Simple DSL for Power-Performance Modeling with Segmented Linear Models*

Yuan HE[1], Yasutaka WADA[2], Guanqin PAN[1], and Masaaki KONDO[3,4]

[1]Shenyang University of Technology, Liaoning, China [2]Meisei University, Tokyo, Japan
[3]The University of Tokyo, Tokyo, Japan [4]RIKEN, Kobe, Japan

Power Management for Power-Constrained HPC Systems

- **Effective Utilization of Limited Power Budget**
  - ✓ to Further Scale the Performance of 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

→ **A Common Power-Performance Optimization Framework**
- ✓ **to Allow the Reuse of Existing Optimization**
- ✓ **to Automate the Power Management Processes**

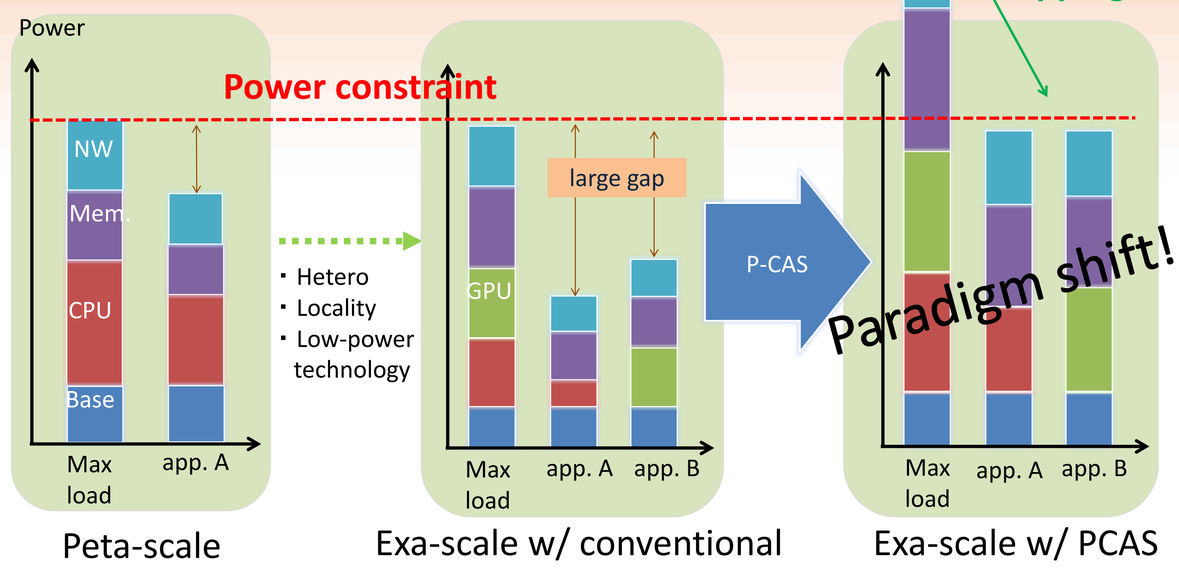

The PomPP Framework for Convenient/Effective Power Management

- **Hardware Configuration and Calibration**
  - ✓ to Set/Modify System Settings
  - ✓ to Calibrate Hardware for Variations in Manufacturing
- **Common I/F to Control Power Knobs**
  - ✓ to Provide Simple APIs to Control Power Knobs
- **Automatic Instrumentation**
  - ✓ to Add Power-Knob Control APIs to User Applications
    - with TAU/PDT based Tools
- **Profiling and Optimization**
  - ✓ for Stats Collection
  - ✓ for Optimized Production Run
- **Simple DSL as the Front-end**
  - ✓ All Features Mentioned on the Left
  - ✓ to Describe Optimization Algorithms
  - ✓ to Specify Power-Performance Models
    - Linear Models
    - Segmented Models
    - Look-Up Tables

Specify system configurations, given power budget, power-performance models, applications to be run, etc.

Provide a script for hardware calibration to handle manufacturing variability

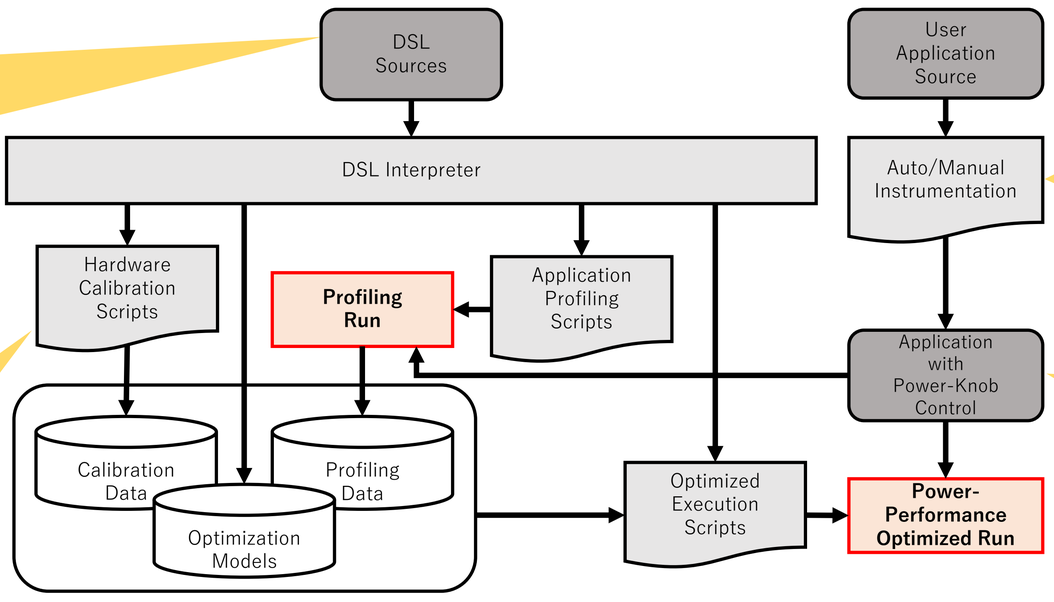

Insert our unified API calls for power knob control and power-performance profiling

Possible to use the same binary for both profiling run and optimized run

Case Study with Simple/Segmented Linear Models

- **Evaluated on a System with 16 Cores**
  - ✓ Intel Xeon E5-2640 v2 per Socket
  - ✓ RAPL as the Power Knob
- **Assume Linear and Segmented Linear Relationships between Power and Performance**
  - ✓ Generated from Profiling Results
  - ✓ 2 Profiling Runs per Application for the Simple Linear Models
  - ✓ 10 Profiling Runs per Application for the Segmented Linear Models
- **2 User-Set Performance Demands**
  - ✓ 50% and 80%, respectively
  - ✓ Power Capping Values for These Performance Demands are Calculated with the Models
  - ✓ Evaluated with Applications from the NPB Suite
  - ✓ Segmented Models Works Better in all 16 Cases

Performance/Power Models of NPB Applications

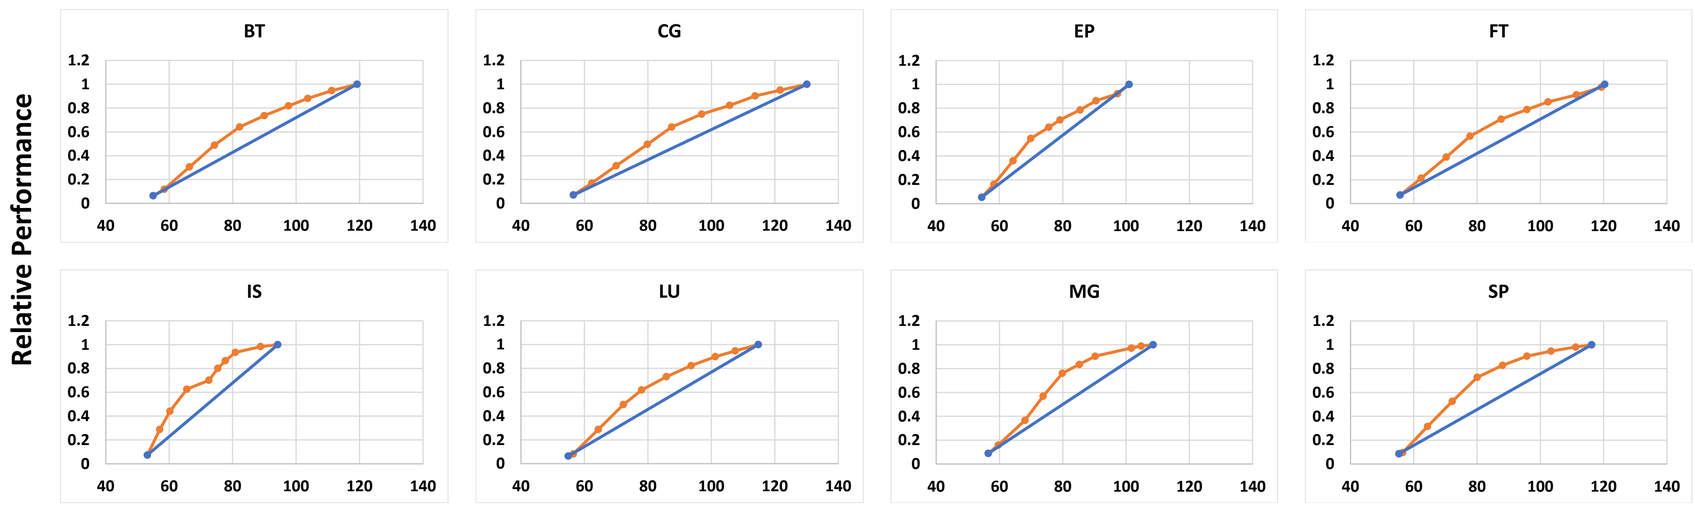

Segmented models are exceedingly more accurate than linear models in obtaining the power caps for the 2 preset performance demands

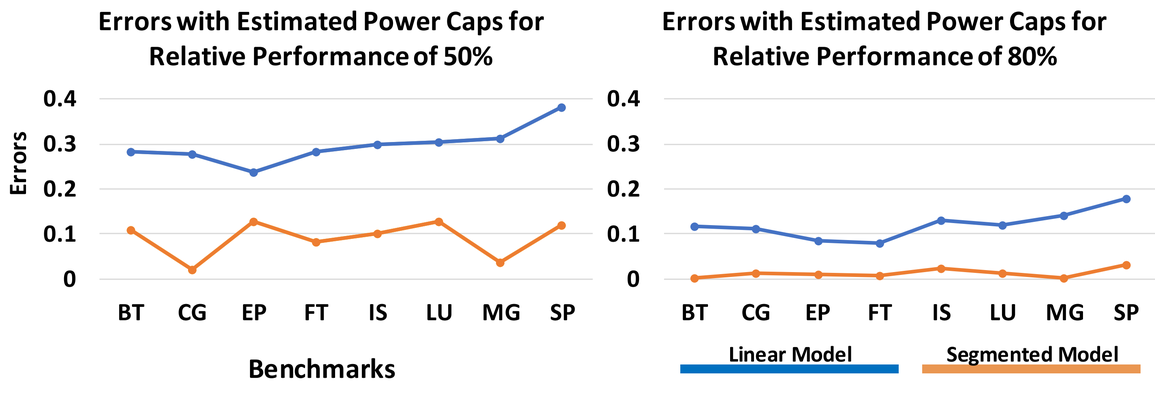

* This work was supported in part by JSPS KAKENHI under Grant Number 17K12665 and by the Natural Science Foundation of Liaoning Province under Grant Number 20180550194.