Master Informatics Eng.

2017/18

A.J.Proença

Concepts from undegrad Computer Systems (1)

(most slides are borrowed, mod's in green)

Concepts from undegrad Computer Systems

– most slides are borrowed from

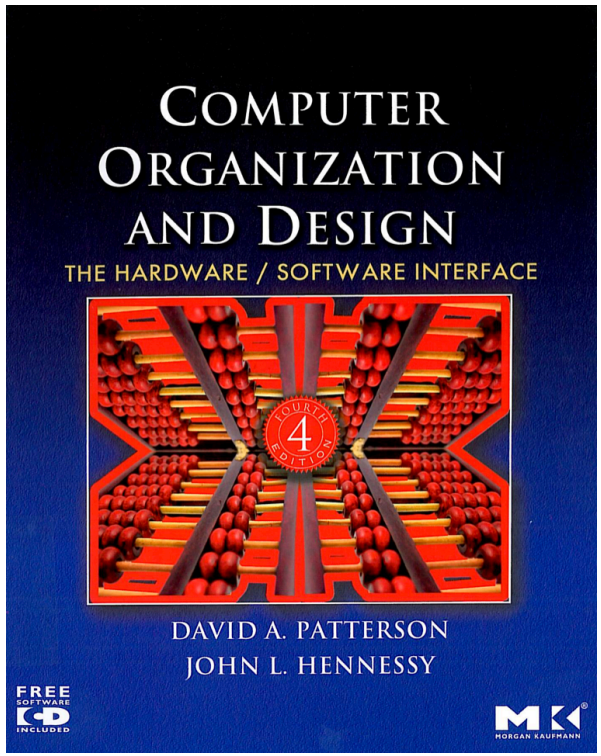

and some from

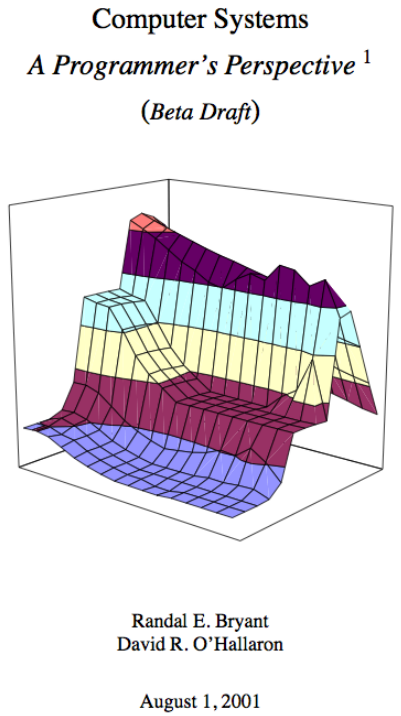

more details at
http://gec.di.uminho.pt/miei/sc/

Background for Advanced Architectures

Key concepts to revise:

- *numerical data representation (for error analysis)*
- *ISA (Instruction Set Architecture)*
- *how C compilers generate code (a look into assembly code)*
  - *how scalar and structured data are allocated*
  - *how control structures are implemented*
  - *how to call/return from function/procedures*
  - *what architecture features impact performance*
- ***Improvements to enhance performance in a single CPU***
  - ***ILP: pipeline, multiple issue, …***
  - ***data parallelism: SIMD/vector processing, ...***
  - ***memory hierarchy: cache levels, ...***
  - ***thread-level parallelism***

As arquiteturas *multicore* mais recentes:

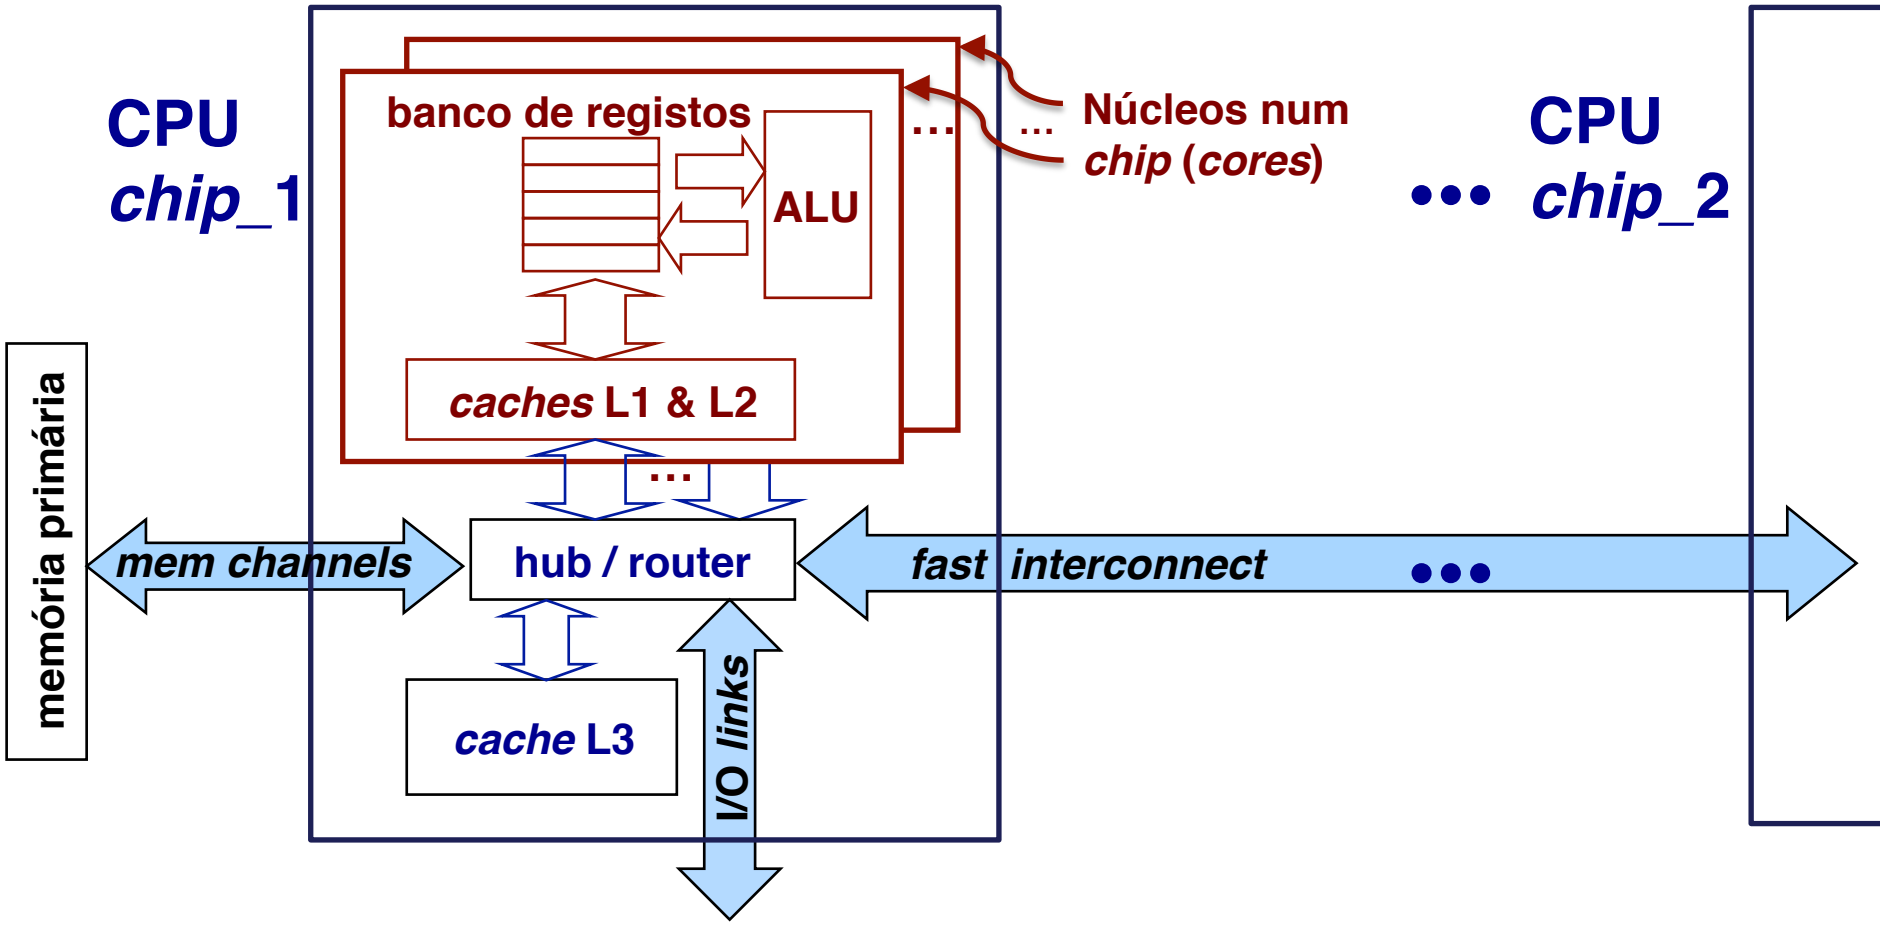

Lançamento da Intel em 2012: Sandy/Ivy Bridge (8-core)

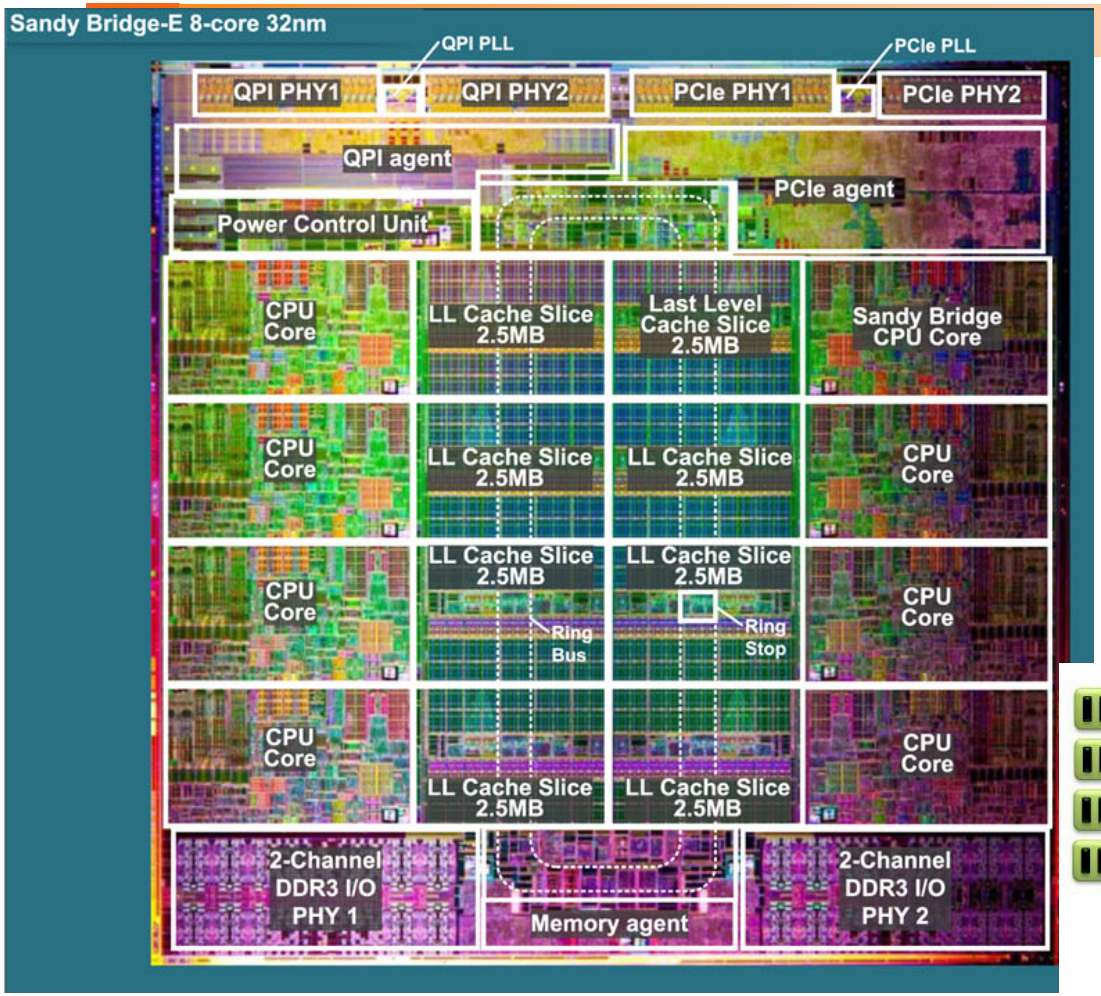

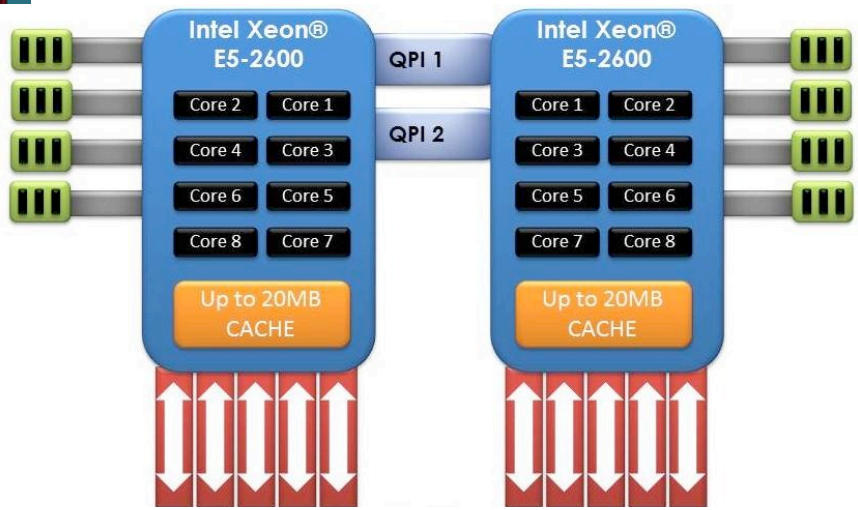

Intel® Xeon® Processor E5 v4 Product Family HCC

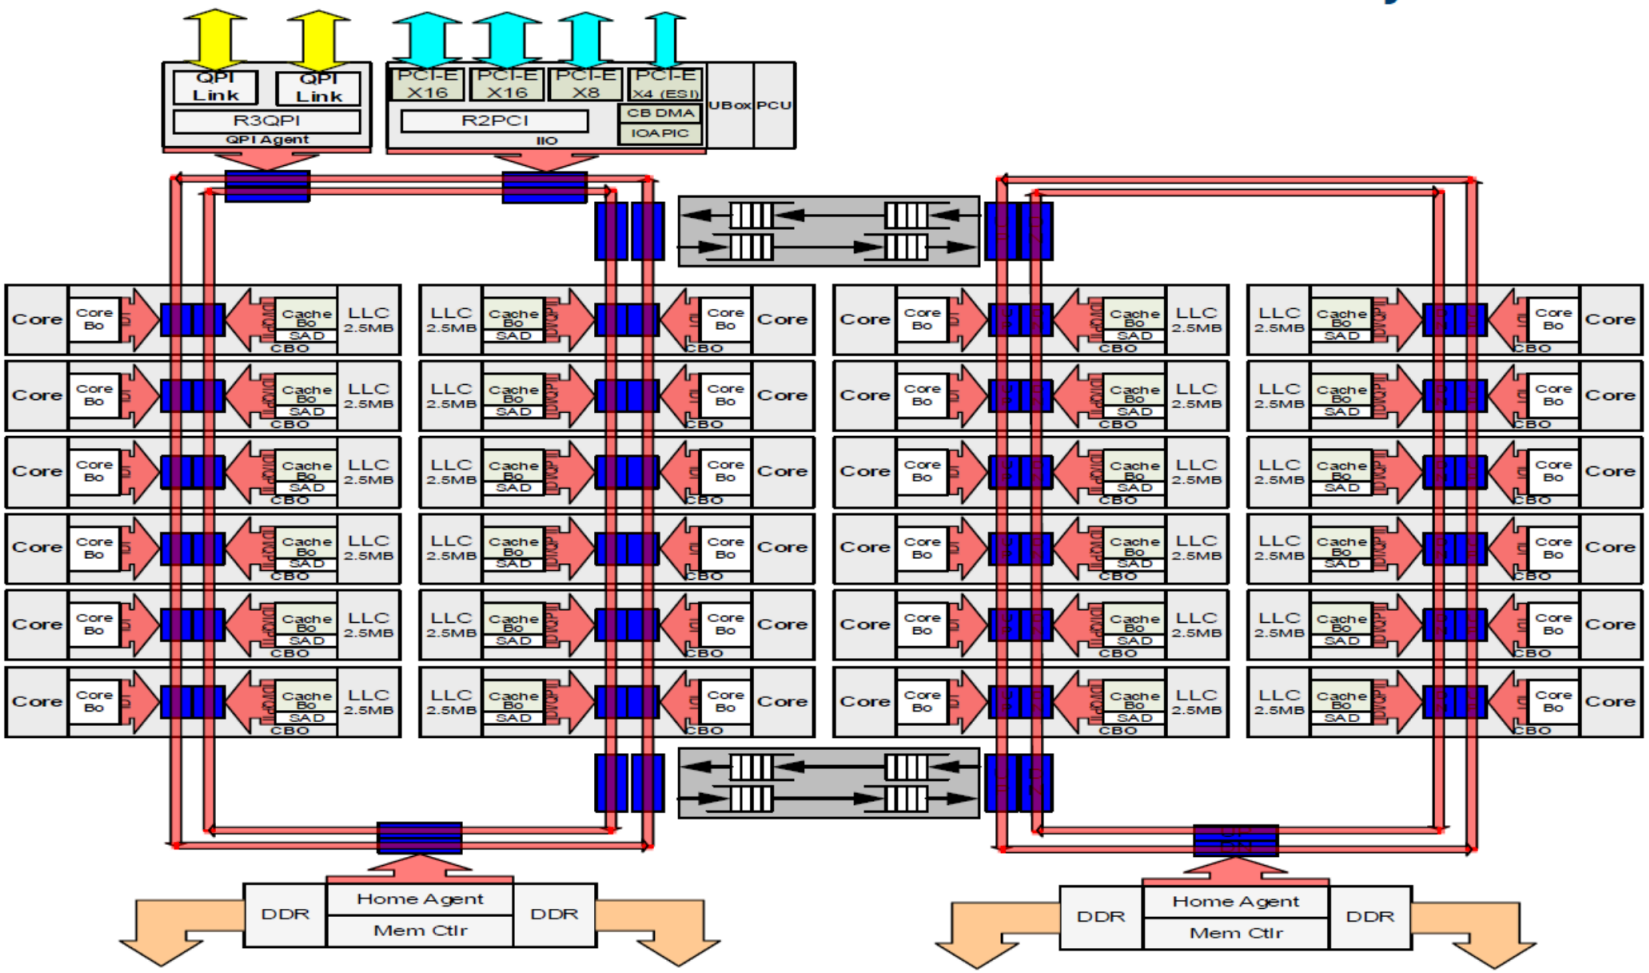

Chips da Intel em 2012/13: Xeon Phi com 60 cores

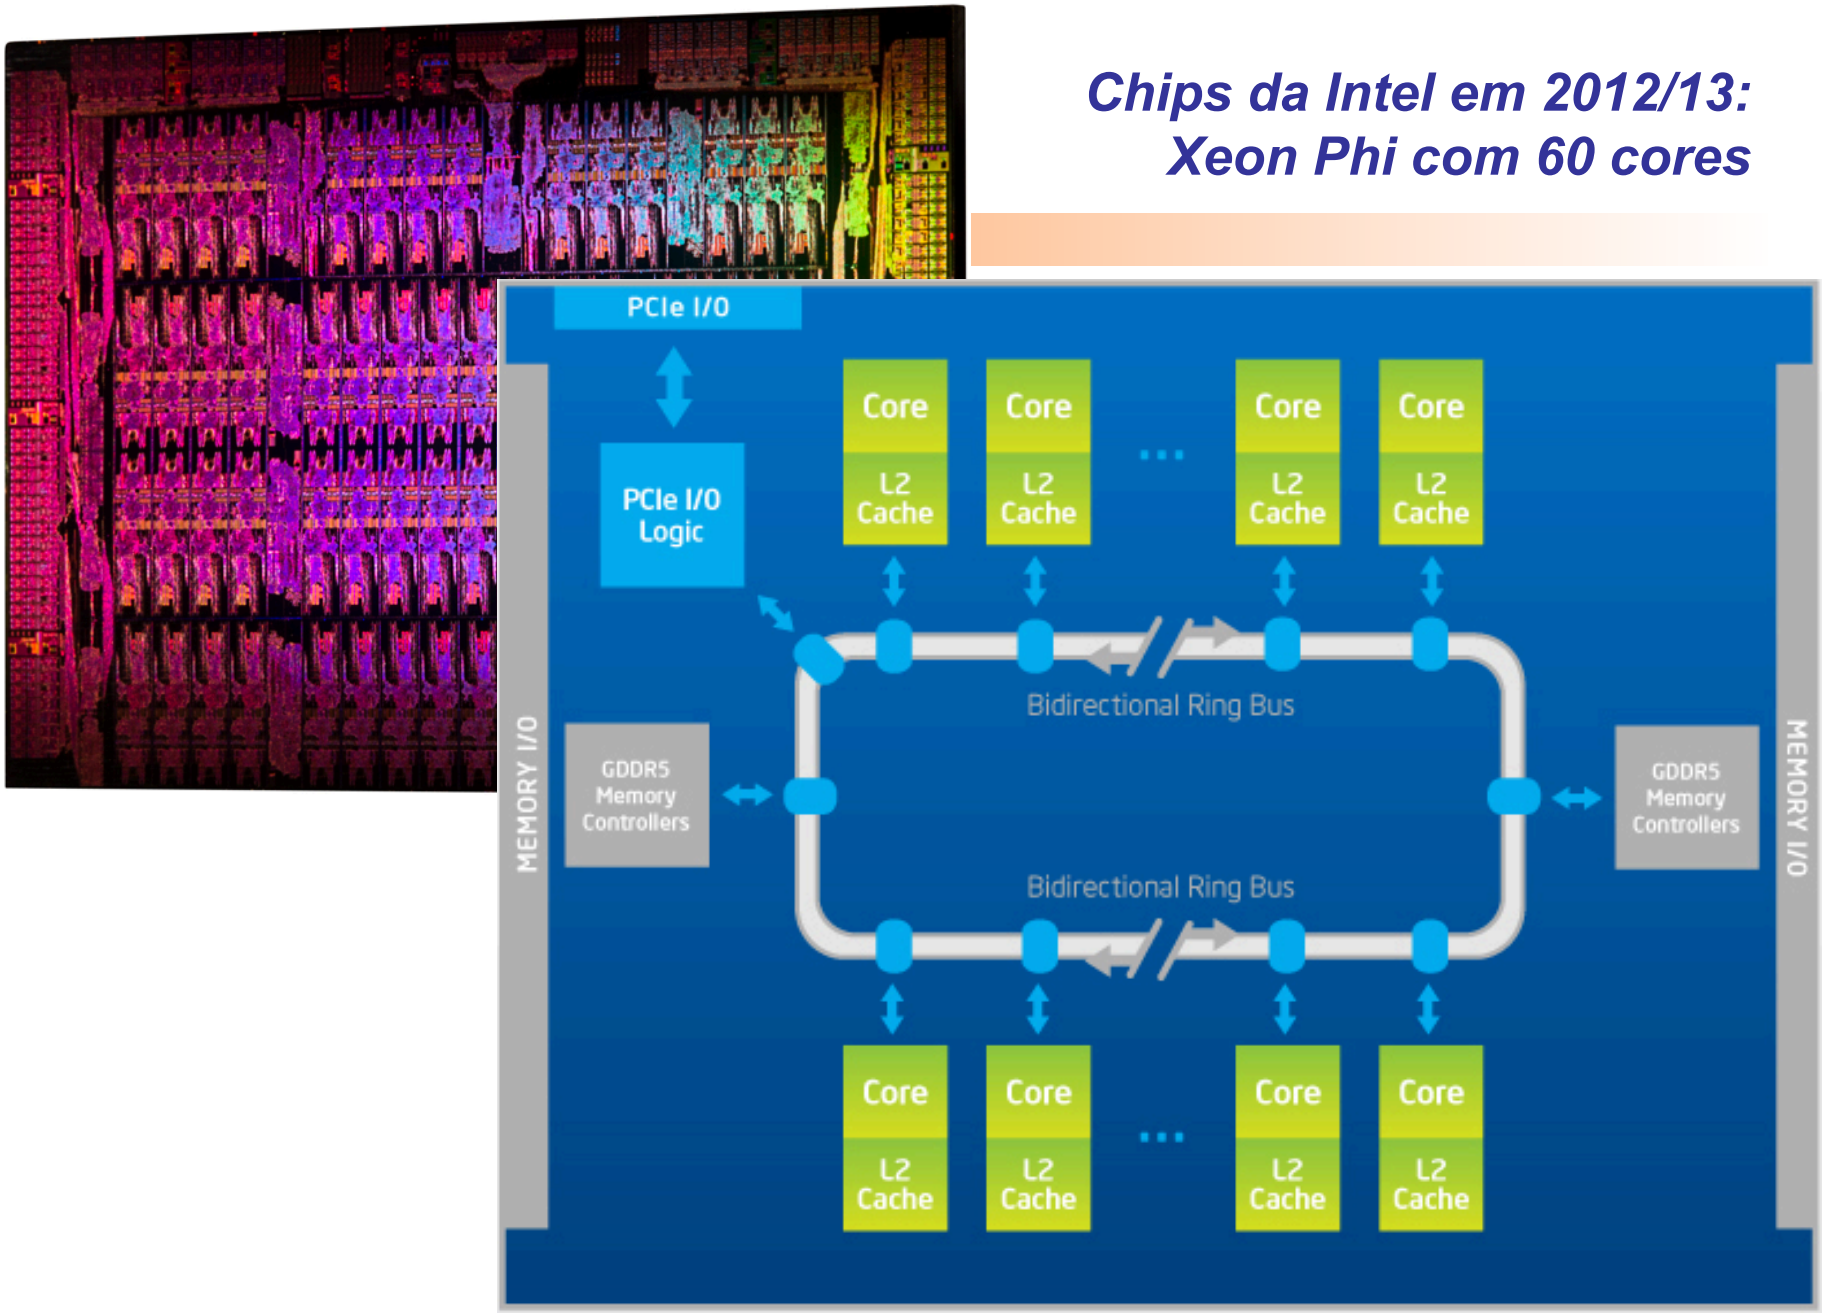

Intel new Phi in 2016: KNL with 72 cores

Knights Landing Overview

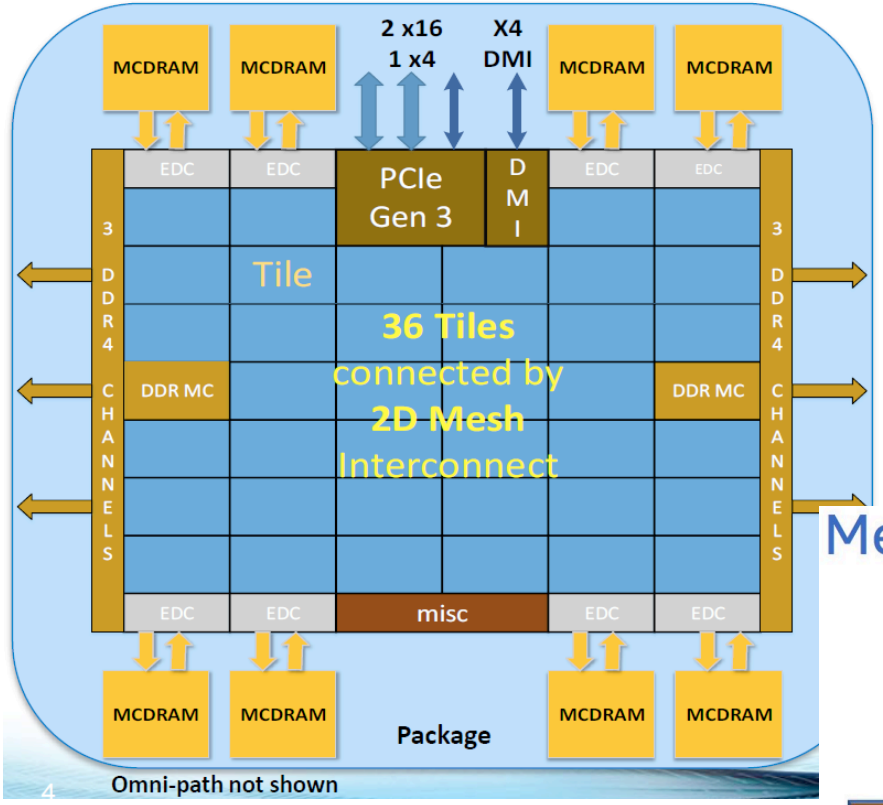

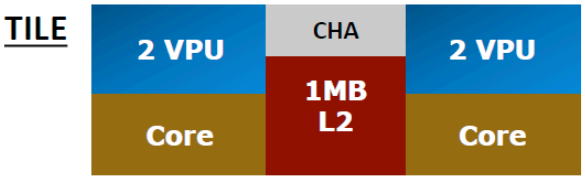

Chip: 36 Tiles interconnected by **2D Mesh**
Tile: 2 Cores + 2 VPU/core + 1 MB L2

Memory: MCDRAM: 16 GB on-package; High BW
DDR4: 6 channels @ 2400 up to 384GB
IO: 36 lanes PCIe Gen3. 4 lanes of DMI for chipset
Node: 1-Socket only
Fabric: Omni-Path on-package (not shown)

Memory Modes

Three Modes. Selected at boot

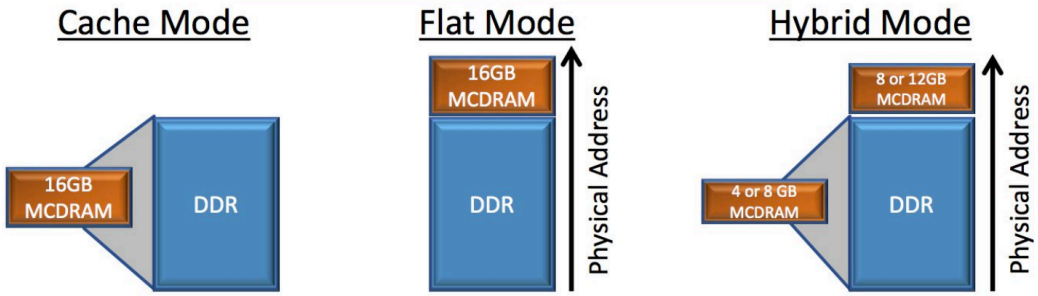

Exemplo de chip com processadores RISC: 2x ARM's no A6 do iPhone 5

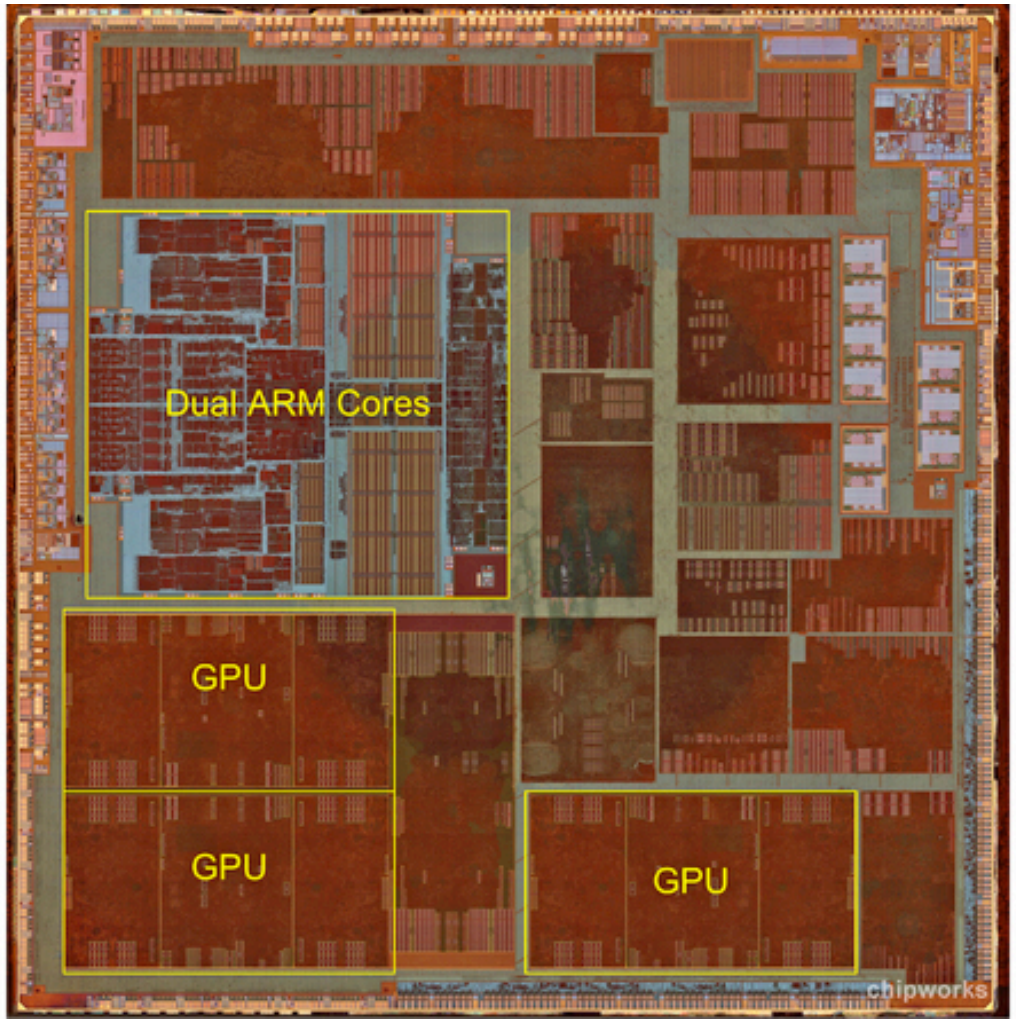

Exemplo de chip com processadores RISC: 4+1 ARM's no Tegra 4i da NVidia

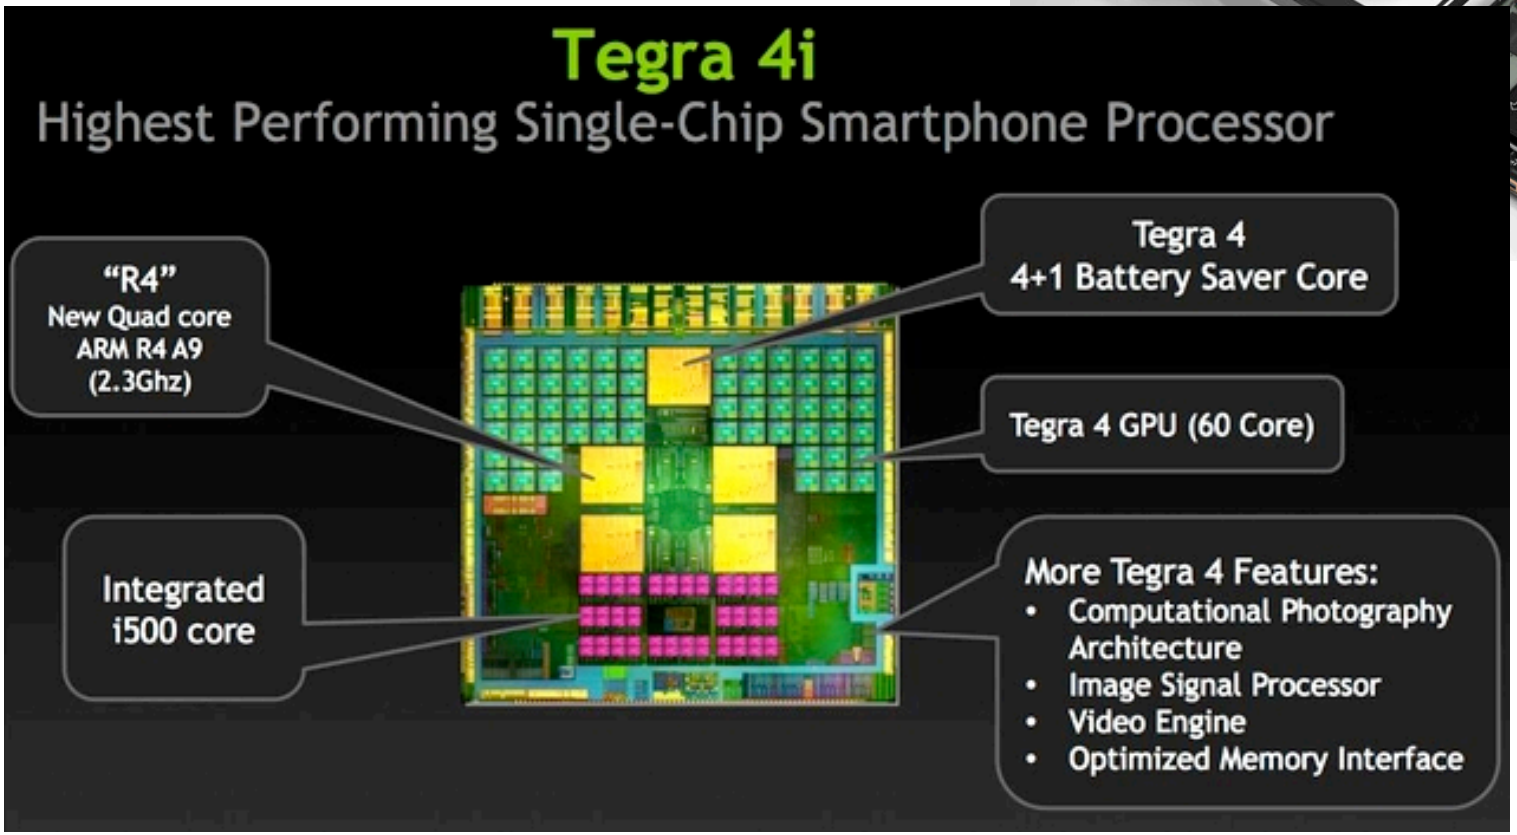

Exemplo de chip com processadores RISC: 4+4 ARM's no Exynos 5 Octa, Galaxy S 4

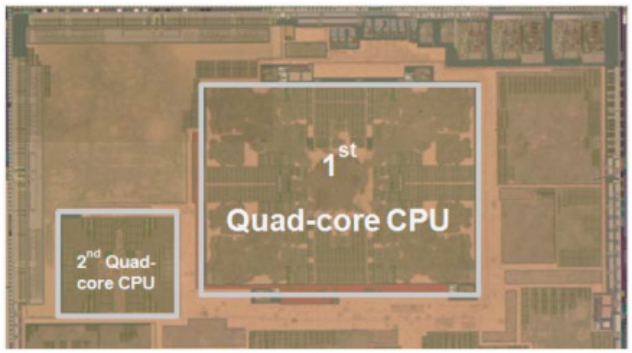

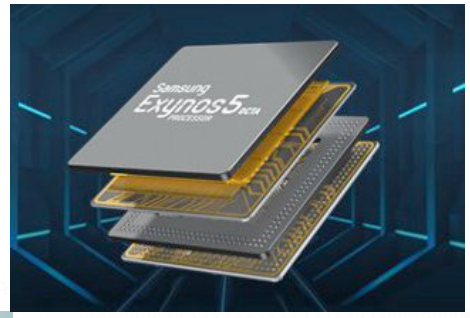

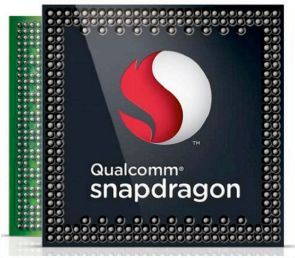

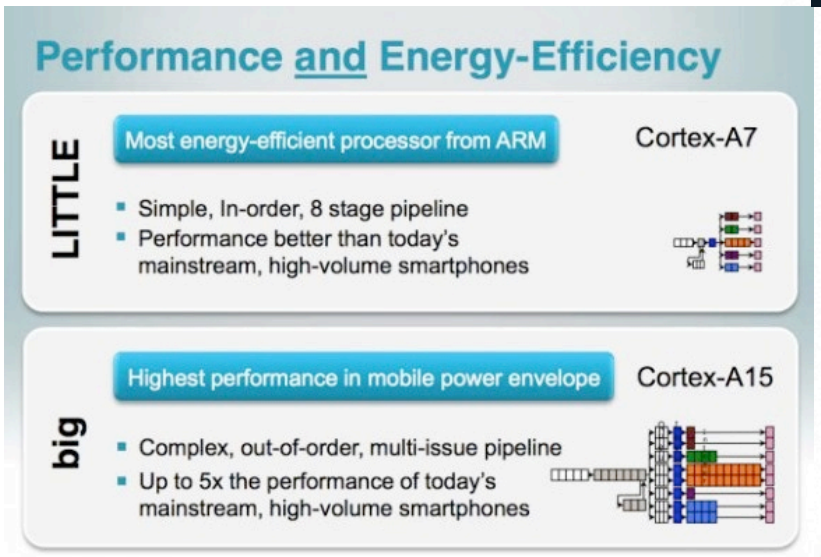

Processadores Intel x86 versus ARM

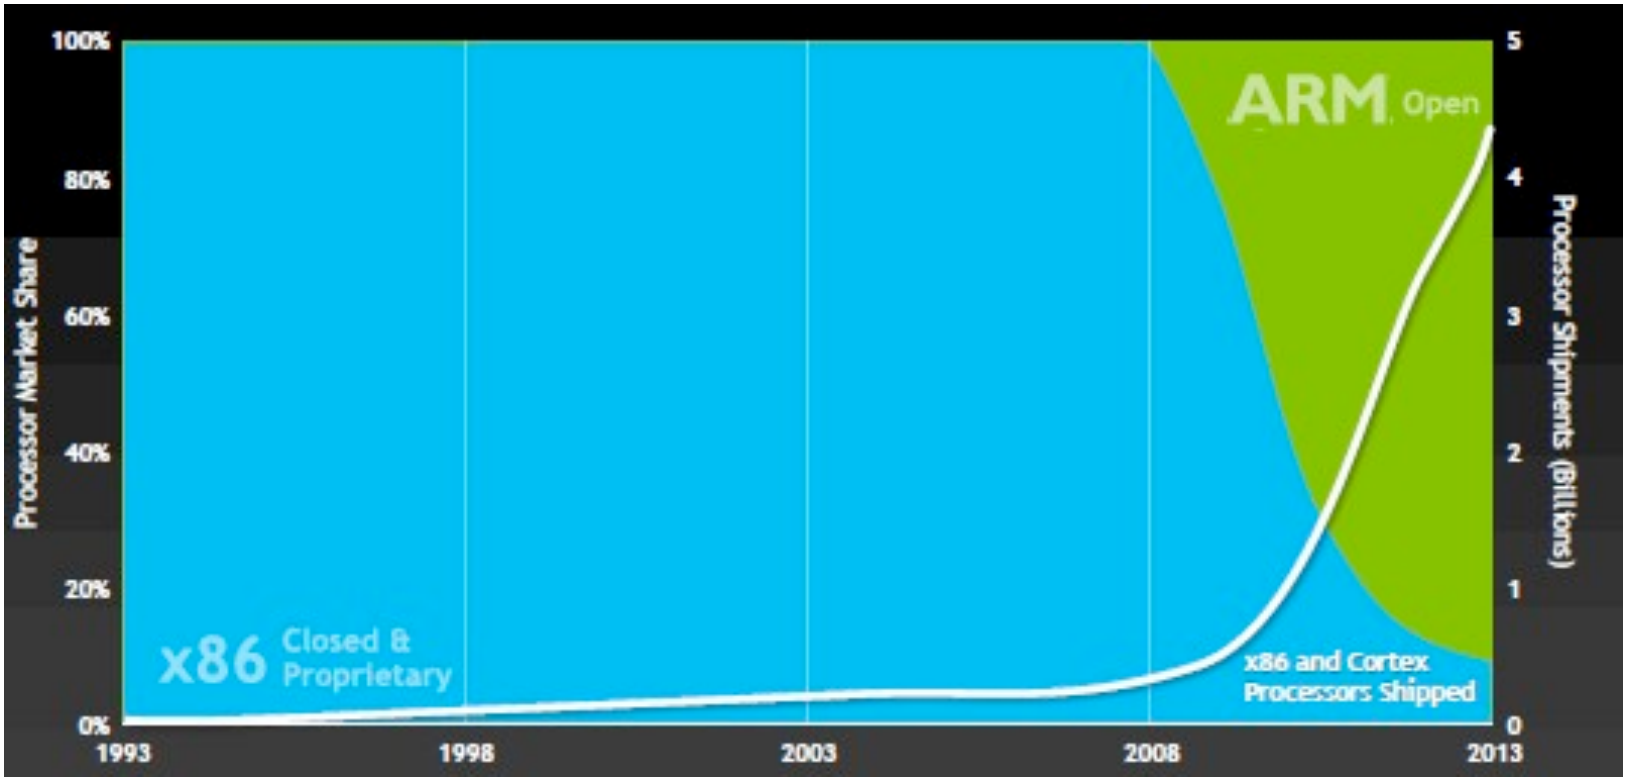

Key textbook for AA

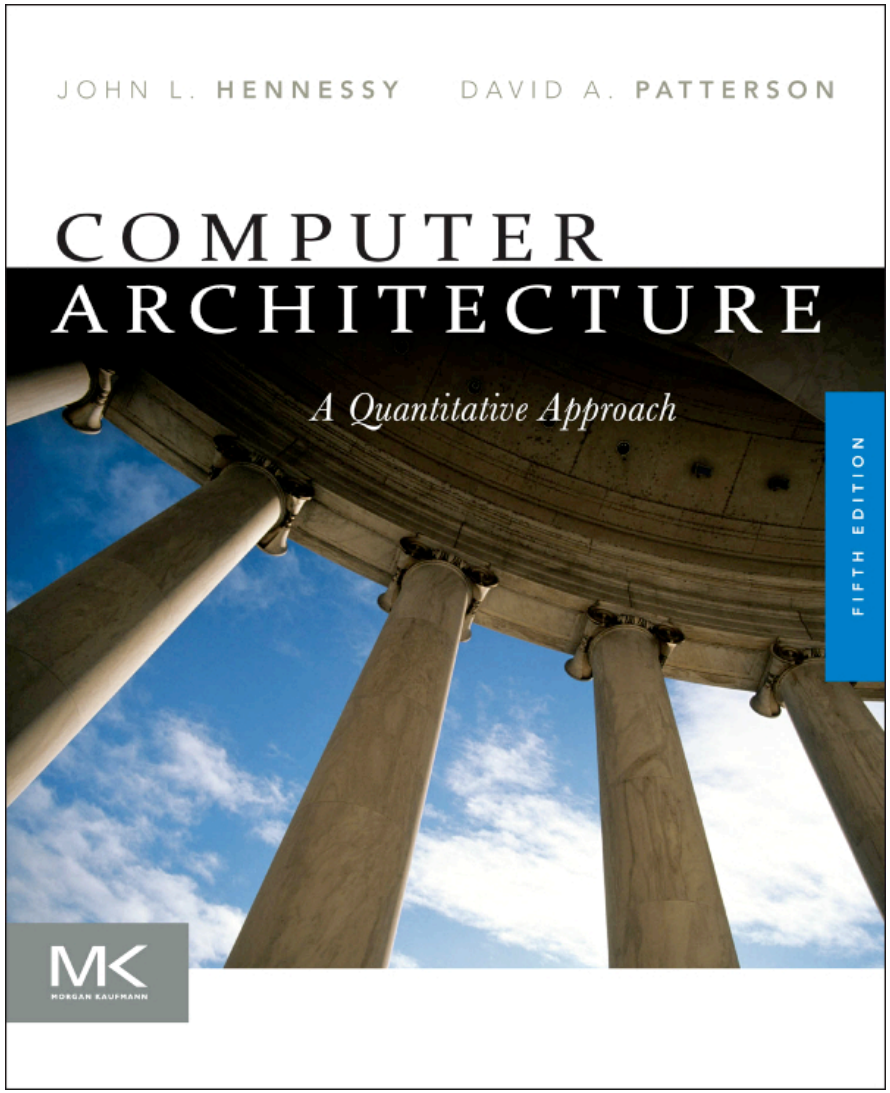

Computer Architecture, 5th Edition

Hennessy & Patterson

Table of Contents

Printed Text

Chap 1: Fundamentals of Quantitative Design and Analysis
Chap 2: Memory Hierarchy Design
Chap 3: Instruction-Level Parallelism and Its Exploitation
Chap 4: Data-Level Parallelism in Vector, SIMD, and GPU Architectures
Chap 5: Multiprocessors and Thread-Level Parallelism
Chap 6: The Warehouse-Scale Computer
App A: Instruction Set Principles
App B: Review of Memory Hierarchy
App C: Pipelining: Basic and Intermediate Concepts

Online

App D: Storage Systems
App E: Embedded Systems
App F: Interconnection Networks
App G: Vector Processors
App H: Hardware and Software for VLIW and EPIC
App I: Large-Scale Multiprocessors and Scientific Applications
App J: Computer Arithmetic
App K: Survey of Instruction Set Architectures
App L: Historical Perspectives

MK®

Recommended textbook (1)

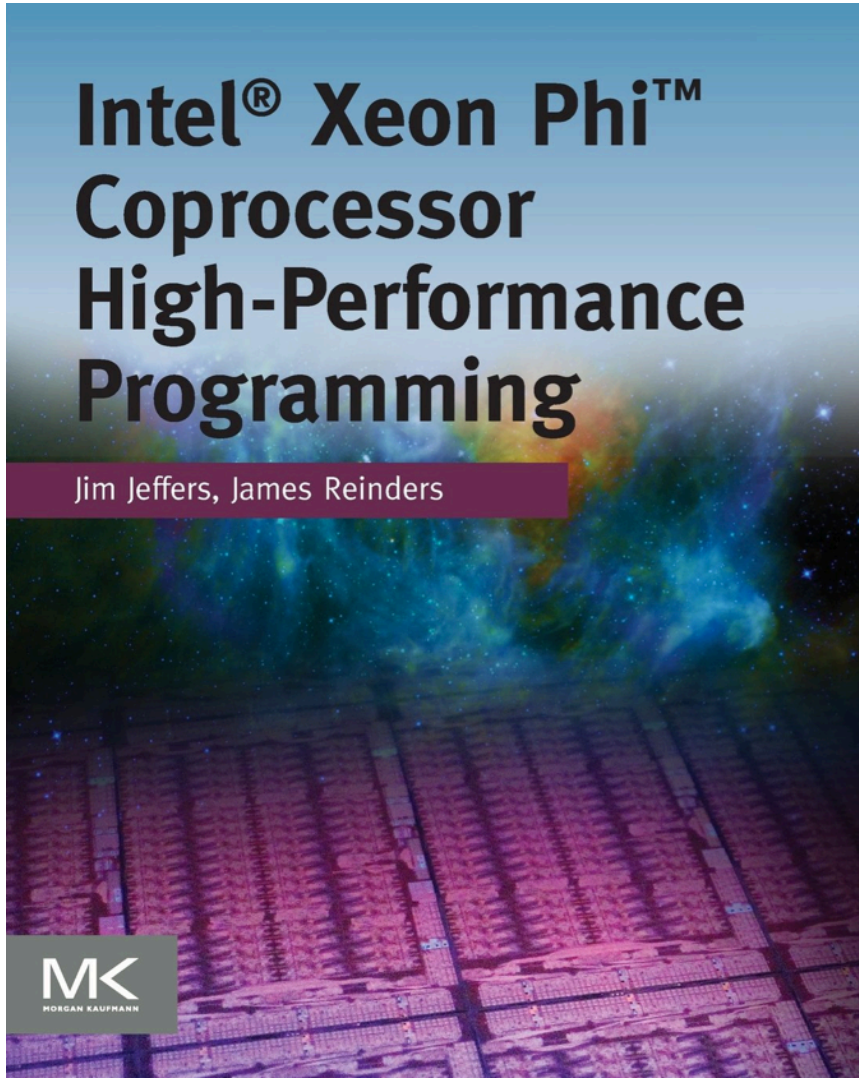

Contents

Recommended textbook (1)

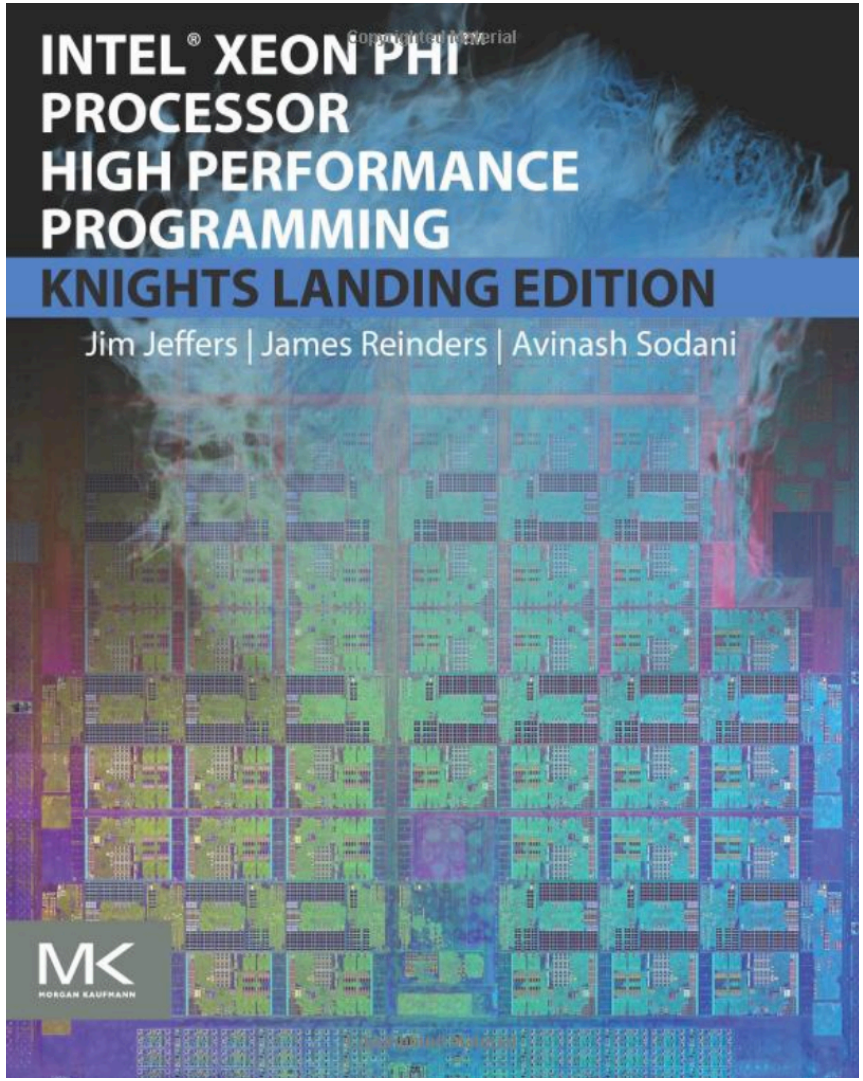

Contents

MK®

Recommended textbook (2)

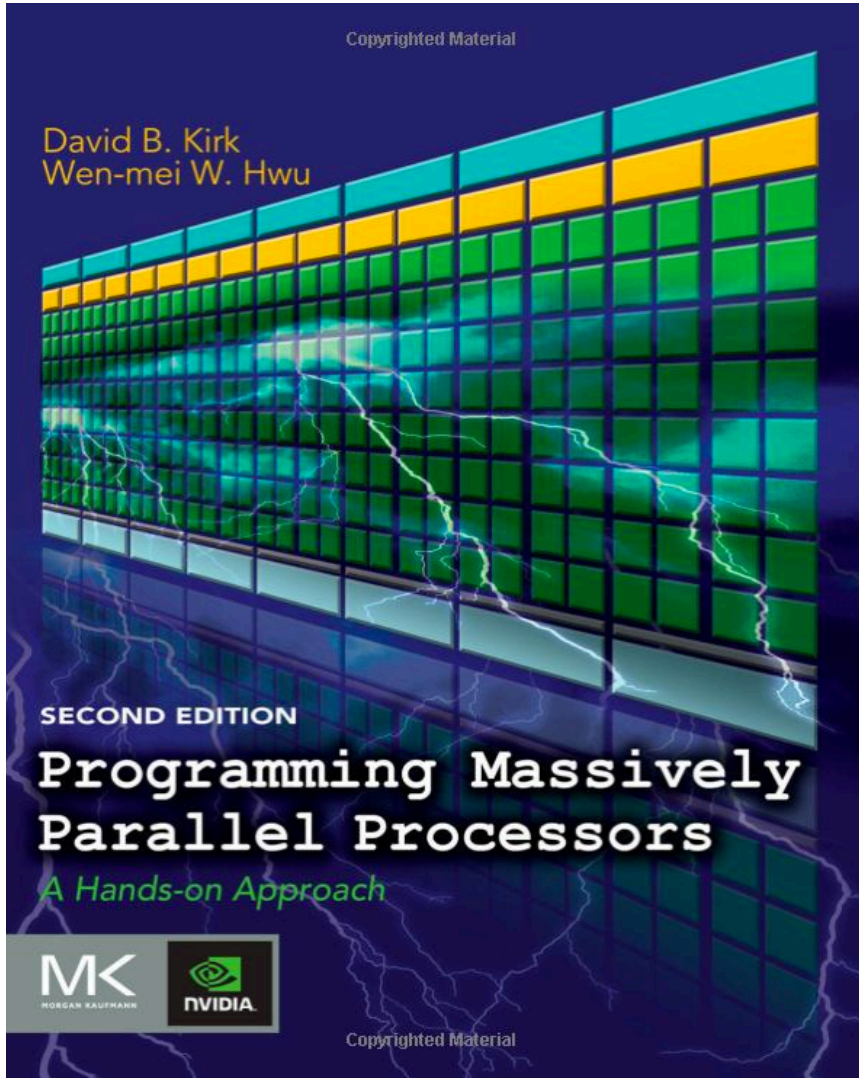

Contents