CS433
AMD ZEN 2

Hyoungwook Nam (hn5)
Anjana Suresh Kumar (anjanas3)
Vibhor Dodeja (vdodeja2)
Namrata Mantri (nmantri2)

Table of Contents

1. Overview
2. Pipeline Structure
3. Memory Hierarchy
4. Security and Power
5. Takeaways

Overview

History of AMD's x86 microarchitectures (1)

K5 - K7 (95~02)

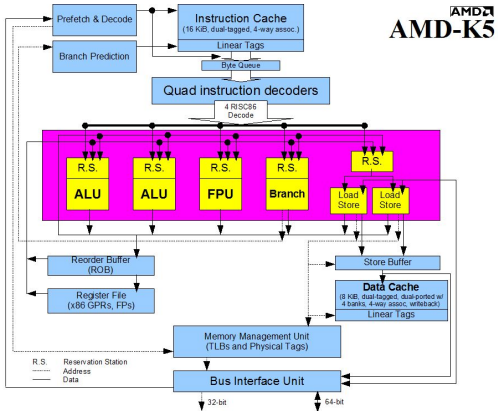

x86 frontend, **RISC backend**
Superscalar, OoO, speculation
SIMD, L2 cache (K6)

K8 (03~08)

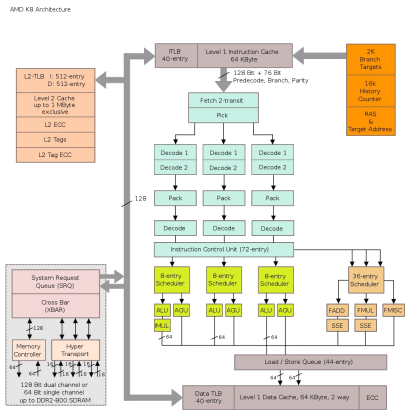

Introduced **x86-64** ISA
Dual-core (Athlon 64 X2)
Integrated memory controller

K10 (09~11)

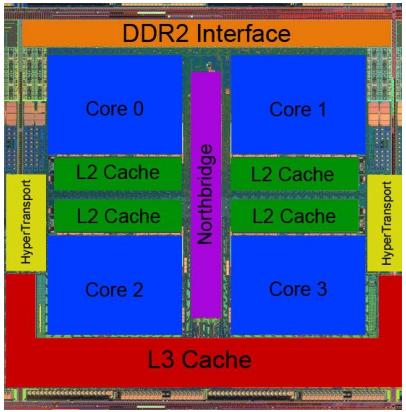

Up to 6 cores
Shared L3 cache
GPU integrated APUs (Fusion)

https://www.tomshardware.com/picturestory/713-amd-cpu-history.html

History of AMD's x86 microarchitectures (2)

Bulldozer (11~16)

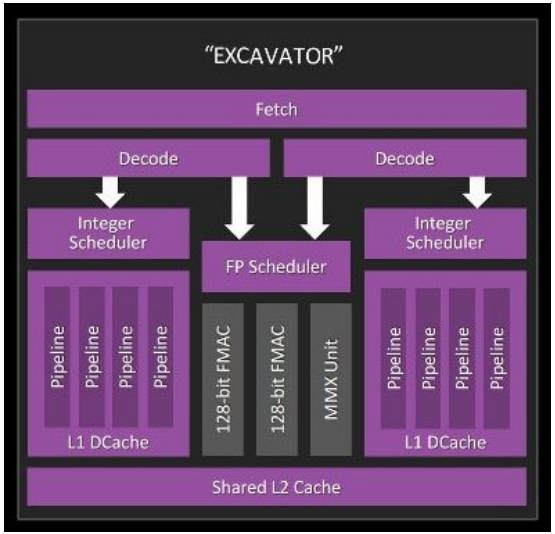

Multi-core module (MCM)
Two cores per **module**
Shared FP and L2 in a module

Zen (17 ~)

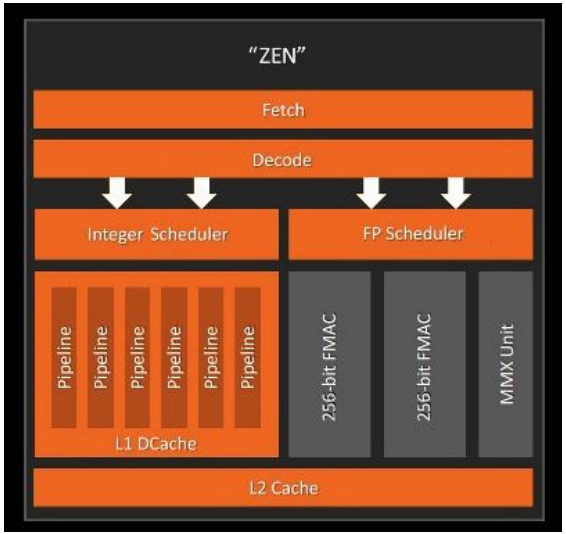

Simultaneous Multi-thread (SMT)
Two threads per **core**
Higher single-thread performance

AMD Financial Analyst Day, May 2015

Multi-core Module (MCM) Structure of Zen

Single EPYC Package

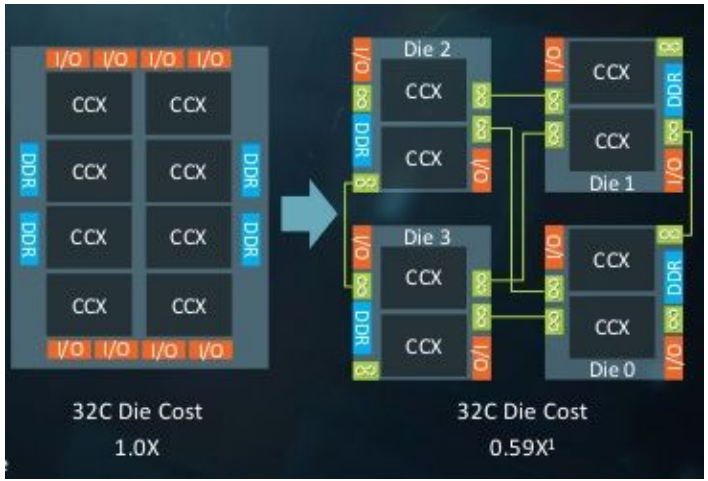

Single Die (Chiplet)

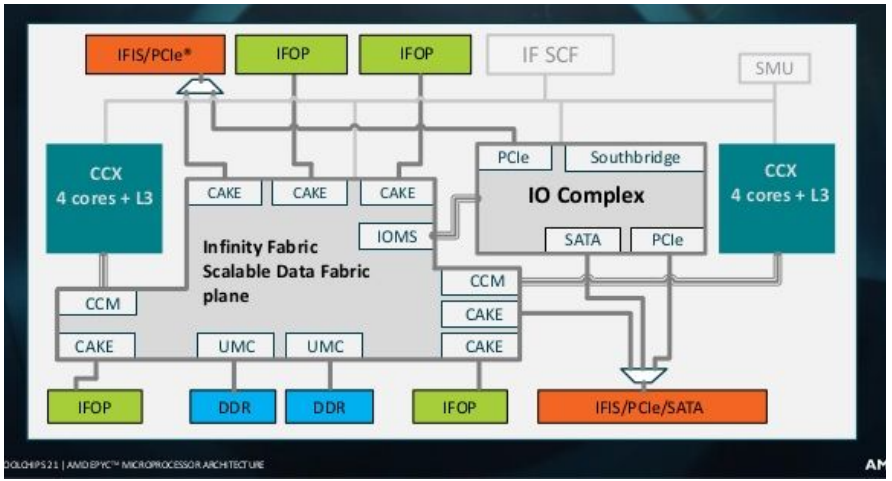

Multiple **dies** in a **package,** 2 **core complexes** (ccx) per die, and up to **4 cores** per ccx. (~4c8t per die)

Fully connected NUMA between dies with **infinity fabric** (IF) which also interconnects ccx.

https://www.slideshare.net/AMD/amd-epyc-microprocessor-architecture

Zen 2 Changes Over Zen 1 and Zen+

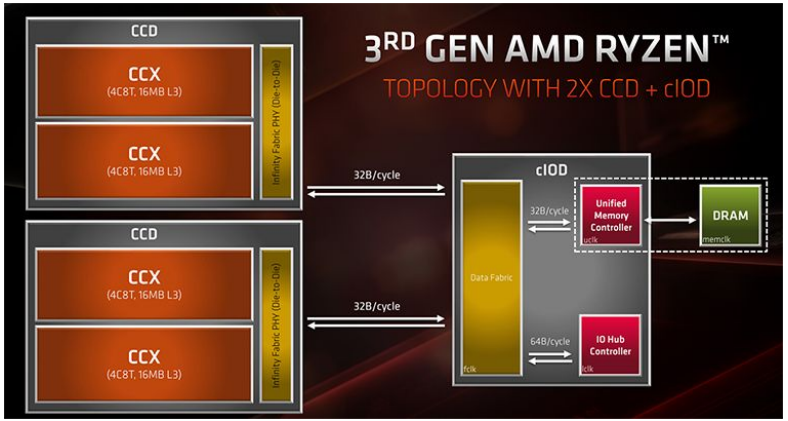

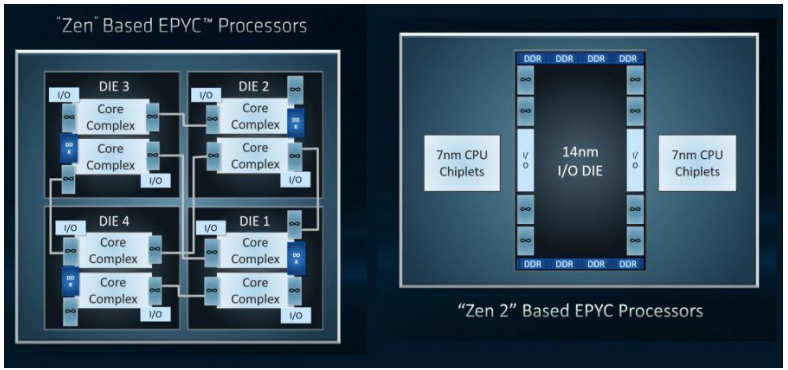

Dedicated IO chiplet using hybrid process

- TSMC 7nm CPU cores + GF 14nm IO chiplet

2x more cores per package

- up to 16 for consumer and 64 for server

More ILP

- Better predictor, wider execution, deeper window, etc.

2x Larger L3 and faster IF2

Extra security features against spectre attacks

https://www.pcgamesn.com/amd/amd-zen-2-release-date-specs-performance

https://hexus.net/tech/news/cpu/131549-the-architecture-behind-amds-zen-2-ryzen-3000-cpus/

Pipeline Structure

Fetch

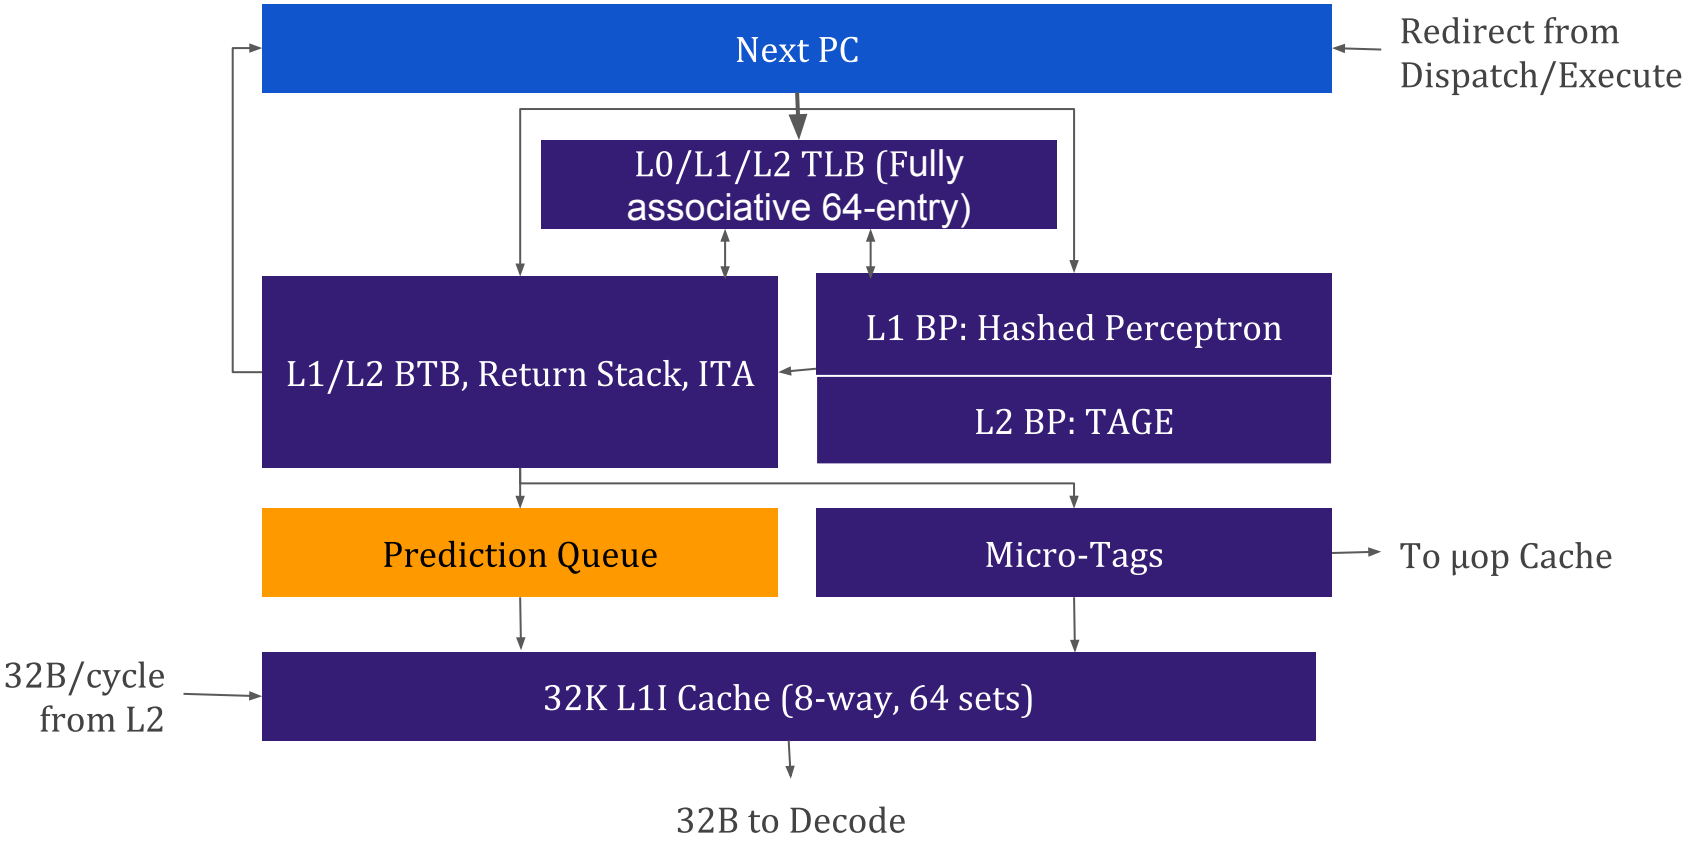

https://fuse.wikichip.org/news/2458/a-look-at-the-amd-zen-2-core/

L1 Hashed Perceptron

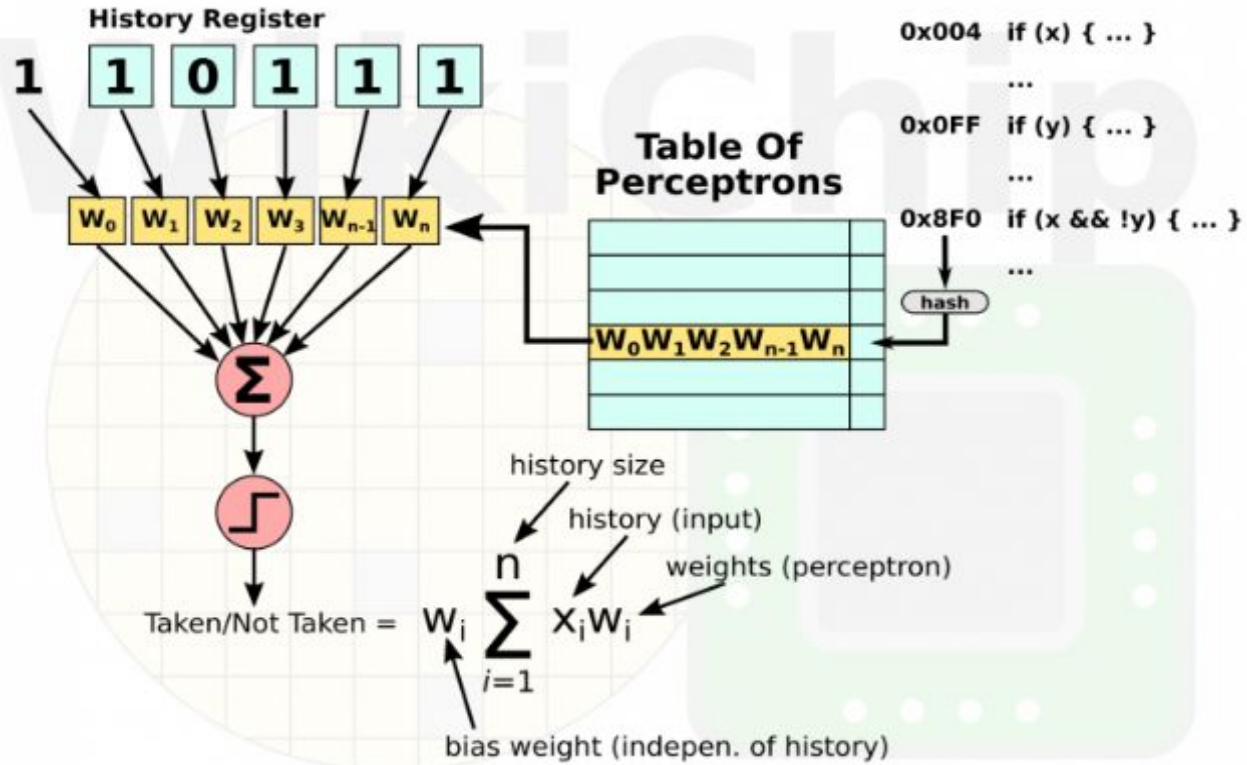

https://fuse.wikichip.org/news/2458/a-look-at-the-amd-zen-2-core/

Back

L2 TAGE

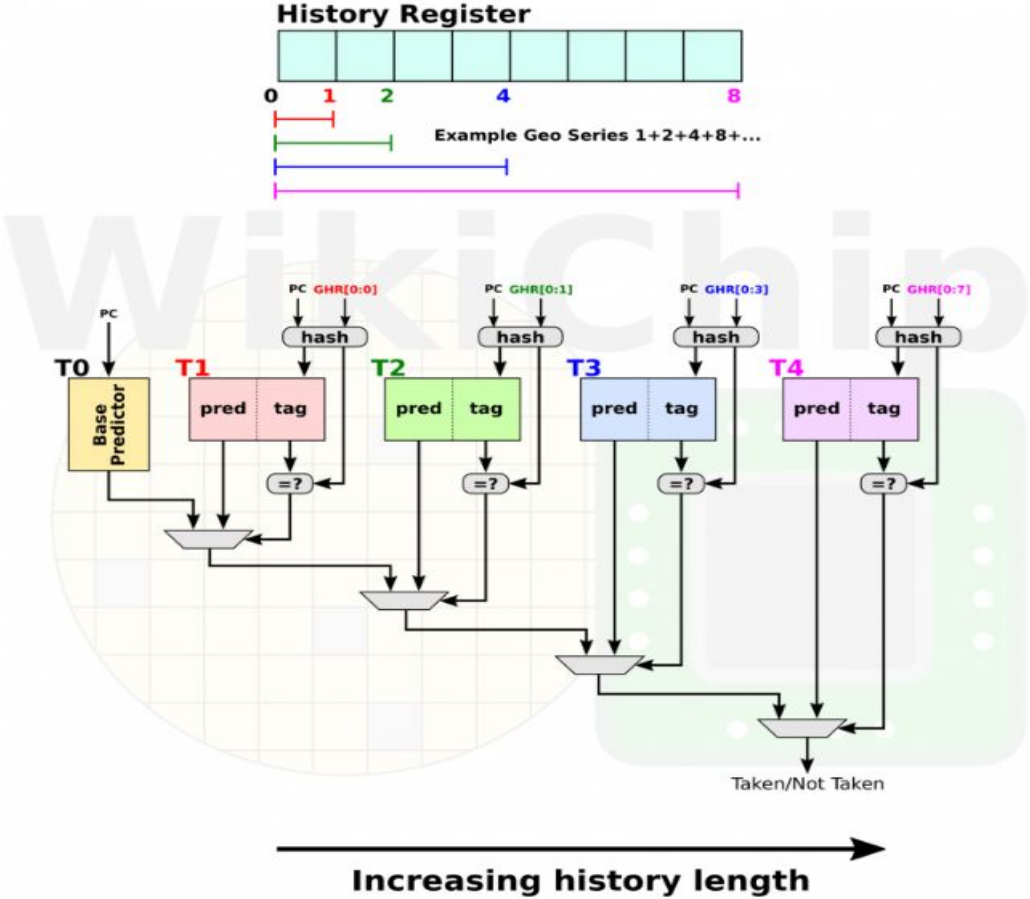

Back

https://fuse.wikichip.org/news/2458/a-look-at-the-amd-zen-2-core/

Decode

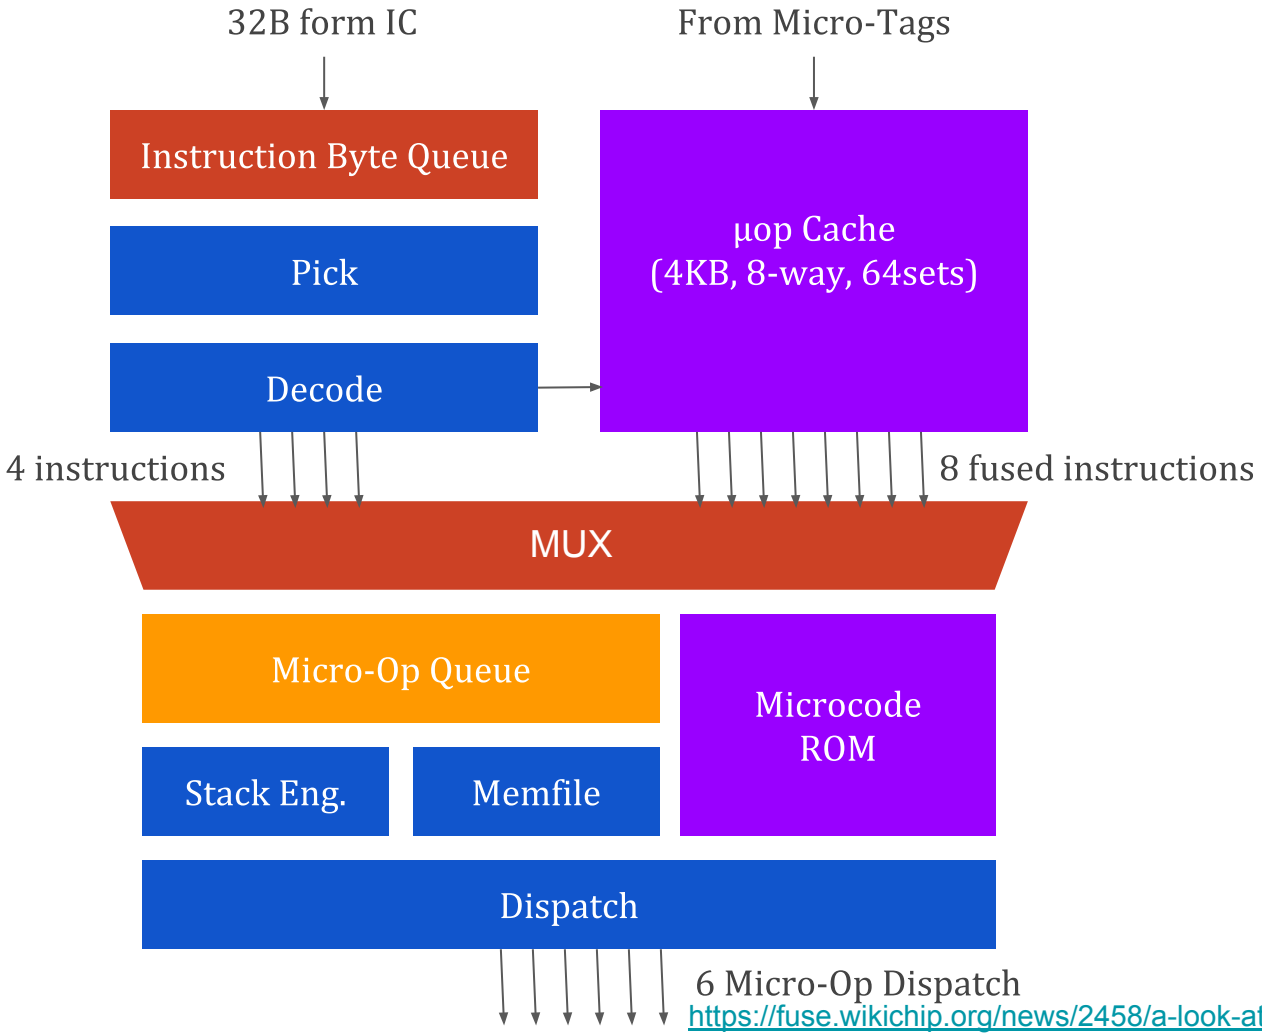

https://fuse.wikichip.org/news/2458/a-look-at-the-amd-zen-2-core/

Modes of Operation

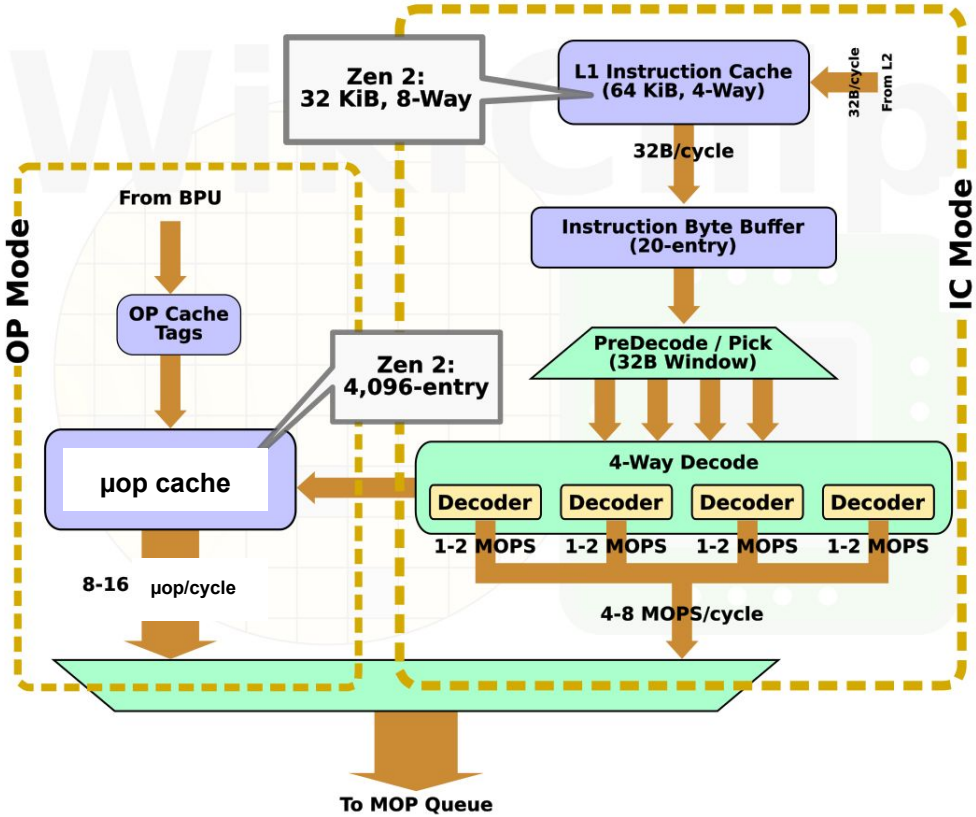

Back

https://fuse.wikichip.org/news/2458/a-look-at-the-amd-zen-2-core/

Integer Execute

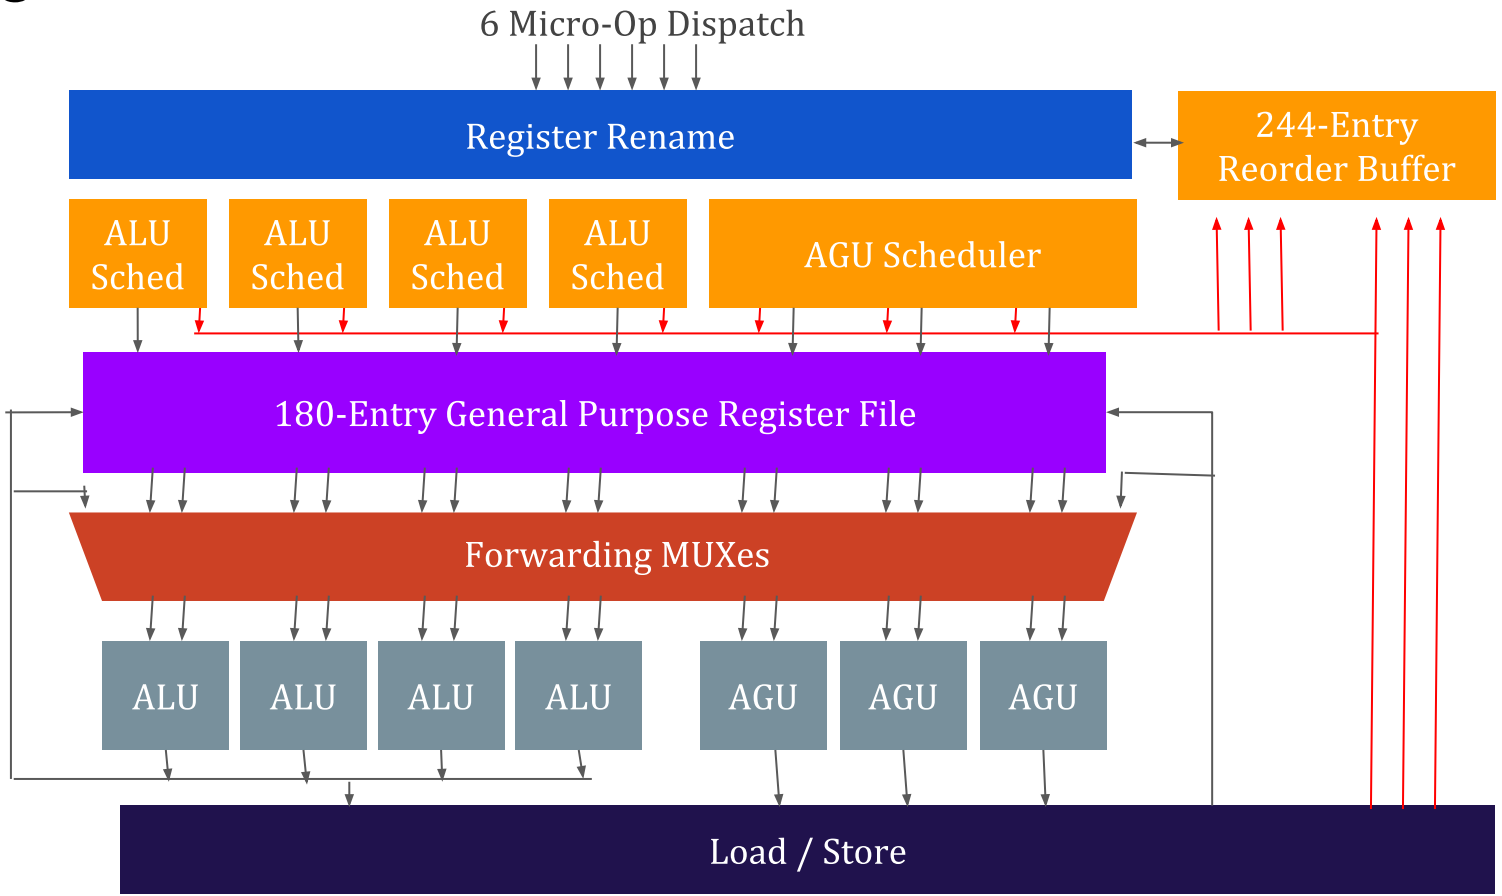

https://fuse.wikichip.org/news/2458/a-look-at-the-amd-zen-2-core/

FP/Vector Execute

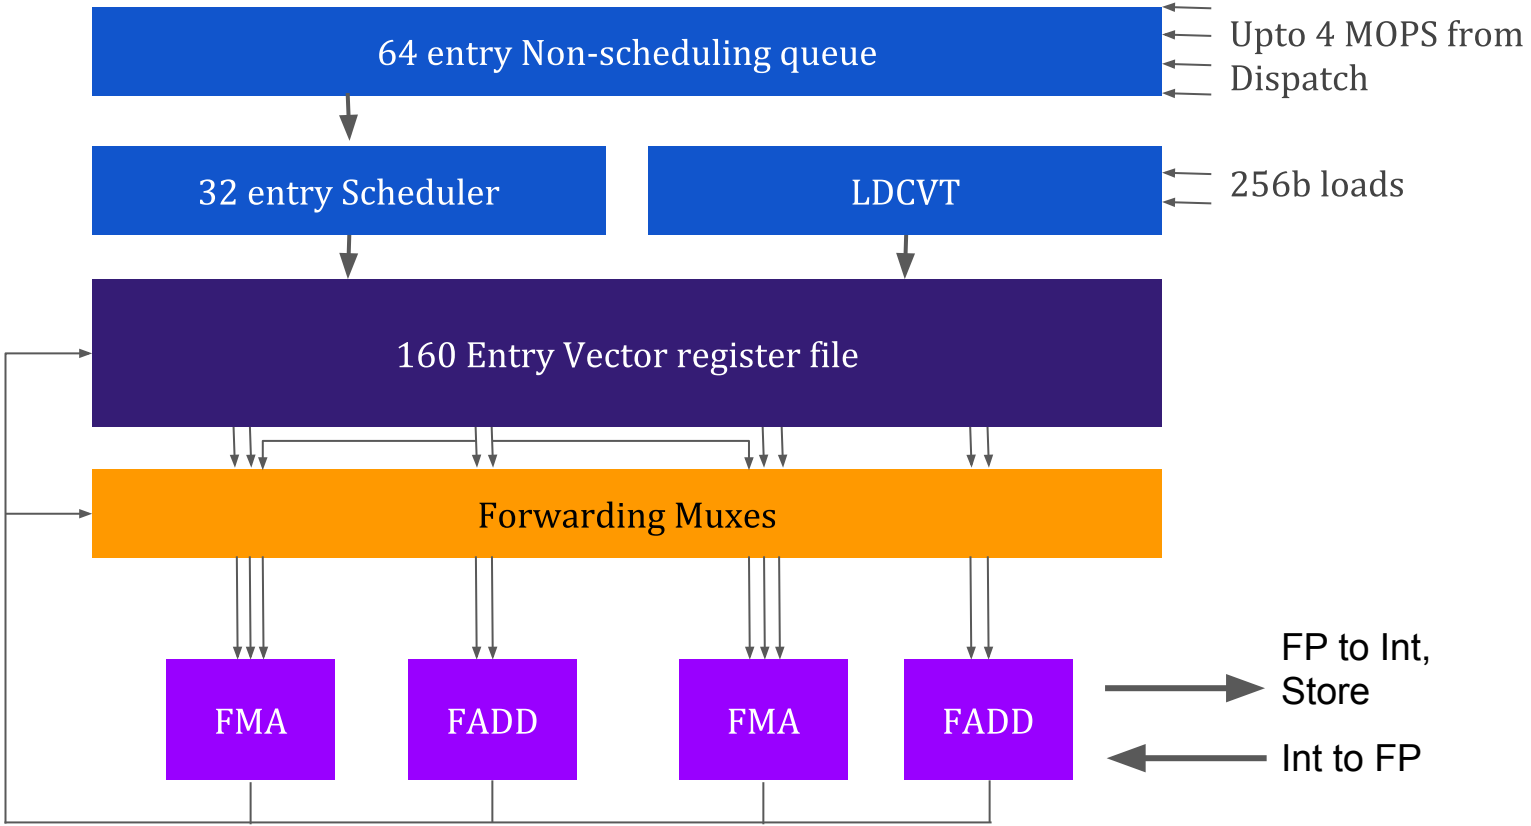

Source: https://www.slideshare.net/AMD/the-path-to-zen-2

Memory Hierarchy

Memory Stage

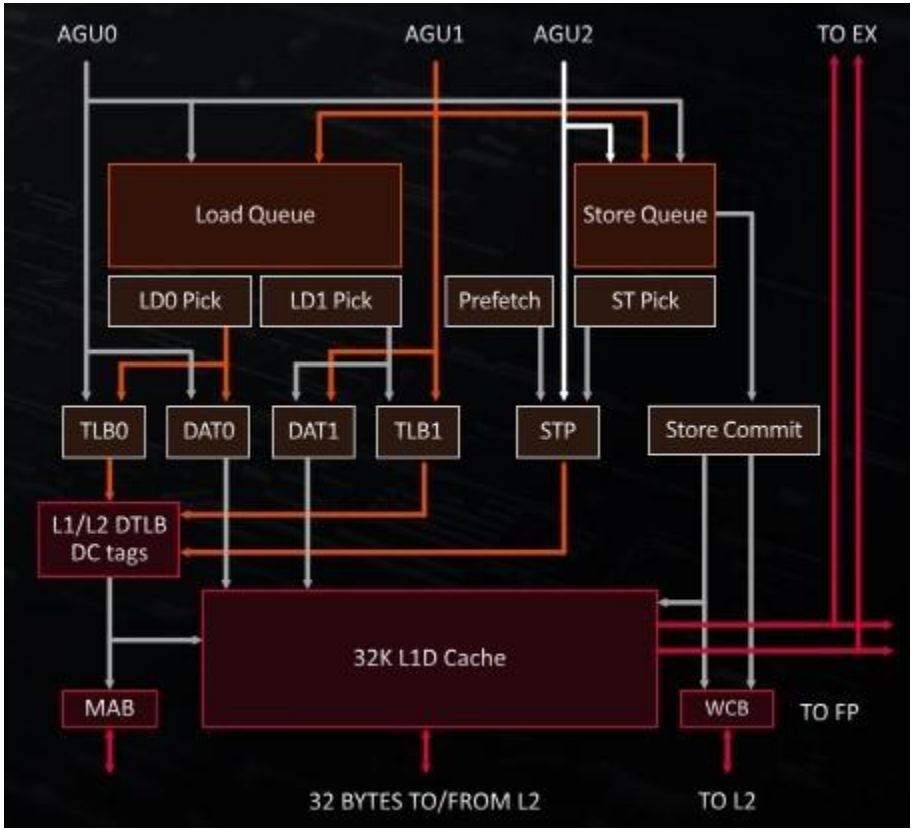

Source: https://www.slideshare.net/AMD/the-path-to-zen-2

Cache Hierarchy

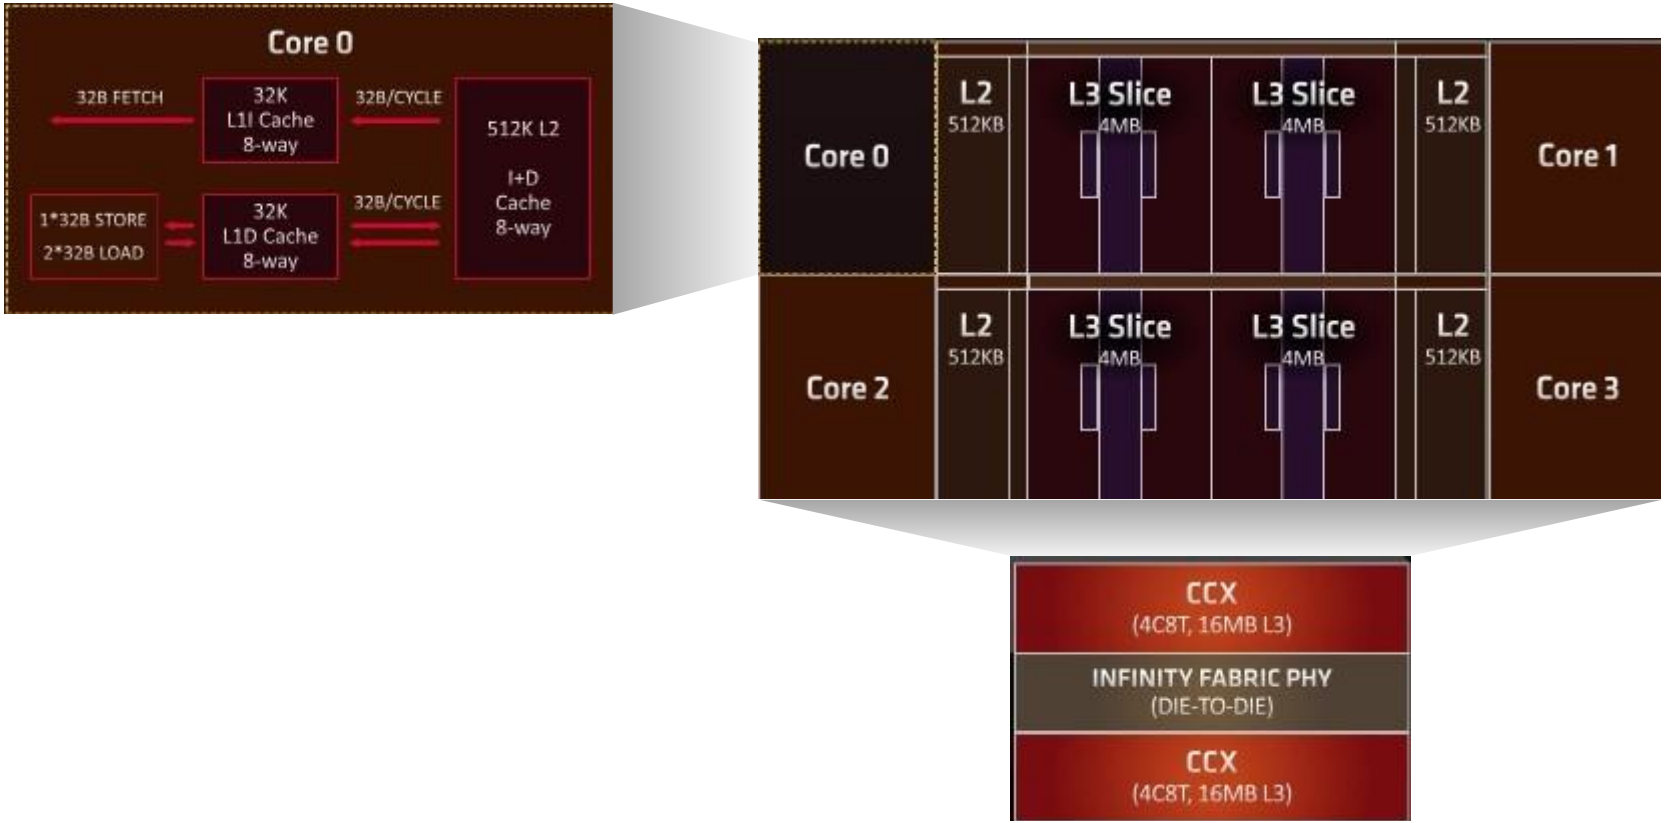

Source: https://www.slideshare.net/AMD/the-path-to-zen-2

Infinity Fabric

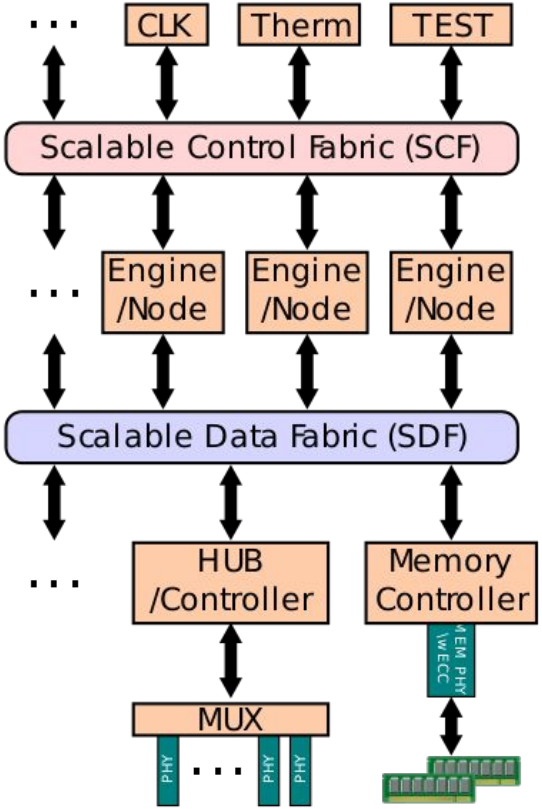

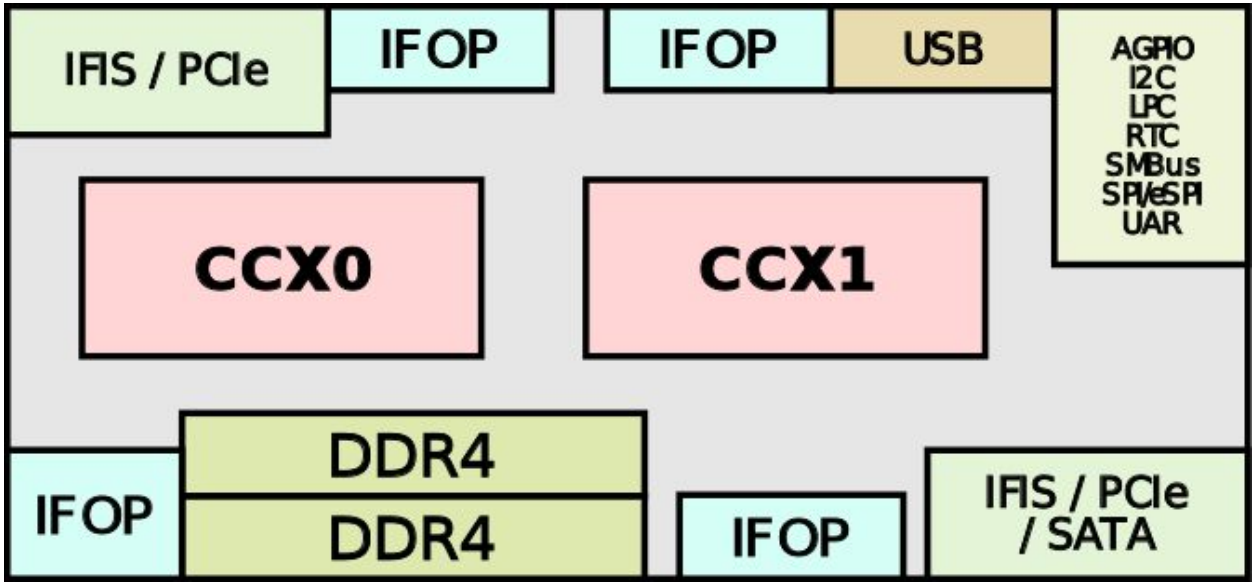

Source: https://en.wikichip.org/wiki/amd/infinity_fabric

Infinity Fabric

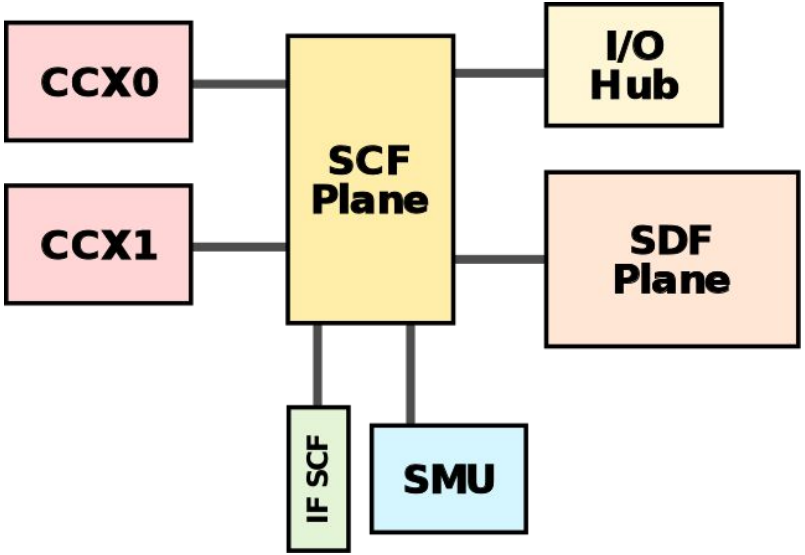

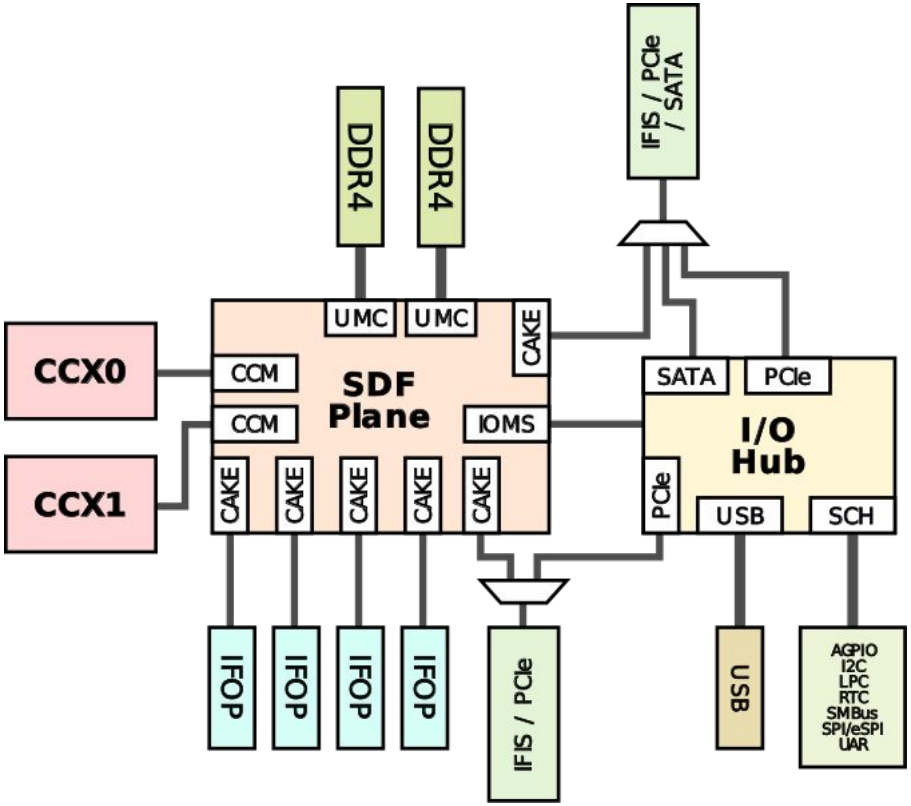

Source: https://en.wikichip.org/wiki/amd/infinity_fabric

Security and Power

SME and SEV

- Need for memory encryption (especially in servers) :
  - Plain text data in memory susceptible
- **Secure Memory Encryption (SME)** : x86 extension for page-granular memory encryption support using a single 128-bit AES key
  - Marked pages are automatically decrypted and encrypted
  - Unmarked pages suffer no overheard

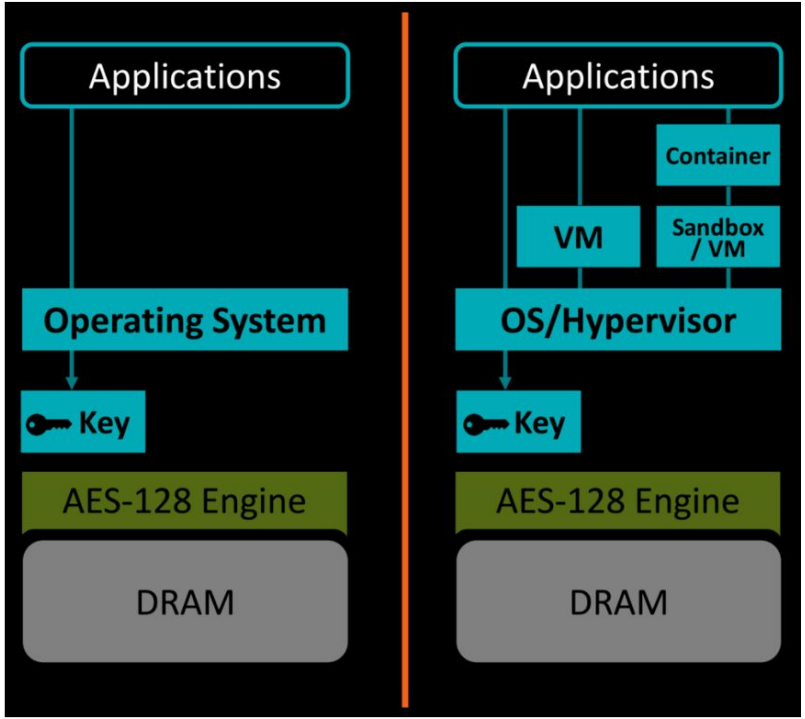

https://en.wikichip.org/wiki/amd/microarchitectures/zen#Security
https://en.wikichip.org/wiki/x86/sme

SME and SEV

- **Secure Encrypted Virtualization (SEV)** : extension of SME that enables a per-virtual machine SME using per-VM key
- Key management by **AMD Secure Processor** making the key inaccessible to software

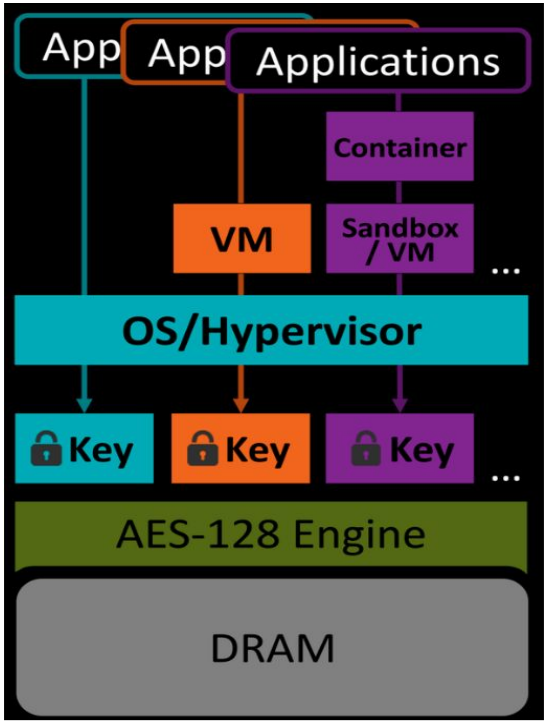

Security for SMT

- **Branch predictor tagged per thread** to prevent interference between threads (both innocent and malicious)
- **Load-Store queues also tagged per thread** to prevent any store-to-load forwarding of data from the wrong thread

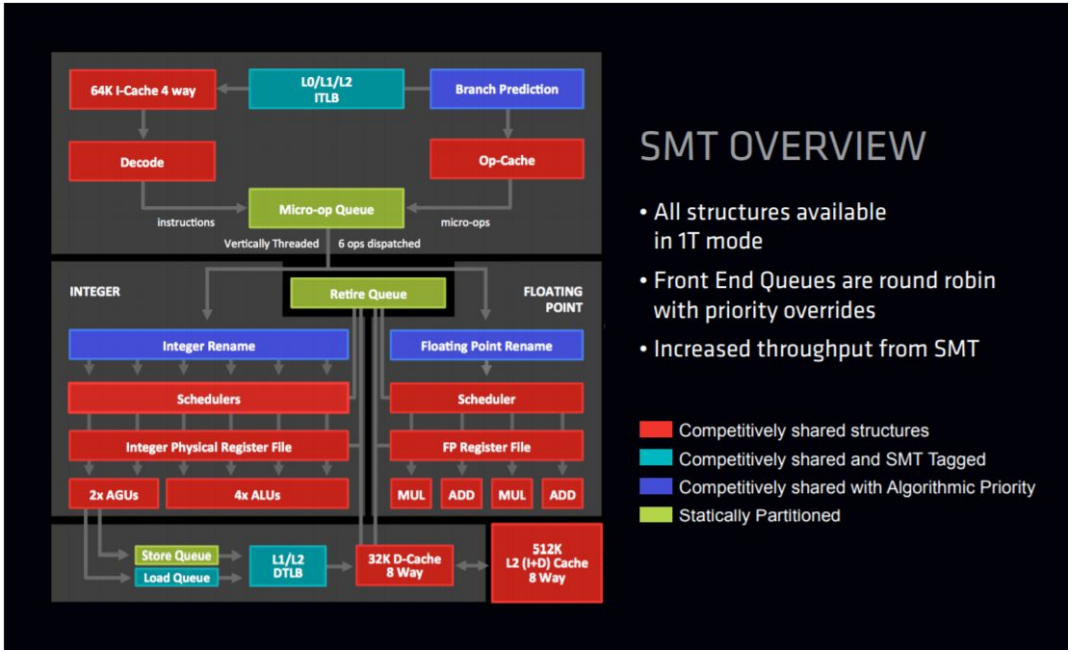

https://www.amd.com/system/files/documents/security-whitepaper.pdf

Security for speculative execution attacks

- **Spectre** : Exploits various kinds of mis-speculation to access data the program was not supposed to
- Observable side effects caused by speculation exploited by attacks
- Zen 2 has hardware mitigation against the Spectre V4 speculative store bypass vulnerability

https://meltdownattack.com/
https://www.amd.com/system/files/documents/security-whitepaper.pdf
https://www.techpowerup.com/256478/amd-zen-2-has-hardware-mitigation-for-spectre-v4

Power Management

- **Dynamic voltage-frequency scaling (DVFS)** used for reducing power
- **Per-core digital LDO regulator** and digital frequency synthesizer (DFS) to vary frequency and voltage across power states on individual core basis
- LDO regulates RVDD for each power domain and create an optimal VDD per core using a **system of sensors** embedded across the entire chip
- Enables **fine power tuning** on a per core level based on information collected from the core and overall chip

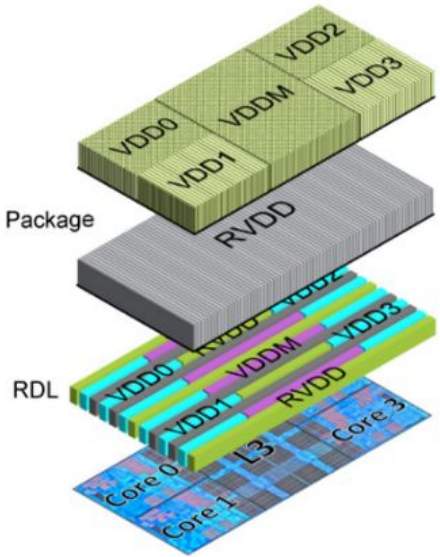

- **RDL** - Redistribution layer
- **LDOs** - Regulate RVDD to create VDD per core
- **RVDD** - Ungated supply
- **VDD** - Gated core supply
- **VDDM** - L2/L3 SRAM supply

https://en.wikichip.org/wiki/amd/microarchitectures/zen#Power

Takeaways

1. Zen2 is an **x86** core architecture which implements **SMT** and scales up with **MCM** packaging
2. Higher IPC due to improved BPU, optimized L1I cache and Op cache
3. Zen2 uses scalable Infinity fabric interconnect for faster, more power efficient memory accesses
4. Zen2 has spectre mitigations in hardware
5. Zen has per-core power tuning

Thank you

Appendix

Performance

According to AMD evaluations :

- **15%** performance gain over Zen+ compared at same frequency
- Zen 2 offers >1.25x performance gain at the same power, or up to half power at same performance
- For select benchmarks, **+75% performance per watt gain** over its previous generation
- 62% performance improvements result of architectural enhancements, and remaining 38% due to 7nm manufacturing process

https://hothardware.com/reviews/amd-zen-2-architecture-explained?page=3

Zen Simultaneous Multi-threading (SMT)

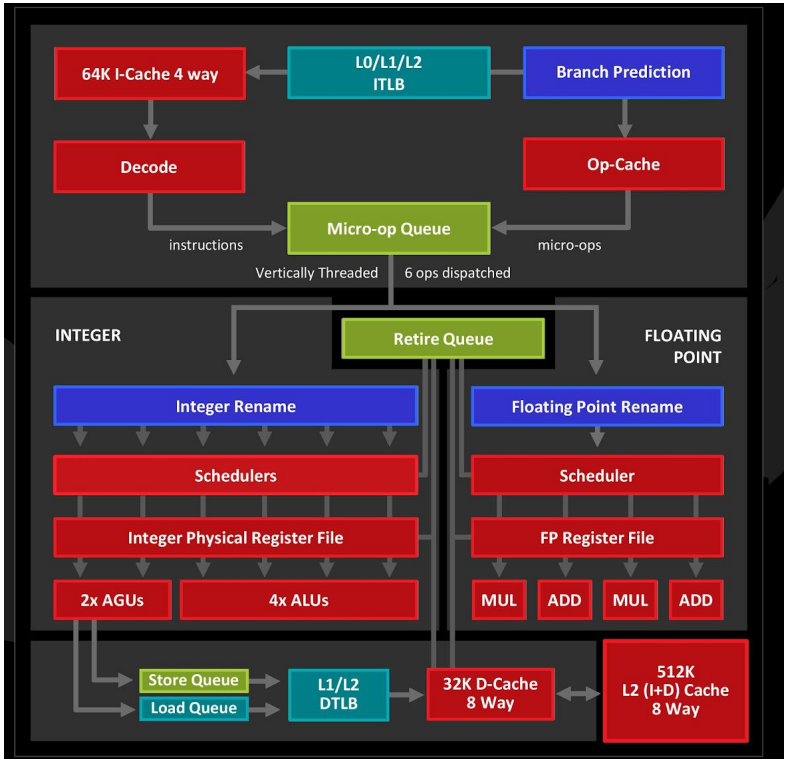

- Shared resources

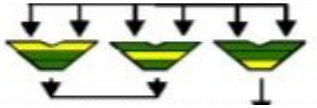

- Shared with **priority**

- Statically **partitioned**

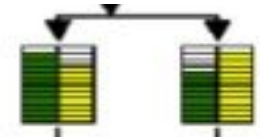

- Shared and **tagged**

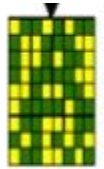

AMD Presentation on Hot Chips, Aug 2016
https://www.slideshare.net/am_sharifian/intel-hyper-threading-technology