Master Program (Laurea Magistrale) in Computer Science and Networking

High Performance Computing Systems and Enabling Platforms

Marco Vanneschi

4. Shared Memory Parallel Architectures

4.4. Multicore Architectures

Multicore examples

- General purpose *vs* special purpose
  - Special purpose: Network Processors, DSP
- Homogeneous *vs* heterogeneous
- SPM *vs* NUMA
- Low parallelism *vs* high parallelism
  - Moore law: exponential growth of core number per chip
- …

General purpose - current (low parallelism) chips

- X86 based
  - Intel Xeon (Core 2 Duo, Core 2 Quad), Nehalem
  - AMD Athlon, Opteron quad-core (Barcelona)
- Power based
  - IBM Power 5, 6
  - IBM Cell
- UltraSPARC based
  - Sun UltraSparc T1
  - Sun UltraSparc T2
- Except IBM Cell: homogeneous, shared cache (C2 / C3) multiprocessors

Intel SMPs

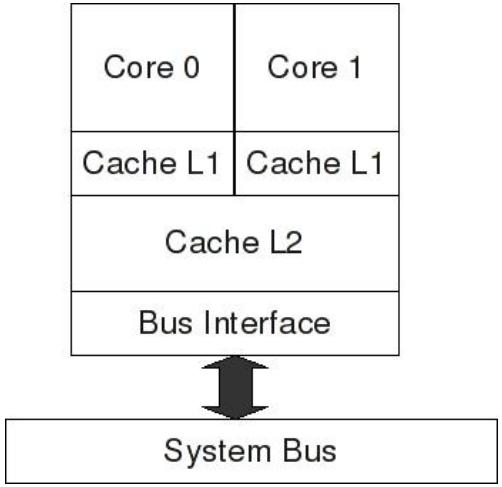

Xeon Core 2 Duo
3 GHz
L1: 32Kb + 32Kb
L2: 6Mb

Off-chip main memory interface
System bus: 10.6 GB/s

One thread per core, 4-superscalar

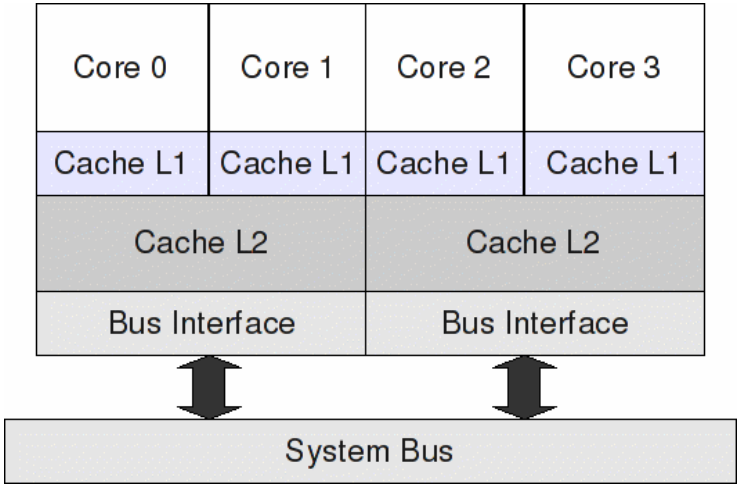

Xeon Core 2 Quad /Harpertown
Two Core 2 in the same chip
External main memory: shared by both L2 caches

Automatic cache coherence
MESI Snooping

Intel SMPs: Nehalem

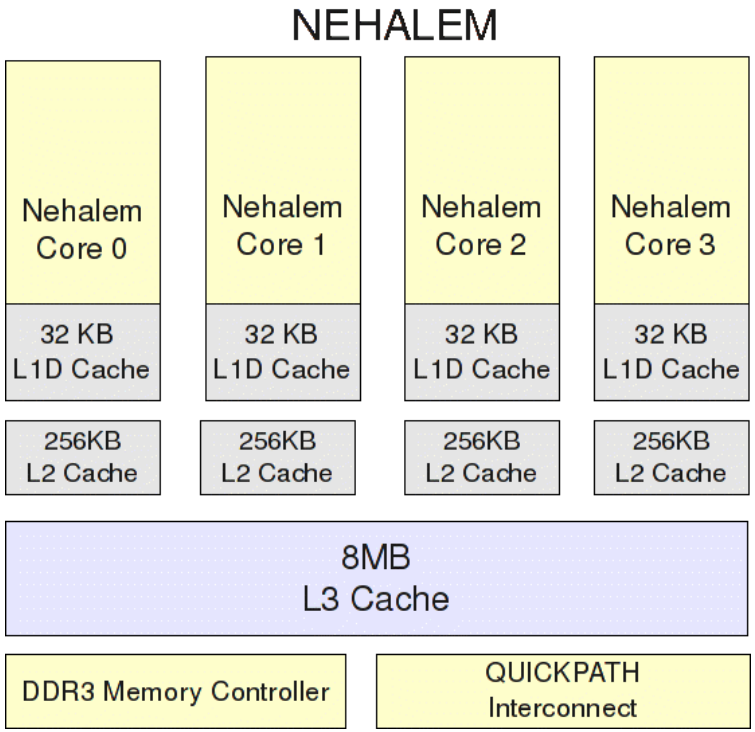

Evolution of SMP Xeon

Private L2 caches
Shared L3 cache, MESIF

Trend: 8, 16 core

Memory interface on chip

Point-to-point interconnection structure, 32 GB/s

Two simultaneous threads per core

AMD Opteron Quad Core

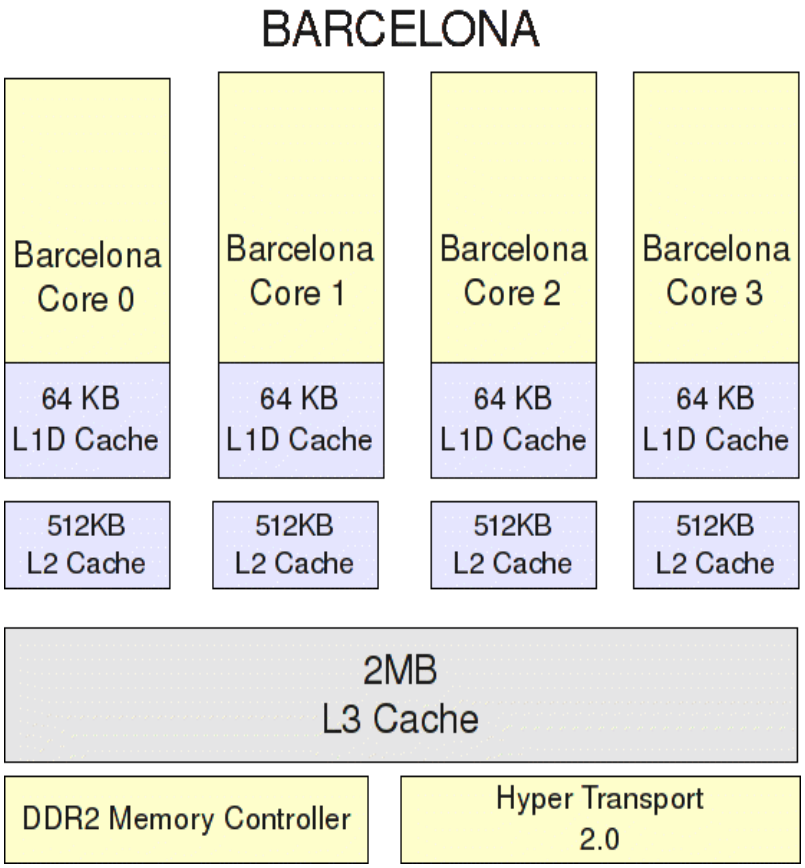

Shared L3 cache

Possible Cache-coherent **NUMA** behaviour:
access to remote L2 caches,
via point-to-point interconnection,
with L3 cache acting (also) as a synchronization agent.

SUN Niagara 2

UltraSPARC T2

SMP, shared L2 cache

8 simple, pipelined, in-order cores, 1 floating point unit per core

8 simultaneous threads per core

Interconnection: crossbar

IBM BlueGene/P

- PowerPC 32-bit 450 quad-core
- NUMA
- Mesh interconnect
- Automatic cache coherence

- 4 simultaneous floating point operations per clock cycle (850 MHz)
- Significantly less power consumption (16 W) compared to > 65 W of x86 quad-core
- BlueGene massively parallel system: > 75000 quad-core chips (> 290000 cores total).

IBM Cell BE

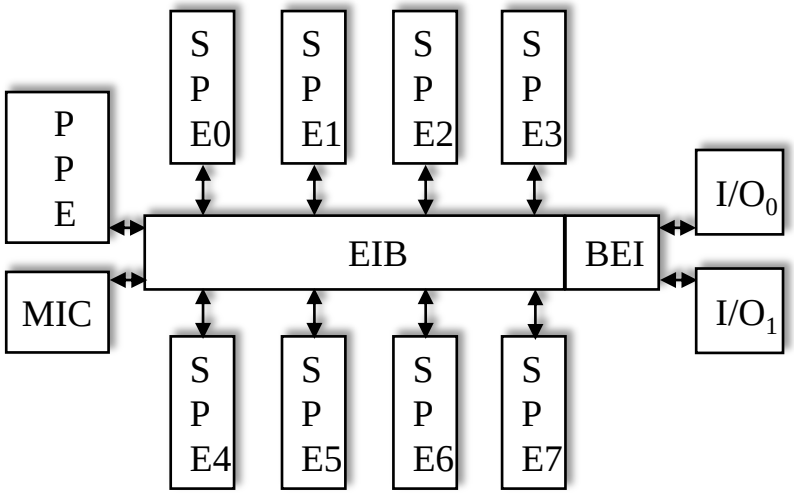

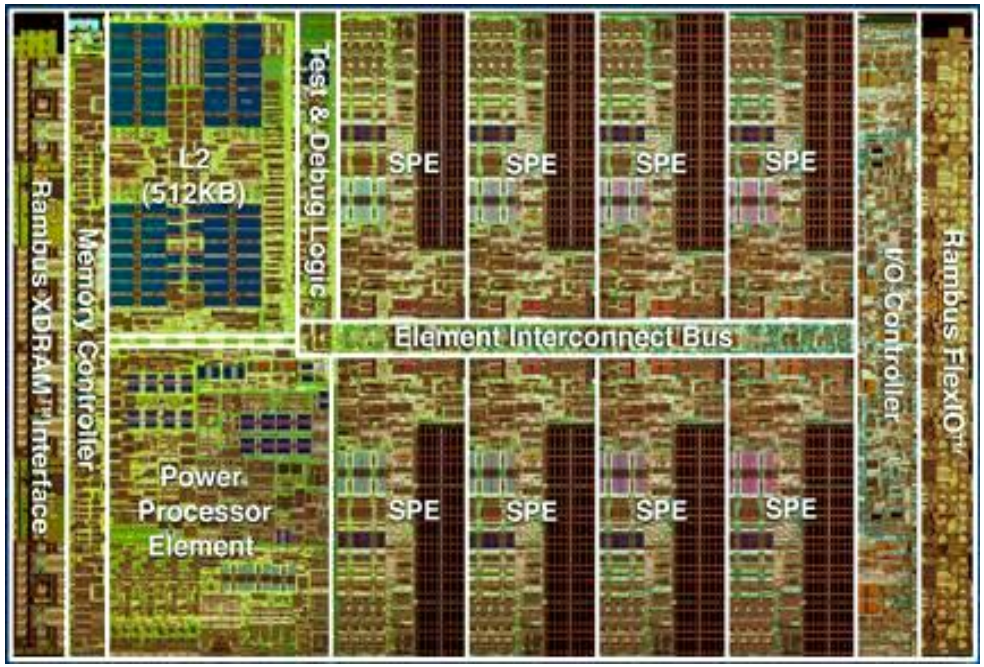

Evolution of uniprocessor (Power PC Processor Element, PPE) with 8 powerful I/O coprocessors
towards a **heterogeneous NUMA** multiprocessor,

Coprocessors evolution towards Processign Cores: Synergistic Processing Element (SPE), with vectorization capabilities

IBM Cell BE

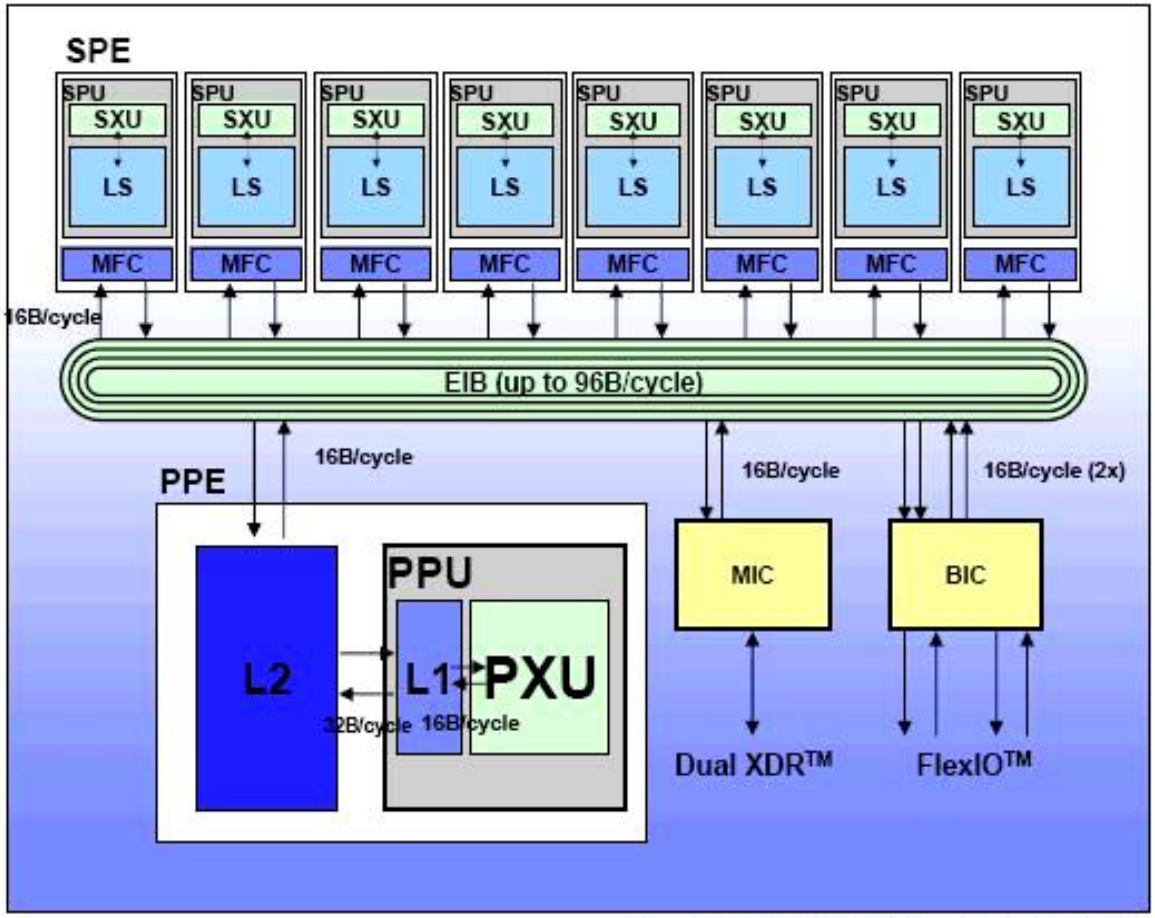

Source: M. Gschwind et al., Hot Chips-17, August 2005

PPE: superscalar, in-order, L2 cache accessible by SPEs

SPE: RISC, 128-bit,pipelined, in-order,
vectorized instructions

SPE Local Memory: 256 Kb, **NOT cache**

SPEs set = **NUMA**, with additional access to the PPE memory (DMA)

Interconnection structure:
4 bidirectional **Rings**, 16 bytes per ring

Robust cost model for memory access and communication.

Intel Terascale project

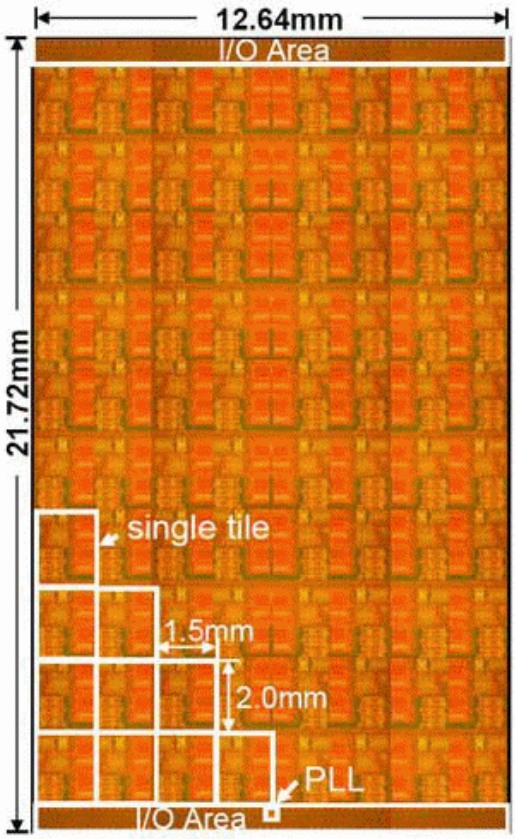

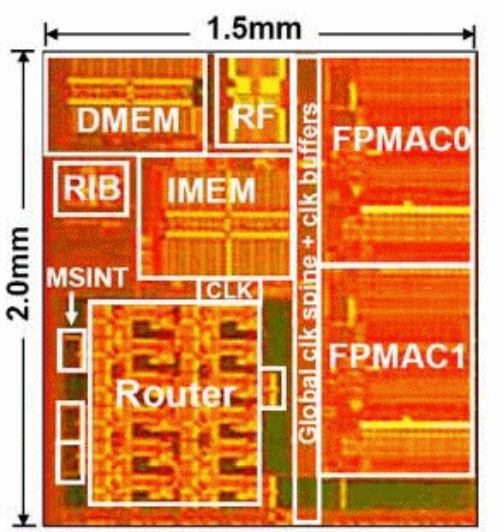

Technology	65nm CMOS Process
Interconnect	1 poly, 8 metal (Cu)
Transistors	100 Million
Die Area	275mm²
Tile area	3mm²
Package	1248 pin LGA, 14 layers, 343 signal pins

80 core

Network on-chip: **bidimensional mesh** (k-ary 2-cube)
Wormohole routing

VLIW processors: 96 bit instruction word

NOT cache

Tilera Tile 64

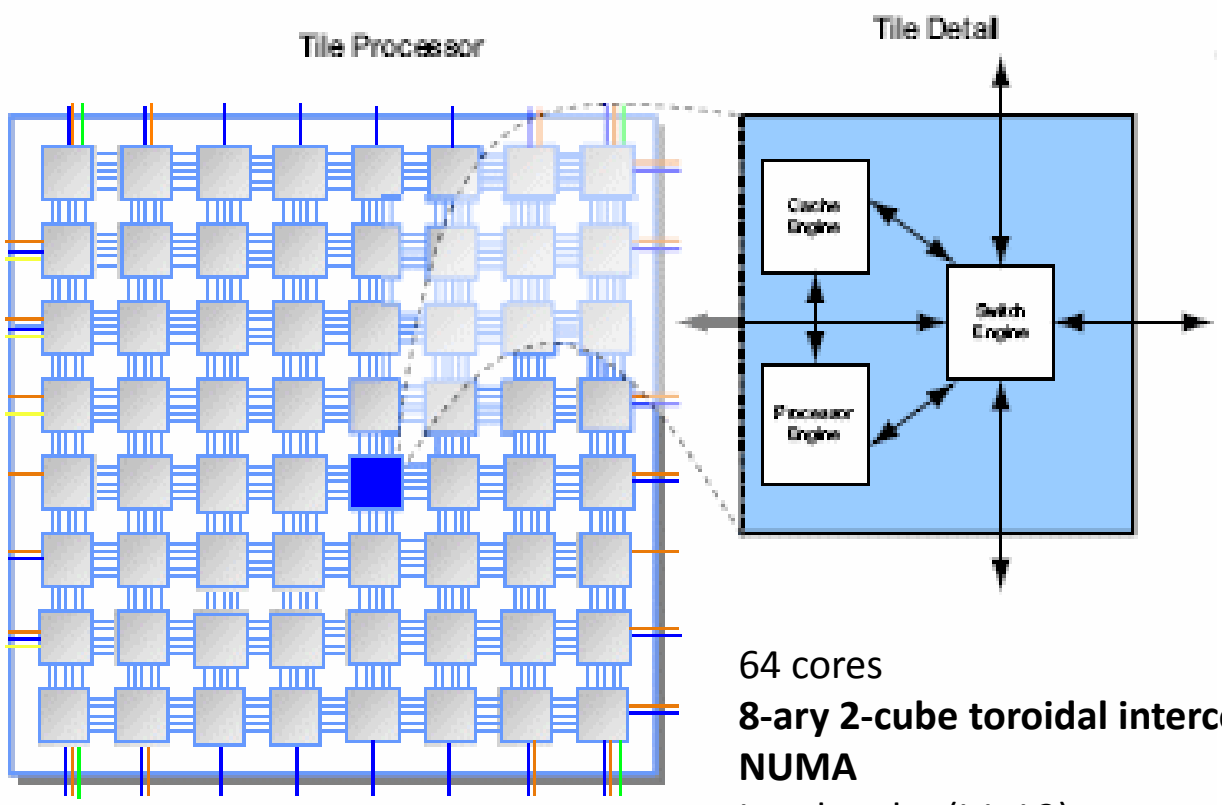

64 cores

8-ary 2-cube toroidal interconnection

NUMA

Local cache (L1, L2): **program-controlled cache coherence**

No floating-point

Oriented to Video encoding, Network Packet processing

Network processors

- Parallel architectures oriented to network processing:
  - Real-time processing of multiple data streams
  - IP protocol packet switching and forwarding capabilities
    - Packet operations
    - Packet queueing
    - Checksum / CRC per packet
    - Pattern matching per packet
  - Tree searches
  - Frame forrwarding
  - Frame filtering
  - Frame alteration
  - Traffic control and statistics
  - QoS control
  - Enhance security
  - Primitive network interfaces on-chip
- Intel, IBM, Ezchip, Xelerated, Agere, Alchemy, AMCC, Cisco, Cognigine, Motorola, …
- Similar features for DSP multicore processors

Network processors and multithreading

- Network processors apply multithreading to bridge latencies during (remote) memory accesses
  - Blocked multithreading (BMT)
  - Multithreading applied to cores that perform the data traffic handling
- Hard real-time events (i.e., deadline soluld never be missed)
  - Specific instruction scheduling during multithreaded execution
- Examples:
  - Intel IXP
  - IBM PowerNP

Network processors: Intel Internet eXchange Processor (IXP)

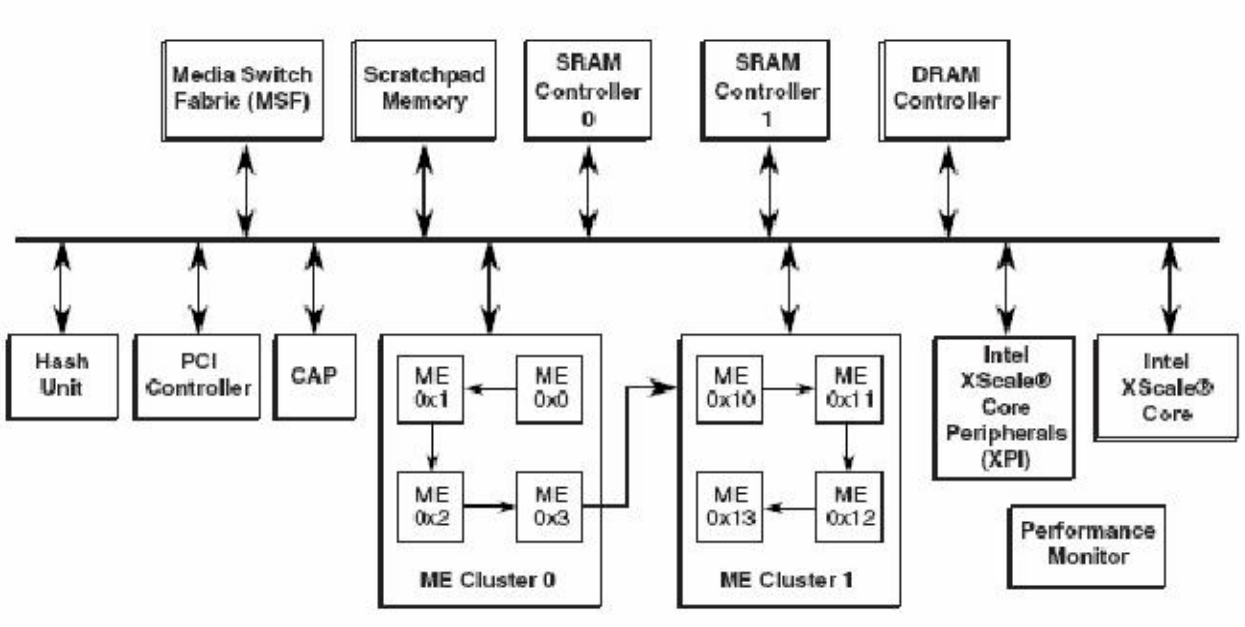

8 cores (IXP2400) or 16 cores (IPX2800), specialized for low-level packet processing,

fifty 40-bit instructions

+ one RISC Intel Xscale, 600-700 MHz: heterogeneous NUMA

Pipelined architecture,

8 threads per core (zero cost context switching thread-thread)

Ring-like core interconnection