

ECE 571 – Advanced Microprocessor-Based Design Lecture 19

Vince Weaver

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

`vincent.weaver@maine.edu`

13 April 2017

Announcements

- HW#10 will be measuring DRAM on the Haswell machine



Notes from last time

- Can you buy phase change ram?
Micron sold from 2012-2014? No one wanted.
Amorphous if you heat and quench, crystal if cook a while
- Millipede memory, tiny bumps, MEMS devices to read
- Can you buy Optane?
April 24th? Special M.2 slot on Gen7 (Kaby lake? motherboards)



For now, 16GB and 32GB modules, using like a cache of your hard disk.

- Hybrid Memory Cube, Micron, 15x as fast as DDR3. Fujitsu Sparc64 2015 has some
- High Bandwidth Memory (AMD+Hynix) AMD Fiji, NVIDIA Pascal, Intel Knight's Landing Interposer (diagram)



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'16. conference. Won an "award".
- Haswell-EP server with 80GB RAM is 13W of power that's not even with all slots full
428GFLOPS incidentally (2.1 GFLOPS/w)
130W CPU/16 cores, DRAM more than a core.



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 Intel tries, but their documentation can be a real pain sometimes, often conflicting and out of date.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.
- Three ways to read RAPL results
- RAPL measured using perf tool
- DRAM RAPL. Parametric model. Genetic algorithms. Calibrated at boot.



Page 3

- Instrumenting the hardware
 - P4 power connector
 - ATX power measurement and previous students
 - Hall effect sensors vs sense resistors
- DIMM extender card
 - Various voltages (how many? how many relevant?)
- PCIe extender cards
 - small resistance. Instrumentation amplifier
 - Data acquisition board.



- Measure with perf.
- 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 4-5-6

- Measurement accuracy concerns temp, etc.
- Does putting the DIMM in make things better/worse?
- Overhead of using perf?



Page 6-7

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



Page 8

- Results
- Do Tables tell full story?
- Figure 8 can see on i5 under-report, plus really bad on Samsung
- Same DIMMs are being used
- Phase Plots. Do they, match well?
- Haswell-EP results are better.



Easy Future Experiments

- Conduct same measurements on other machines
- 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.

