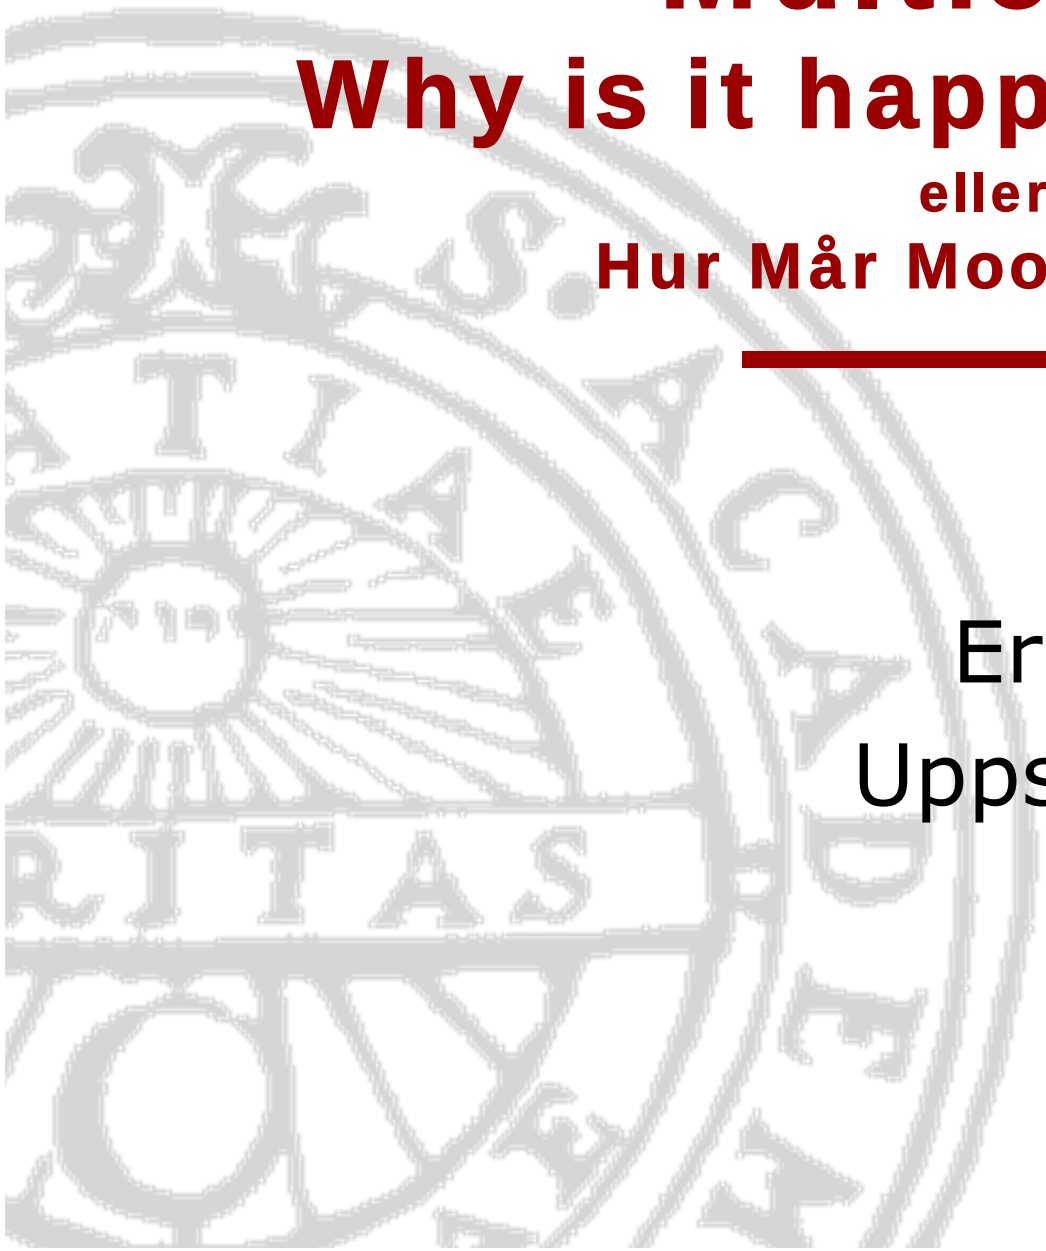

Multicore:
Why is it happening now?
eller
Hur Mår Moore's Lag?

Erik Hagersten
Uppsala Universitet

Darling, I shrunk the computer

Mainframes

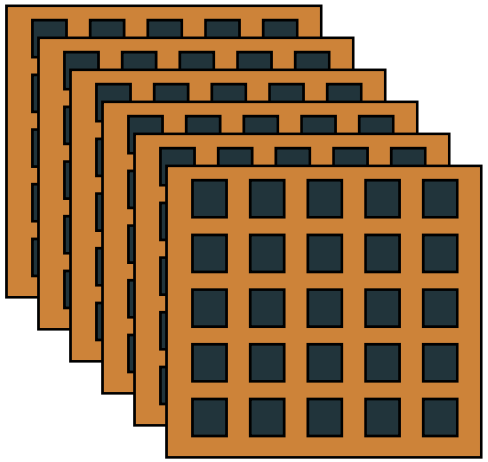

Super Minis:

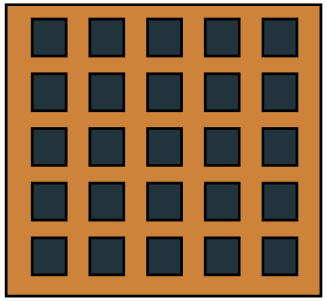

Microprocessor:

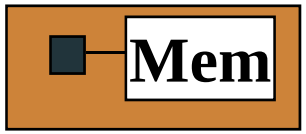

Multicore: Many CPUs on a chip!

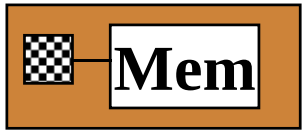

UPPSALA UNIVERSITET

Outline

- Why multicore now?
- Performance bottlenecks in MCs
- Commercial offerings
- Reflection for the future

PDC Summer School 2013

Everybody is doing it! But, why now?

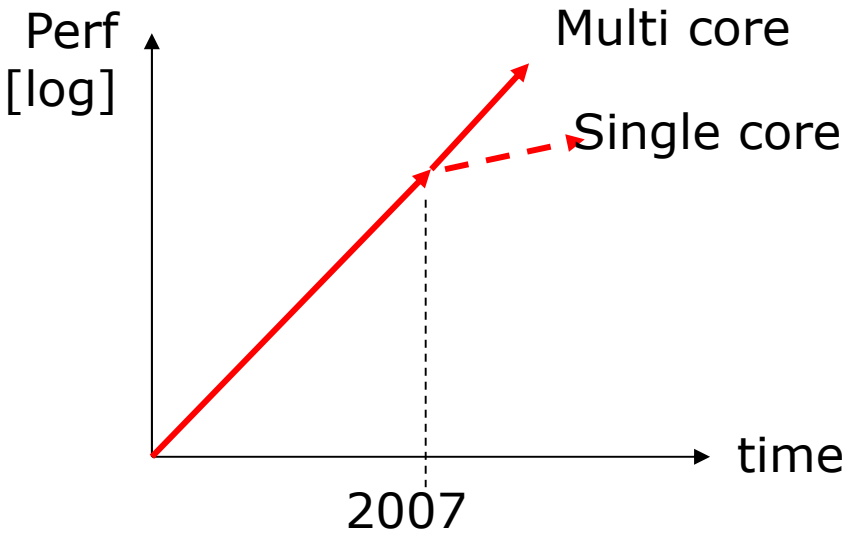

1. Not enough ILP to get payoff from using more transistors
2. Signal propagation delay » transistor delay
3. Power consumption $P_{dyn} \sim C \bullet f \bullet V^2$

UPPSALA UNIVERSITET

Multi-core CPUs:

- **Ageia** PhysX, a multi-core physics processing unit.
- **Ambric** Am2045, a 336-core Massively Parallel Processor Array (MPPA)
- **AMD**
  - Athlon 64, Athlon 64 FX and Athlon 64 X2 family, dual-core desktop processors.
  - Opteron, dual- and quad-core server/workstation processors.
  - Phenom, triple- and quad-core desktop processors.
  - Sempron X2, dual-core entry level processors.
  - Turion 64 X2, dual-core laptop processors.
  - Radeon and FireStream multi-core GPU/GPGPU (10 cores, 16 5-issue wide superscalar stream processors per core)
- **ARM** MPCore is a fully synthesizable multicore container for ARM9 and ARM11 processor cores, intended for high-performance embedded and entertainment **applications**.
- **Azul** Systems Vega 2, a 48-core processor.
- **Broadcom** SiByte SB1250, SB1255 and SB1455.
- **Cradle** Technologies CT3400 and CT3600, both multi-core DSPs.
- **Cavium** Networks Octeon, a 16-core MIPS MPU.
- **HP** PA-8800 and PA-8900, dual core PA-RISC processors.
- **IBM**
  - POWER4, the world's first dual-core processor, released in 2001.
  - POWER5, a dual-core processor, released in 2004.
  - POWER6, a dual-core processor, released in 2007.
  - PowerPC 970MP, a dual-core processor, used in the Apple Power Mac G5.
  - Xenon, a triple-core, SMT-capable, PowerPC microprocessor used in the Microsoft Xbox 360 game console.
- **IBM**, Sony, and Toshiba Cell processor, a nine-core processor with one general purpose PowerPC core and eight specialized SPUs (Synergystic Processing Unit) **optimized** for vector operations used in the Sony PlayStation 3.
- **Infineon** Danube, a dual-core, MIPS-based, home gateway processor.
- **Intel**
  - Celeron Dual Core, the first dual-core processor for the budget/entry-level market.
  - Core Duo, a dual-core processor.
  - Core 2 Duo, a dual-core processor.
  - Core 2 Quad, a quad-core processor.
  - Core i7, a quad-core processor, the successor of the Core 2 Duo and the Core 2 Quad.
  - Itanium 2, a dual-core processor.
  - Pentium D, a dual-core processor.
  - Teraflops Research Chip (Polaris), an 3.16 GHz, 80-core processor prototype, which the company says will be released within the next five years[6].
  - Xeon dual-, quad- and hexa-core processors.
- **IntellaSys** seaForth24, a 24-core processor.
- **Nvidia**
  - GeForce 9 multi-core GPU (8 cores, 16 scalar stream processors per core)
  - GeForce 200 multi-core GPU (10 cores, 24 scalar stream processors per core)
  - Tesla multi-core GPGPU (8 cores, 16 scalar stream processors per core)
- Parallax Propeller P8X32, an eight-core microcontroller.
- picoChip PC200 series 200-300 cores per device for DSP & wireless
- Rapport Kilocore KC256, a 257-core microcontroller with a PowerPC core and 256 8-bit "processing elements".
- Raza Microelectronics XLR, an eight-core MIPS MPU
- **Sun Microsystems**
  - UltraSPARC IV and UltraSPARC IV+, dual-core processors.
  - UltraSPARC T1, an eight-core, 32-thread processor.
  - UltraSPARC T2, an eight-core, 64-concurrent-thread processor.
- Texas Instruments TMS320C80 MVP, a five-core multimedia video processor.
- Tilera TILE64, a 64-core processor
- XMOS Software Defined Silicon quad-core XS1-G4

[source: Wikipedia 2008]

Example: Freq. Scaling

$$P_{dyn} = C * f * V^2 \approx \text{area} * \text{freq} * \text{voltage}^2$$

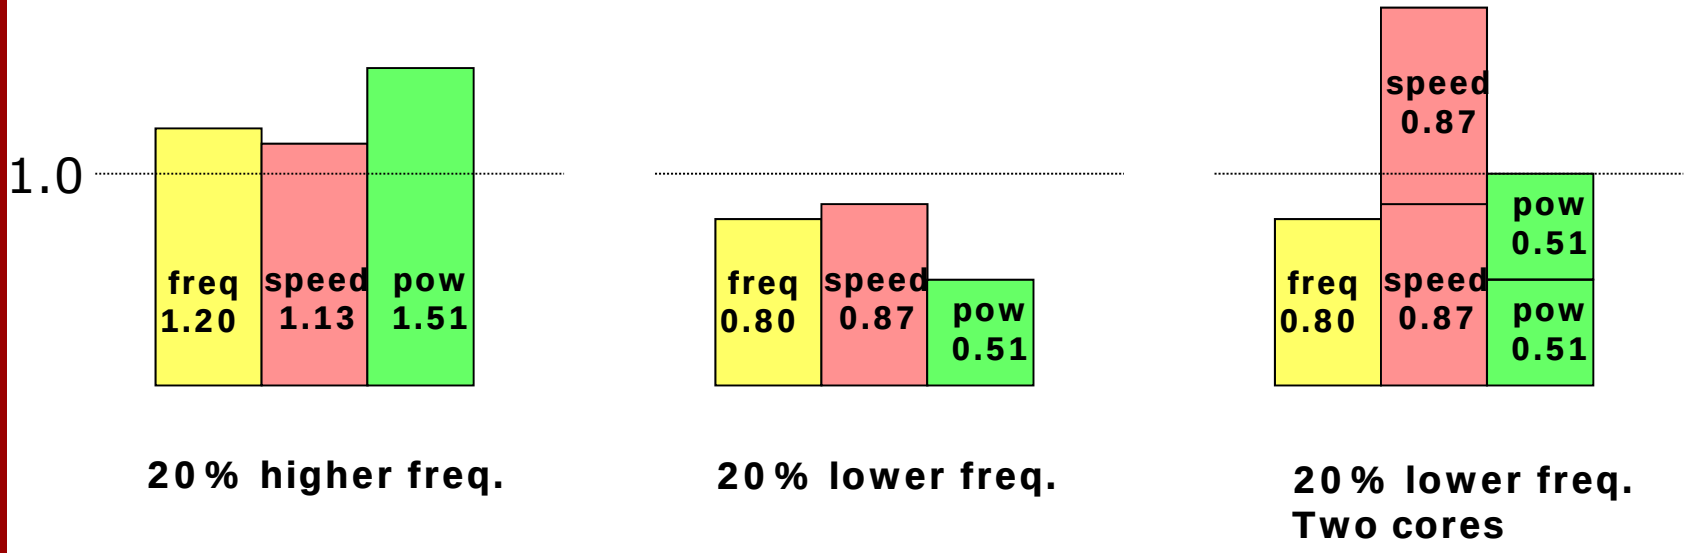

20% higher freq.

20% lower freq.

20% lower freq. Two cores

Darling, I shrunk the computer

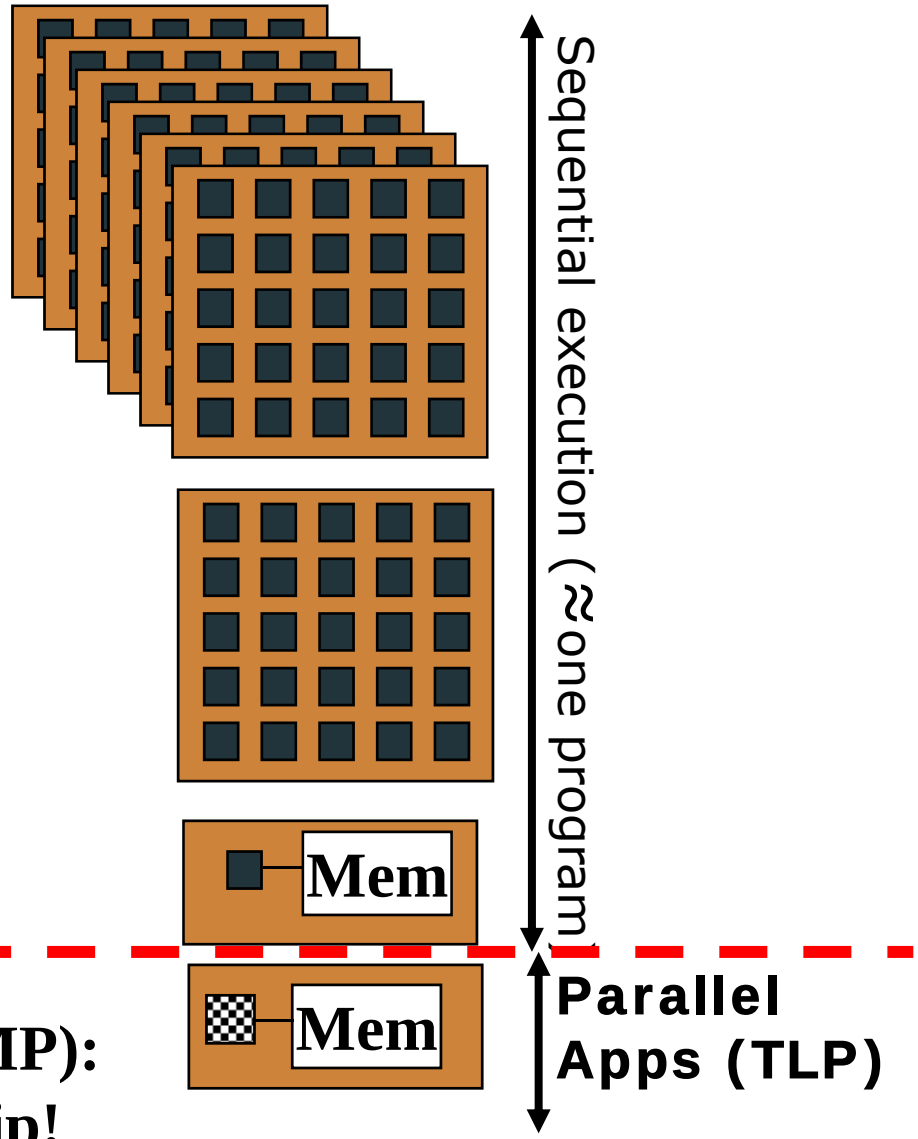

Paradigm Shift

Chip Multiprocessor (CMP): A multiprocessor on a chip!

Shared Bottlenecks

(the MCs on these slides are generalized)

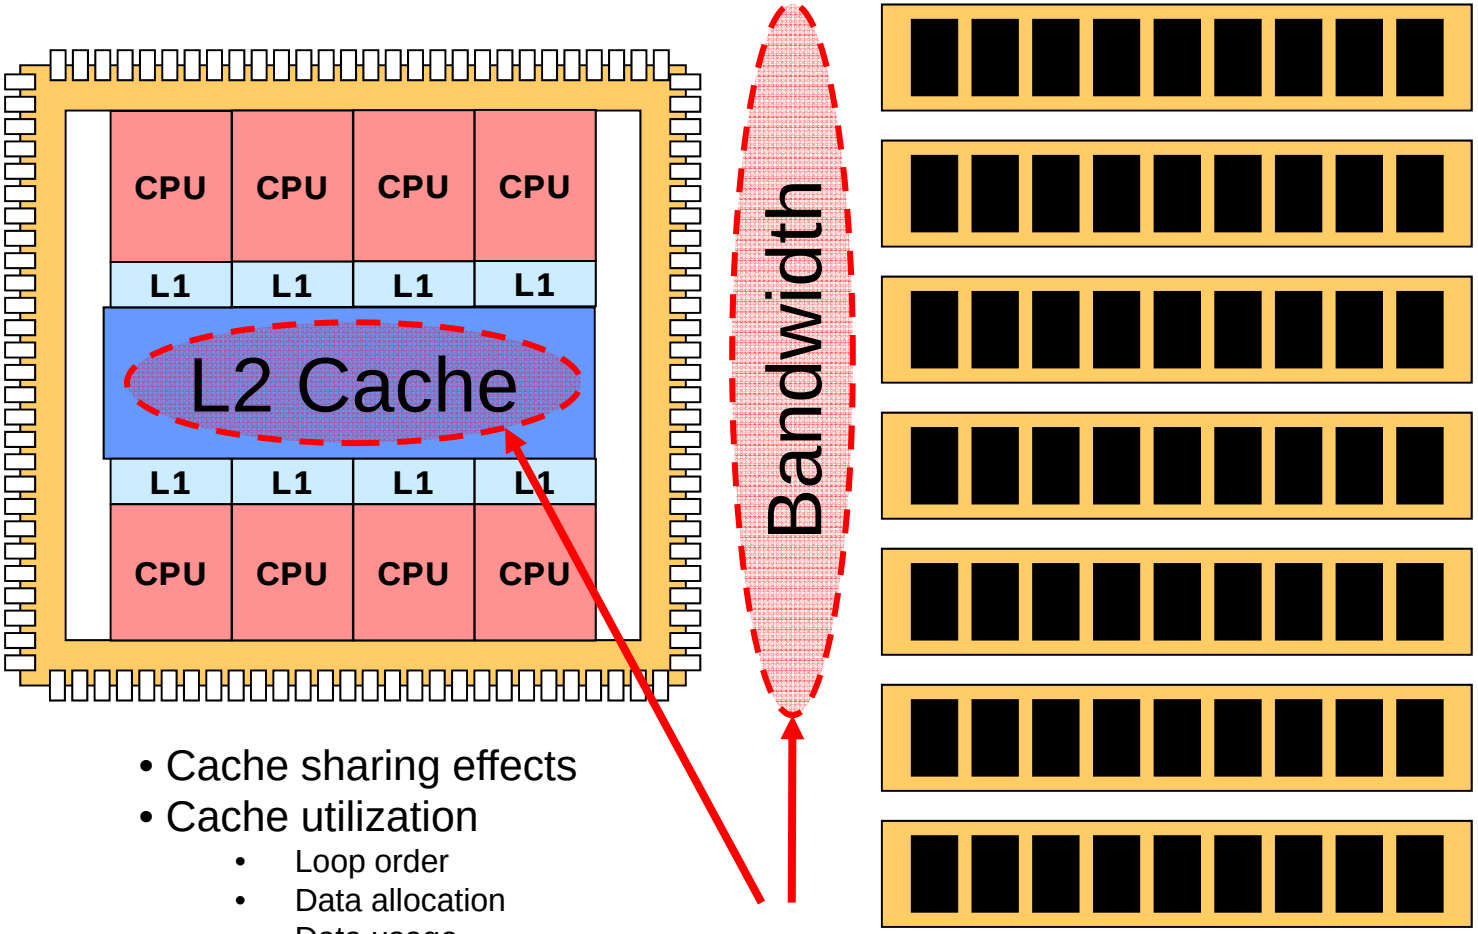

- Cache sharing effects
- Cache utilization
  - Loop order
  - Data allocation
  - Data usage…
- Data reuse
  - Tiling (aka Blocking)
  - Fusion…

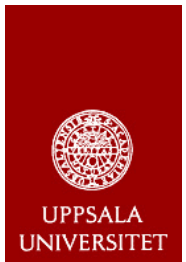

UPPSALA UNIVERSITET

Example: Poor Throughput Scaling!

Example: 470.LBM
"Lattice Boltzmann Method" to simulate incompressible fluids in 3D

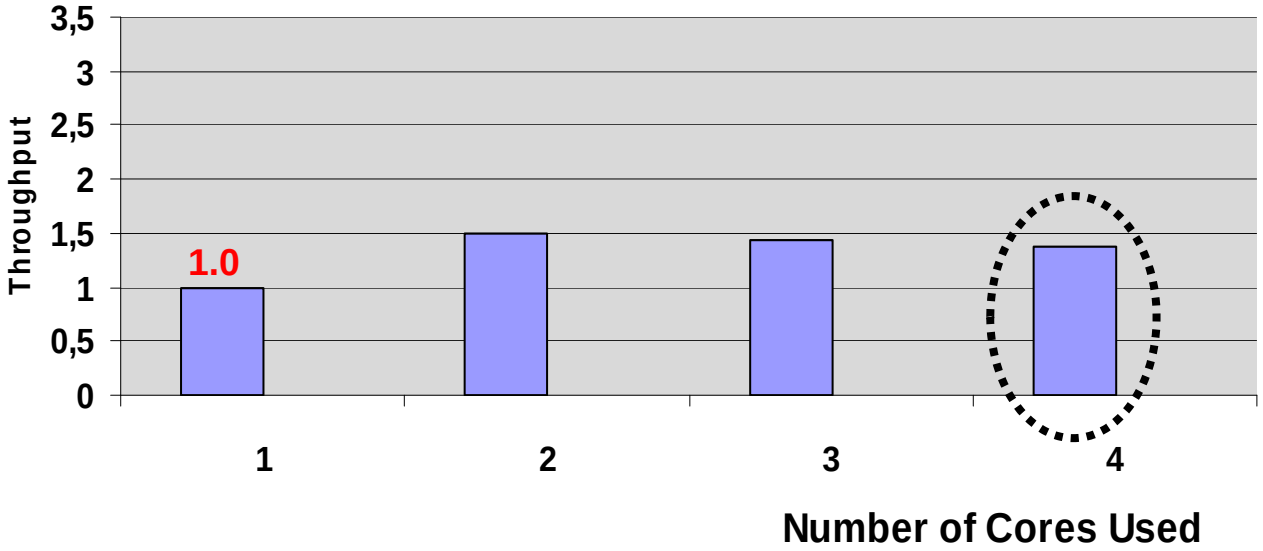

Throughput (as defined by SPEC):
Amount of work performed per time unit when several instances of the application is executed simultaneously.
Our TP study: compare TP improvement when you go from 1 core to 4 cores

UPPSALA UNIVERSITET

Throughput Scaling, More Apps

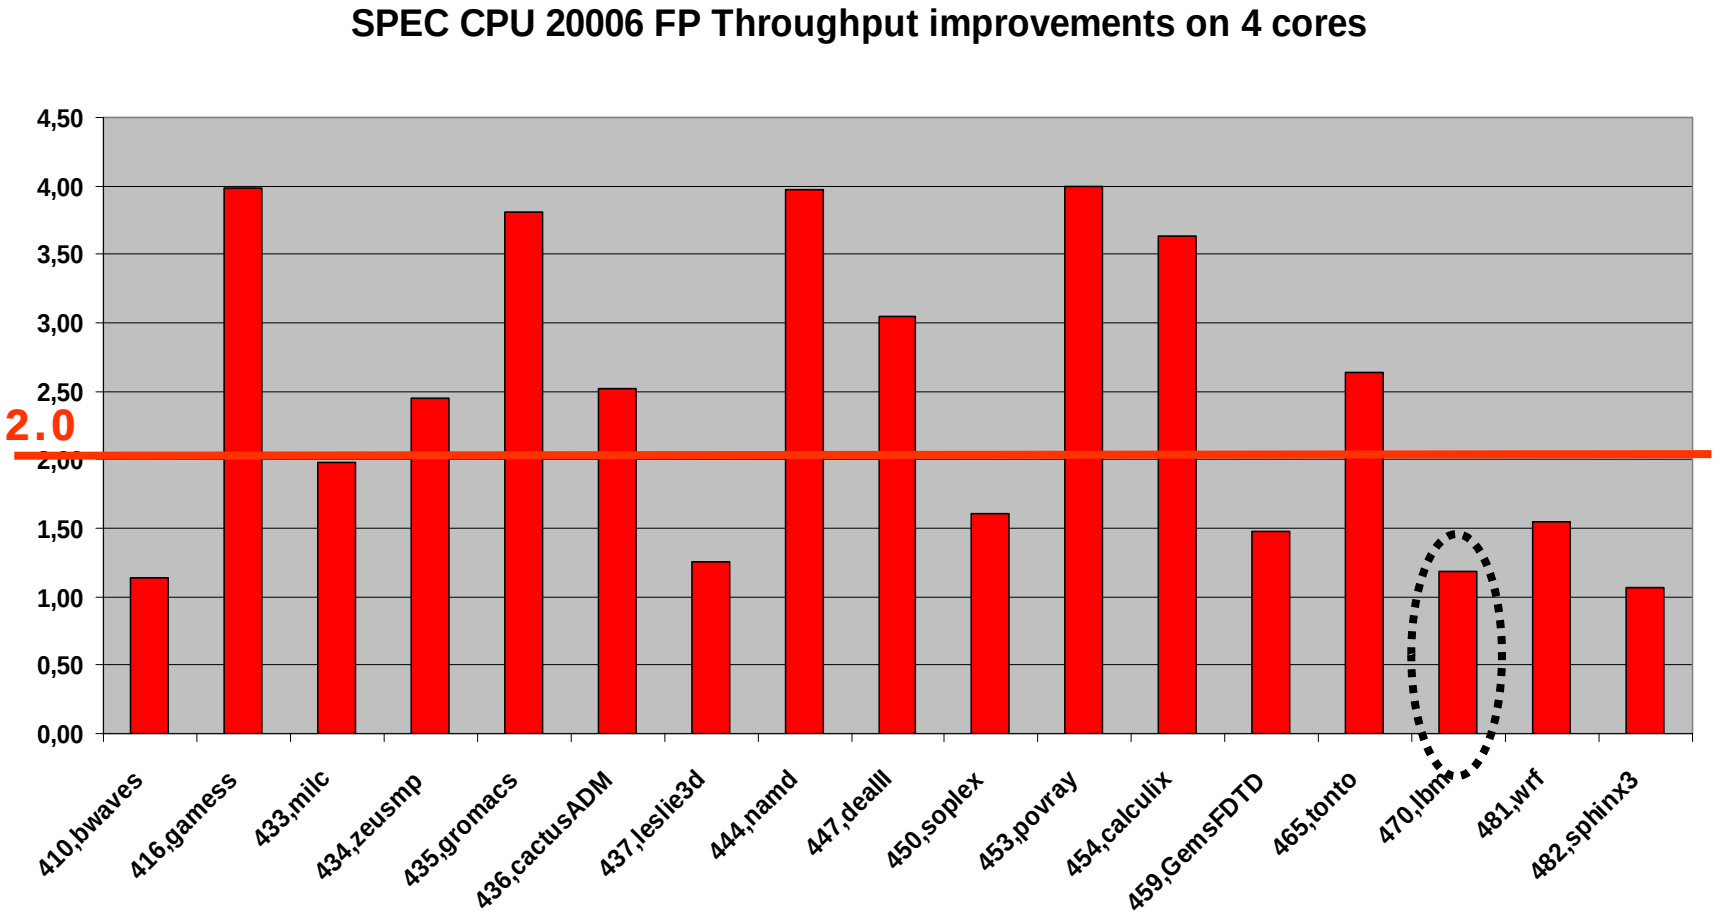

Intel X5365 3GHz "Core2", 1333 MHz FSB, 8MB L2.
(Based on data from the SPEC web)

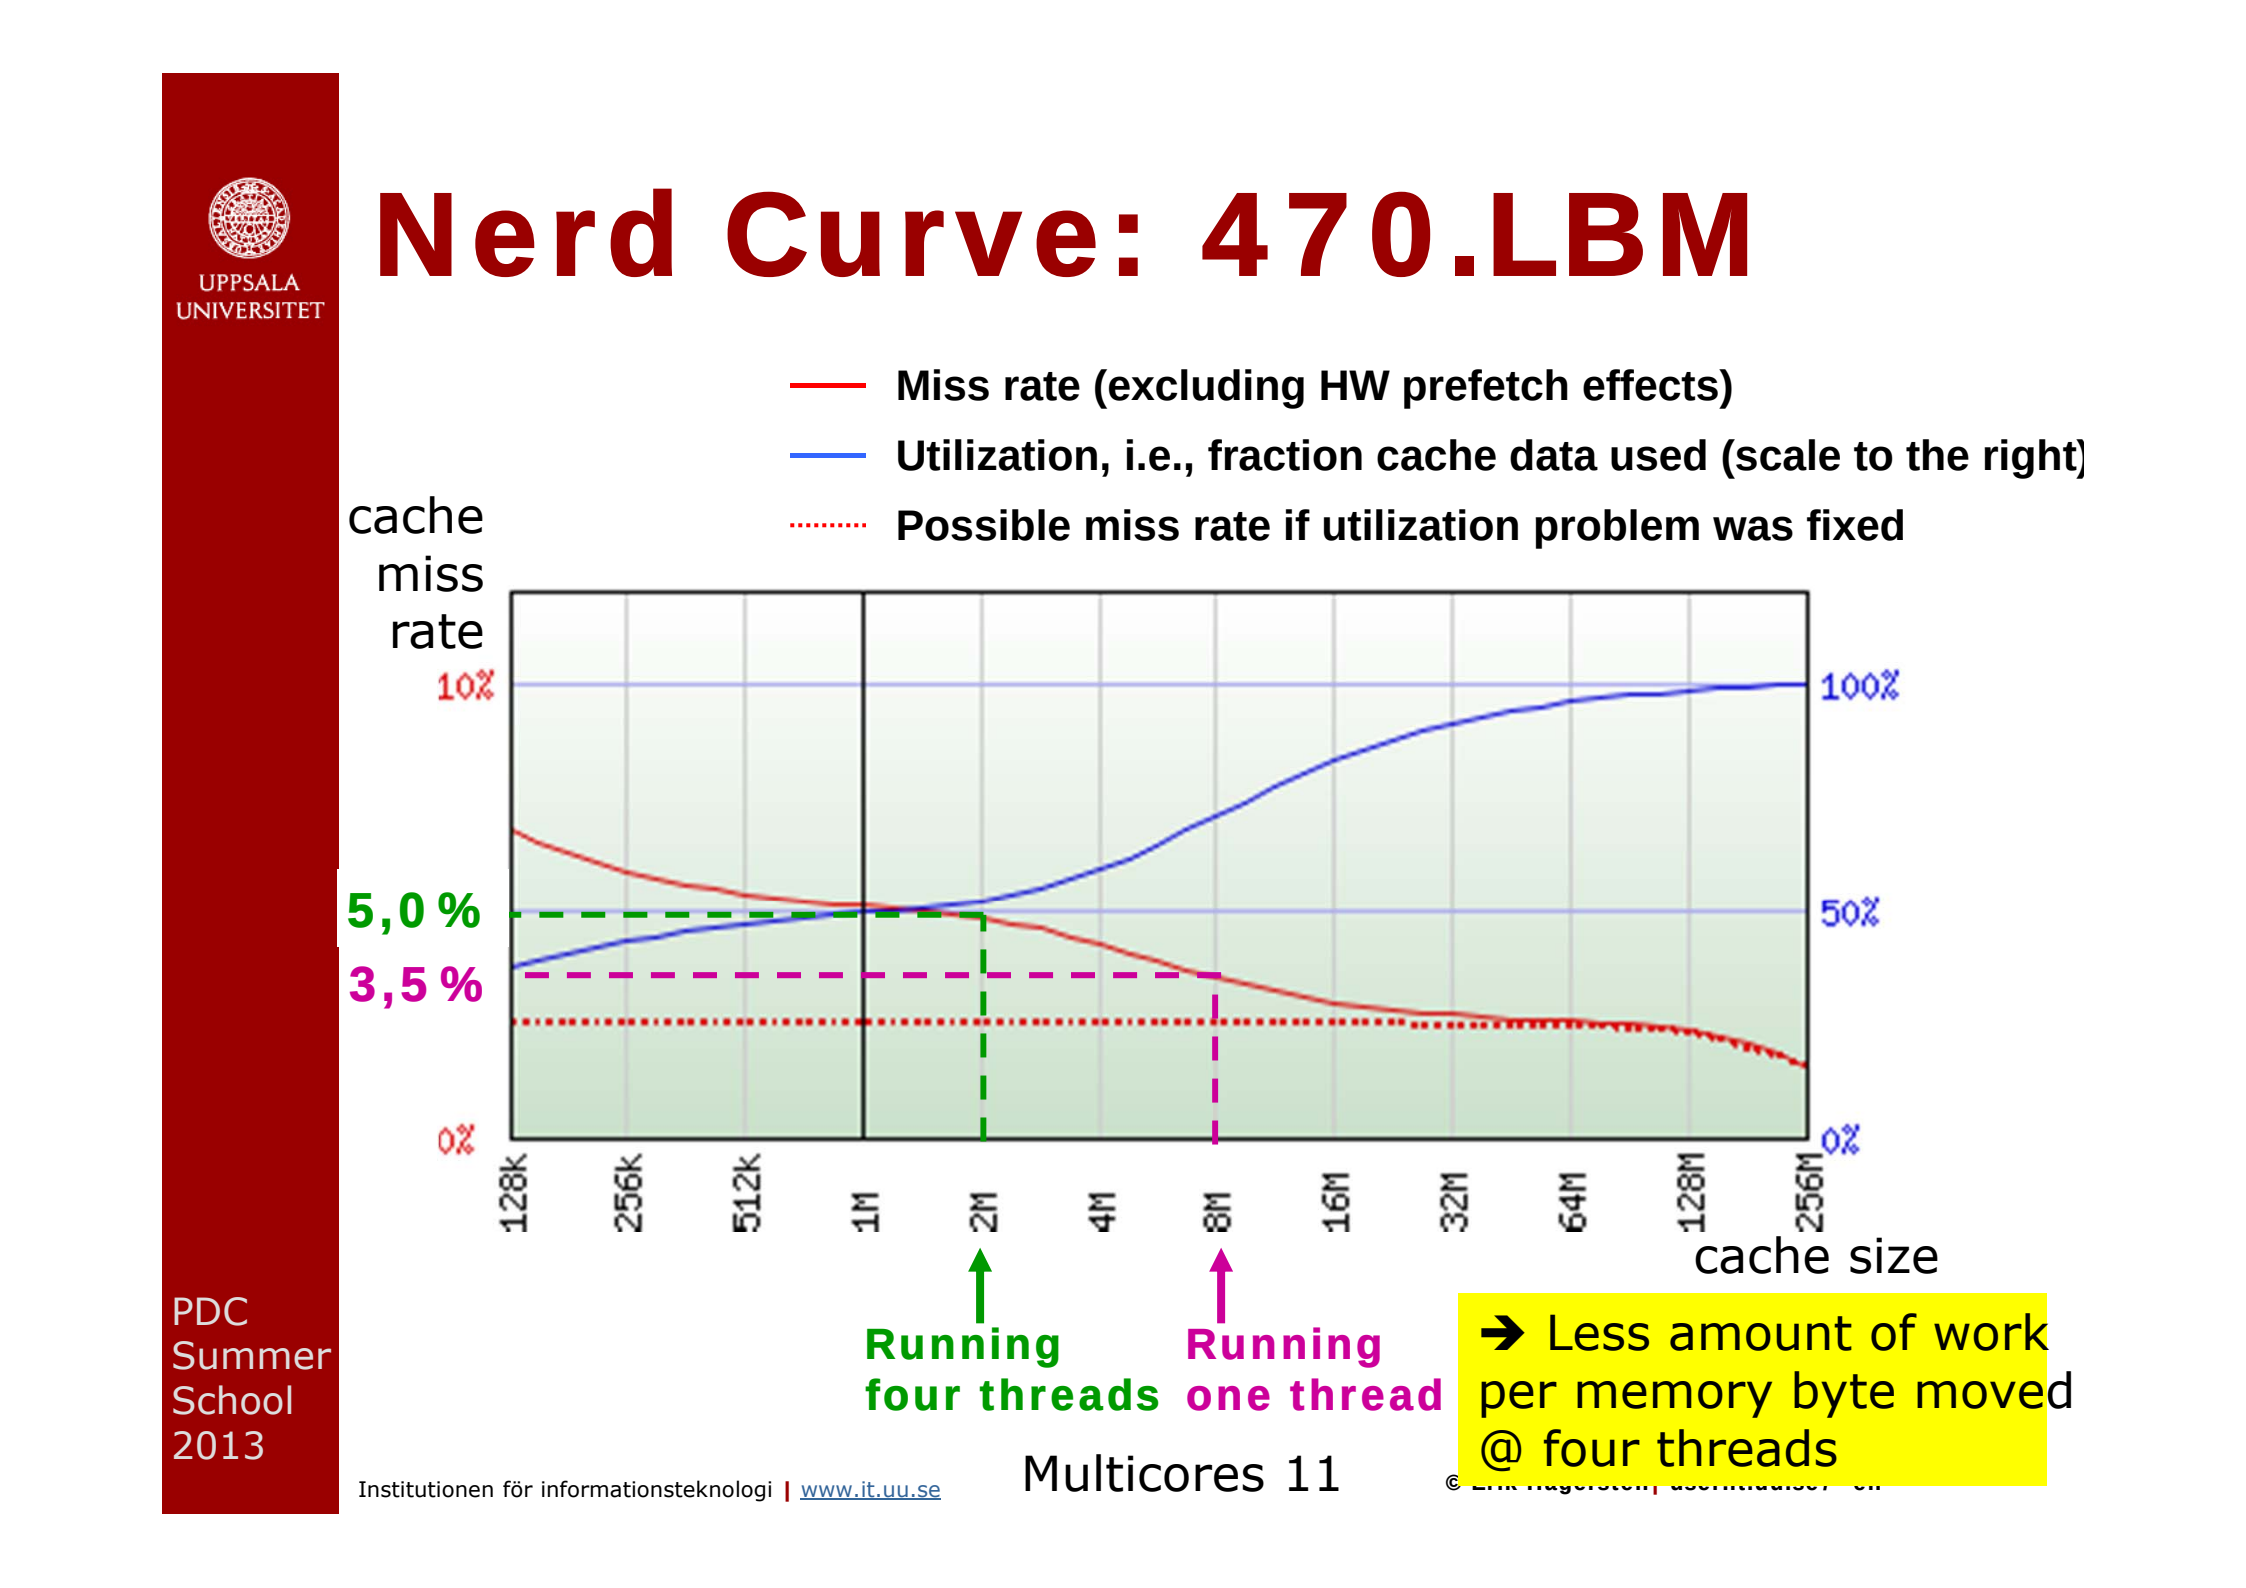
Nerd Curve: 470.LBM
Miss rate (excluding HW prefetch effects)
Utilization, i.e., fraction cache data used (scale to the right)
Possible miss rate if utilization problem was fixed
cache miss rate
10%
5,0 %
3,5 %
0%
100%
50%
0%
128k
256k
512k
1M
2M
4M
8M
16M
32M
64M
128M
256M
cache size
Running four threads
Running one thread
➔ Less amount of work per memory byte moved @ four threads
UPPSALA UNIVERSITET
PDC Summer School 2013
Institutionen för informationsteknologi | www.it.uu.se
Multicores 11

Nerd Curve (again)

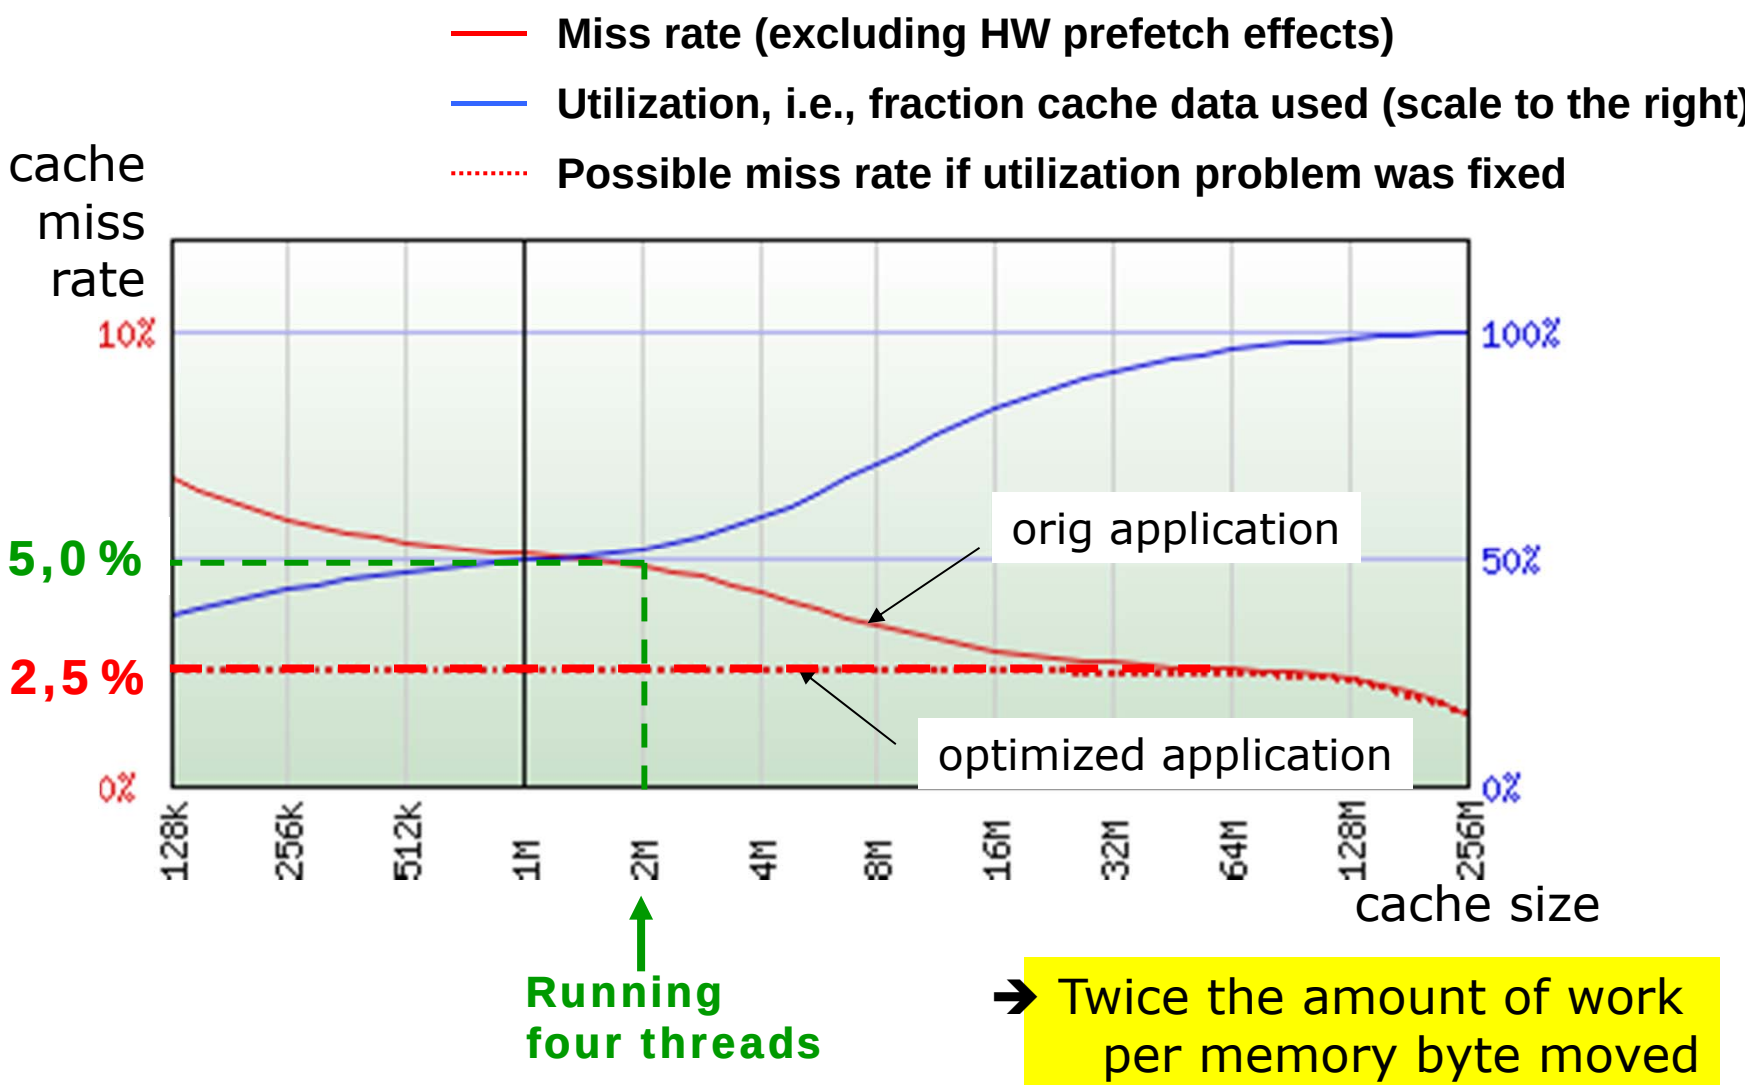

➔ Twice the amount of work per memory byte moved

➔ Better Memory Usage!

Example: 470.LBM

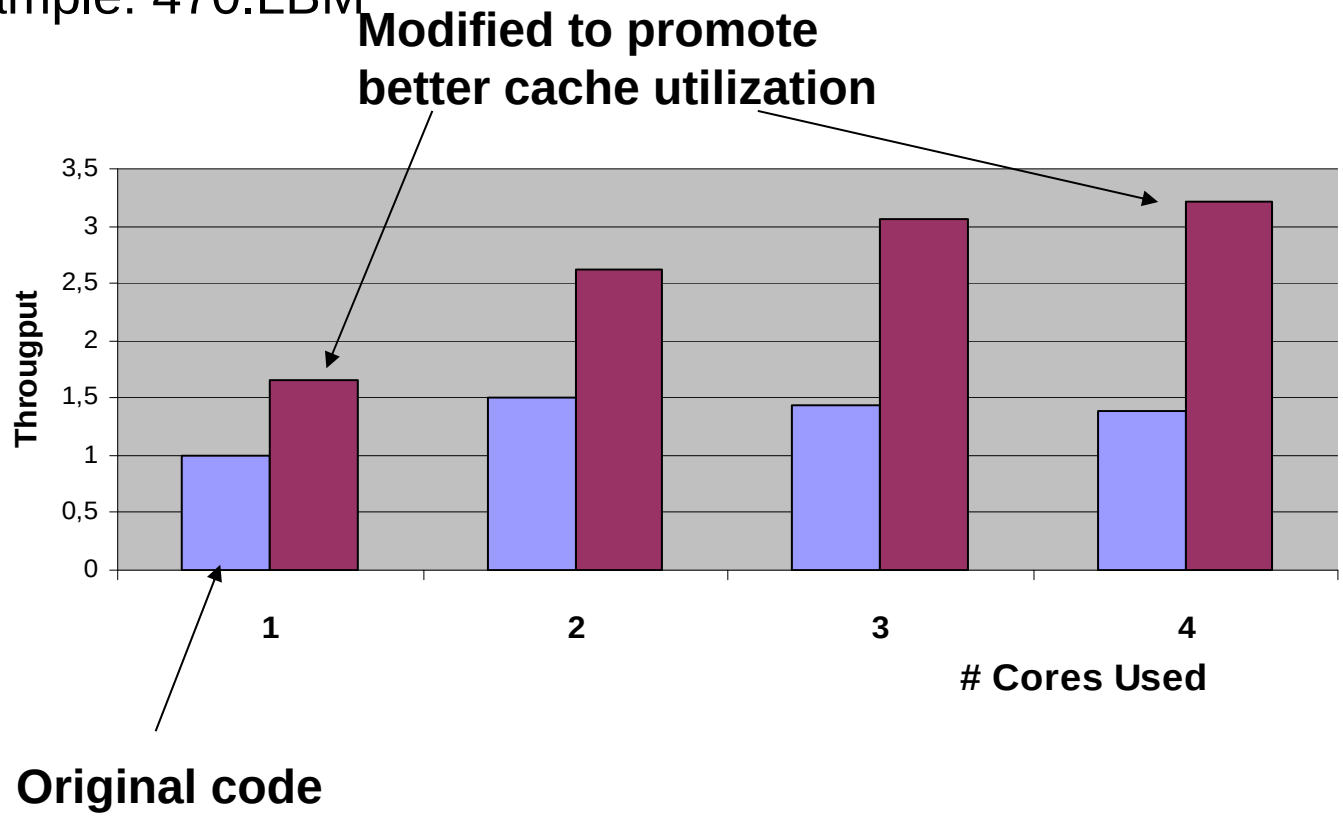

UPPSALA UNIVERSITET

BW in the Future?

#Cores ~ #Transistors

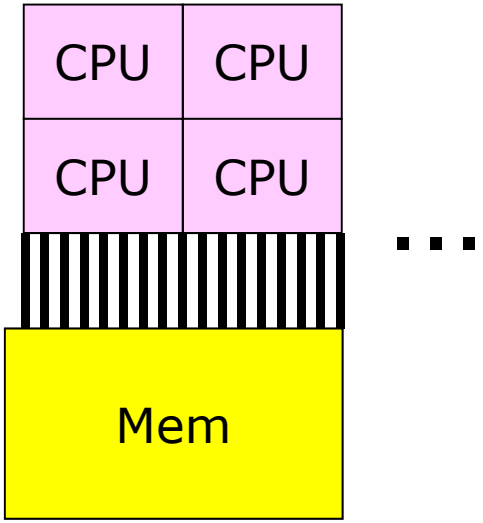

Computation vs Bandwidth

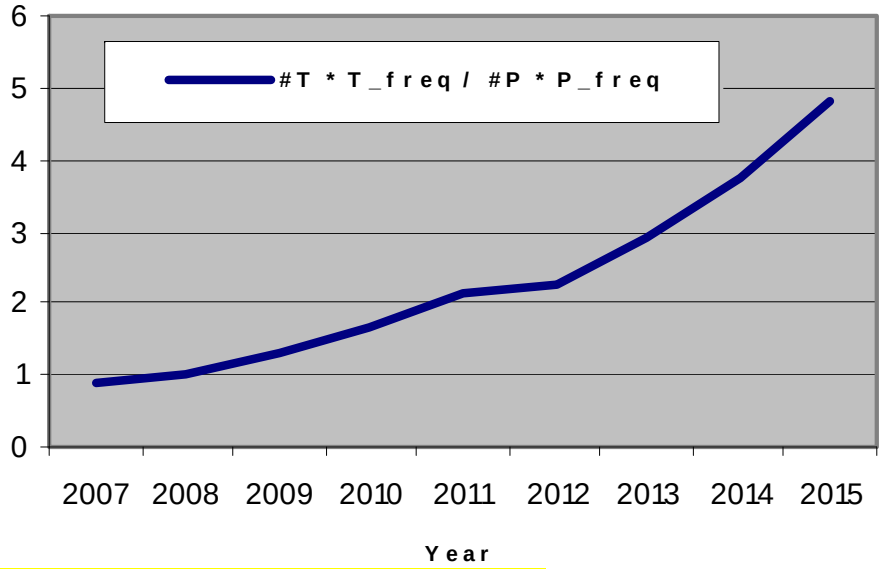

ap for Semiconductors (ITRS)

Chip Multiprocessors

HPCWire.com this morning:
Up Against the Memory Wall
”Nevermind the cores. Just hand over the cache”

HPCWire December 07:
More Than 16 Cores May Well Be Pointless
[by Sandia Labs]

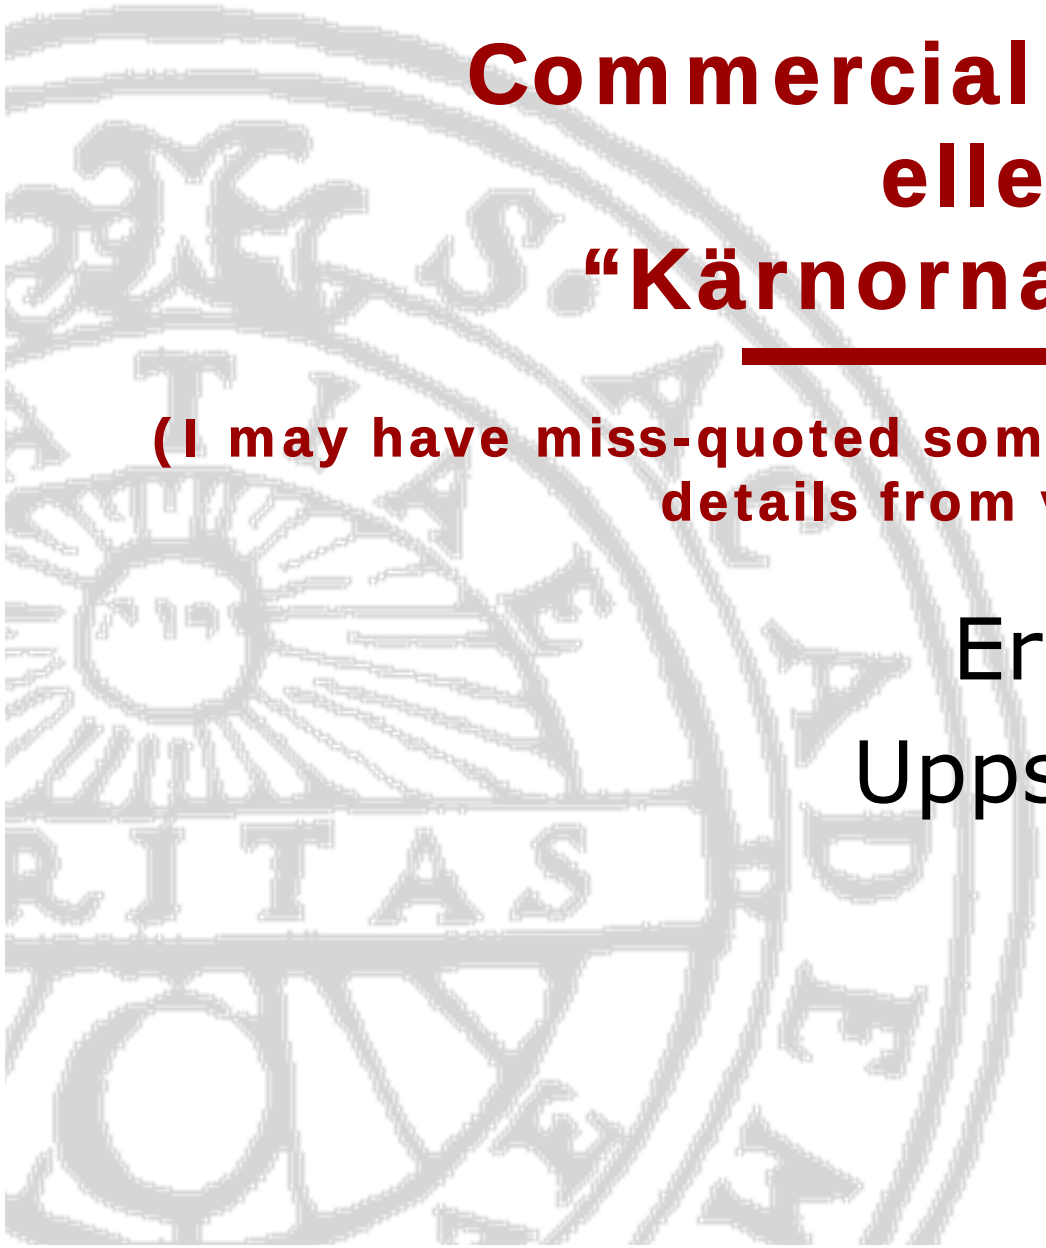

Commercial snapshot eller "Kärnornas krig"

(I may have miss-quoted some details, get architecture details from vendors)

Erik Hagersten

Uppsala Universitet

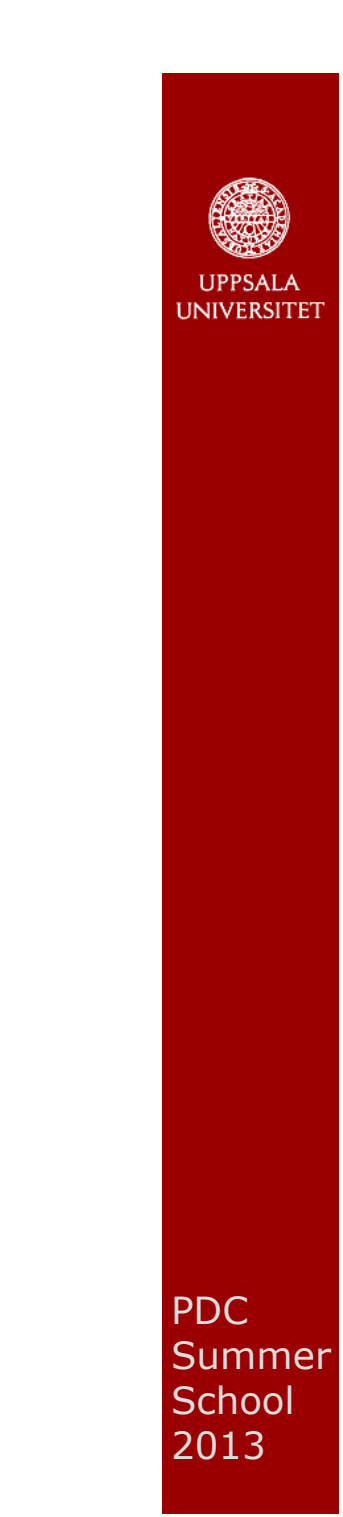

Intel Core2 Quad, 2006

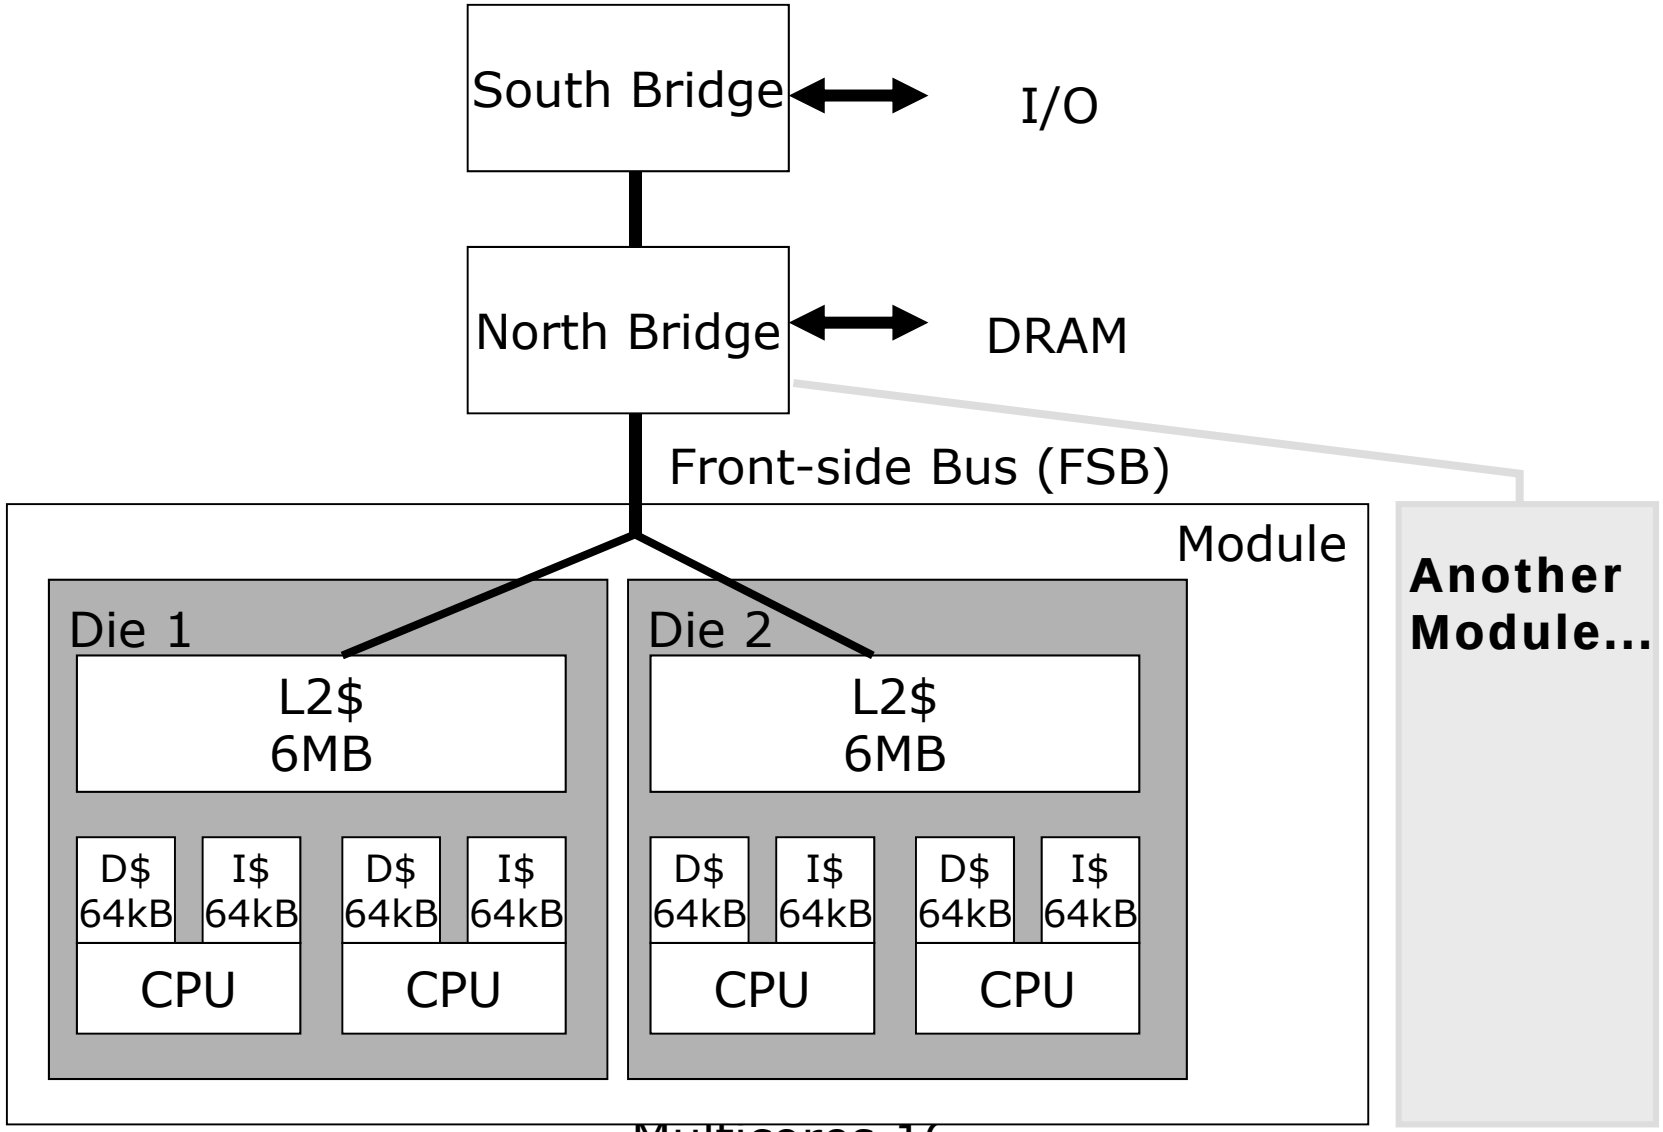

AMD Shanghai, 2007

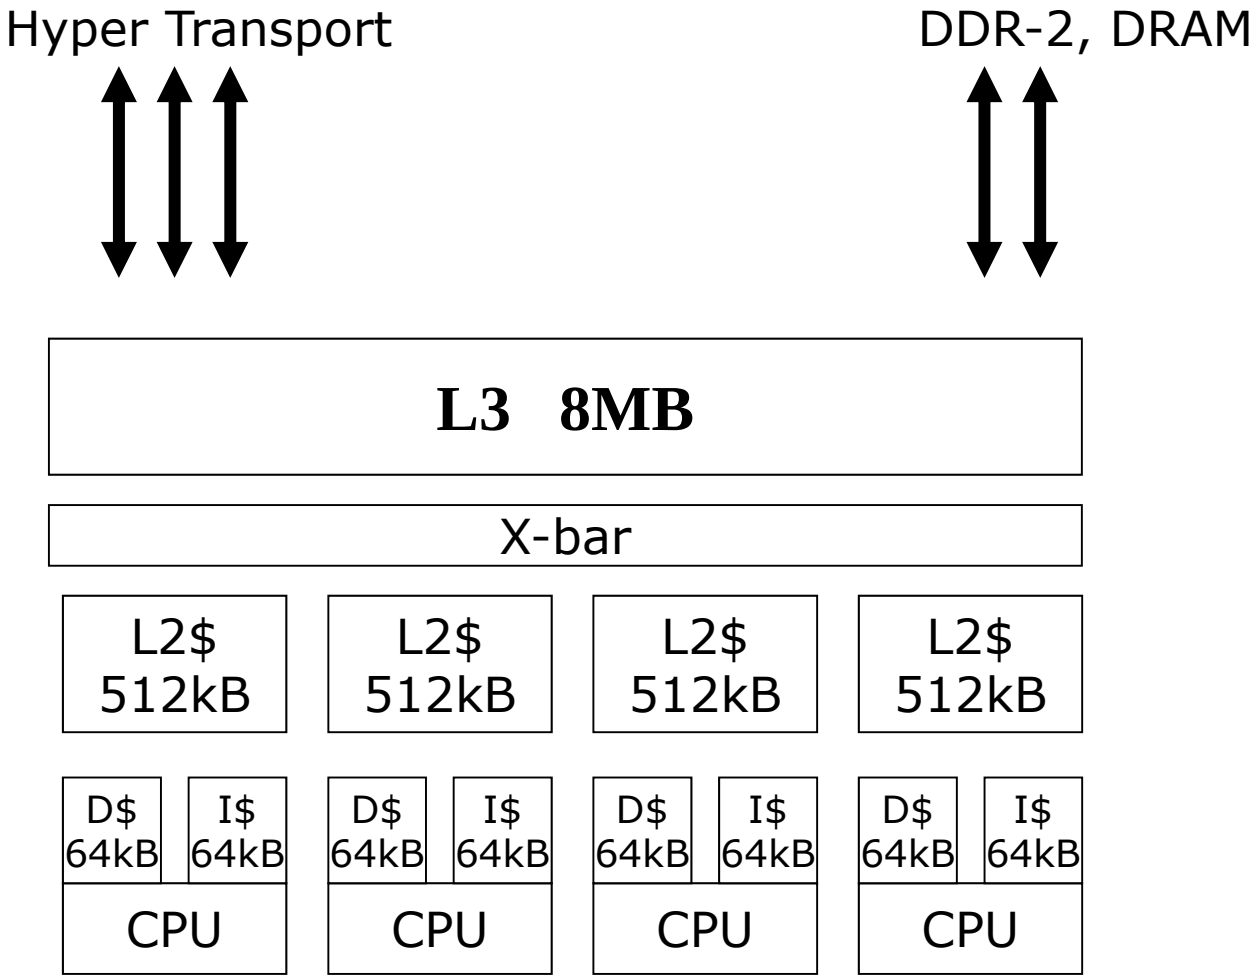

AMD MC System Architecture

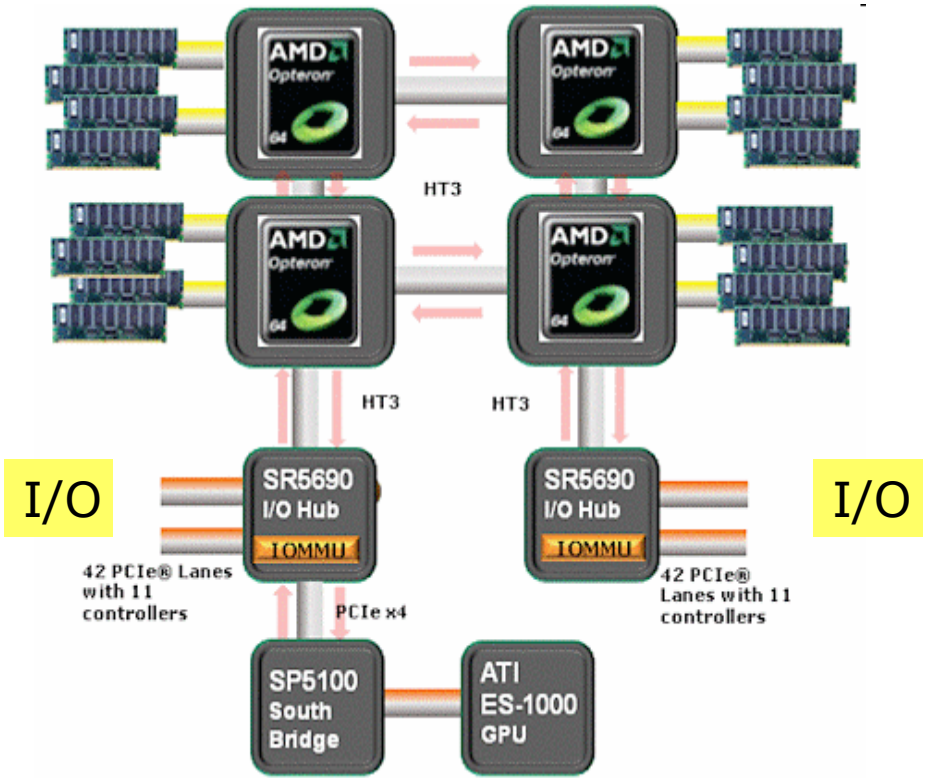

Intel: Nehalem, Core i7 Q1 2009 (4 cores)

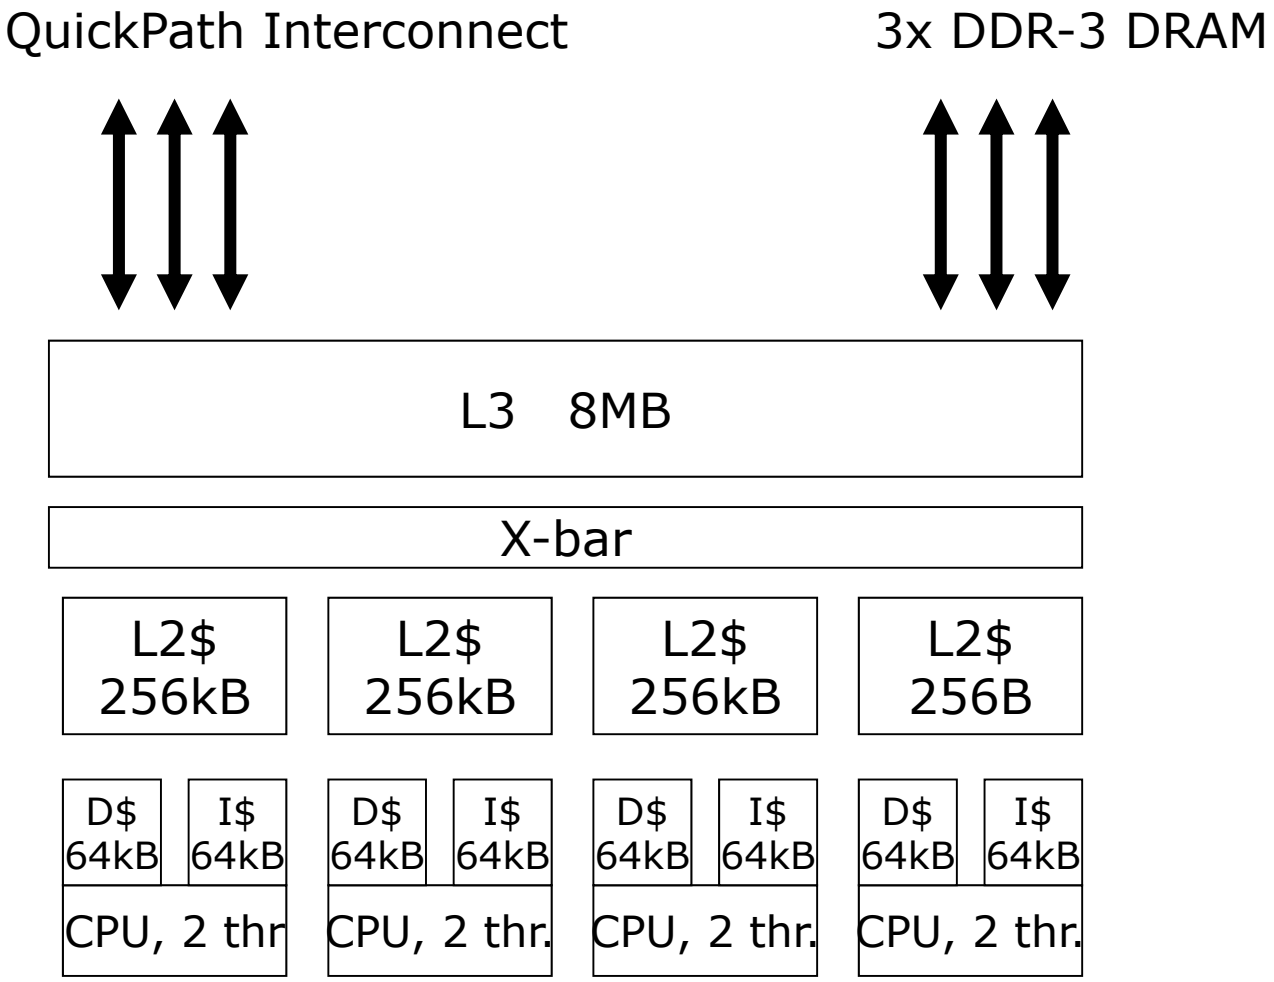

Up to 4 cores x 2 threads

Nehalem ”Core i7”

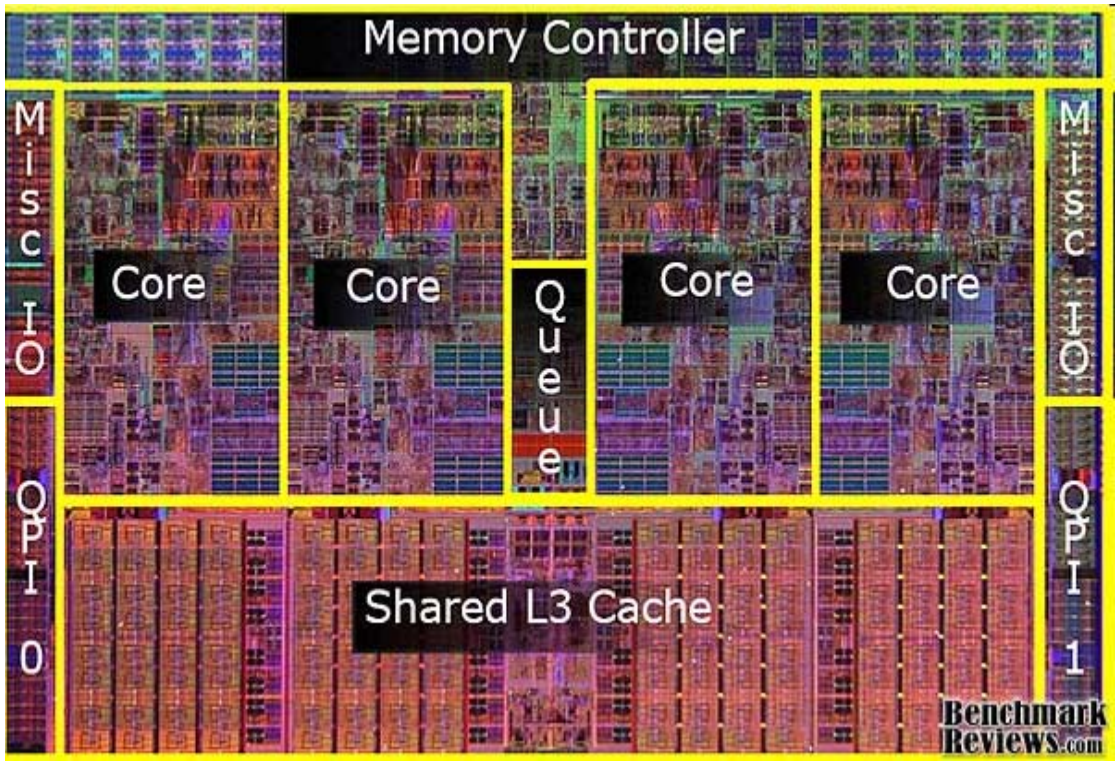

UPPSALA UNIVERSITET

Intel: ”Nehalem-Ex” (i7)

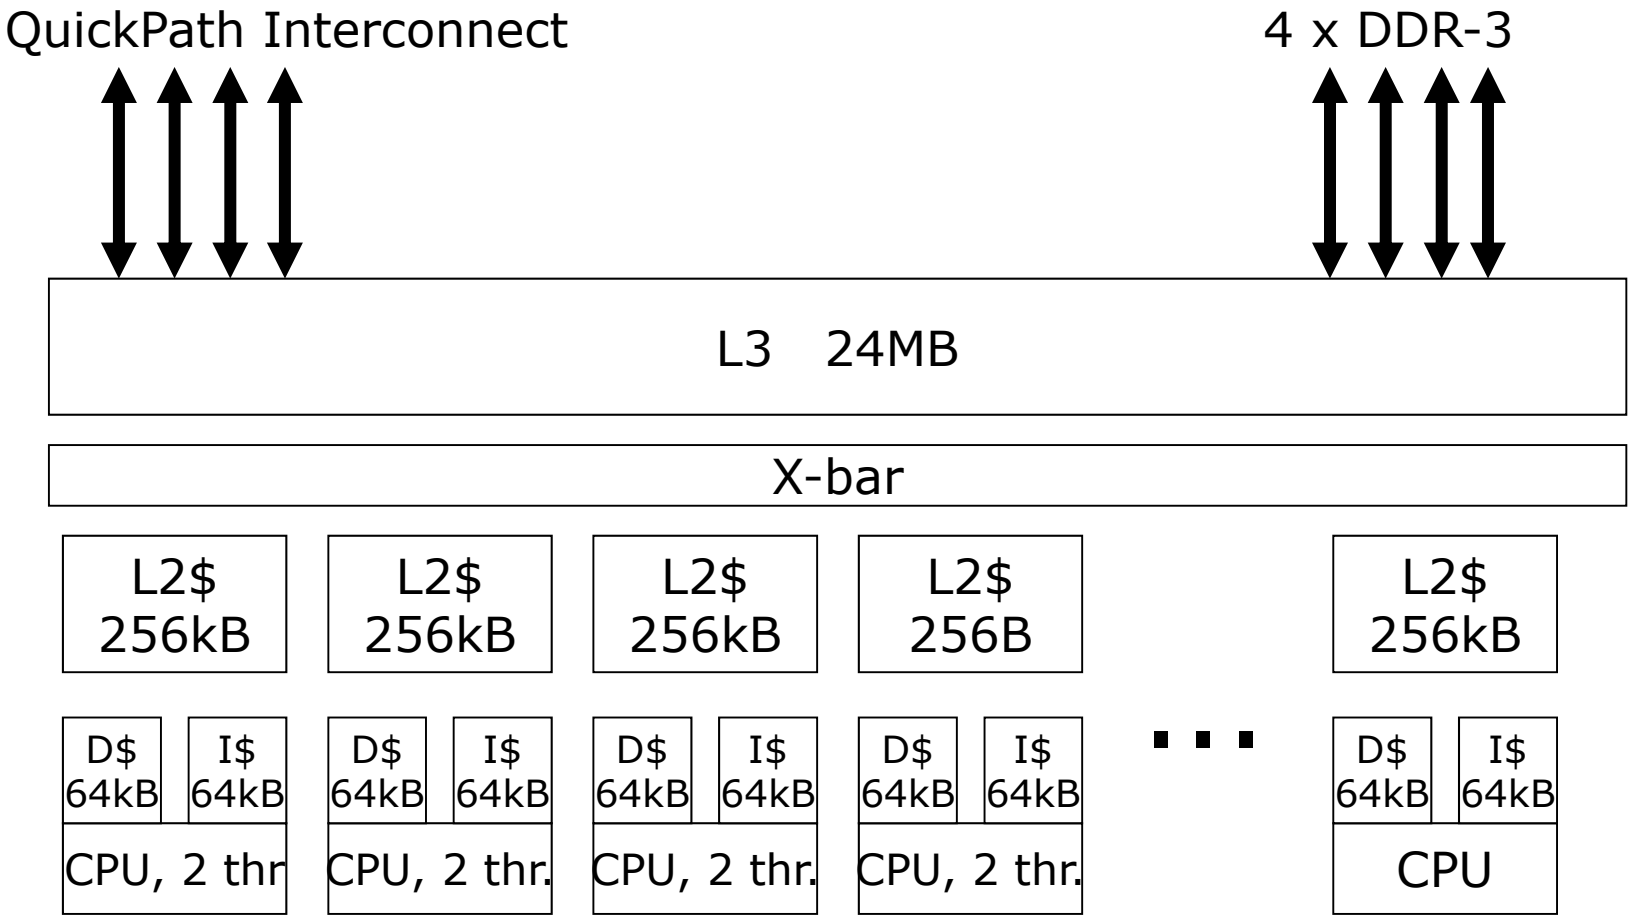

8 cores x 2 threads

Multicores 21

How is the silicon used (i7-Ex)?

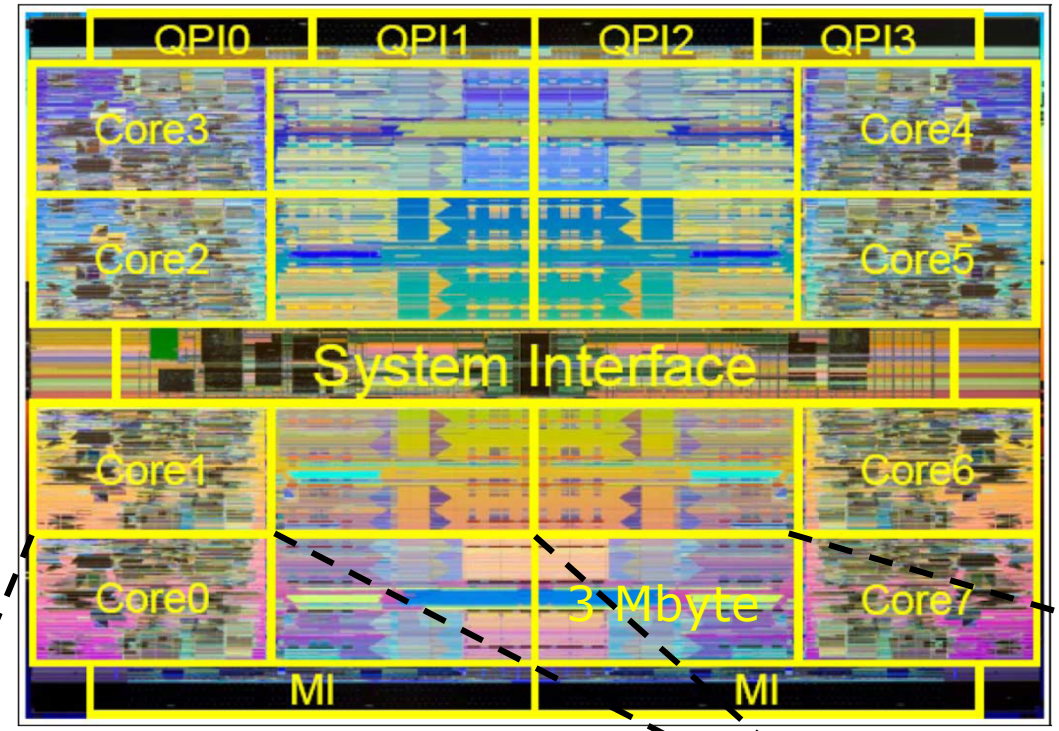

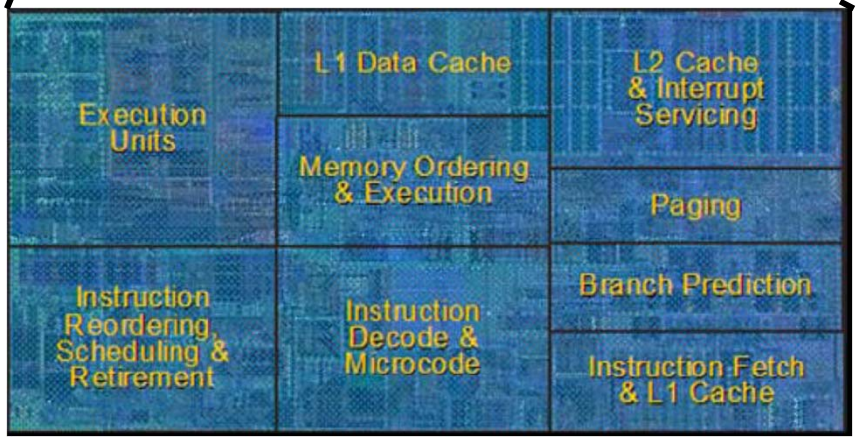

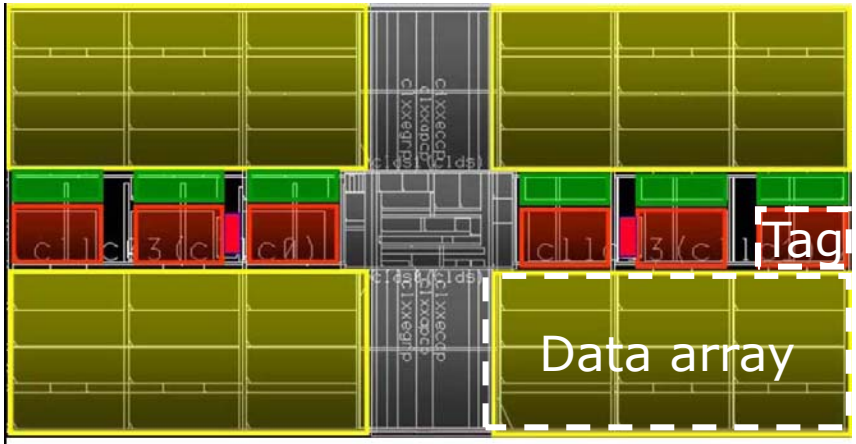

Source: JSSC Jan 2010, Rusu et. al

Multicore

UPPSALA UNIVERSITET

How is the silicon used?

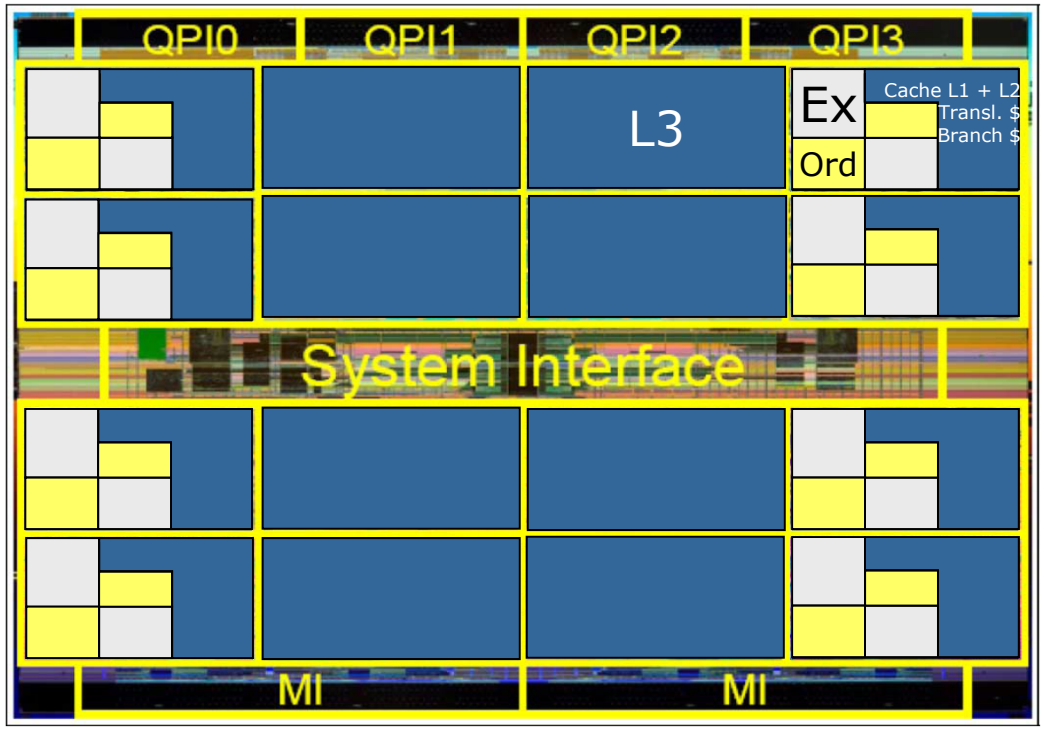

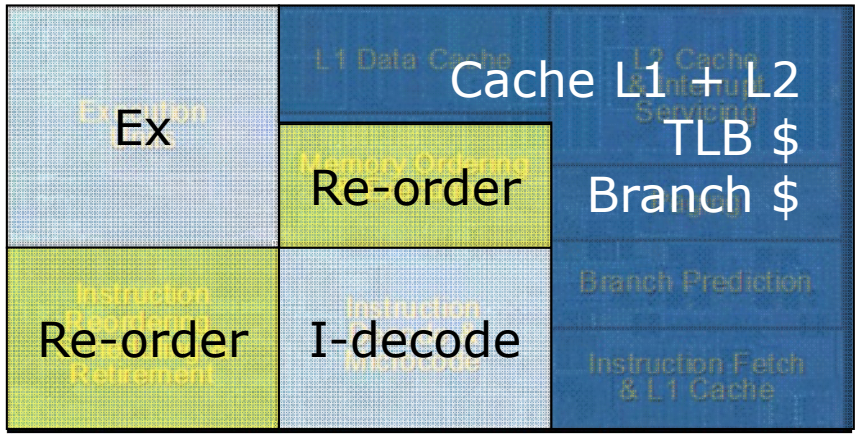

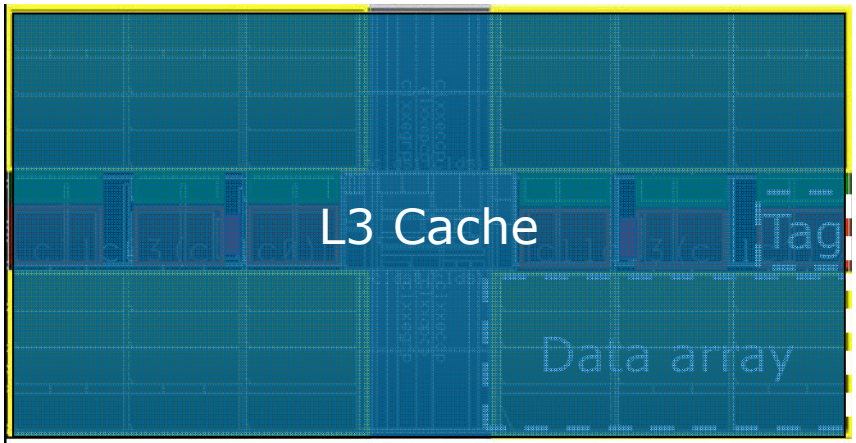

Multicore

Source: JSSC Jan 2010, Rusu et. al

More recent:
Sandy Bridge, 32 nm & Ivy Bridge, 22nm

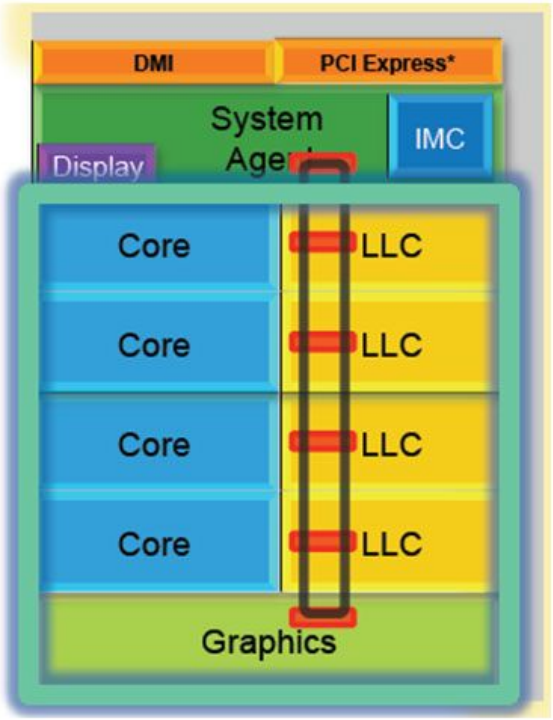

- Integrated graphics!
- Ring bus for coherent LLC
- Up to 12 cores (24 threads)
- ~3GHz-ish

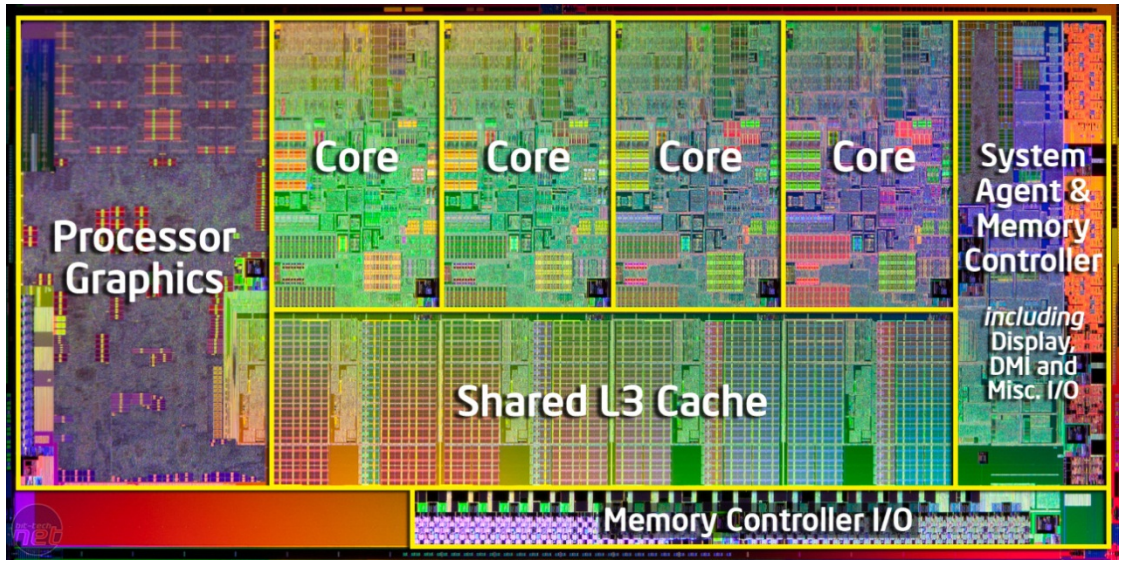

UPPSALA UNIVERSITET

General Purpose:
Is there a demand for more cores?

- Intel released their quad core at Supercomputing 2006.
- 7 years later we are only(!!) at 12 cores and 24 threads.
- Core complexity and caches are favoured over number of cores

Still, certain market segments and special applications need cores...

UPPSALA UNIVERSITET

Hot Chips Conference 2012

" 16-core SPARC T5 CMT Processor with glueless 1-hop scaling to 8-sockets"

" Knights Corner, Intel's first Many Integrated Core (MIC) Architecture Product"

Sun Niagara T1, 2005 Targeting transactions and DB

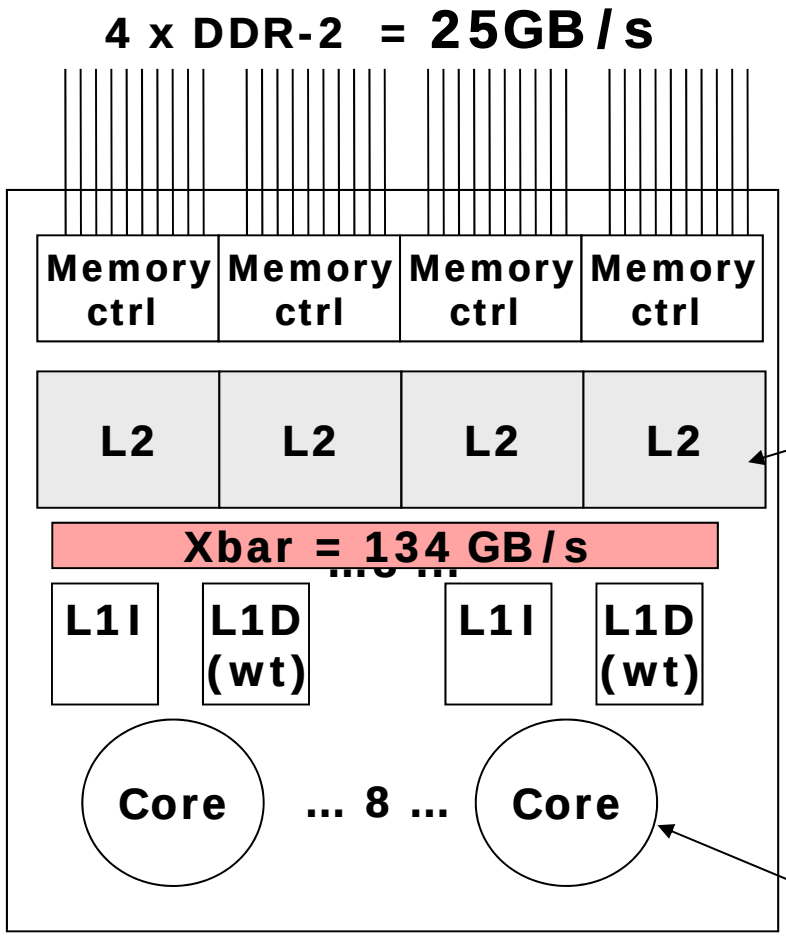

In-order 5 stage pipeline
1.2 GHz
Four interleaved threads

UPPSALA UNIVERSITET

2012: SPARC Niagra T5

- 16 cores
- 8 threads per core (128 threads/chip)
- 2x superscalar out-of-order @3.6GHz
- 8MB L3 cache (64kB/thread)
- Glueless 8 multisockets (1-hop)
- Coherent global NUMA memory
- Lots of HW accelerations

UPPSALA UNIVERSITET

Oracle marketing slide: Speeds and feeds, 8 sockets

DDR3-1066
1+TB/sec

Coherency Bisection Bandwidth 840 GB/sec

PCI Gen3 Bandwidth 256 GB/sec

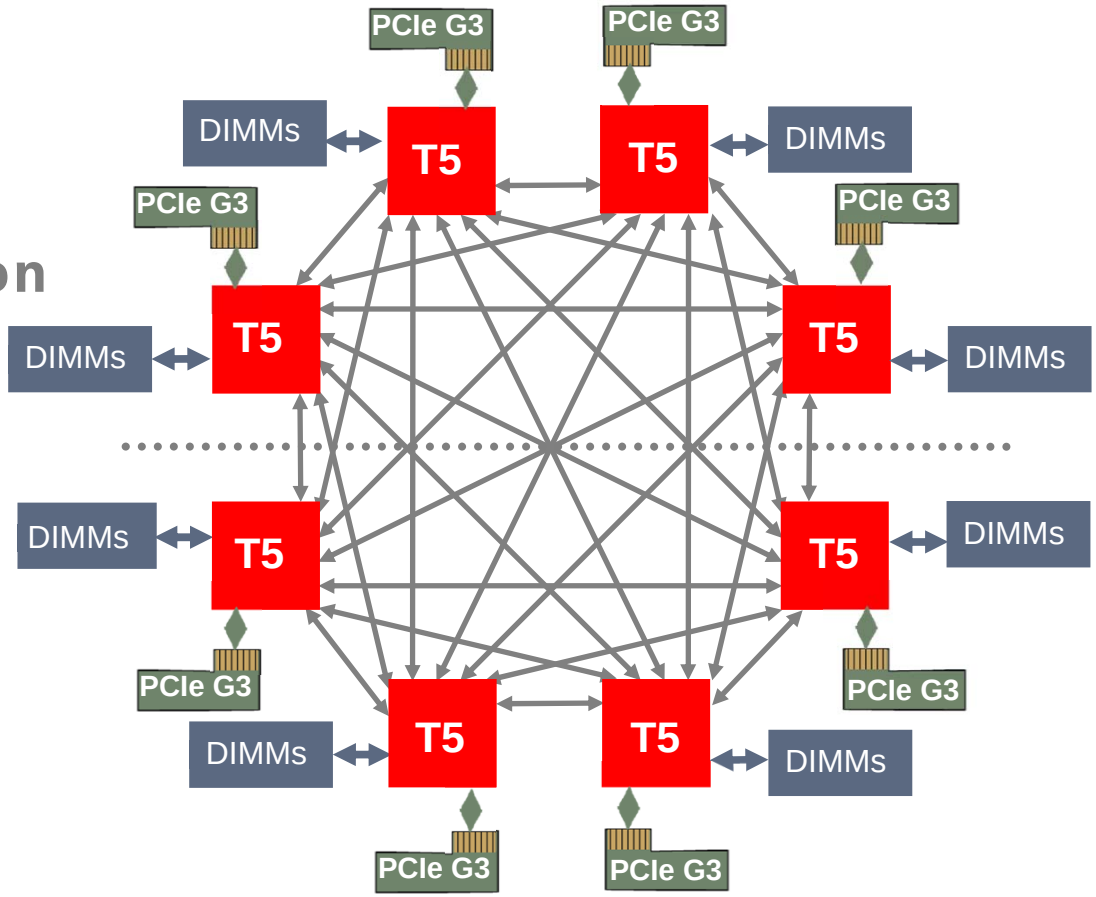

PDC Summer School 2013

UPPSALA UNIVERSITET

Oracle marketing slide: Special purpose HW

On-Chip Accelerators	SPARC T5	IBM Power7, HP	Intel Westmere/ Sandybridge
Asymmetric /Public Key Encryption	RSA, DH, DSA, ECC	none	RSA, ECC
Symmetric Key / Bulk Encryption	AES, DES, 3DES, Camellia, Kasumi	none	AES
Message Digest / Hash Functions	CRC32c, MD5, Sha-1, SHA-224, SHA-256, SHA-384, SHA-512	none	none
Random Number Generation	Supported	none	Supported

T5 Processor Overview

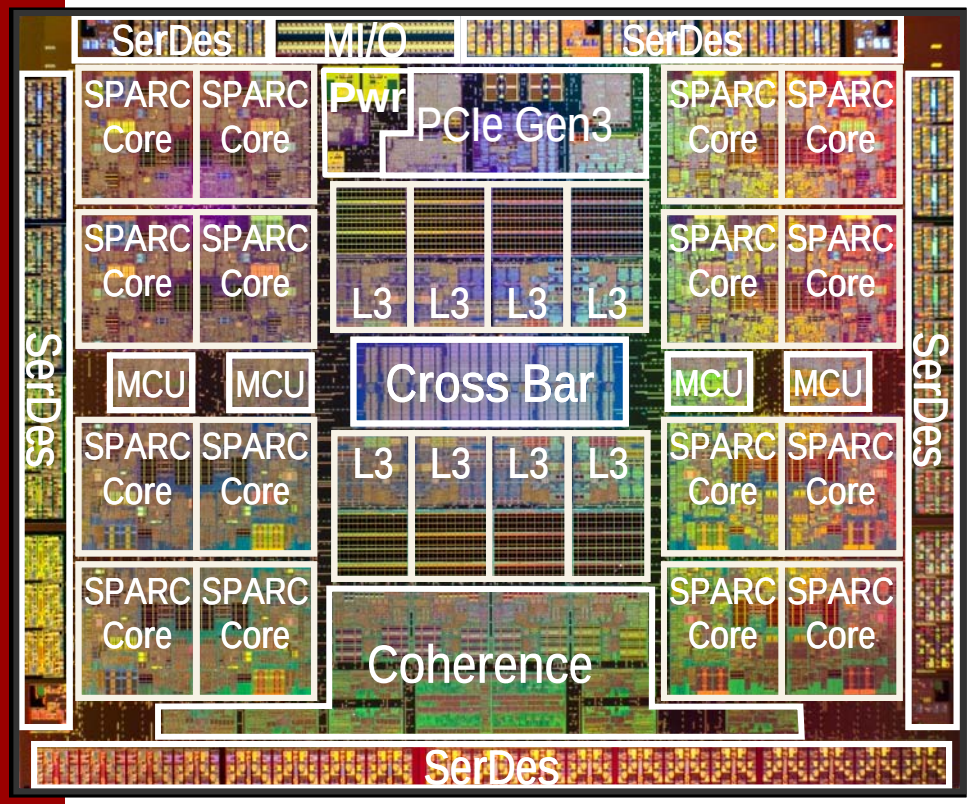

- 16 S3 cores @ 3.6GHz
- 8MB shared L3 Cache
- 8 DDR3 BL8 Schedulers providing 80 GB/s BW
- 8-way 1-hop glue-less scalability
- Integrated 2x8 PCIe Gen 3
- Advanced Power Management with DVFS

TILERA Architecture Targeting Embedded

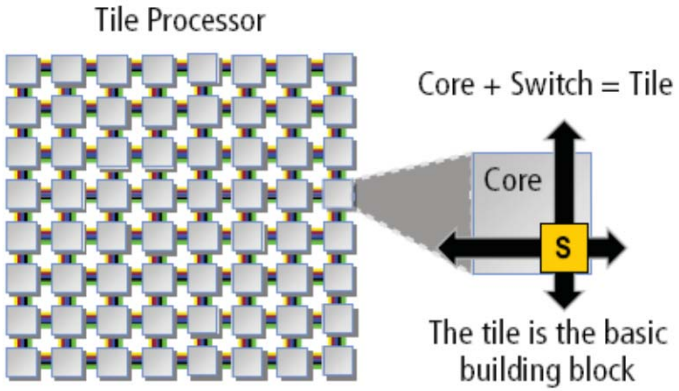

64 cores connected in a mesh
Each core:

- **MIPS core + Local L1 + L2 caches**
- **A slice of the shared distributed L3 cache**

Linux + ANSI C
New Libraries
New IDE
Stream computing
…

UPPSALA UNIVERSITET

GPUs and Accelerators: Lots of hype in HPC!

Fermi from nVIDIA -- a huge step in the right direction

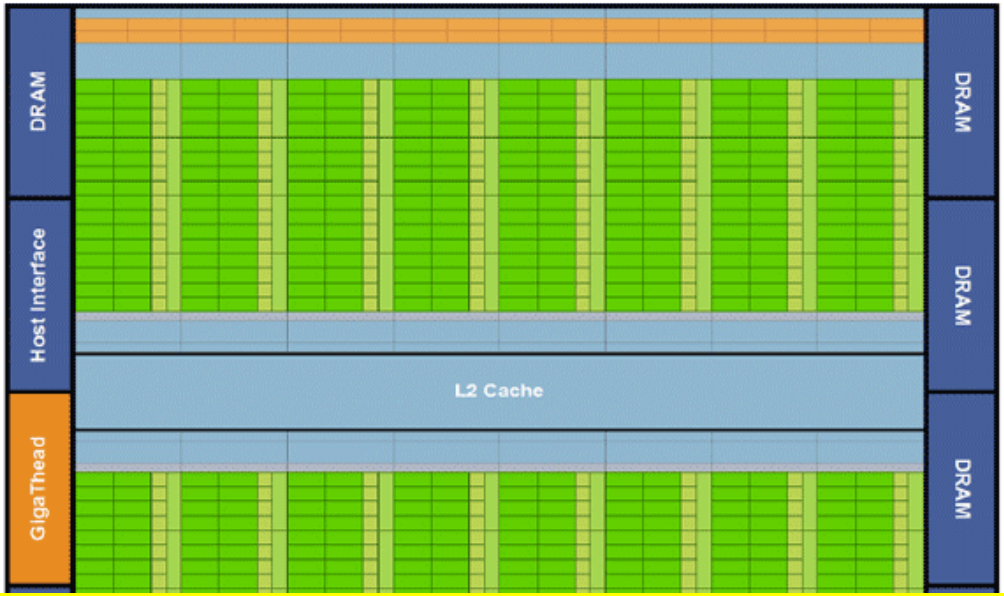

- 512 ”cores” (***P***)
- 16 ***P***/StreamProcessor (***SP***)
- Full DP-FP IEEE support
- 64kB L1 cache /***SP***
- 768kB global shared cache
- Atomic instructions
- ECC correction
- Debugging support
- …

- **CUDA language. Requires ”heroic programming” [M.Bull, EPCC]**
- **How much is an nVIDIA core worth vs. an x86 core?**
- **nVIDIA used to say 100x performance [over some old x86]**
- **At SC2011, nVIDIA’s banner said ”2x peformance guaranteed”**
- **With what productivity?**

++ New interestring programming approach announced: OpenACC
David BS to tell you more…

MIC = Multiple Integrated Cores
Xeon Phi (Knights Corner)
first commercial MIC implementation

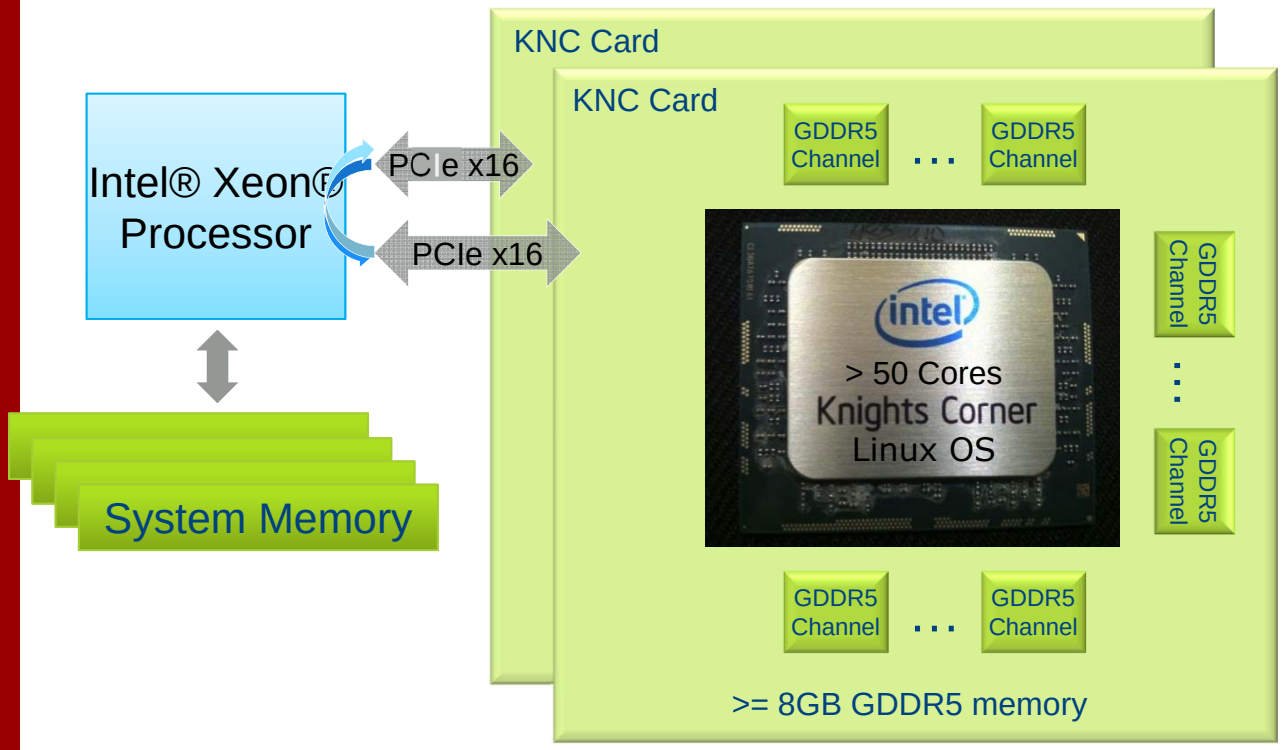

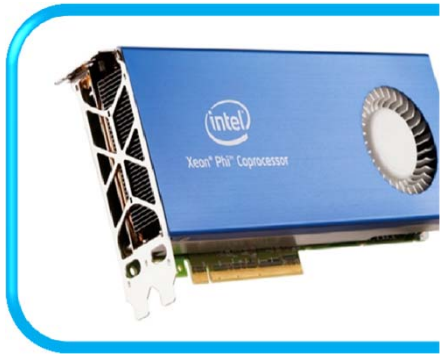

UPPSALA UNIVERSITET

Intel's MIC HPC architecture

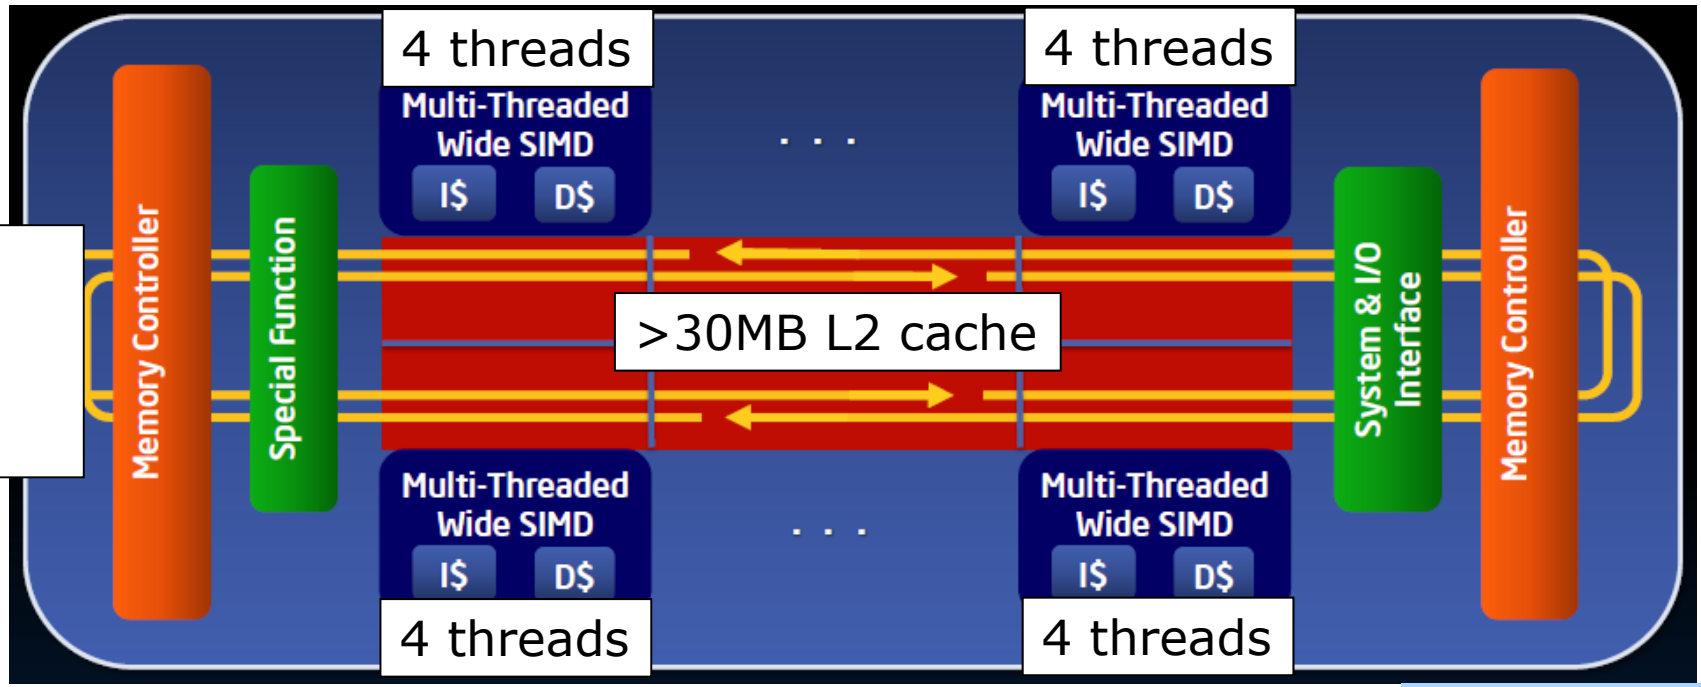

??= Unknown. Check the Web for updates

- Knights Corner (Intel Xeon Phi), 22nm
- >60 cores, X86 instructions ("enhanced"), 4 Threads
- Coherent caches. Runs Linux. "Just recompile and run"
- > 1TFLOPS (DP).
- 2% of the transistors devoted to x86 instructions

UPPSALA UNIVERSITET

The MIC core ”tile”

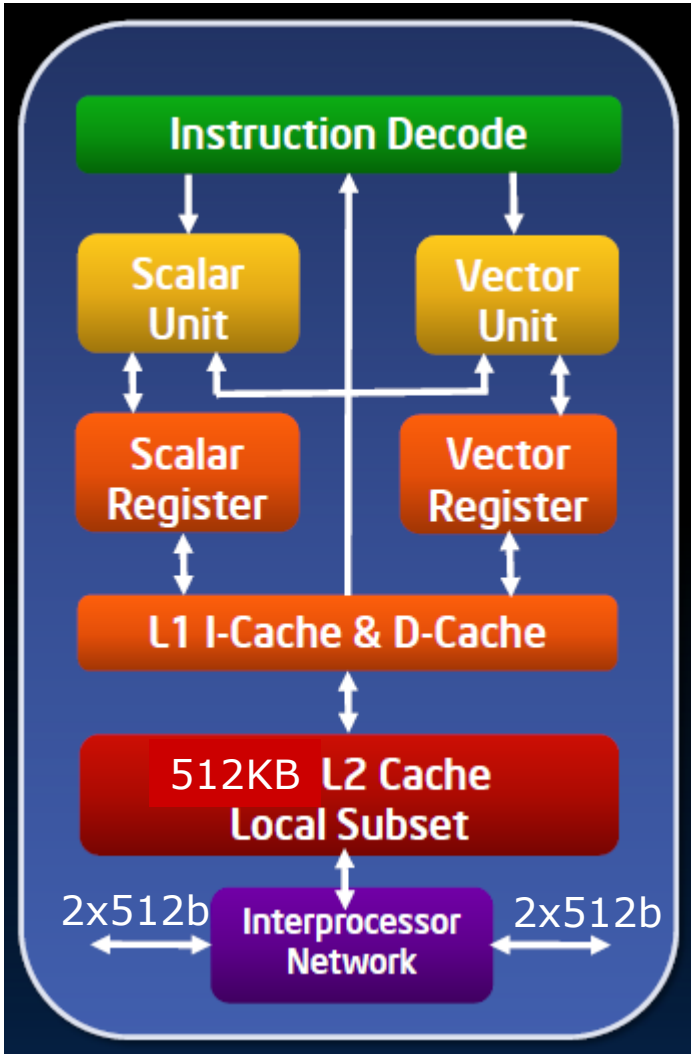

- 2x superscalar in-order
- 4 Threads SMT/core
- 512b SIMD instr. (16x SP)
- Scatter/gather operations
- 32k L1 I and D cache
- 512kB coherent L2
- 16 HW prefetch streams
- 2x512bit ringbus
- …

UPPSALA UNIVERSITET

MIC Tool Stack

- It is ”just an x86 Linux box”
- Use your favorit language
- Intel compiler, tuning tool etc.
- 3rd party tools (for example mine:-)
- A huge alpha program since a year back
- Already on the Top500 list (#150)
- Low power/FLOP
- Many strategic sales already
- Question: can we afford 200+ threads/chip

Multicores: A different beast

Dr Dobb's 2005: The Free Lunch Is Over
A Fundamental Turn Toward Concurrency in Software
By Herb Sutter

"The biggest change in software development since the object-oriented revolution is knocking at the door, and its name is Concurrency…

Dr Dobb's 2012: Welcome to the Jungle
By Herb Sutter

"Mainstream computers […] are being permanently transformed into heterogeneous supercomputer clusters. Henceforth, a single compute-intensive application will need to harness different kinds of cores, in immense numbers, to get its job done.".

UPPSALA UNIVERSITET

Trends 1 (2)

Varies with application areas:

- Number of cores,
- Cache capacity,
- Bandwidth,
- Threads/core,
- Cache/thread varies.

- Important SW functions moved to HW

➔ HW diversity is larger than ever

➔ Intel has never released so many different chips per year

UPPSALA UNIVERSITET

Trends 2

- Active power management (DVFS)
- Non-uniformity
- Heterogeneity
  - tiny/fat cores in the same system
  - different ISA
  - specialized HW
- Dark silicon is around the corner (!!)
  - Cannot have all transistors turned on
  - Reconfigurable architectures

PDC Summer School 2013

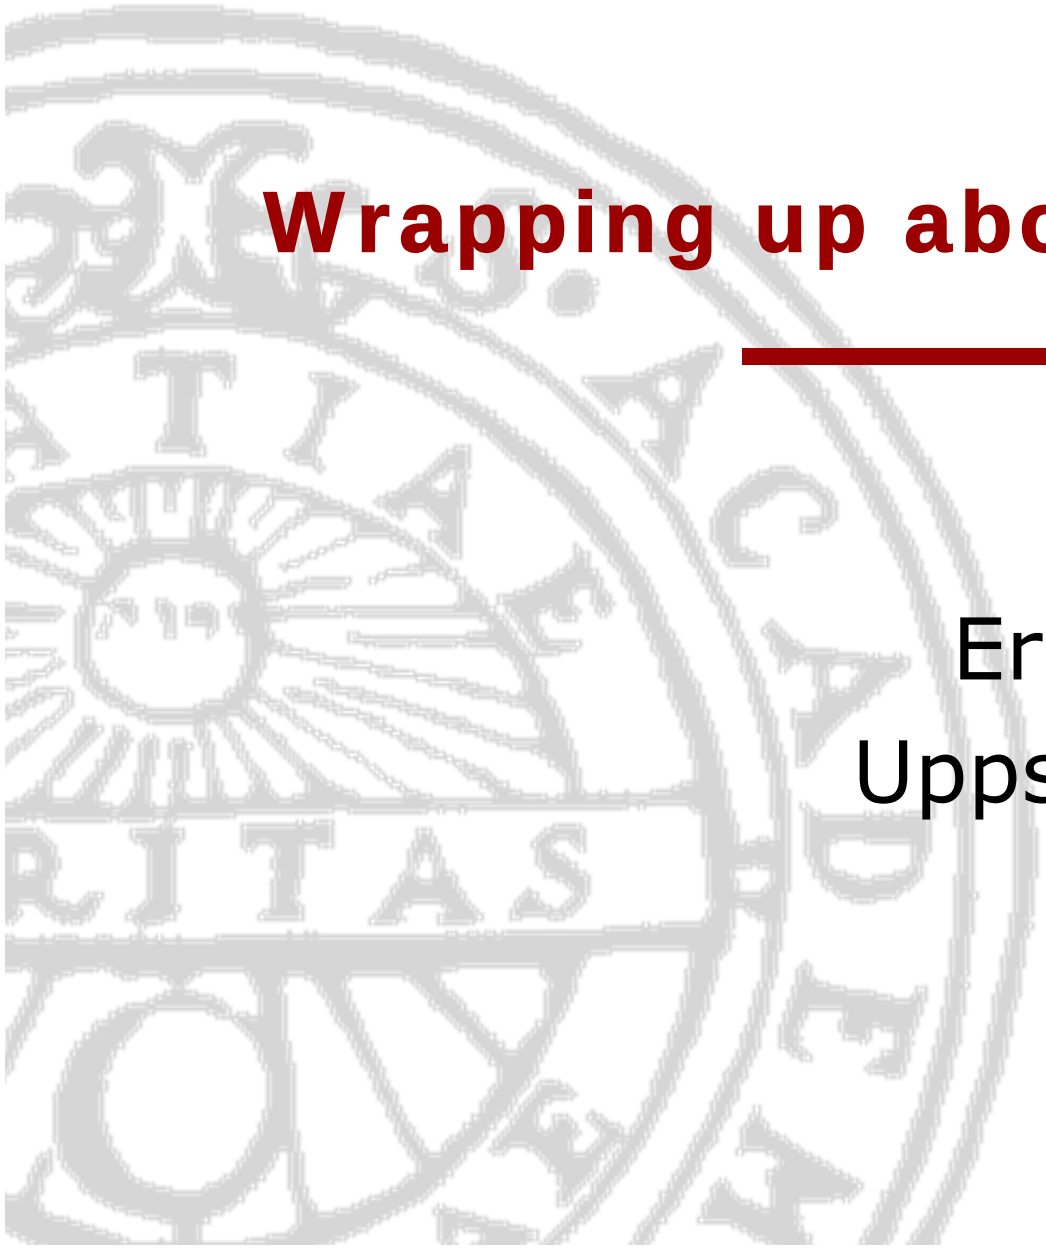

Wrapping up about multicores

Erik Hagersten
Uppsala Universitet

UPPSALA UNIVERSITET

What matters for multicore performance?

- Are we buying…
  - CPU frequency?
  - Number of cores?
  - MIPS and FLOPS?
  - Memory bandwidth?
  - Cache capacity?
  - Memory capacity?
  - Performance/Watt?

UPPSALA UNIVERSITET

MC Questions for the Future

- How to get parallelism?
- How to get performance/data locality?
- How to debug?
- A case for new funky languages?
- A case for automatic parallelization?
- Are we buying:
  - compute power,
  - memory capacity, or
  - memory bandwidth?
- Will 1000 cores be mainstream in 5 years?
- Will the CPU market diverge into desktop/capacity/capability/special-purpose CPUs again?
- **A non-question: will it happen?**