# ECE 571 – Advanced Microprocessor-Based Design Lecture 16

Vince Weaver

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

vincent.weaver@maine.edu

29 March 2018

# Announcements

- HW8 will be assigned, another reading
- Project Topics were due

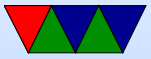

# The AMD Zen Processor

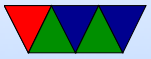

# The AMD Zen and Ryzen 7 Review: A Deep Dive on 1800X, 1700X and 1700

- By Ian Cutress
- Intel vs AMD
  - Zen - 14nm + FinFets, Intel bin there since Broadwell (2014)
- Perceived issues with Bulldozer
  - Good at integer workloads
  - High power consumption
  - Write-through L1 and shared L2 not good

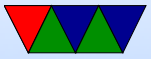

  - OSes not good at dual-thread scheduling
- New 14nm Fin-FET technology
- Goal is small cores
- Split AVX instructions in two and use saved space for L2 cache
- Using AM4 Socket
- 3.2-3.6GHz with turbo speeds of 4GHz
- Four core "core complexes" (CCX) connected via Infinity Fabric
- Writeback L1 cache
- Ryzen 3/5/7 and Threadripper

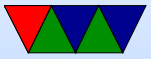

- DDR4 Memory
- Low Power Design
  - Clock gating
  - Write-back L1
  - Micro-op cache – saves re-decoding on every fetch
  - Stack engine
  - Move elimination
- Throughput
  - Enhanced pre-fetcher
  - 8MB of shared L3 – pure victim cache, not prefetched into

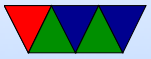

  - Large 512k L2 (intel only 256k)
  - 64kB L1 icache 4-way
  - 32kB L1 dcache 8-way – two 16-byte loads and one 16-byte store per cycle
- SMT – Simultaneous Multi-threading
- 40% IPC increase, 270% performance-per-watt increase
- Will scale to laptops too
- Goal: core, cache and power
- Core improvements
  - Two threads/core
  - Branch predictor improved

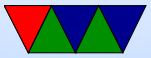

  - 6-wide uop dispatch (vs 4)
  - Larger queues
- Fused ops
- 72 out-of-order loads
- Four FP pipes that can handle 128-bit Fused Multiply Add, can be combined for 256-bit AVX, more than than (512) must be split
- "decoupled" branch predictor, can predict down a path
- TLB in the branch predictor (L0/L1/L2)
- Decode 4 instructions at once, can be fused
- Stack engine

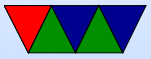

- Dispatch 6 per cycle, each has 14 entry queue
  4 ALU, 2 AGU (address generation)
  2 branch pipelines, one CRC, one IMUL
  2 branches per cycle, 8-wide retire
- Caches
  64 entry L1 TLB, 1.5k entry L2 TLB
- Floating Point
  2 level scheduling? 8-wide retire? Two AES units?
- SMT, one thread can be tagged as more important
- New instructions
  - ADX – extended add

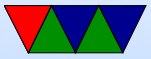

  - RDSEED – for random number
  - SMAP – supervisor access mode prevention
  - SHA1/SHA256 – hash
  - CLFLUSHOPT – flush
  - XSAVEC – compact XSAVE
  - CLZERO – zero cache line (AMD only)
  - PTE Coalescing (AMD only) – merge 4k pages into 32k page
- "Pure Power" – monitor power/temp/speed and adjust in real time
- "Precision Boost" – can jump frequency in 25MHz

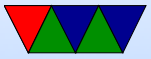

increments
1000 sensors per chip

- "SenseMi" Extended Frequency Range – automated overclock
- Smart neural-network branch predictor?
- Smart neural-network prefetcher?
- SoC, PCIe lanes
- AM4 socket – 1331 pins
- DDR4-2666 UDIMMs?
  - Registered vs Unregistered
  - Registered has a buffer on board. More expensive but

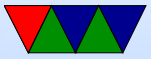

can have more DIMMs on a channel
  - Registered may be slower (if it buffers for a cycle)
  - RDIMM/UDIMM
- Did it win at the benchmarks?

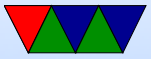

# Zen: An Energy Efficient High-Performance x86 Core

- By Singh, Schaefer, Rangarajan, Josh, Henrion, Schrieber, Rodriguez, Kosonocky, Naffziger, and Novak
- Zen replaces Excavator (high-end) and Jaguar (low-power)
- 52% increase in IPC, reducing $C_{AC}$ by 15% switching capacitance
- 12-layer telescoping metal stack, 56x top metal layer, metal-insulator-metal capacitors (MIMCap)

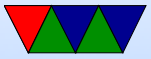

- Three $V_t$ options
- CCX L3 is a cross-bar
- RVVD (real voltage) distributed to package. high efficiency volt regulator at 6.125mV increments LDO (low-drop off regulators) adjust with 2-3mV per step.
- Separate VDDM
- Cache: separate macros for tag, data, state. Parity is checked.
- Protected by DECTED (Double Error Correct/Triple error detect), not much more overhead than SECDED
- L3 cache is 16 mm2

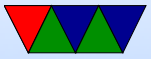

- Level-shifting FIFO to handle the different voltage or clock frequencies between processors and L3
- Shadow tags.
- Lots of level converters
- L1 arrays use 8T macro
- Write-back more energy efficient
- WWL boost. Boosts on write if 1, not if 0
- Recombinant clock mesh, four level.
- Clock gating. Clock power saves 30% over jaguar and 25% over excavator
- Bulldozer, Piledriver, Steamroller, Excavator

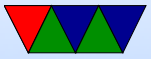

- Fine-grained digital frequency synthesizer
- MIM decoupling capacitors (metal-insulator-metal)
- Silicon schmoo plot (what's a schmoo plot?)