Power Shifting Opportunities on BG/Q Using Memory Throttling

Bo Li
Virginia Tech

Edgar A. León
Lawrence Livermore National Laboratory

Overview

- **Shift power to resources on the critical path**
  - Characterize code phases/regions/physics
  - Leverage IBM Blue Gene/Q's memory throttling
  - Minimize impact on time-to-solution
- **Contributions**
  - Demonstrated significant power shifting opportunities
  - First to employ real throttling on a supercomputer
  - Linear regression model to guide per-region throttling

Memory Throttling on BG/Q

- **BG/Q node architecture**
  - 16 A2 cores @ 1.6 GHz, 16 GB memory @ 1.33 GHz
- **Kernel_SetPowerConsumptionParam(X)**
  - Add X *DDR idle cycles* for each memory access
  - Node granularity
- **Power measurement infrastructure**
  - EMON2 high-resolution monitoring
  - Node-board granularity (32 nodes)

Memory Bandwidth and Throttling Latency

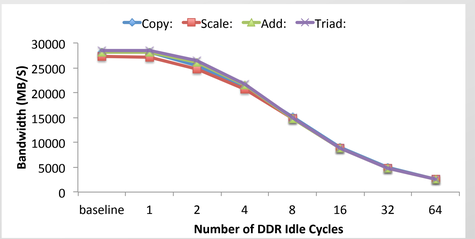

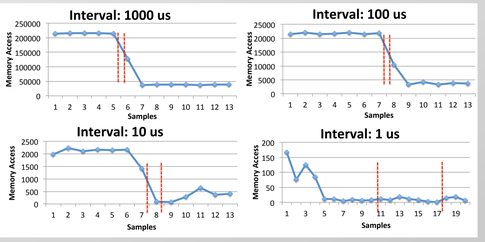

LULESH

- **Explicit hydrodynamics mini-application**
- **Five code regions**

pF3D Kernels

- **Simulates laser beams and plasma interaction**
- **Kernels:** absorb, acadv, couple4, apcpft

Optimal memory speed is a function of problem size, concurrency, and code region / kernel

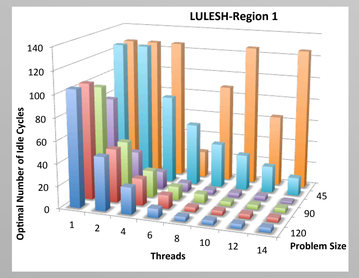

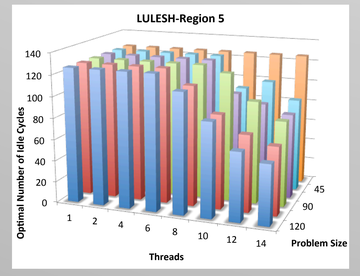

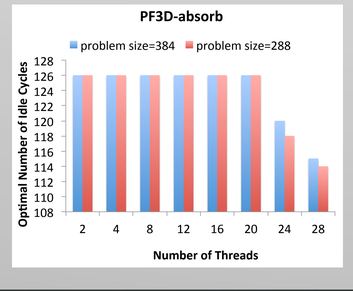

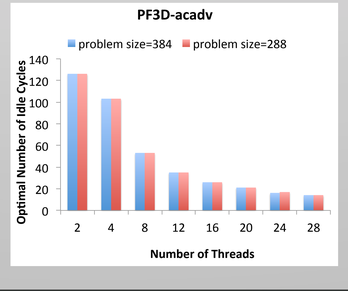

Predicting the *Optimal* Memory Throttling

- **Linear regression model based on HW counters**
  - CPU cycles, instructions
  - LSU load/store misses
  - L2 cache misses, prefetching, loads/stores
- **Initial model**

$$f_{\min} = \sum_{i=1}^{n} \left(w_i \bullet \frac{counter_i}{CPUcycles} \right)$$

- **Revised model**

$$f_{\min} = \sum_{i=1}^{n} \left(w_i \bullet \frac{counter_i}{CPUcycles} \right) + w_{n+1} \bullet \frac{comp}{mem}$$

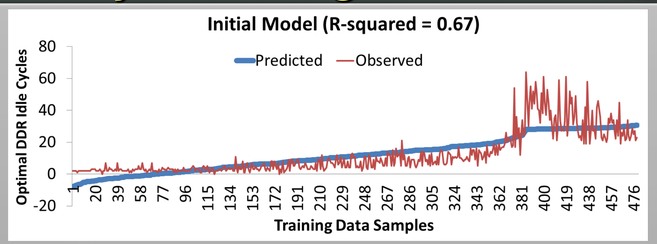

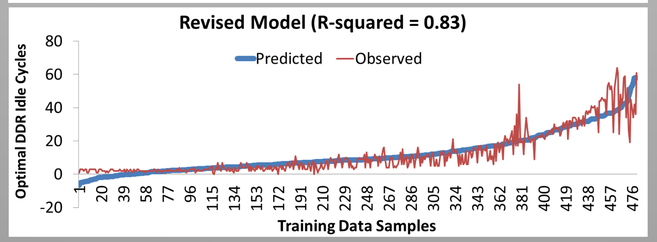

Impact on Performance & Power

Performance loss compared to the maximum memory speed

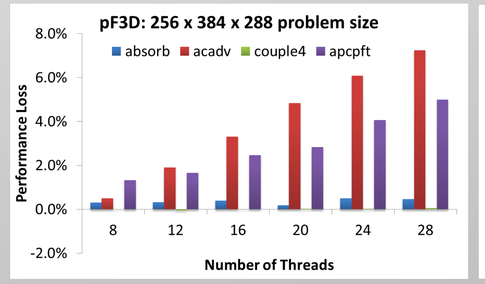

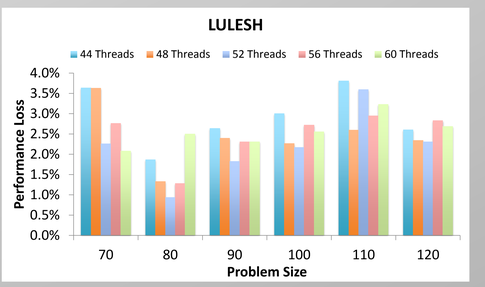

- Low performance degradation with model-based throttling
- Memory bound regions highly sensitive to memory speed
- Improved dynamic energy and power at low performance cost in LULESH

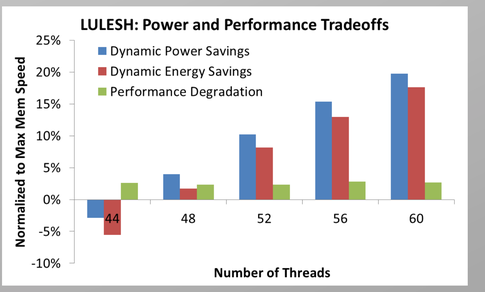

Conclusions & Future Work

- **Significant power shifting opportunities with model-based throttling**
  - Memory throttling applied on a per-region or per-kernel basis
  - Model predictions result in low performance degradation
  - Savings of up to 20% dynamic power with 3% performance cost
- **Future work**
  - Study a broader set of applications
  - Shift power on architectures with other throttling capabilities
  - Evaluate non-linear models based on machine learning

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-POST-658188.