural challenges
  - Processor, memory, interconnect technologies
  - Increased use of accelerators
  - Power-aware micro-architectures
- Compilers, OS, runtime support
  - Power-aware code generation
  - Power-aware scheduling
  - DVFS, programming models, abstractions
- Considerations for system level energy efficiency
  - Optimizing CPU alone is not sufficient, need to look at entire system/ cluster
  - Application/workload aware-optimizations
- Power-aware algorithm design

Approach

- Abnormal operational state detection (e.g., poor performance, hotspots)
- Reactive and proactive approaches
  - Reacting to anomalies to return to steady state
  - Predict anomalies in order to avoid them
- Workload/application-awareness
  - Application profiling/characterization

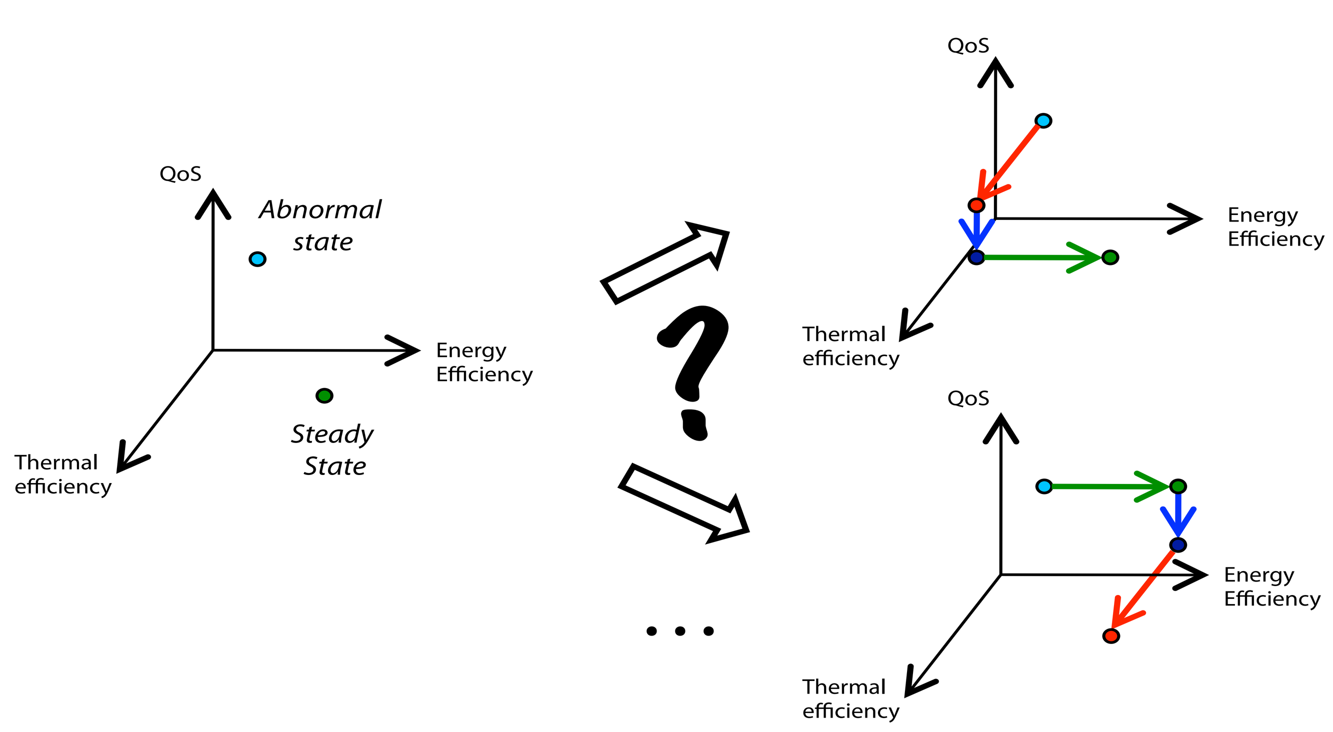

Goals

- Autonomic (self-monitored and self-managed) computing systems
- Optimizing (minimizing):
  - Energy efficiency
  - Cost-effectiveness
  - Utilization

while ensuring (maximizing):
  - Performance/quality of service delivered
- Addressing both "traditional" and virtualized system (i.e., data centers and **Green**HPC in the Cloud)

Cross-layer Energy-Power Management

- Architecture

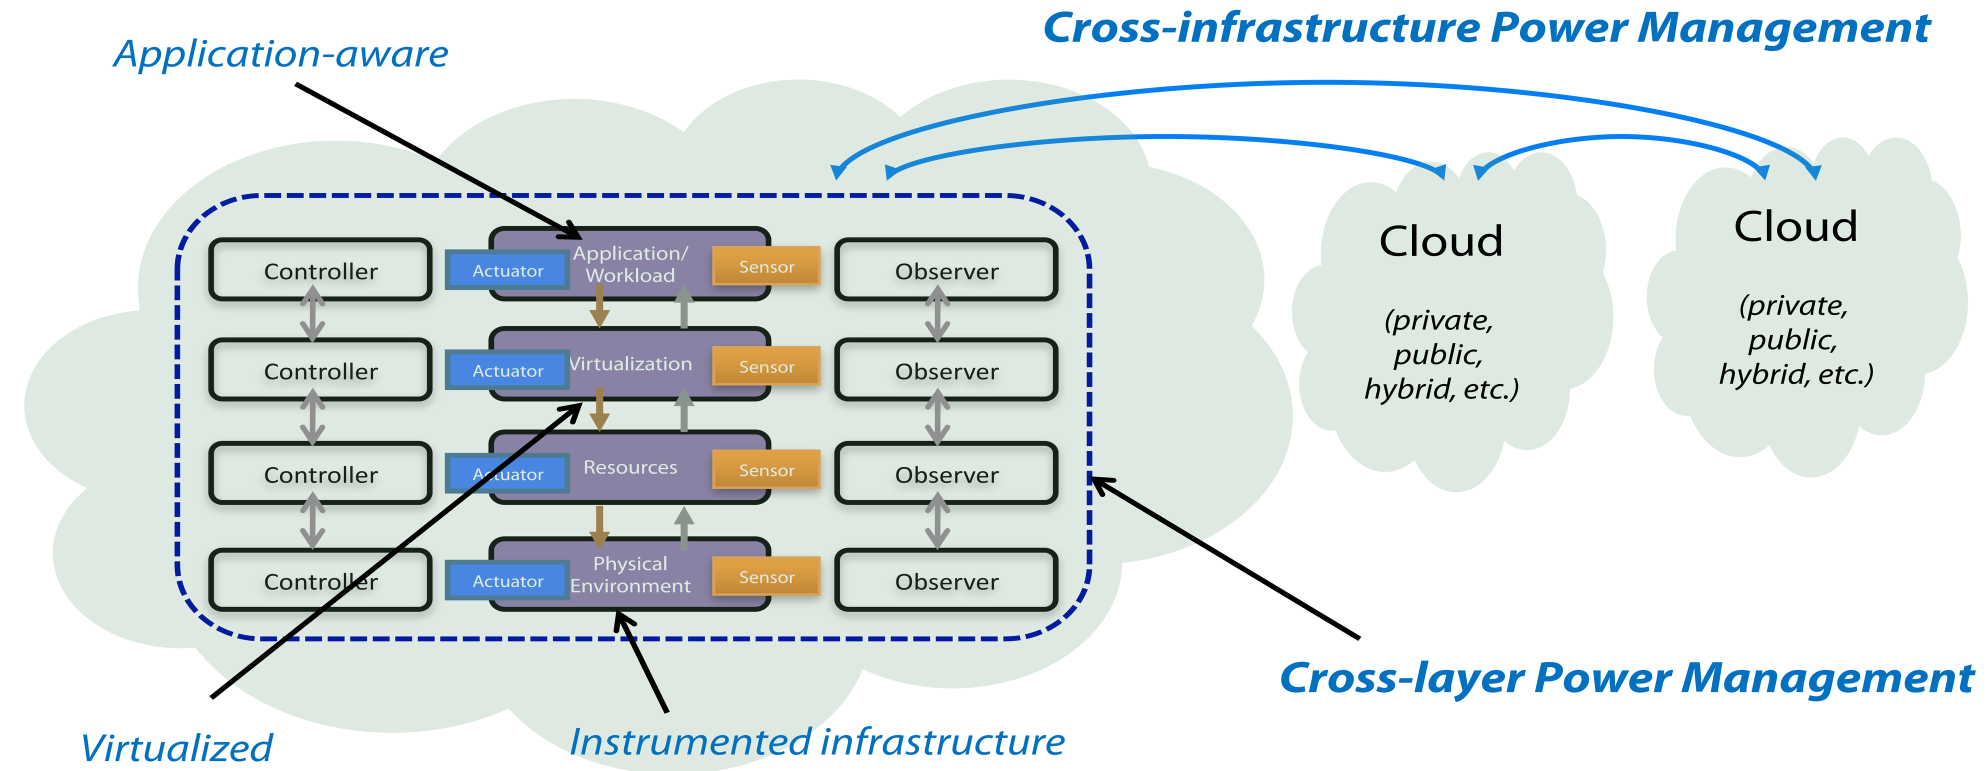

Component-based Power Management [HiPC'10/11]

- **Application-centric** aggressive power management at component level
  - Workload profiling and characterization
- HPC workloads (e.g., HPC Linpack)
- Use of low power modes to configure subsystems (i.e., CPU, memory, disk and NIC)
- Energy-efficient memory hierarchies
  - Memory power management for multi- many-core systems
  - Multiple channels to main memory
  - Application-driven power control (i.e., ensure bandwidth to main memory leveraging channel affinity)

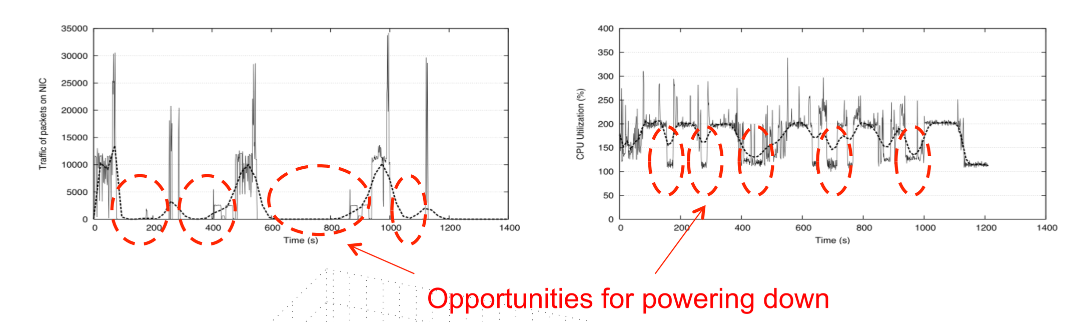

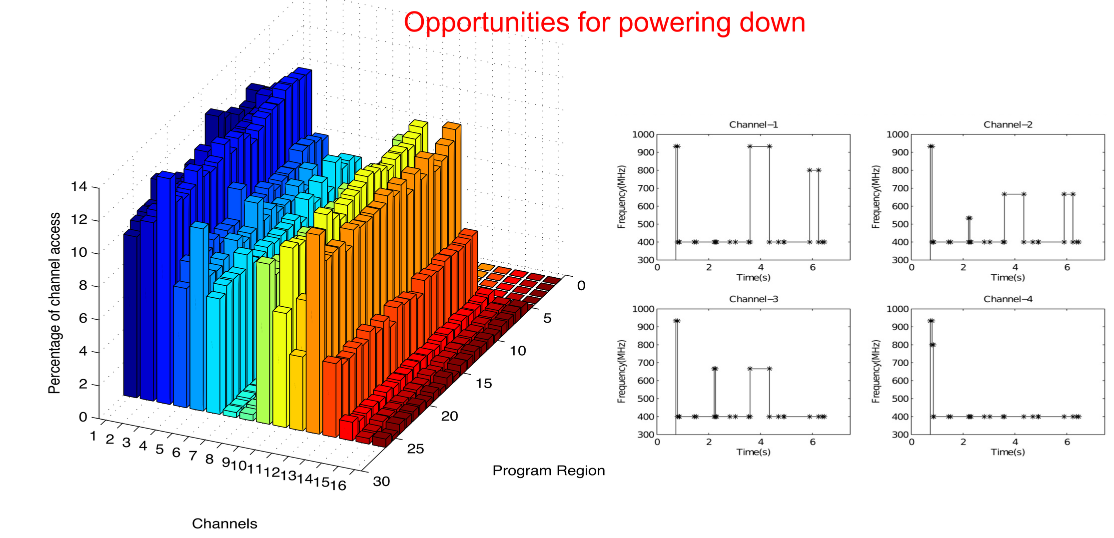

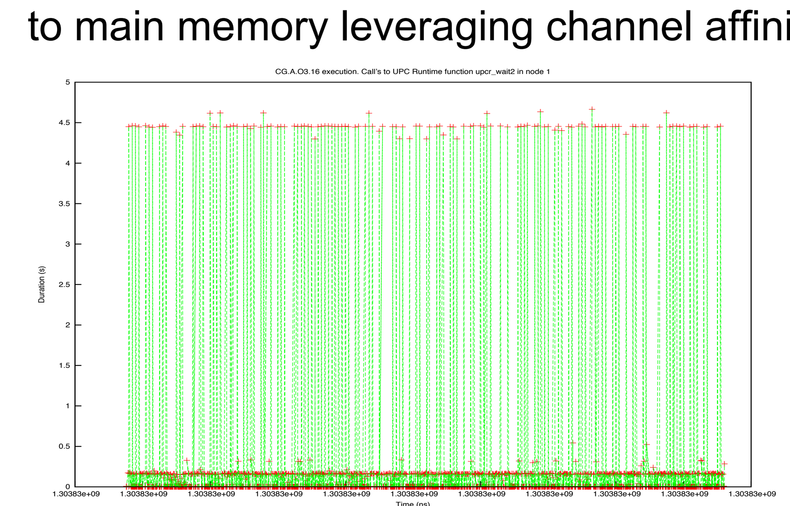

Example: Potential saving during collective operations (e.g., barriers) in the SCC platform

Runtime Power Management

- Partitioned Global Address Space (PGAS)
  - Implicit message-passing
  - Unified Parallel C (UPC) so far
- Target platforms
  - Many-core (i.e., Intel SCC)
  - HPC clusters

Energy-aware Autonomic Provisioning [IGCC'10]

- Virtualized Cloud infrastructures with multiple geographical **distributed** entry points.
- Workloads composed of HPC applications
- Distributed Online Clustering (DOC)
  - Cluster job requests in the input stream based on their resource requirements (multiple dimensions, e.g., memory, CPU, network requirements)
- Optimizing energy efficiency in the following ways:
  - Powering down subsystems when they are not needed
  - Efficient, just-right VM provisioning (reduce over-provisioning)
  - Efficient proactive provisioning and grouping (reduce re-provisioning)

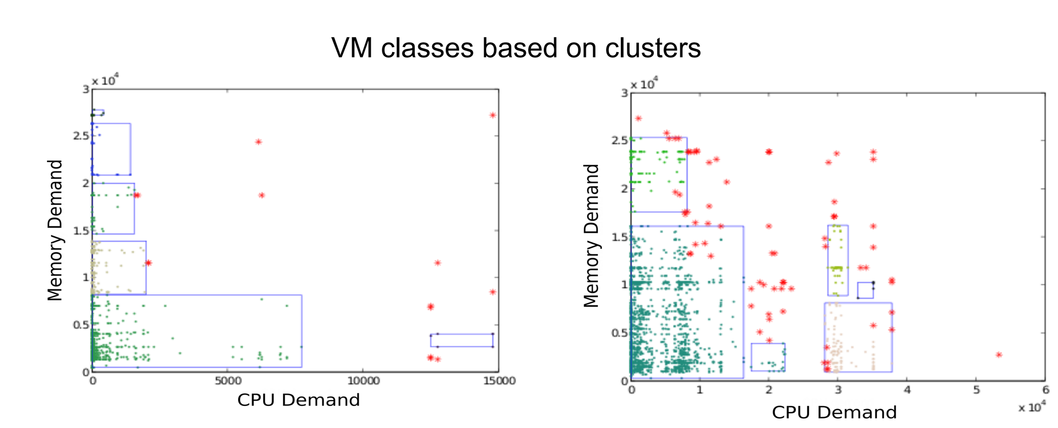

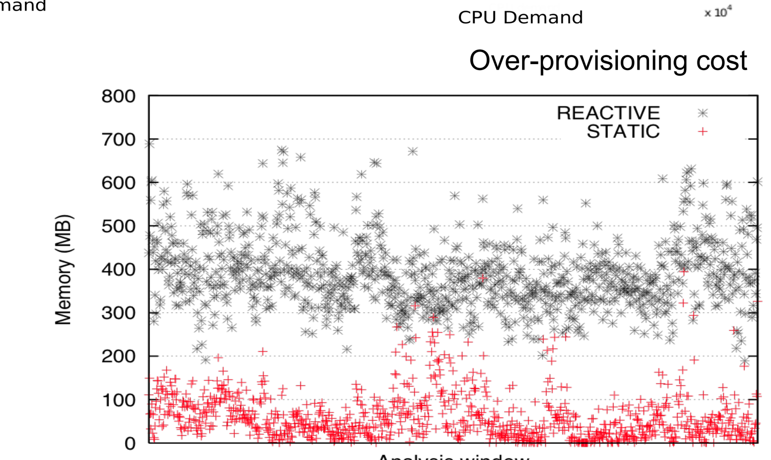

Energy and Thermal Autonomic Management

- Reactive thermal and energy management of HPC workloads
  - Autonomic decision making to react to thermal hotspots considering multiple dimensions (i.e., energy and thermal efficiency) using different techniques: VM migration, CPU DVFS, CPU pinning
- Proactive energy-aware application-centric VM allocation for HPC workloads
  - Strategy for proactive VM allocation based on VM consolidation while satisfying QoS guarantees
  - Based on application profiling (profiles are known in advance) and benchmarking on real hardware

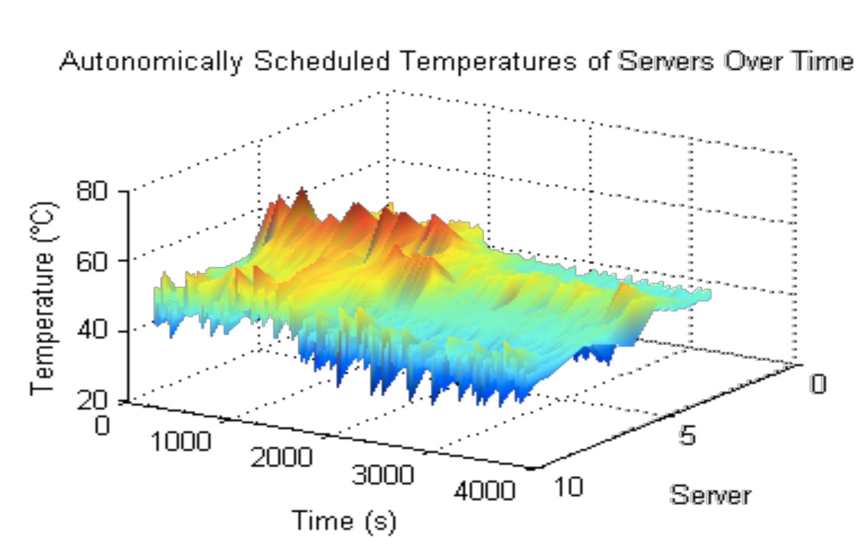

China – US Software Workshop, September 2011

The Center for Autonomic Computing is supported by the National Science Foundation under Grant No. 0758596.

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY