



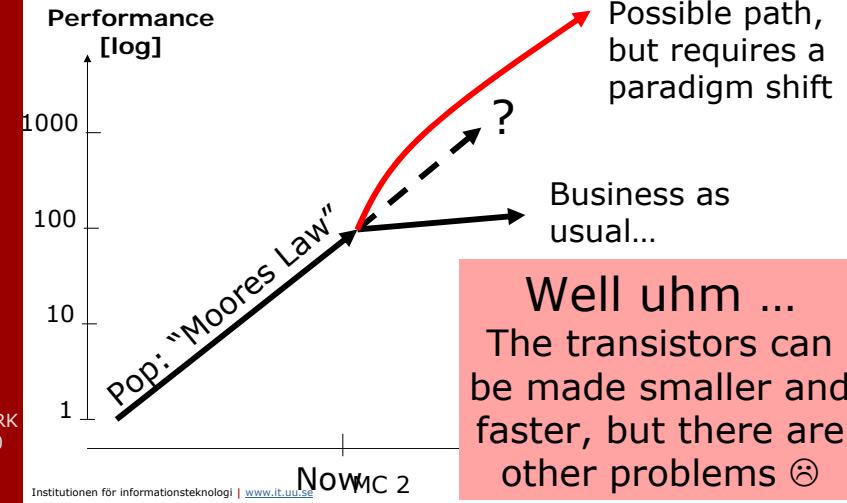
Multicore: Why is it happening now? eller Hur Mår Moore's Lag?

Erik Hagersten
Uppsala Universitet



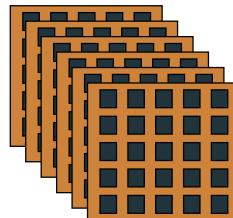
Are we hitting the wall now?

Pop: Can the transistors be made even smaller and faster?

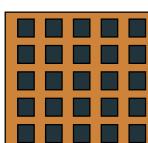


Darling, I shrunk the computer

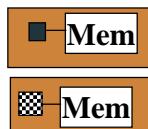
Mainframes



Super Minis:



Microprocessor:



Multicore: Many CPUs on a chip!

AVDARK
2010



Multi-core CPUs:

- Agila PhysX, a multi-core physics processing unit.
- 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 Mali, a programmable multi-core 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.
 - Xeon, 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 (Synergistic 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 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.
 - Teraflop 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 10 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 PB8K32, an eight-core, 32-thread microcontroller
 - picochip RC200 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 proc
 - UltraSPARC T2, an eight-core, 64-concurrent core
 - Texas Instruments T32KC400 MVP, a five-core microcontroller
 - Tilera TILE64, a 64-core processor
 - XMOS Software Defined Silicon quad-core XS1-G4

[source: Wikipedia]

AVDARK
2010

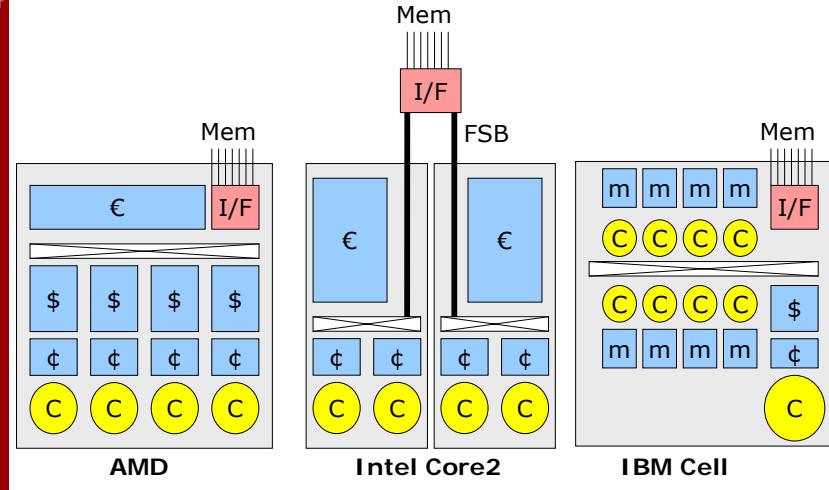
High-performance single-core CPUs:

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 5

© Erik Hagersten | user.it.uu.se/~eh

Many shapes and forms...

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 6

© Erik Hagersten | user.it.uu.se/~eh

Outline

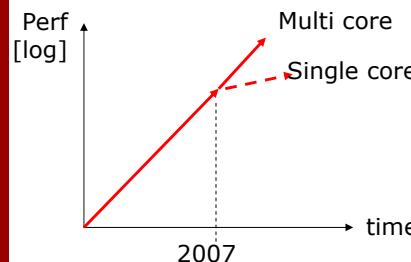
- Why multicore now?
- Technology reasons
- Performance bottlenecks in MCs
- Commercial offerings
- Reflection for the future

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 7

© Erik Hagersten | user.it.uu.se/~eh

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 $P_{dyn} \sim C \cdot f \cdot V^2$

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 8

© Erik Hagersten | user.it.uu.se/~eh

Microprocessors today: Whatever it takes to run one program fast.

Exploring ILP (instruction-level parallelism):

- Faster clocks → Deep pipelines
- Superscalar Pipelines
- Branch Prediction
- Prefetching
- Out-of-Order Execution
- Trace Cache
- Speculation
- Predicate Execution
- Advanced Load / Store
- Return Address
-

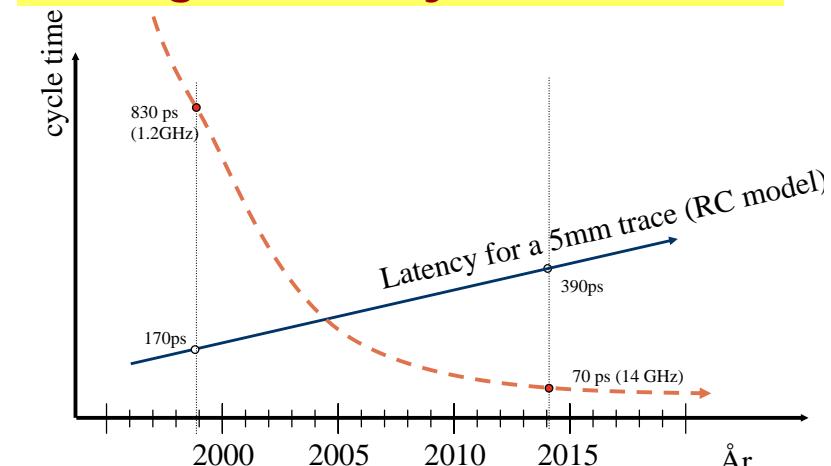
AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 9

© Erik Hagersten | user.it.uu.se/~eh

Bad News #1:
We have already explored most ILP
(instruction-level parallelism)

Bad News #2 Looong wire delay → slow CPUs



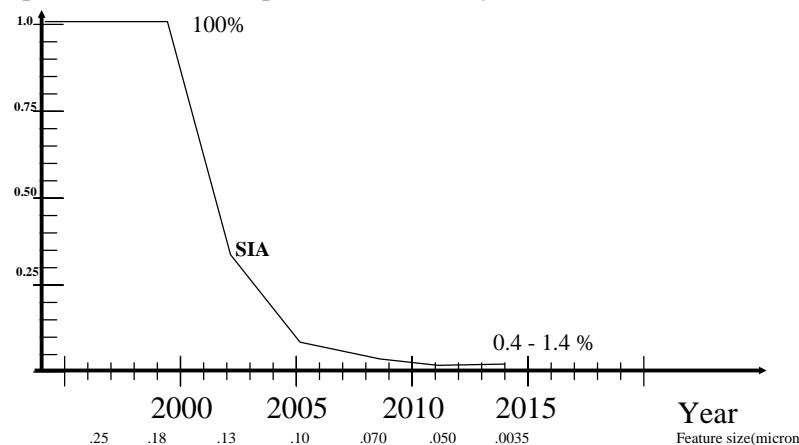
Quantitative data and trends according to V. Agarwal et al., ISCA 2000
Based on SIA (Semiconductor Industry Association) prediction, 1999

MC 10

© Erik Hagersten | user.it.uu.se/~eh

Bad News #2 Looong wire delay → slow CPUs

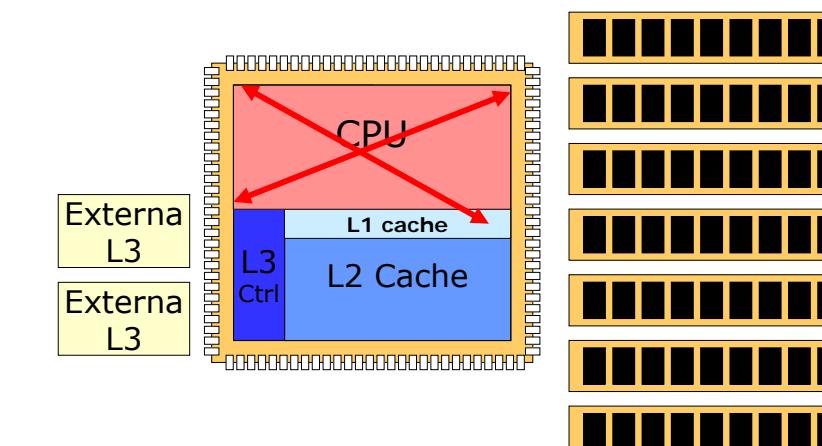
Span -- Fraction of chip reachable in one cycle

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 11

© Erik Hagersten | user.it.uu.se/~eh

Bad News #2: wire delay → Multiple small CPUs with private L1\$



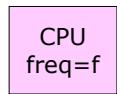
MC 12

© Erik Hagersten | user.it.uu.se/~eh

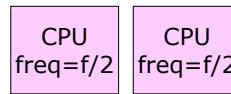
Bad News #3: Power consumption

→ Lower frequency → lower voltage

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

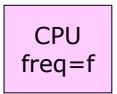


VS.

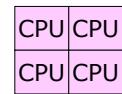


$$P_{\text{dyn}}(C, f, V) = CfV^2$$

$$P_{\text{dyn}}(2C, f/2, <V) < CfV^2$$



VS.



freq = f/2

$$P_{\text{dyn}}(C, f, V) = CfV^2$$

$$P_{\text{dyn}}(C, f/2, <V) < \frac{1}{2} CfV^2$$

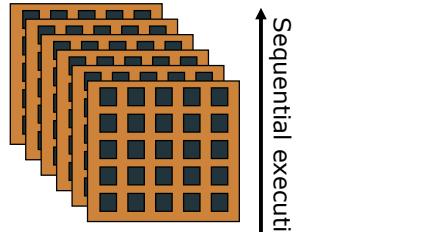
AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 13

© Erik Hagersten | user.it.uu.se/~eh

Darling, I shrunk the computer

Mainframes

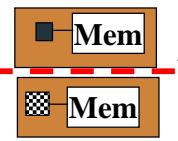


Sequential execution (\approx one program)

Super Minis:



Microprocessor:



Parallel Apps (TLP)

Paradigm Shift

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

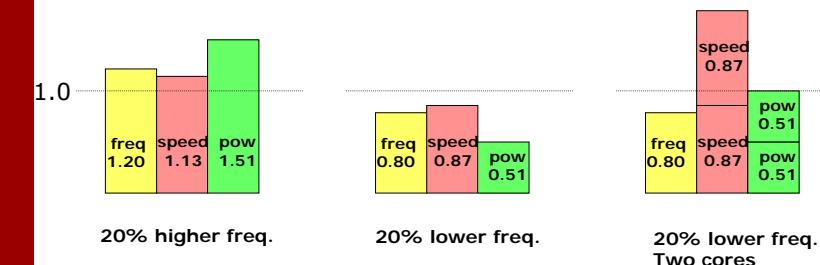
AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 15

© Erik Hagersten | user.it.uu.se/~eh

Example: Freq. Scaling

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

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 14

© Erik Hagersten | user.it.uu.se/~eh

How to create threads?

Not ideal for multi cores

Prog1, Prog2, OS

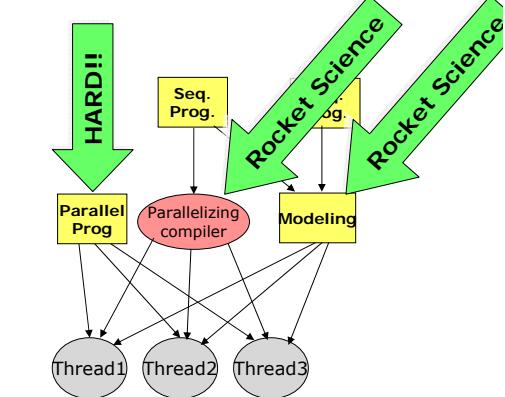
Thread1, Thread2, Thread3

Throughput Computing

- Multitasking
- Virtualization
- Concurrency
- ...

Parallel Prog, Parallelizing compiler

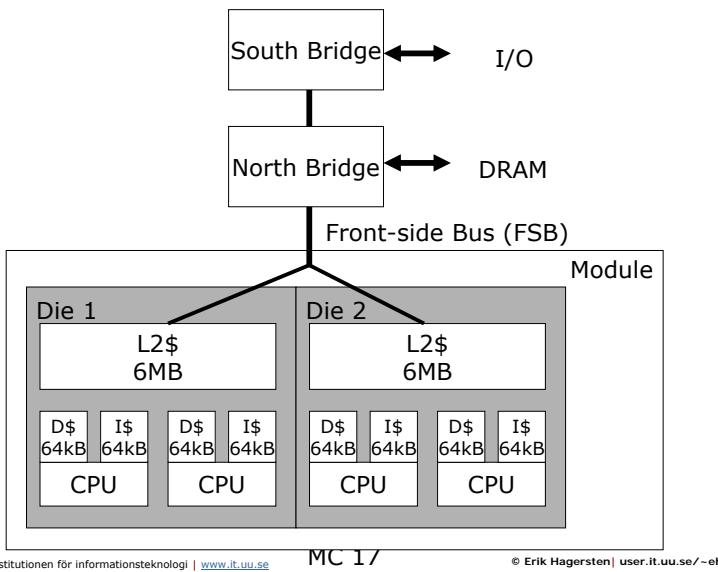
Seq. Prog., Modeling

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

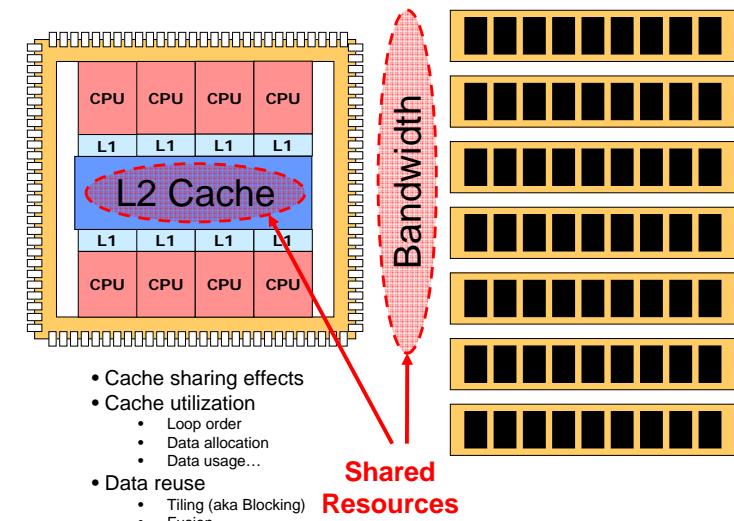
MC 16

© Erik Hagersten | user.it.uu.se/~eh

Intel Core2 Quad, 2006

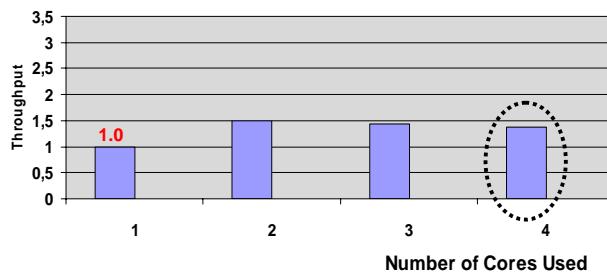


Shared Bottlenecks (the MCs on these slides are generalized)



Example: Poor Throughput Scaling!

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



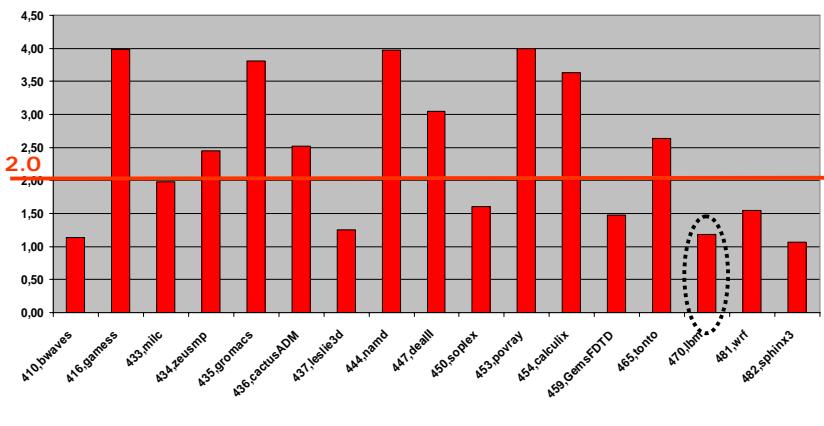
Throughput (as defined by SPEC):

Amount of work performed per time unit when several instances of the application is executed simultaneously.

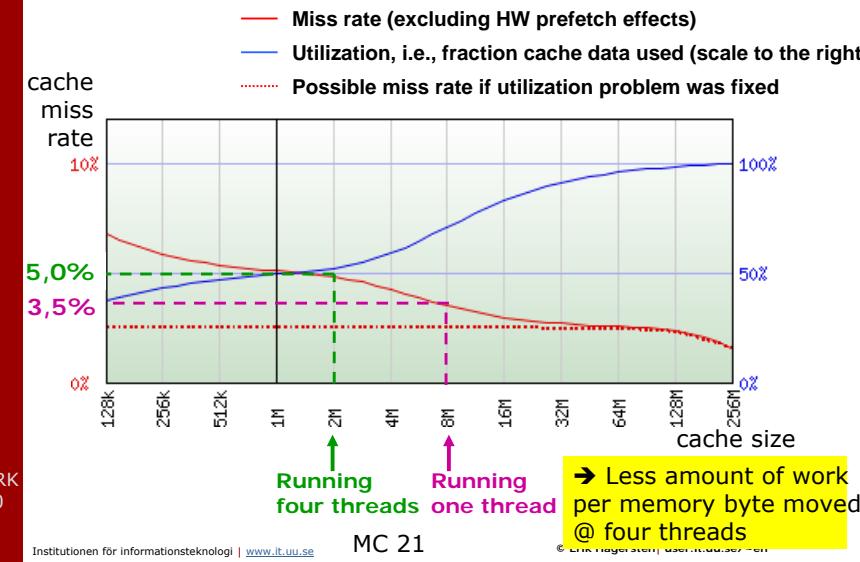
Our TP study: compare TP improvement when you go from 1 core to 4 cores

Throughput Scaling, More Apps

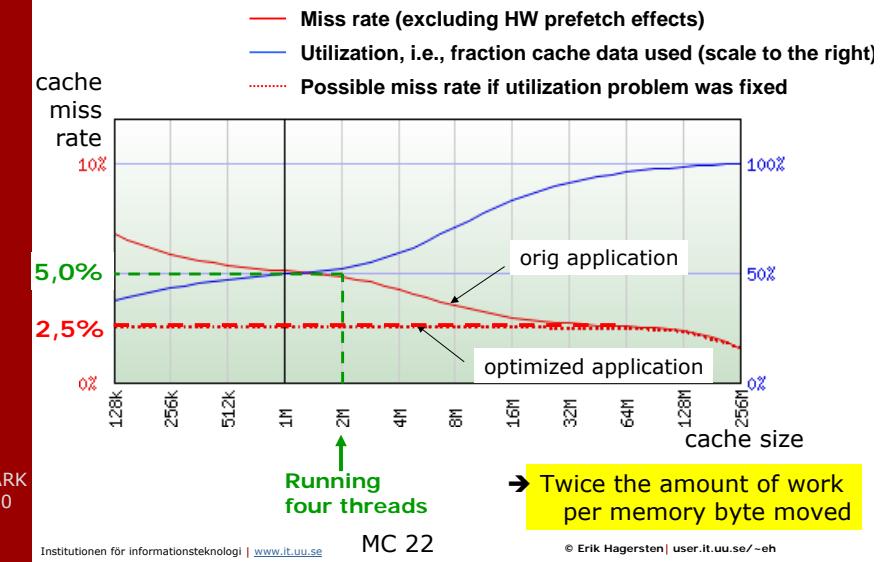
SPEC CPU 20006 FP Throughput improvements on 4 cores



Nerd Curve: 470.LBM

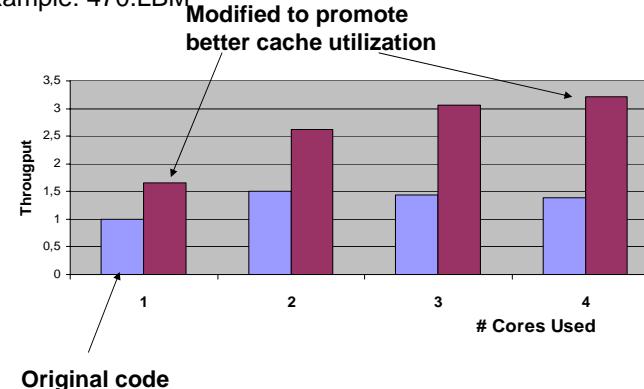


Nerd Curve (again)



→ Better Memory Usage!

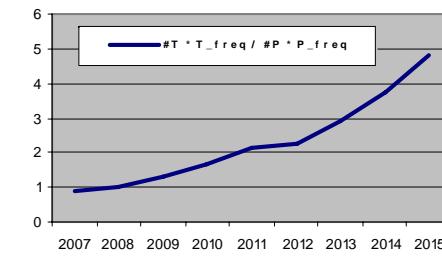
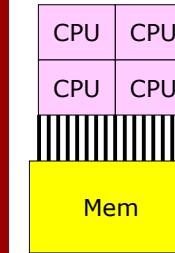
Example: 470.LBM



BW in the Future?

Computation vs Bandwidth

#Cores ~ #Transistors



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

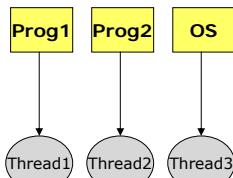
Map for Semiconductors (ITRS)

Chip Multiprocessors

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

Throughput Computing:

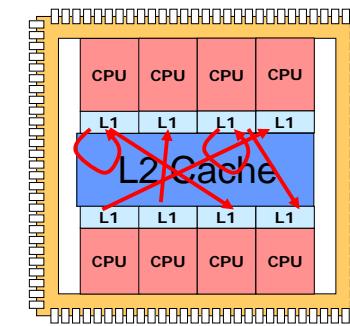
- Does not explore the shared caches well
(cache capacity/thread: $\sim 100\text{-}500\text{ kB}$)
- Requires N times more memory
- If limited by memory capacity or memory bandwidth
→ go parallel

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 25

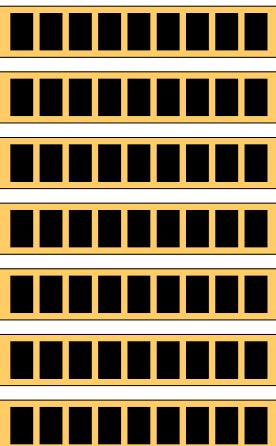
© Erik Hagersten | user.it.uu.se/~eh

Thread Interaction



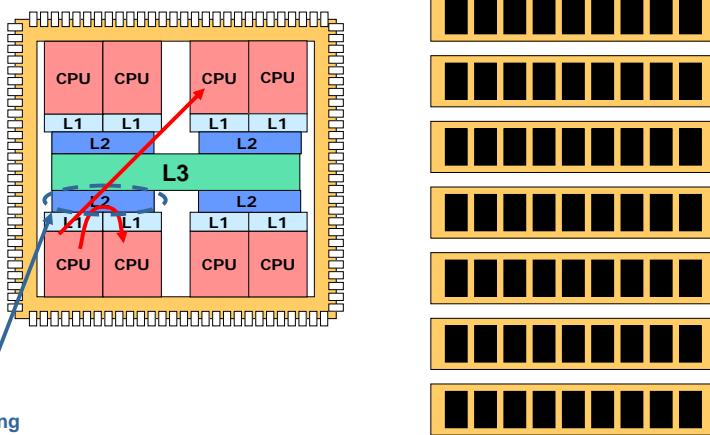
- Coherence traffic
- Communication utilization
- Load imbalance
- Synchronization
- False sharing
- ...

MC 26



© Erik Hagersten | user.it.uu.se/~eh

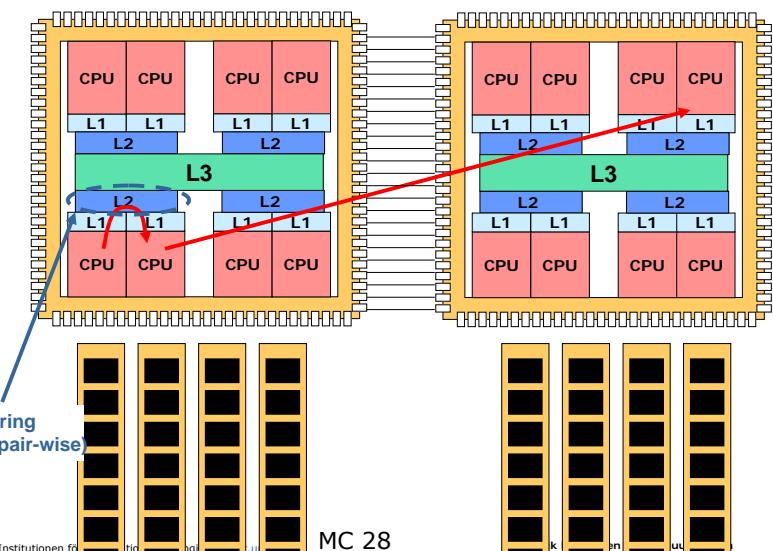
Multicore Challenges: Non-uniformity Communication

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 27

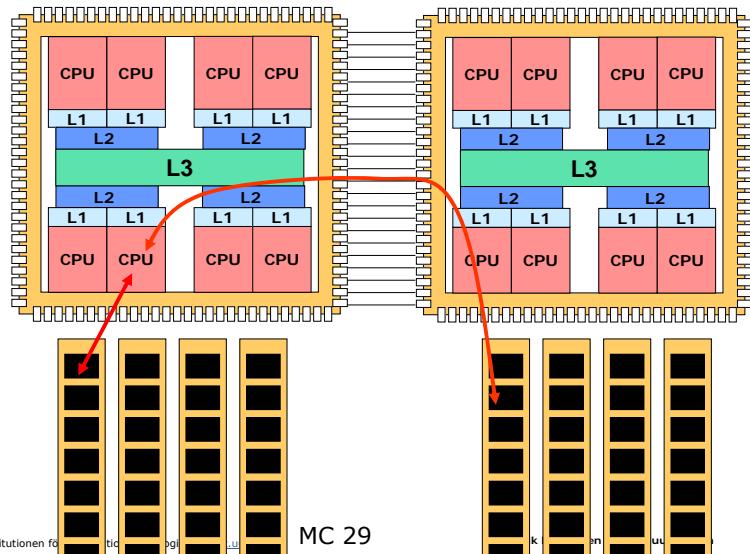
© Erik Hagersten | user.it.uu.se/~eh

Multicore Challenges (Multisocket) Non-uniformity Communication

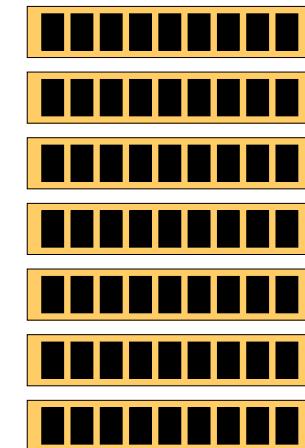
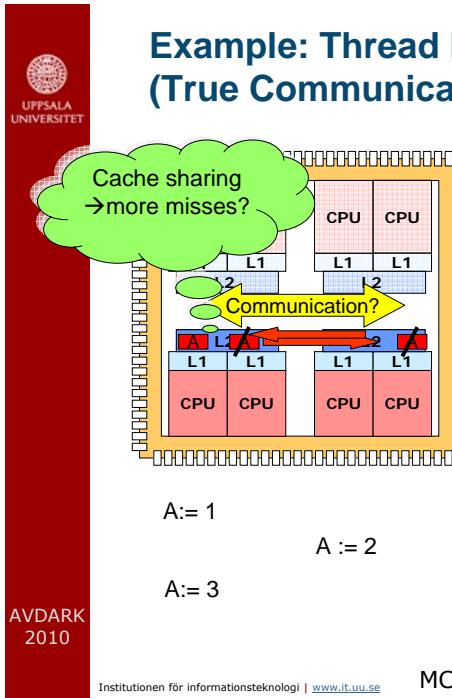
AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 28

Multicore Challenges (Multisocket) Non-uniformity Memory



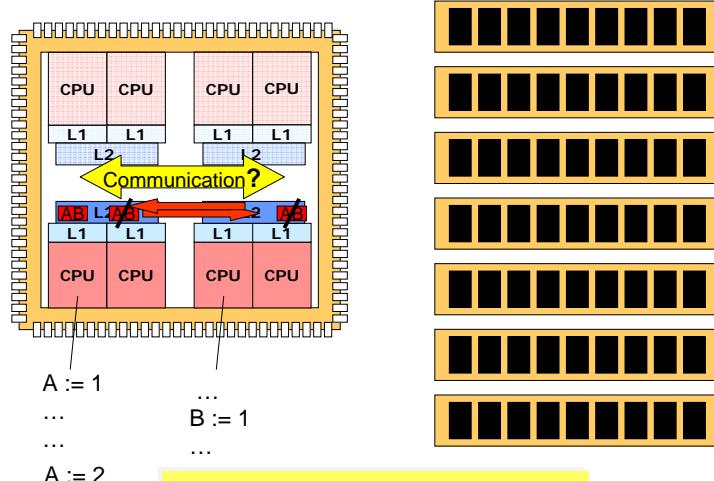
Example: Thread Interaction (True Communication)



© Erik Hagersten | user.it.uu.se/~eh

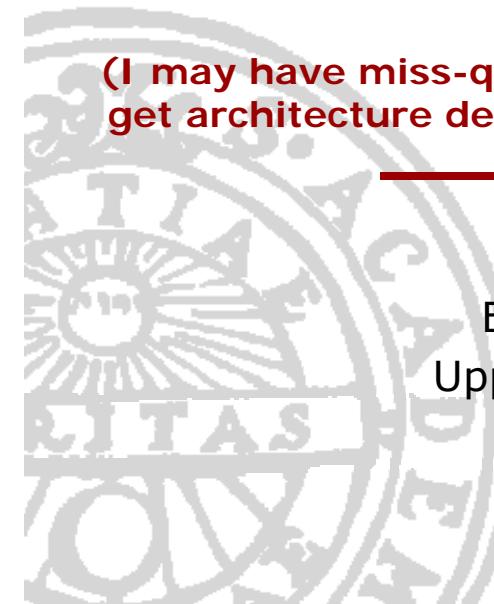
30

Example: Thread Interaction (False sharing)



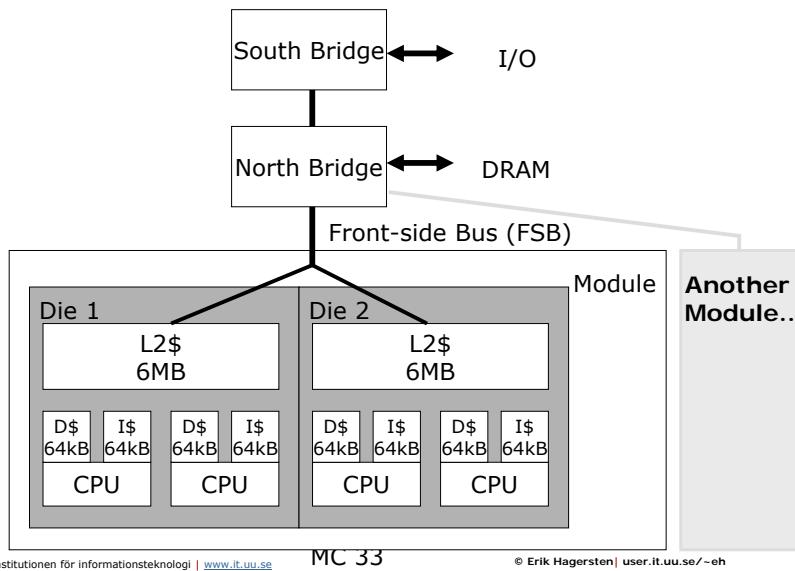
Commercial x86 snapshot

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

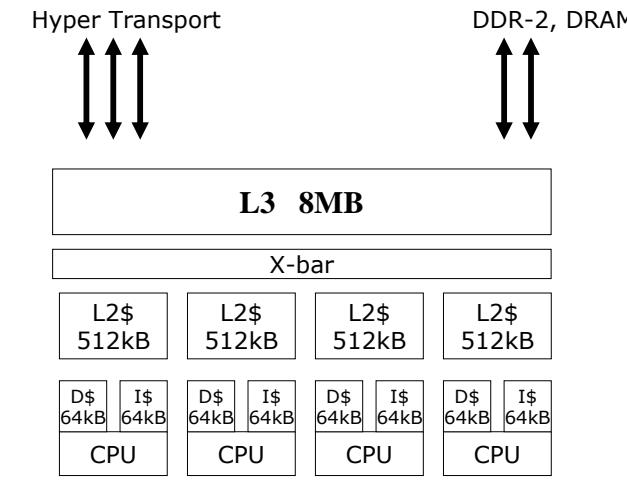


Erik Hagersten
Uppsala Universitet

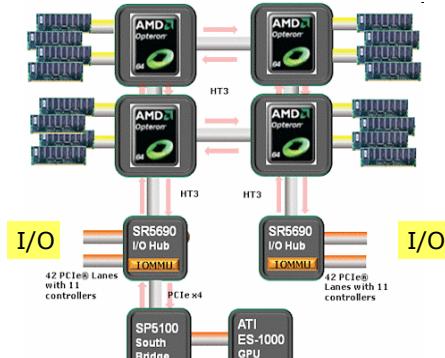
Intel Core2 Quad, 2006


 AVDARK
2010

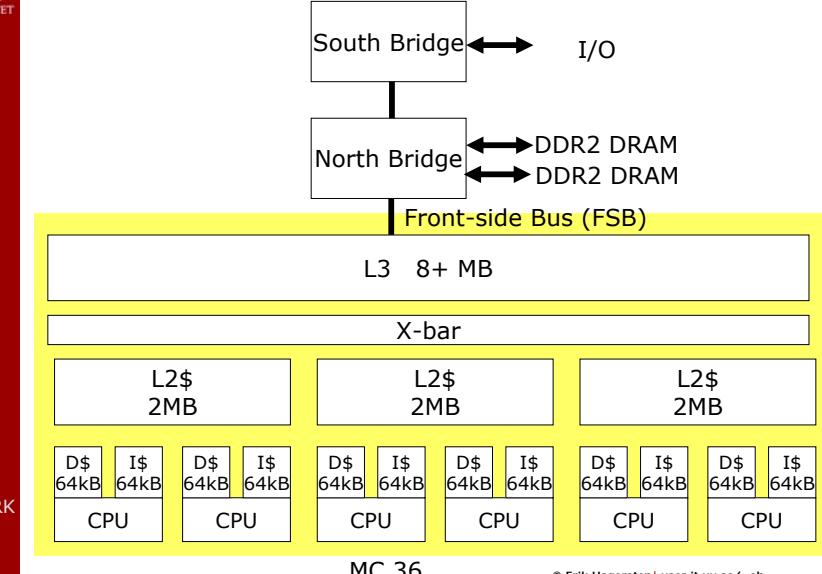
AMD Shanghai, 2007


 AVDARK
2010

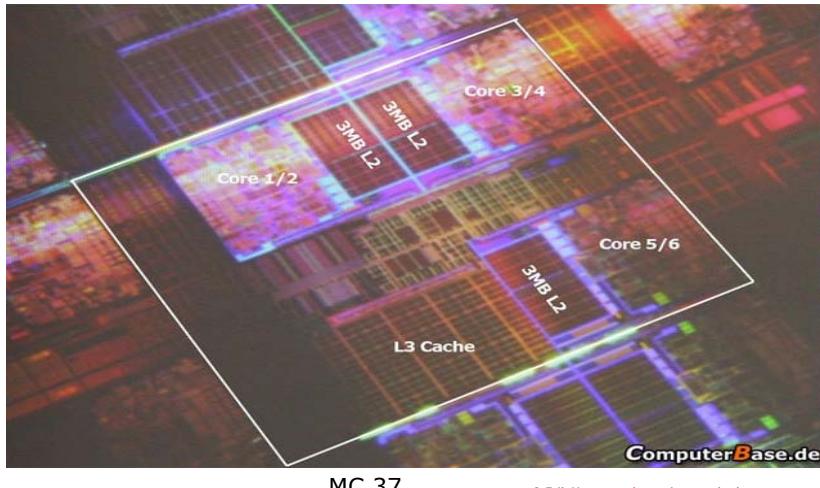
AMD MC System Architecture


 AVDARK
2010

Intel: Dunnington


 AVDARK
2010

Intel Dunnington

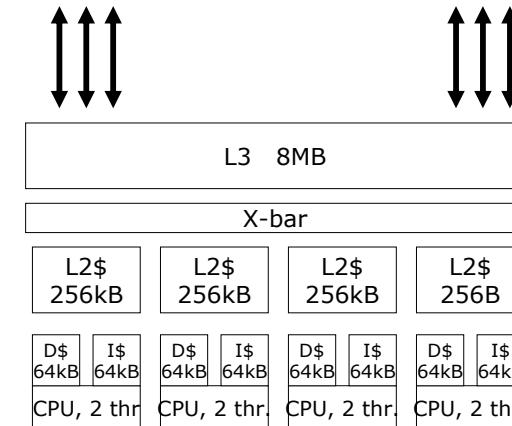
AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

© Erik Hagersten | user.it.uu.se/~eh

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

QuickPath Interconnect

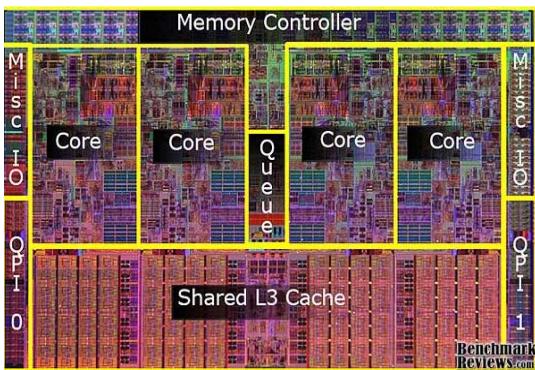
3x DDR-3 DRAM

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

MC 38

© Erik Hagersten | user.it.uu.se/~eh

Nehalem "Core i7"

AVDARK
2010Institutionen för informationsteknologi | www.it.uu.se

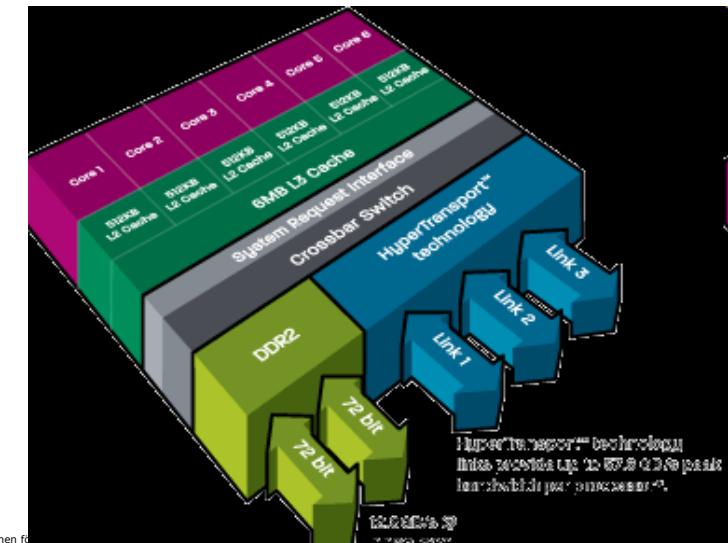
MC 39

© Erik Hagersten | user.it.uu.se/~eh

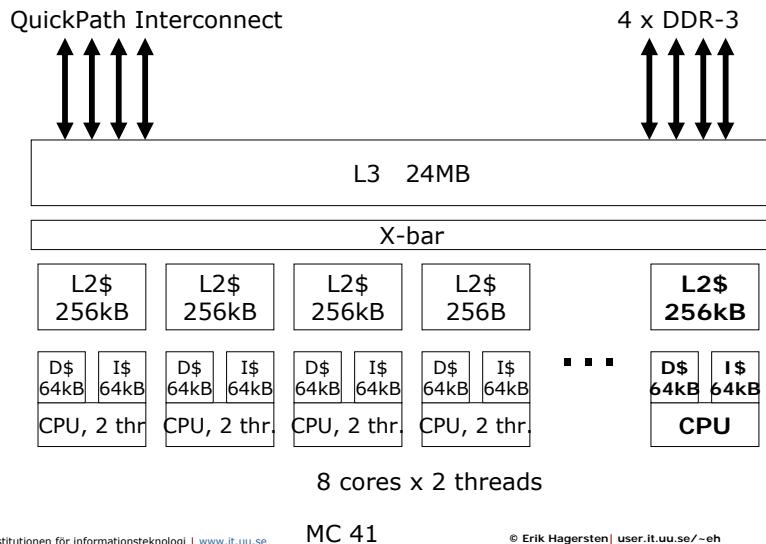
AMD Istanbul, 6 cores

AVDARK
2010

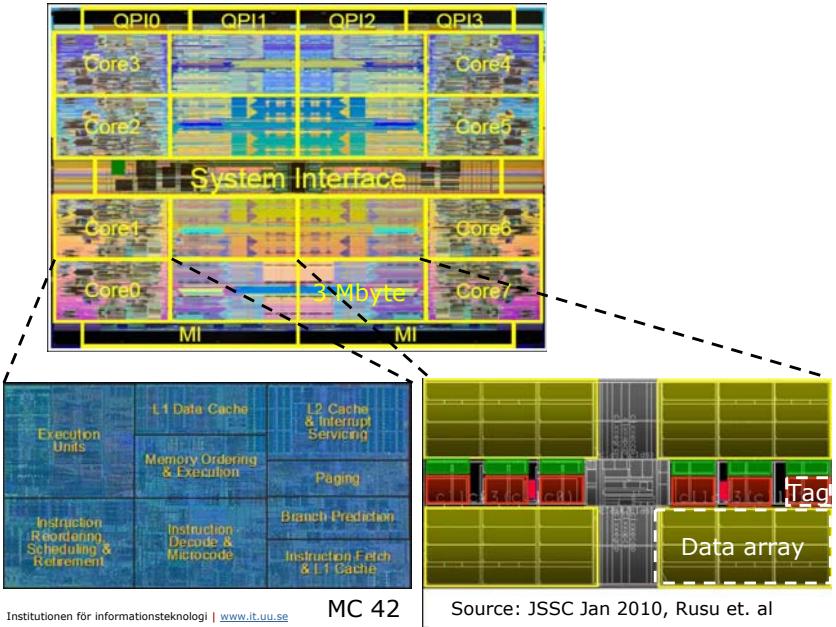
Institutionen f



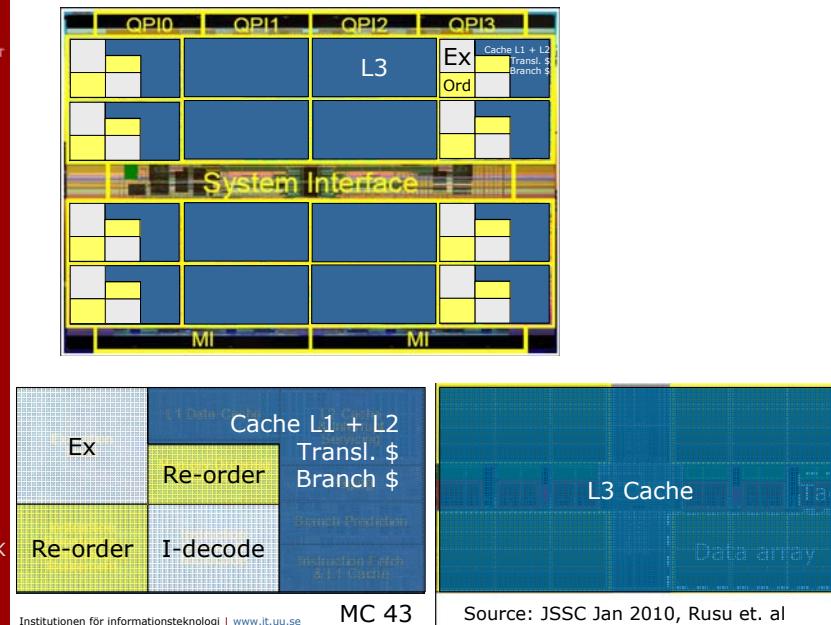
Intel: "Nehalem-Ex" (i7)



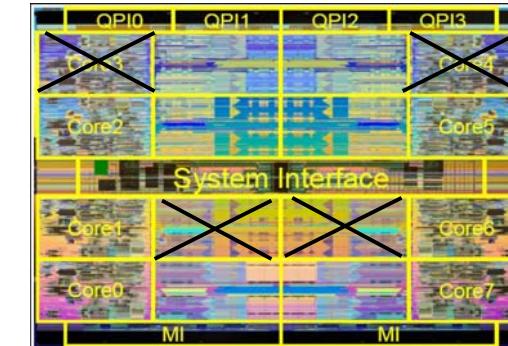
How is the silicon used (i7-Ex)?



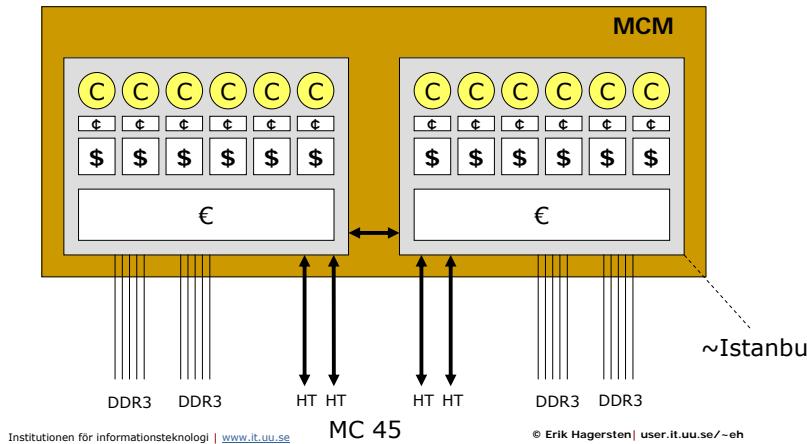
How is the silicon used?



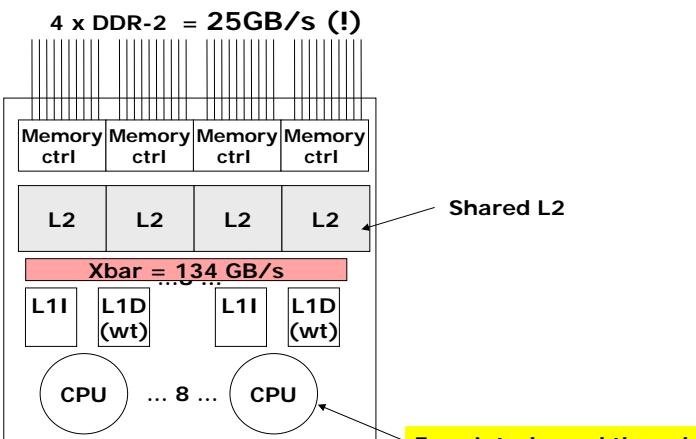
Handling silicon defects?



AMD Magny-Cours

AVDARK
2010

Sun Niagara, 2005!!

AVDARK
2010

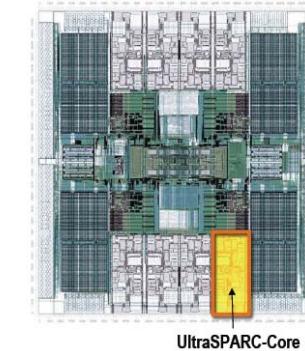
Some other multicores

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

Erik Hagersten
Uppsala Universitet

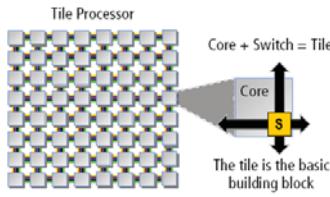
UPPSALA
UNIVERSITET

Niagara Chip

AVDARK
2010

Sun Microsystems

TILER A Architecture



64 cores connected in a mesh
Local L1 + L2 caches
Shared distributed L3 cache
Linux + ANSI C
New Libraries
New IDE
Stream computing
...

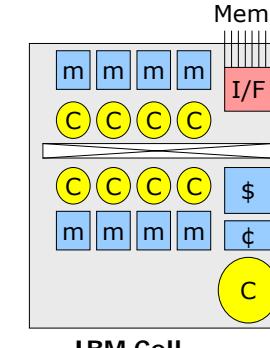
AVDARK
2010

Institutionen för informationsteknologi | www.it.uu.se

MC 49

© Erik Hagersten | user.it.uu.se/~eh

IBM Cell Processor



IBM Cell

AVDARK
2010

Institutionen för informationsteknologi | www.it.uu.se

MC 50

© Erik Hagersten | user.it.uu.se/~eh

So-called accelerators

- Sits on the IO bus (!!)
- GP Graphics processors, aka GPGPU [e.g. NVIDIA, AMD/ATI]
- Specialized accelerators [e.g., FPGA/Mitrionics, ASIC/ClearSpeed]
- Specialized languages for the above [CUDA, Ct, Rapid Mind, Open-CL, ...]

AVDARK
2010

Institutionen för informationsteknologi | www.it.uu.se

MC 51

© Erik Hagersten | user.it.uu.se/~eh

So-called accelerators

- Sits on the IO bus (!!)
- GP Graphics processors, aka GPGPU? [e.g. NVIDIA, AMD/ATI]
- Specialized accelerators? [e.g., FPGA/Mit]
 - My view: Not very general purpose yet!
 - Fits well for a few VERY IMPORTANT app domains!
 - Limited applicability?
 - Programmer productivity?
 - Application life time?
 - New generation devices will be more useful...
- Specialized languages [CUDA, Ct, Rap]

AVDARK
2010

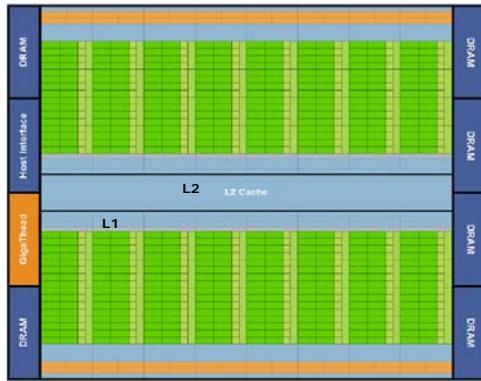
Institutionen för informationsteknologi | www.it.uu.se

MC 52

© Erik Hagersten | user.it.uu.se/~eh

Fermi from nVIDIA

a huge step in the right direction



- 512 "processors" (P)
- 16 P /StreamProcessor (SP)
(i.e., 8k P total)
- SP is SIMD-ish (sort off)
- Full DP-FP IEEE support
- 64kB L1 cache / SP
- 768kB global shared cache
(less than the sum of L1:s)
- Atomic instructions
- ECC correction
- Debugging support
- Giant chip/high power

David Black-Schaffer to give you the full story

AVDARK
2010

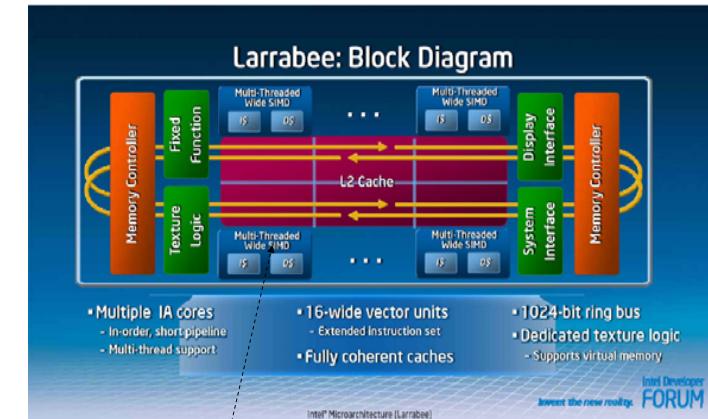
Institutionen för informationsteknologi | www.it.uu.se

MC 53

© Erik Hagersten | user.it.uu.se/~eh

Scaling the x86 Manycore GP computing Larrabee (now called MIC) from Intel 2011??

"more than 50 cores"



SIMD instructions (16-way??)

MC 54

© Erik Hagersten | user.it.uu.se/~eh

Wrapping up about multicores

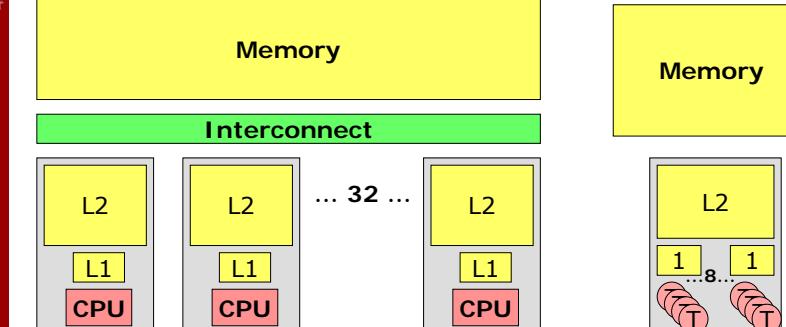
Erik Hagersten
Uppsala Universitet

Looks and Smells Like an SMP (aka UMA)?

SMP system



Multicore system



Well, how about:

- Cost of parallelism?
- Cache capacity per thread?
- Memory bandwidth per thread?
- Cost of thread communication? ...

AVDARK
2010

Institutionen för informationsteknologi | www.it.uu.se

MC 56

© Erik Hagersten | user.it.uu.se/~eh

What matters for multicore performance?

- Are we buying...

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

AVDARK
2010

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 128 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?**

AVDARK
2010