CS 152 Computer Architecture and Engineering

Lecture 23: Putting it all together: Intel Nehalem

Krste Asanovic
Electrical Engineering and Computer Sciences
University of California, Berkeley

http://www.eecs.berkeley.edu/~krste
http://inst.cs.berkeley.edu/~cs152

Intel Nehalem

- Review entire semester by looking at most recent microprocessor from Intel
- Nehalem is code name for microarchitecture at heart of Core i7 and Xeon 5500 series server chips
- First released at end of 2008

- Figures/Info from Intel, David Kanter at Real World Technologies.

Nehalem System Example: Apple Mac Pro Desktop 2009

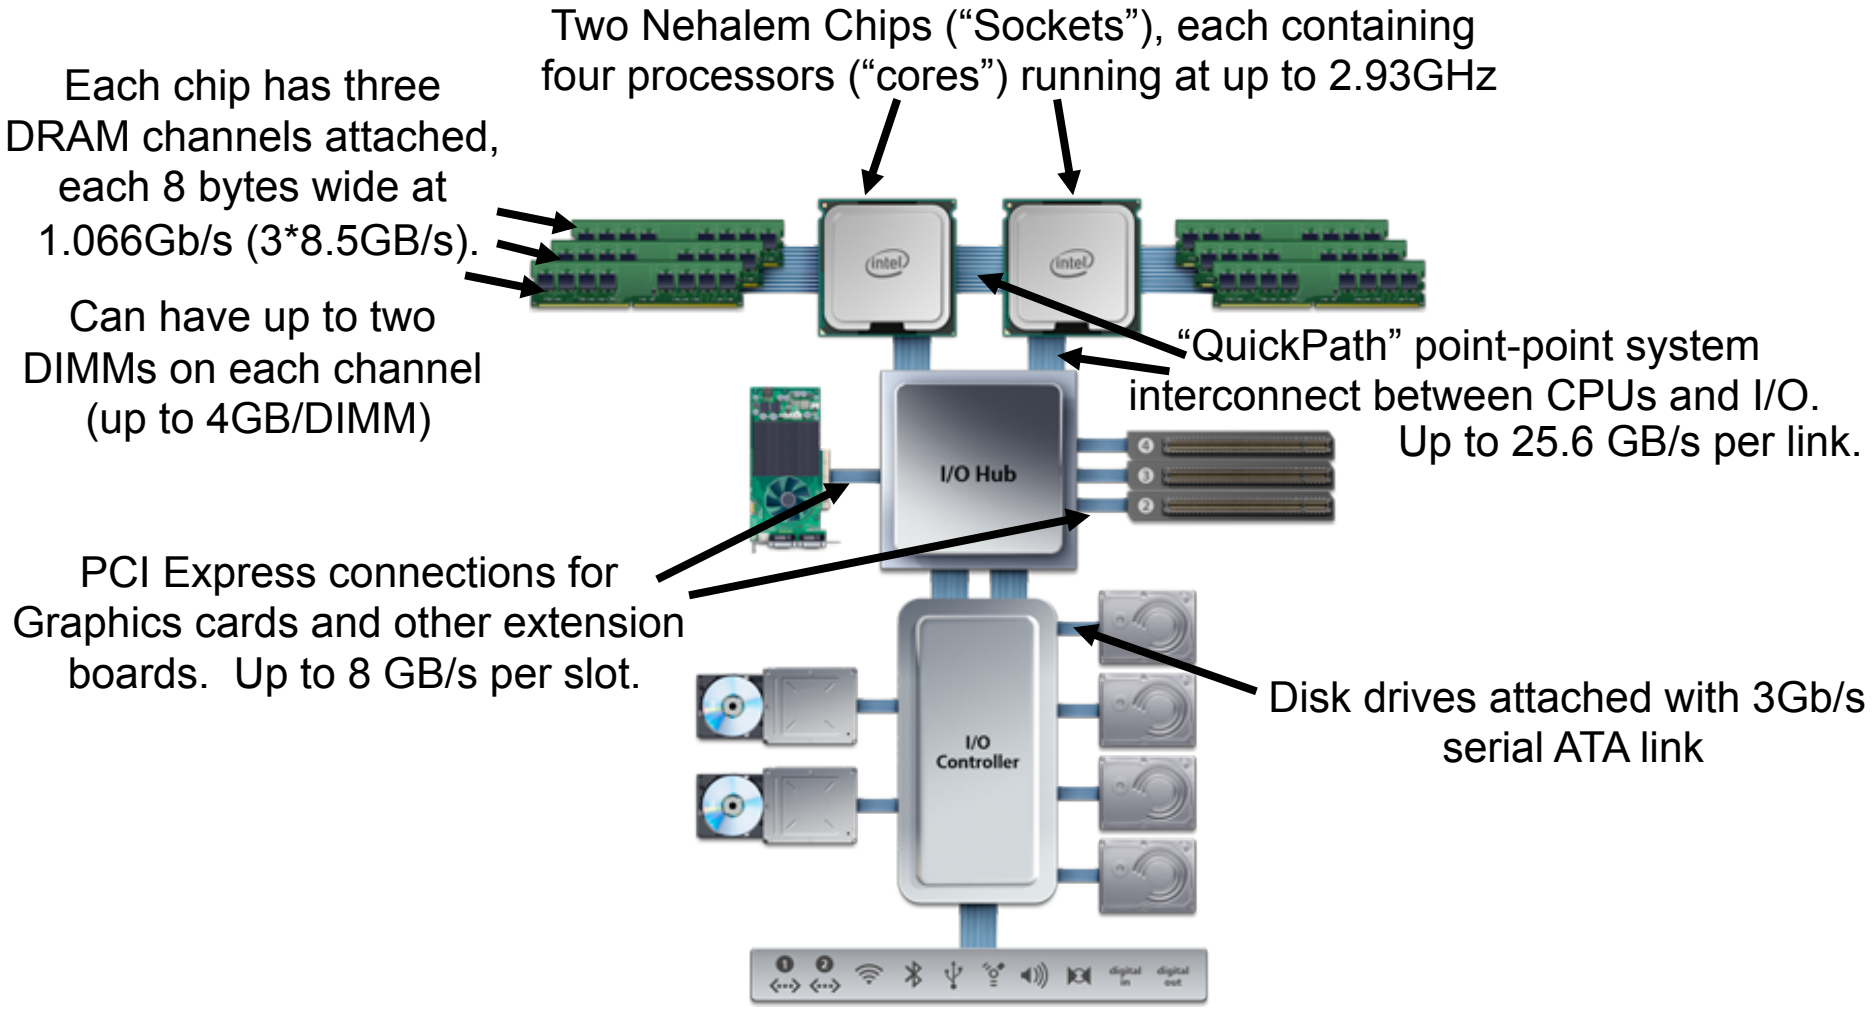

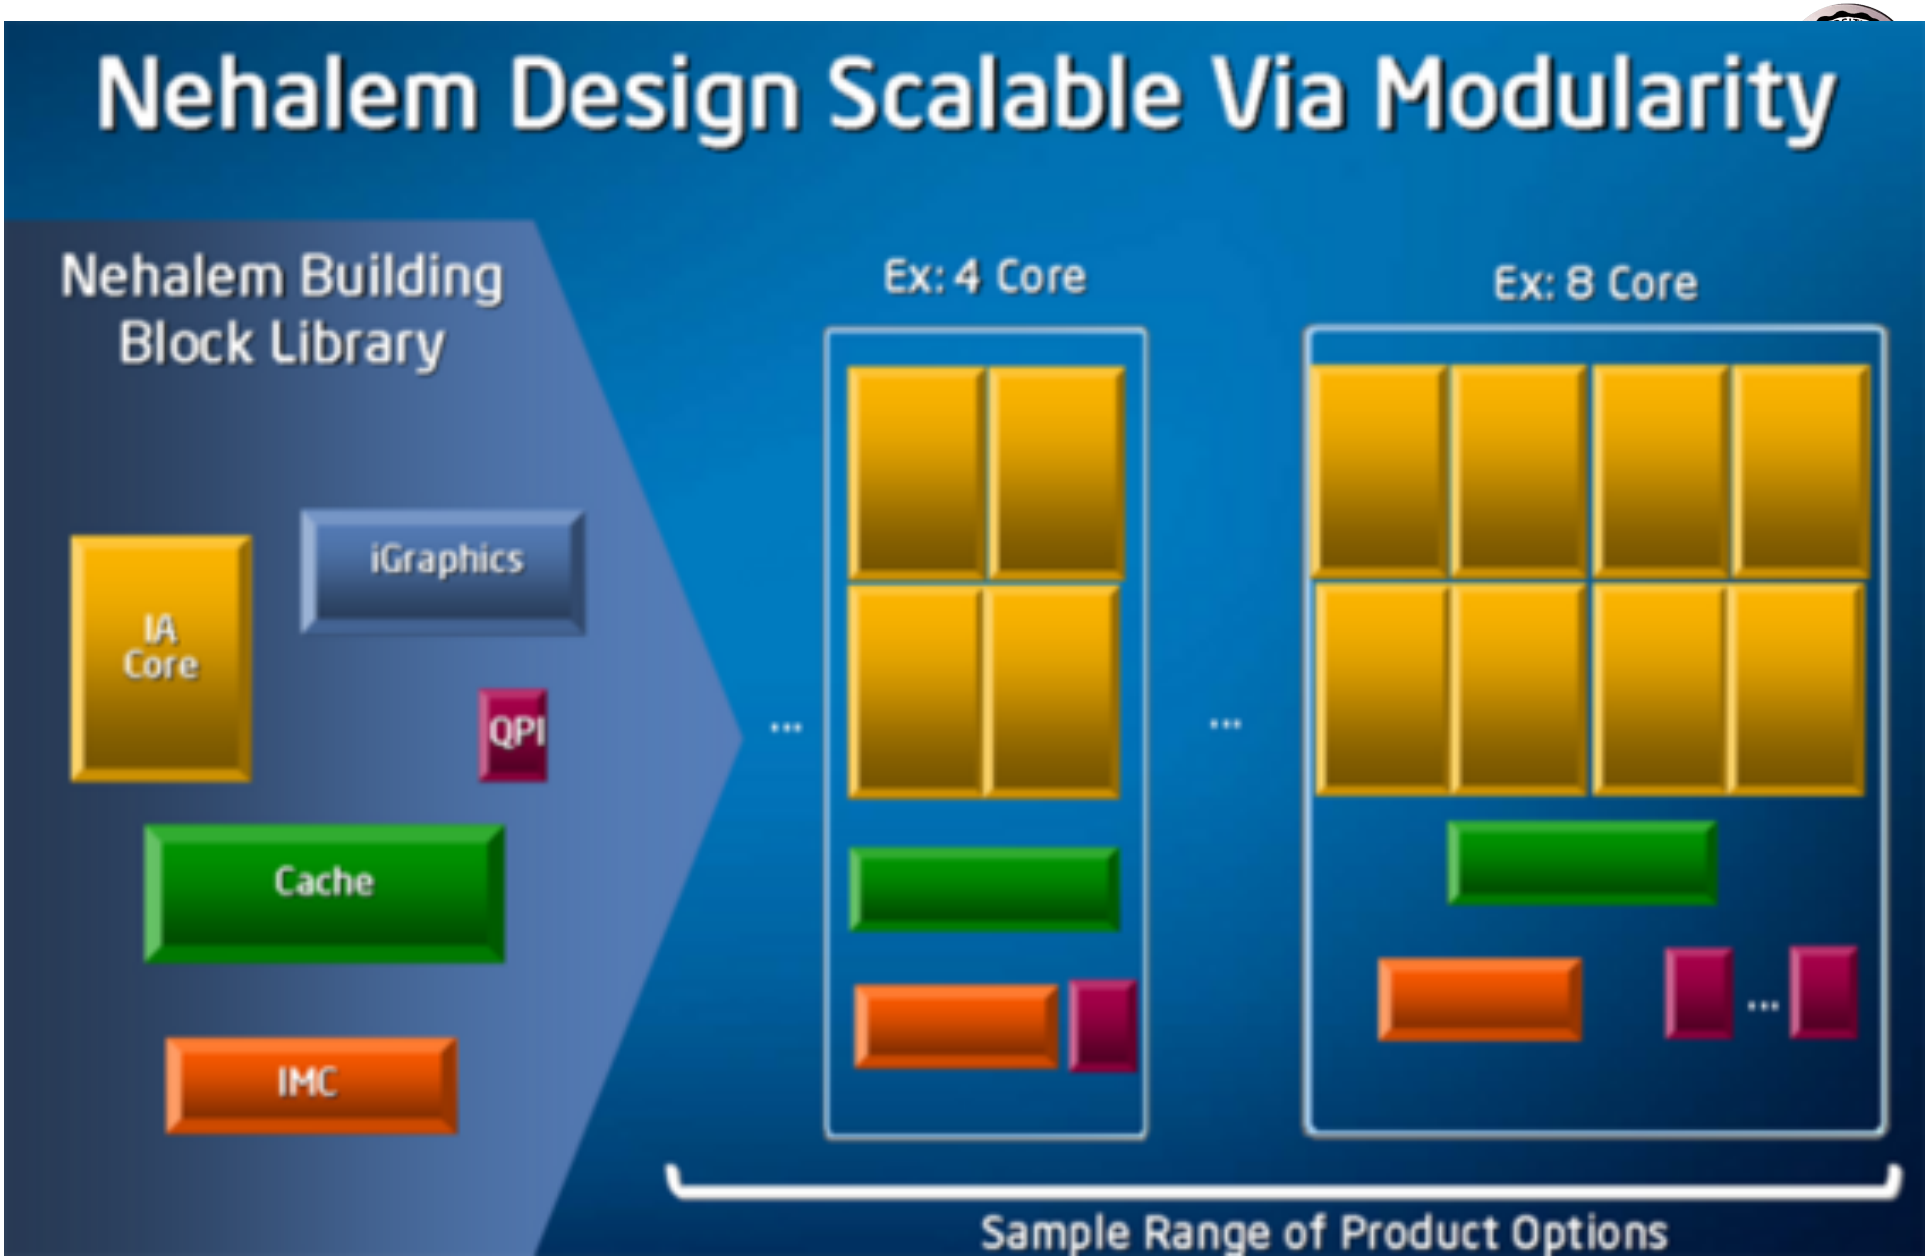
Nehalem Design Scalable Via Modularity
Nehalem Building Block Library
IA Core
iGraphics
QPI
Cache
IMC
Ex: 4 Core
Ex: 8 Core
Sample Range of Product Options

Nehalem Die Photo

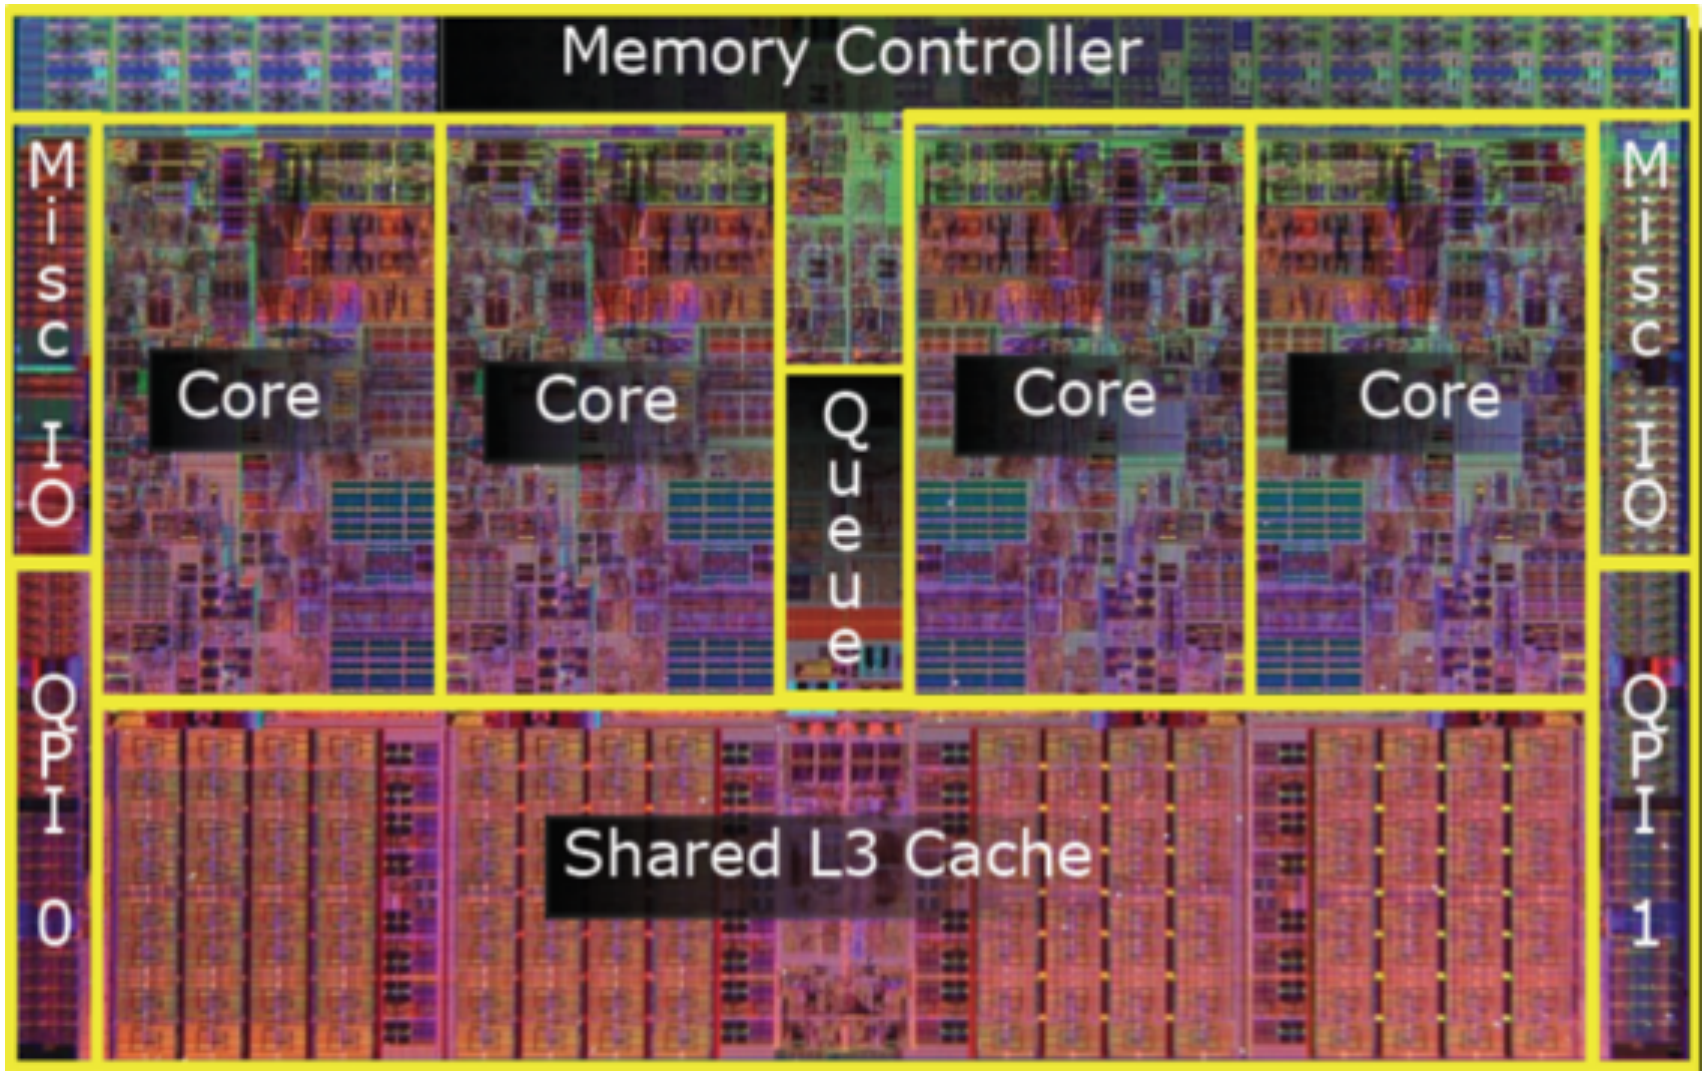

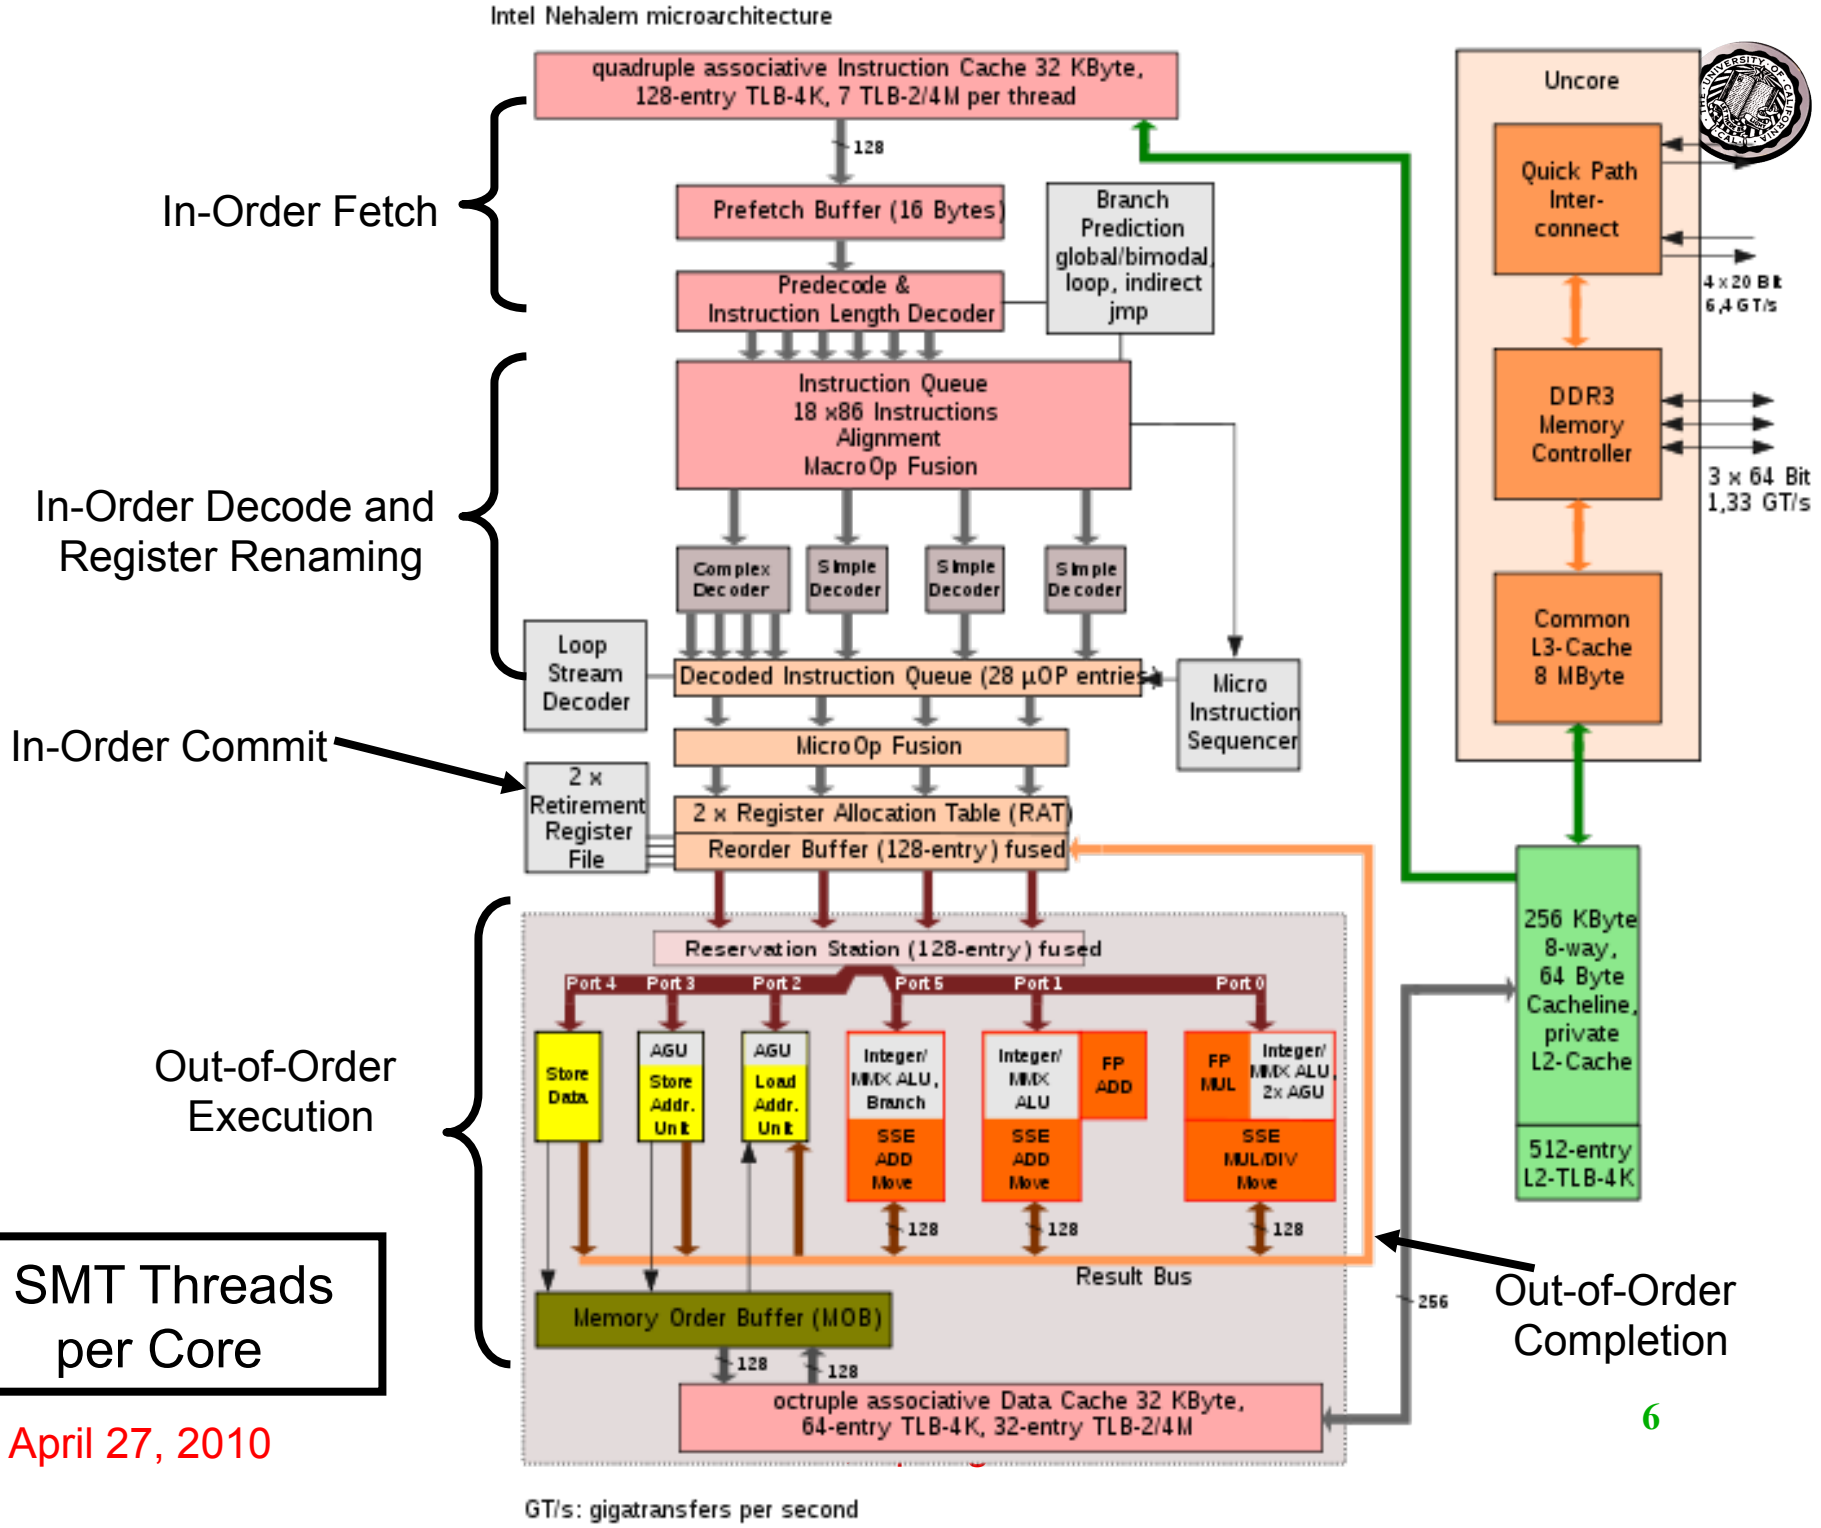
Intel Nehalem microarchitecture
quadruple associative Instruction Cache 32 KByte, 128-entry TLB-4K, 7 TLB-2/4M per thread
128
Prefetch Buffer (16 Bytes)
Branch Prediction global/bimodal, loop, indirect jmp
Predecode & Instruction Length Decoder
Instruction Queue 18 x86 Instructions Alignment Macro Op Fusion
Complex Decoder
Simple Decoder
Simple Decoder
Simple Decoder
Loop Stream Decoder
Decoded Instruction Queue (28 µOP entries)
Micro Instruction Sequencer
Micro Op Fusion
2 x Retirement Register File
2 x Register Allocation Table (RAT)
Reorder Buffer (128-entry) fused
Reservation Station (128-entry) fused
Port 4
Port 3
Port 2
Port 5
Port 1
Port 0
Store Data
AGU Store Addr. Unit
AGU Load Addr. Unit
Integer/ MMX ALU, Branch
SSE ADD Move
Integer/ MMX ALU
FP ADD
SSE ADD Move
FP MUL
Integer/ MMX ALU, 2x AGU
SSE MUL/DIV Move
128
Result Bus
Memory Order Buffer (MOB)
octruple associative Data Cache 32 KByte, 64-entry TLB-4K, 32-entry TLB-2/4M
256
Uncore
Quick Path Inter-connect
4 x 20 Bit 6,4 GT/s
DDR3 Memory Controller
3 x 64 Bit 1,33 GT/s
Common L3-Cache 8 MByte
256 KByte 8-way, 64 Byte Cacheline, private L2-Cache
512-entry L2-TLB-4K
In-Order Fetch
In-Order Decode and Register Renaming
In-Order Commit
Out-of-Order Execution
Out-of-Order Completion
2 SMT Threads per Core
April 27, 2010
GT/s: gigatransfers per second

Front-End Instruction Fetch & Decode

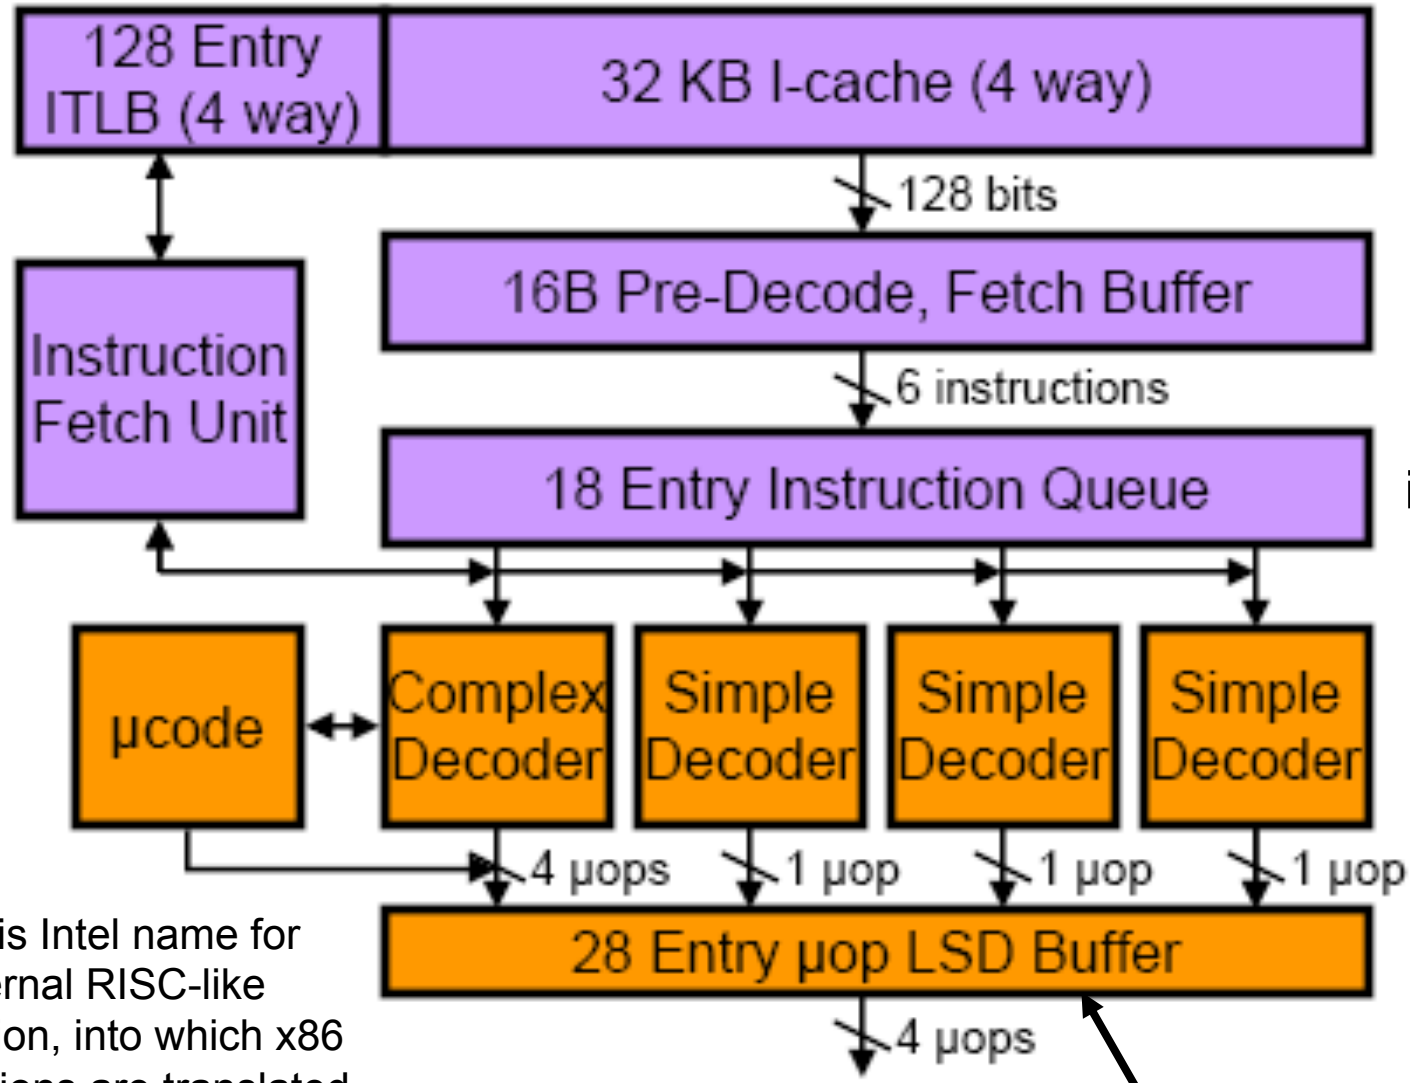

µOP is Intel name for internal RISC-like instruction, into which x86 instructions are translated

Loop Stream Detector (can run short loops out of the buffer)

Branch Prediction

- Part of instruction fetch unit

- Several different types of branch predictor
  - Details not public
- Two-level BTB
- Loop count predictor
  - How many backwards taken branches before loop exit
  - (Also predictor for length of microcode loops, e.g., string move)
- Return Stack Buffer
  - Holds subroutine targets
  - Renames the stack buffer so that it is repaired after mispredicted returns
  - Separate return stack buffer for each SMT thread

x86 Decoding

- Translate up to 4 x86 instructions into uOPS each cycle
- Only first x86 instruction in group can be complex (maps to 1-4 uOPS), rest must be simple (map to one uOP)
- Even more complex instructions, jump into microcode engine which spits out stream of uOPS

Split x86 in small uOPs, then fuse back into bigger units

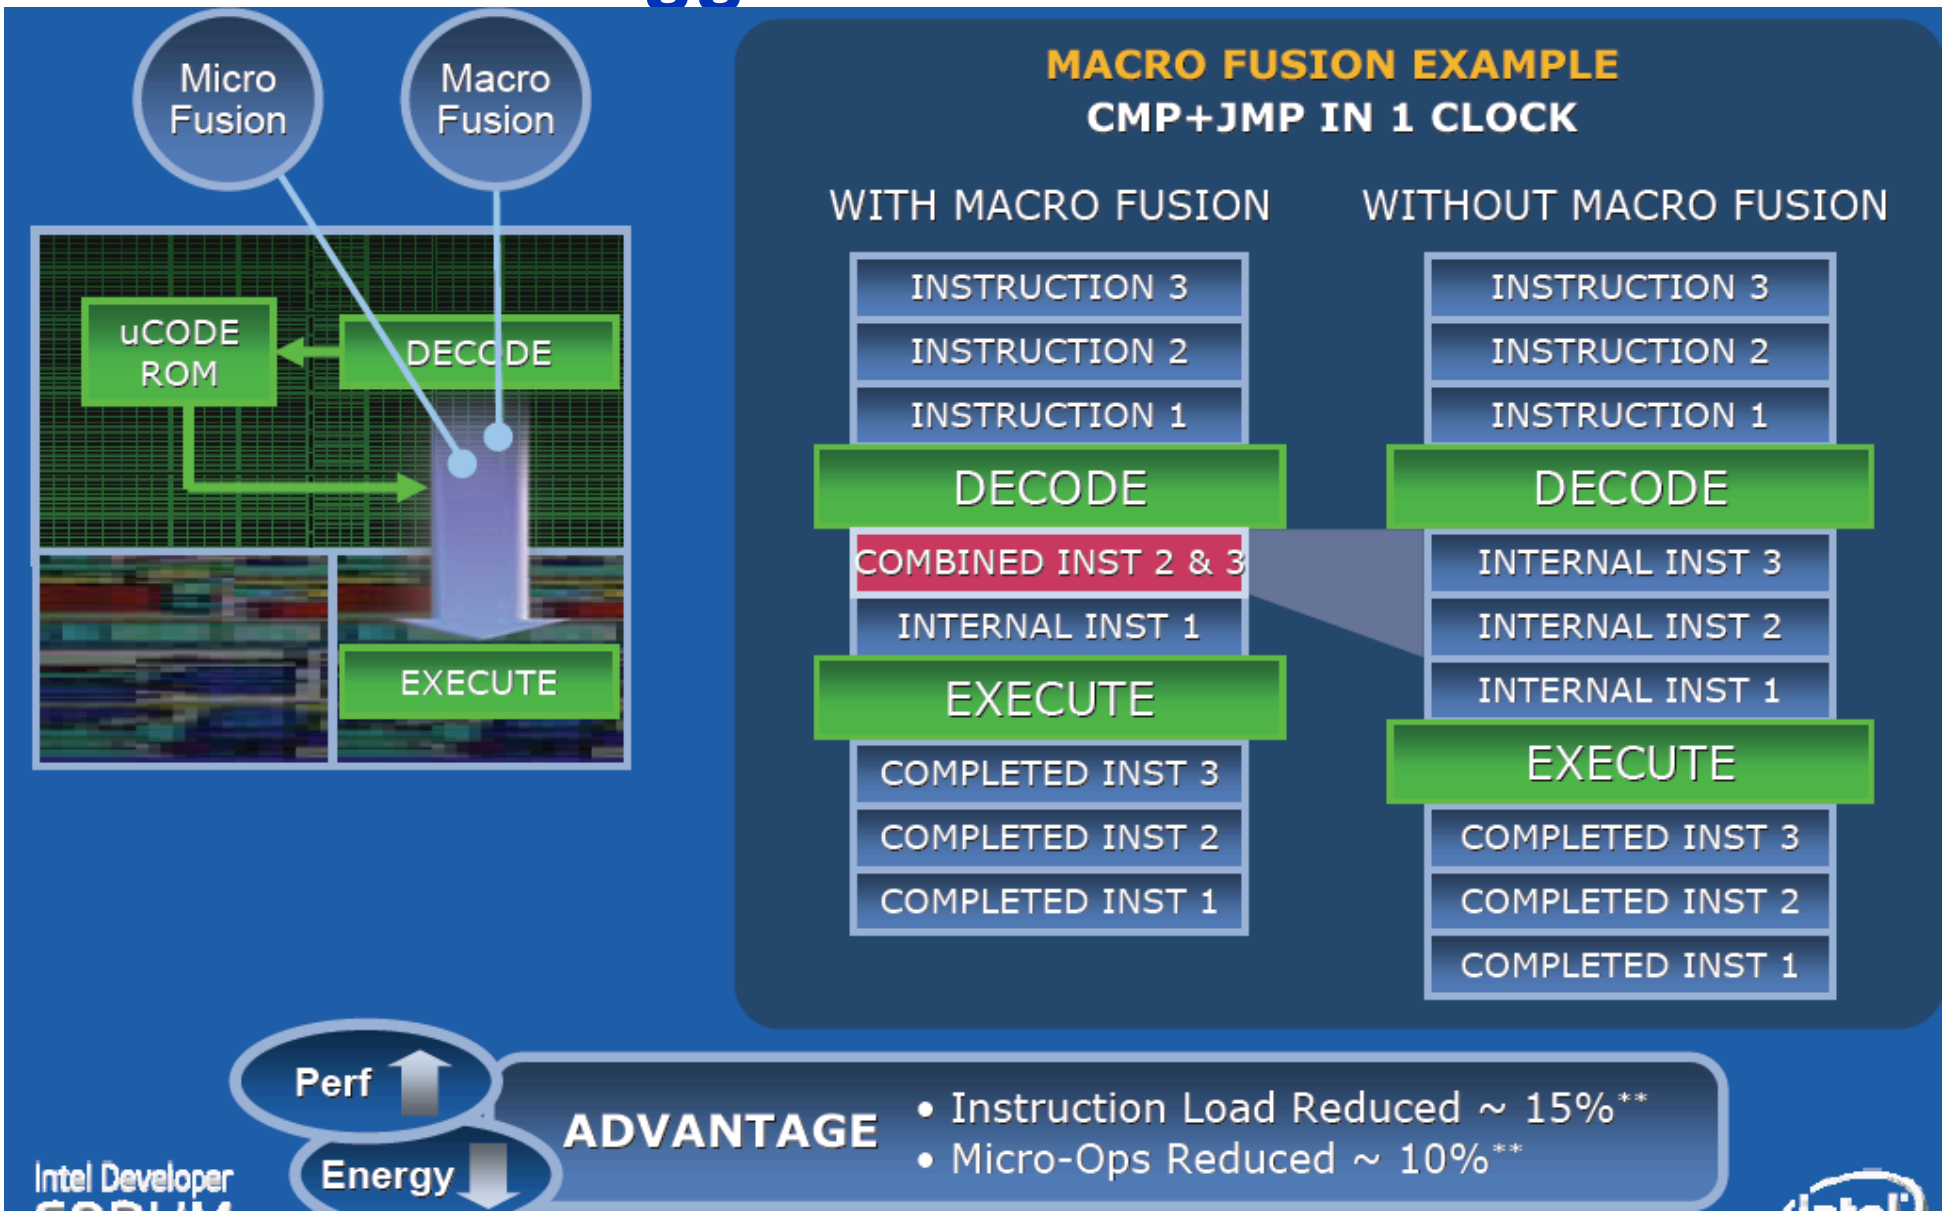

Loop Stream Detectors save Power

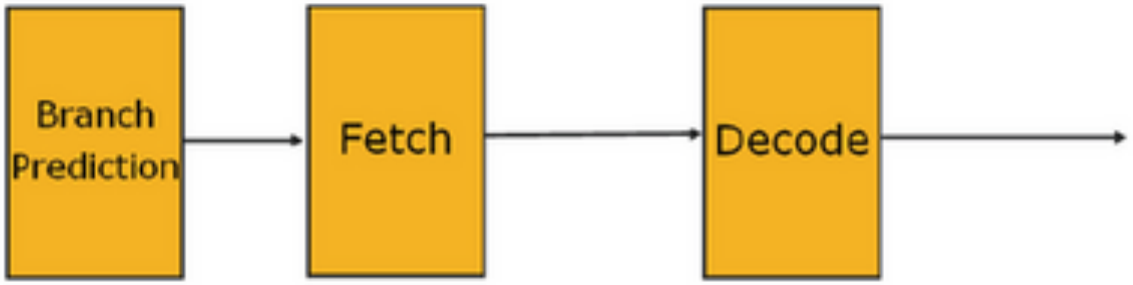

Intel® Core™2 Loop Stream Detector

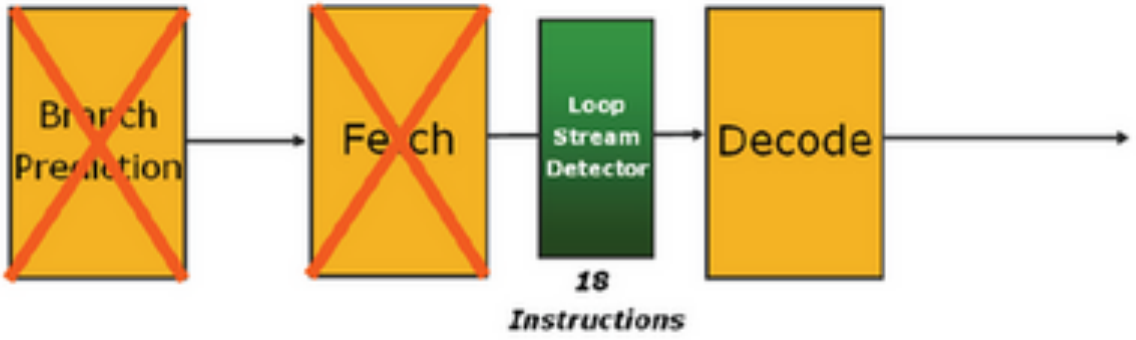

Intel Core Microarchitecture (Nehalem) Loop Stream Detector

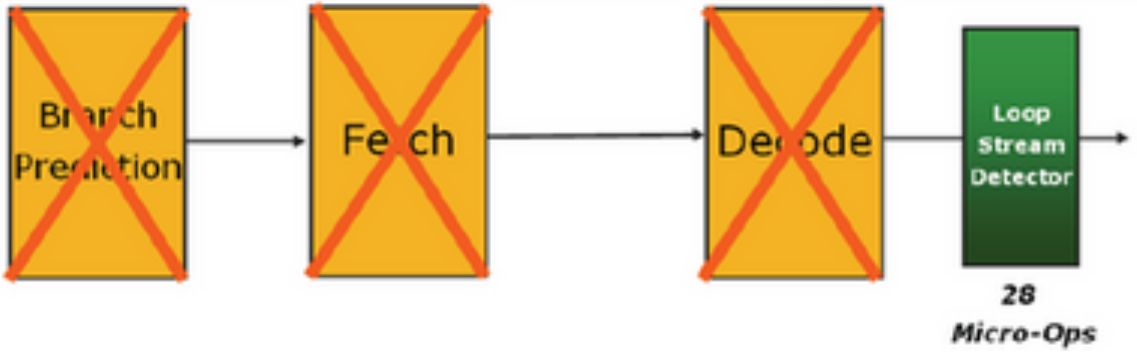

Out-of-Order Execution Engine

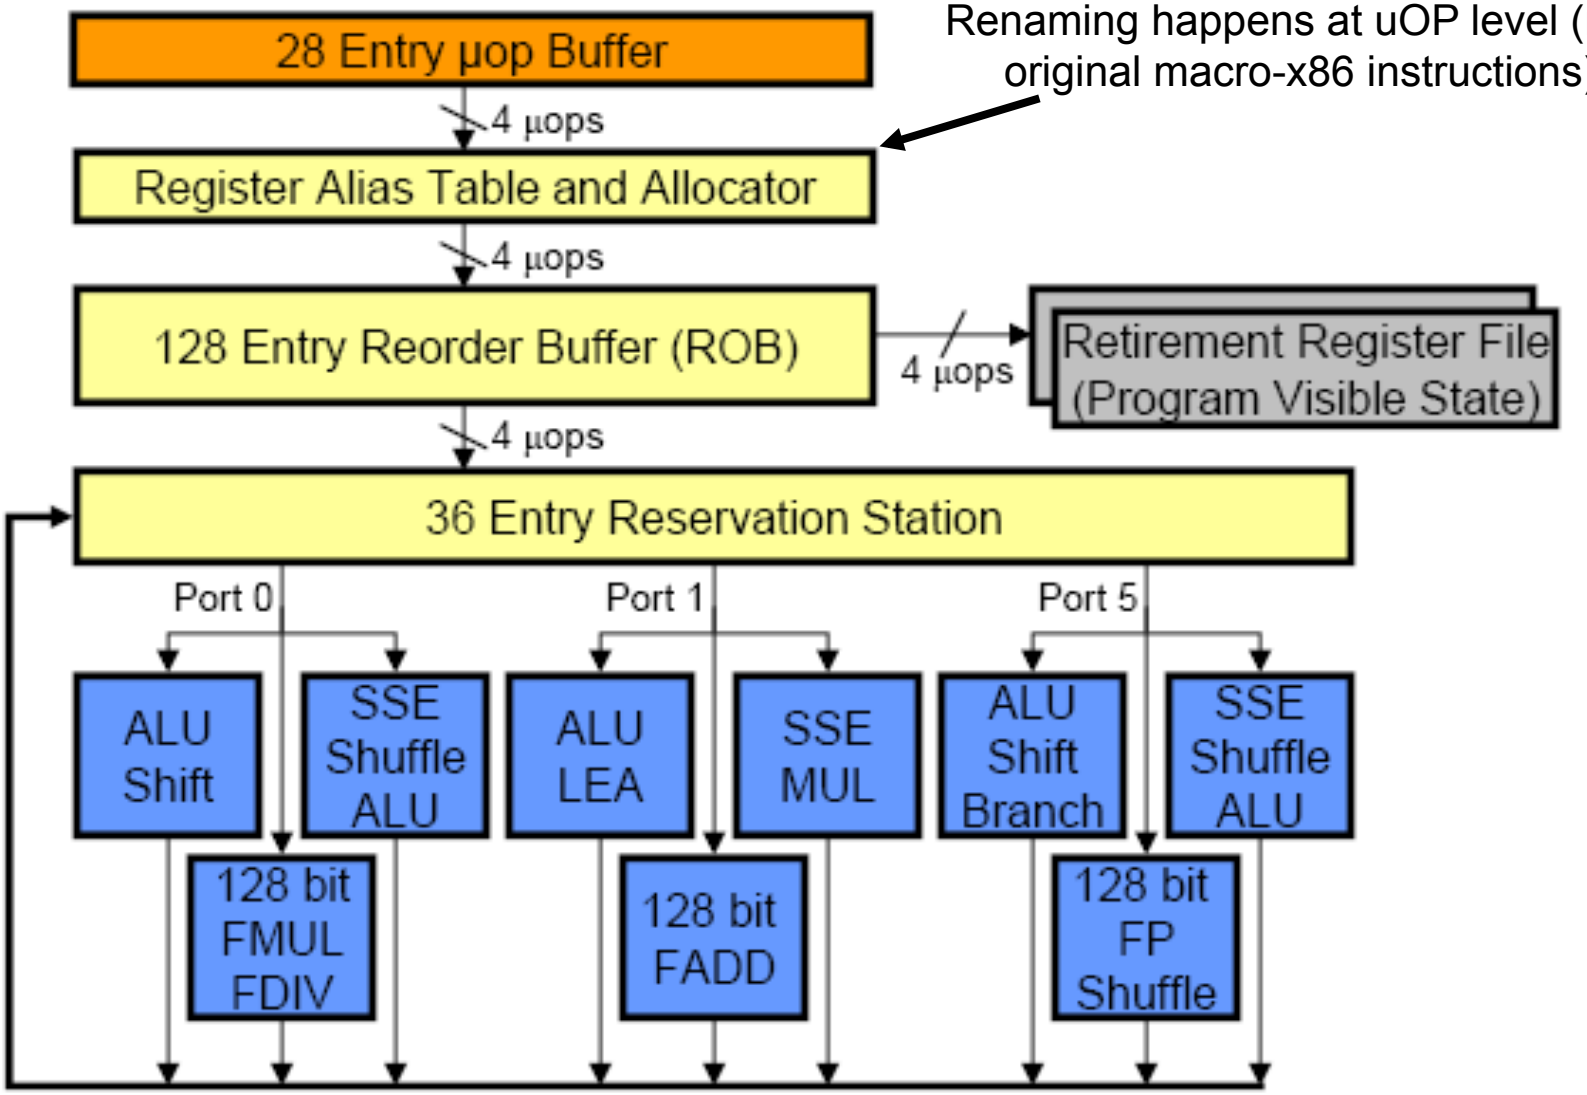

SMT effects in OoO Execution Core

- Reorder buffer (remembers program order and exception status for in-order commit) has 128 entries divided statically and equally between both SMT threads
- Reservation stations (instructions waiting for operands for execution) have 36 entries competitively shared by threads

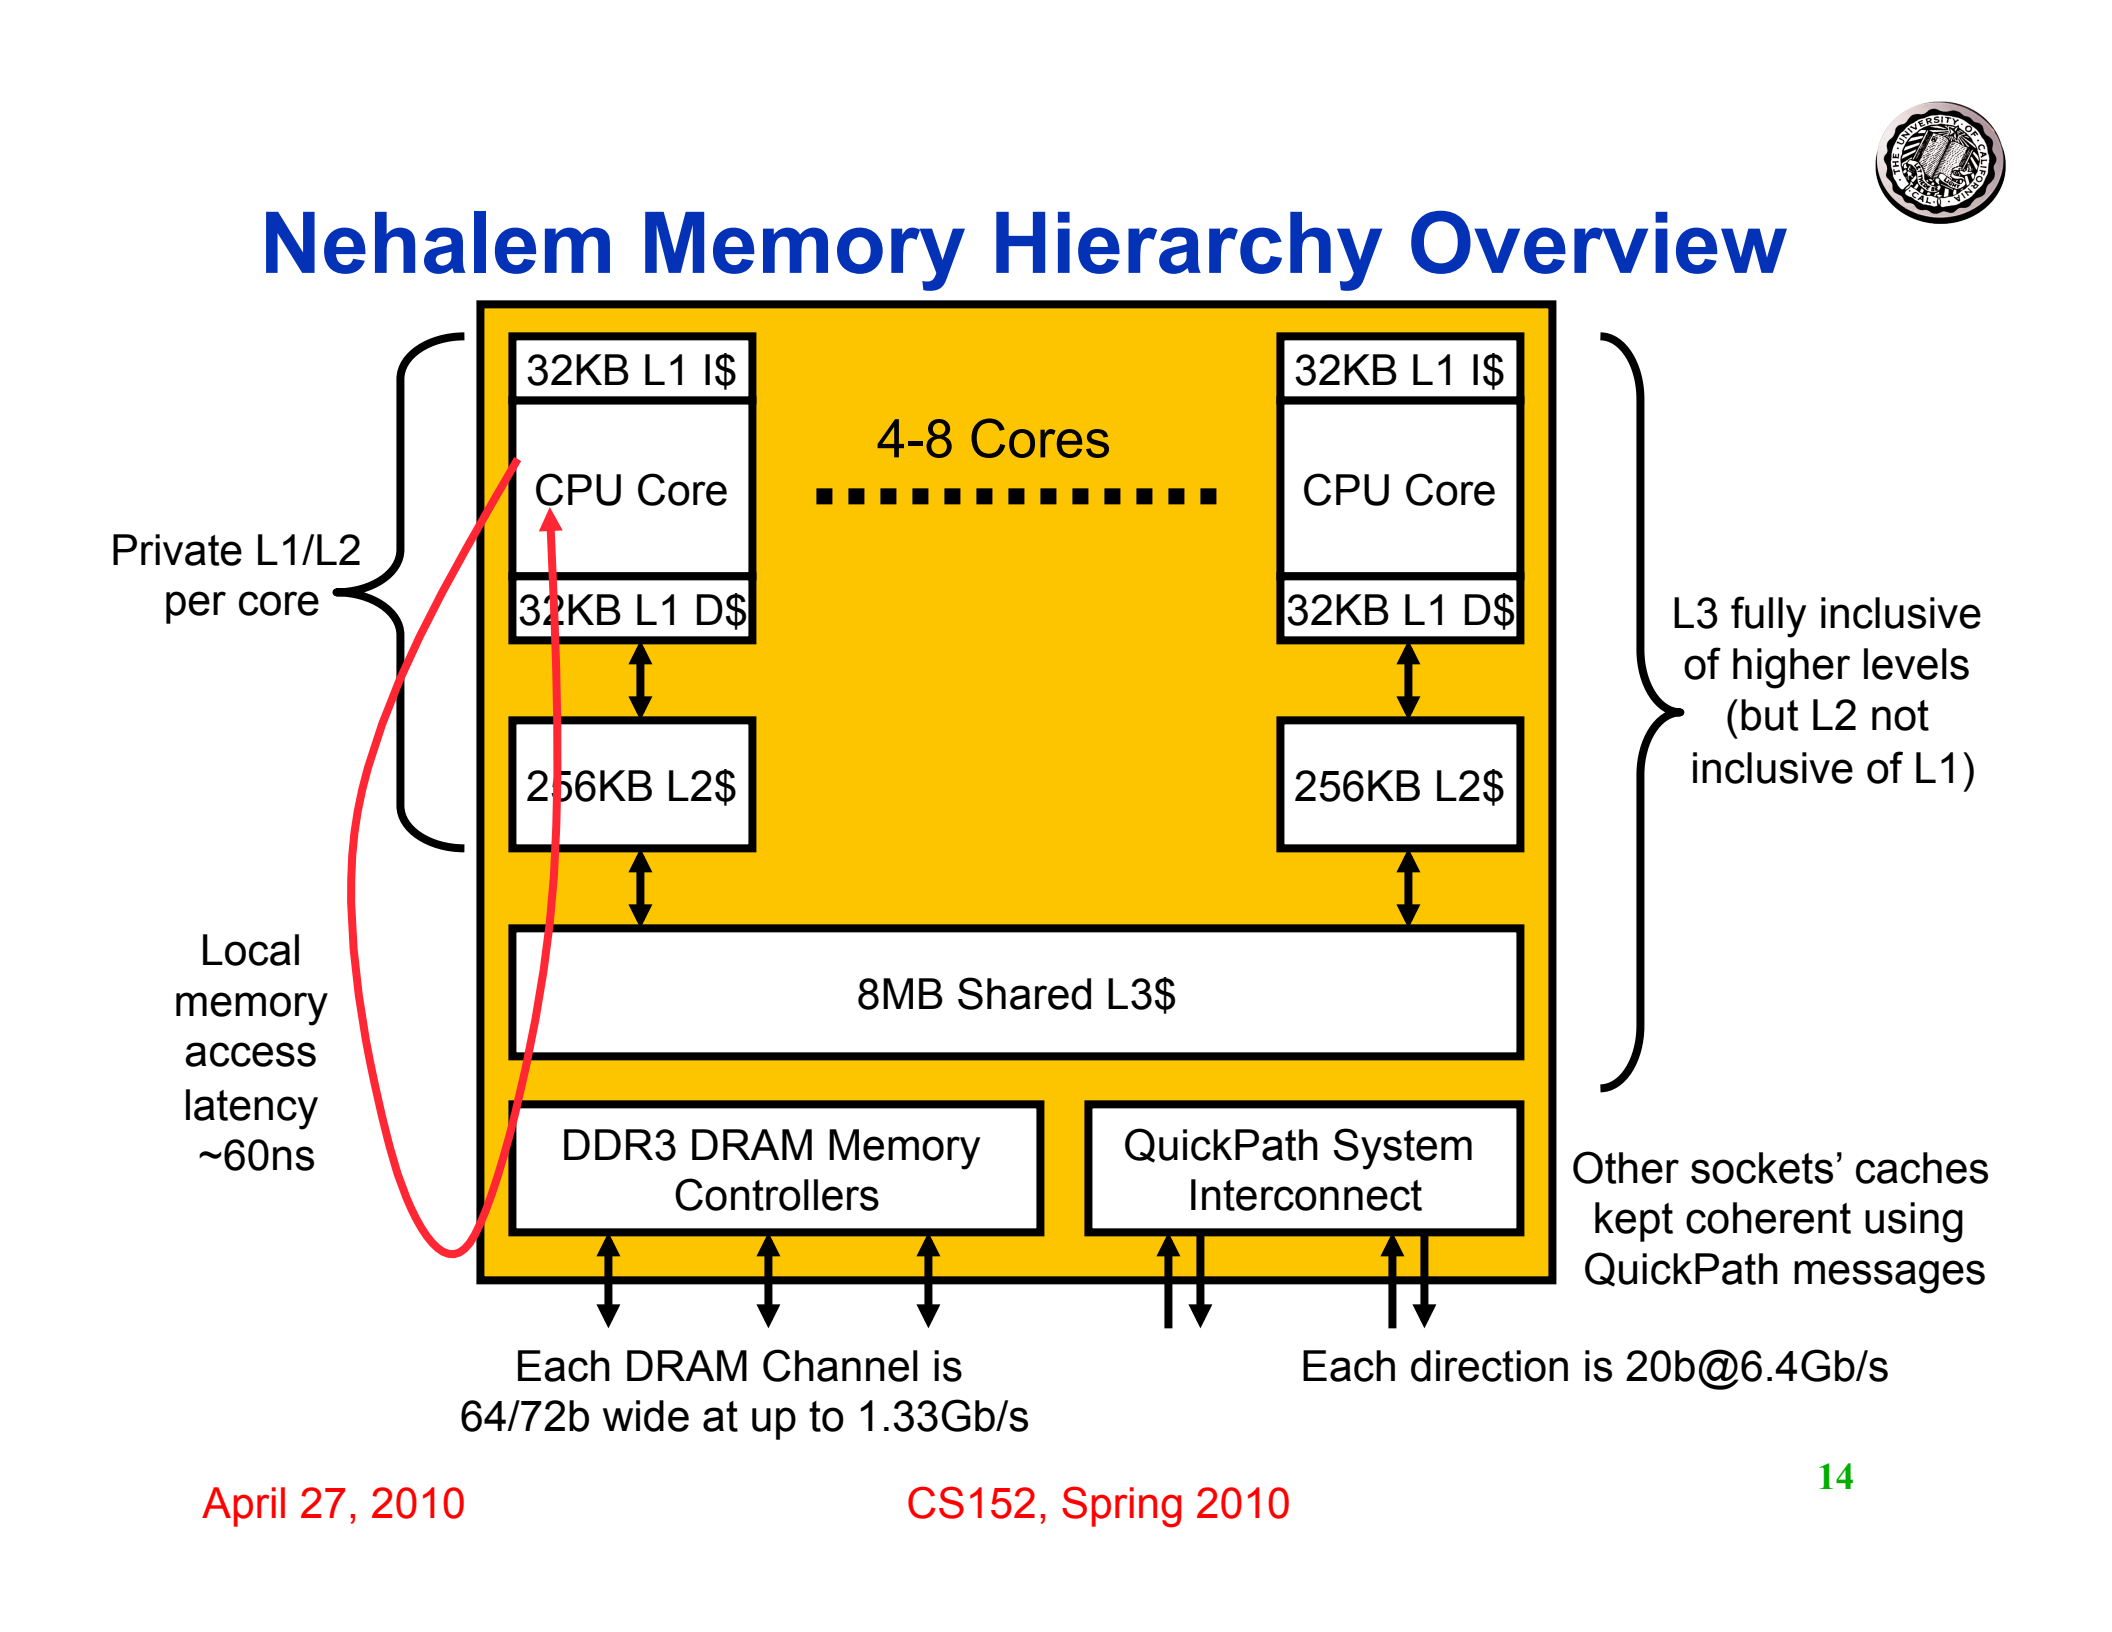

Nehalem Memory Hierarchy Overview

All Sockets can Access all Data

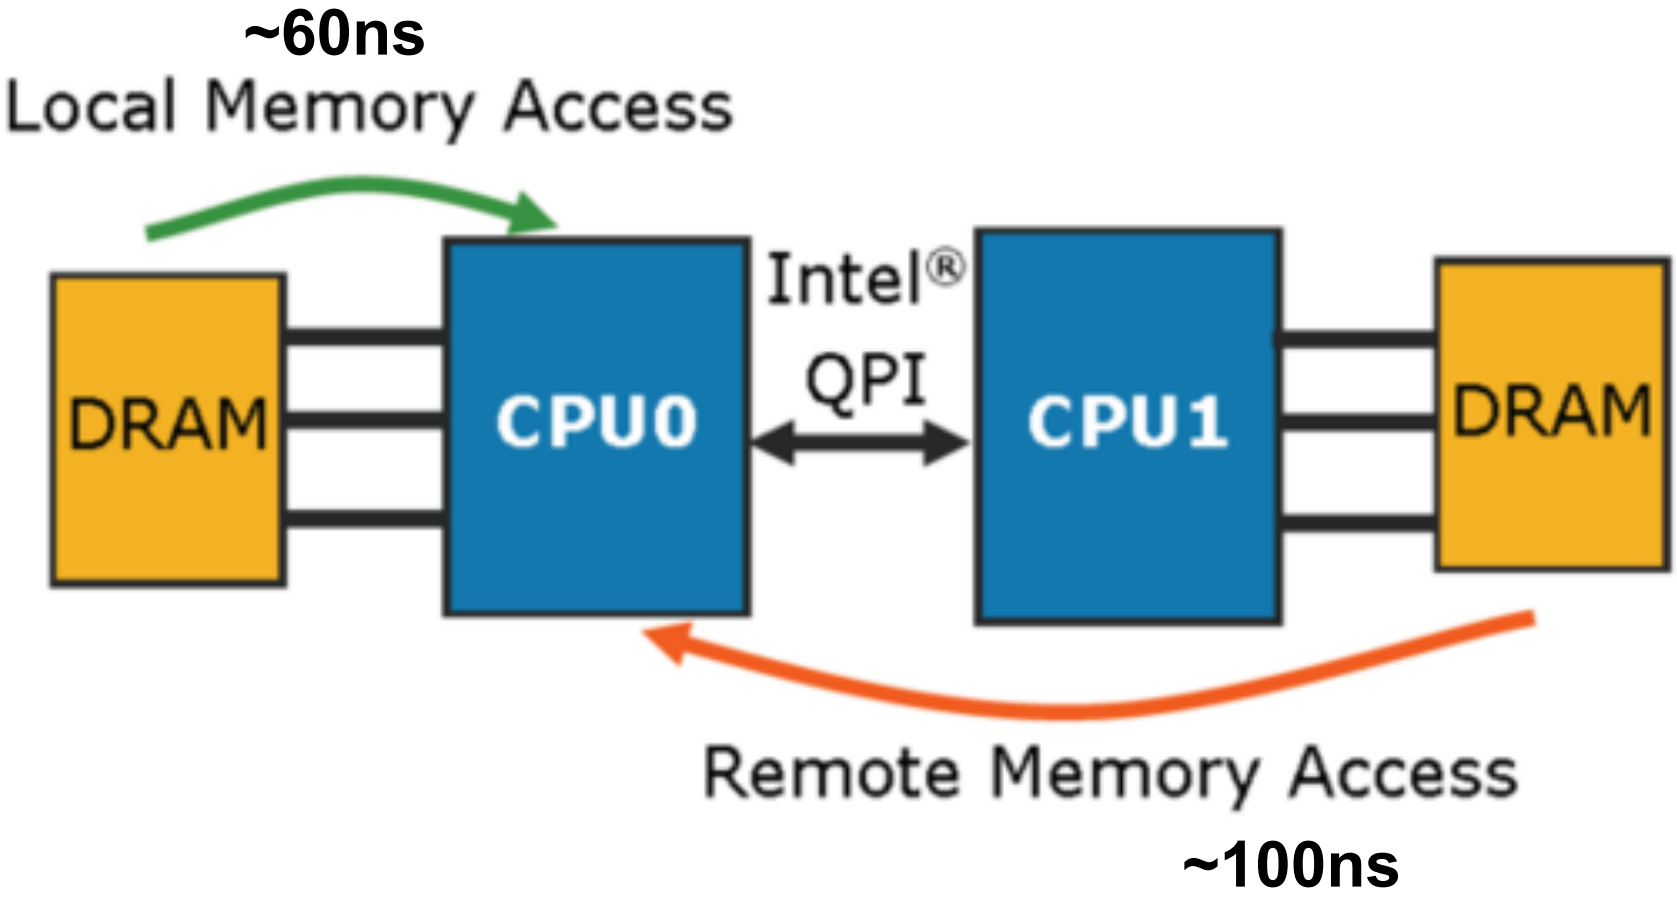

Core's Private Memory System

Load queue 48 entries

Store queue 32 entries

Divided statically between SMT threads

Up to 16 outstanding misses in flight per core

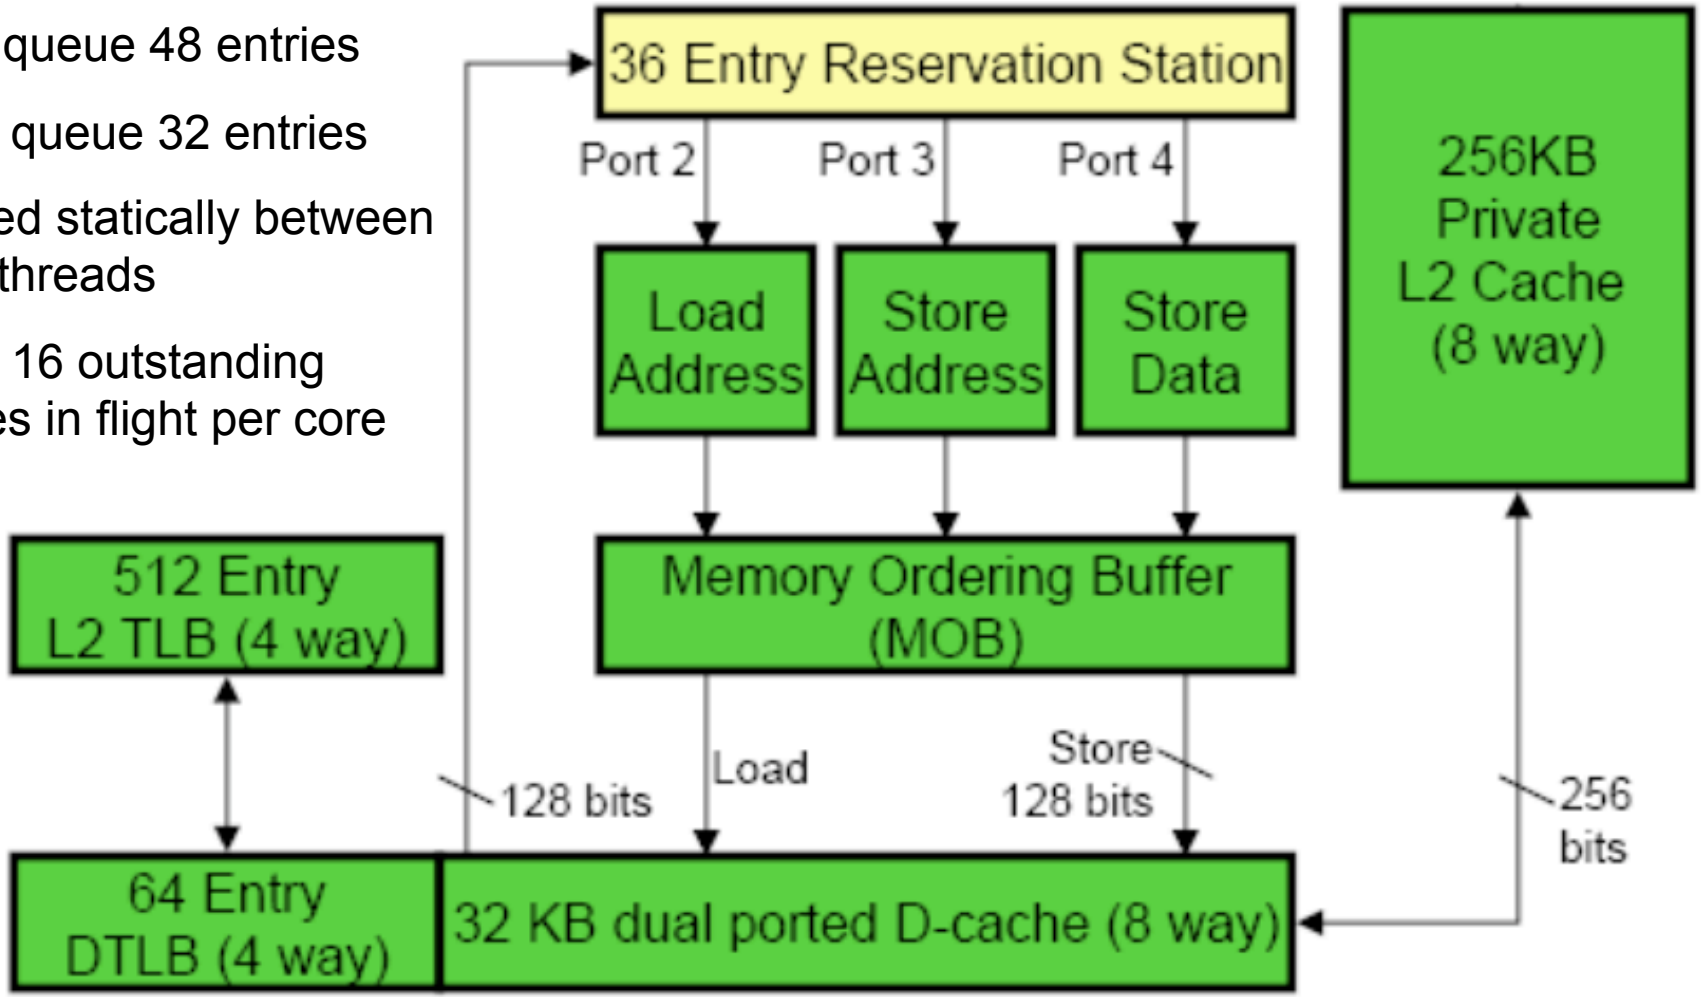

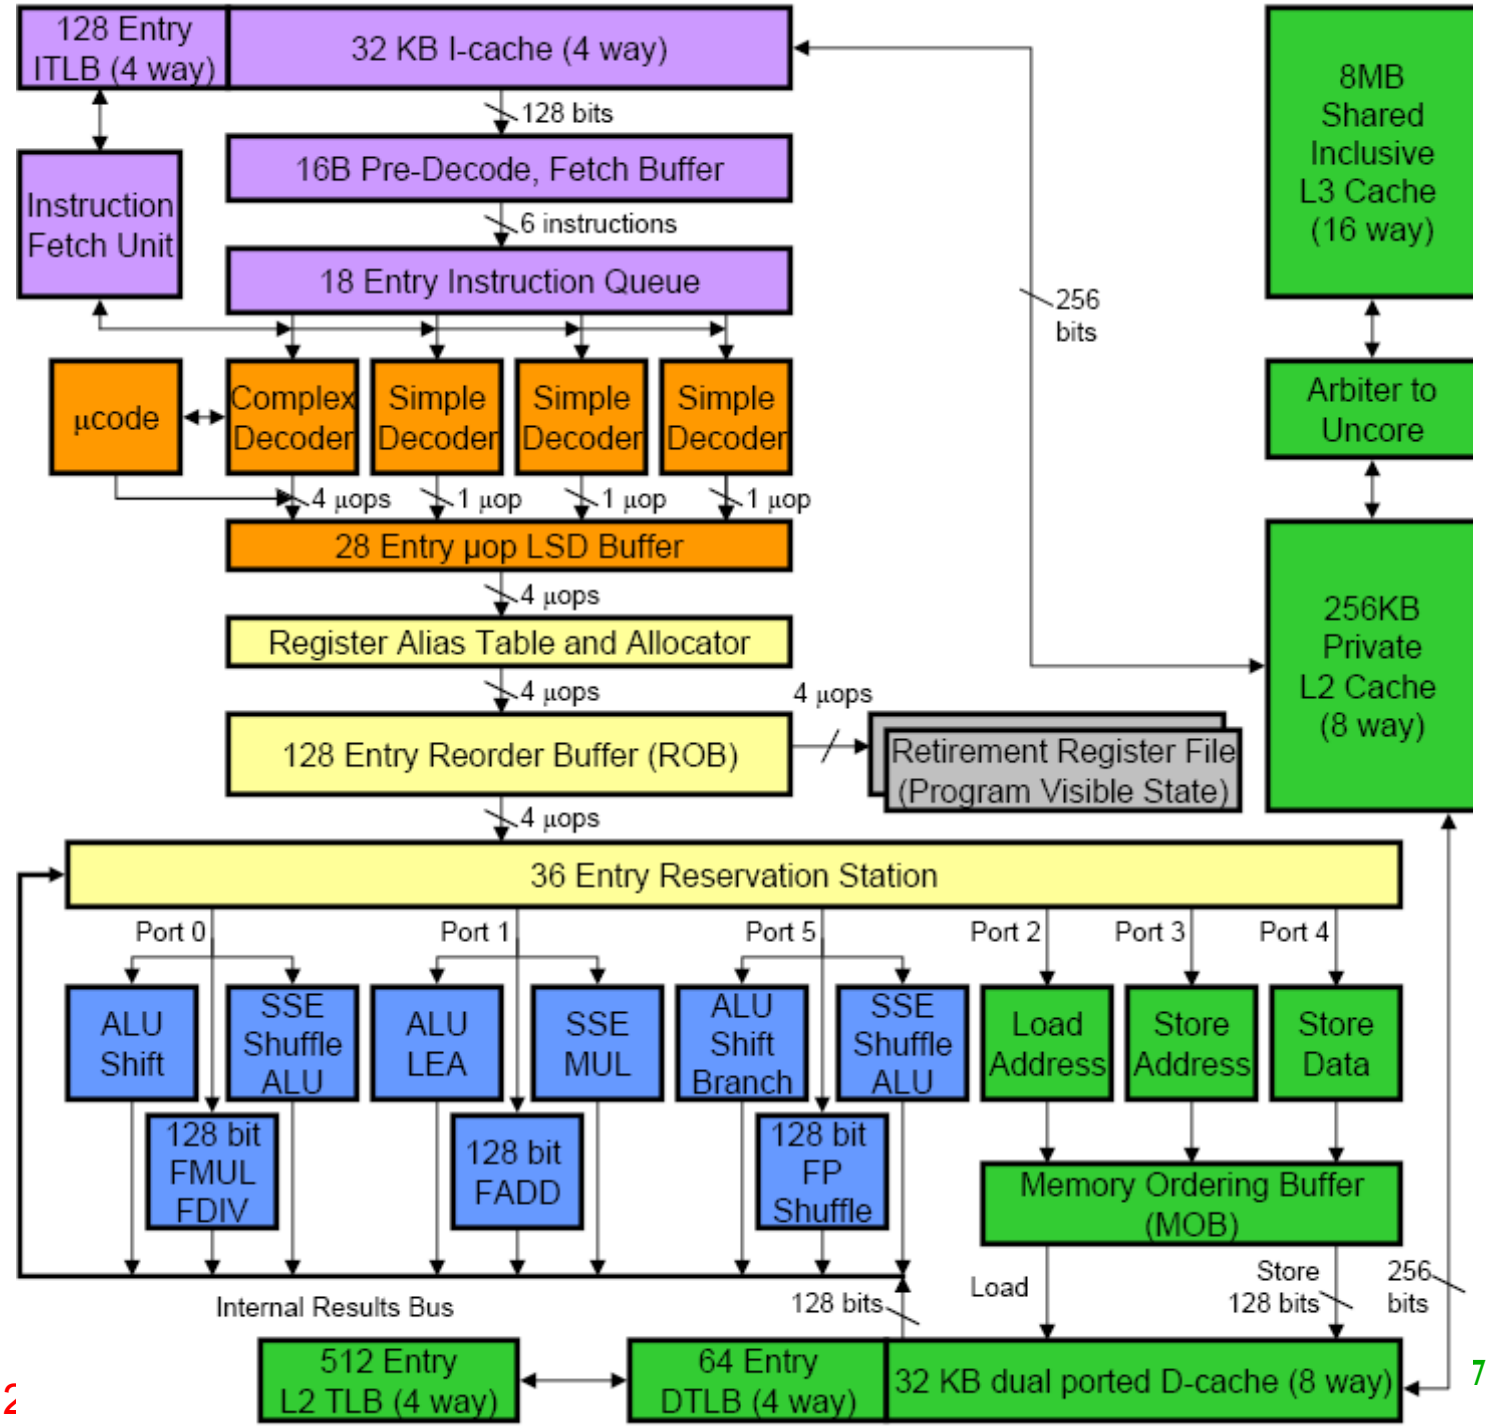
128 Entry ITLB (4 way)
32 KB I-cache (4 way)
128 bits
16B Pre-Decode, Fetch Buffer
6 instructions
Instruction Fetch Unit
18 Entry Instruction Queue
256 bits
8MB Shared Inclusive L3 Cache (16 way)
µcode
Complex Decoder
Simple Decoder
Simple Decoder
Simple Decoder
Arbiter to Uncore
4 µops
1 µop
1 µop
1 µop
28 Entry µop LSD Buffer
4 µops
Register Alias Table and Allocator
4 µops
4 µops
128 Entry Reorder Buffer (ROB)
Retirement Register File (Program Visible State)
256KB Private L2 Cache (8 way)
4 µops
36 Entry Reservation Station
Port 0
Port 1
Port 5
Port 2
Port 3
Port 4
ALU Shift
SSE Shuffle ALU
ALU LEA
SSE MUL
ALU Shift Branch
SSE Shuffle ALU
Load Address
Store Address
Store Data
128 bit FMUL FDIV
128 bit FADD
128 bit FP Shuffle
Memory Ordering Buffer (MOB)
Load
Store 128 bits
256 bits
Internal Results Bus
128 bits
512 Entry L2 TLB (4 way)
64 Entry DTLB (4 way)
32 KB dual ported D-cache (8 way)

Cache Hierarchy Latencies

- L1 32KB 8-way, latency 4 cycles
- L2 256KB 8-way, latency <12 cycles
- L3 8MB, 16-way, latency 30-40 cycles
- DRAM, latency ~180-200 cycles

Nehalem Virtual Memory Details

- Implements 48-bit virtual address space, 40-bit physical address space
- Two-level TLB
- I-TLB (L1) has shared 128 entries 4-way associative for 4KB pages, plus 7 dedicated fully-associative entries per SMT thread for large page (2/4MB) entries
- D-TLB (L1) has 64 entries for 4KB pages and 32 entries for 2/4MB pages, both 4-way associative, dynamically shared between SMT threads
- Unified L2 TLB has 512 entries for 4KB pages only, also 4-way associative
- Additional support for system-level virtual machines

Virtualization Support

- TLB entries tagged with virtual machine and address space ID
  - No need to flush on context switches between VMs
- Hardware page table walker can walk guest-physical to host-physical mapping tables
  - Fewer traps to hypervisor

Core Area Breakdown

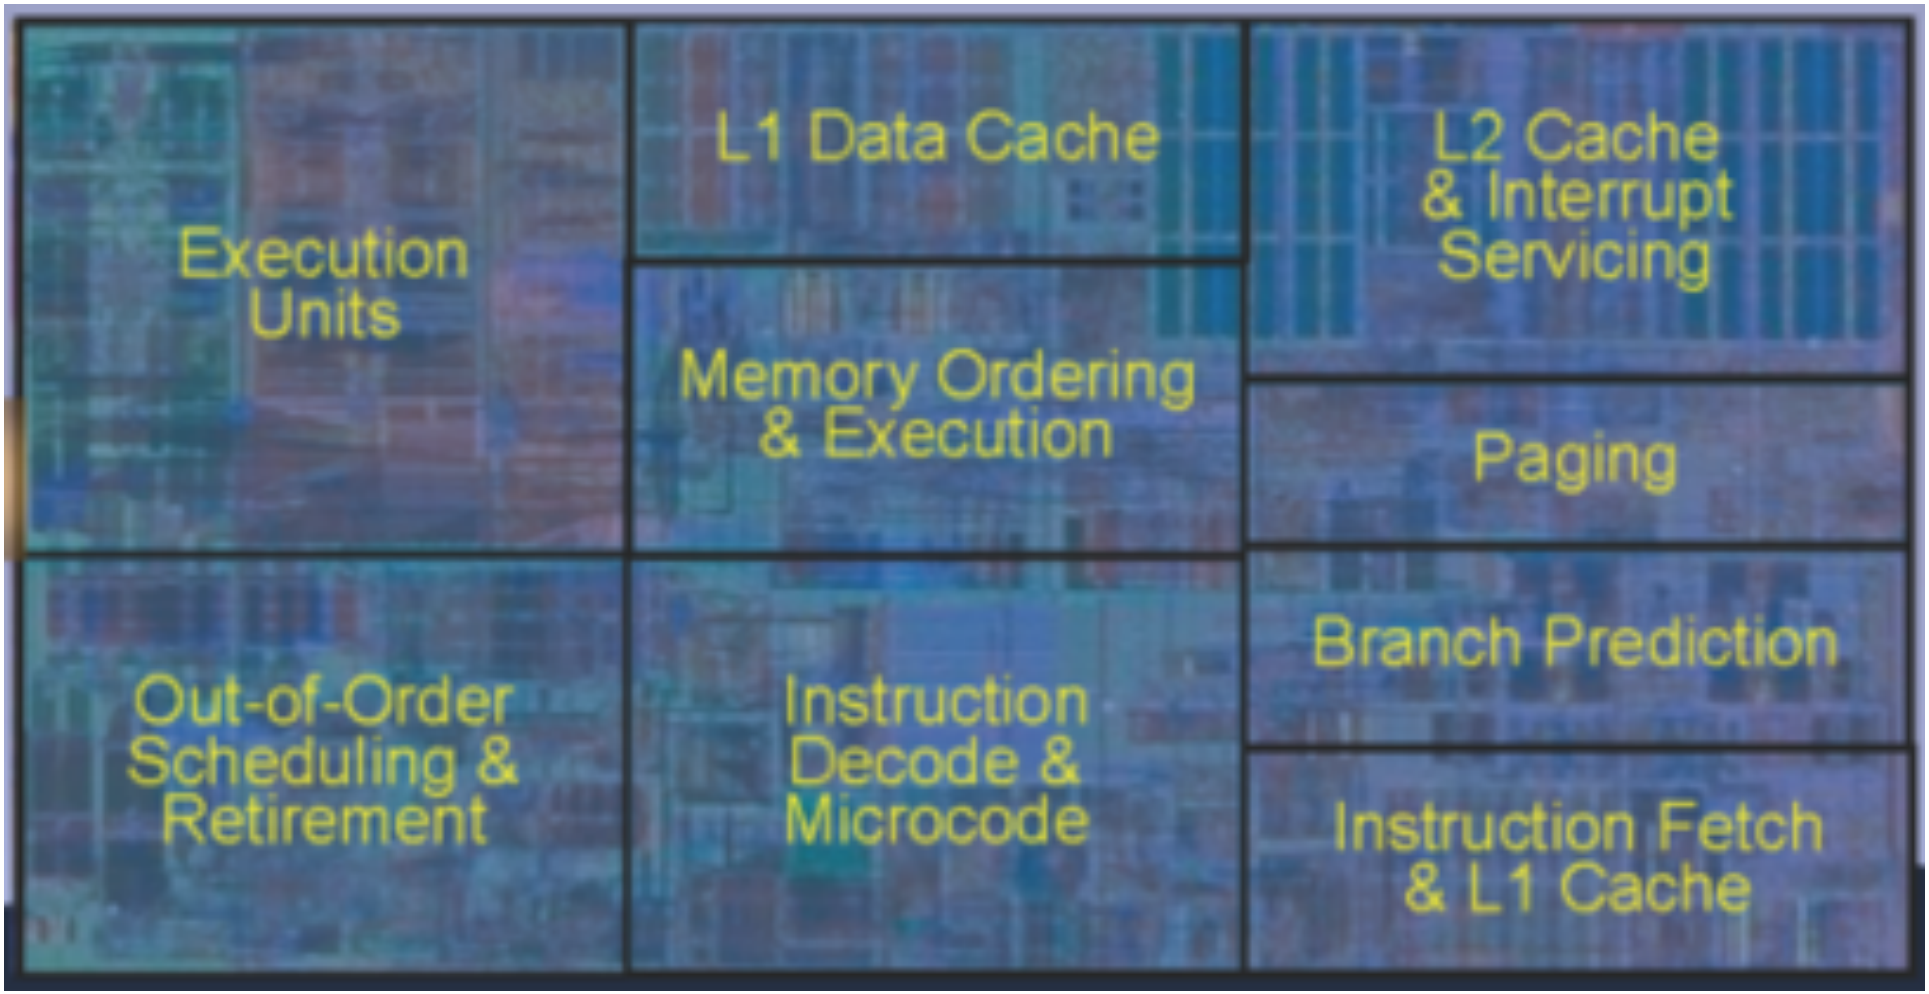

(Nehalem) Turbo Mode

Power Gating

Zero power for inactive cores

Turbo Mode

In response to workload adds additional performance bins within headroom

No Turbo

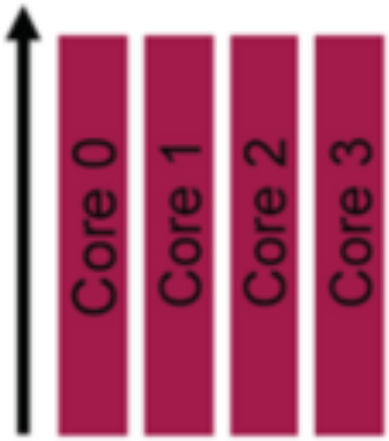

Workload Lightly Threaded or < TDP

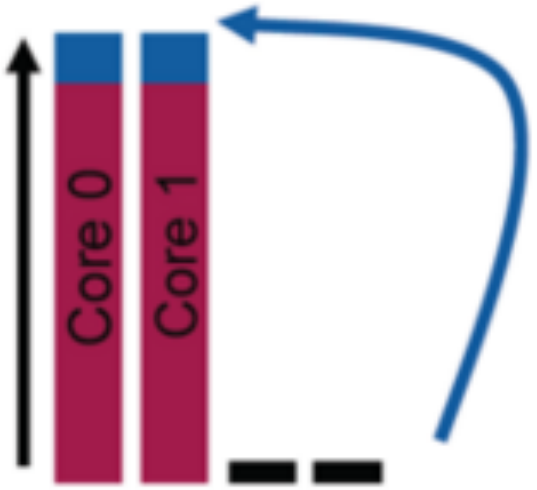

Related Courses

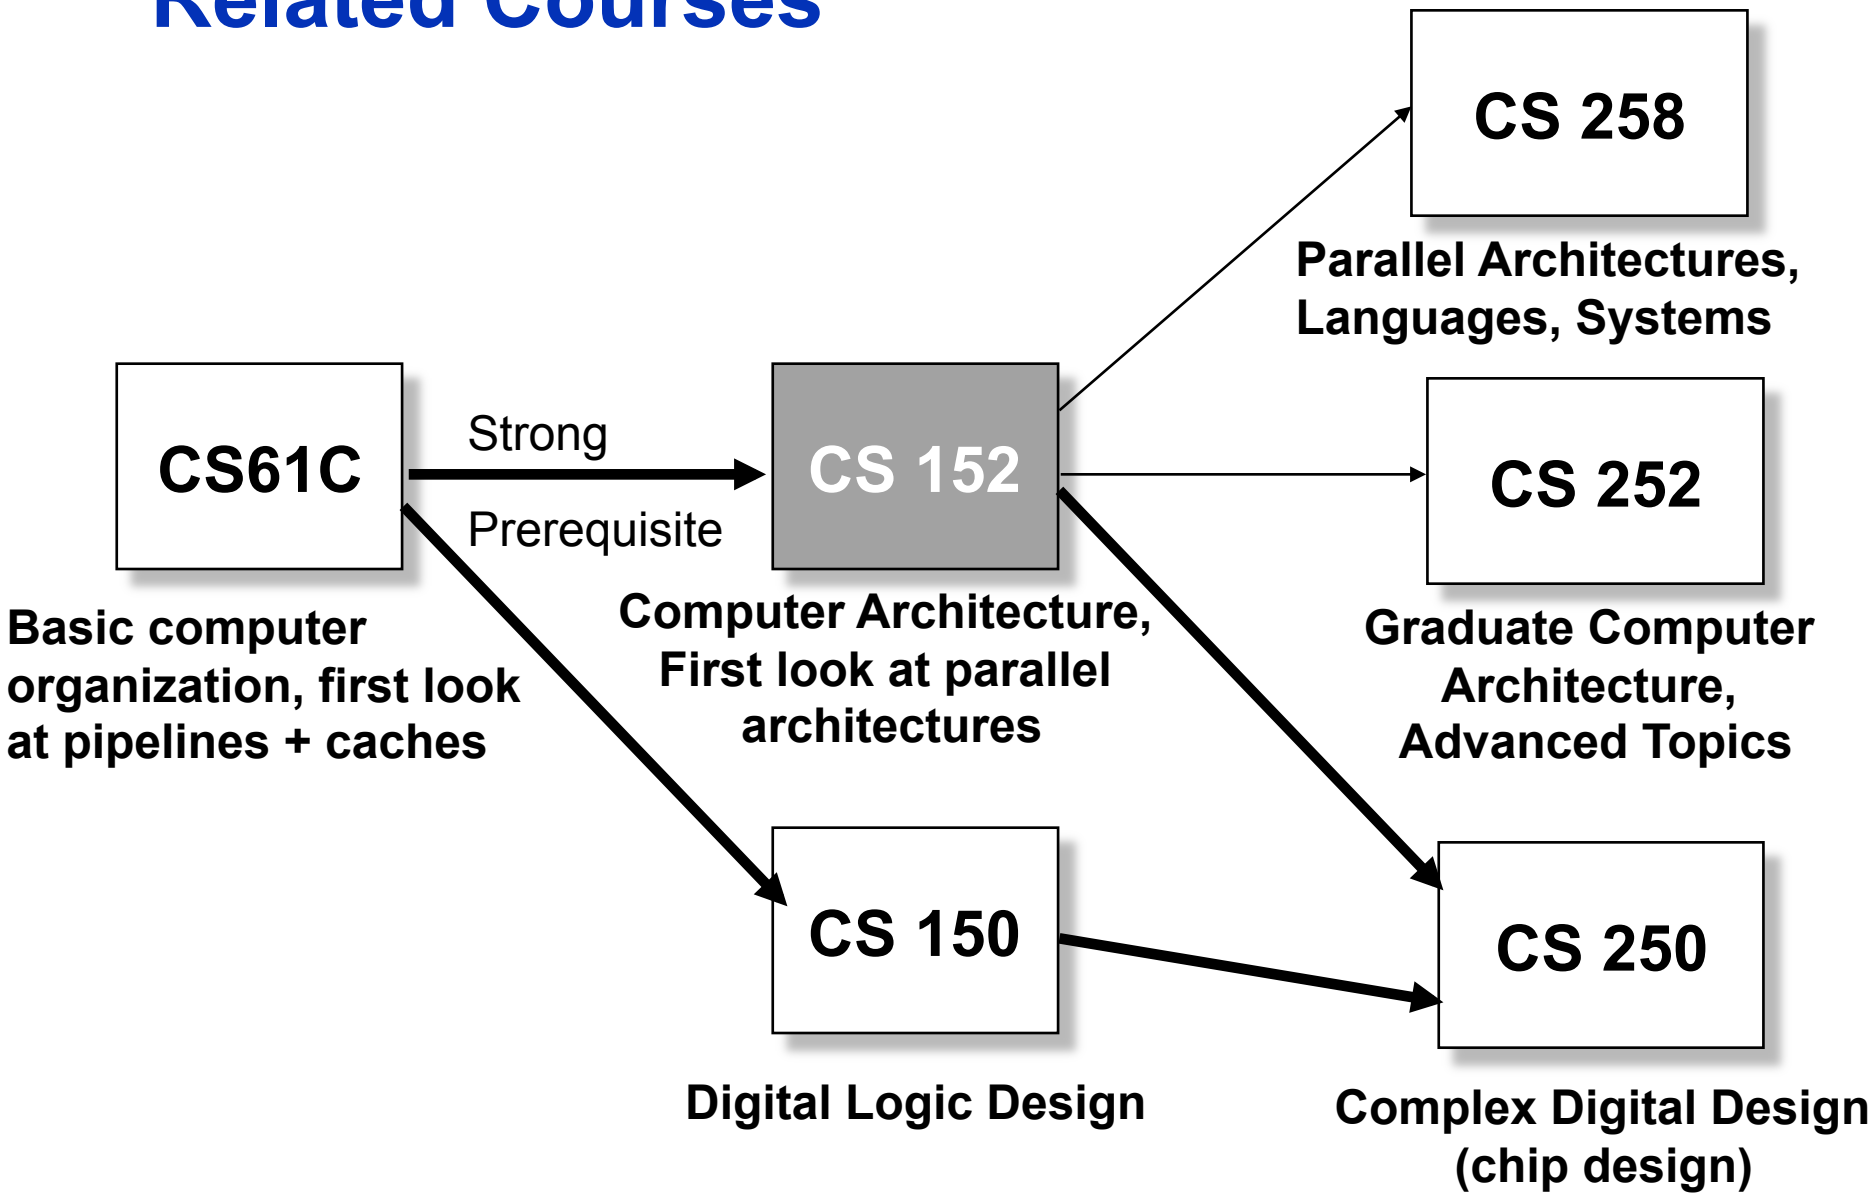

Advice: Get involved in research

E.g.,

- RAD Lab - data center
- Par Lab - parallel clients
- AMP Lab – algorithms, machines, people
- LoCAL – networking energy

- Undergrad research experience is the most important part of application to top grad schools, and fun too.

End of CS152

- Final Quiz 5 on Thursday (lectures 19, 20, 21)
- HKN survey to follow.
- Thanks for all your feedback - we'll keep trying to make CS152 better.