Process Variation-Aware Power Scheduling for HPC Applications

Neha Gholkar

Consortium for Advanced Simulation of LWRs
North Carolina State University
Advisor: Prof. Frank Mueller
Advisor: Dr. Barry Rountree

Motivation

➢ US DoE's power budget for exascale computing is 20MW
- → **Power constraint**
- → Hardware **overprovisioning** wrt power supply

➢ Need to optimize parallel jobs for performance under a hard power constraint

➢ Application's Performance can be quantified in terms of Instructions Per Second (IPS)

➢ **GOAL : Maximize IPS/Watt**

Problem

➢ Naive Strategy – Uniform Power Capping

➢ Power Scheduling decisions enforced by capping power at processor level
- → Intel's **R**untime **A**verage **P**ower **L**imit (**RAPL)**

➢ Under Power Bounds
- → variation in performance

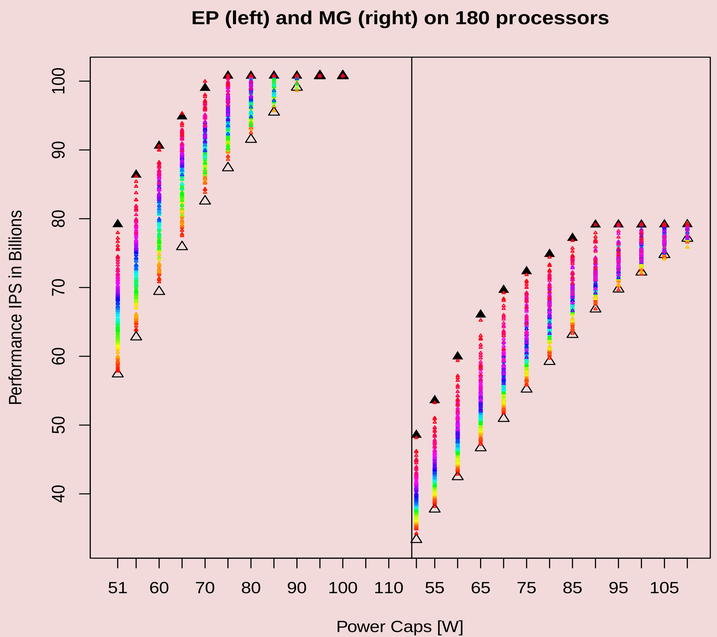

➢Under Power Bounds
- → Variation in peak power efficiency
- → Uniform Power Capping fails

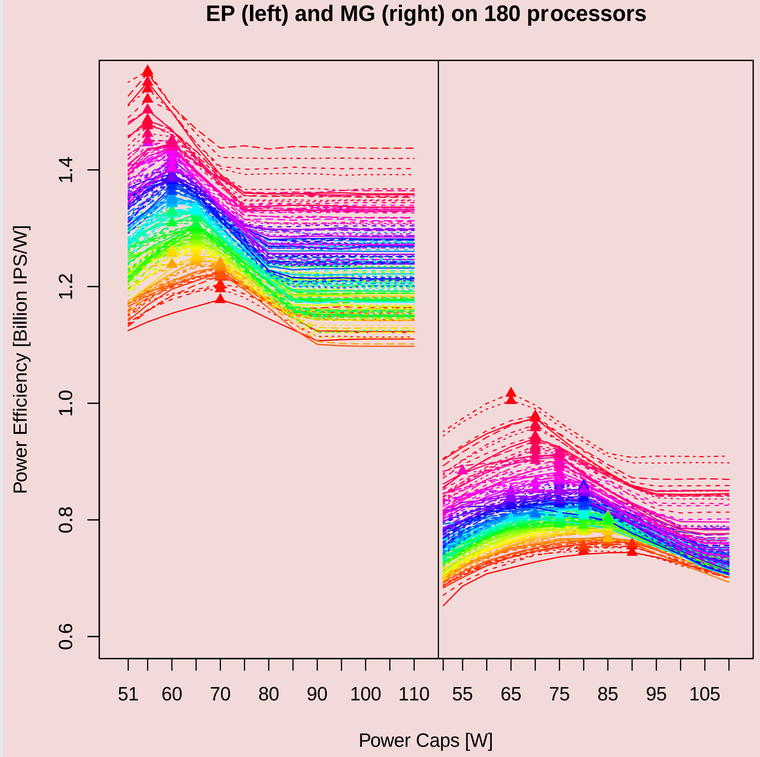

Solution : Process Variation-Aware Power Balancing

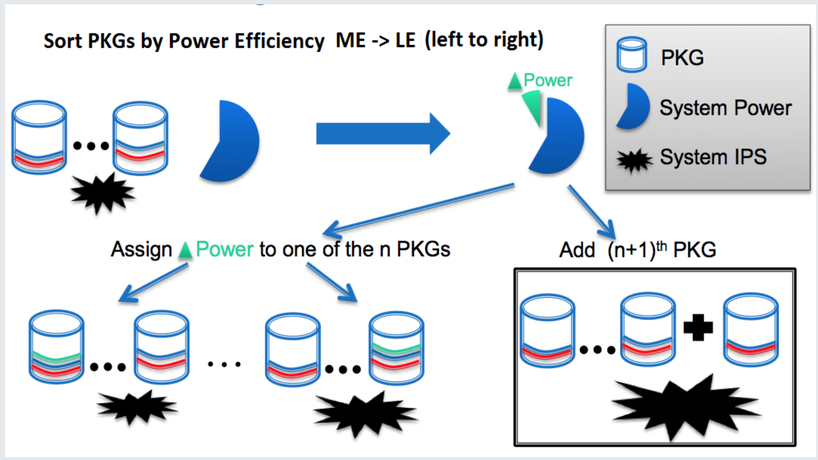

Experimental Results

➢ Results for Embarrassingly Parallel (EP) and Scalar Penta-diagonal solver (SP) from the NAS Parallel Benchmark Suite.

➢ Application Power budget : 4KW, 8KW and 16KW on 8, 16 and 32 processors respectively..

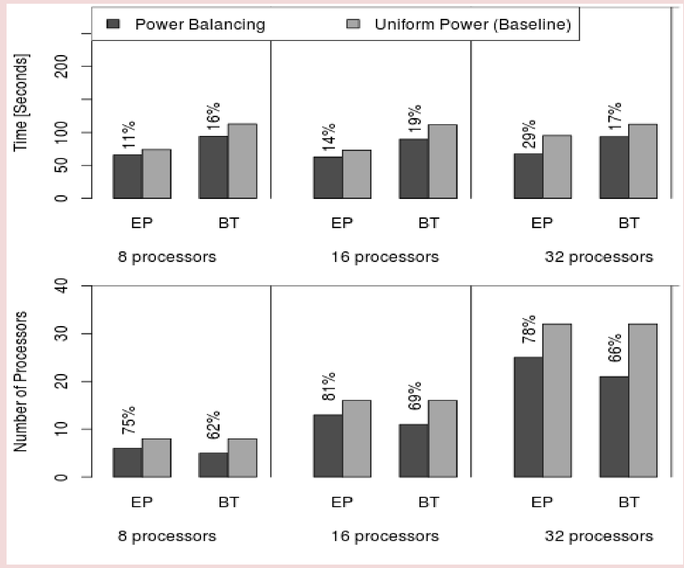

Conclusions

➢Our model improves performance by up to 29% as compared to the uniform power capping (naïve) approach.

➢ Optimal configurations are found at lower processor counts as compared to the naïve approach.

Acknowledgements

We would like to extend our thanks to Livermore Computing and CASL.