ECE 571 – Advanced Microprocessor-Based Design Lecture 17

Vince Weaver

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

vincent.weaver@maine.edu

6 April 2017

Announcements

- Project topics
- HW#9 will be posted

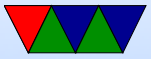

HW#7 Review

- ```
  objdump --disassemble-all ./bzip2 | grep prefetch
  (nothing)
  objdump --disassemble-all ./bzip2.swprefetch | grep prefetch
  ./bzip2.swprefetch:     file format elf64-x86-64
    401397: 0f 18 0b              prefetcht0 (%rbx)
    401e8c: 0f 18 88 a0 00 00 00  prefetcht0 0xa0(%rax)
  ... * 20 lines
  ```

- ```
  objdump --disassemble-all ./equake_l | grep prefetc
  (nothing)
  objdump --disassemble-all ./equake_l.swprefetch | grep prefetch
  ./equake_l.swprefetch:     file format elf64-x86-64
    40192a: 0f 18 0a              prefetcht0 (%rdx)
    401b8d: 0f 18 4d 48           prefetcht0 0x48(%rbp)
    401be8: 0f 18 4d 48           prefetcht0 0x48(%rbp)
  ```

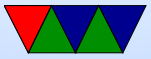

```
402036: 0f 18 09                prefetcht0 (%rcx)
40479a: 0f 18 4d 00             prefetcht0 0x0(%rbp)
4047a9: 0f 18 0b                prefetcht0 (%rbx)
```

- BZIP

		l2-cache-misses	prefetches	time
1a:	bzip2:	33%	165M	3.4s
2a:	SW prefetch:	33%	165M	3.4s
5a:	HWdisable	44%	155k	3.5s
5a:	HWdisable+Sw	44%	151k	3.4s

- Equake

		l2-cache-misses	prefetches	time
3a:	equake_l:	15%	37B	138s
4a:	equale_l swpref	16%	37B	134s
5a:	hwdisable	68%	3M	160s
5a:	hwdisable swpref	68%	3M	160s

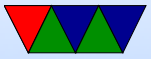

- Summary: disabling prefetch hurt, dramatically so on equake.
  Unclear what exactly the prefetch perf counter is measuring
  Enabling SW prefetch does not seem to do much, even with HW prefetch disabled.

- Why? Lots of possible reasons. compiler bug. hardware bug. hardware engineers not enable SW prefetch (is it incorrect to ignore?) other.

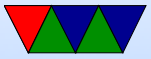

Reading of the Article

http://anandtech.com/show/9582/intel-skylake-mobile-desktop-launch-architecture-analysis/

The Intel Skylake Mobile and Desktop Launch, with Architecture Analysis by Ian Cutress

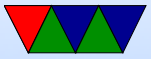

Background on where info comes from

Intel Developer Forum This one was in August 2015

Name	tech		Year
Conroe/Merom	65nm	Tock	2006
Penryn	45nm	Tick	2007
Nehalem	45nm	Tock	2008
Westmere	32nm	Tick	2010
Sandy Bridge	32nm	Tock	2011
Ivy Bridge	22nm	Tick	2012
Haswell	22nm	Tock	2013
Broadwell	14nm	Tick	2014
Skylake	14nm	Tock	2015
Kaby Lake	14nm	Tock	2016
CoffeeLake	14nm	Tock	2017?
CannonLake	10nm	Tick	2017?
IceLake			?

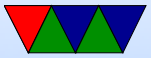

Notes

- Clock: tick-tock. Upgrade the process technology, then revamp the uarch.
- 14nm technology? Finfets?
- What technology are Pis at? 40nm?
- 14nm yields getting better.
- hard to get, even with electron beam lithography
- plasma damage to low-k
- silicon only 0.111nm
- finfet. Intel has plants Arizona, etc.

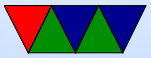

- Delay to 10nm
- 7nm? EUV?

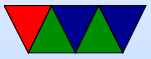

Skylake Processor – Page 1

- "6th generation" – Intel naming is awful
- Confusing naming i3, i5, i7, Xeon, Pentium, m3, m5, m7, etc.
- 4.5W ultra-mobile to 65W desktop
- They release desktop first these days. For example just today releasing "Xeon E5-2600 v4" AKA Broadwell-EP
- YUHS – number of cores, level of graphics

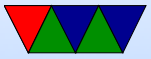

- Number of pins important. Low-power stuck with LPDDR3/DDR3L instead of DDR4 possibly due to lack of pins?
- 128M of eDRAM? For GPU?
- Intel no longer releasing info on how many transistors/transistor size? "no longer relevant for the end-user experience?"

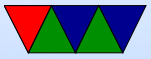

Skylake Processor – Page 2

- "mobile first" design. Easier to scale up than down?
- Want to use in tablets. What % of market does intel have on tablets?
- Up to 60% lower active power (win8.1 video playback, i7-6920HQ vs i7-4910MQ)
- Up to 60% faster estimated SPECint 2006 (i5-6300HQ vs i5-4300M)

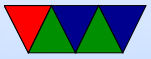

- Up to 40% better graphics On 3DMARK Sky Diver overall (m76Y75 vs m5Y71)
- Overtook AMD performance on integrated graphics (though costing more)
- Power features:
  - Fine grain power gating (core, system agent, graphics, chipset)
  - Digital PLL
  - HW and Cdyn reduction(?) (cdyn is capacitance)
  - Lower Vmin

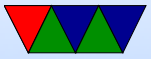

- DDR, Serial I/O power optimization
- Speed shift
- C9 display with Self Refresh Panel
- “modern standby”
- Configurable power management (i.e. the OEMs can customize, it’s off chip again)

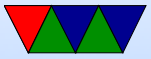

Skylake Processor – Page 3

- Microarchitecture
- Lanes increased to 128B between cores
- Sound DSP, faster DDR memory, advanced integrated graphics, wider core, enhanced LLC
- DDR3L and DDR4
- 16 PCIe lanes (what is PCI typically used for? external GPUs and disk)

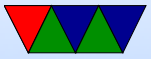

- DMI 3.0 (8.0GT/s, 3.93GB/s), but motherboard traces down to 7in from 8in?
  DMI=direct media interface, connection between north and south bridge multiple lanes, differential signalling

- Chipset on package, low power I/O OPIO (on-package I/O) 1pJ/bit

- Haswell/Broadwell had FIVR (fully integrated voltage regulator)
  simpler design, higher temperatures, made overclocking harder

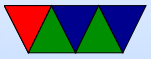

Z-height problem, too big transistors, holes in motherboard for capacitors and inductors (for tablets want as flat as possible)

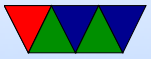

Skylake Processor – Page 4

- New socket, LGA 1151 (soldered use BGA 1515, BGA 1356 and BGA 1440), not back compatible
- Intel RST on PCIe storage?
- USB-A vs USB-C, thunderbolt
- DRAM: DDR3L vs DDR4. SO-DIMM?
- Two memory controllers per channel. Max of 4*16GB = 64GB

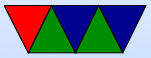

- Memory speed of 15-15-15

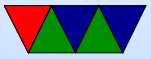

Skylake Processor – Page 5

- Performance clock-per-clock only 6% over Haswell and 3% over Broadwell
- Larger instruction window, extract more parallelism
- Can dispatch 6 uops at once (over 4uops for Haswell)
- "Higher-capacity and improved" branch predictor
- "Faster prefetch"

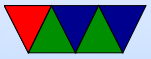

- Shorter latencies
- more units
- Power down when not in use
- Speedup AES-GCM and AES-CBC
- Better branch predictor, no details
- Improved Divider
- FMultiplier regress to Haswell numbers because it allows better performance for enterprise silicon?

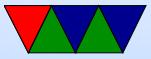

- FMUL only disclosed when asked about a rumor?
- Deeper store buffer
- Improved page miss handling (probably mem not tlb?)
- Better L2 cache miss bandwidth
- New instructions for cache management (for NVRAM?)
- Better hyperthreading
- L2 reduced from 8-way to 4-way. Saves power and area? Similar to Haswell performance with lower power

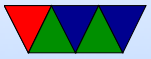

- Intel sells custom configured silicon if you are a bit enough customer
- Security Tech – SGX software guard extensions? Encrypted memory (why is that a good idea?)
- Intel MPX (memory stack/heap protection)
- Trusted execution, protected enclaves
- eDRAM – embedded DRAM – L4 cache? For graphics card? 64MB and 128MB

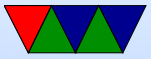

Skylake Processor – Page 6

- Graphics
- Multiple graphics sizes, can scale for power on smaller
- Lossless data compression – save power and perf by compressing before sending across bus
- 16-bit float support – less accuracy, lower power
- MPO – multiplane overlay

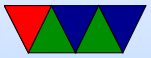

Skylake Processor – Page 7

- Speed shift
- System Agent
- PCU – power control unit, possibly even embedded intel CPU RAPL?
- Up to 4 independent power domains
- High granularity power gating, at least 12 gates

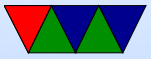

- Several frequency domains: core, uncore, 2 GPU, eDRAM

- Intel Speed Shift – let CPU control power, rather than the OS. OS can take up to 30ms to adjust, chip can do it in 1ms.

- HAVE CMOS power equation, but say dynamic plus static is approximately f**3

- Psys, knows total system power total power provided to CPU

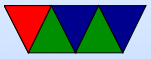

- Race to Idle
- Diminishing returns. Reintroduce idea of duty cycle (see old p4 implementation)
- Can turn off CPU cores as often as 800us

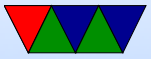

Skylake Processor – Page 8

- Kaby Lake (kah-bee lake)

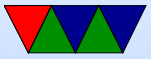

Atom Article

- Good practice using IEEE explore
- Intel Atom C2000 Processor Family: Power-Efficient Datacenter Processing
- From HOTCHIPS. Conference every year, highlights invited to write up paper for IEEE Micro

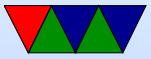

Atom Article p1

- Avoton/Rangely
- Atom is Intel's embedded line, but this one designed to be used in servers
- Seven times performance and six times power saving than previous? (they mean against the very first atom design S1200 a long time ago)
- 22nm tri-gate

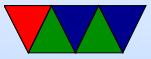

- “cold storage” segment?
- Microservers

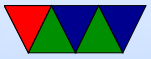

Atom Article p2-3

- Tradeoff single-thread perf for larger threads, higher energy density
- Avoton – 2,4,8 threads, 6-20W, 22nm
- Low leakage power. Takes the "Silvermont Tock"
- Each pari of Silvermont cores 1MB L2 cache
- 25.6GB/s the two channels of DDR3/3L can provide

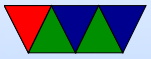

- North Complex, SSA, P-Unit

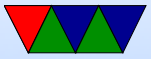

Atom Article p4-5

- South Complex
- PCH – legacy hardware like i8259, io-apic
- 4 Ethernet ports, at least 1GB. fast, also removes need for additional Ethernet chip on motherboard (saves space and power?)
- PCIe
- SATA drives, USB2

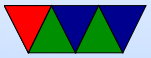

- Why no USB3? not really used on server (kb, mouse, debug)

- Rangeley – communications segment

- four workloads — application processing, control processing, packet processing, and signal processing
  Rangely can do all but the last (need higher end chip with AVX for that)

- Rangeley adds QAT crypto hardware, reliability, warranty guarantees (see recent news on failing Atom chips)

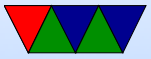

- Fast IP Forwarding

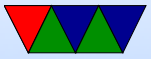

Atom Article p6

- Silvermont
- 35% IPC improvement
- now an out-of-order processor
- 2x better than previous saltwell
- Fig 5: large branch predictor (gshare?), L1 inst prefetch, more single uop, OOO, larger l2 cache, lower l2 latency,

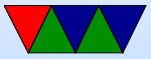

faster mul/div/fp-add, larger data TLB, OoO loads, better mem copy perf, better cache miss tolerance