TAPEMS, 2017

2nd International Workshop on Theoretical Approaches to Performance Evaluation, Modeling and Simulation

CCGrid2017

Empirical Mode Decomposition for Modeling of Parallel Applications on Intel Xeon Phi Processors

Masha Sosonkina

Joint work with: Gary Lawson, Tal Ezer, Yuzhong Shen

Old Dominion University
Department of Modeling, Simulation and Visualization Engineering
Norfolk, Virginia, USA

MSVE modeling simulation visualization engineering OLD DOMINION UNIVERSITY

Acknowledgements

AFOSR
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

U.S. DEPARTMENT OF ENERGY | Office of Science

THE Ames Laboratory
Creating Materials & Energy Solutions
U.S. DEPARTMENT OF ENERGY

National Science Foundation
WHERE DISCOVERIES BEGIN

Outline

1. Motivation
2. Empirical Mode Decomposition
3. Model Construction
4. Experiment Overview
5. Measured vs. Modeled Results
6. Future Directions

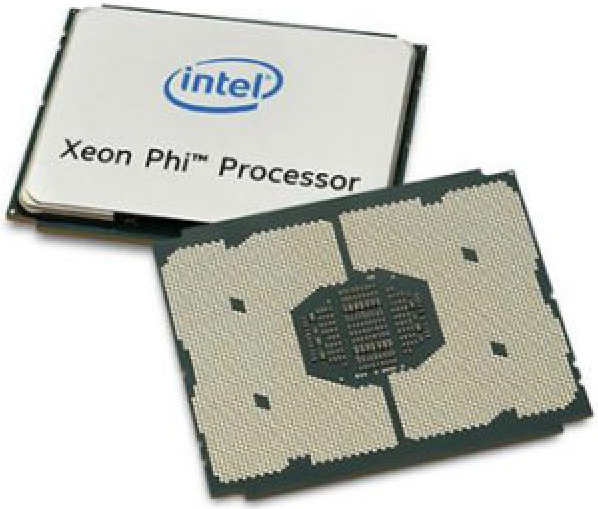

Motivation

- Existing models primarily focus on performance
  - Number crunching was the bottleneck until the mid-2000's
  - Current hardware bottlenecks include
    - Memory bandwidth
    - Power usage

- There is a need for a new model which includes both power and performance
  - High-level abstraction of hardware and application
  - Focus on power and energy consumption

Typical cluster power draw when fully occupied

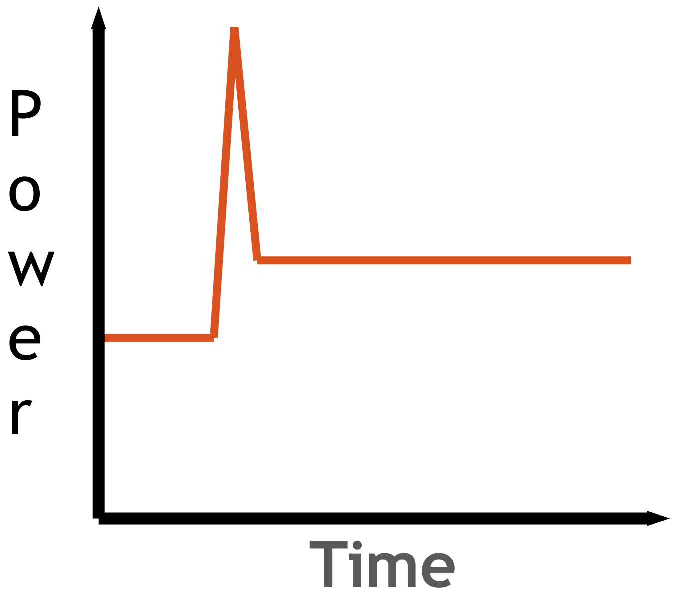

Ellsworth et. al., 2016

Modeling Multi-Core Energy

Not easily extendable to real-world applications

- Roofline
  - Operational intensity
  - Support for power and energy
    - Vuduc *et al.*, 2014
- Device Specific
  - Advanced knowledge of the architecture
  - Instruction-level modeling
    - Shao and Brooks, 2009
- McPAT
  - Expert knowledge of the architecture
  - State-of-the-art modeling framework
    - Li *et al.*, 2009

Empirical Mode Decomposition (EMD)

Decompose a non-parametric non-stationary time-series into intrinsic mode functions and the residual trend

[N. Huang, et al. *The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis*, 1998.]

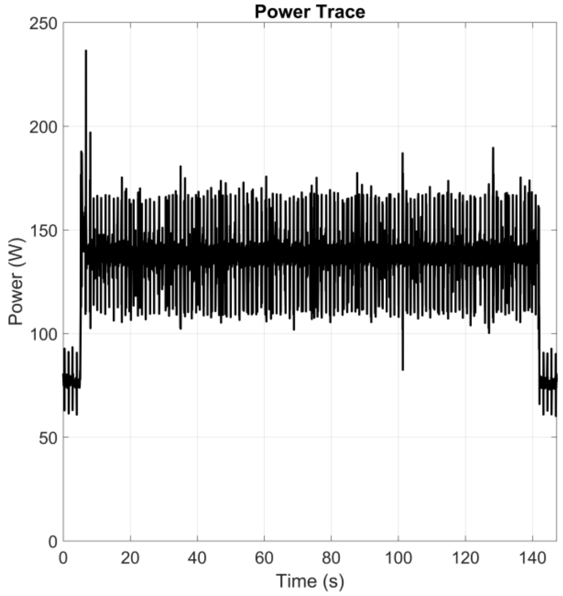

(a) Power Trace for CoMD

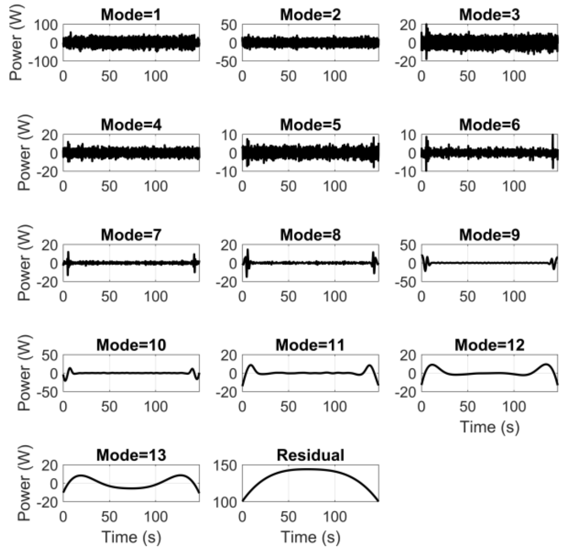

(b) IMF Amplitudes for CoMD

- Used in many scientific fields:
  - Medicine, Finance, Geosciences
- Used to uncover correlations in physical phenomena

Empirical Mode Decomposition (cont'd)

- Power draw is a complex physical response to the workload of a given application on the hardware

- Final result of EMD is the residual trend
  - For power draw, this trend shows power consumption over time when high-frequency interactions have been removed
  - Can be modeled using quadratic equation

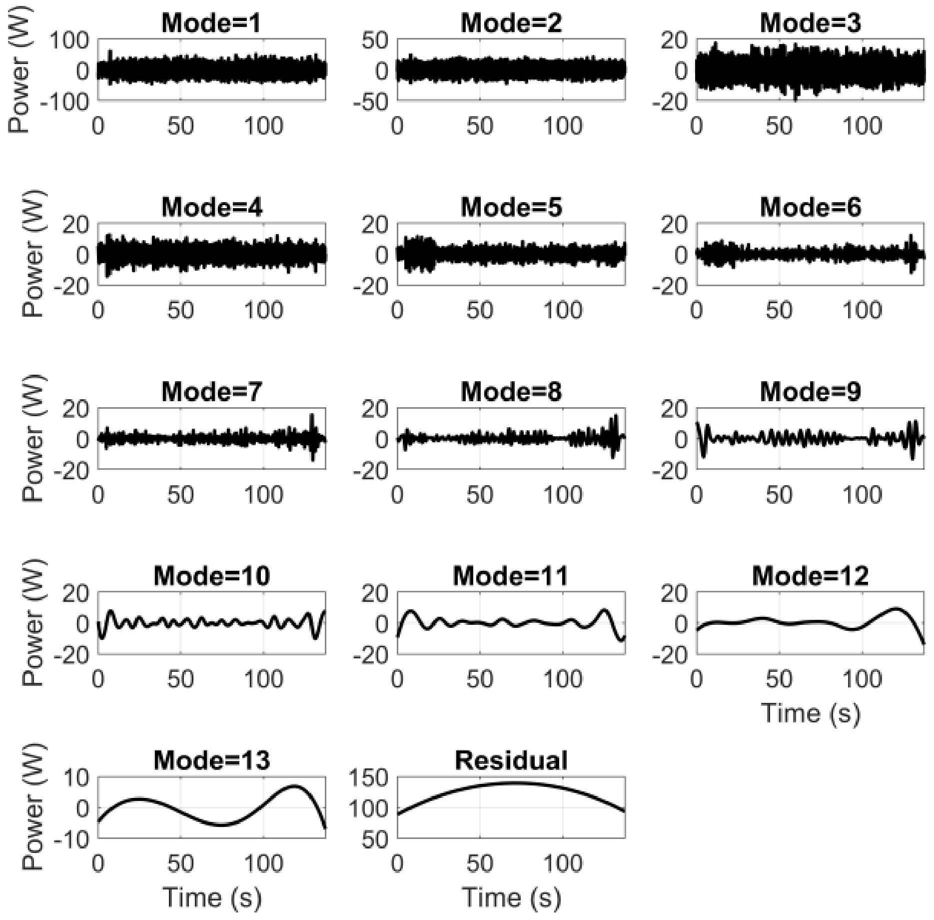

Ensemble Empirical Mode Decomposition (EEMD)

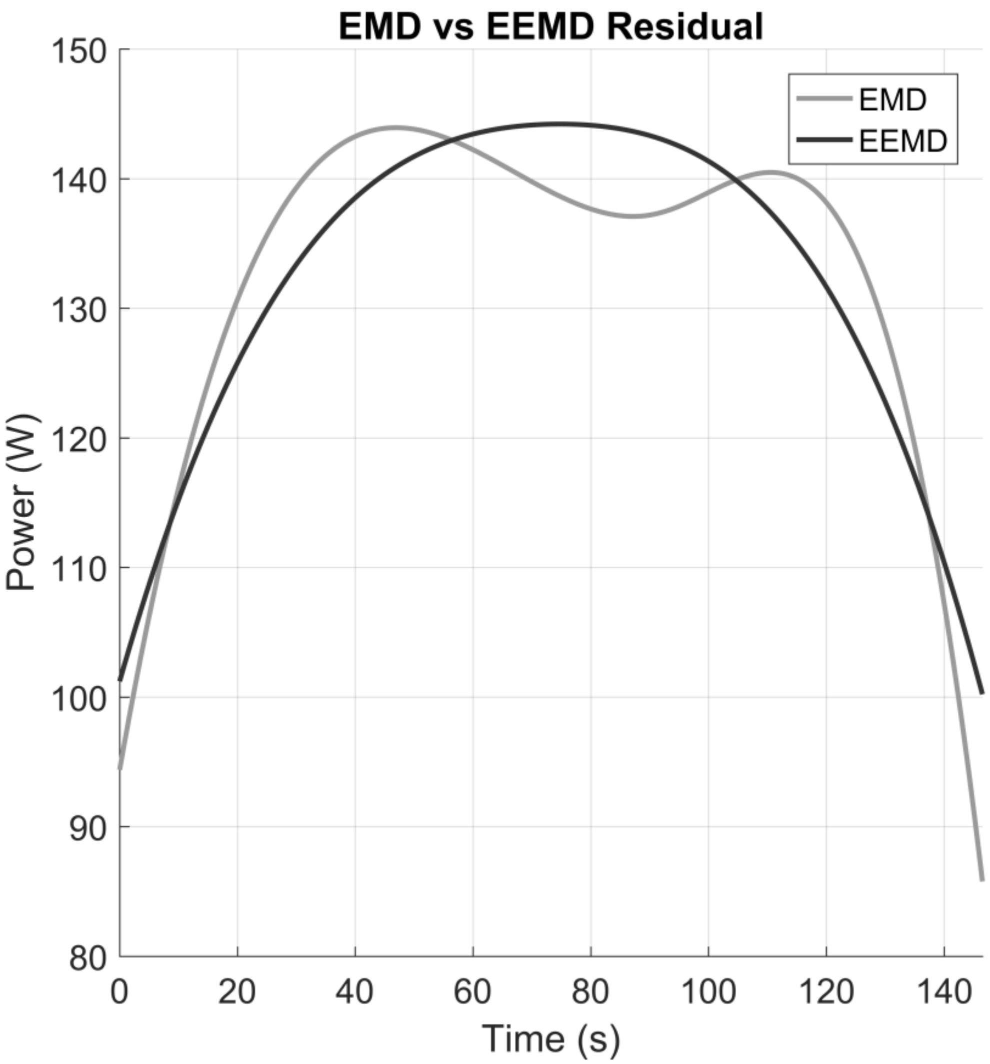

- EMD strongly depends on variations in frequency and amplitude of the time-series to decompose all IMFs
  - EEMD alleviates this dependency

- EEMD [Wu & Huang, 2009]
  - Uniform white noise is added to the source signal before IMFs are decomposed
  - EMD is performed on this new signal many times
  - The resulting IMFs are the mean of all the IMFs

Model Construction

1. Choose an execution configuration
   - E.g., Number of cores, power limit
2. Run a few times (*n*)
   1. Obtain Power Trace using Measurement Procedure
   2. Apply EEMD
      - 5 W @ *N* = 50
   3. Collect Residual Trend
3. Fit quadratic model to the *n* residuals
   - Including idle power measurements in the power trace enables a quadratic fit

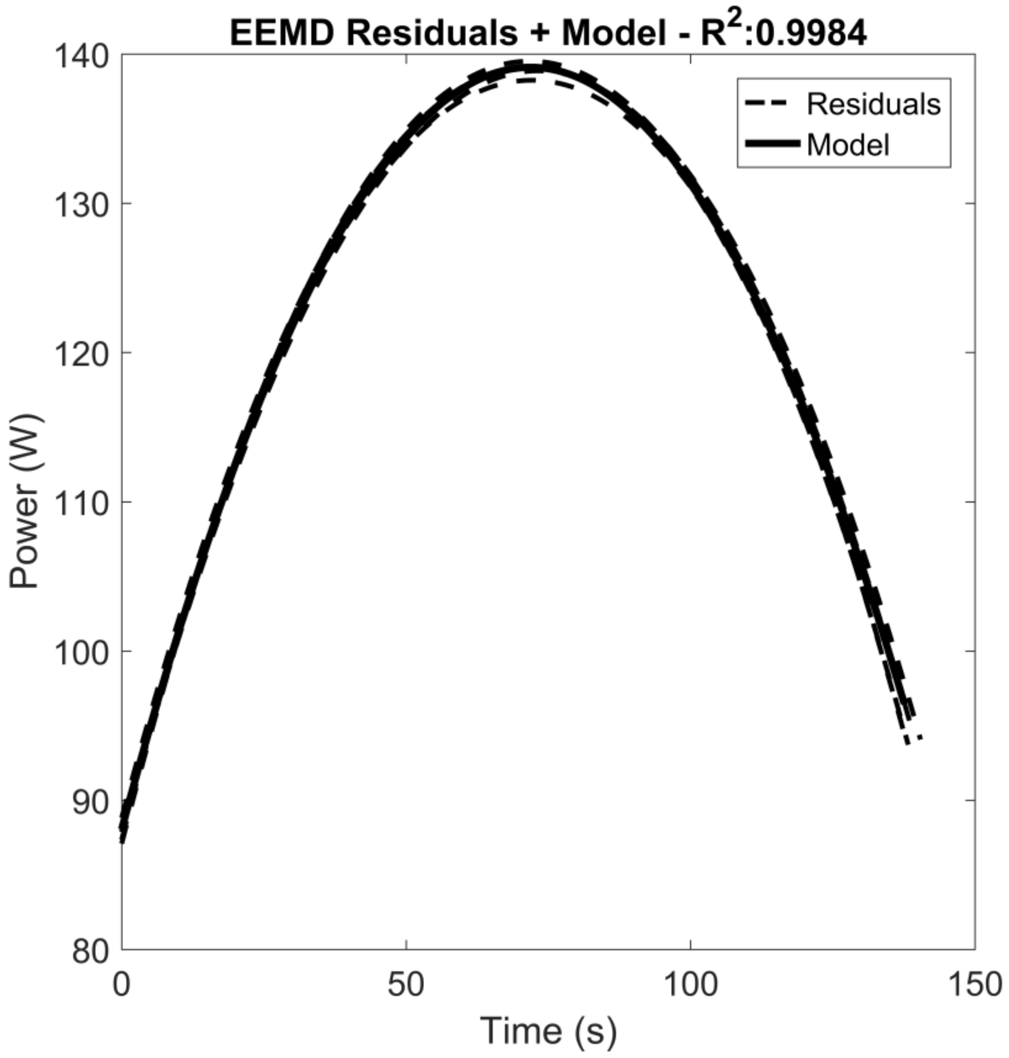

Model Construction (cont'd)

Suppose no idle power measurements

Same Configuration

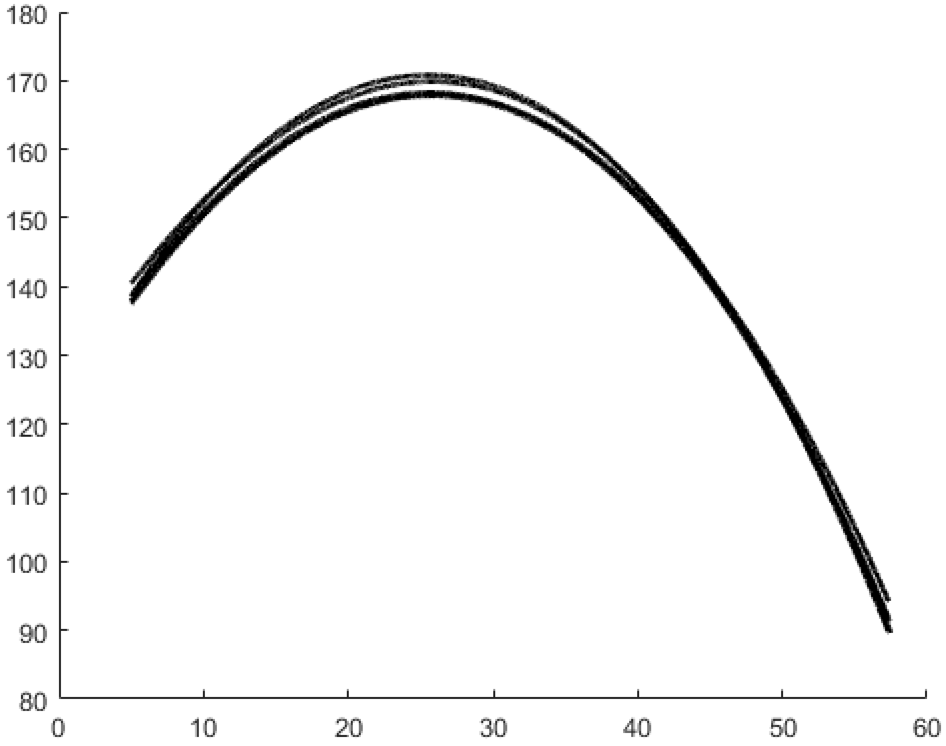

Varying # Cores

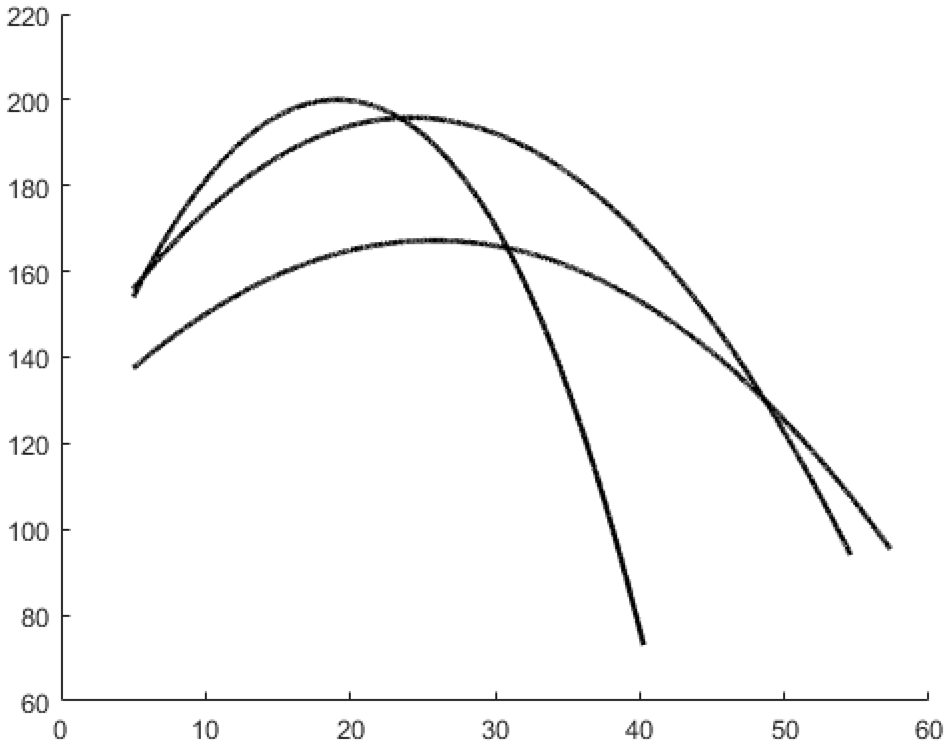

Measurement Procedure

1. Begin power sampling
   - Sandia Power API
     - RAPL via Linux Power Capping Framework
     - 5ms sampling rate
     - Power = DRAM + Core
2. Wait 5 seconds
   - Capture idle power
3. Begin application
   Wait for end
4. Wait 5 seconds
   - Capture idle power
5. End power sampling

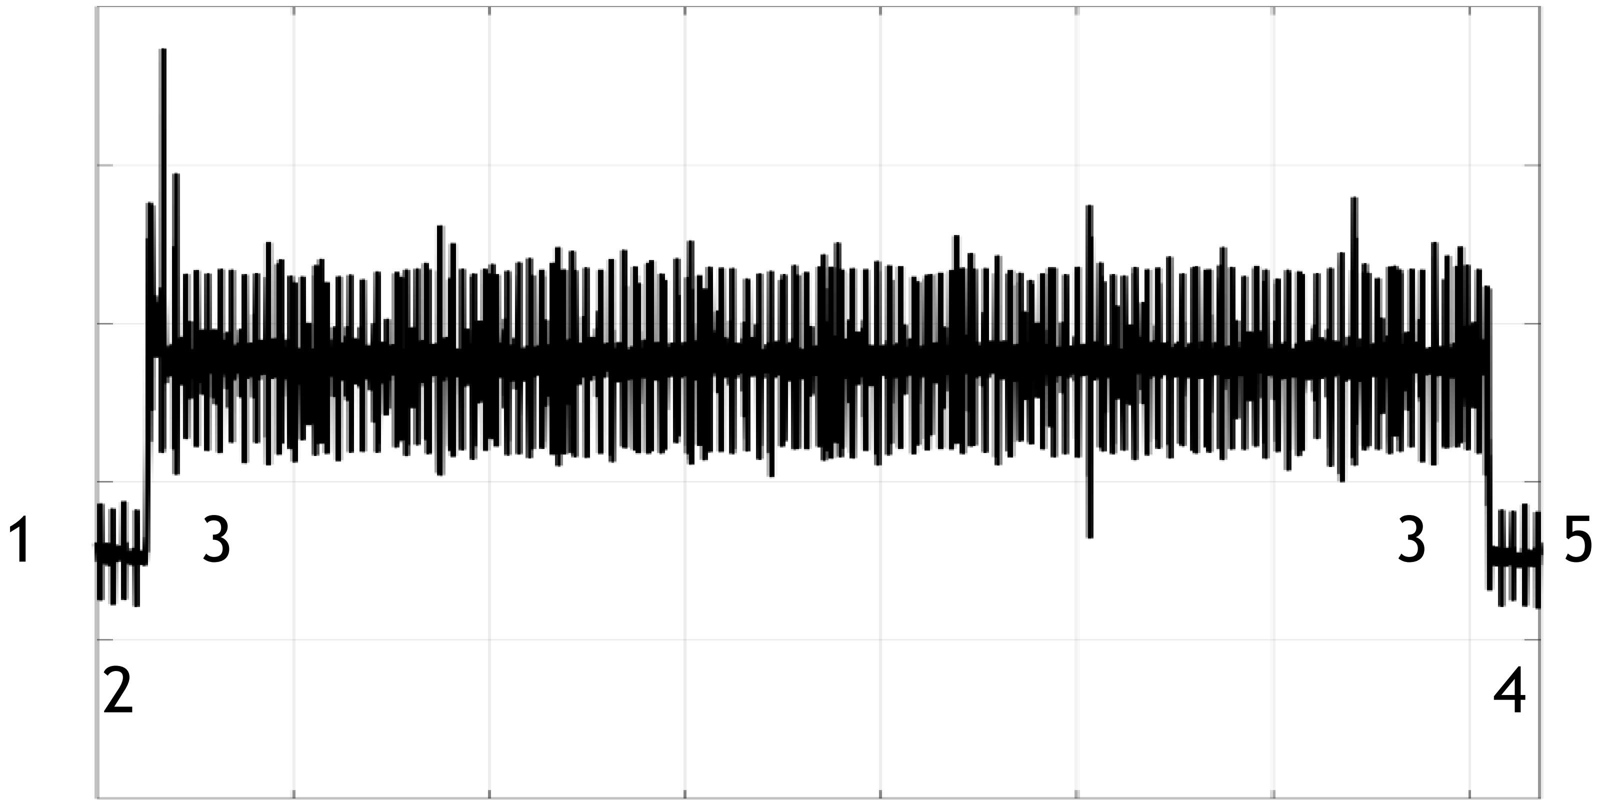

Experiment Specifications

Computing Platform

- *Rulfo* single-node system at Old Dominion University
- Xeon Phi Processor 7210
  - 64 cores @ 1.3 GHz
    per core:
    - 4 HW threads
    - 2 512-bit VPU
  - 32 MB L2 cache
  - 16 GB MCDRAM
  - 215 Watt TDP

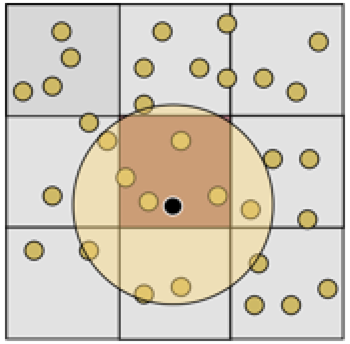

- **Application** configurations
  - **CoMD**
    - Force Kernel: LJ, EAM
    - Problem Sizes: 60, 80, 100
  - **GAMESS**
    - Problems:
      1L2Y, 20w, S265, S301

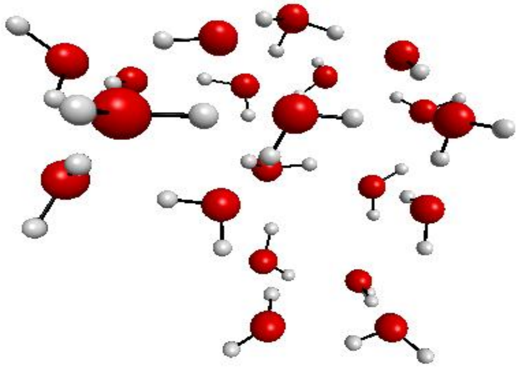

- **Hardware** configuraions
  - Number of Cores: 32, 48, 63
    - 1 core dedicated to power measurements
  - DRAM Memory: DDR, MCDRAM (flat)
  - Power Limiting: 140, 130, 120, 110, 100, 90

Modeled vs. Measured Results (CoMD)

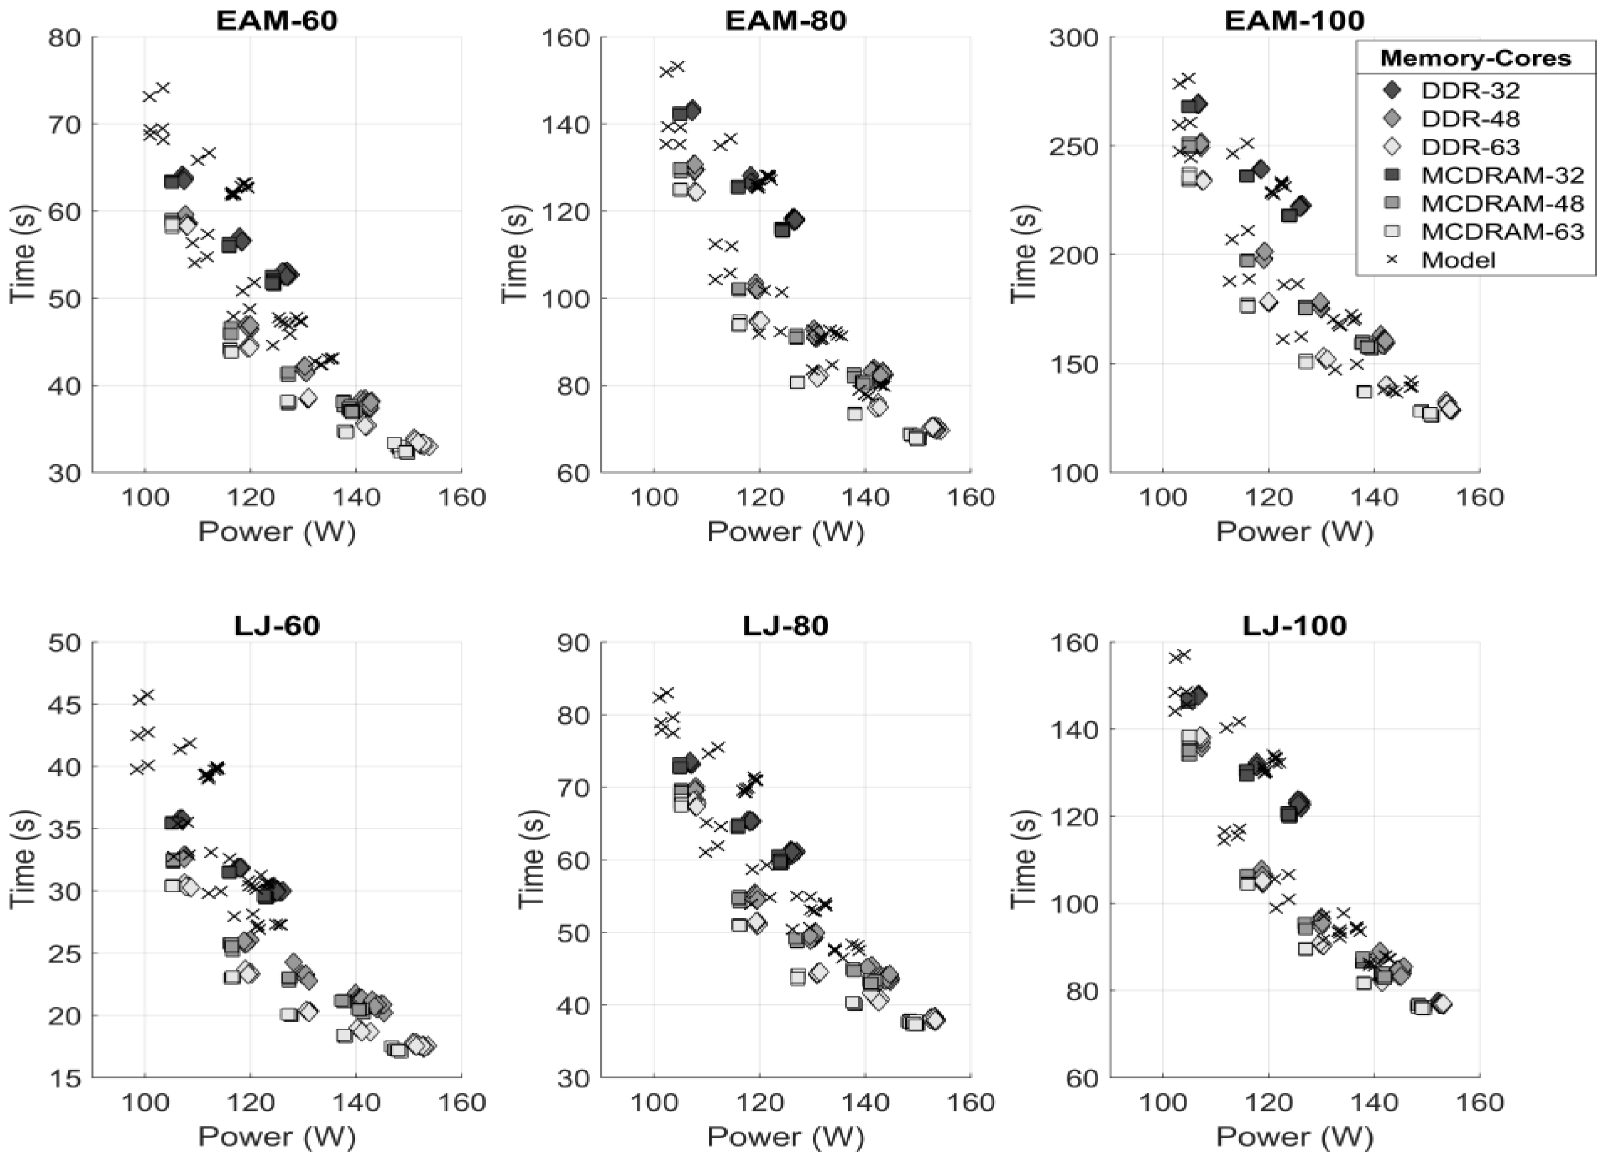

(a) CoMD

Modeled vs. Measured Results (GAMESS)

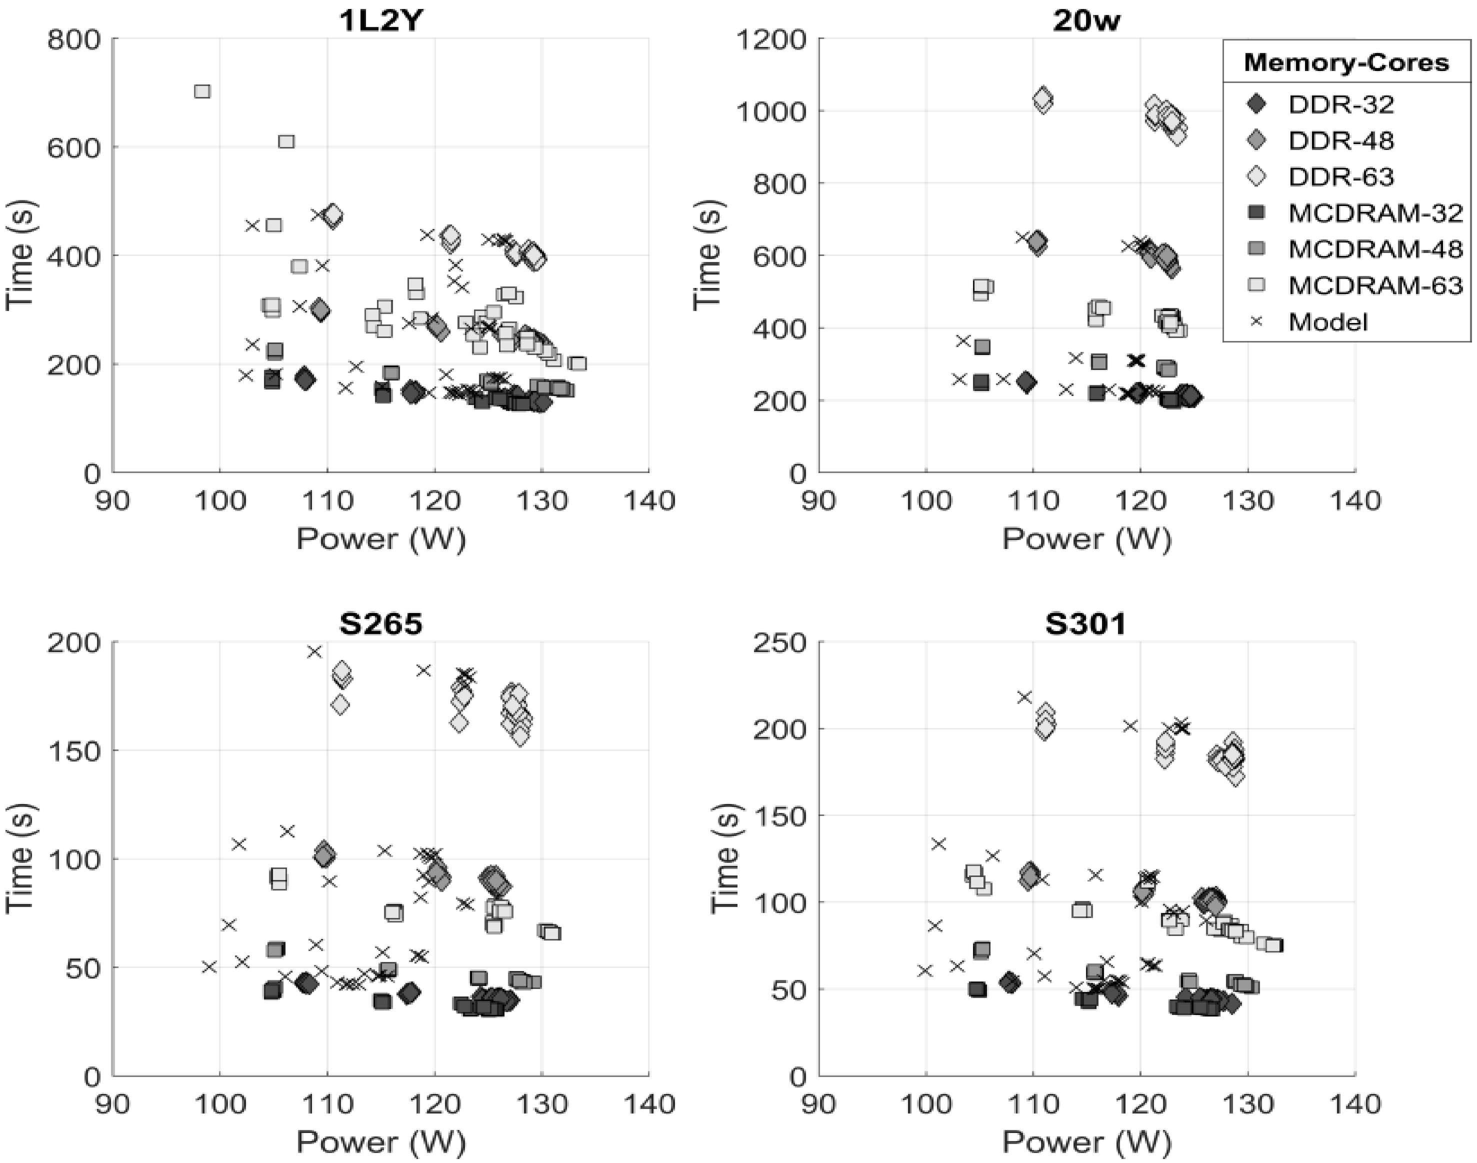

(b) GAMESS

Measured Results (MCDRAM)

MCDRAM (flat-mode) provides best performance vs. DDR

- 1-2% energy savings
- 1-3W power reduction
- 1-2% time reduction

Measured Results (CoMD & GAMESS)

CoMD

- Max cores (63)
  - To max performance
- 120W power limit

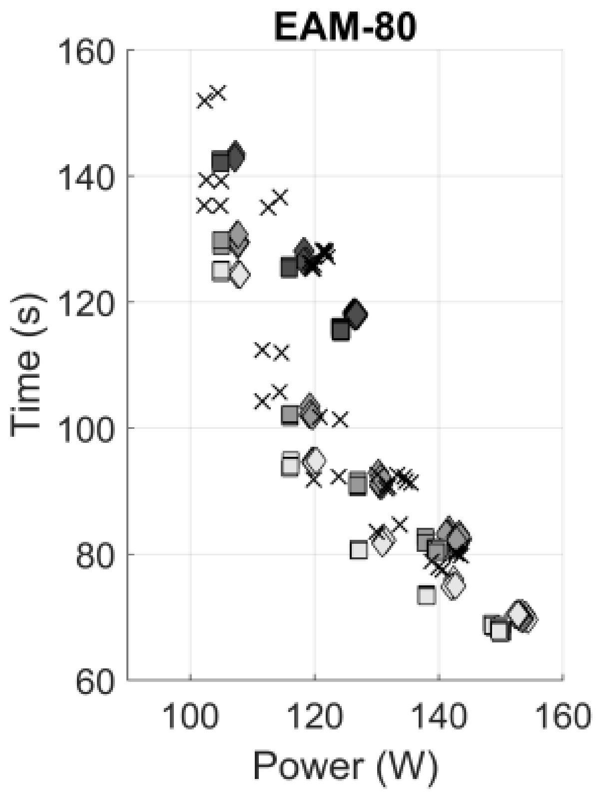

GAMESS

- Min cores (32)
  - To max memory bandwidth
- Default power (215W)

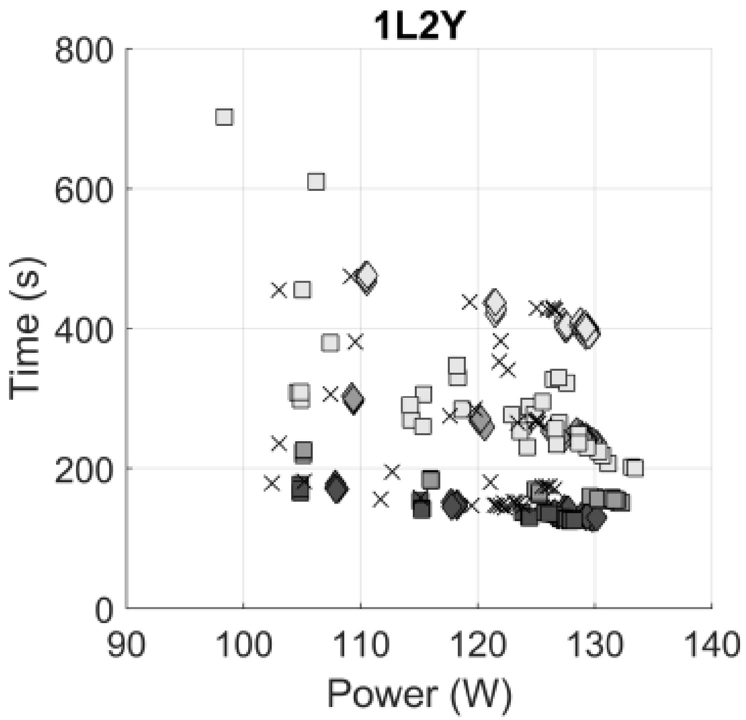

Modeled Results

Quadratic Model

- Time
  - Coefficient $a < 0$
  - Direct relationship to $|a|$
- Power
  - Coefficient $b > 0$ (max power)
  - Coefficient $c > 0$ (min power)
  - Direct relationship with $|b|$ & $|c|$

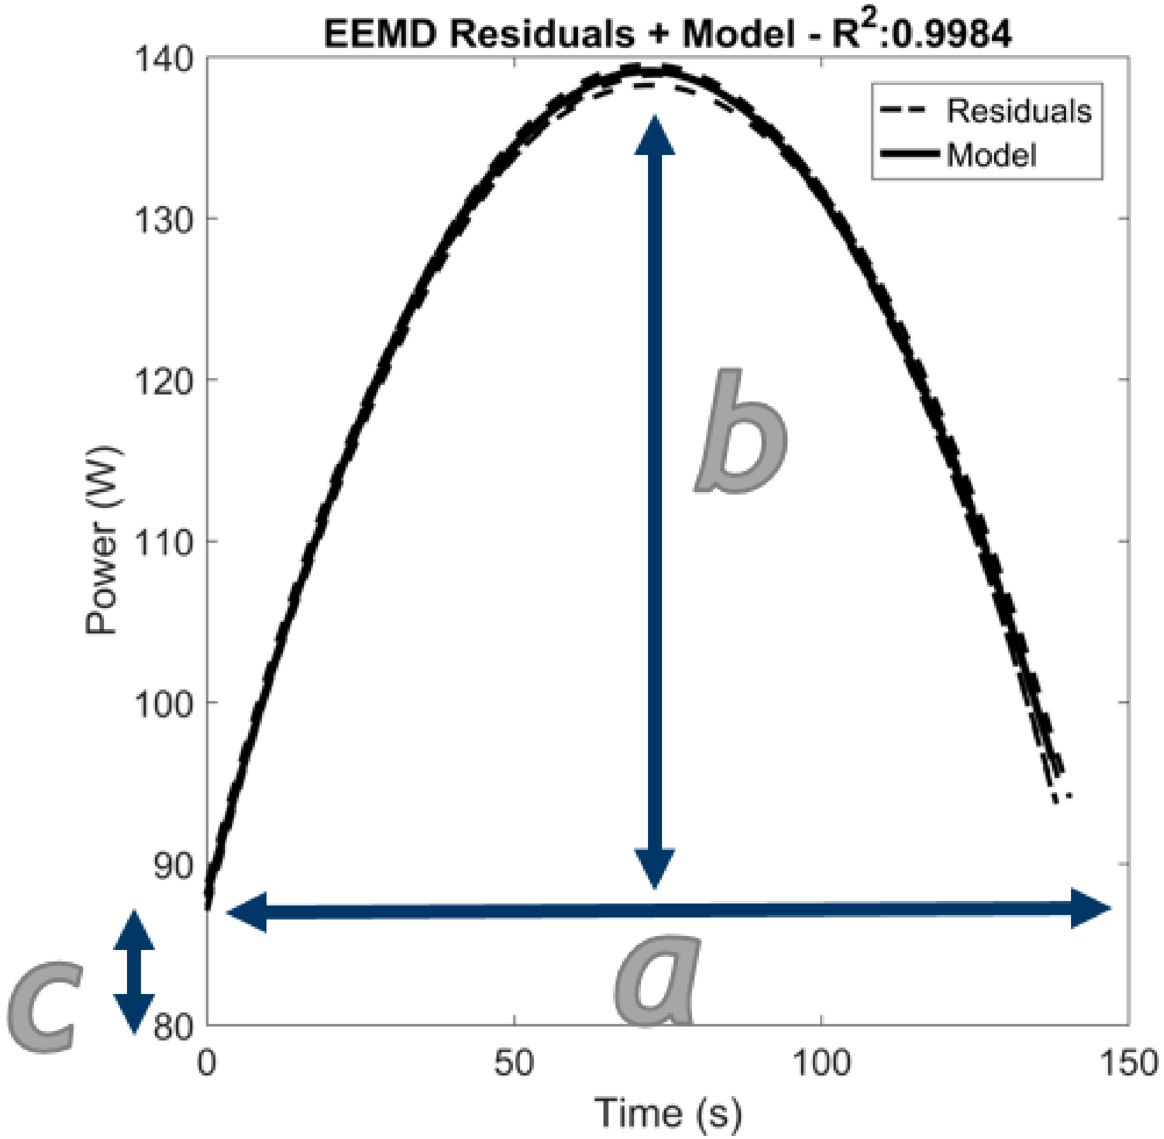

Model Error

$$Error = \frac{measured - modeled}{measured}$$

Power

- No more than 10% error
- Is under-estimated
  - Residual does not go up to max power draw

Time

- No more than 15% error
- Is over-estimated
  - Model time assumes power draw at the start and end of the trace is equal

CoMD

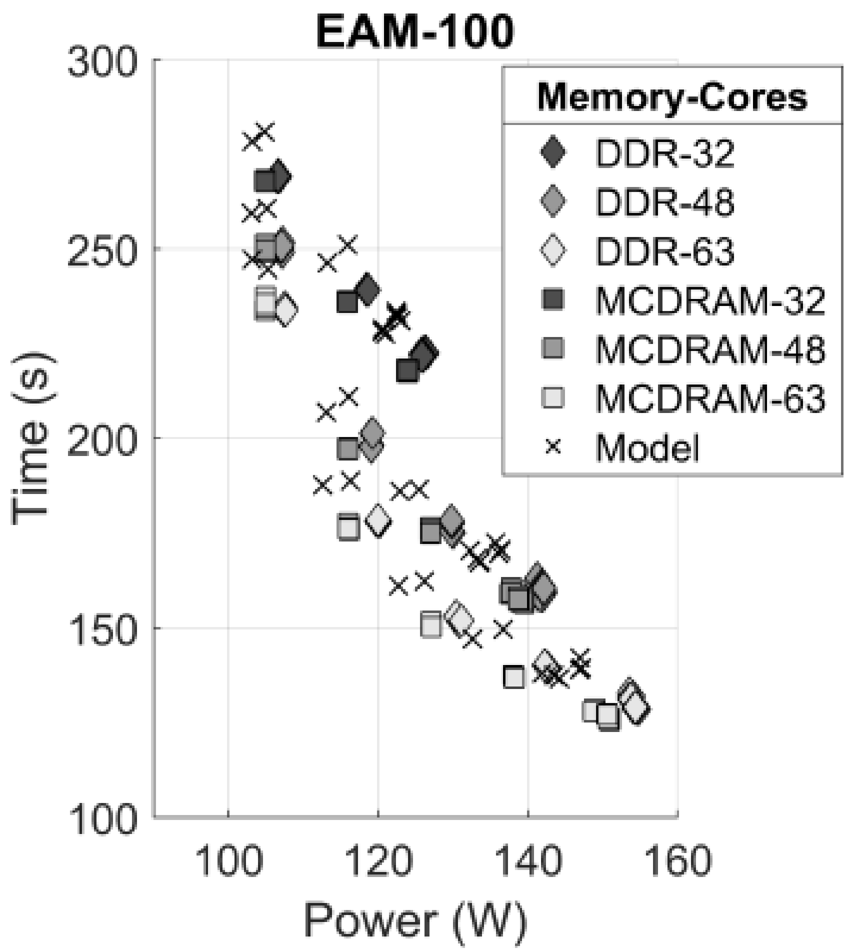

Model Error (cont'd)

$$Error = \frac{measured - modeled}{measured}$$

Energy

- 5% – 30% error
- Shorter traces incur higher error than longer traces
  - Short trace < 100 seconds
  - Long trace > 100 seconds

GAMESS

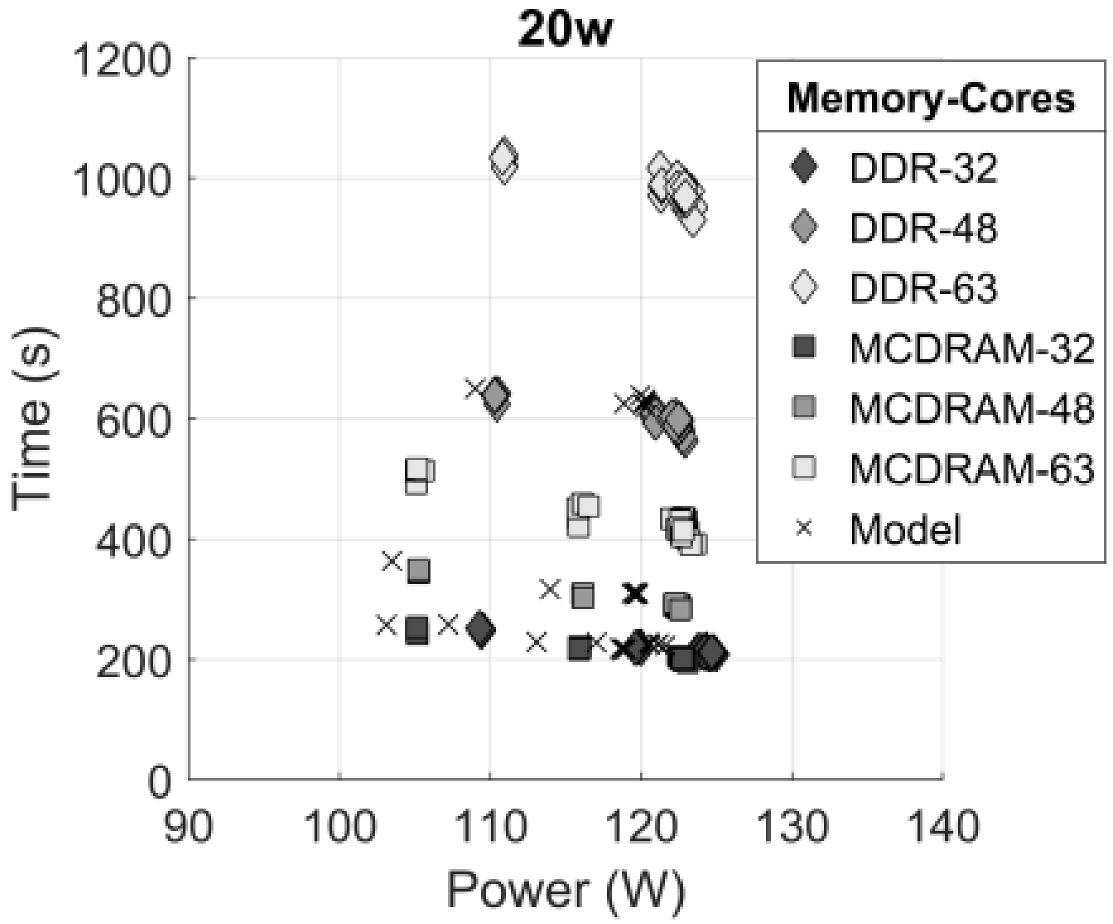

Advantages of Modeling Energy with EMD

Generality

- Applies to any hardware architecture
  - Multi-core CPU, Accelerator, or Hybrid
  - Single-node, Multi-node
- Applies to any software application
  - Benchmark
  - Kernel
  - Mini-App
  - Real-World

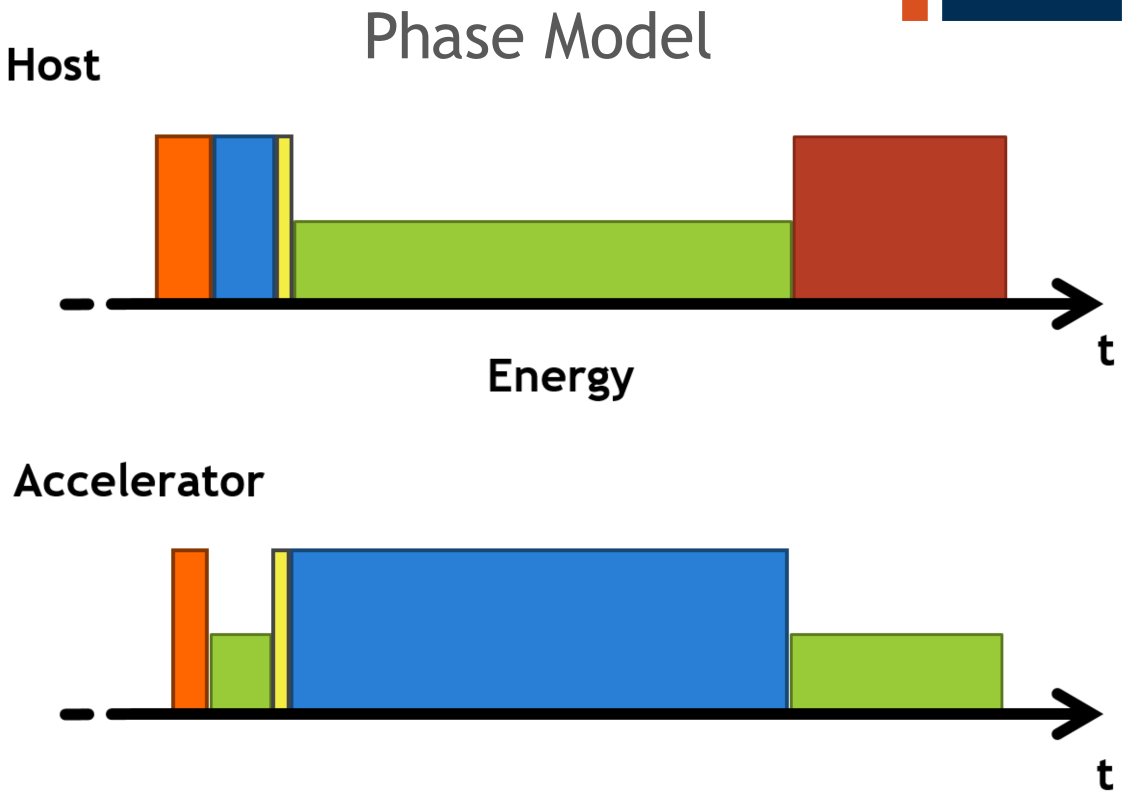

Lawson, Sundriyal, MS, and Shen. 2015.
Modeling performance and energy for applications offloaded to Intel Xeon Phi.

Advantages of Modeling Energy with EMD (cont'd)

Rigorous

- Guided specific measurement procedure
- EMD residual fits to quadratic function

Ease-of-Use

- Only requires a power trace to produce a model
- No advanced knowledge of application or hardware required

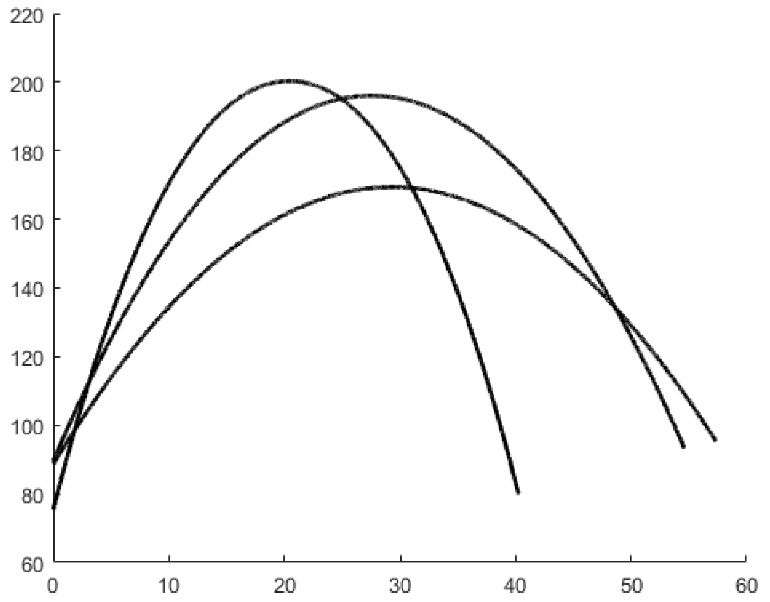

Performance

- Manageable overhead
  - Known to increase with trace size
  - Alleviated by running EMD on segments

Future Directions

- Refine model construction and definition
- Investigate applicability of IMF modes
- Enhance the model with predictive capabilities

References

D. Ellsworth, T. Patki, M. Schulz, B. Rountree and A. Malony, "A Unified Platform for Exploring Power Management Strategies," *2016 4th International Workshop on Energy Efficient Supercomputing (E2SC)*, Salt Lake City, UT, 2016, pp. 24-30.

Z. Wu and N. Huang. Ensemble empirical mode decomposition: A noise-assisted data analysis method. Advances in Adaptive Data Analysis, 01(01):1–41, 2009.

G. Lawson, V. Sundriyal, M. Sosonkina, and Y. Shen. 2015. Modeling performance and energy for applications offloaded to Intel Xeon Phi. In *Proceedings of the 2nd International Workshop on Hardware-Software Co-Design for High Performance Computing* (Co-HPC '15). ACM, New York, NY, USA, , Article 7 , 8 pages. DOI: https://doi.org/10.1145/2834899.2834903

n

n

Questions?

MCDRAM

Multi-Channel DRAM

- Stackable Memory
- 16 GB
- 3 Modes of Operation
  - Flat
  - Cache
  - Hybrid

- Flat-Mode (this work)
  - Use MCDRAM over traditional DDR
  - `numactl –membind 1 ./<application command>`
- Cache-Mode
  - L3 Cache Level
  - Specific variables (in code) assigned to MCDRAM
    - Requires explicit code changes
- Hybrid-Mode
  - Split MCDRAM 25/75 or 50/50 between Flat and Cache modes

EMD / HHT Analysis

CoMD on Ivy-Bridge (max cores)

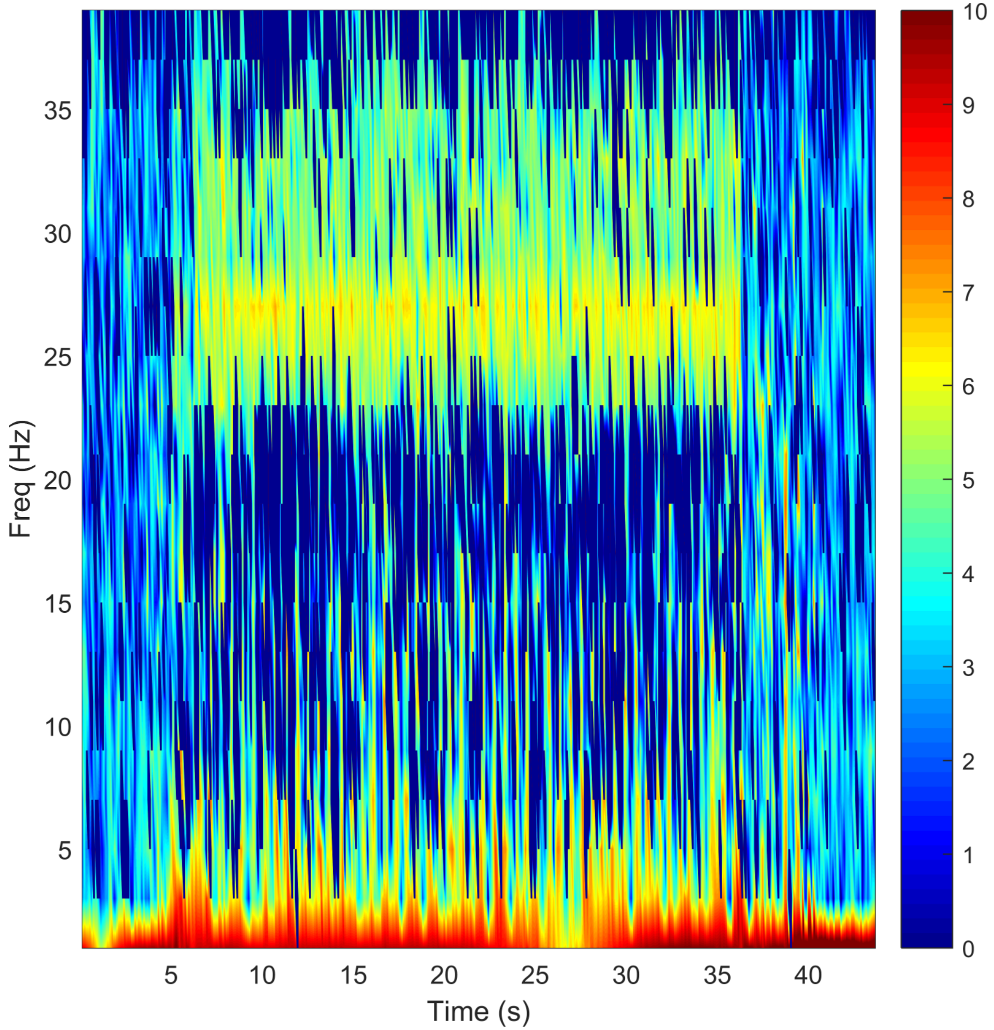

GAMESS on KNL (max cores)