

ECE 571 – Advanced Microprocessor-Based Design Lecture 5

Vince Weaver

<http://web.eece.maine.edu/~vweaver>

vincent.weaver@maine.edu

2 February 2017

Announcements

- HW#1 graded
- HW#2 due
- HW#3 will be posted



HW#1 Review

- bzip2 benchmark – what does it do?
- 19 billion instructions +/- 400 or so
(this is test input maybe?)
- 13 billion cycles +/- 6 million?
- Reversed: similar – HW2 will show you why I asked that
- Perf record: 3.5s,

66.48%	bzip2	bzip2	[.] mainSort
17.45%	bzip2	bzip2	[.] BZ2_compressBlock
6.42%	bzip2	bzip2	[.] mainGtU.part.0
6.12%	bzip2	bzip2	[.] handle_compress.isra.2
0.70%	bzip2	bzip2	[.] add_pair_to_block



- Valgrind, 1m10.189s == roughly 20 times slower?

```

11,291,448,187   ???:mainSort [/opt/ece571/401.bzip2/bzip2]
 3,381,835,437   ???:BZ2_compressBlock [/opt/ece571/401.bzip2/bzip2]
 2,138,813,059   ???:handle_compress.isra.2 [/opt/ece571/401.bzip2/bzip2]
 1,958,107,443   ???:mainGtU.part.0 [/opt/ece571/401.bzip2/bzip2]
 165,396,105    ???:BZ2_blockSort [/opt/ece571/401.bzip2/bzip2]
 140,068,091    ???:add_pair_to_block [/opt/ece571/401.bzip2/bzip2]

```

- Gprof, also 3.5s
Different results, using function entry instead of exact instruction count for sampling?

time	seconds	seconds	calls	s/call	s/call	name
71.77	2.16	2.16	53	0.04	0.04	mainSort
18.94	2.73	0.57	53	0.01	0.05	BZ2_compressB
6.98	2.94	0.21	12223	0.00	0.00	default_bzallo
1.00	2.97	0.03	1272	0.00	0.00	BZ2_hbMakeCode
0.66	2.99	0.02	1856468	0.00	0.00	add_pair_to_b



- Skid instructions – mov is more likely than sub?

perf annotate:

```
1.14 5f0:  mov    (%r10),%edx
0.56      lea    (%rdx,%r13,1),%eax
0.80      movzbl (%r15,%rax,1),%eax
3.29      sub    %r9d,%eax
```

instructions:pp

perf annotate:

```
0.78 5f0:  mov    (%r10),%edx
0.88      lea    (%rdx,%r13,1),%eax
3.14      movzbl (%r15,%rax,1),%eax
0.99      sub    %r9d,%eax
0.52      cmp    $0x0,%eax
0.58      jne    689
0.90      movslq %ebx,%rax
```



Paper Discussion

Producing Wrong Data Without Doing Anything Obviously Wrong! by Mytkowicz, Diwan, Hauswirth and Sweeney, ASPLOS'09.



Leftover Notes from Last Time



Multi-core

- More's law gives you lots of transistors. Hit limit of how fast to make a single processor, so why not just put more on the die?
- Exploits multi-programmed parallelism rather than instruction-level parallelism



Multi-threaded

- SMT (simultaneous multithreading), Intel Hyperthreading
- Hybrid of multi-core and multi-pipeline
- Your pipelines might not always be full, especially if waiting on memory
- Why not duplicate fetch/decode logic, and have two programs execute at once on same set of pipelines.
- If one is idle/stalled, run instructions from other thread



- Looks to OS as if you have two cores, but really just one with two instruction dispatch stages
- Extra logic to make sure that pipelines used fairly, the results get committed to the right register file, etc.



Power and Energy



Definitions and Units

People often say Power when they mean Energy

- Energy – Joules, kWh (3.6MJ), Therm (105.5MJ), 1 Ton TNT (4.2GJ), eV (1.6×10^{-19} J), BTU (1055 J), horsepower-hour (2.68 MJ), calorie (4.184 J)
- Power – Energy/Time – Watts (1 J/s), Horsepower (746W), Ton of Refrigeration (12,000 Btu/h)
- Volt-Amps (for A/C) – same units as Watts, but not same thing
- Charge – mAh (batteries) – need V to convert to Energy



Metrics to Optimize

- Power
- Energy
- MIPS/W, FLOPS/W (don't handle quadratic V well)
- *Energy * Delay*
- *Energy * Delay²*



Power Optimization

- Does not take into account time. Lowering power does no good if it increases runtime.



Energy Optimization

- Lowering energy can affect time too, as parts can run slower at lower voltages

Which is better?

