How Much Power Does your Server Consume? Estimating Wall Socket Power Using RAPL Measurements

Kashif Nizam Khan

Zhonghong Ou, Mikael Hirki, Jukka K. Nurminen, Tapio Niemi

Aalto University
School of Science

HELSINKI INSTITUTE OF PHYSICS

Motivation

- The Large Hadron Collider produces 30 petabytes of data every year
- CERN uses 1.3 terawatt hours of electricity annually.
- Datacenters in the U.S. used 91 billion kilowatt-hours of electricity in 2013

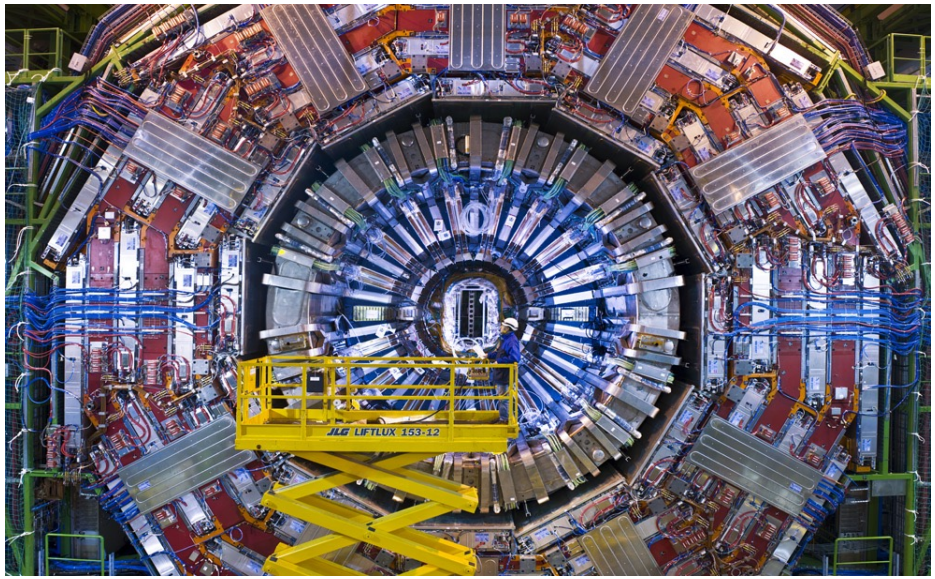

Motivation

- The Large Hadron Collider produces 30 petabytes of data every year
- Datacenters in the U.S. used 91 billion kilowatt-hours of electricity in 2013
- CERN uses 1.3 terawatt hours of electricity annually.

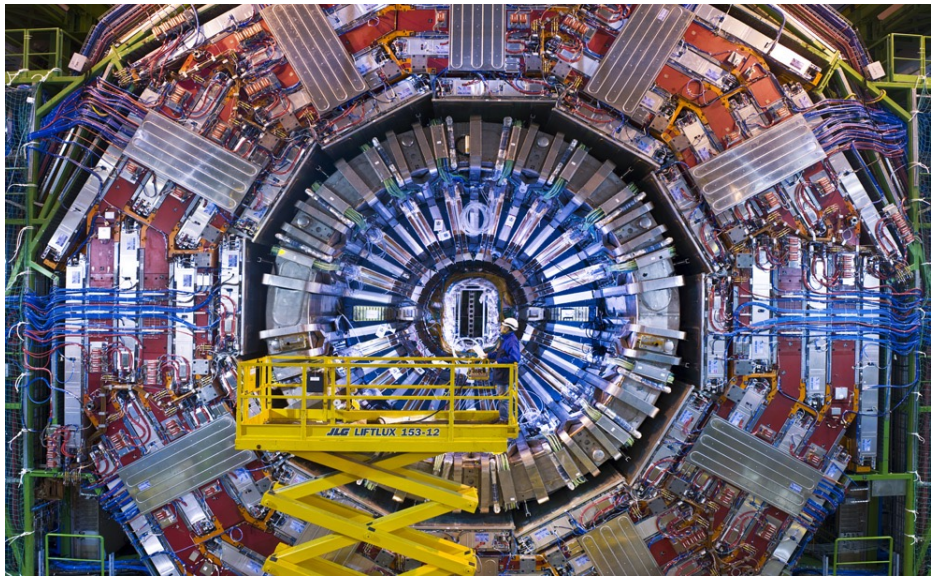

How much energy is consumed?

Motivation

- The Large Hadron Collider produces 30 petabytes of data every year
- Datacenters in the U.S. used 91 billion kilowatt-hours of electricity in 2013
- CERN uses 1.3 terawatt hours of electricity annually.

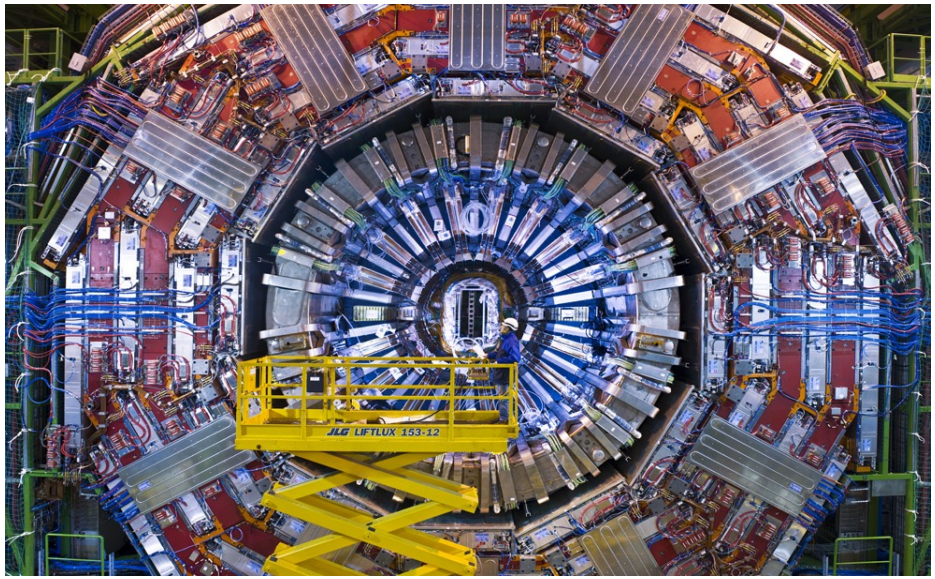

How much energy is consumed?

Where is the energy spent?

How to measure the energy consumption?

- External devices/energy meters
- Energy sensors
- Modeling power consumption with performance counters

Problems

- Instrumentation can be expensive
- Hinders normal operation of the system
- Accuracy is relative to performance degradation

Idea

To predict the wall socket power consumption without minimal interruption and high accuracy

Methodology

- Leverage RAPL to predict full system power consumption from the wall socket

- Our method:
  - Carefully designed experiments reveal the correlation between processor package power and wall socket power
  - Propose a model to predict the wall socket power
  - Verify the model using a gamut of diversified benchmarks and applications

Contributions

- We propose a predictive model to estimate wall socket power from processor package power, with high accuracy.
- Our prediction model achieves 5.6% error rate
- Advantages:
  - Minimal interruption
  - Easily executable
  - Allocate proper energy budget
  - Power limit to best utilize electricity pricing variations

Intel RAPL

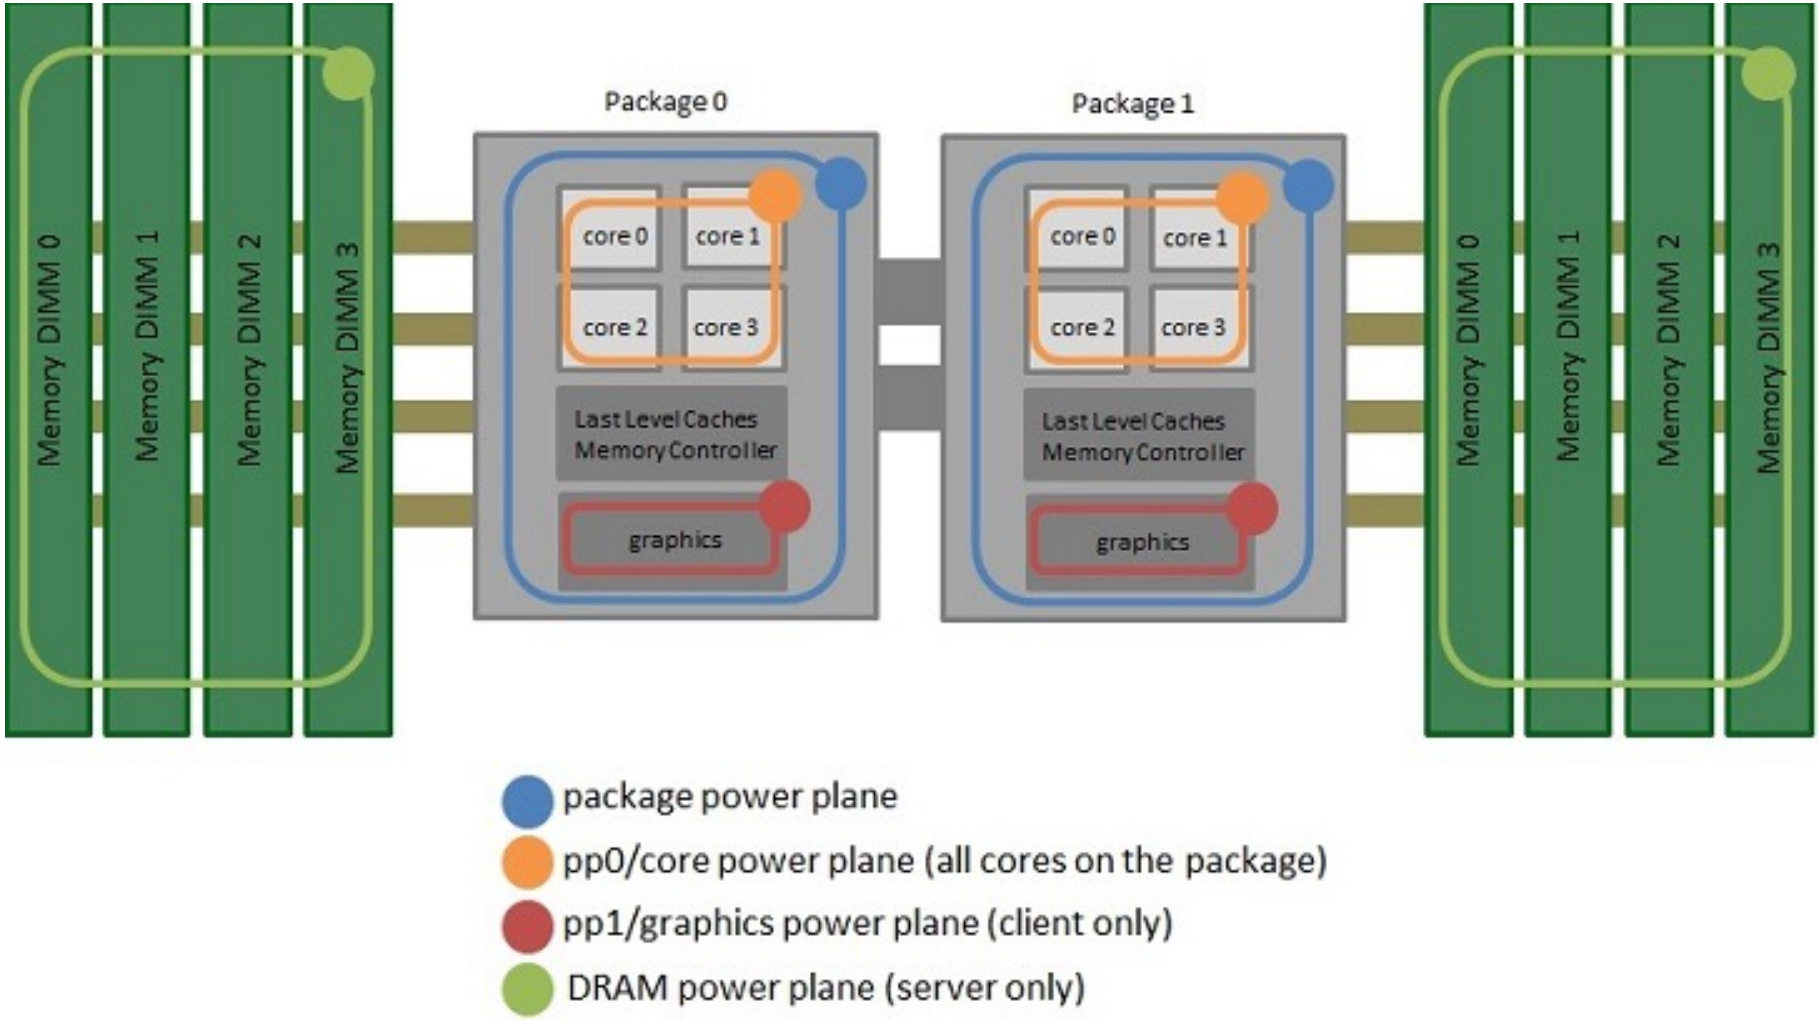

https://software.intel.com/en-us/articles/intel-power-governor

System and benchmark specifications

Processor (Intel)	Sockets	Cores	Hyperthreads	Frequency Range	L3 Cache	Memory	Tag
Core i7-4770	1	4	4	0.8–3.4 GHz	8 MB	16 GB	Machine 1
Xeon E3-1230	1	4	4	0.8–3.3 GHz	8 MB	16 GB	Machine 2
Xeon E5-2650	2	16	16	1.2–2.6 GHz	40 MB	64 GB	Machine 3

Benchmarks

- Stress-ng
- Stream
- ParFullCMS
- Parsec

System and benchmark specifications

Stress-ng: Stress the CPU cores with 100% work- load

Web site. http://kernel.ubuntu.com/~cking/stress-ng/.

System and benchmark specifications

Stream: Understand the characteristics of different systems in terms of power consumption when running a memory intensive task.

Web page. http://www.cs.virginia.edu/stream/.

System and benchmark specifications

ParFullCMS: A Geant4 benchmark, multi-threaded high energy physics workload. Employs complex geometry for simulation and essentially exhibits similar properties like Compact Muon Solenoid (CMS) experiments in CERN.

System and benchmark specifications

Parsec: A non-synthetic benchmark. Diverse instruction mix, memory access and network operations.

Application domains: Financial, computer vision, deduplication etc.

Webpage : http://parsec.cs.princeton.edu/

Results

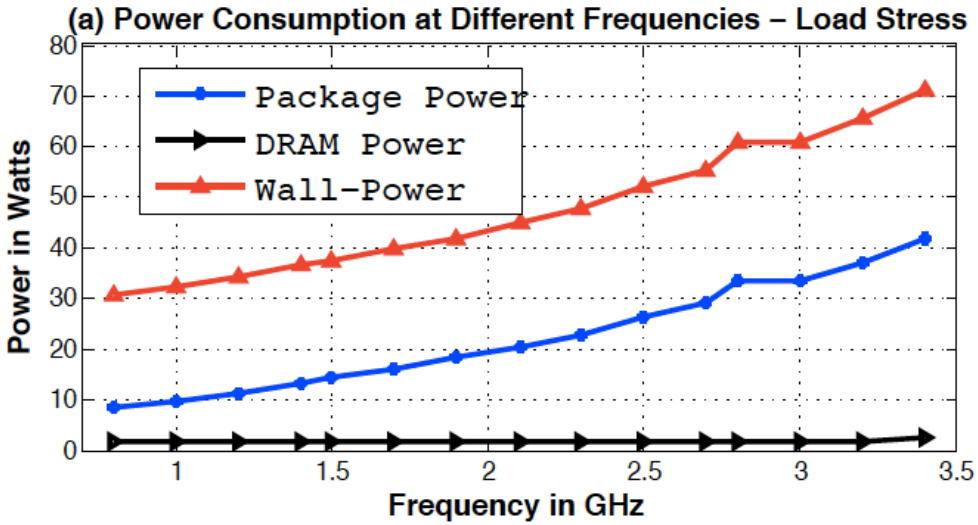

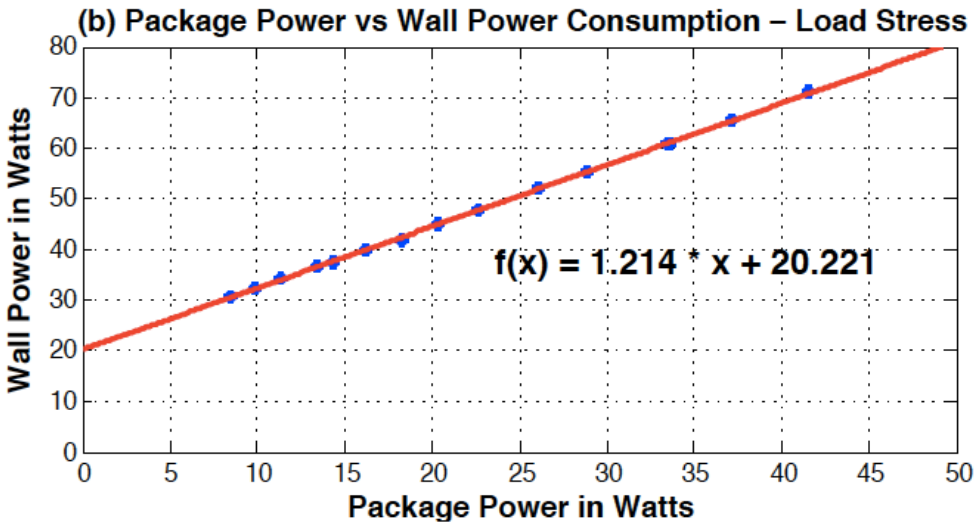

Experimental results of Machine 1 - Stress-ng

Results

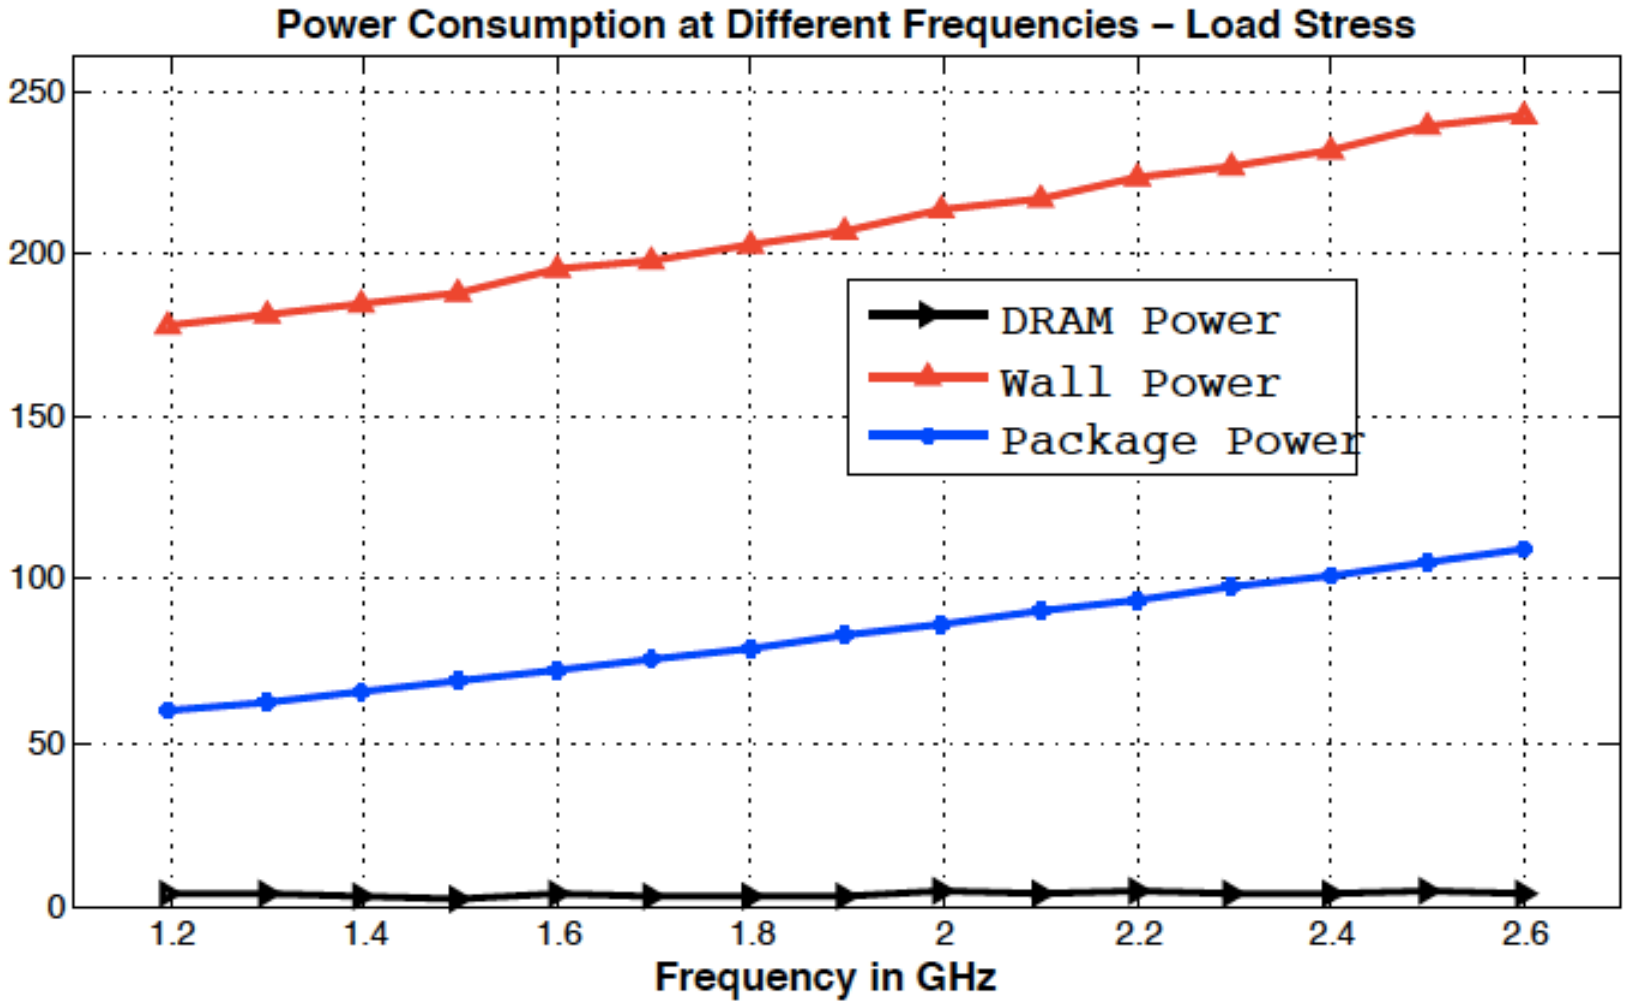

Experimental results of Machine 3 - Stress-ng

Results

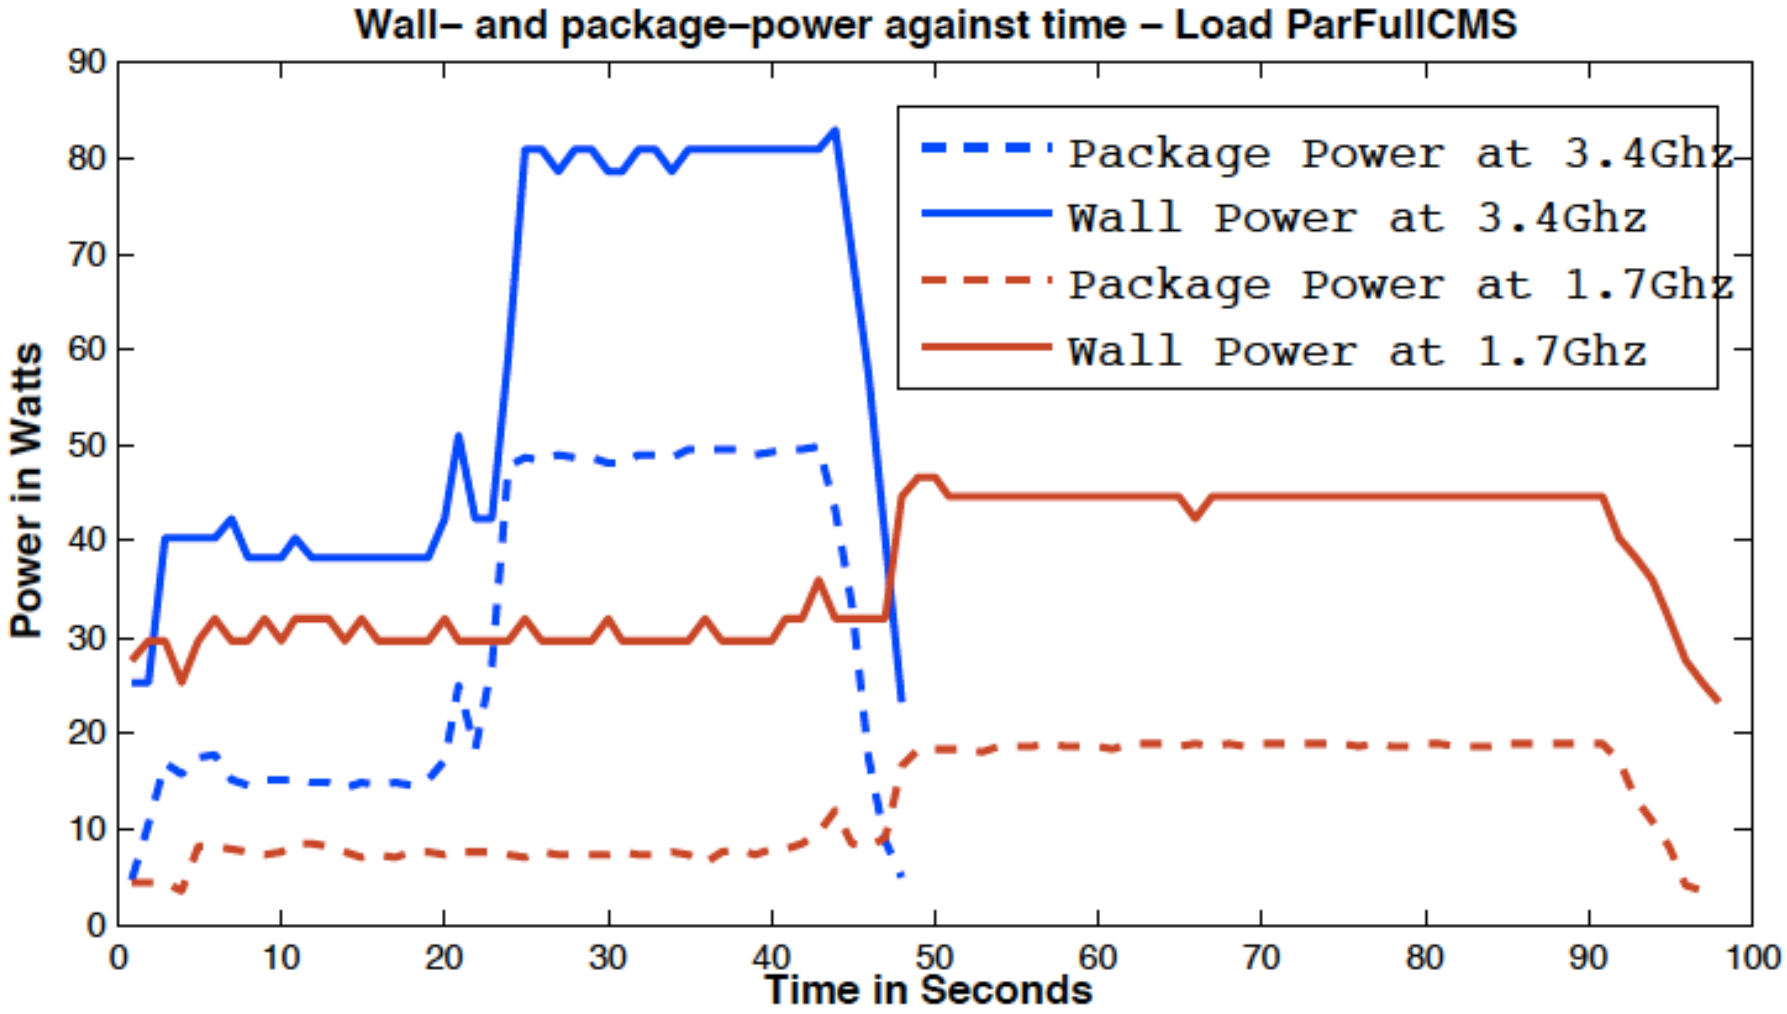

Wall and package power consumption with time - ParFullCMS

Results

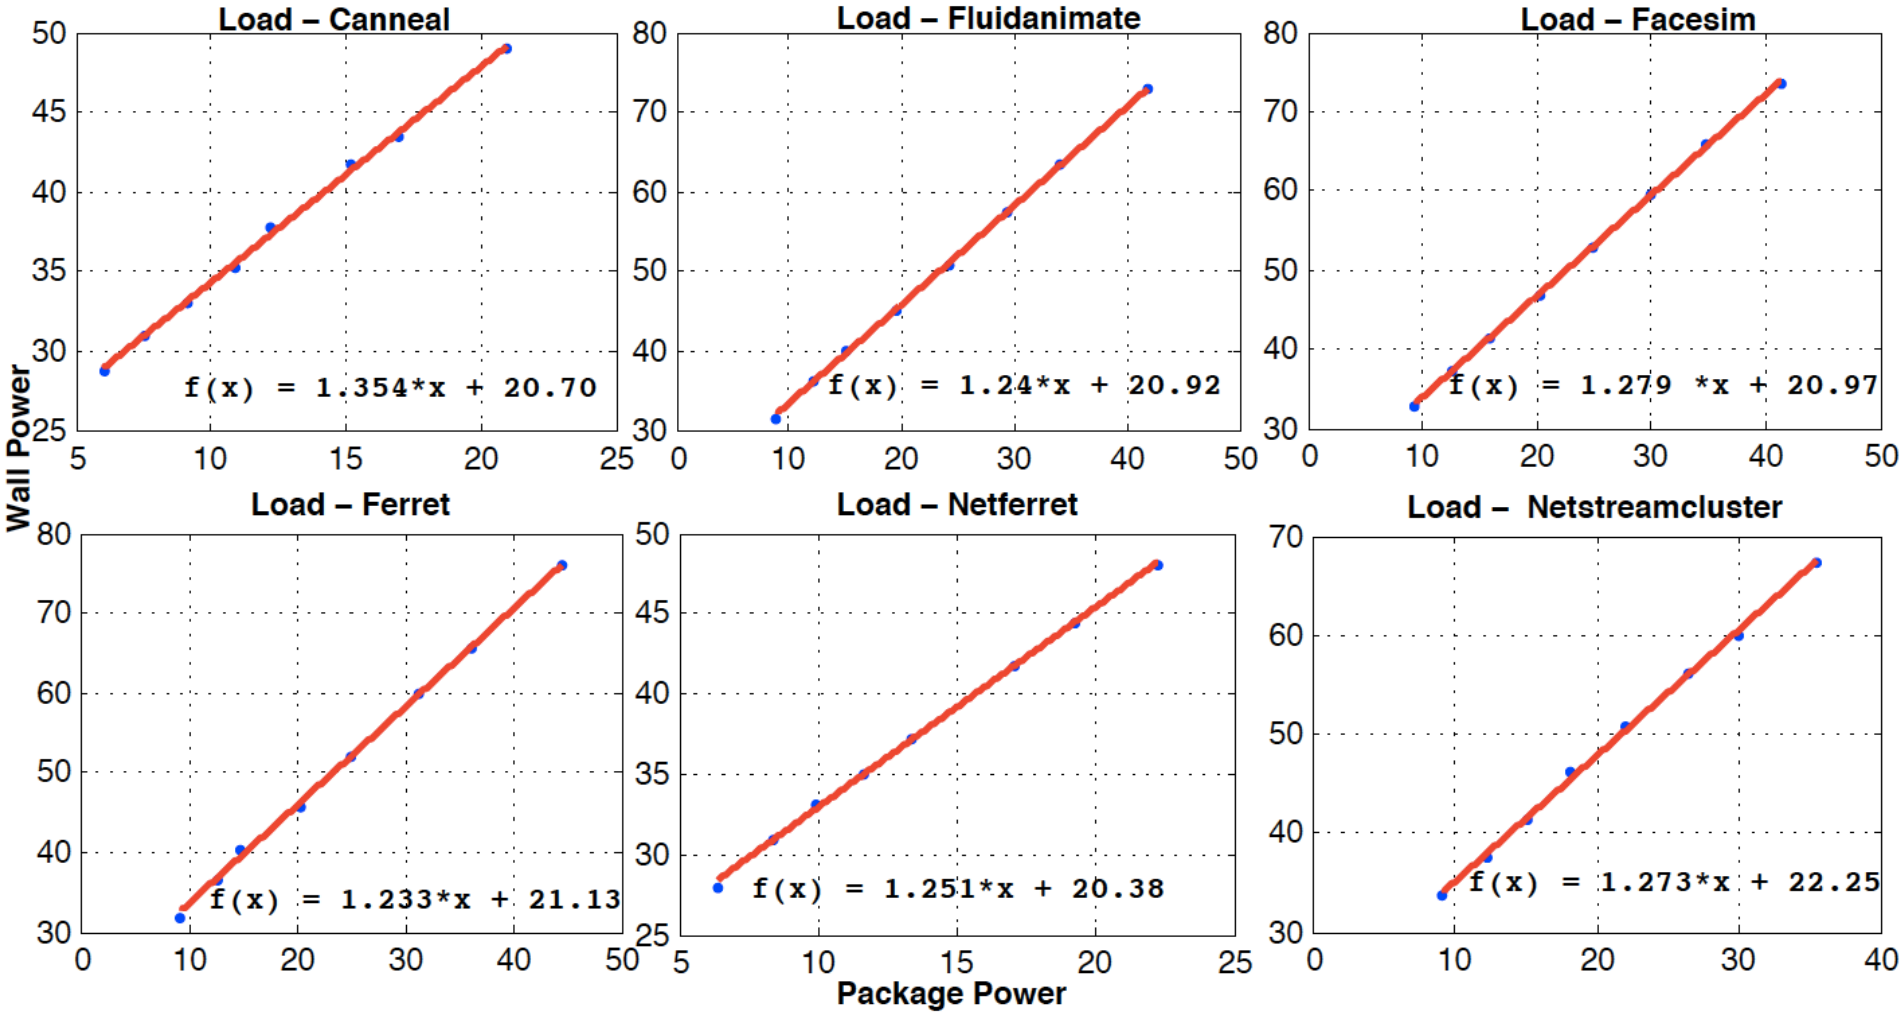

Package power vs. wall power - Parsec

Model Formulation

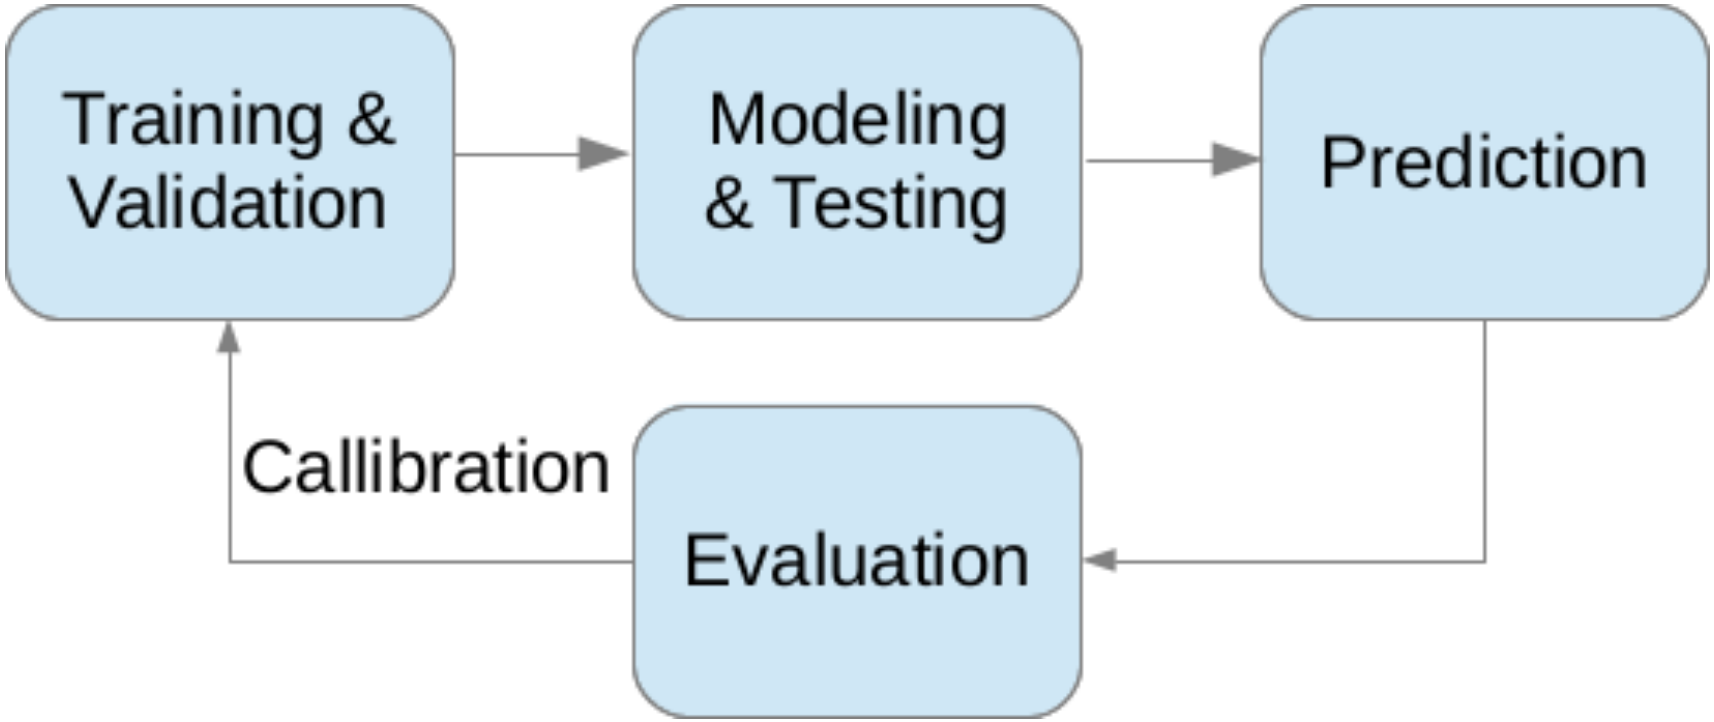

Prediction Errors of the Model

k	E_T	E_V	E_{Test}
1	4.87	6.42	5.59
2	4.38	7.89	7.08
3	4.37	8.34	7.13
4	4.37	8.17	7.18

Prediction Errors of the Model

k	E_T	E_V	E_{Test}
1	4.87	6.42	5.59
2	4.38	7.89	7.08
3	4.37	8.34	7.13
4	4.37	8.17	7.18

Prediction Errors of the Model

k	E_T	E_V	E_{Test}
1	4.87	6.42	5.59
2	4.38	7.89	7.08
3	4.37	8.34	7.13
4	4.37	8.17	7.18

$$P_{wall} = 1.227 * P_{package} + 22.084$$

Discussion and Conclusion

- There are cases when RAPL measurements are not enough to measure the wall power consumption,
  - server with multiple disks is performing a disk intensive task,
  - a server where the processing is done by the GPU rather than the CPU.
- For the disk example, the wall power consumption can be estimated using the following equation:

$$P_t = P_i + P_{RAPL} + P_{disk}$$

Discussion and Conclusion

- System that we use are relatively small scale
- Data-sets has to more diverse and rich in numbers
- We are currently enhancing our work with more data-sets and we plan to test the model on bigger scale servers
- Preliminary results show promising low error rates
- We also plan to extend our work for other processor architectures - ARM and AMD

Thank you!

Questions?