On the Road to Concurrency...

Core™ Architecture

Victoria Zhislina
Senior Application Engineer
Intel

Core™ Architecture

- **Introduction**
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - Smart Memory Access
  - Enhanced Smart Cache
  - Intelligent Power Capability
- Software Impact
- Benchmark Results
- 45nm Technology

Intel® Architecture History

Architecture:
Instruction set definition and compatibility

Microarchitecture:
Hardware implementation maintaining instruction set compatibility with high-level architecture

Processors:
Productized implementation of Microarchitecture

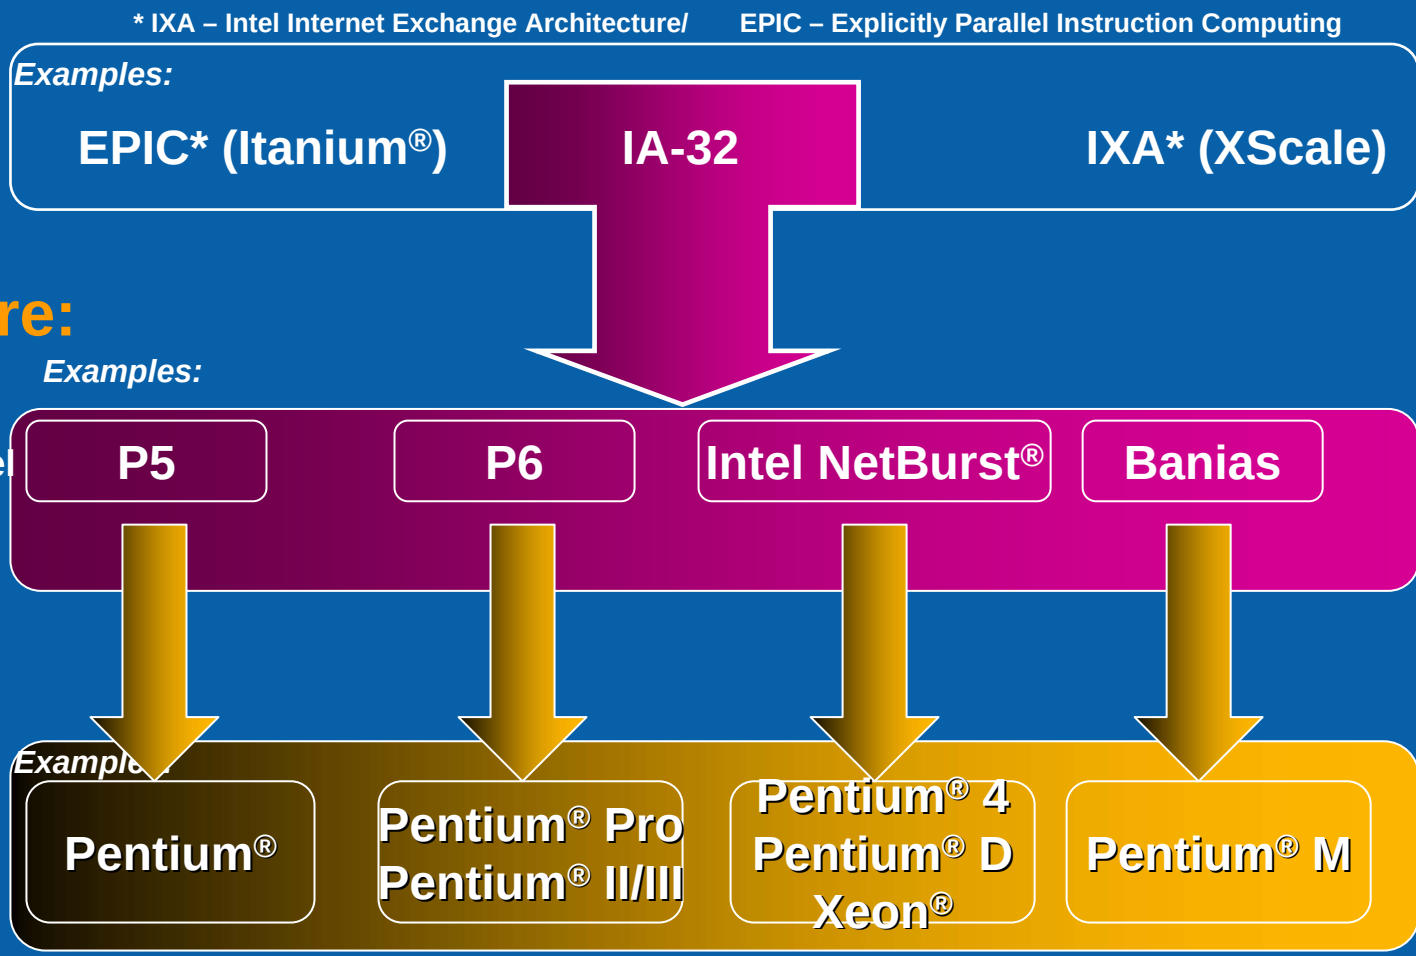

Overview

Intel® NetBurst®

+ New Innovations

Mobile Microarchitecture

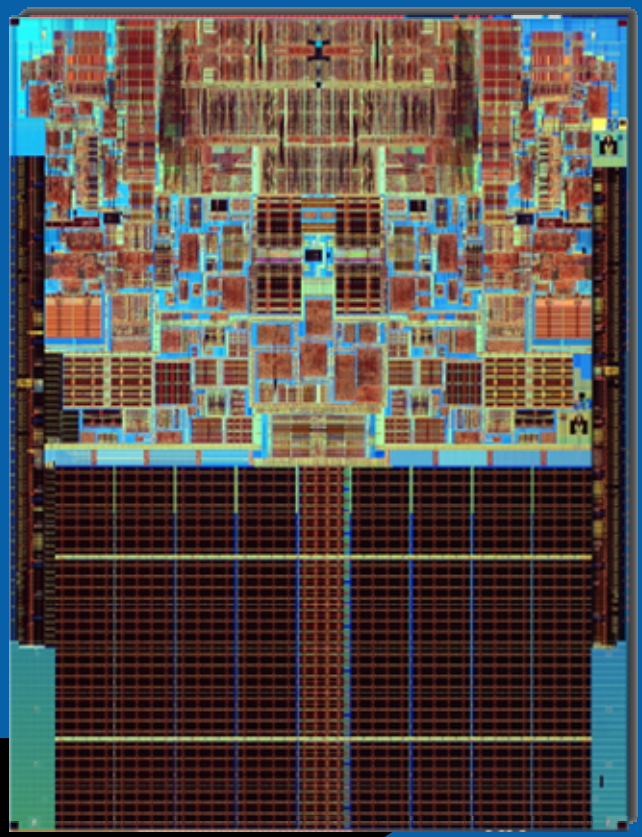

Intel® Core™ 2 Duo/Quad/Extreme processors

A Note on N

- The Core Microarchitecture
  - Is used in the core of Core 2 Duo an Core 2 Quad
  - As opposed to the Core Duo which does *not* use the Core Microarchitecture
- By the way, Core is the first Intel converged core
  - As opposed to Nehalem, which will be a converged core different form the Core Microarchitecture
  - Or to Core Duo, which is not a converged core

Dual Core Block Diagram

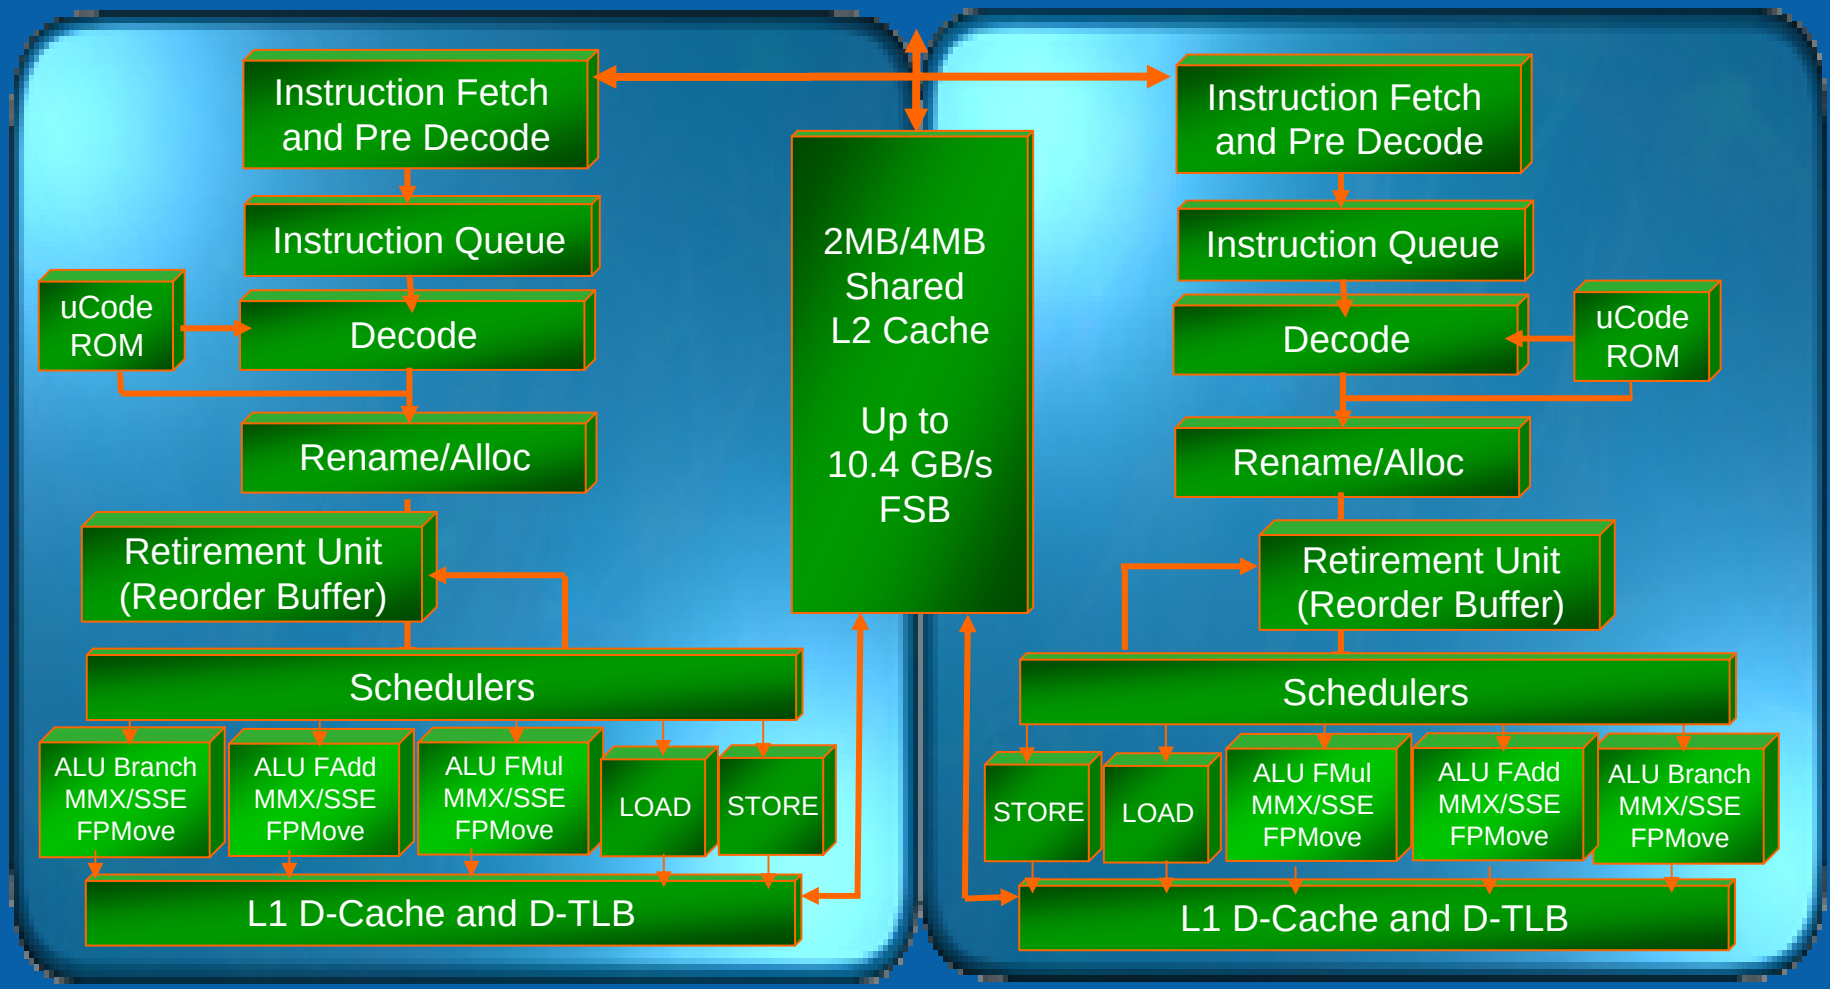

Micro Architecture Comparison

	Intel NetBurst™++	Core Architecture
Pipeline Stages	31	14
Threads per core	2	1
L1 Cache Org.	2 x (12K uop Trace Cache/16K Data)	2 x (32K I/32K Data)
L2 Cache Org.	2 x 2MB	1 x 4MB (shared)
Instr. Decoders	1	4
Integer Units	2 (2x core freq)	3 (1x core freq)
SIMD Units	2 x 64-bits	3 x 128-bits
SIMD Inst. Issued per Clock	1	3
FP Units	3 (Add/Mul/Div)	3 (Add/Mul/Div)
FP Inst. Issued per clock	1	Up to 2 (Add + Mul or Div)
Power	135W	80W

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - **Enhanced Execution Capabilities**
  - Smart Memory Access
  - Enhanced Smart Cache
  - Intelligent Power Capability
- Software Impact
- Benchmark Results
- 45nm Technology

Wide Dynamic Execution: Macro-Fusion

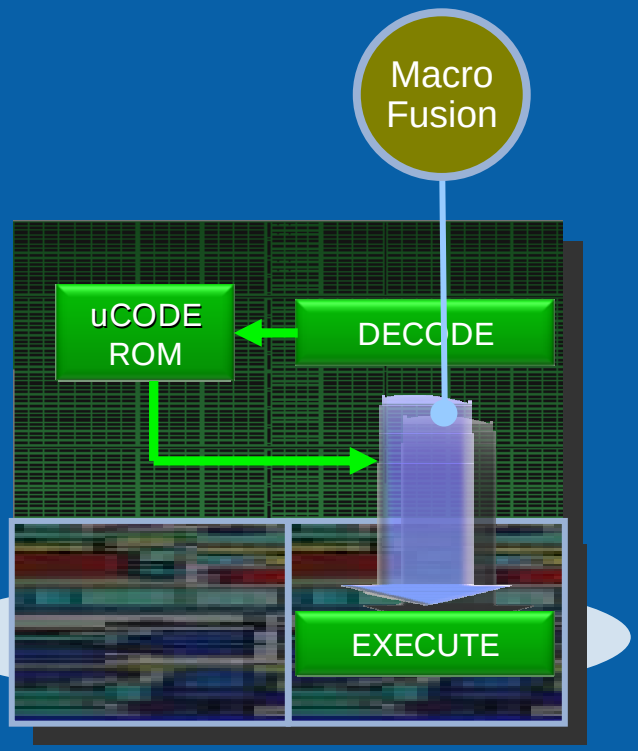

MACRO FUSION EXAMPLE

CMP+JMP IN 1 CLOCK

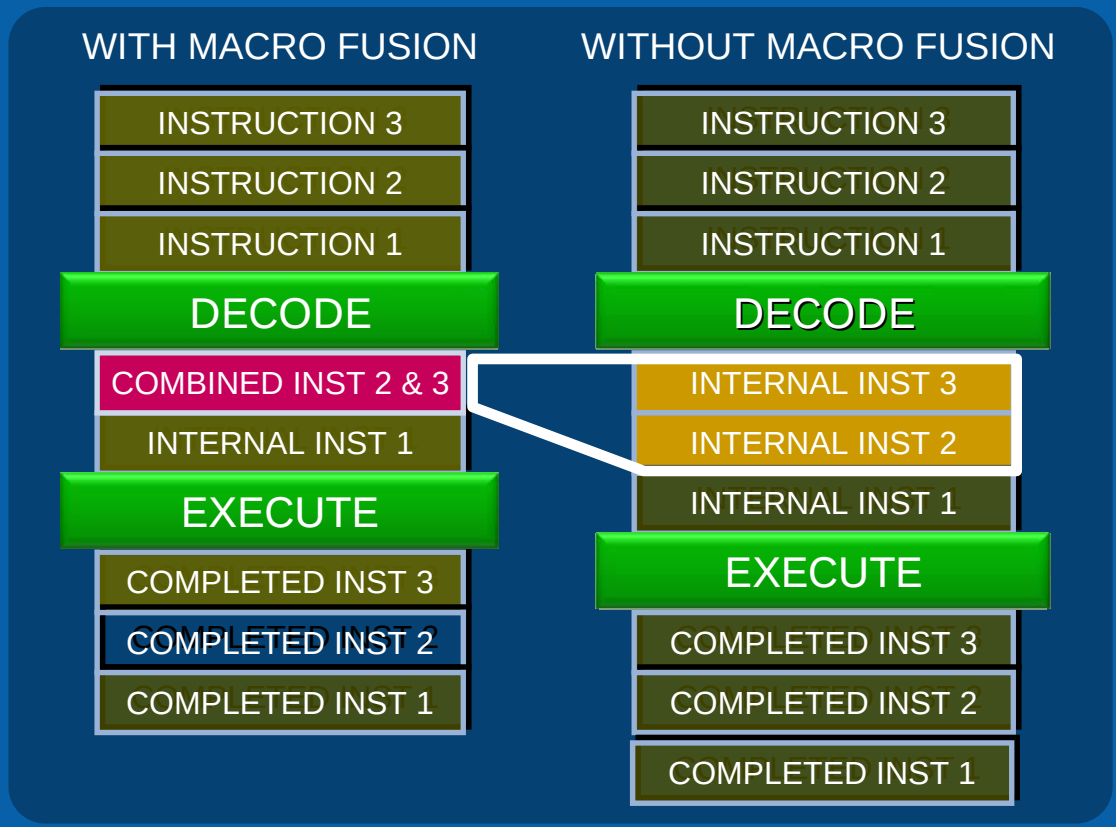

Perf ↑

Energy ↓

ADVANTAGE

- Instruction Load Reduced ~ 15%
- Micro-Ops Reduced ~ 10%

*Graphics not representative of actual die photo or relative size

Without Macro-Fusion

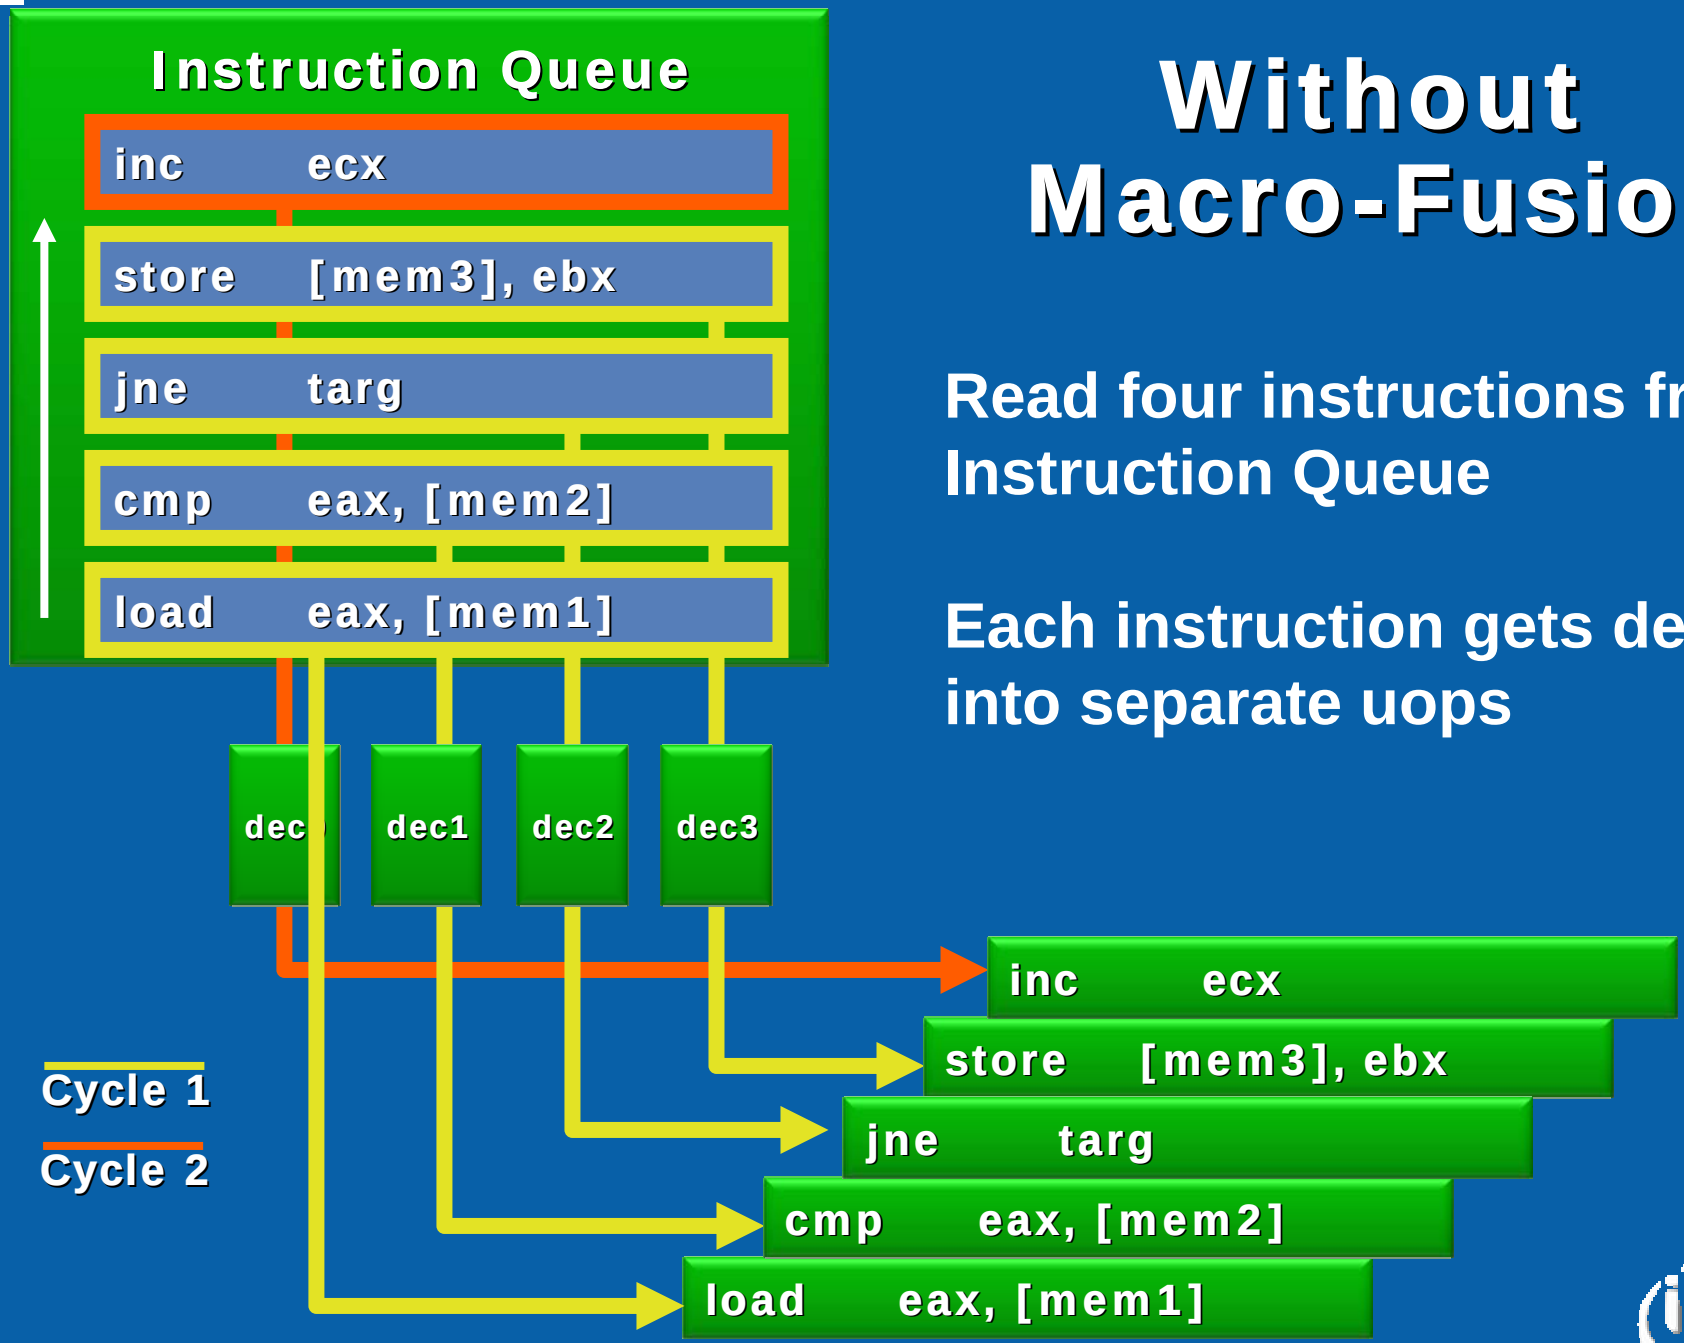

Read four instructions from Instruction Queue

Each instruction gets decoded into separate uops

intel

With Macro-Fusion

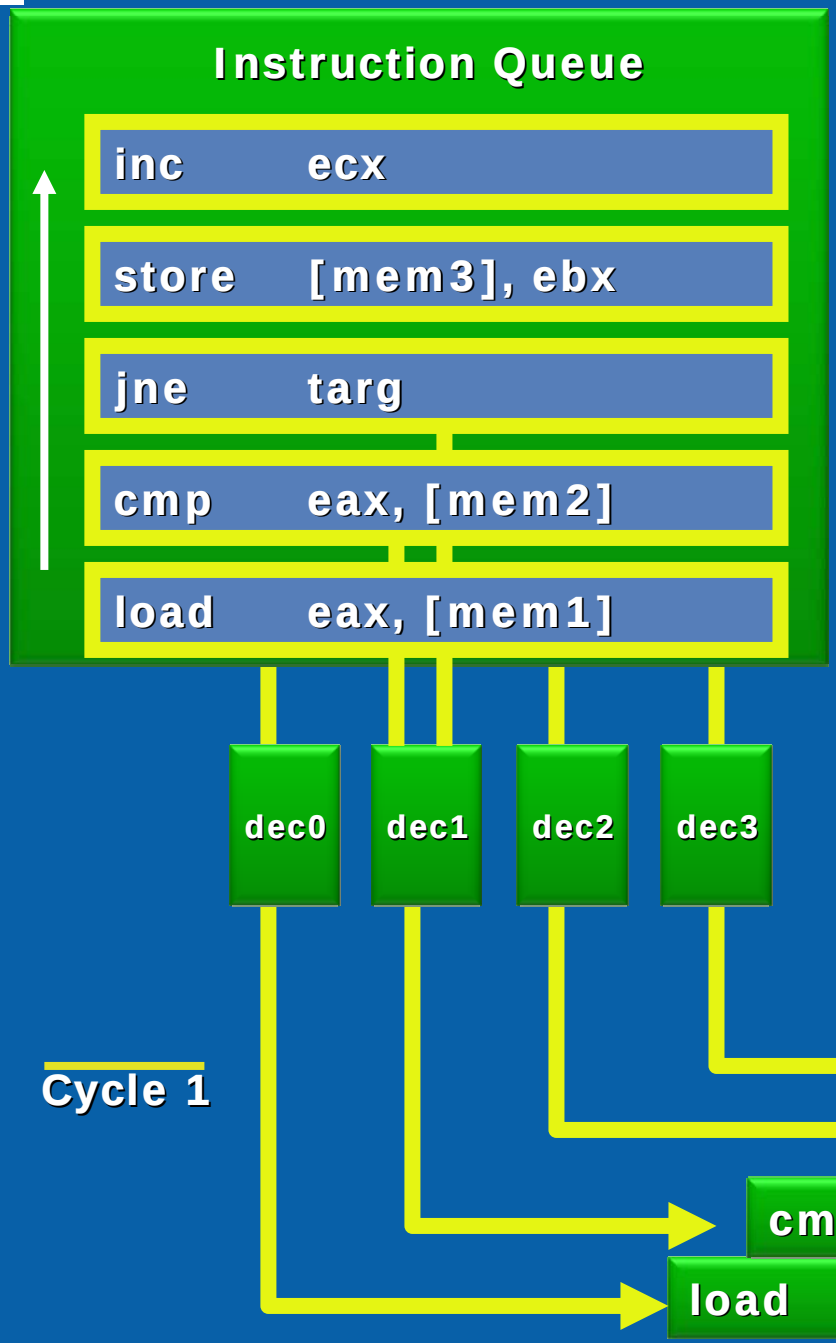

Read five Instructions from Instruction Queue

Send fusable pair to single decoder

Single uop represents two instructions

Intel® Advanced Digital Media Boost

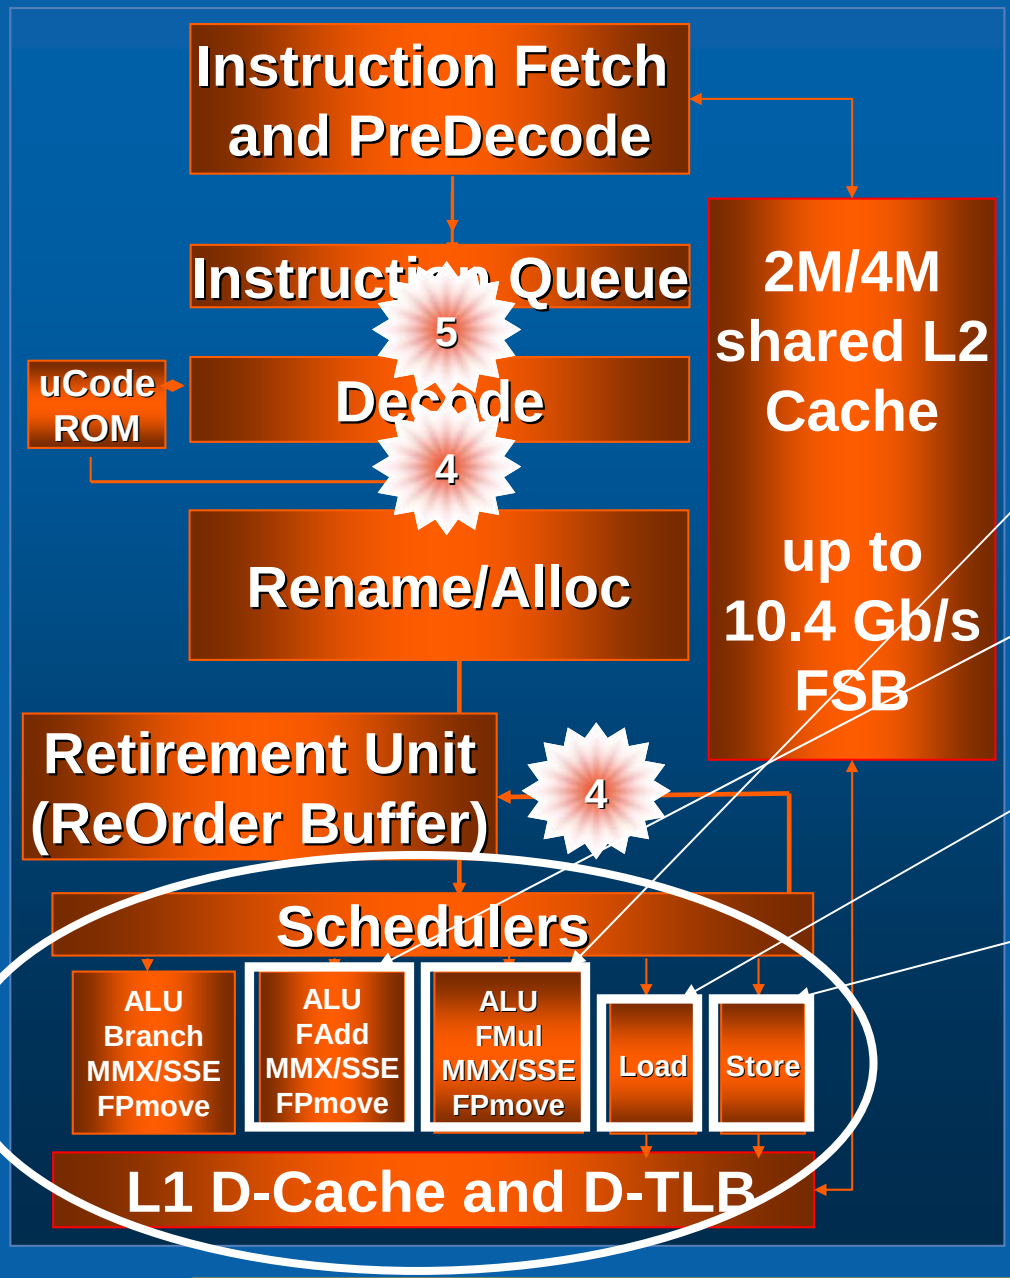

128-bit packed Multiply

plus

128-bit packed Add

plus

128-bit packed Load

plus

128-bit packed Store

2x Compute Throughput / Clock

intel

Intel® Advanced Digital Media Boost

In Each Core

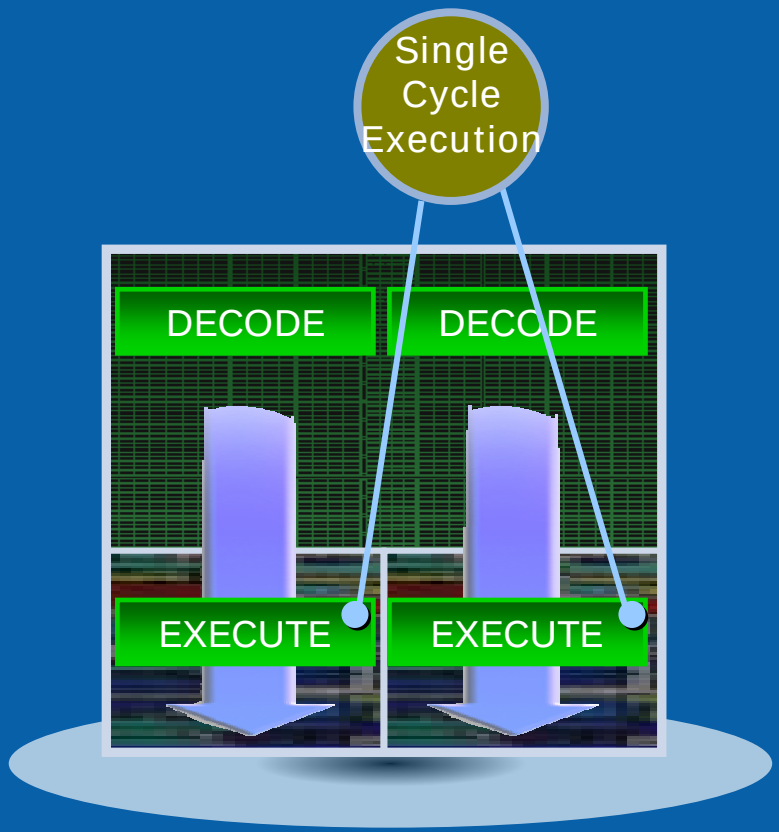

SSE Operation

(SSE/SSE2/SSE3)

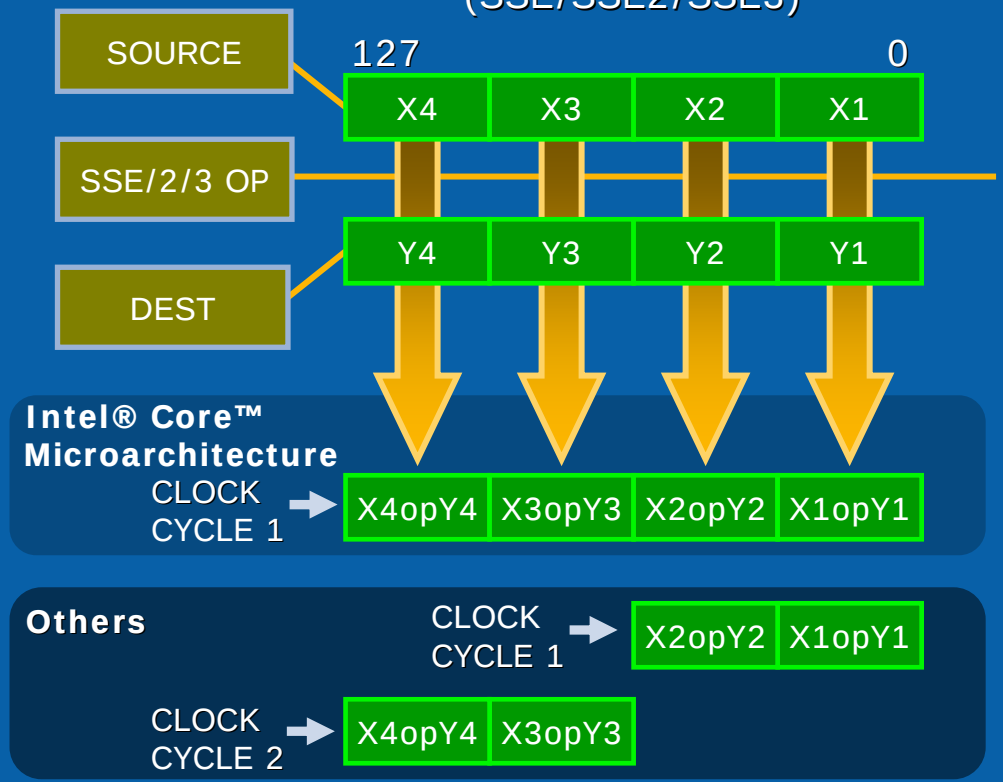

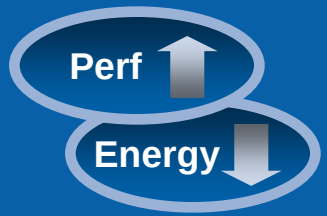

ADVANTAGE

- Increased Performance
- 128 bit Single Cycle In Each Core
- Improved Energy Efficiency

*Graphics not representative of actual die photo or relative size

intel

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - **Smart Memory Access**
  - Enhanced Smart Cache
  - Intelligent Power Capability
- Software Impact
- Benchmark Results
- 45nm Technology

Intel® Smart Memory Access

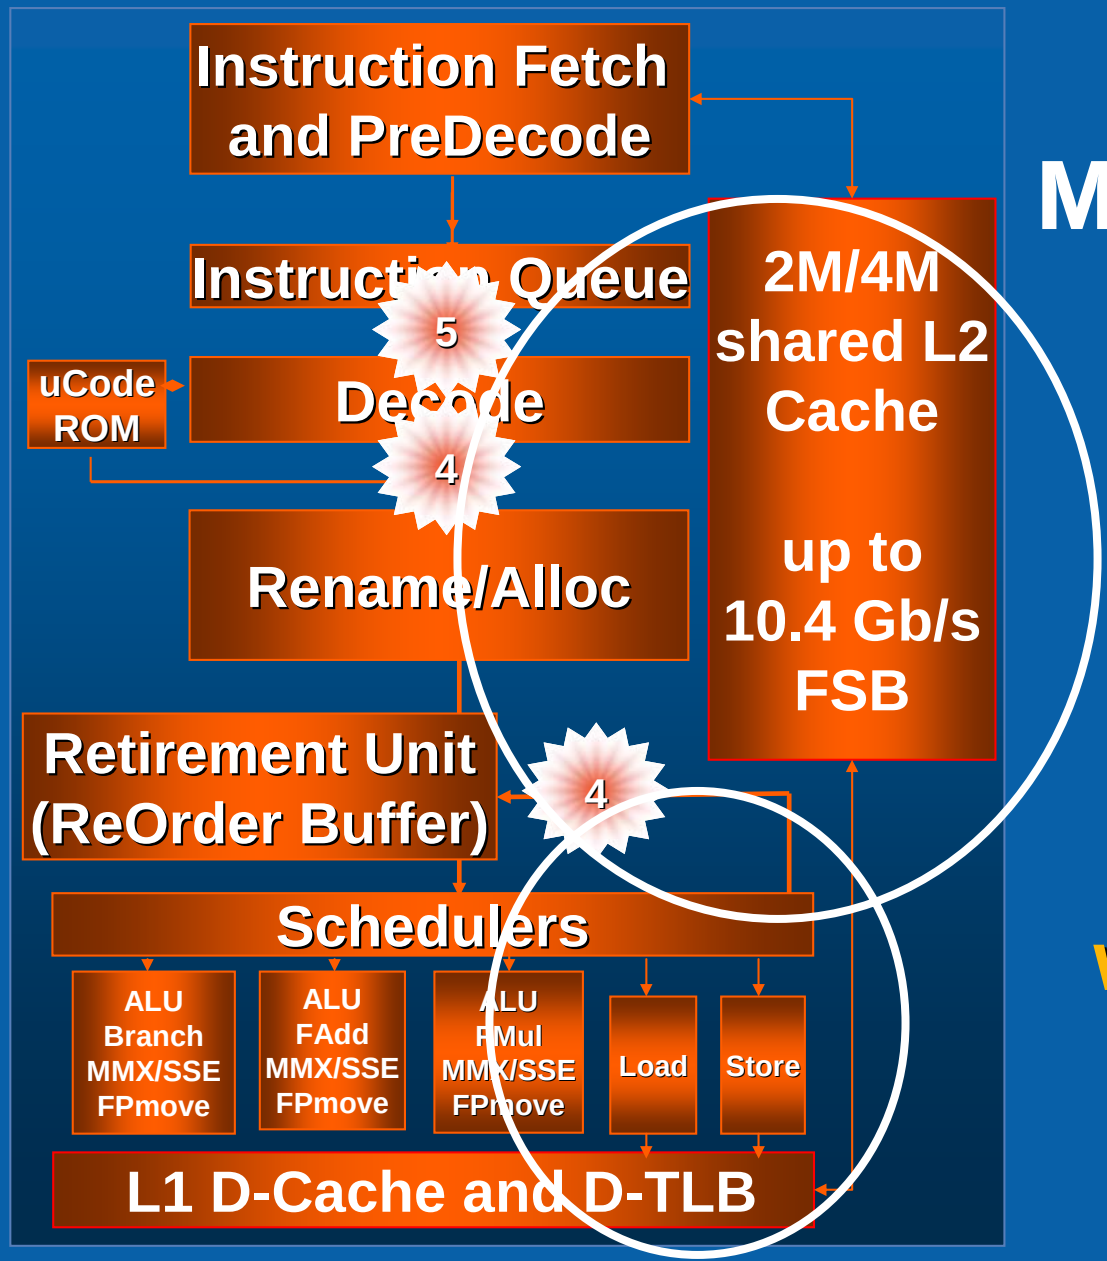

- Improved Prefetchers
- Memory Disambiguation

Data Is Where You Need It, When You Need It

Hiding Latency to Memory Subsystem

Intel Smart Memory Access: Prefetchers

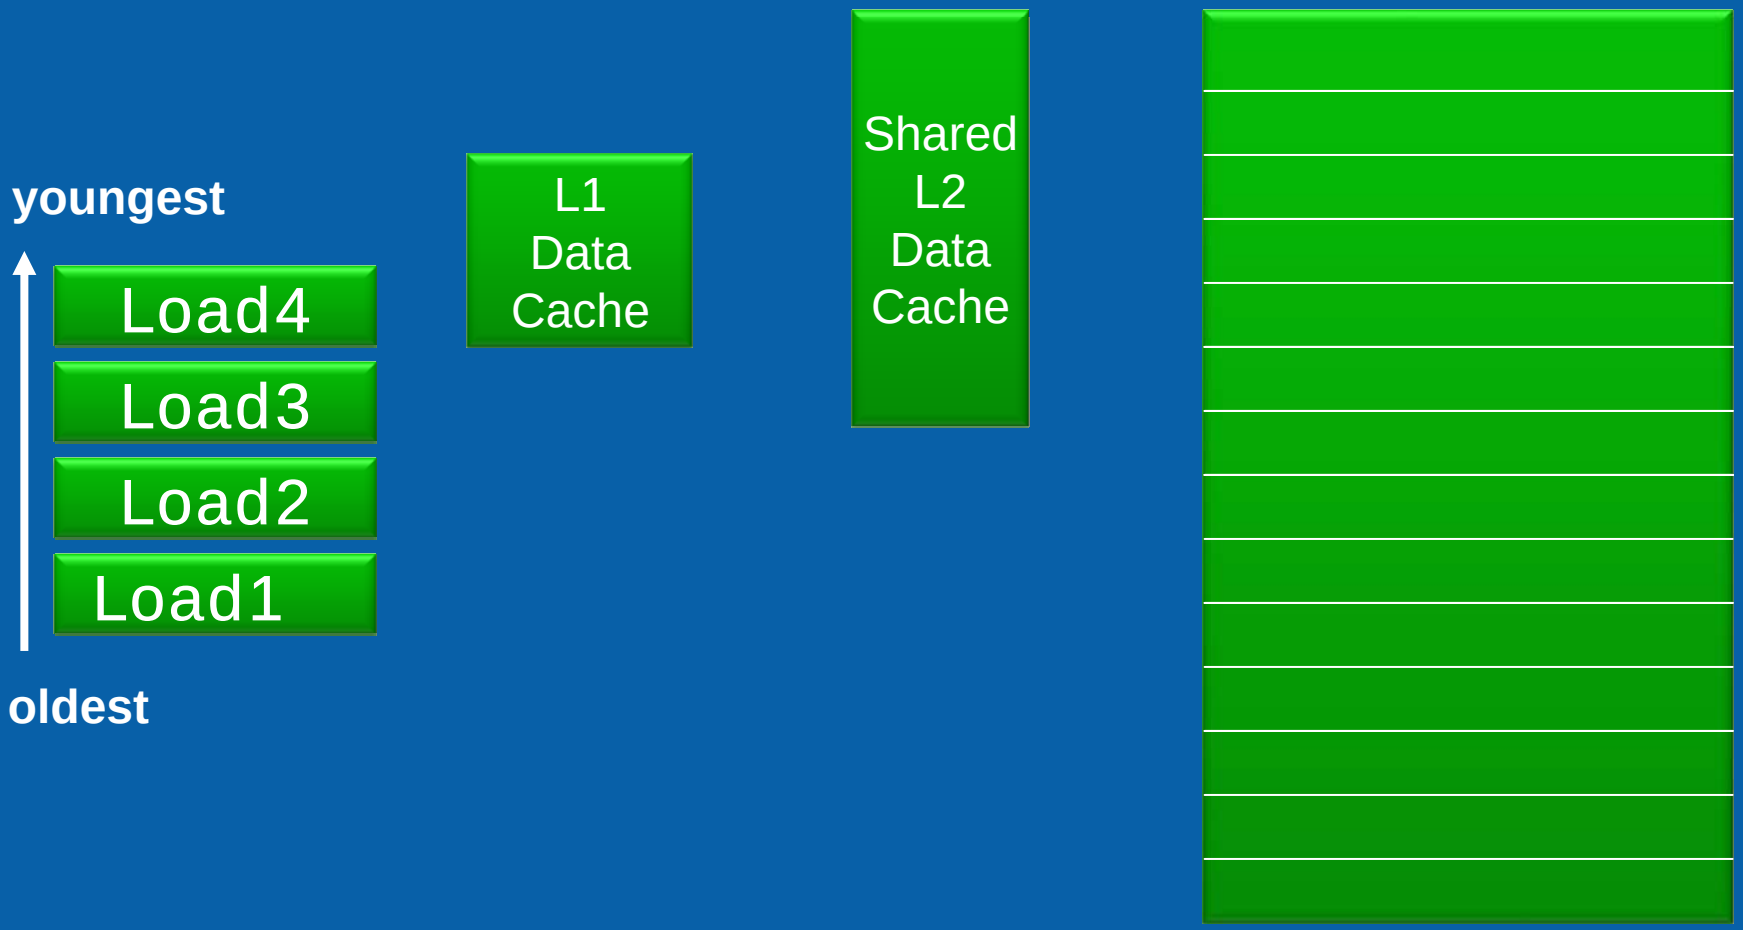

Intel Smart Memory Access: Prefetchers

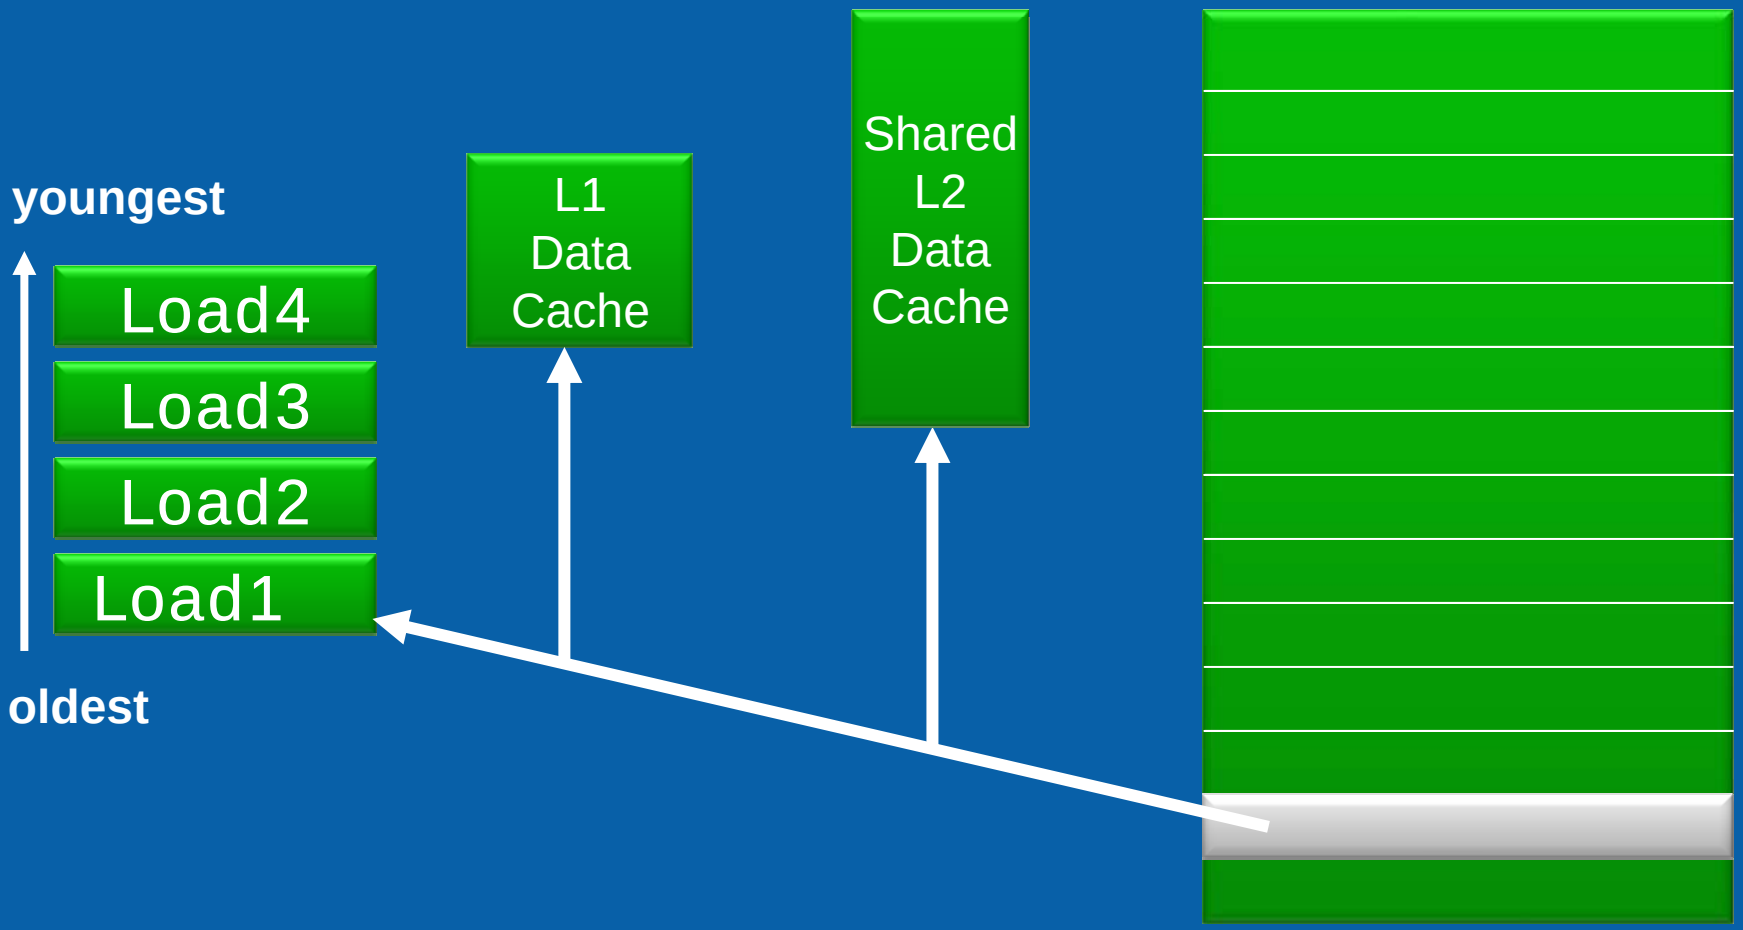

Memory is too far away

Intel Smart Memory Access: Prefetchers

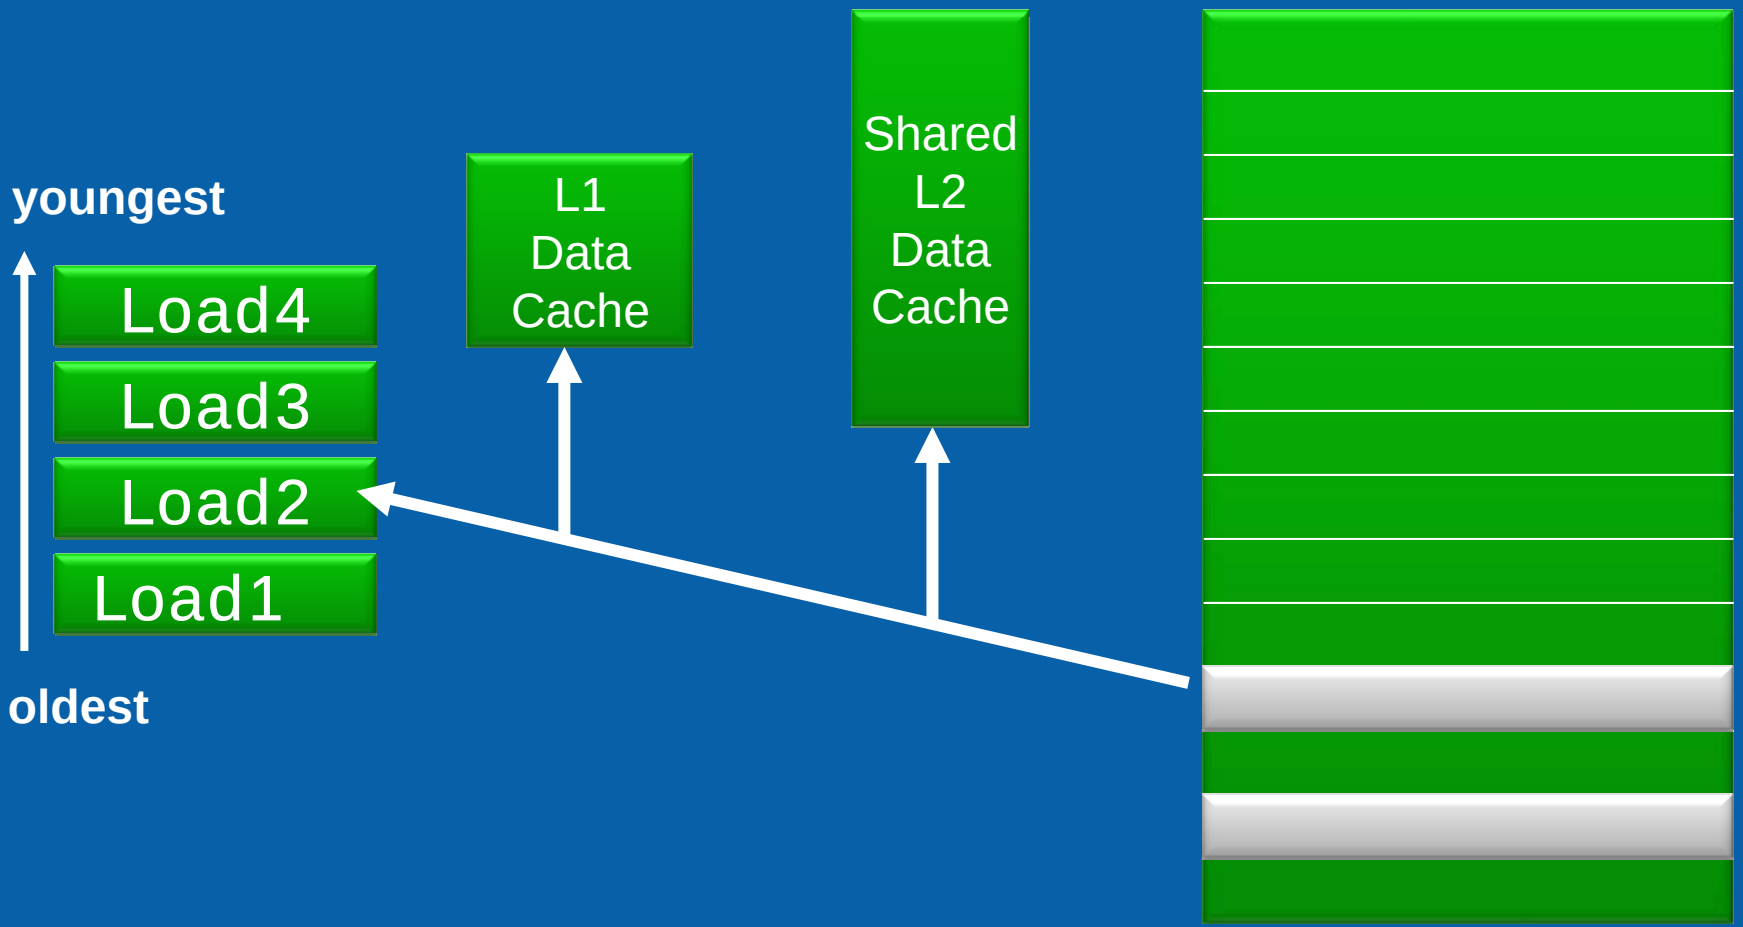

Caches are closer when they have the data

Intel Smart Memory Access: Prefetchers

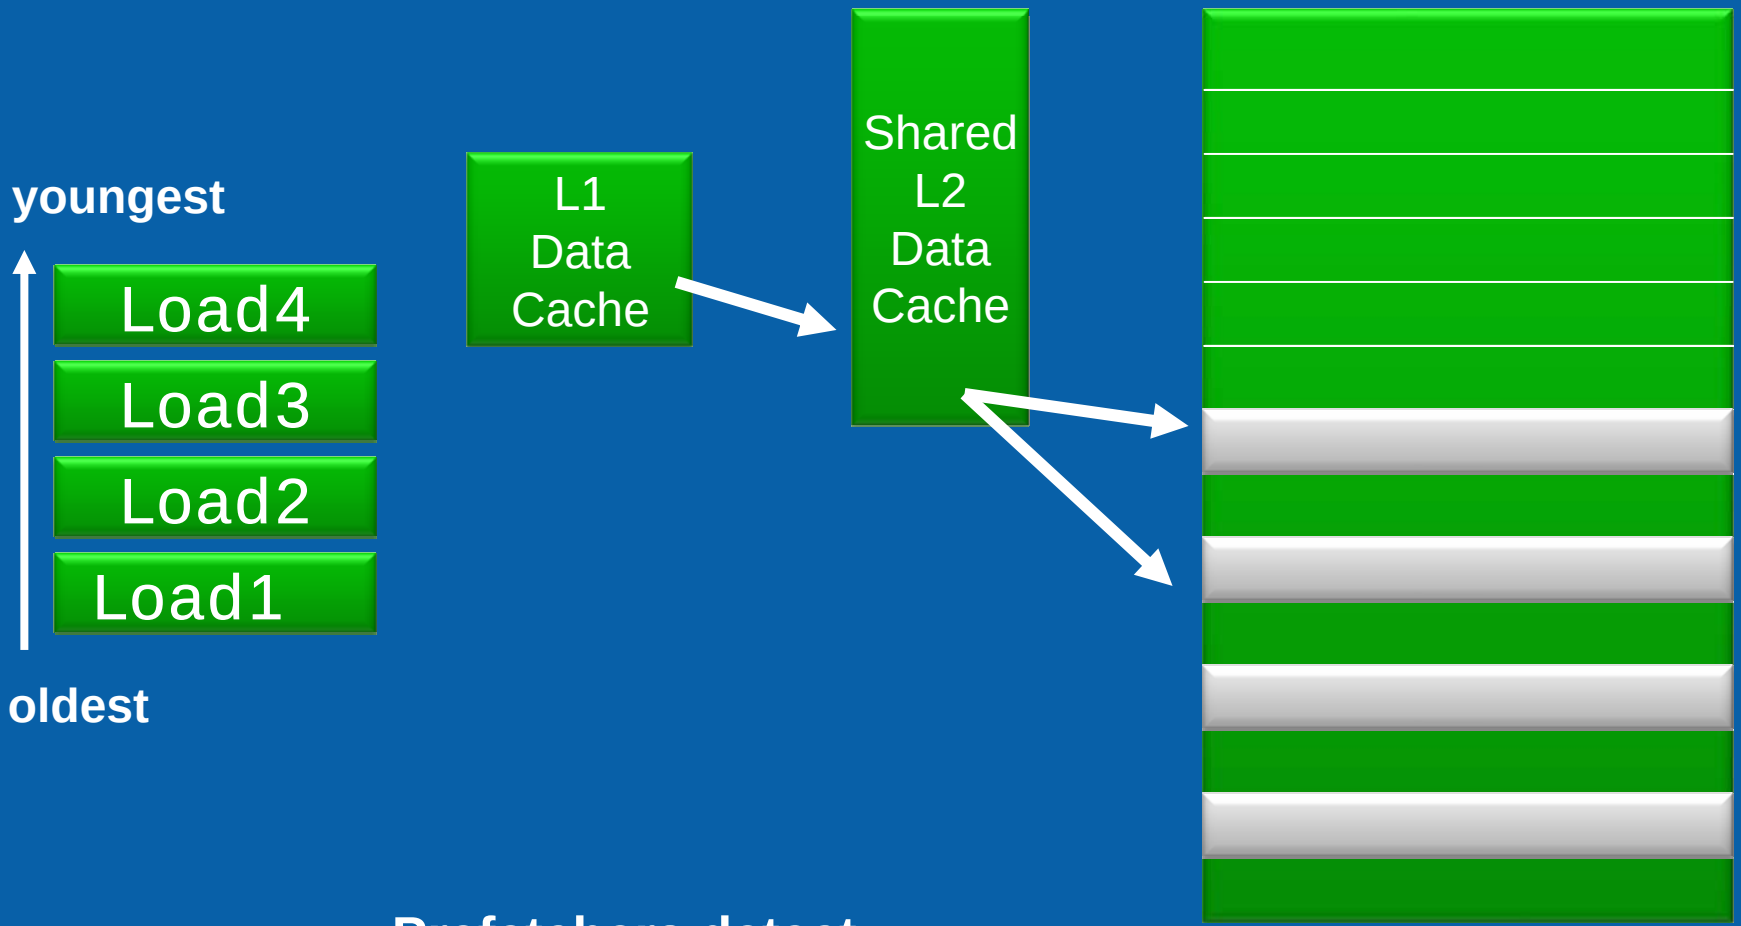

Prefetchers detect applications data reference patterns

Intel Smart Memory Access: Prefetchers

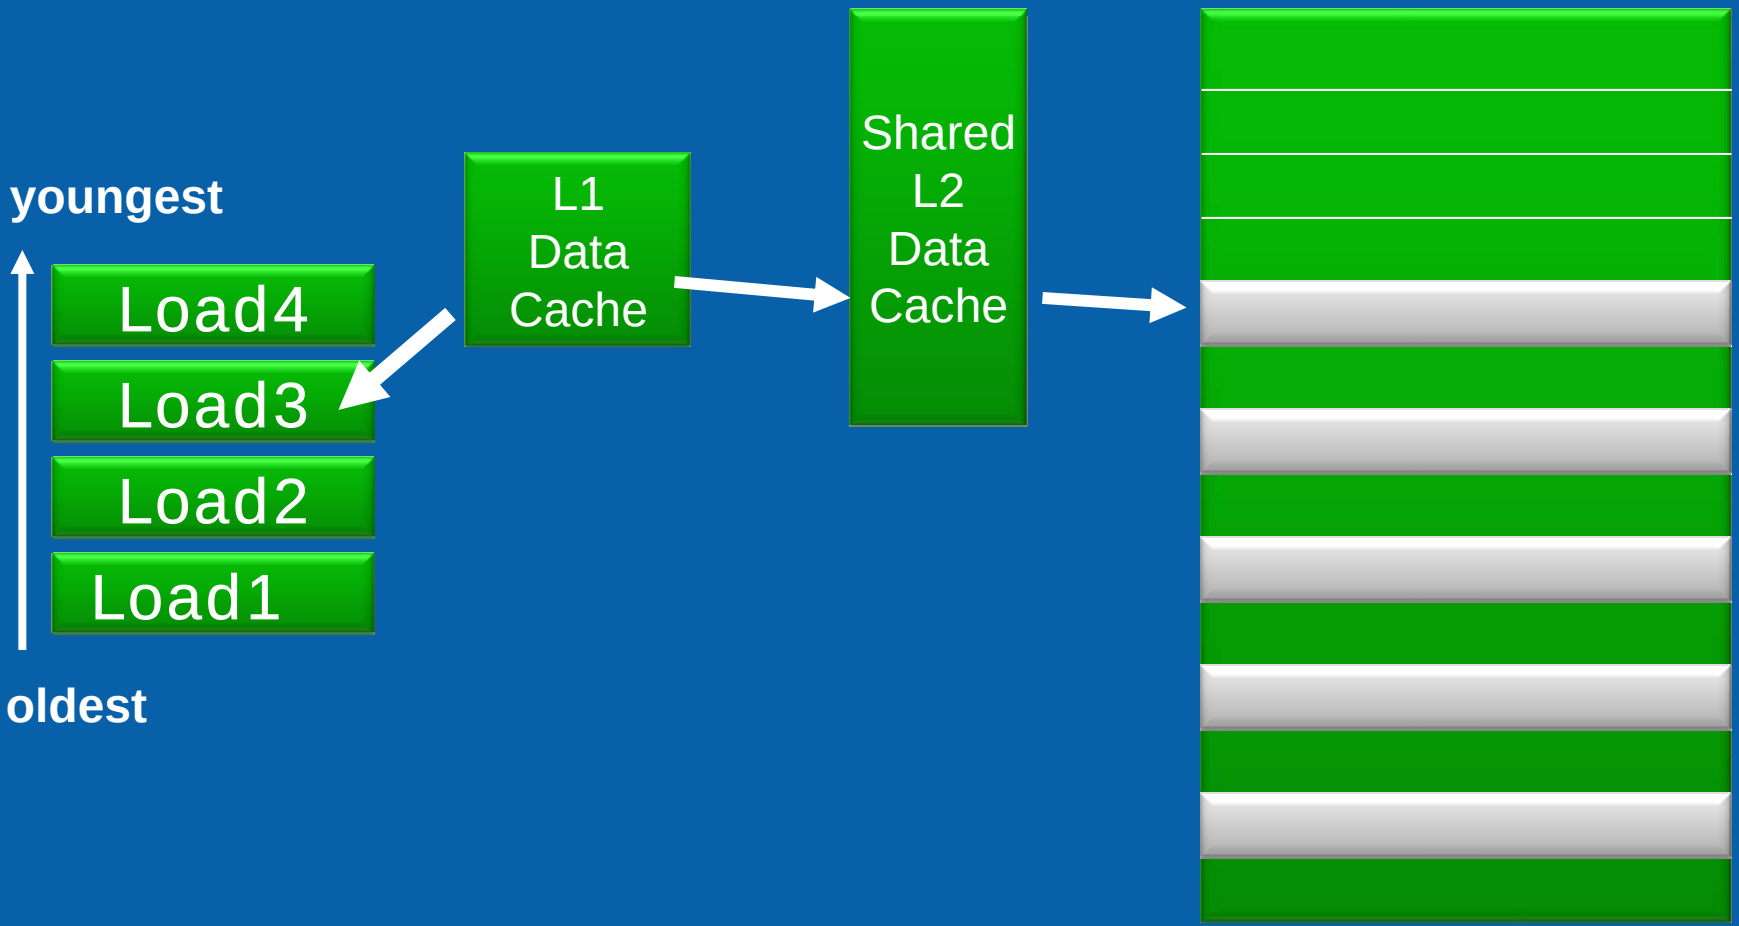

And bring the data closer to data consumer

Intel Smart Memory Access: Prefetchers

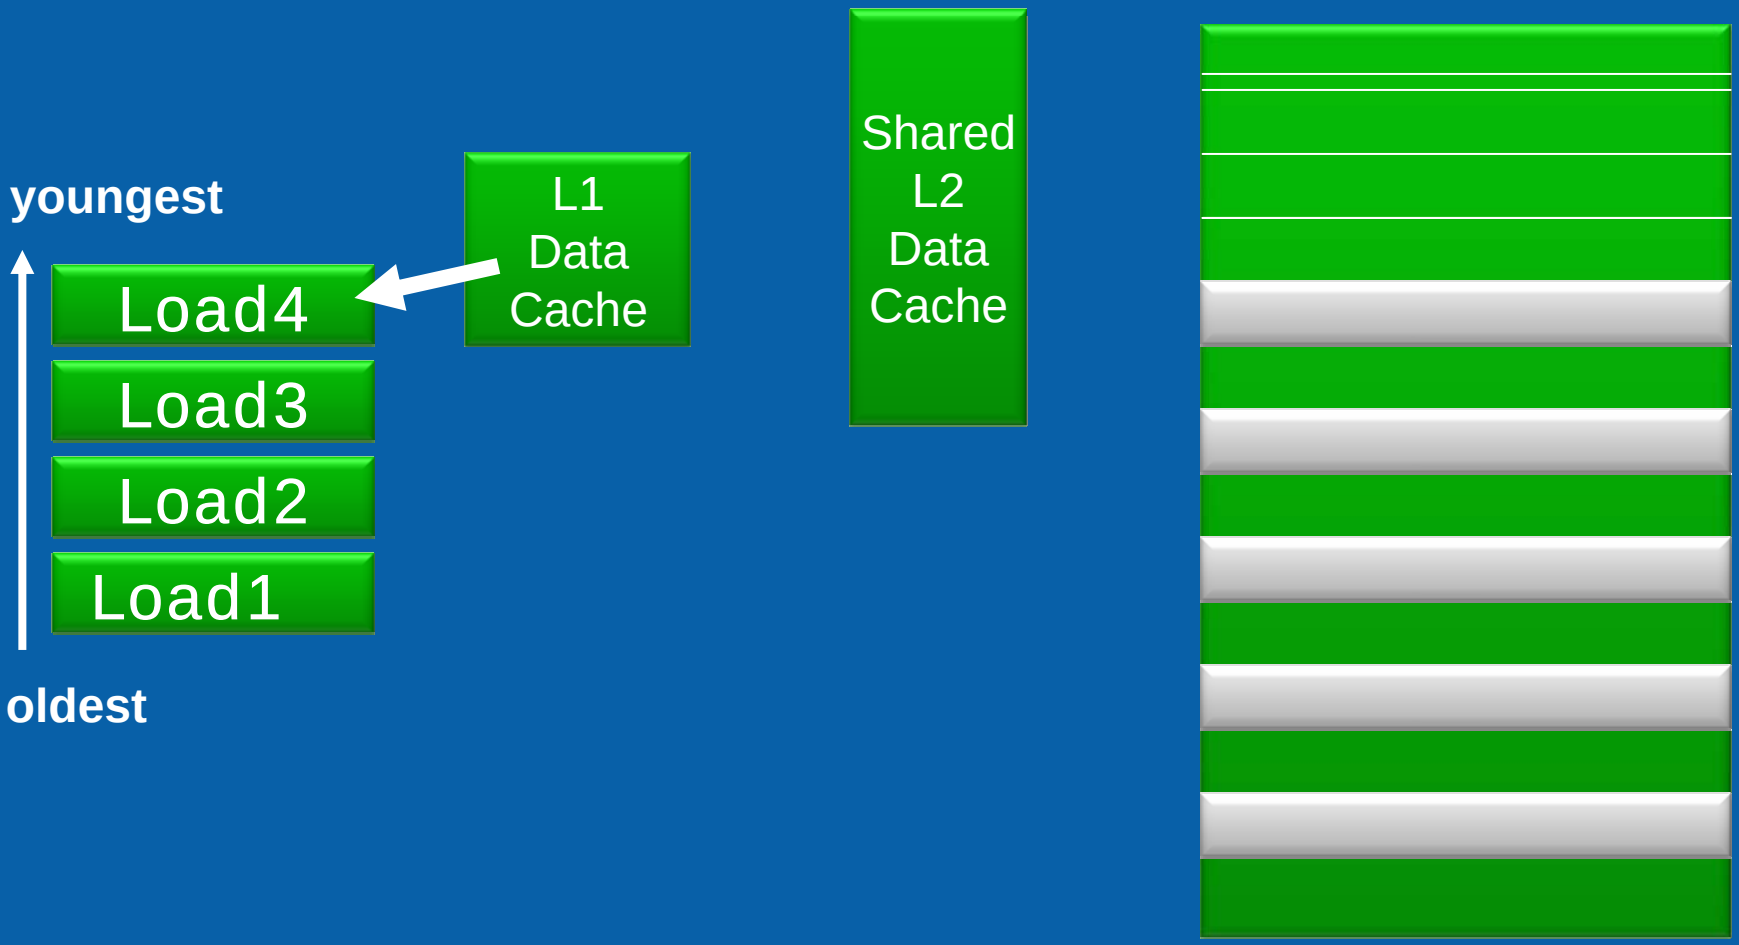

Solving the Problem of Where

Prefetchers and Multi Core

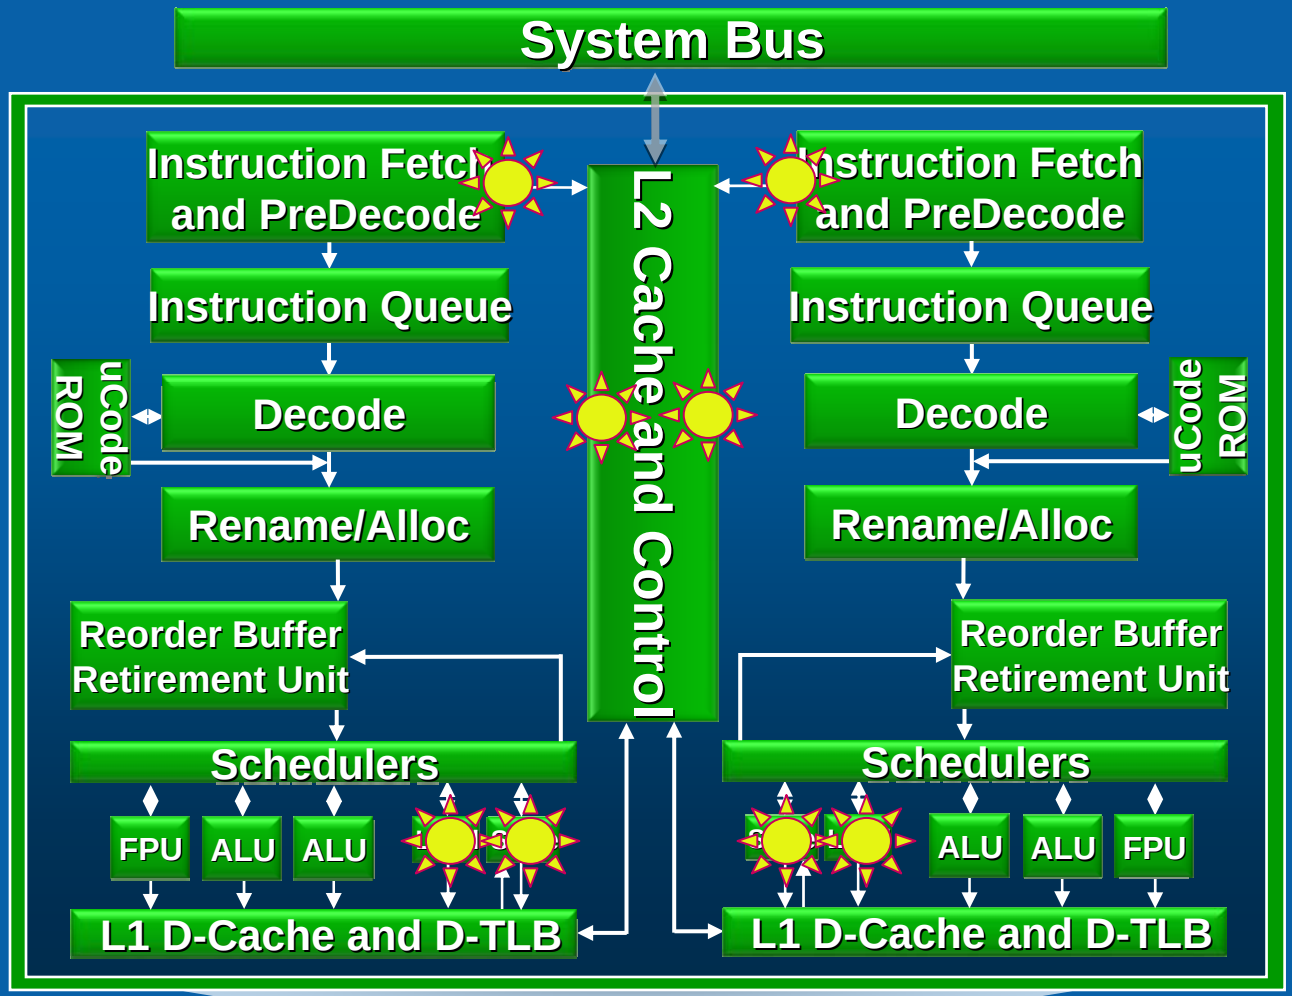

2 data and 1 instruction prefetcher per core

→ able to handle multiple simultaneous patterns

2 prefetchers in the L2 cache

→ tracking multiple patterns per core

Prefetchers monitor demand traffic and regulate "aggression"

Implementation "knobs" allow platform and segment specific settings tailored to applications and usage models

8 Prefetchers per two-core processor

Smart Memory Access: Memory Disambiguation

- Memory Disambiguation predictor
  - Loads that are predicted NOT to forward from preceding store are allowed to schedule as early as possible
    - increasing the performance of OOO memory pipelines
- Disambiguated loads checked at retirement
  - Extension to existing coherency mechanism
  - Invisible to software and system

Solving the Problem of When

Smart Memory Access: Memory Disambiguation

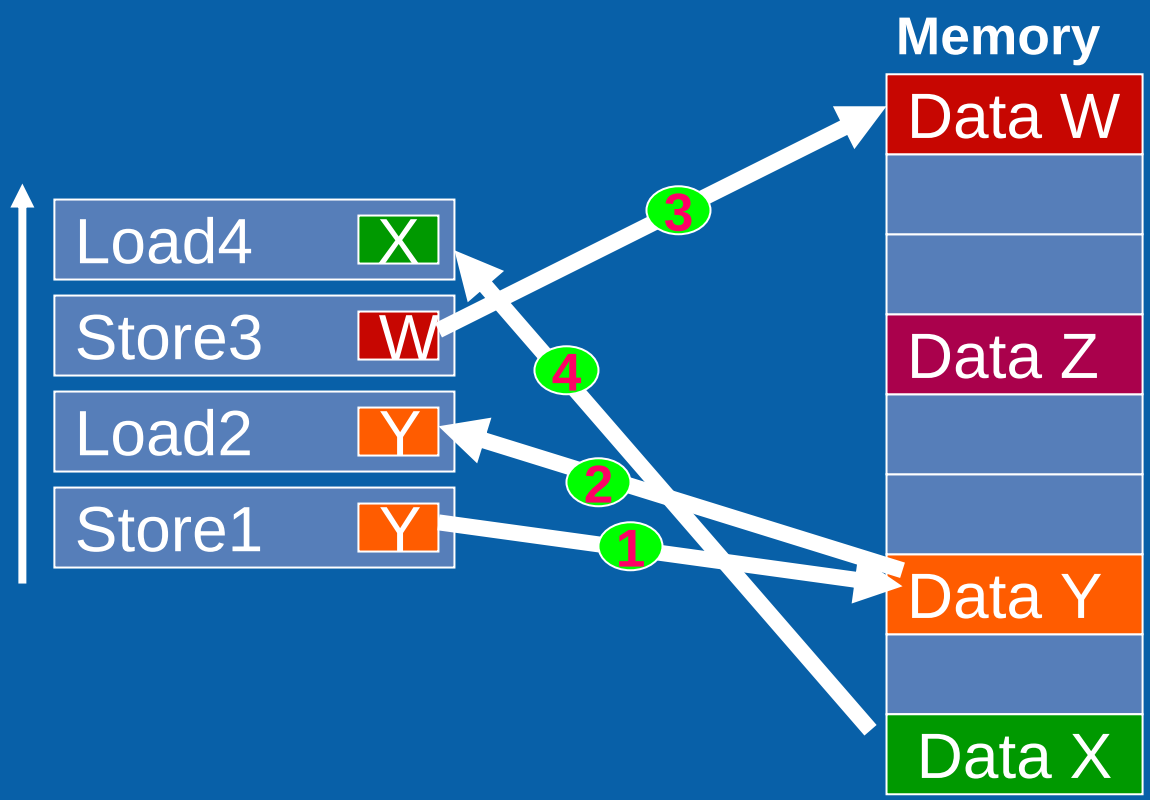

Load4 must WAIT until previous stores complete

Waits for Data X before can execute

Without Memory Disambiguation

Subsequent Loads Must Wait

Smart Memory Access: Memory Disambiguation

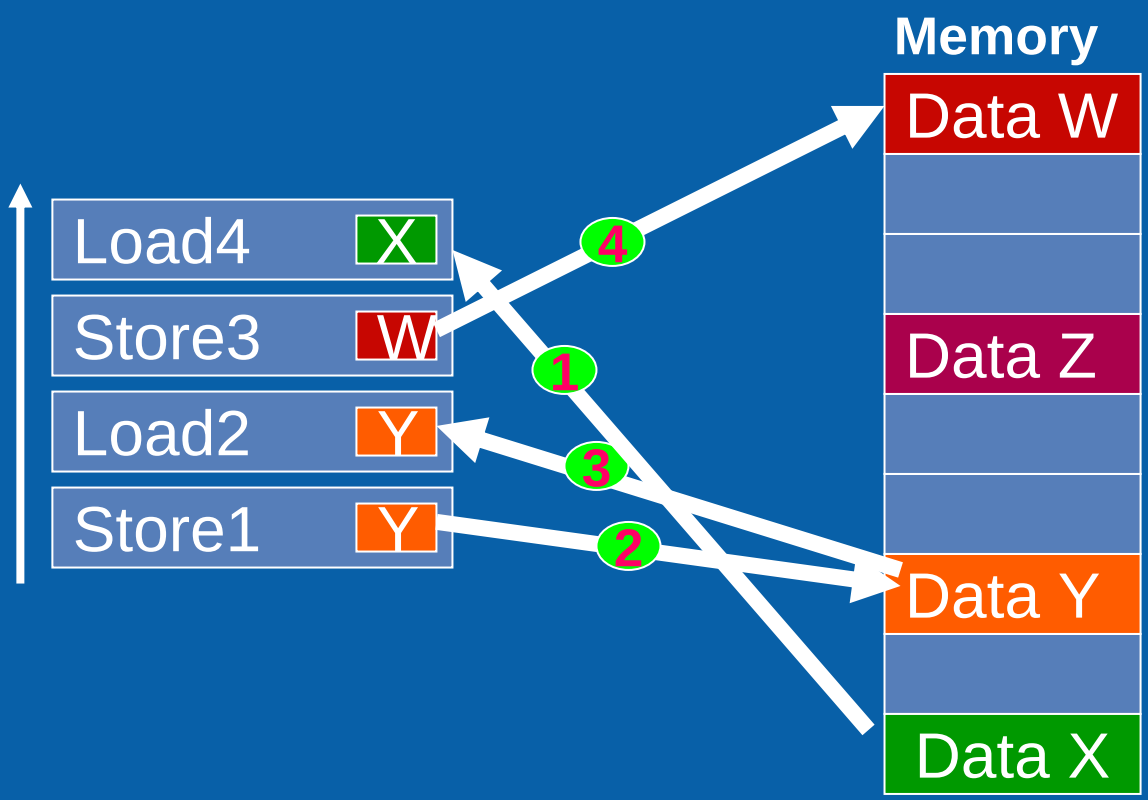

Load4 can get its data FIRST

WITH Intel's new Memory Disambiguation

Loads can decouple from Stores

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - Smart Memory Access
  - **Enhanced Smart Cache**
  - Intelligent Power Capability
- Software Impact
- Benchmark Results
- 45nm Technology

Efficient Data Sharing

Intel® Advanced Smart Cache

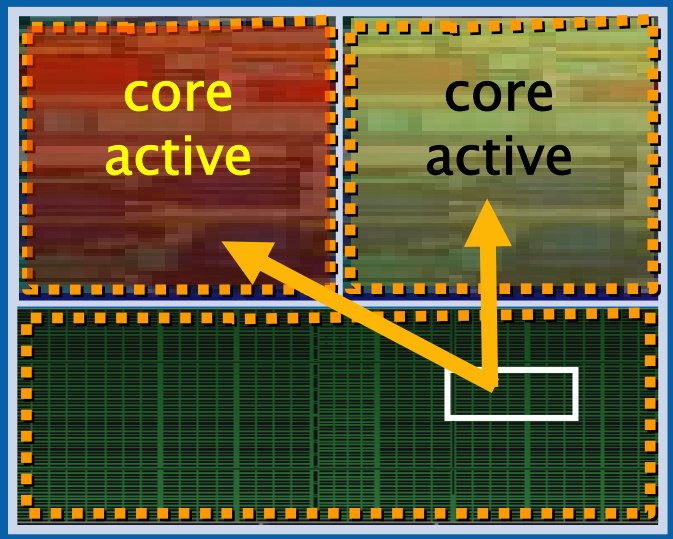

Shared L2 minimizes front side bus traffic and reduces control logic complexity

Independent L2

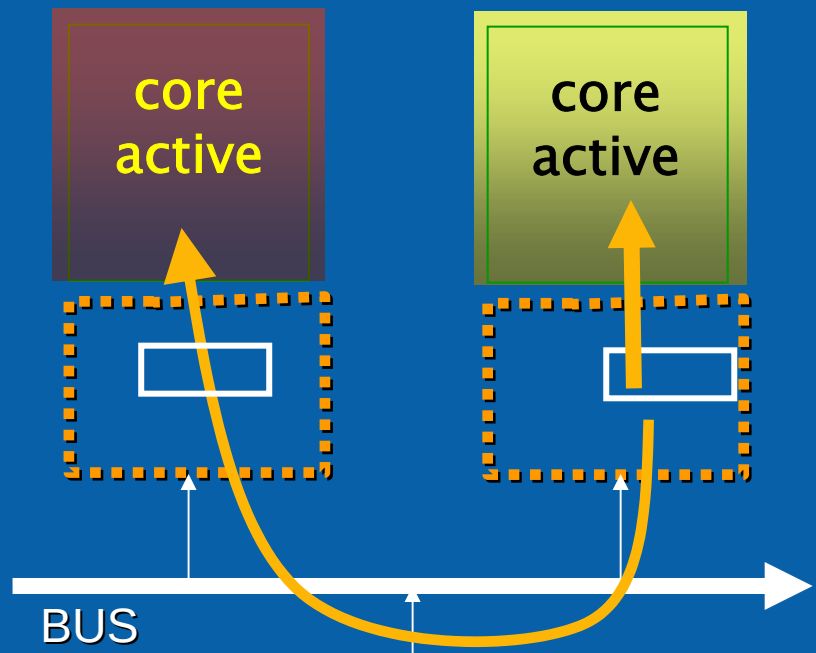

Performance loss and increase in power

*Graphics not representative of actual die photo or relative size

Additional Parallel Design Features

- Intel® Advanced Smart Cache
  - Design features such as shared cache can share data between cores increasing performance over 2 single core CPU's
  - One core can use more of the cache than the second core according to the workload

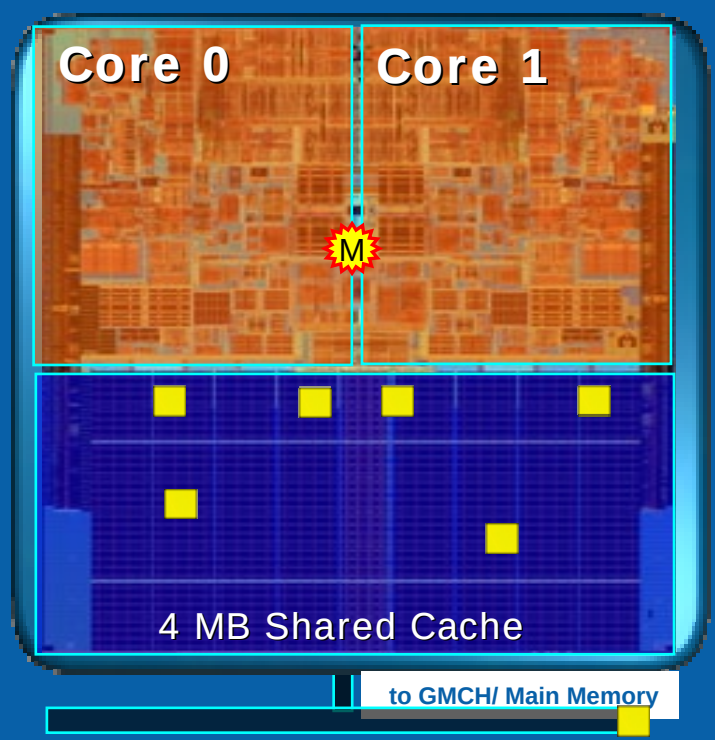

Dynamic Cache Allocation

Intel® Advanced Smart Cache

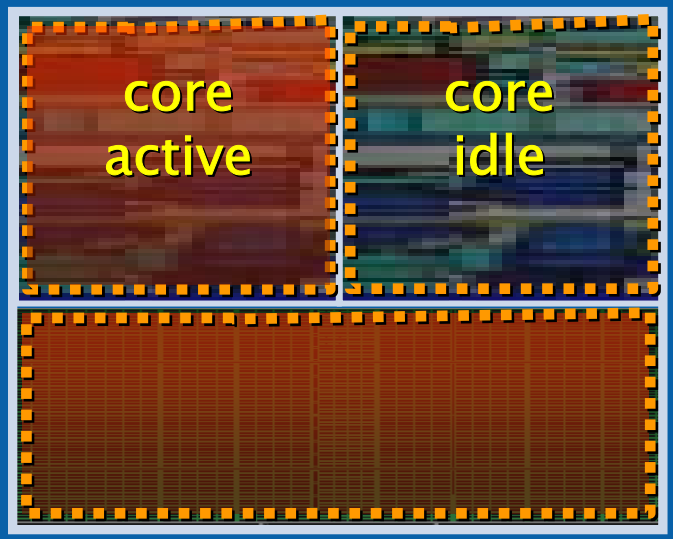

Shared L2 enables active execution core to access full cache when one other execution core is idle

Independent L2

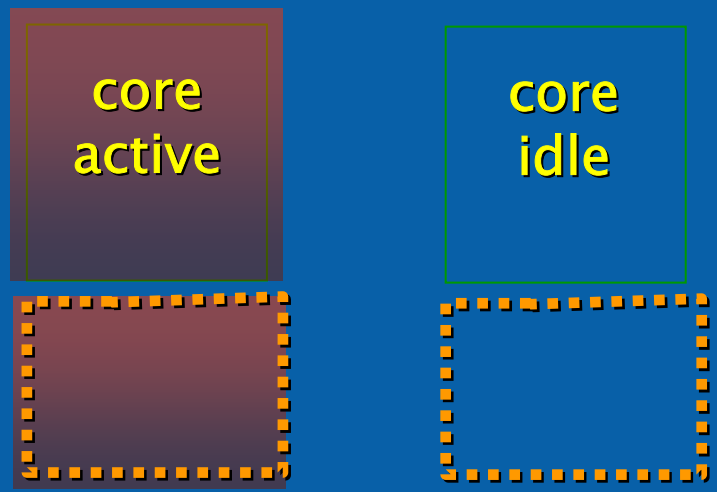

Performance loss due to increased cache misses

Cache not utilized

*Graphics not representative of actual die photo or relative size

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - Smart Memory Access
  - Enhanced Smart Cache
  - **Intelligent Power Capability**
- Software Impact
- Benchmark Results
- 45nm Technology

Intelligent Power Capability

Extending the power management architecture

- Intel® Pentium® M processor innovated a new power management architecture
- Intel® Core™ Duo processor extended the Pentium® M processor capability to multi-core

New Power Features within each processor core

- Ultra fine-grained power control

Enhancing Energy Efficiency

Intel® Dynamic Power Coordination

Shared power management logic to coordinate Enhanced Intel® Speedstep Technology (p-state) and idle power management state (C-state) transitions to efficiently manage voltage and frequency

Supports Intel® Enhanced Deeper Sleep mode unique to mobile

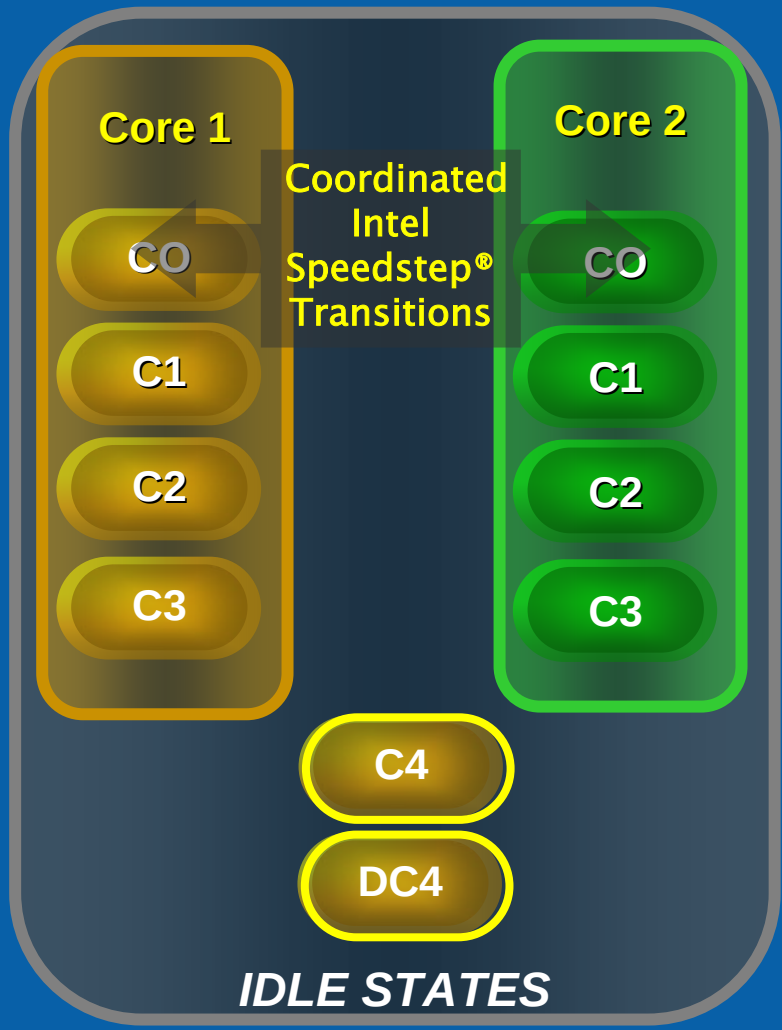

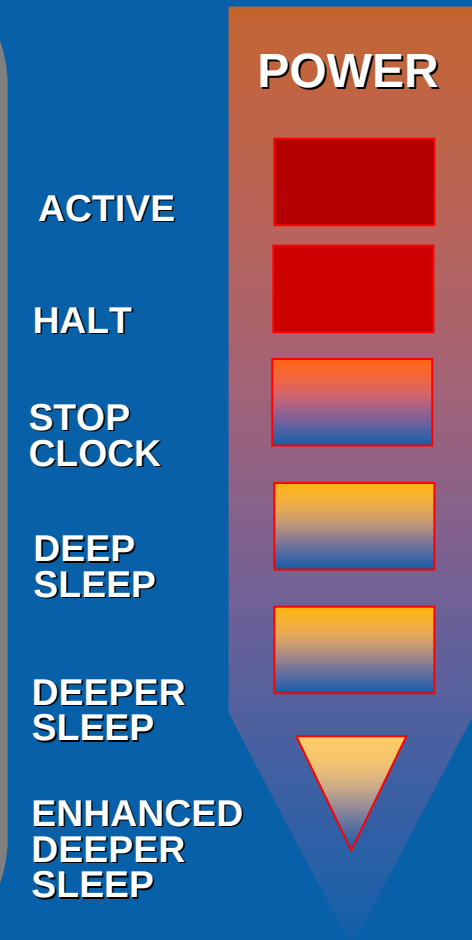

Power Efficient Dual-Core Performance on Demand

intel

Ultra Fine Grained Power Control

Even during periods of high performance execution, many parts of the chip core can be shut off.

Example could be a SW memory initialization executing from front end with IQ operating as loop cache.

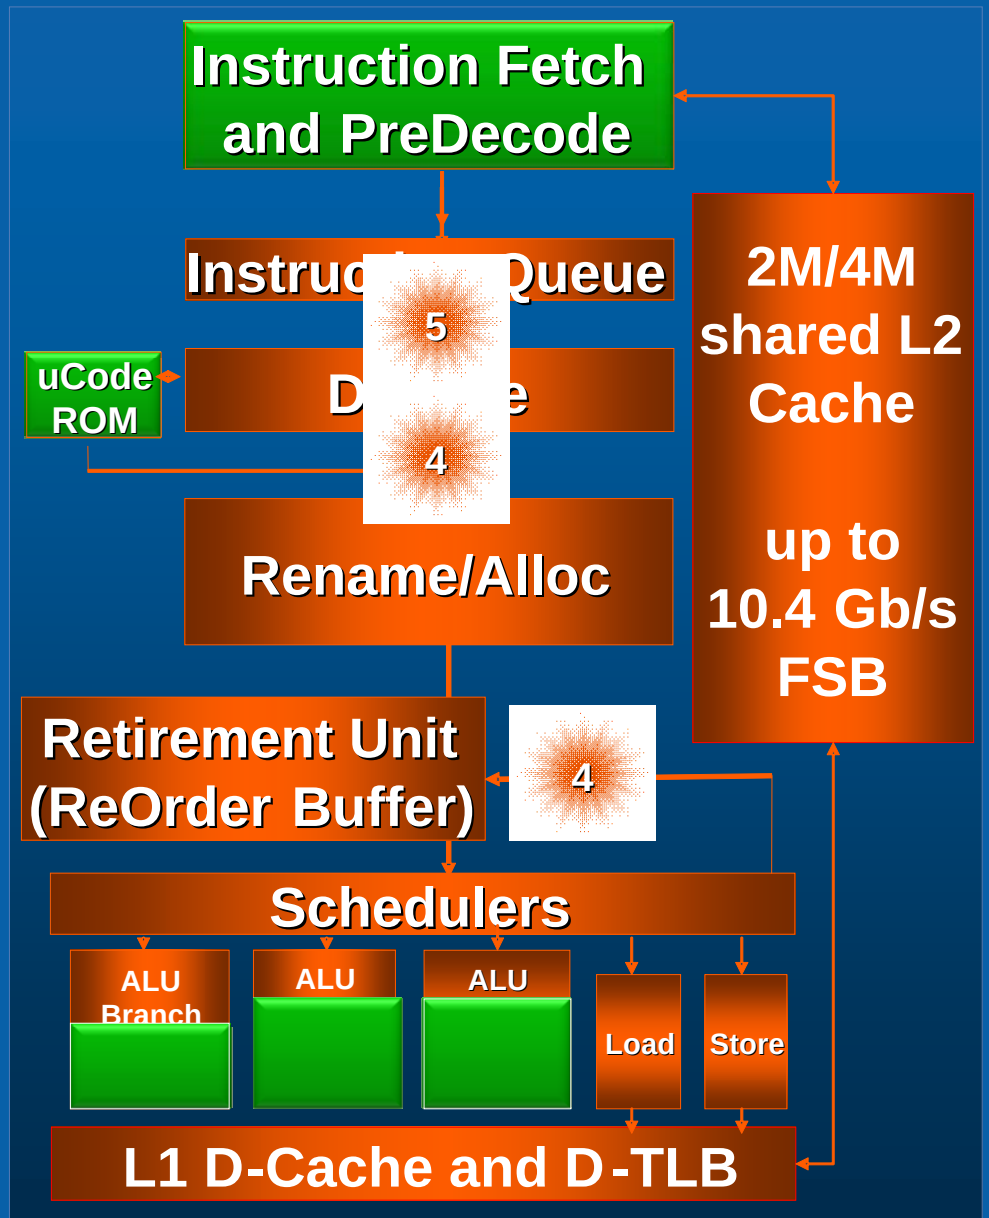

Unit OFF

Unit ON

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - Smart Memory Access
  - Enhanced Smart Cache
  - Intelligent Power Capability
- **Software Impact**
- Benchmark Results
- 45nm Technology

Software Impact

- "Wide Dynamic execution" & "Digital Media Boost"
  - Deliver additional performance with no impact or changes needed for software developers
- "Advanced Smart Cache"
  - Additional gains can be achieved by intelligent scheduling via software developer.
    - Two threads which "communicate" frequently should be scheduled to same two cores sharing L2 cache.

Software Impact

- "Smart Memory Access"
  - Delivers additional performance out of the box.
  - Further performance can be gained by use of software pre-fetching and BIOS tuning.
- "Intelligent power capability"
  - Delivers performance/watt advantage with no impact or changes needed for software developers

Optimizing for Intel® Core™ Microarchitecture

- Use CMP = employ both Cores
  - Go to multithreading!
- Prefer SSE as much as possible. If you didn't do it so far, vectorize the code now!!
  - Intel Compiler has very good vectorization engine

Optimizing for Intel® Core™ Microarchitecture (advanced)

- Use Intel *VTune™ Performance Analyzer* for performance problems revealing
  - CPI
  - Specific CPU events for Core-arch:

    RESOURCE_STALLS.RS_FULL,
    L2_IFETCH.SELF.MESI,
    RESOURCE_STALLS.RS_FULL,
    RESOURCE_STALLS.ROB_FULL etc- see VTune help

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - Smart Memory Access
  - Enhanced Smart Cache
  - Intelligent Power Capability
- Software Impact
- **Benchmark Results**
- 45nm Technology

Real Application Comparison

System Power

- **Power measured is total system watts during the benchmark run**

Lower is better

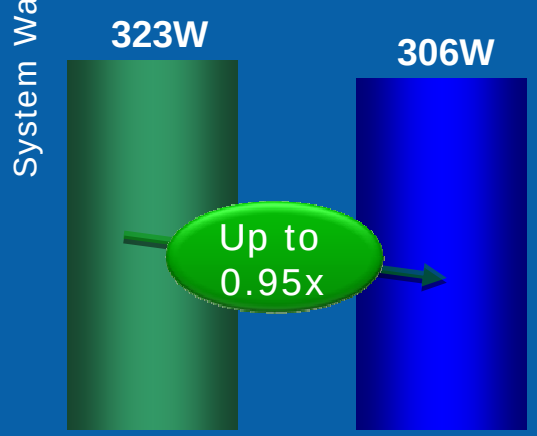

Sun X4200* 2P Opteron* 2.4 GHz

HP DL380G5* 2P Woodcrest 3.0 GHz

Performance

SunGard ACR* financial application

Higher is better

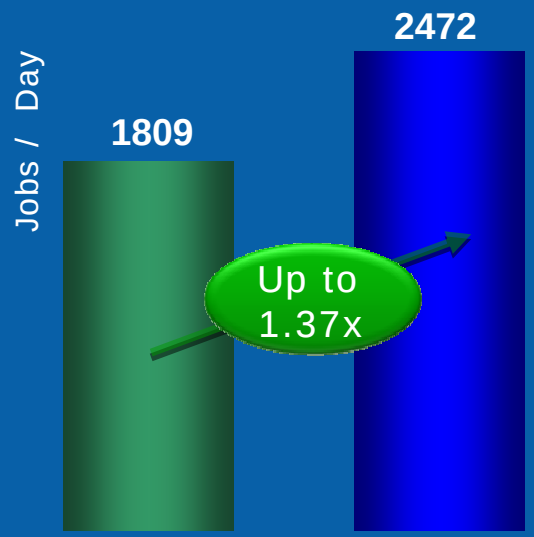

Sun X4200* 2P Opteron* 2.4 GHz

HP DL380G5* 2P Woodcrest 3.0 GHz

Performance/ Watt

Higher is better

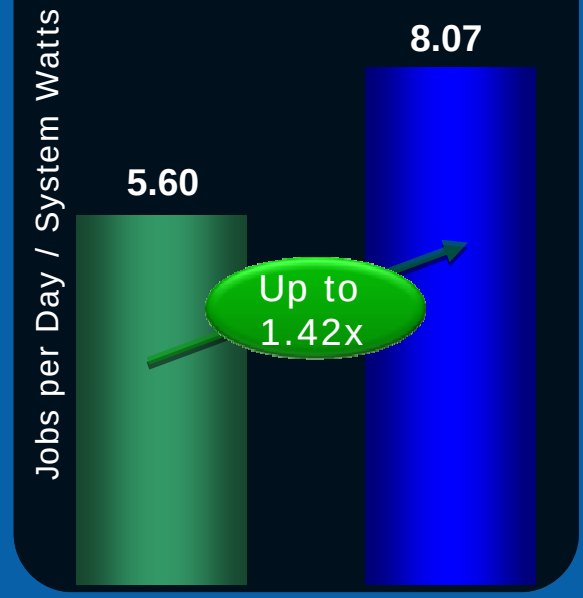

Sun X4200* 2P Opteron* 2.4 GHz

HP DL380G5* 2P Woodcrest 3.0 GHz

Woodcrest delivers leading Performance and Performance/ Watt

Bensley/Clovertown Platform Performance

Competitive Landscape based on Published AMD results

SPECfp_rate_ base2000

Higher is better

Windows Server

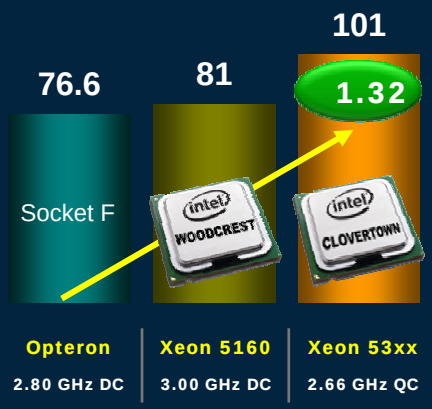

SAP*-SD

Higher is better

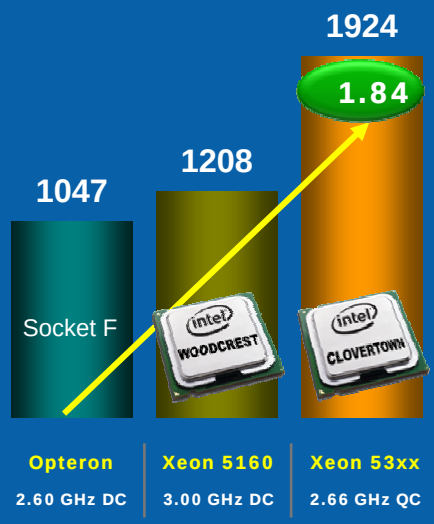

SPECjbb*2005

Higher is better

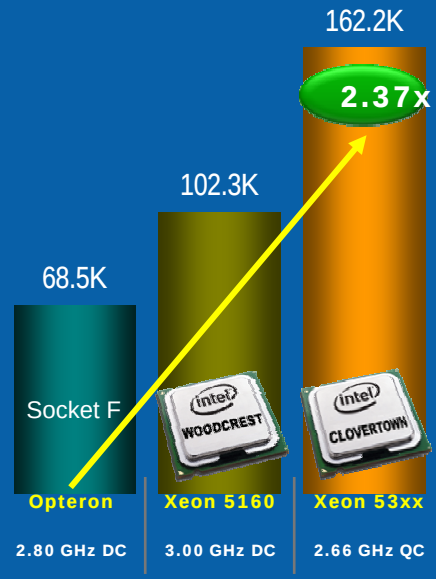

SPECint*_rate_ base2000

Higher is better

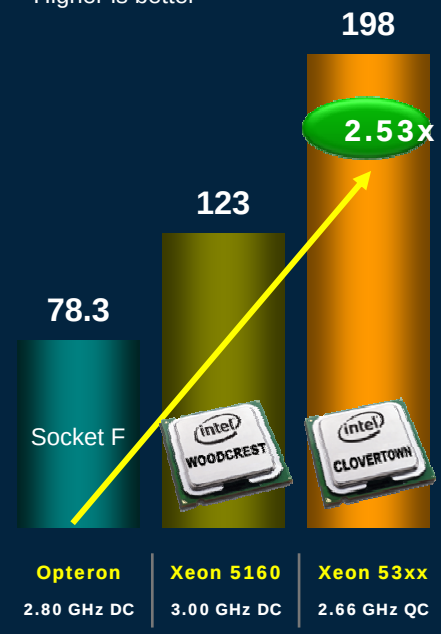

Data Source: AMD data based on Published results as of September 19, 2006. Intel data based on published and measured data. See backup for details

Clovertown extends Bensley platform leadership

Xeon 53xx – Quad-Core Intel® Xeon® Processor 53xx; ("Clovertown 2.67 GHz");
Xeon 5160 – Dual-Core Intel® Xeon® Processor 5160; ("Woodcrest 3.00 GHz");

Opteron 2.80 DC – Dual-Core AMD Opteron™ Model 2220 (2.80 GHz);

Core™ Architecture

- Introduction
- New Core™ Advanced Features:
  - Enhanced Execution Capabilities
  - Smart Memory Access
  - Enhanced Smart Cache
  - Intelligent Power Capability
- Software Impact
- Benchmark Results
- **45nm Technology**

45nm Technology

- Penryn – code name for an enhanced Intel® Core™ microarchitecture at 45 nm
  - Industry's first 45 nm High-K processor technology
  - ~2x transistor density
  - >20% gain in transistor switching speed
  - ~30% decrease in transistor switching power
  - Dual core, quad core
  - Shared L2 cache
  - Intel 64 architecture
  - 128-bit SSE

"Penryn"/"Wolfdale"/"Wolfdale DP"
Dual Core Package

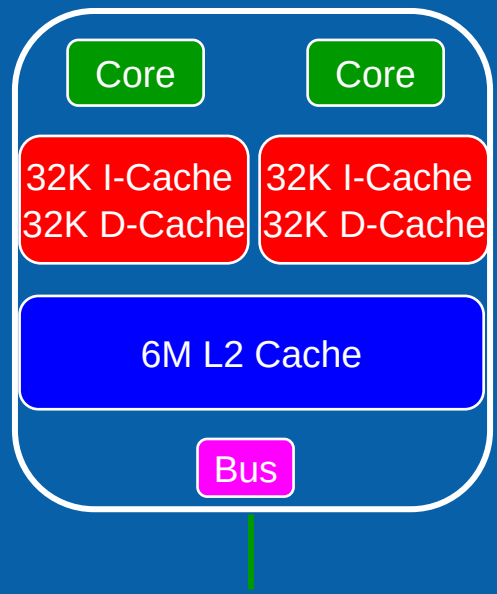

2 Threads, 1 Package
(similar to Intel® Core™ 2 Duo processor)

Penryn Family Optimized Microarchitecture

- Intel SSE4 instructions
  Fast Super Shuffle Engine
- Fast Radix-16 Divider Enhanced
  Intel Virtualization Technology
- Larger Caches: 6MB, 12MB
  24-way Set Associativity
- Split Load Cache Enhancement
  Higher Bus Speeds
- Deep Power Down Technology
  Enhanced Intel Dynamic Acceleration Tech

> 3 GHz

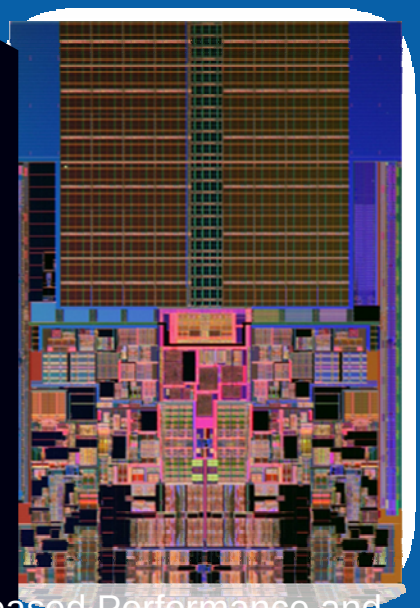

Increased Performance and
Energy Efficiency

intel

Scaleable Across the Enterprise

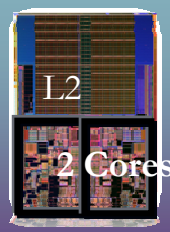

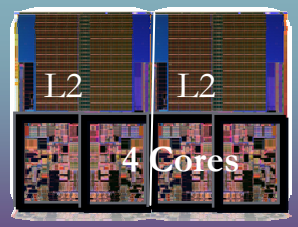

Server / Workstation

Intel Xeon®
processor family
UP, DP, MP

>3GHz
Up to 1600MHz Bus
Quad-Core – up to 12MB L2
50/80W/120W TDP
Dual-Core – up to 6MB L2
40W/65W/80W TDP

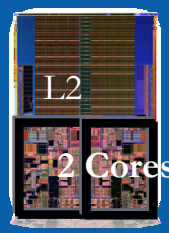

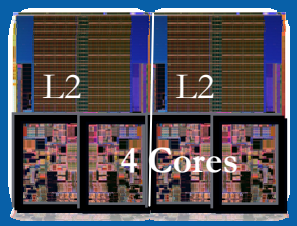

Desktop

Intel Core™ 2 Extreme

Intel Core 2 Processor Family

>3GHz
Quad-Core – up to 12MB L2
95W/130W TDP
Dual-Core – up to 6MB L2
65W TDP

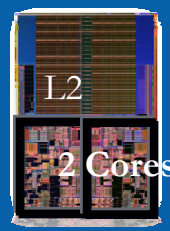

Mobile

Intel Core 2 Processor Family

Dual-Core-up to 6MB L2

Deep Power Down Tech

Enhanced Dynamic Acceleration Technology

intel

Summary

- The new core™ architecture has advanced features which give performance advantage to your software
- This new architecture has a clear performance leadership across a broad spectrum of applications and benchmarks
- The new architecture manufactured using the 45nm process provides yet a higher level of performance at lower power consumption