The PomPP Framework: From Simple DSL to Sophisticated Power Management for HPC Systems*

Yasutaka WADA†, Yuan HE‡, Thang CAO‡, and Masaaki KONDO‡
†Meisei University, Tokyo, Japan, ‡The University of Tokyo, Tokyo, Japan

Power Management for Power-Constrained HPC Systems

- **Effective Utilization of Limited Power Budget**
  - to Realize Exa-Scale HPC Systems
    - ➢ HW Overprovisioning and Power Management
- **Power-Performance Optimization** requires
  - User Effort to Modify Apps for Power Capping
  - Good Understanding of both SW/HW
  - Consideration of Various Systems and Apps

→ (Semi-)Automatic Framework for Power-Performance Optimization

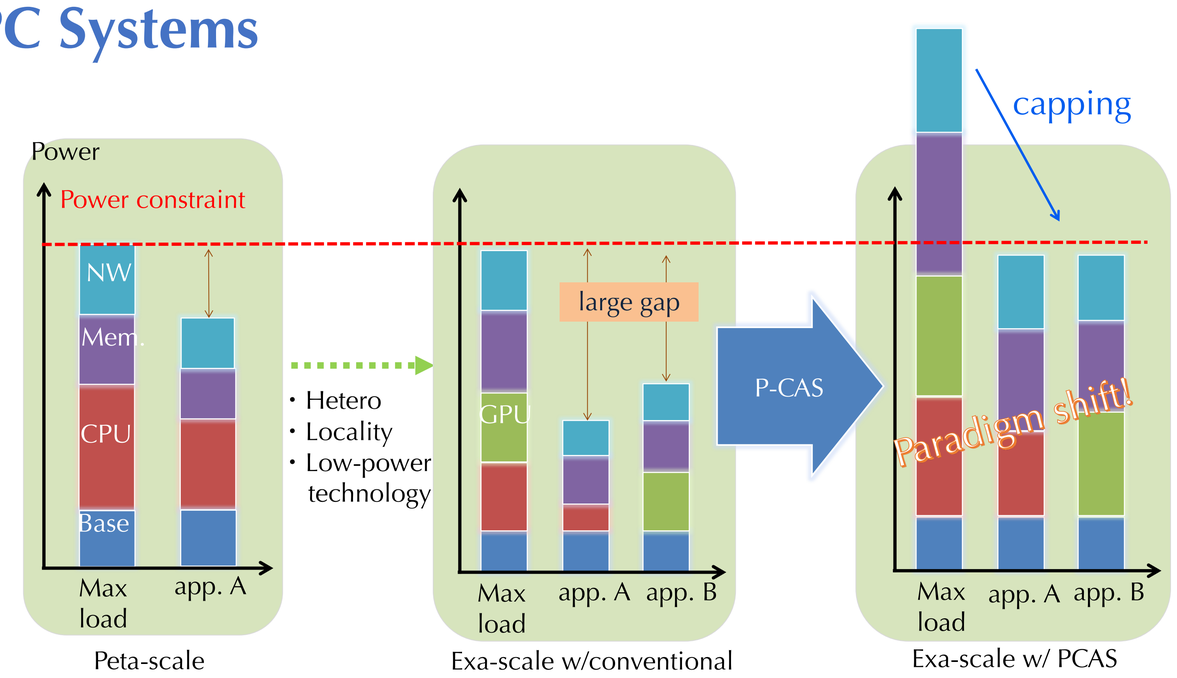

The PomPP Framework for Easy Power Management

- **Common I/F to Control Power-Knobs**
  - Provide Simple APIs to Control Power-Knobs
- **Automatic Instrumentation**
  - Add Power-Knob Control APIs to User Applications
    - ➢ with TAU/PDT based Tools
- **Simple DSL as a Frontend**
  - Generates Job Submission Scripts, etc.
  - Control Power Management Workflow
    - ➢ Optimization Algorithms
    - ➢ Power-Performance Models
    - ➢ Available Power-Knobs

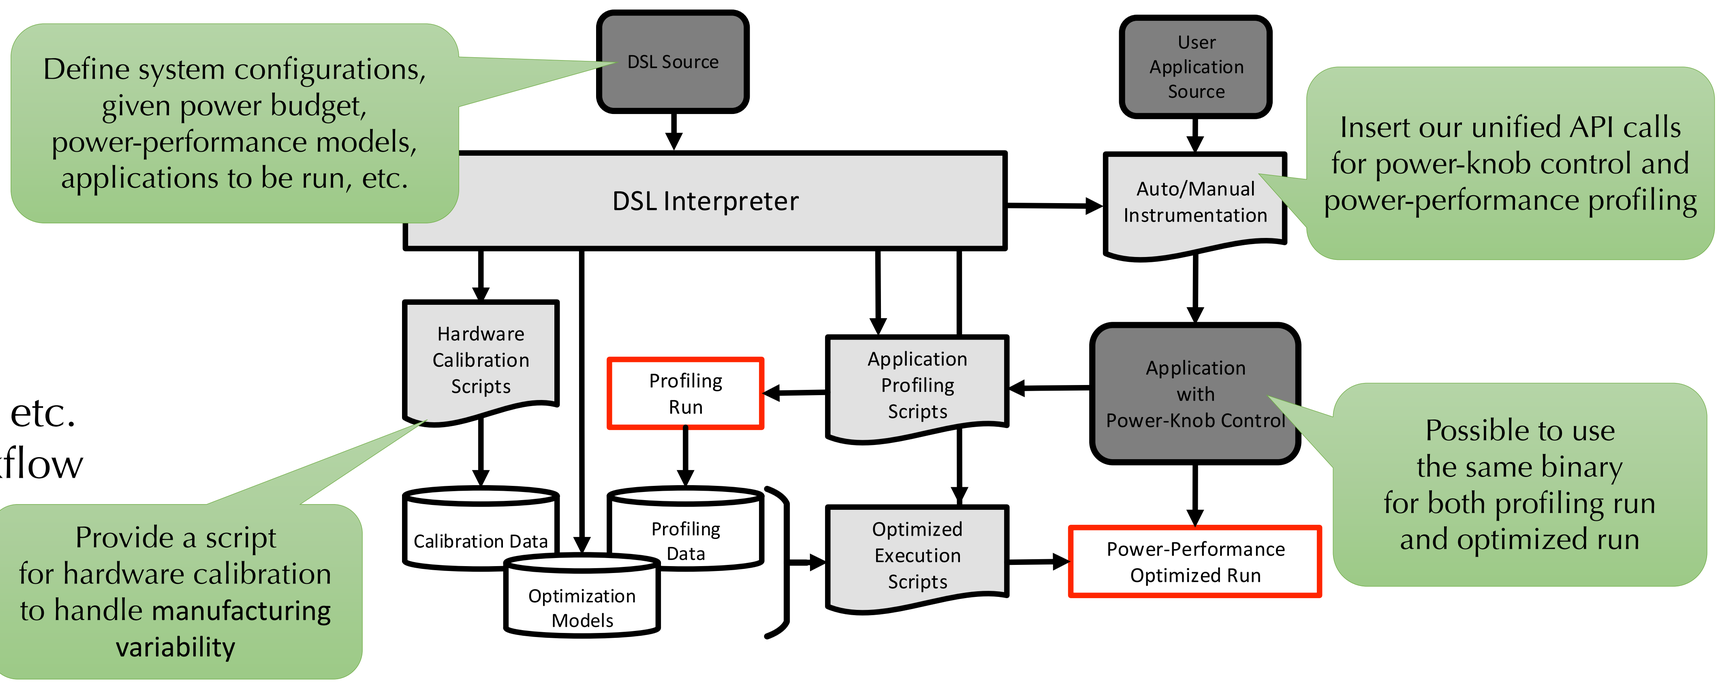

Case Study with Simple Linear Models

- **Run on A HPC System with 16nodes**
  - two Intel Xeon E5-2680 per node
  - RAPL Interface as Power-Knob
- **Assume Linear Relationship between Power and Performance**
  - Generated from Power Profiling Results
  - Simple Model to Reduce # of Profiling Run
- **Allow Slowdown but within the Given Deadline**
  - Power Capping Value is Decided based on Given Deadline and the Mode
  - Evaluated with IS and EP from NPB

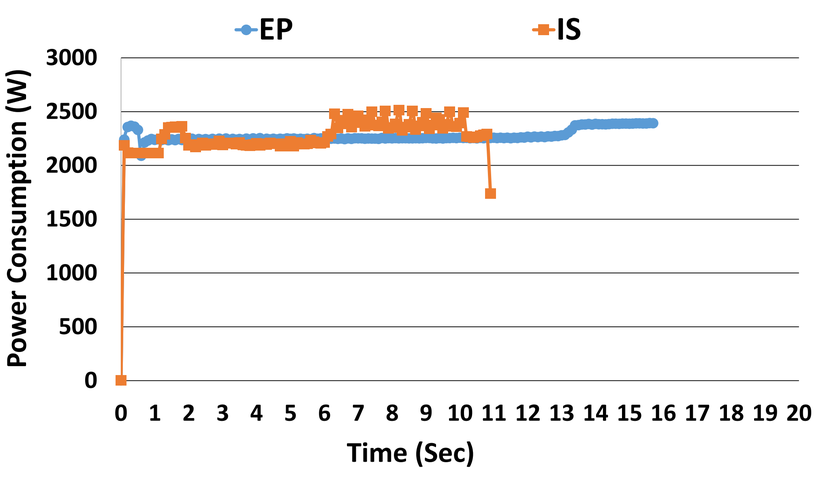

Power Profile with IS and EP from NPB

Profiling runs to extract the power-performance models

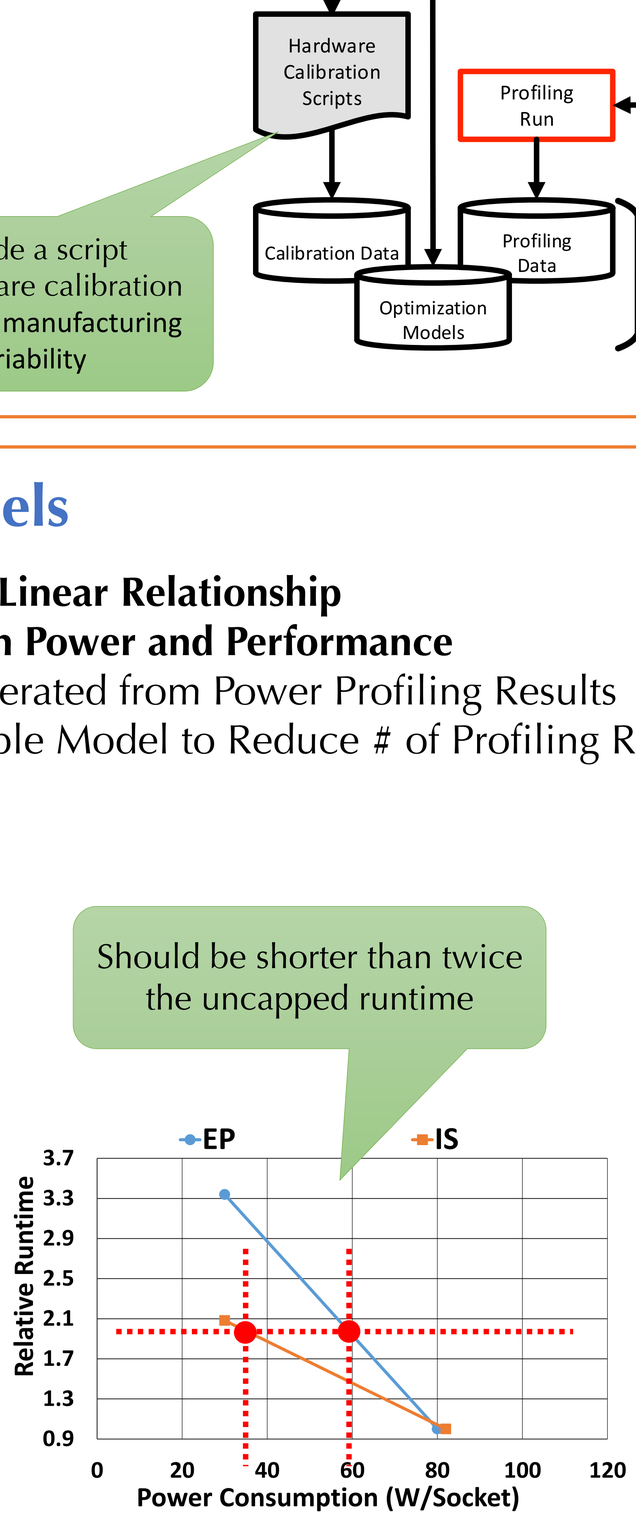

Power-Performance Model generated from the Power Profile

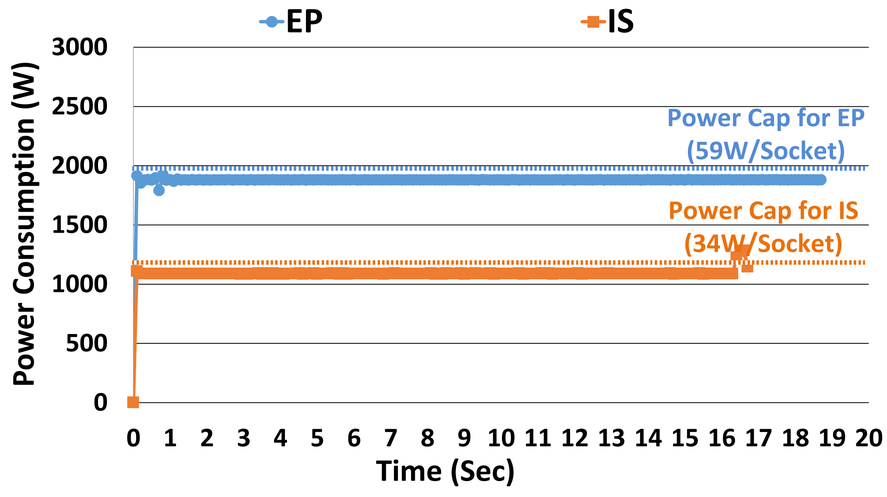

Optimized Execution with a Given Deadline

Guaranteed not to exceed the given deadline while decreasing power

Current and Future Works

- **For Various HPC Systems**
  - Support Other External Tools
  - Support Other Power Knobs than RAPL
- **With System Softwares**
  - Collaboration with Job Scheduler
  - Support for Inter-Job Optimization

* This work was partially supported by the Japan Science and Technology Agency (JST) CREST program for the research project named "Power Management Framework for Post-PetascaleSupercomputers".