Perf Events and Validation

Vince Weaver

ICL Retreat Talk

12 August 2010

Linux Perf Events

The Good

- Available since Linux 2.6.31 (no more patching!)
- Rapidly reaching parity with other interfaces
- Can provide kernel (non-hardware events)
- The combination of a recent kernel (2.6.34+) and the upcoming PAPI patch gives good results.

Linux Perf Events

Limitations

- No support for Uncore counters
- Advanced Sampling (AMD IBS and Intel PEBS) not well supported
- Being developed by kernel devs with no HPC background
- Takes a long time for improvements to filter back to distributions

Validating Performance Counters

- Are we measuring what we think we are?
- Do the results match reality?
- Does it matter?

Current Status

- 135 : Nehalem Native Events
- 107 : Predefined PAPI Events
- 40 : Predefined PAPI Events valid on Nehalem
- 3 : Events currently validated by PAPI tests

`PAPI_FP_OPS`, `PAPI_FP_INS`, `PAPI_BR_INS`

Roofline Model

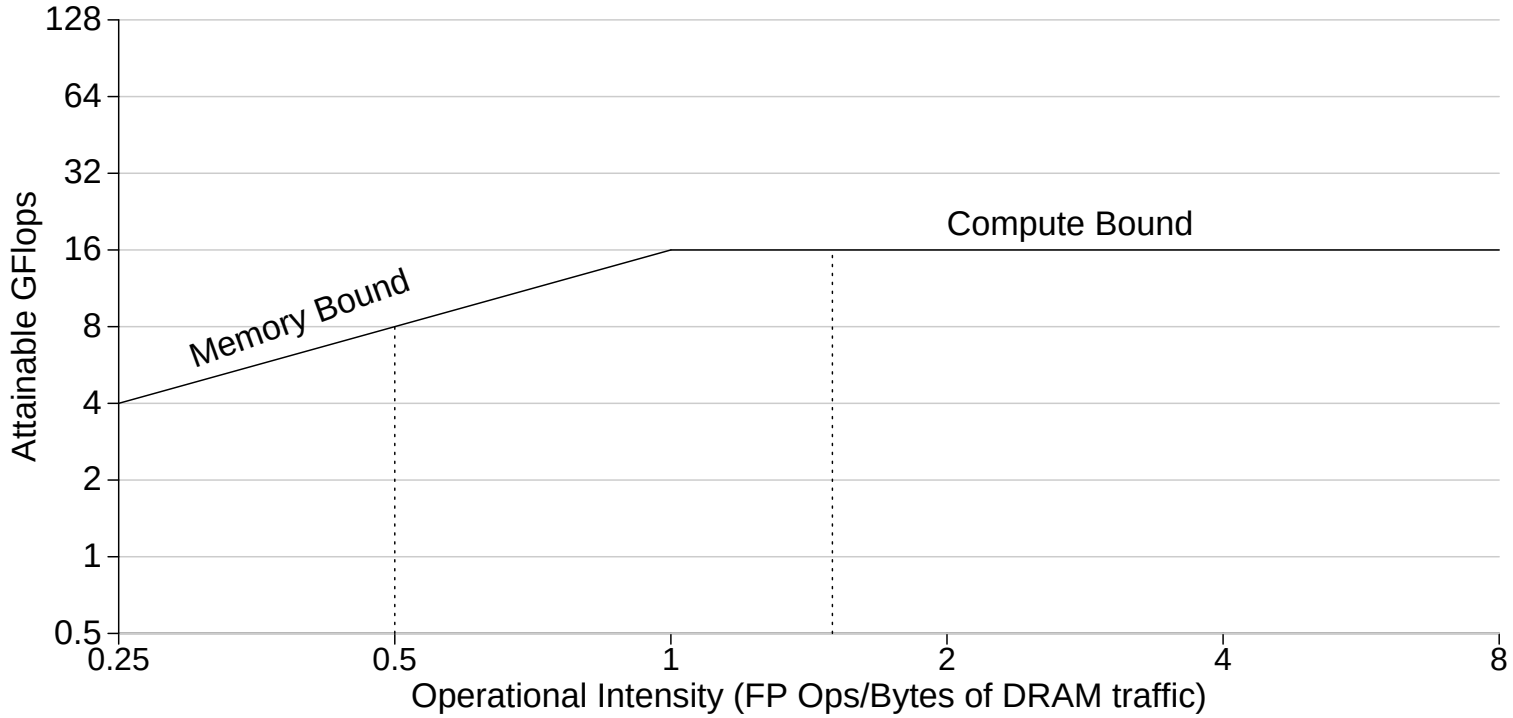

Roofline Model – Metrics

- Floating Point Operations
- Accesses to DRAM (bytes)

Deterministic Count Summary

machine	Instr	Branches	Loads	Stores
Core2	12	13	23	0
Atom	-43	-43	n/a	n/a
Nehalem	2	4	-400k	410k
Nehalem-EX	6	8	-400k	410k
Pentium D	-52	-57	-12m	-12m
Phenom	11	9	n/a	n/a
Istanbul	9	8	n/a	n/a

PAPI_L2_TCM – Total L2 Cache Misses

Strided Forward/Backward/Random

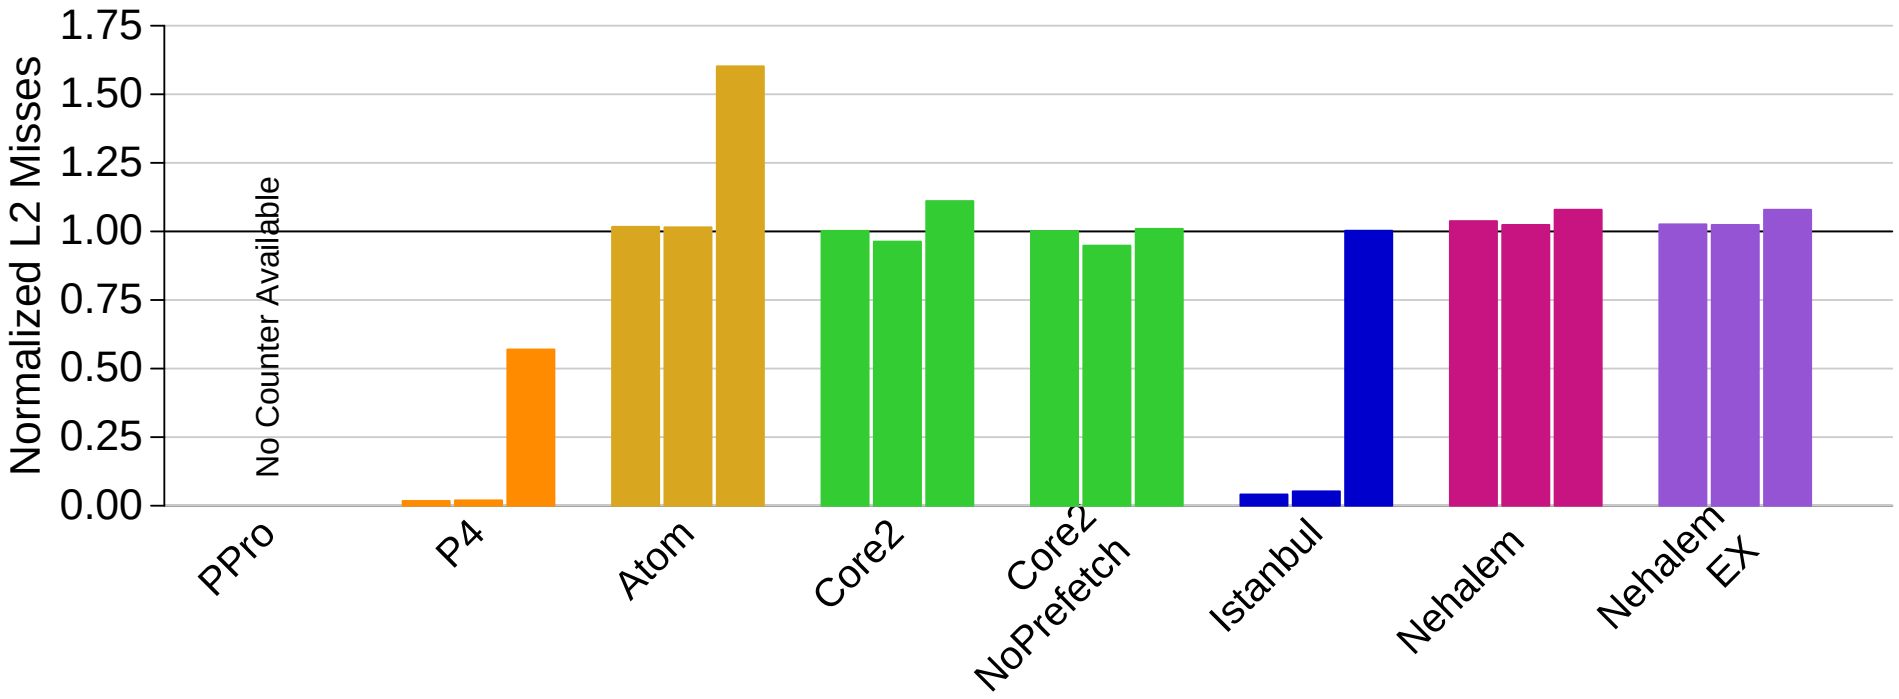

Failure Modes

- Expected
- Relatively Correct
- Wrong Counter
- Random Noise
- Always Reads Zero