ECE 571 – Advanced Microprocessor-Based Design Lecture 31

Vince Weaver

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

vincent.weaver@maine.edu

25 November 2024

Announcements

- Project status update due

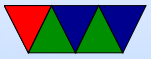

First, Zen4

- This wasn't a reading this year, but the Zen5 reading didn't have a lot going on so just a refresher on Zen4 for comparison

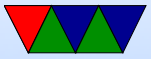

AMD Zen4 Ryzen 9 7950x and Ryzen 5 7600X Review: Retaking The High-End

Article by Ryan Smith and Gavin Bonshor

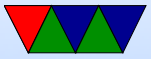

Background

- Zen4 announced end of August 2022
- "Ryzen 9" and "Ryzen 5": What's this 9/7/5/3 stuff about?
- Improved IPC by 13%. AVX-512 and larger caches?
- TDP – 170W is higher than previous chips
- Ryzen 9 7950X – 16 cores, 32 threads, 170W TDP 4.5GHz, boost to 5.7GHz
- Ryzen 9 7900X 12 cores
- Ryzen 7 7700X 8 cores (150W TDP)

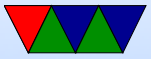

- Ryzen 5 7600X 6 cores (105W TDP)
- Big boost in power efficiency
  TMSC 5nm node, can boost freq w/o power
- Tjmax 95 degrees C (throttle threshold)
- Precision Boost Overdrive?

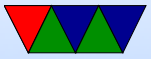

Chiplets

- Has an IOD (I/O Die) in 6nm (much smaller than prev)
- One or two core-complex dies (CCD)
- IOD has PCIe 5.0, DDR5 integrated memory controller
- Integrated GPU based on Radeon RDNA2

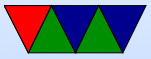

Socket AM5

- AMD doesn't update socket often
- (AM4 was from 2016/2017 and for Zen/Zen2/Zen3 (1331 pins))
- Supports DDR5 and only DDR5 (intel's new chips can also do ddr4)
- PCIe 5.0
- 1718 pins (intel high-end is 1700)
- LGA – Land Grid Array, "land" of pins on bottom, actual pins are on motherboard socket (or can be soldered)

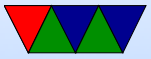

pins don't have to be square or round
higher pin density
on PGA had to fix bent pins with mechanical pencil?

- Backwards compat with AM4 coolers
- Higher power delivery, AM4 only up to CPU TDP 105W (socket 142W)
  170W (socket 230W)
- SVI3 – scalable voltage interface – high speed 2-way communication between CPU and voltage regulators, CPU and VR can work together instead of making conservative assumptions

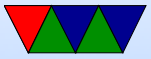

3 variable power rails, cores+graphics, SoC, misc VDD for infinity fabric

- BIOS flashing. Forward compatible motherboard, but how can you flash to new BIOS if doesn't support your CPU? Loaner CPUs?
  new: allow flashing BIOS w/o a CPU, USB BIOS flash-back
- IO
  - AM4 supported 24 PCIe lanes and 4 Superspeed USB
  - AM5 supports 28 PCIe lanes and in theory PCIe5 which has 4GB/sec/lane

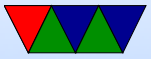

- PCIe5 intended for NVMe but some vendors using it for an x4 slot?
- 3 superspeed type-C ports, a superspeed type-A, and one USB 2
- No native SATA
- Need motherboard support to get PCIe 5
  double speed, so tight tolerances
  also short “throw”
- Note recent NVIDIA GPUs not support PCIe 5 yet
- Support up to 4 displays, HDMI + 3 Type-C USB Display Port

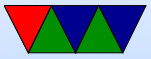

- Hybrid graphics (?)
- MIPI's I3C vs i2c for chip communication
  "improved i2c"
  vaguely backwards compat but faster
  have to pay license
- MIPI – mobile industry processor interface
  technical specs for mobile systems
- ESPI? Replacement for the LPC (low-pin count) bus
  used for legacy PC compatibility
- Digital MIC / MIPI Soundwire

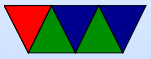

Motherboard

- New chipsets
- High-end one supports PCIe 5.0 to PEG slot?
  PCI Express Graphics, 16x slot, only graphics cards need 16x?
- PCIe 5.0 NVME storage
- SuperSpeed USB? Two 10GBps or 1 20GBps port
- Motherboards must support one 4x PCIe 5.0
- “Extreme” motherboards PCIe must also go to one PCIe slot, usually x16 PEG slot

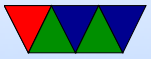

- Daisy-chained chipsets?
- Chipsets talk over PCIe?
- Most stuff at PCIe4 speeds?

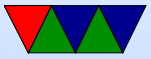

DRAM

- DDR5, can do DDR5-5200
- DDR5-5200 if one-dimm-per-channel (DPC), DDDR5-3600 if two dimms-per-channel (DPC)
- In system with 4 DIMM slots (filling all 4 will actually reduce performance if have to drop to 2-DPC mode)
- Costs more, but trying to be future-proof
- CPU supports ECC memory but motherboard has to support it
- Only unbuffered DIMMs (no RDIMM/LDIMM)

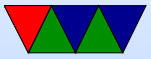

- Overclocking – issues keeping in sync when CPU overclocks
  EXPO – extended profiles for overclocking (similar to Intel XMP but Intel keeps that a secret)

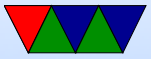

IOD

- Chiplets
- Only CPU on one chiplet, I/O etc on other
- Big improvement, down to 6nm from 14/12nm
- TSMC vs GloFo (Global Foundaries was AMD's old fab that got spun off)
- Jump to 3.4Billion transistors
- ARM security processor, Microsoft Pluton

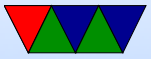

Integrated Graphics

- Low-end integrated graphics included
- Single RDNA 2 compute unit
- Is full RDNA including raytracing
- Compete with Intel that has had integrated GPU with desktop/laptop chips for ages
- AV1, HEVC, H.264 video decoding and HEVC and H.264 video encoding, and this can be used even if a discrete GPU in system
- Drive 4 4k@60Hz displays. HDMI 2.1 up to 48Gbps or

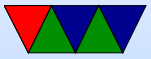

DisplayPort 2.0 up to UHBR10 (?)

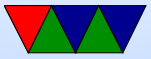

Infinity Fabric

- Die-to-die interconnect
- Improve power efficiency
- They found running narrower link faster had better power, mostly from running lower voltage
- Low power modes
- Cache Coherent Master,

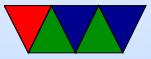

CPU Cores

- Zen4 not a complete re-design (they are planning that for Zen5)
- TSMC 5nm top of the class, Intel and Samsung can't really match yet
- 13% IPC improvement
  - L2 Cache
  - Execution engine
  - Branch prediction
  - Load/Store

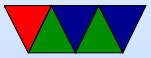

  - Front End
- 5.7GHz clock speeds. This was the old way of increasing performance that had stalled for a really long time
- More experience with TSMC. Had been hoping to break 6GHz
- AVX-512
  per-lane-masking
  neural-network: bfloat16 (16-bit floating point format, truncated 32-bit IEEE with 8-bit exponent rather than 5-bit exponent IEEE 16-bit), VNNI (vector neural network instructions) for deep learning

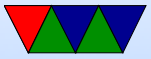

complex history across Intel chips. Knights Landing. Ice Lake and Rocket Lake added it, but Alder Lake dropped it because the little chips in their big/little can't handle
AMD doesn't implement full wide but instead on top of AVX2 circuitry (would take a lot of die area for full implementation)
Still a win as AVX-512 code is denser and can take fewer instructions than equivalent AVX2
Big problem with other AVX-512 implementations is lighting up a full 512-bit datapath at once has large power draws and can hit thermal limits and need to be

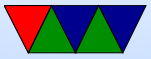

throttled

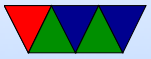

Microarchitecture

- Branch Predictor
  - Predicts two branches per cycle (same zen3)
  - L1 BTB increased to 2x1.5k entries
  - L2 BTB to 2x7k entries
  - op cache? now 6.75k ops?

    caches decoded uops so you don't have to re-decode?
- Pipelines, still schedule 10 integer and 6 FP per clock

  3 cycles for FADD/FMUL and 4 for FMA
- Retire/re-order queue now 320 instructions

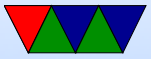

- integer and FP register files now 224 and 192
- load queue now 88 loads, 3 loads / 2 stores per cycle
- L2 cache– Doubling size per core from 512k to 1M
- L3 cache is 32MB for an 8 core CCX, acts as victim cache

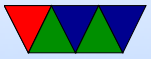

Security

- IBRS (indirect branch restricted speculation) for Spectre, it automatically does this when you enter kernel mode so you don't have to manually do it in OS

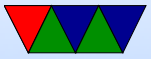

Notes on benchmarks

- L3 cache on core 15-19s, but inter-core 80ns?
- High end chip beats Intel chip

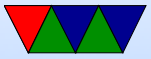

AVX-512

- 32-register ZMM0 - ZMM31, each 512 bits. Lower YMM0 and XMM0
- Long and complex, whole class on this
- Introduces new software prefetch instructions...

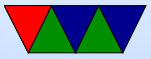

Zen3 vs Zen4 differences

- From `https://en.wikichip.org/wiki/amd/microarchi`
  `zen_4`
- L1 and L2 DTLB increased from 64 to 72, 2048 to 3072 entries
  (L1 fully assoc, L2 12-way, PDE to speed table walks)
- Max physical/linear address size from 48 to 52/57
- In theory up to 92 cores
- AVX-512
  aside, Galois Field New Instructions

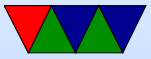

VGF2P8AFFINEINVQB

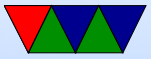

EPYC Server

- Genoa
  - 96 Cores DDR5
  - 6TB RAM, 128 lanes, PCIe 5.0 64 lanes
  - CXL1.1 remote RAM disaggregated Type 3 Memory Pooling
- Everyone thought it would come out as same time as Intel Granite Rapids, but intel has pushed that off to 2024 (from 2024, finally released September 2024, 128 p-cores, 500W)

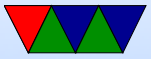

AMD Zen5

- *The AMD Zen5 Microarchitecture: Powering Ryzen AI 300 Series For Mobile and Ryzen 9000 for Desktop* by Gavin Bonshor
- 50 TOPS (Trillions of Ops (integer and other) per second)
- Zen5 16% better IPC than Zen4
- XDNA2 NPU designed for MS Copilot(!)
- Dual-pipe fetch, advanced branch prediction
- 8-wide displatch/retire

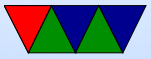

- 6 ALUs, 3 multipliers
- 48k 12-way l1 data cache, 4-cycle load
- Better data prefetcher
- 512-byte AI datapath, AVX-512 (fully 512)
- 2-cycle latency FADD (Zen4 double-pumped)
- Benchmarks
- Full cores TSMC 4nm, Zen 5c (power eff) 3nm
- EPYC/ Turin with 192 Zen5 cores