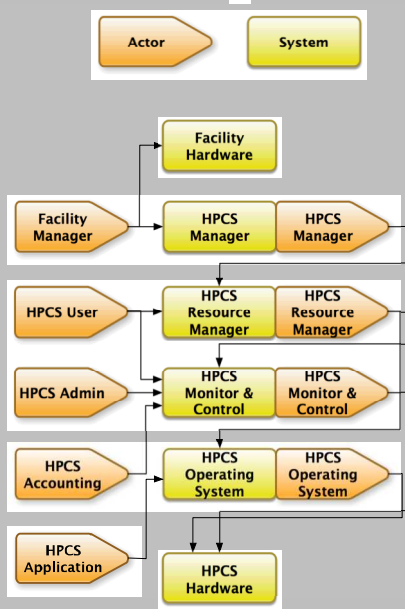


Power and Energy

Coordinator: Jim Laros (SNL)

Participants: *GaTech, LBNL, NWU, ORNL, SNL, UA, UTEP*



*Exceptional
service
in the
national
interest*



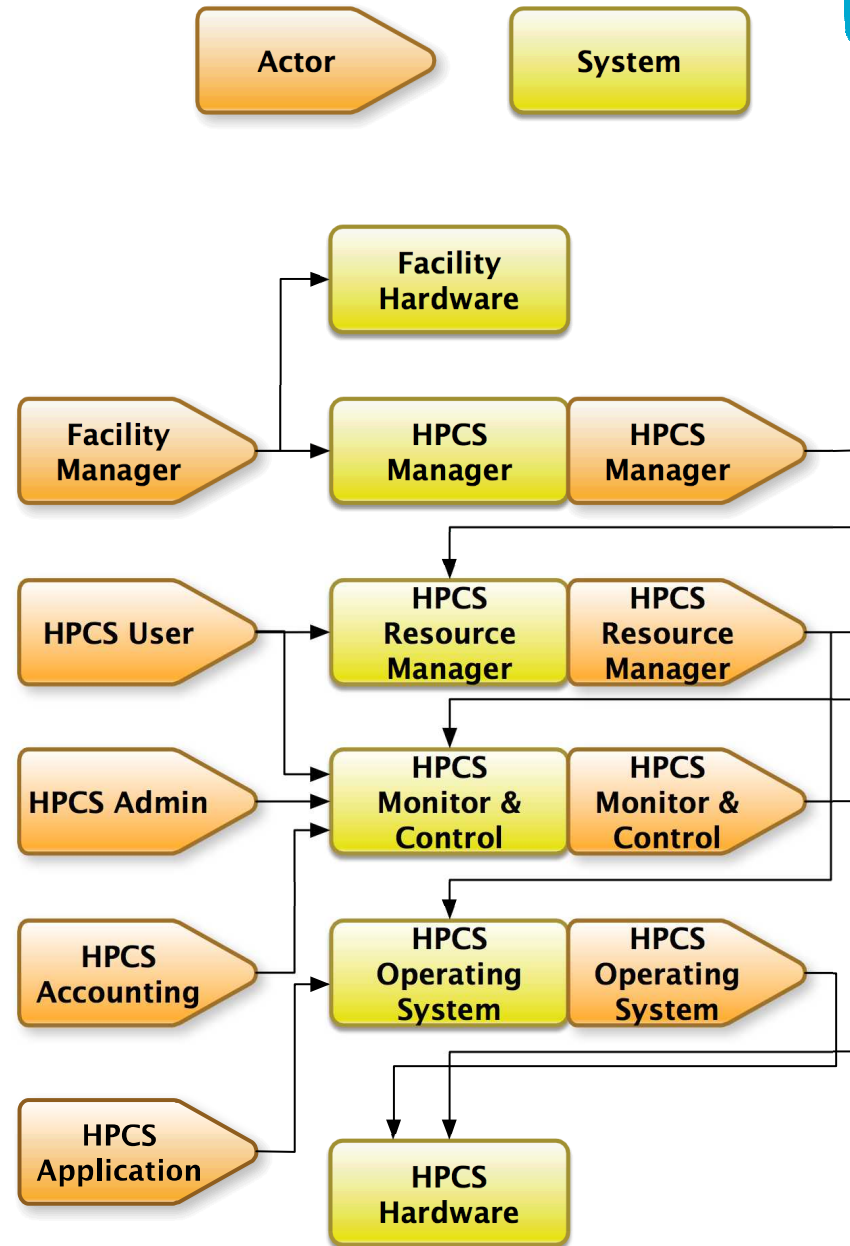
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Power and Energy are Cross-Cutting Sandia National Laboratories

- Objective: Create APIs to enable Power and Energy Research
 - Influences design of Enclave OS/R, NVL, GIB, scheduling, programming models, resilience mechanisms
 - Sandia is coordinating overall API development
 - Individual efforts described in this presentation inform APIs
- Power and Energy constrained future
 - System Level: Annual \$20M - \$40M utilities budget
Large power swings problematic, increase cost
 - Node Level: 100 W power budget per socket
Engineered to support large power swings
Shift power around chip, turbo-boost

Power API

- Widely scoped Power API in development at Sandia Labs
- Provides structure and cohesiveness to proposed work
- Provides links to larger community
 - Bidirectional
- Prototypes will be developed for interactions between targeted components as part of focused efforts:
 - Scheduler
 - Runtime
 - Operating System
 - Application
 - Monitor and Control



HRHI Areas

- Adaptive Virtualization (Northwestern)
- Power/Resilience Modeling (Arizona and University of New Mexico)
- Application Centric Power Management (Georgia Tech)
- Global Power/Energy management (LBNL)

HRHI Adaptive Virtualization

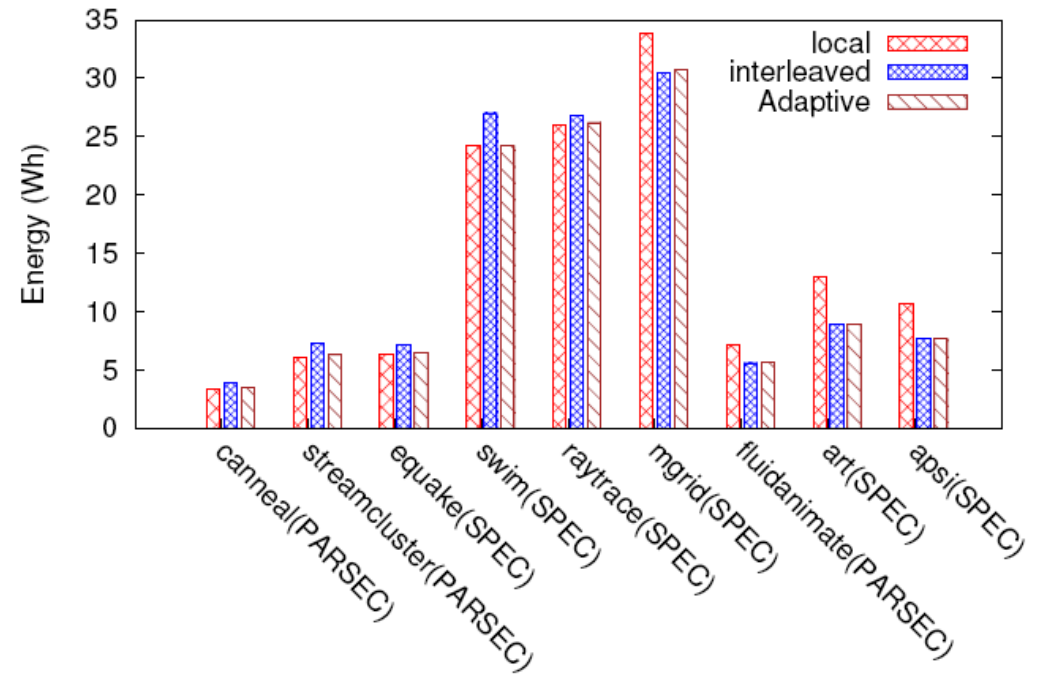
Northwestern University – Peter Dinda

- Can Adaptive Virtualization enable positive power/energy/performance tradeoffs for existing and/or unmodified “black box” applications and OSeS?

- Mechanisms

- Vcore-> hardware thread mapping
- Guest memory-> NUMA domain mapping
- Vcore scheduling / coscheduling
- Dynamic Voltage and Frequency Scaling (DVFS)

- Results compared with “white box” approaches



Adaptive system chooses best mapping with <0.05% overhead

- Minimize execution time: up to 66%
- Minimize energy: up to 31%
- Minimize power: up to 17%

HRHI: Power/Resilience Modeling

University of Arizona – Dave Lowenthal


University of New Mexico – Patrick Bridges

- Predict execution time and power consumption given a hardware configuration
 - Number of nodes, cores/node, power/node
 - Further consider dividing power between different per-node hardware resources
- Extend model to consider power/resilience interactions
 - Study how to combine power and resilience models
 - Determine what OS/runtime information needed to drive and tune these models
 - Use model to handle inter-enclave interactions and ultimately optimize application performance
- Helps define interfaces, requirements and impact in many areas of Power API

HRHI: Application Centric Power Management

Georgia Tech – Karsten Schwan and Ada Gavrilovska

- **What Manageability to expose and how to coordinate across entire application?**
- Fine grain resource scaling:
 - Hide core heterogeneity (e.g., big/little cores) from the application by introducing new power states mapped to use of hetero cores
 - Manage power states to minimize power usage within application's performance constraints: fine grain resource scaling
 - Consider both Core and Uncore (e.g., memory)
- Explicitly managed heterogeneous cores:
 - e.g., accelerators
 - Determine power/performance profiles for applications on cores and with different types of memory (e.g., DRAM vs. NVM)
 - Manage application execution via core/memory remapping
- Explicit examples studied:
 - S3D and other high end codes, co-running with online analytics



Energy Management (LBNL)

- Demonstrate application level management at scale using “global” states and control framework
 - Develop “algebra” for reasoning about energy/power in hardware hierarchies
 - Provide “voting” and arbitration heuristics
 - Develop control mechanisms
 - Integrate energy management with scheduling
 - Provide for performance isolation
- ❖ HR/Hi – Open Loop or Closed Loop control?
 - Open loop
 - Exploration of energy usage for application behavior patterns
 - Pattern classification
 - Declarative pattern behavior at the application level

Questions and Discussion