



Intel® Nehalem Micro-Architecture

Mohammad Radpour

Amirali Sharifian

Outline

- Brief History
- Overview
- Memory System
- Core Architecture
- Hyper- Threading Technology
- Quick Path Interconnect (QPI)



Brief History about Intel Processors

8086. Real Mode.

286. 1st generation, 16-bit Protected Mode added.

386. 2nd generation 32-bit Protected Mode + 1st generation Paging + SMM + VM86 Mode + 32-bit GP registers.

Pentium P55C. MMX/SIMD paradigm + MMX instruction and register sets added + Local APIC.

Pentium Pro. 2nd generation Paging (PAE-36) permits physical memory addressing above 4GB boundary.

Pentium III. Expansion of SIMD paradigm. SSE instruction set + XMM register set added.

Pentium 4. Netburst microarchitecture + Hyper-Threading.

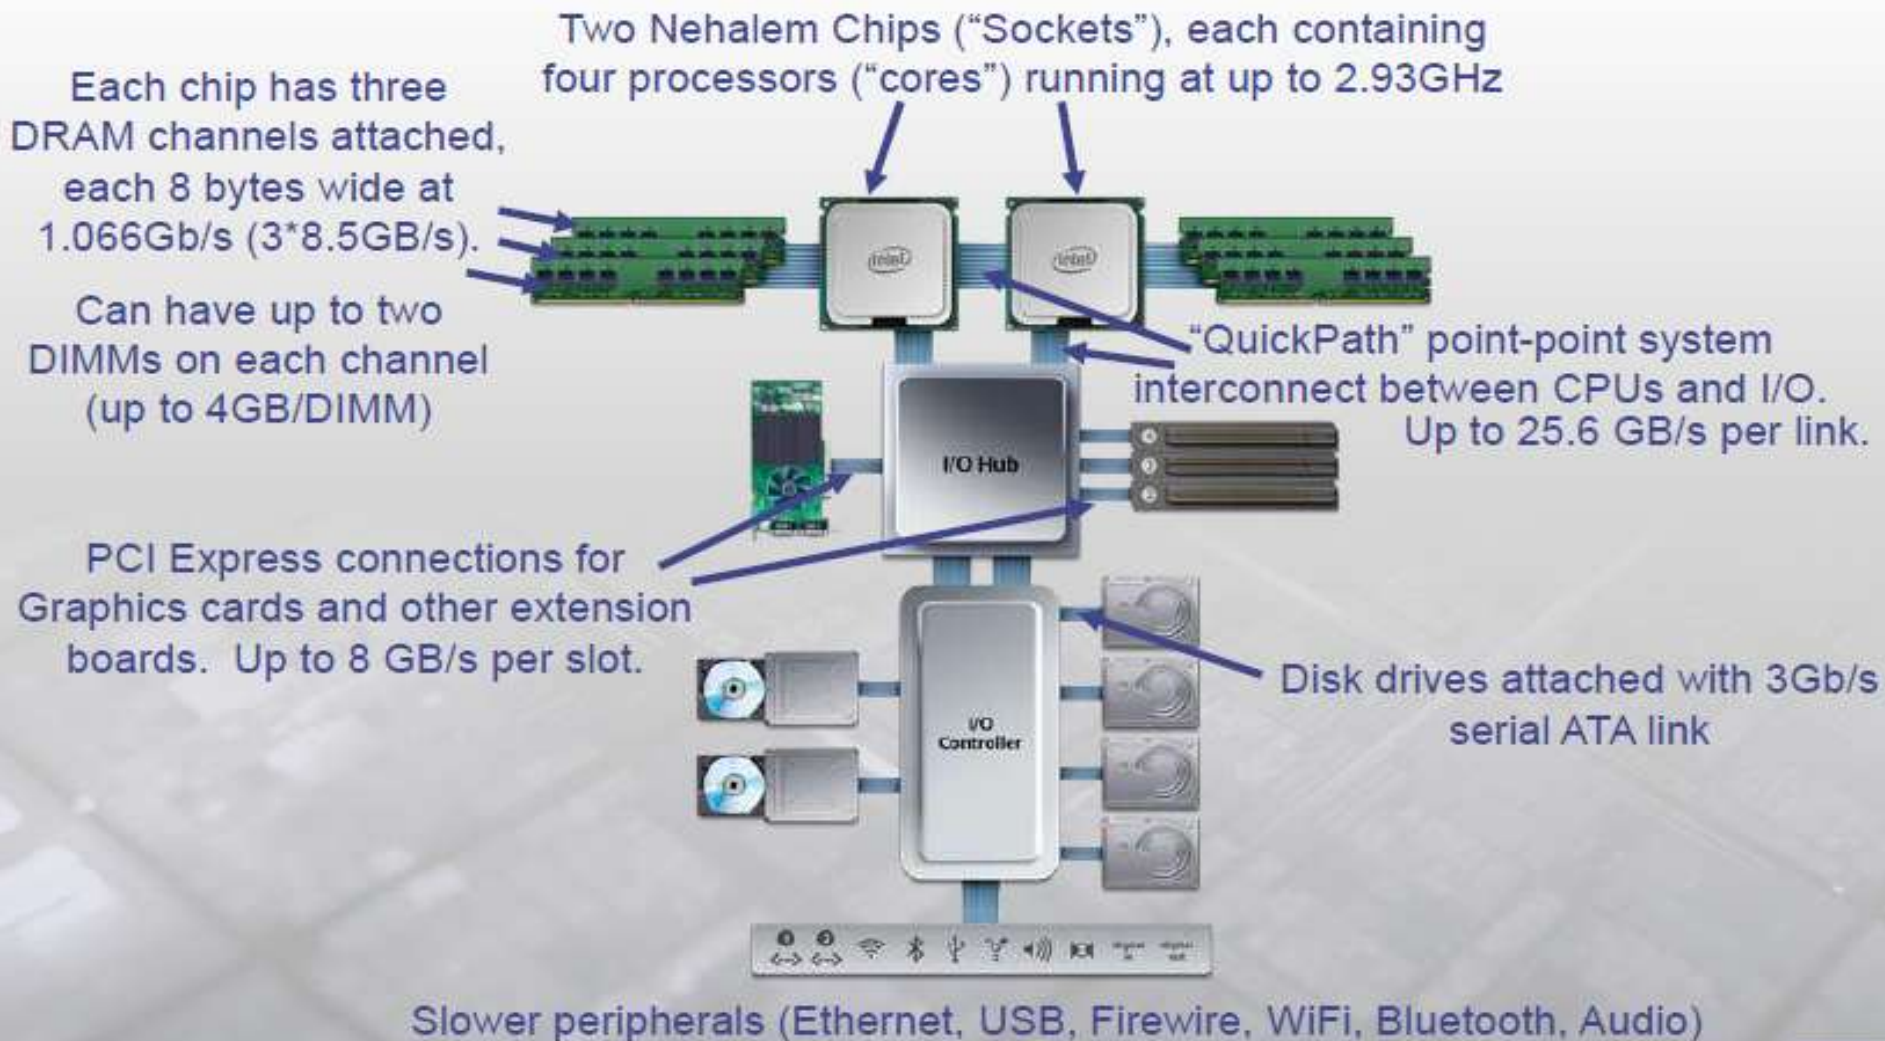
Core Solo/Duo. Virtualization Technology added.

Core 2 Solo/Duo. 3rd generation Paging + IA-32e Mode (64-bit register set + 64-bit addressing added).

Core i7 (first Nehalem-based product). Hyper-Threading back + up to 8 cores + Integrated DRAM controller + QPI.

Note: Smaller, incremental evolutionary steps not included (e.g., SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, etc.)


Nehalem System Example:



Building Blocks

Nehalem Design Scalable Via Modularity

Nehalem Building Block Library



Ex: 4 Core

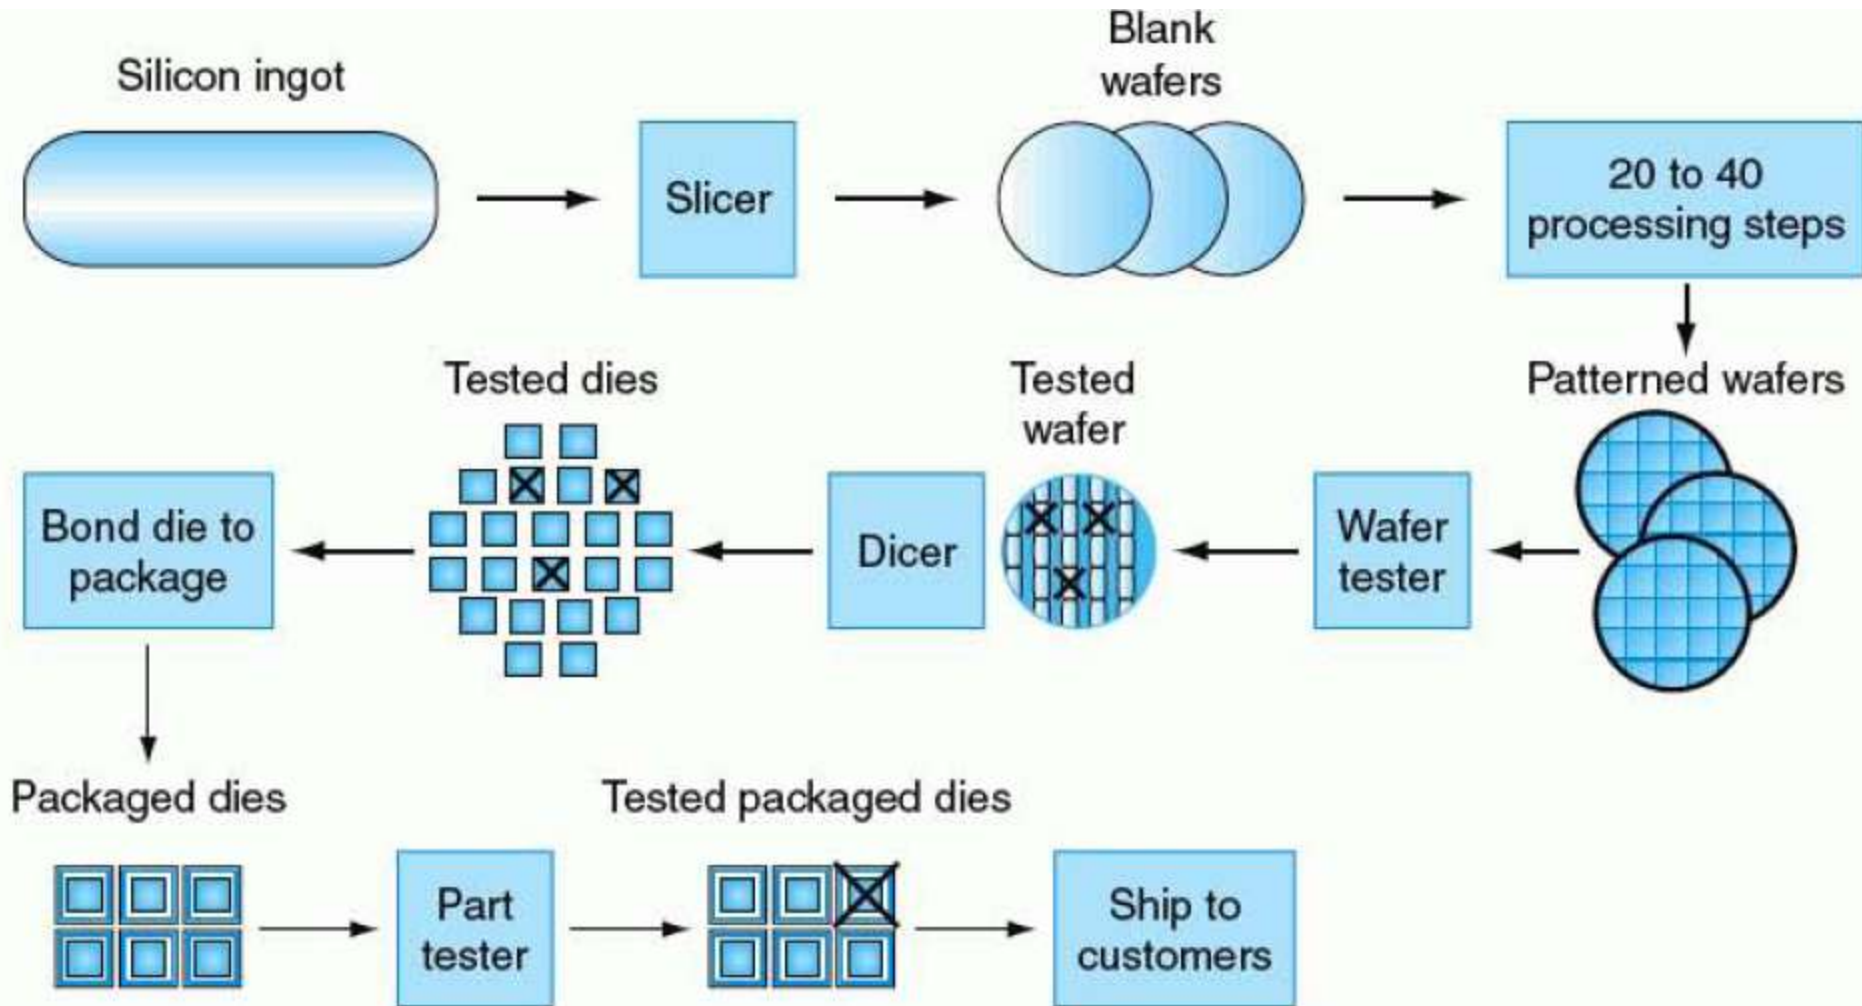


Ex: 8 Core




Sample Range of Product Options

How to make Silicon Die?



Overview of Nehalem Processor Chip


- Four identical compute core
- UIU: Un-core interface unit
- L3 cache memory and data block memory



B. Nehalem micro-photograph.

Overview of Nehalem Processor Chip(cont.)

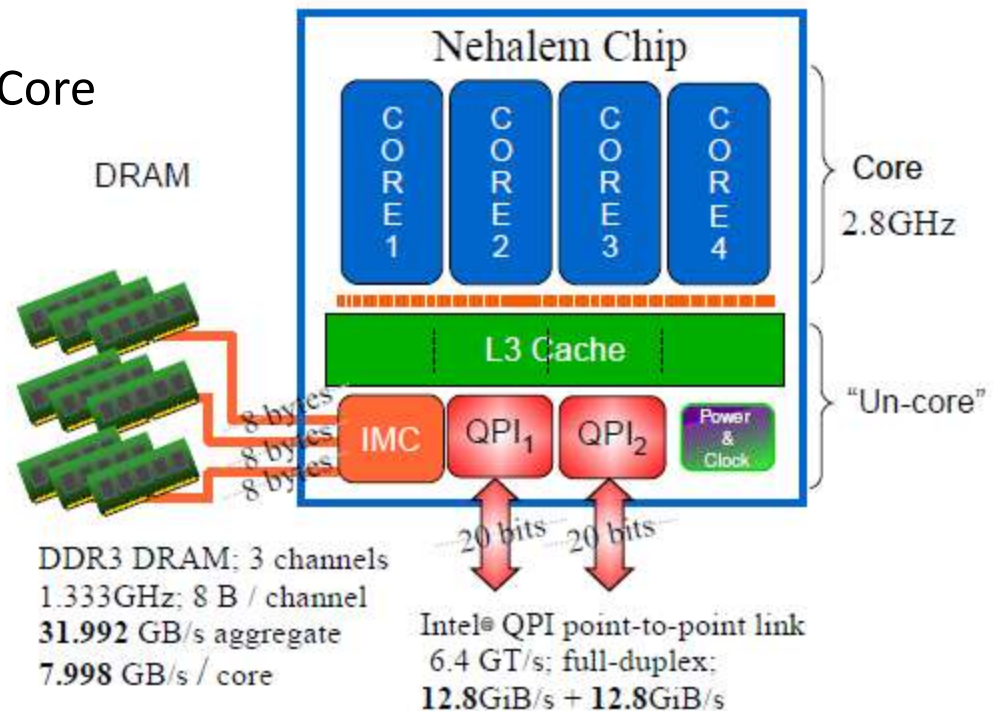
- IMC : Integrated Memory Controller with 3 DDR3 memory channels
- QPI : Quick Path Interconnect ports
- Auxiliary circuitry for cache-coherence, power control, system management, performance monitoring



B. Nehalem micro-photograph.

Overview of Nehalem Processor Chip(cont.)

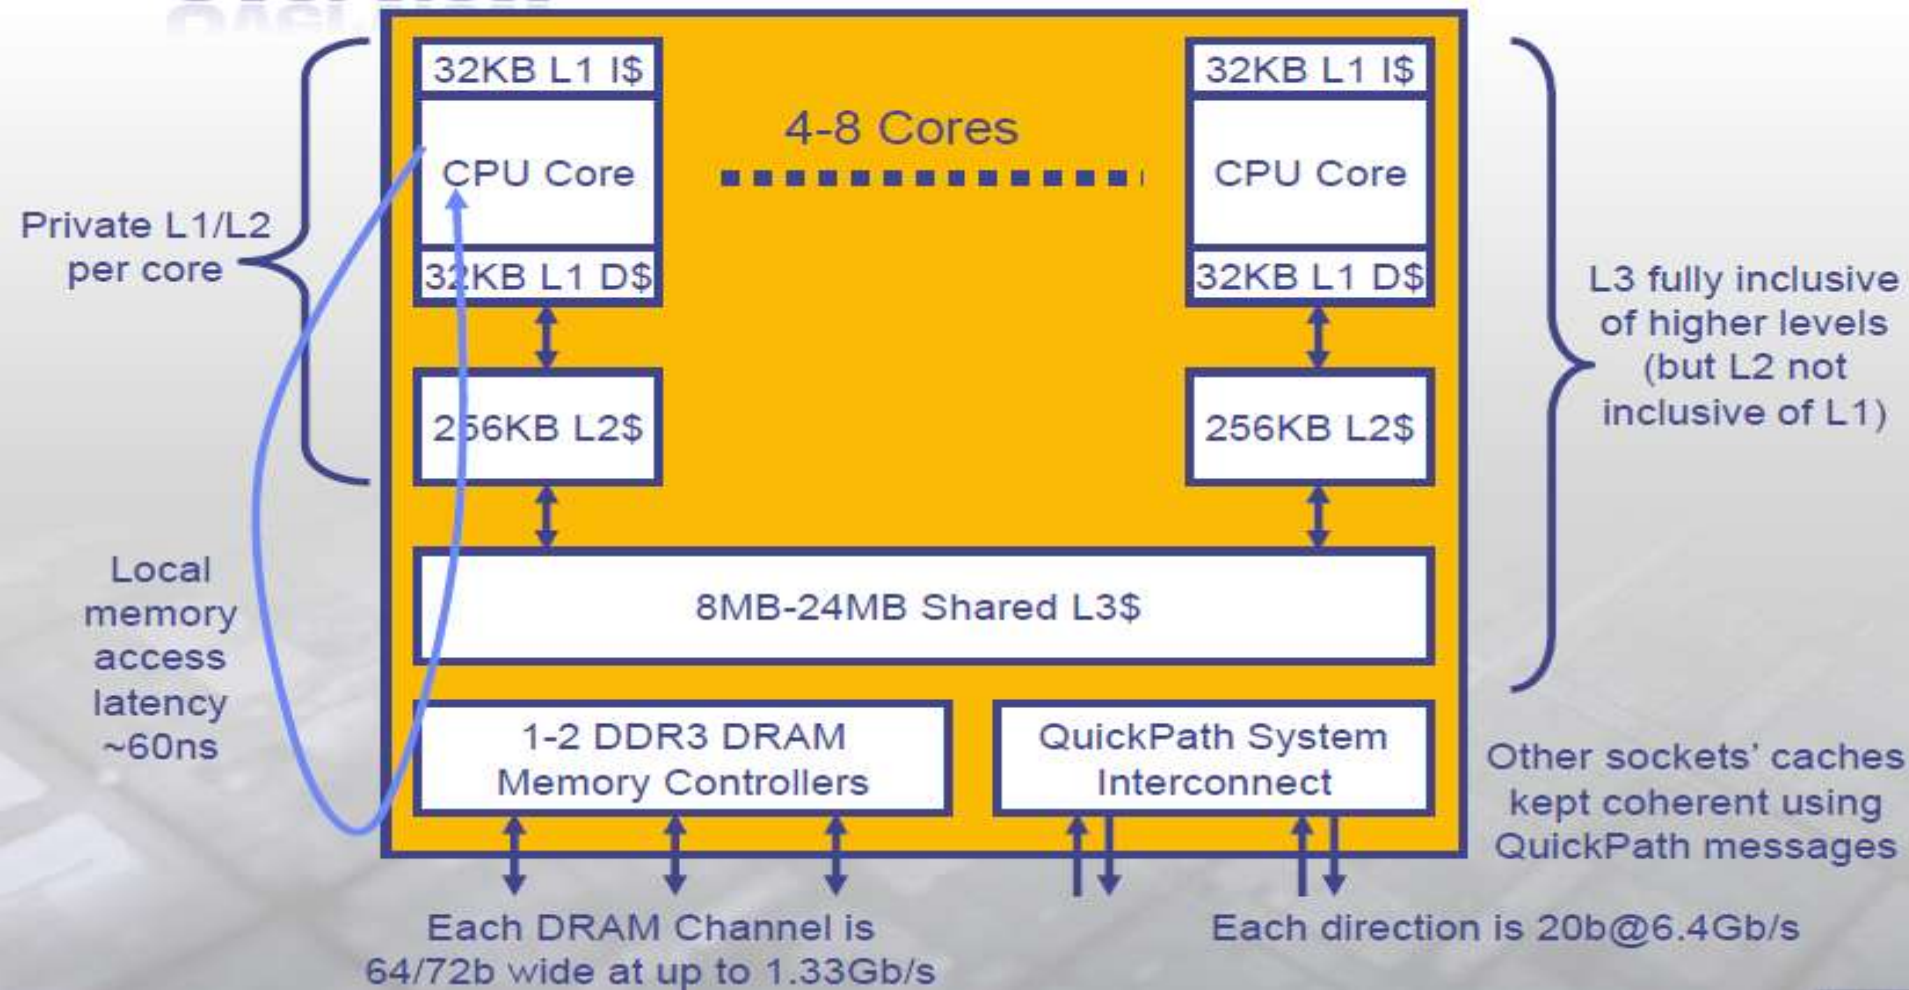
- Chip is divided into two domains:
 - “Un-core” and “core”
- “Core” components operate with a same clock frequency of the actual Core
- “Un-Core” components operate with different frequency.





Memory System and Core Architecture

Nehalem Memory Hierarchy Overview

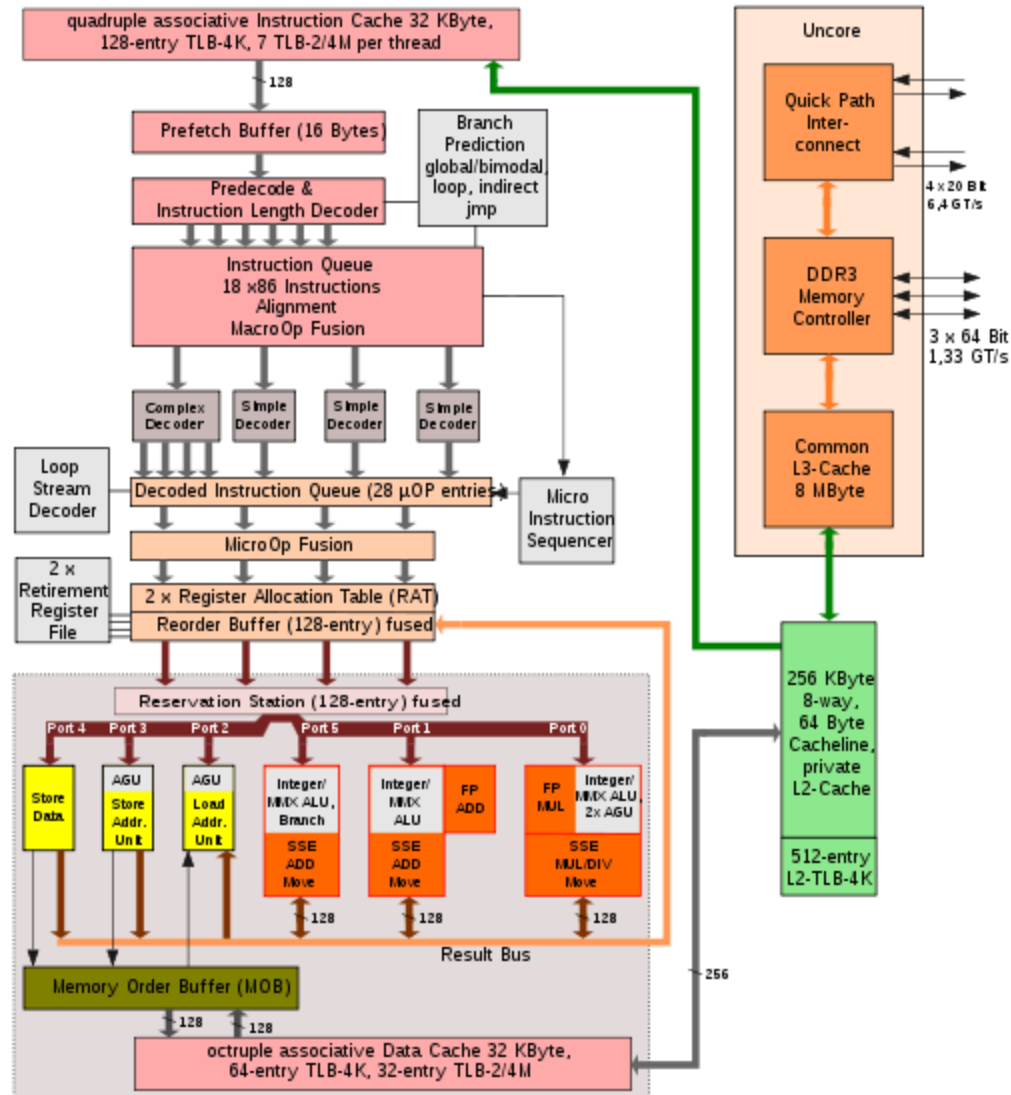


Cache Hierarchy Latencies

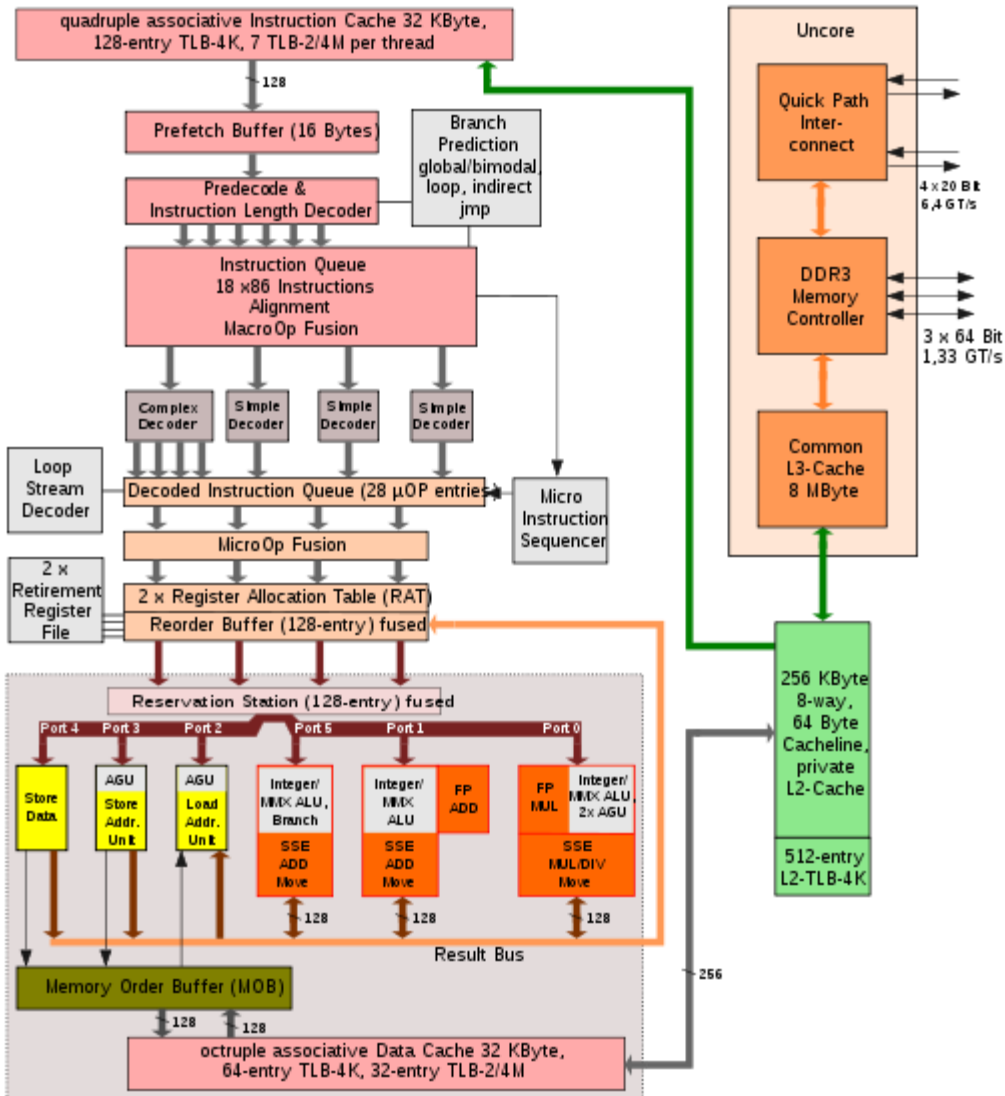
- L1 32KB 8-way, Latency 4 cycles
- L2 256KB 8-way, Latency < 12 cycles
- L3 8MB shared , 16-way, Latency 30-40 cycles (4 core system)
- L3 24MB shared, 24-way, Latency 30-60 cycles(8 core system)
- DRAM , Latency ~ 180 – 200 cycles

Intel® Smart Cache – Level 3

Nehalem Microarchitecture

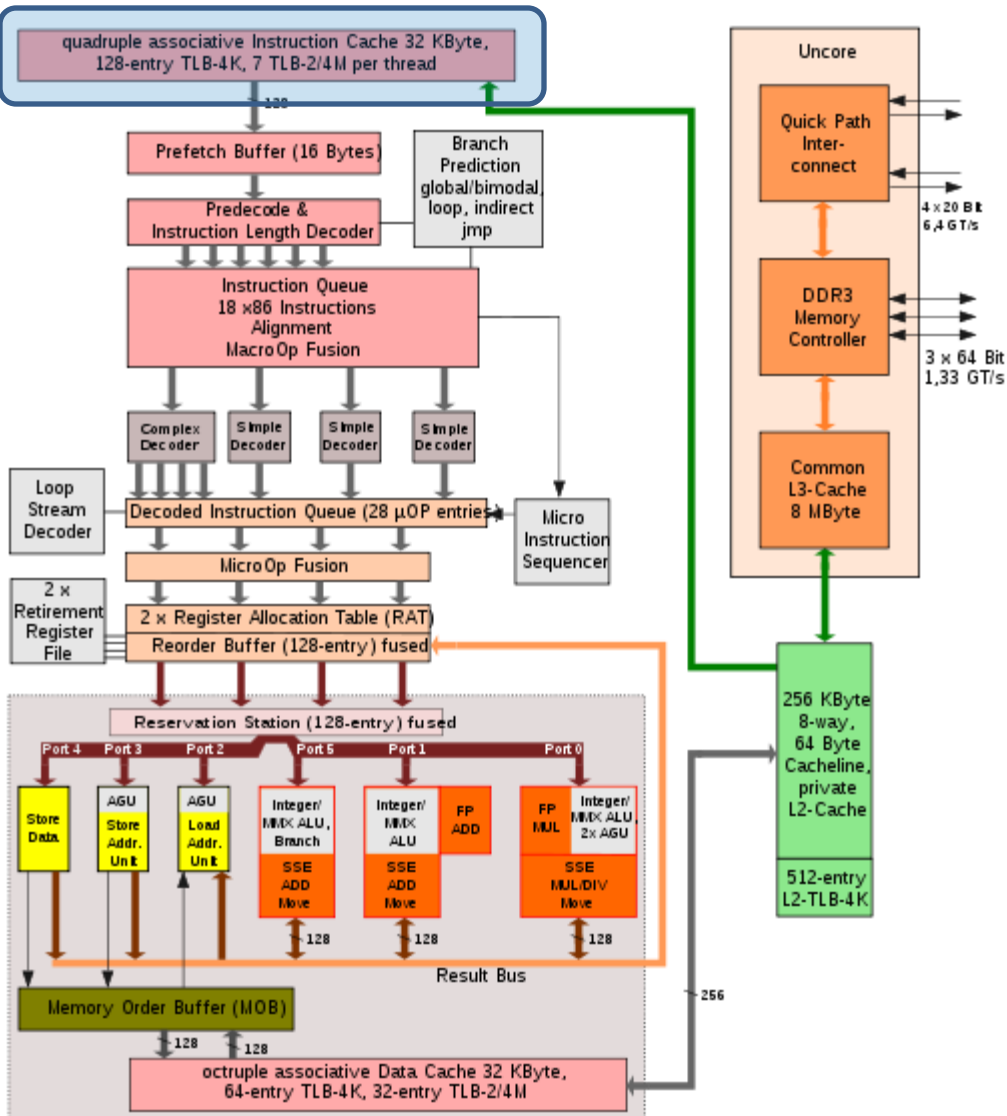


Instruction Execution



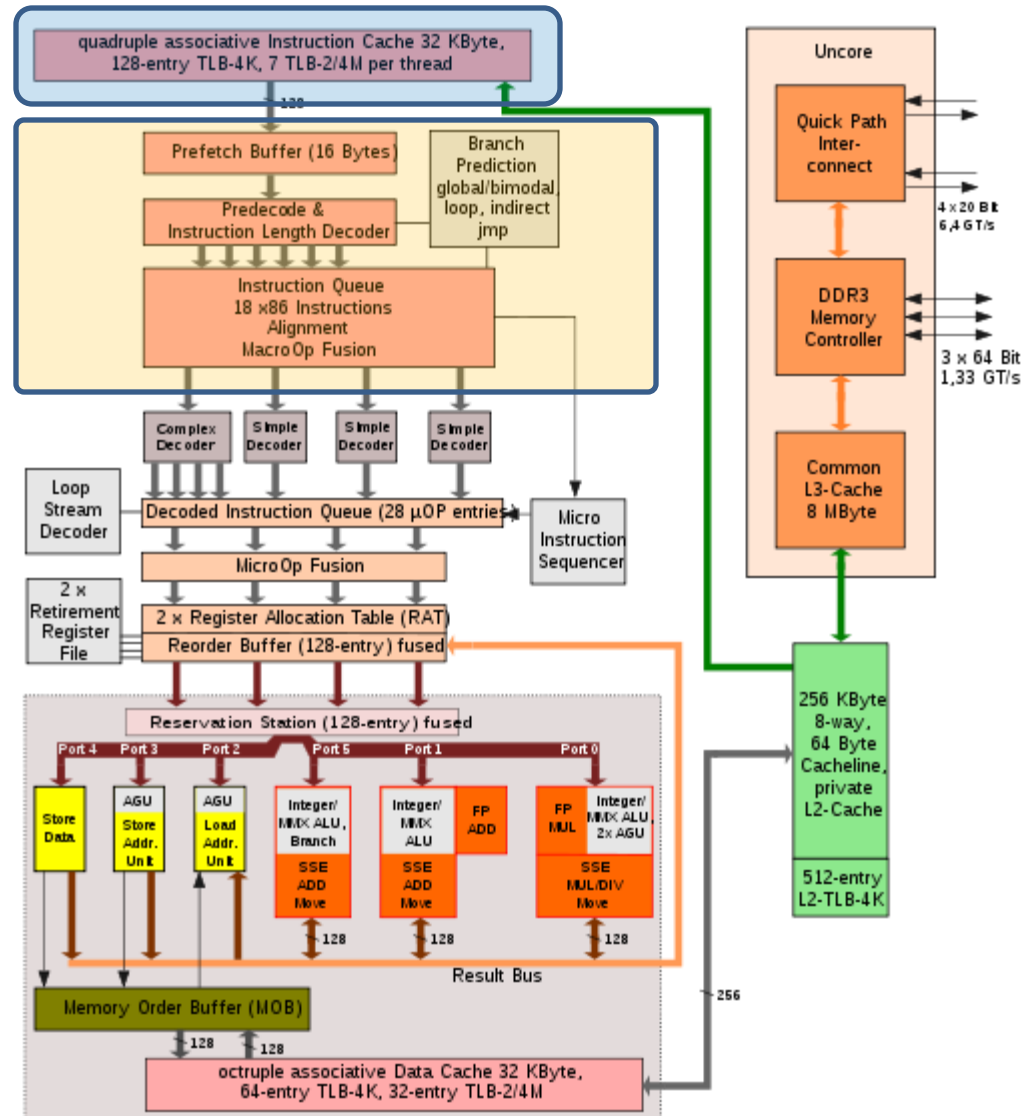
Instruction Execution (1/5)

1. Instructions fetched from L2 cache



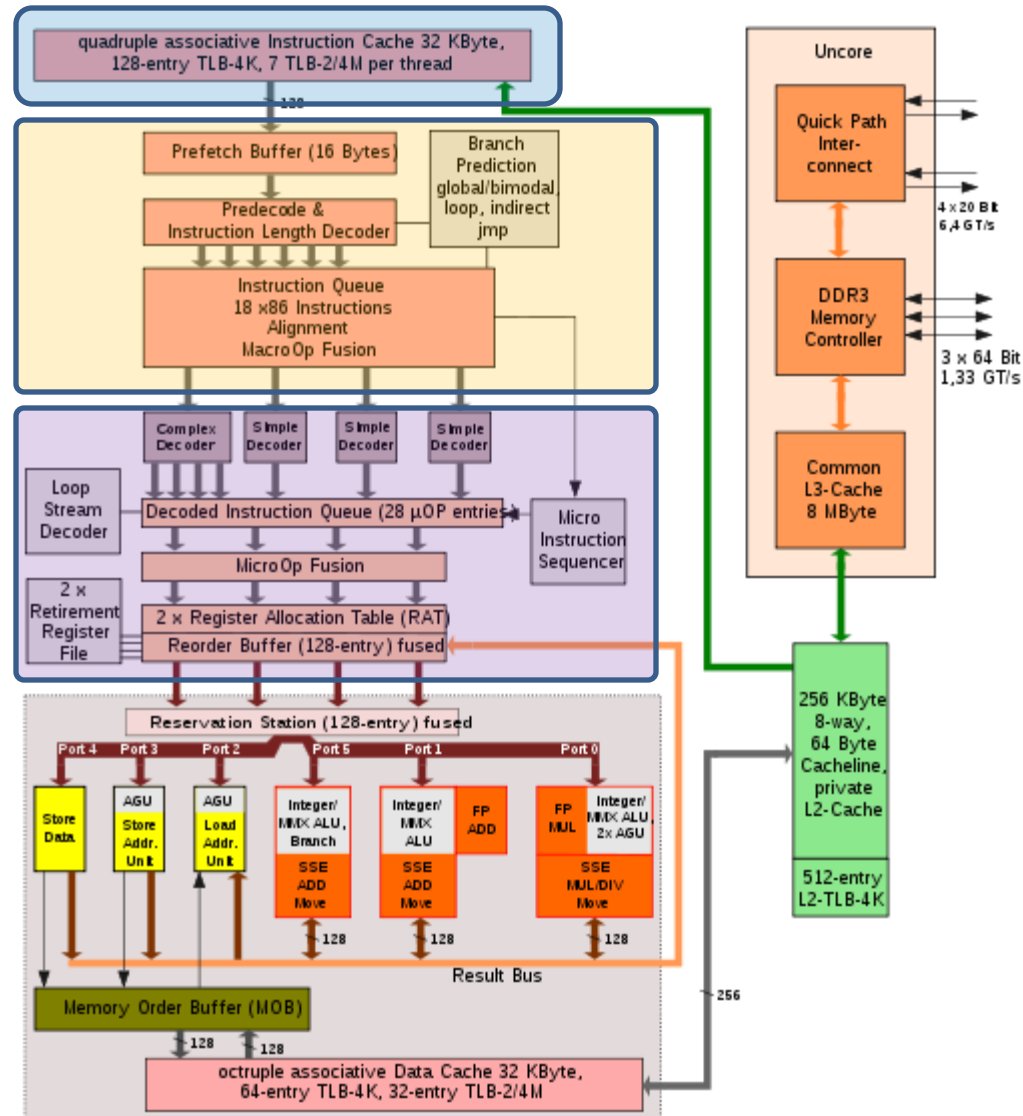
Instruction Execution (2/5)

1. Instructions **fetched** from L2 cache
2. Instructions **decoded, prefetched and queued**

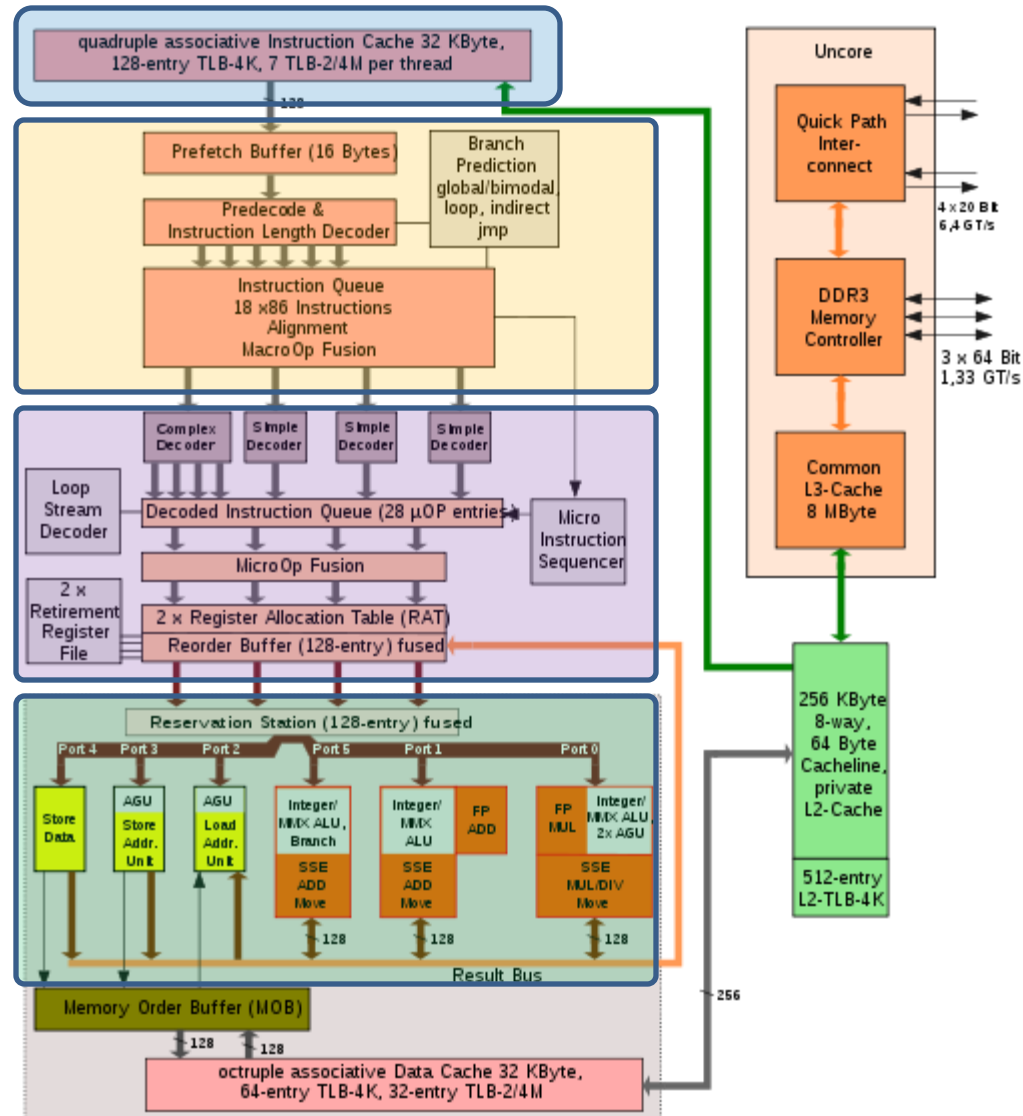


Instruction Execution (3/5)

1. Instructions fetched from L2 cache
2. Instructions decoded, prefetched and queued
3. Instructions optimized and combined

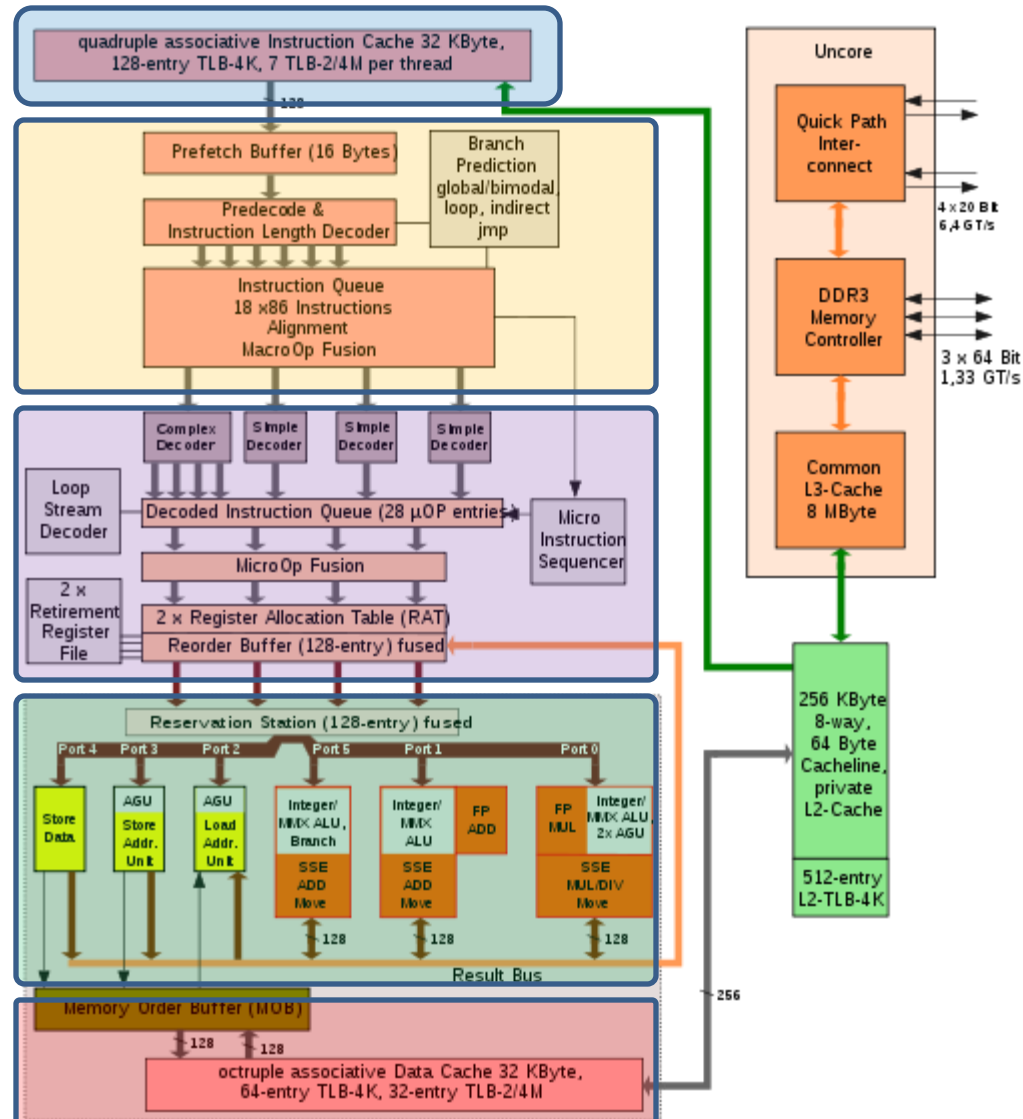


Instruction Execution (4/5)



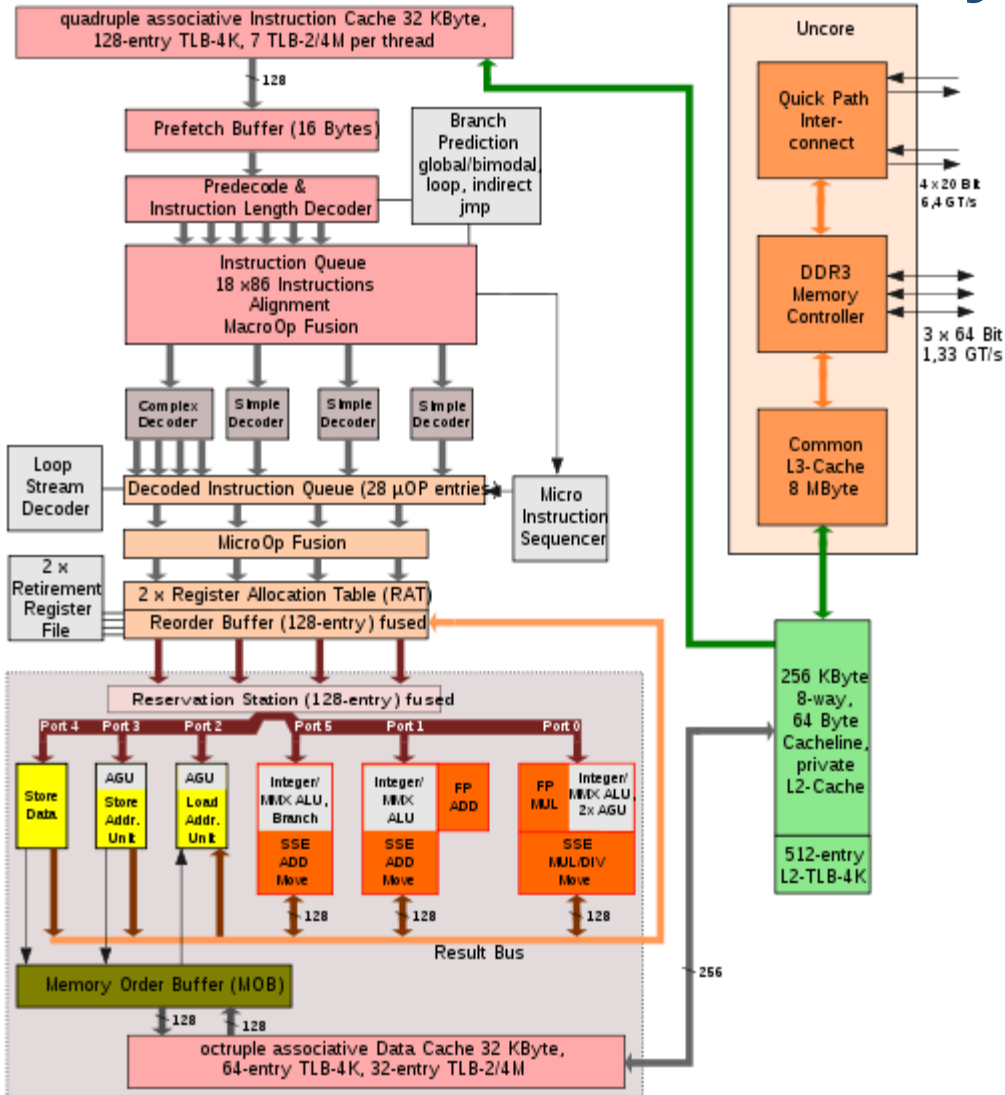
1. Instructions **fetched** from L2 cache
2. Instructions **decoded, prefetched and queued**
3. Instructions **optimized and combined**
4. Instructions **executed**

Instruction Execution (5/5)



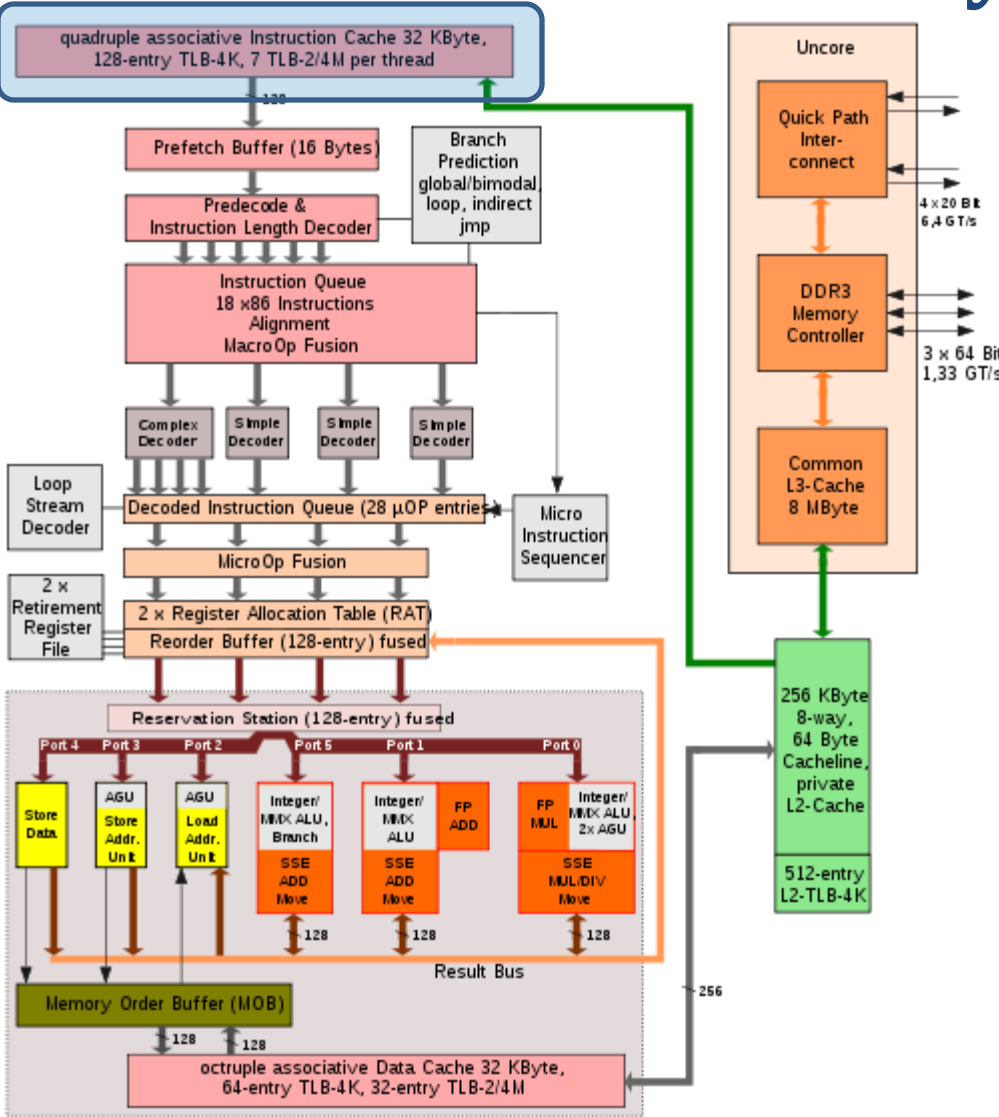
1. Instructions **fetched** from L2 cache
2. Instructions **decoded, prefetched and queued**
3. Instructions **optimized and combined**
4. Instructions **executed**
5. Results **written**

Caches and Memory

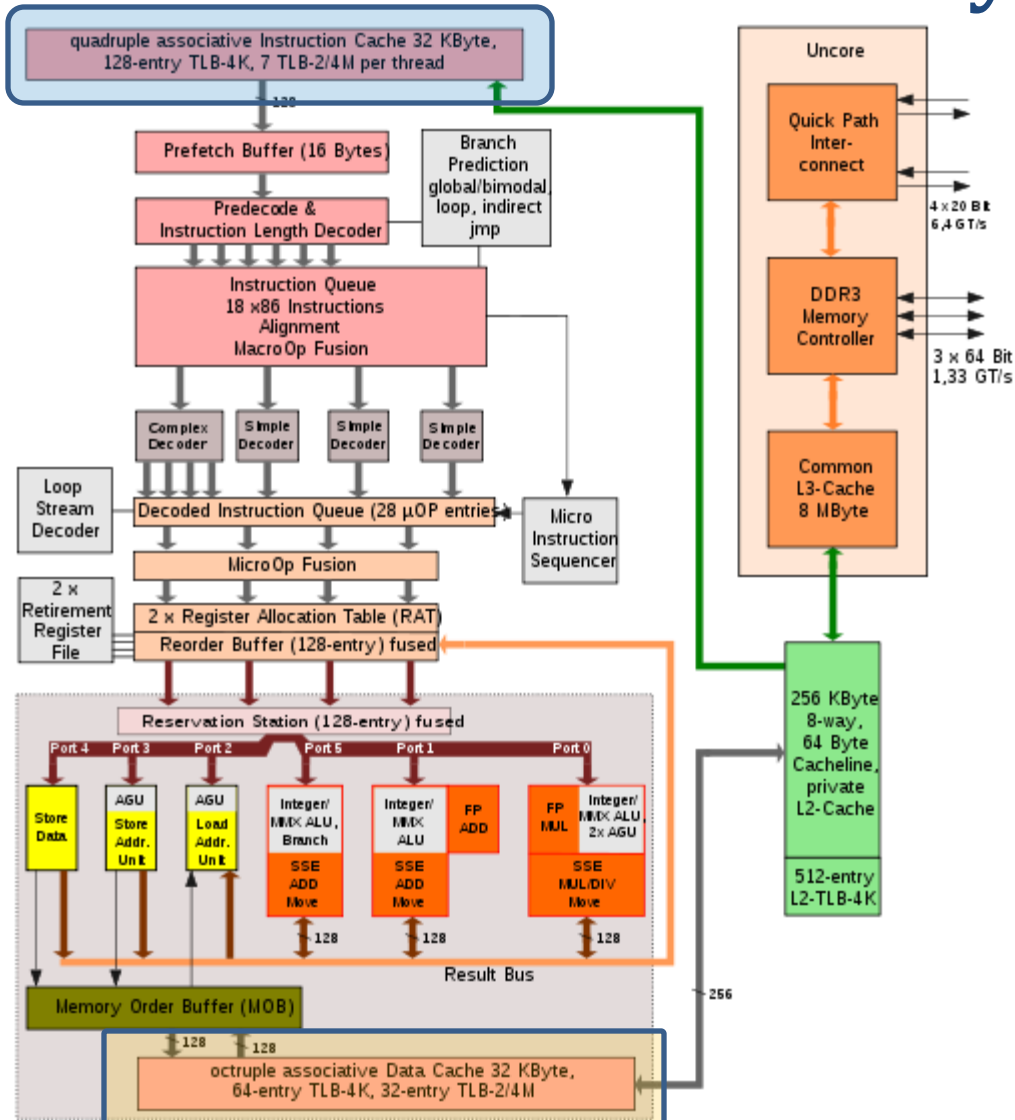


Caches and Memory (1/5)

1. 4-way set associative instruction cache



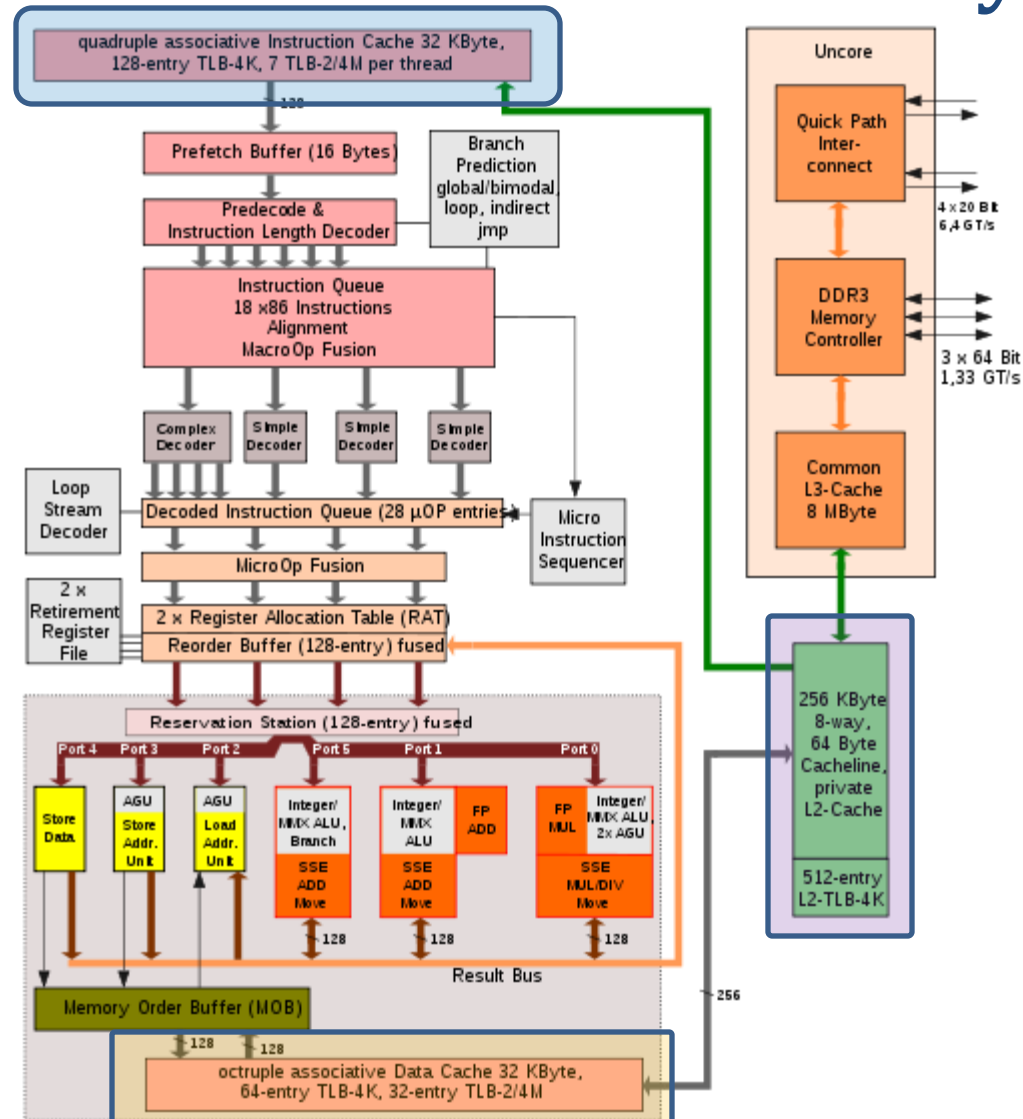
Caches and Memory (2/5)



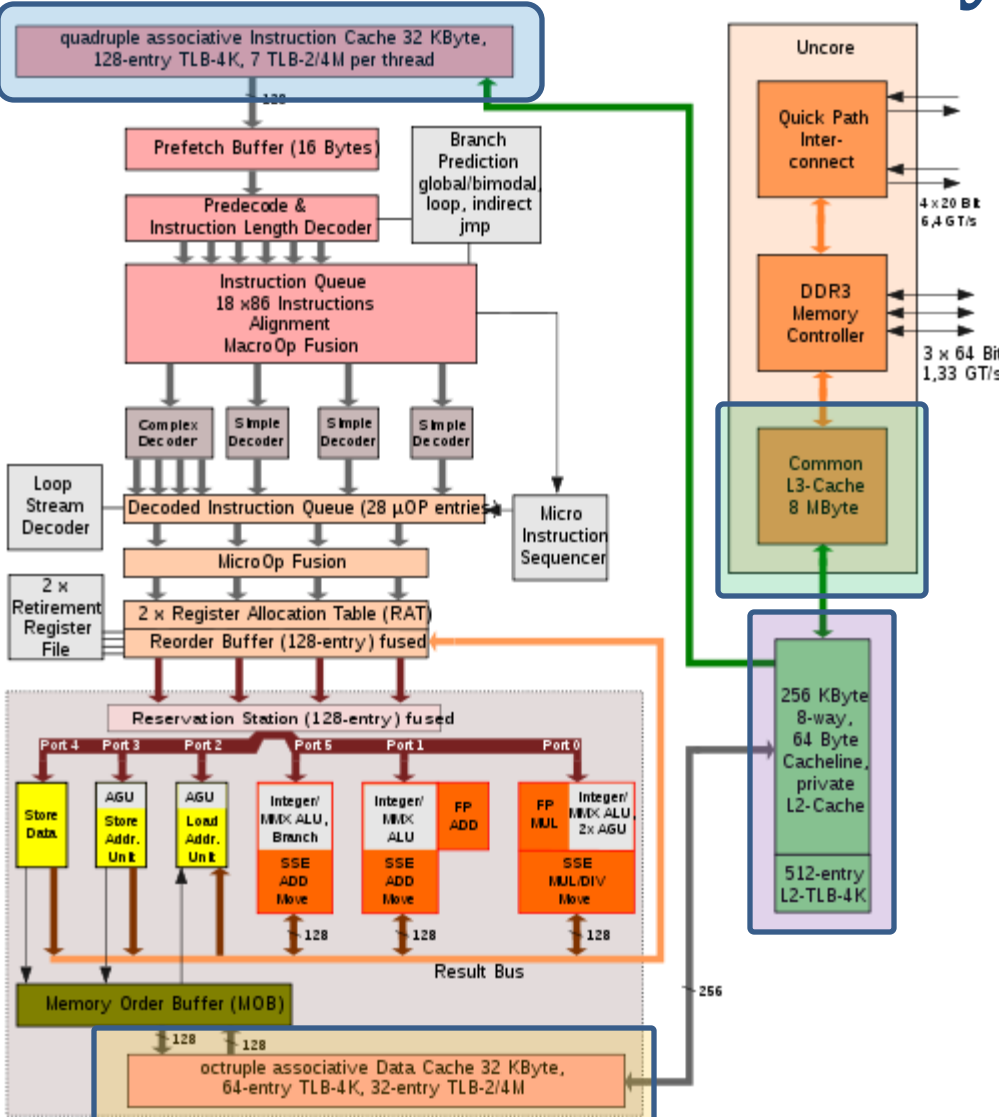
1. 4-way set associative instruction cache
2. 8-way set associative L1 data cache (32 KB)

Caches and Memory (3/5)

1. 4-way set associative instruction cache
2. 8-way set associative L1 data cache (32 KB)
3. 8-way set associative L2 data cache (256 KB)

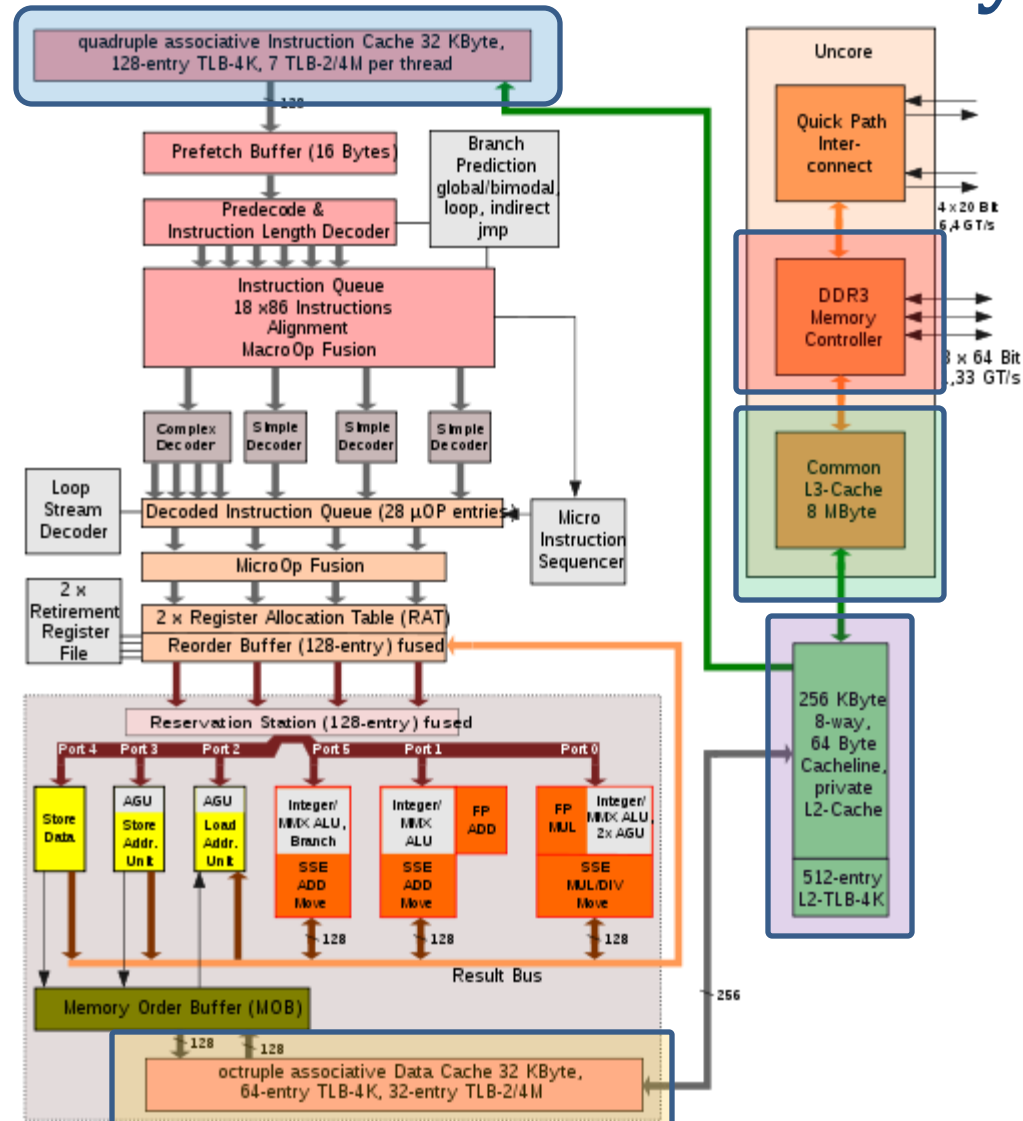


Caches and Memory (4/5)



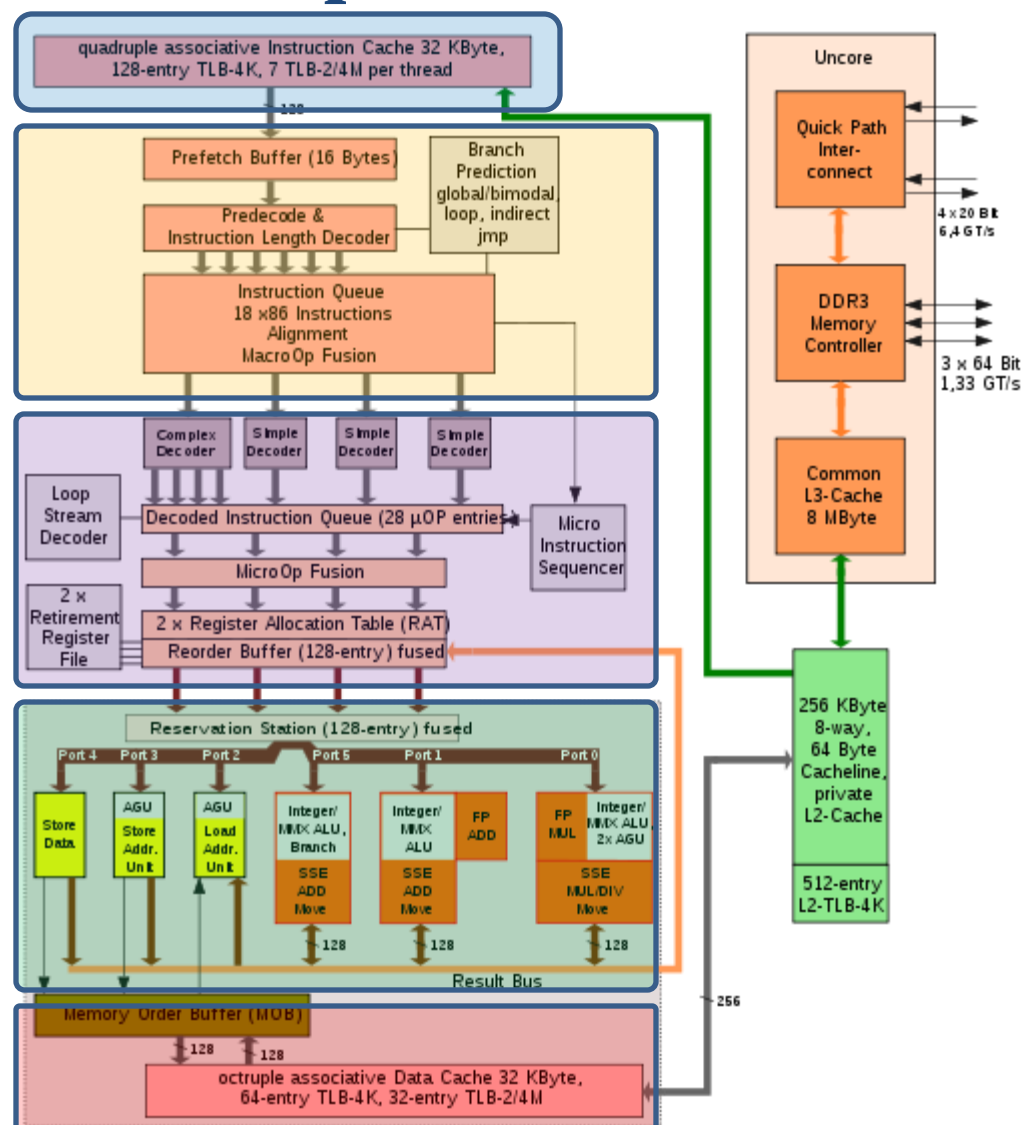
1. 4-way set associative instruction cache
2. 8-way set associative L1 data cache (32 KB)
3. 8-way set associative L2 data cache (256 KB)
4. 16-way shared L3 cache (8 MB)

Caches and Memory (5/5)



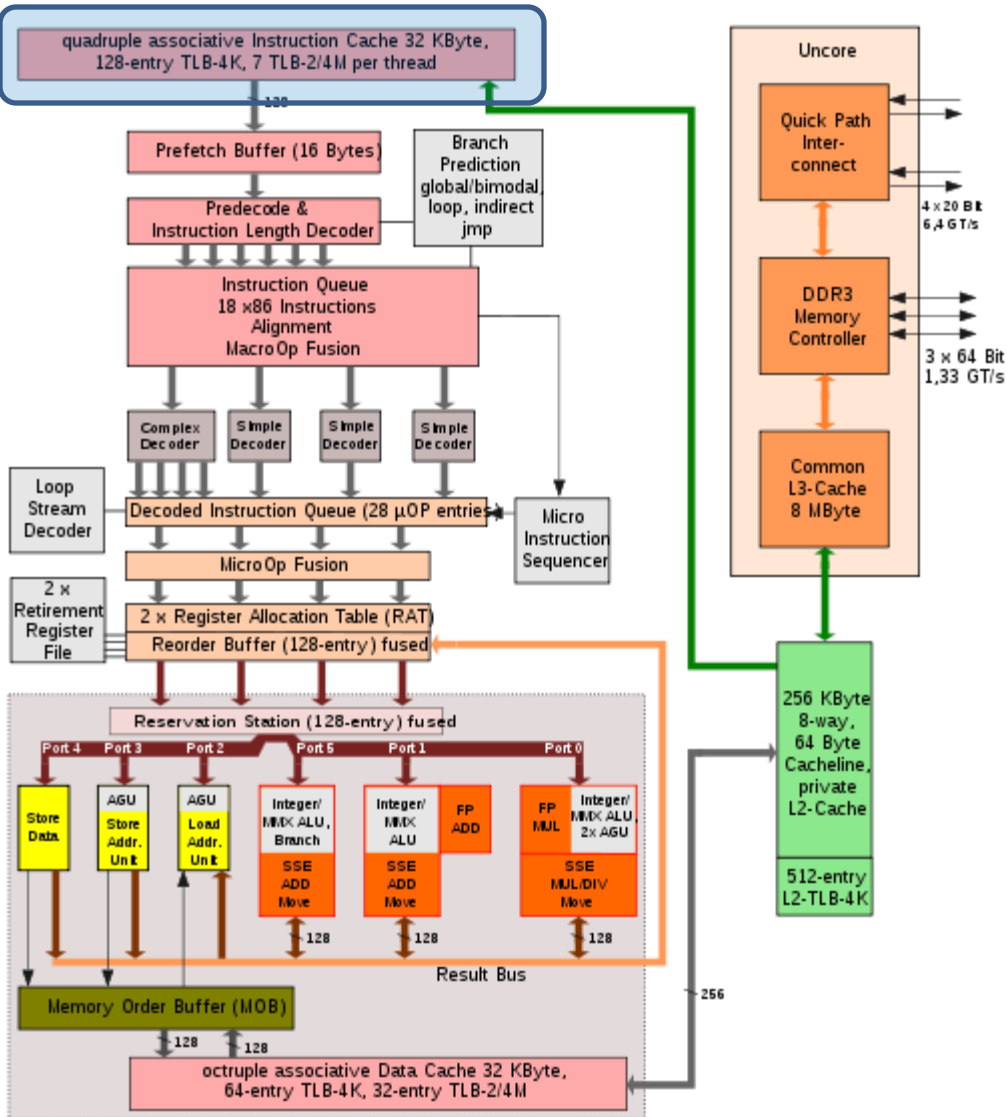
1. 4-way set associative instruction cache
2. 8-way set associative L1 data cache (32 KB)
3. 8-way set associative L2 data cache (256 KB)
4. 16-way shared L3 cache (8 MB)
5. 3 DDR3 memory connections

Components



1. Instructions **fetched** from L2 cache
2. Instructions **decoded, prefetched and queued**
3. Instructions **optimized and combined**
4. Instructions **executed**
5. Results **written**

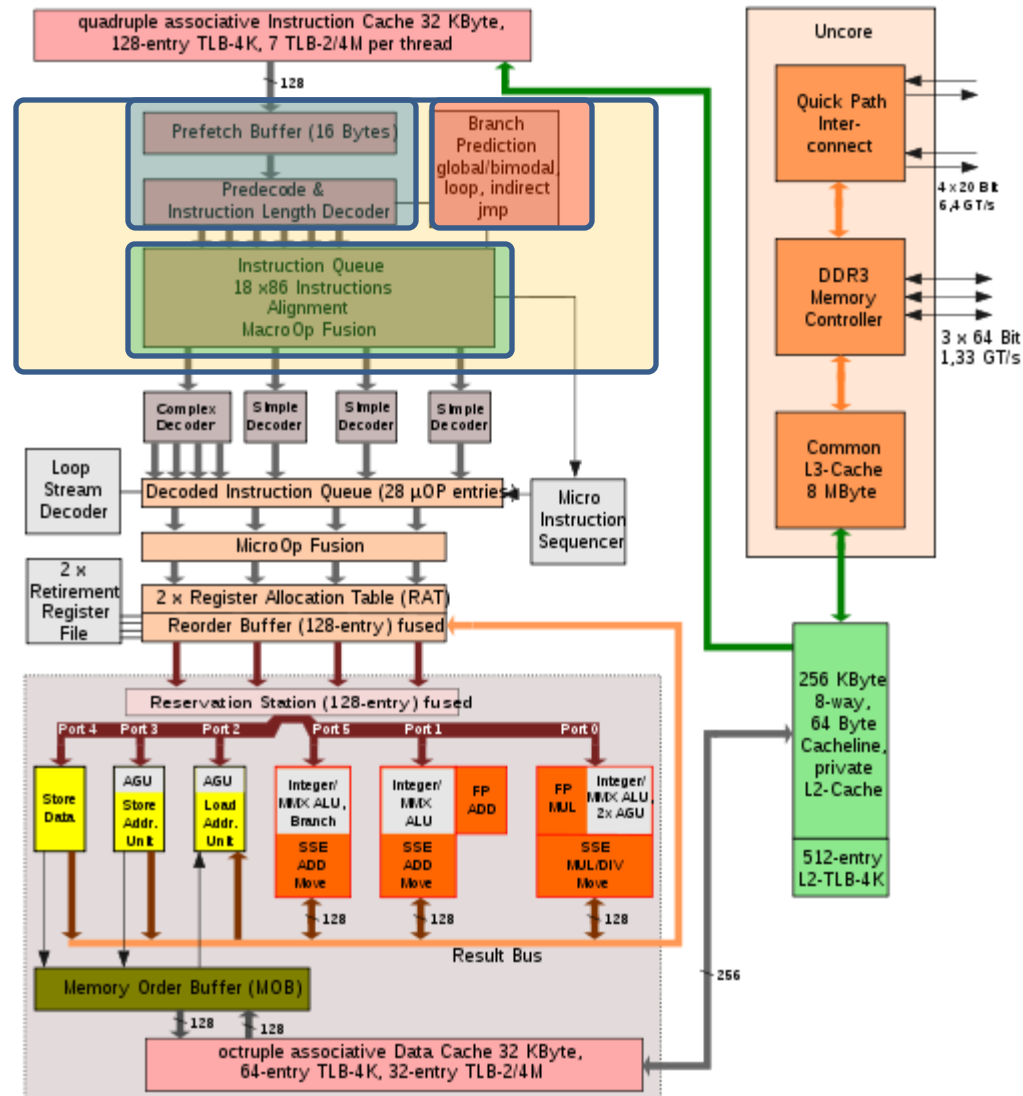
Components: Fetch



1. Instructions fetched from L2 cache

- 32 KB instruction cache
- 2-level TLB
 - L1
 - Instructions: 7-128 entries
 - Data: 32-64 entries
 - L2
 - 512 data or instruction entries
- Shared between SMT threads

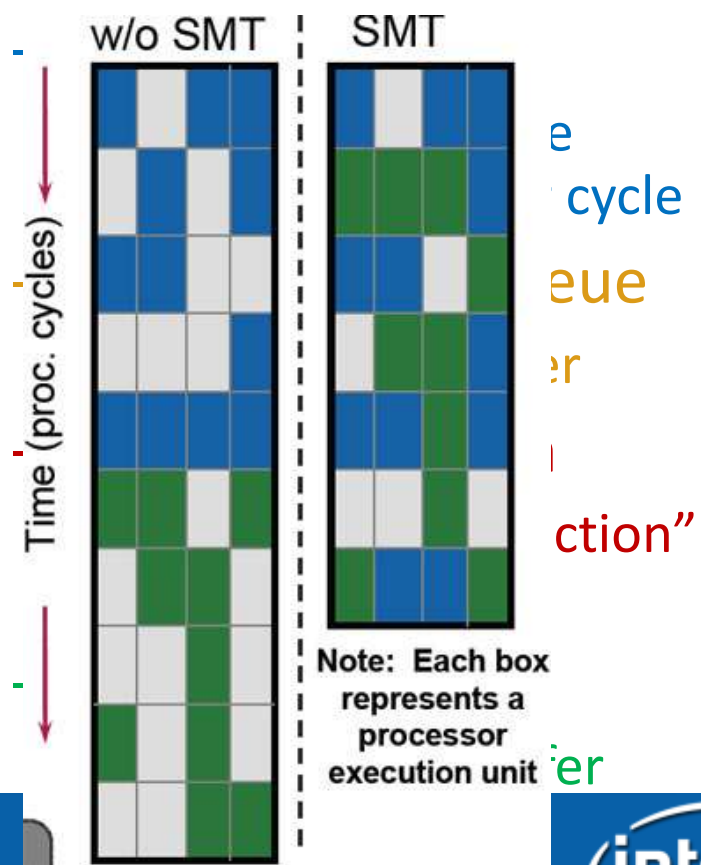
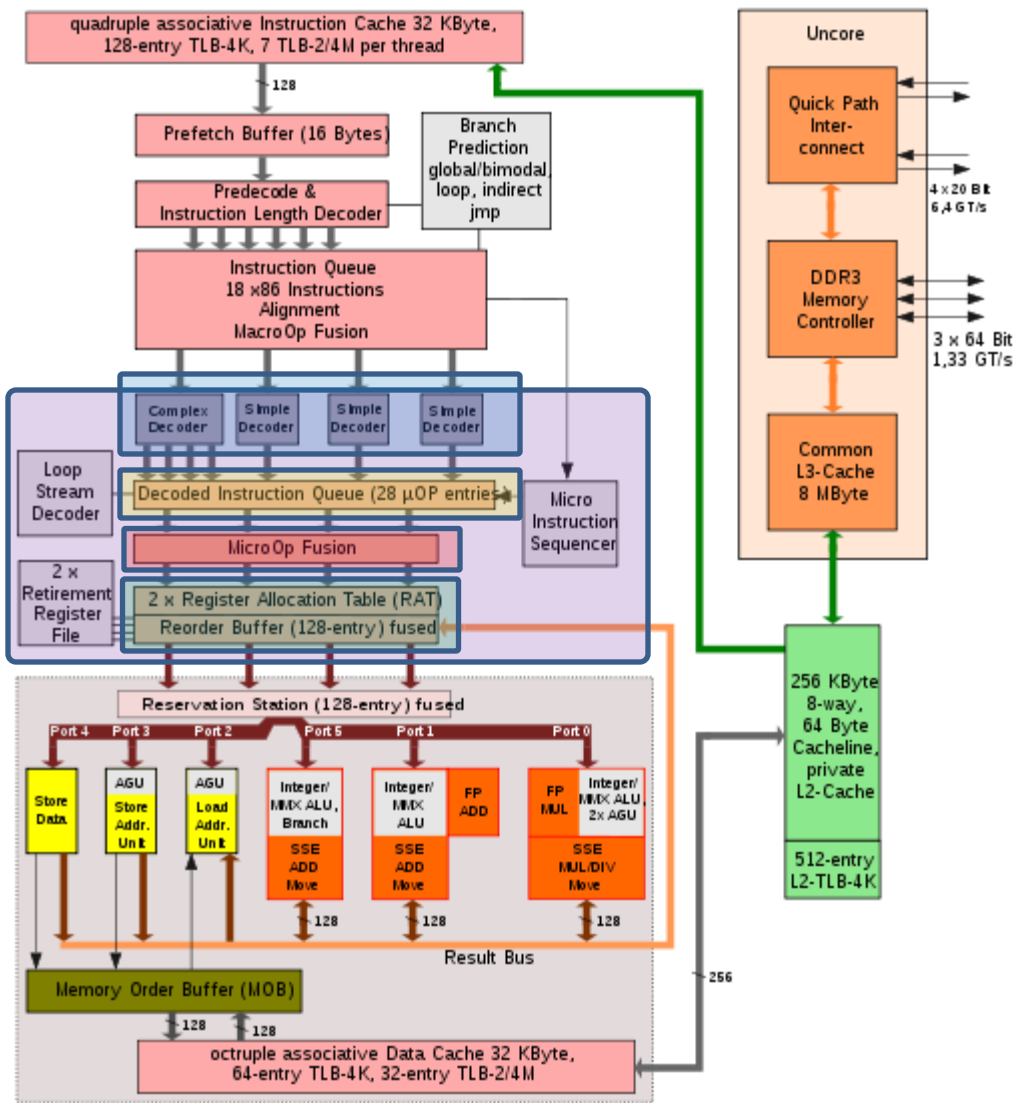
Components: Decode



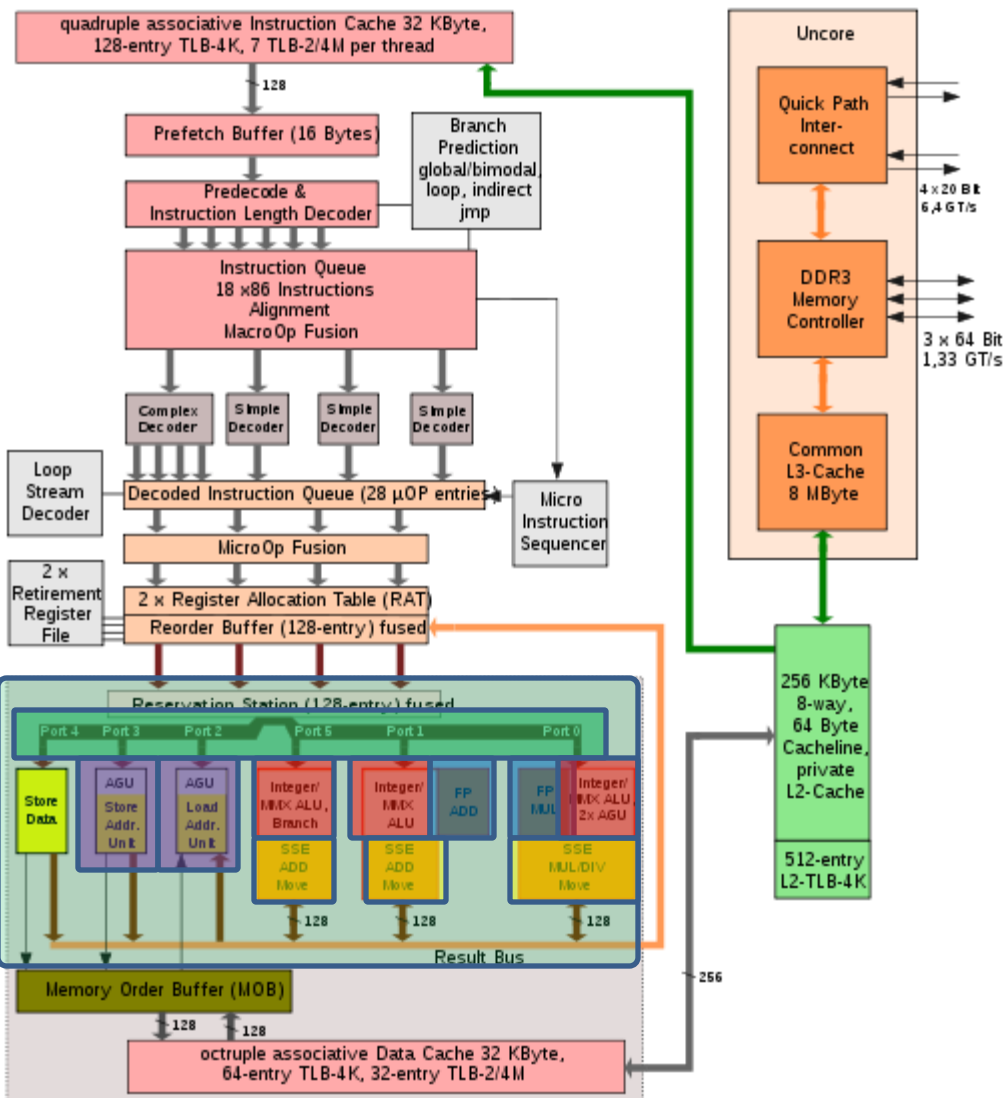
- ## 2. Instructions decoded, prefetched and queued
- 16 byte prefetch buffer
 - 18-op instruction queue
 - MacroOp fusion
 - Combine small instructions into larger ones
 - Enhanced branch prediction

Components: Optimization

3. Instructions optimized and combined



Components: Execution



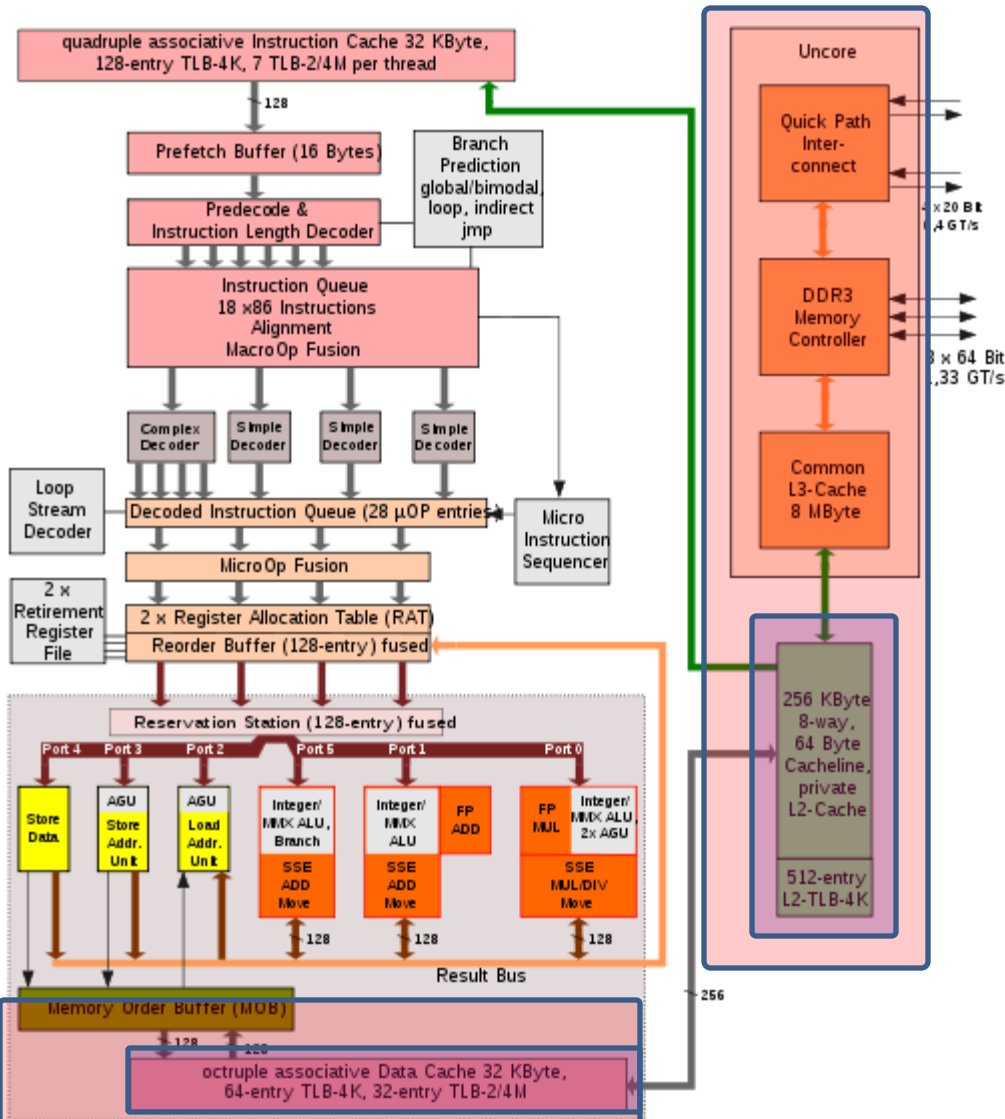
4. Instructions executed

- 4 FPUs
 - MUL, DIV, STOR, LD
- 3 ALUs
- 2 AGUs
 - Address generation
- 3 SSE Units
 - Supports SSE4
- 6 ports connecting the units

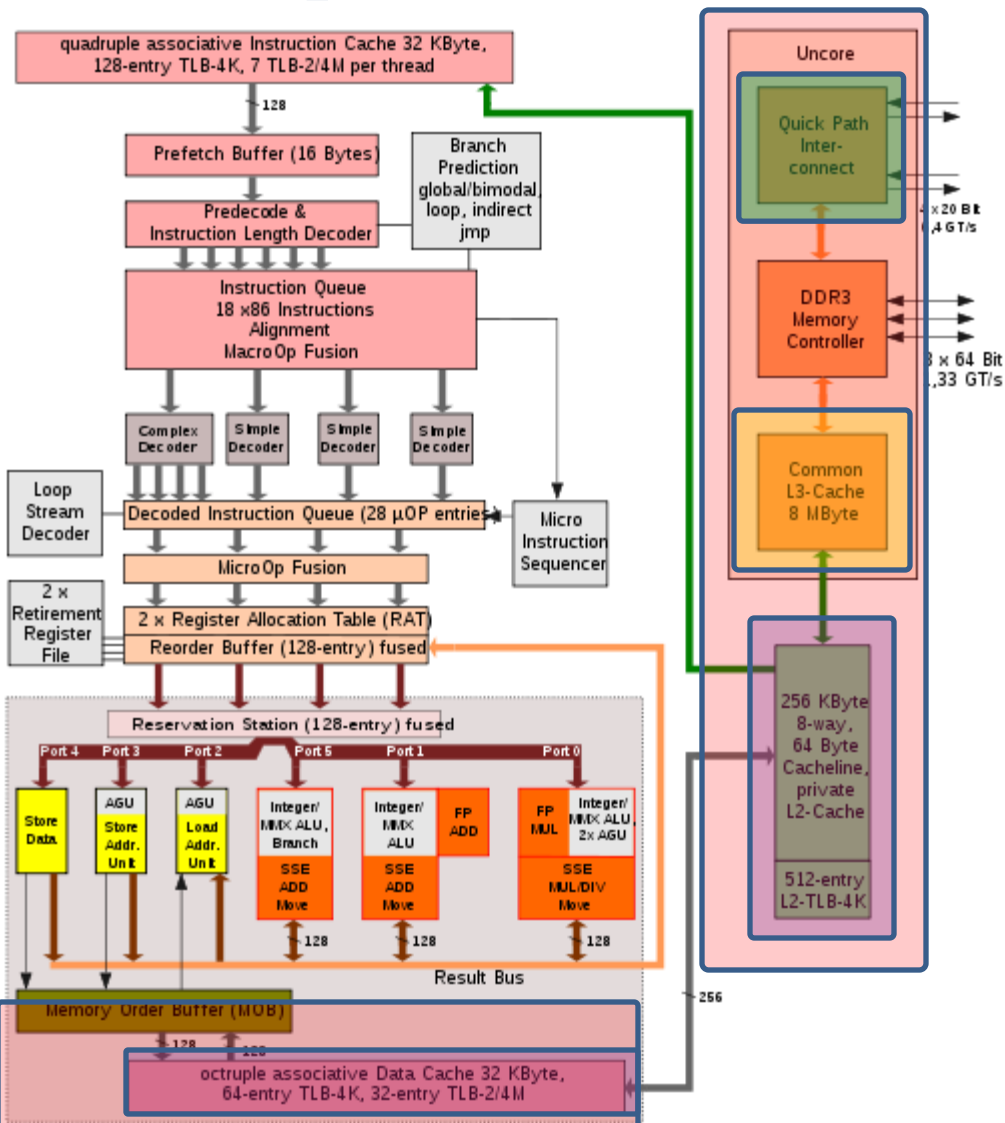
Components: Write-Back

5. Results written

– Private L1/L2 cache



Components: Write-Back



5. Results written

- Private L1/L2 cache
- Shared L3 cache
- QuickPath
 - Dedicated channel to another CPU, chip, or device
 - Replaces FSB

End