

# **ECE 571 – Advanced Microprocessor-Based Design Lecture 21**

Vince Weaver

`http://www.eece.maine.edu/~vweaver`

`vincent.weaver@maine.edu`

14 April 2016

# Announcements

- HW#10 will be another reading



# Project Stuff

- Status result/Literature search due on Tuesday
- Would like a brief update on how things are going.
- Do you need to borrow any hardware (only 1 spare wattsup-pro)
- Willing to volunteer for Tuesday rather than Thursday
- Literature search: 5 items for alone, 8 in group  
Prefer if academic, but some things those might not



exist so books, web-pages, blogs acceptable too.

If academic and can't find paper, if IEEE, ACM, etc you can log in on UMaine library website and download for free if on campus.

- Cite your references.
- You can fold this into your larger project writeup.
- It's OK if you find out someone else has done your exact project before. It's good to validate results.



# Reading

*A Validation of DRAM RAPL Power Measurements*  
by Desrochers, Paradis and Weaver



# Digression on Academic Papers



# Page 1

- Work I've been doing with some students.
- MEMSYS conference
- Haswell-EP server with 80GB RAM is 13W of power that's not even with all slots full  
428GFLOPS incidentally (2.1 GFLOPS/w)



# Page 2

- Notes on the documentation. Intel tries, but their documentation can be a real pain sometimes, often conflicting and out of date. Also their terminology can be really confusing.
- Instrumenting the hardware
  - P4 power connector
  - ATX power measurement and previous students
  - Hall effect sensors vs sense resistors





# Page 3

- DIMM extender card  
PCIe extender cards  
small resistance. Instrumentation amplifier  
Data acquisition board.
- RAPL measured using perf tool
- Synchronizing the measurements. Hard at high frequencies.  
Other ways to do it? Use serial port and data acquisition



board

On green500 list/wattsup just use NTP to make sure within a second.

- RAPL overhead, only measure at 10Hz.  
Overhead of too many interrupts, writing to disk. Also power overhead.



# Page 4

- Benchmark choice.  
idle  
stream  
BLAS: ATLAS, OpenBLAS, MLK
- GPU: OpenCL ray-tracer  
KSP



# Page 5

- Results
- Figure 2: Idle. Is a system truly idle? It is measuring the perf counters and such.
- Figure 3: Stream benchmark. Package power reads a bit low. DRAM very close when busy, low when idle.
- Figure 4: HPL Atlas Bursty. Note that when LLC miss happens, CPU power goes down (CPI gets worse) but memory power goes up.



- Figure 5: HPL OpenBLAS
- Figure 6: HPL MKL
- Figure 7: Raytracer  
DRAM behavior not well captured
- Figure 8: KSP: more CPU as at least one CPU is running the physics engine. Again DRAM power not captured well.
- Table results. Why no FLOPS/w for stream? Do ratios hold up?



# Easy Future Experiments

- Conduct same measurements on other Haswell machines (we have at least two others)
- See if the memory extender causes any slowdown/different RAPL results
- Get another memory extender and see how it works with two DIMMs
- Measure RAPL overhead, can we run at 1kHz if we read



MSR directly too a buffer w/o any other overhead? Still need a timer of some sort.



# More Difficult Future Experiments

- Measure on server system (our SNB-EP is dead, and HSW-EP is DDR4 which requires two voltage planes and DDR4 extender)

