INF5063:
Programming heterogeneous multi-core processors

Introduction

Håkon Kvale Stensland

August 28th, 2012

IAD
INFORMATION ACCESS DISRUPTIONS

UiO : University of Oslo

INF5063


INF5063
ONE CORE TO RULE THEM ALL, ONE CORE TO FIND THEM,
ONE CORE TO BRING THEM ALL AND IN THE DARKNESS ALIGN THEM

IXA: Internet Exchange Architecture

- **IXP2400** basic features:
  - 1 embedded 600 MHz Intel XScale (ARMv5)
  - 8 packet 600 MHz *μengines*
  - onboard memory
  - 3 x 1 Gbps Ethernet ports
  - multiple, independent busses
  - low-speed serial interface
  - interfaces for external memory and I/O busses
  - …

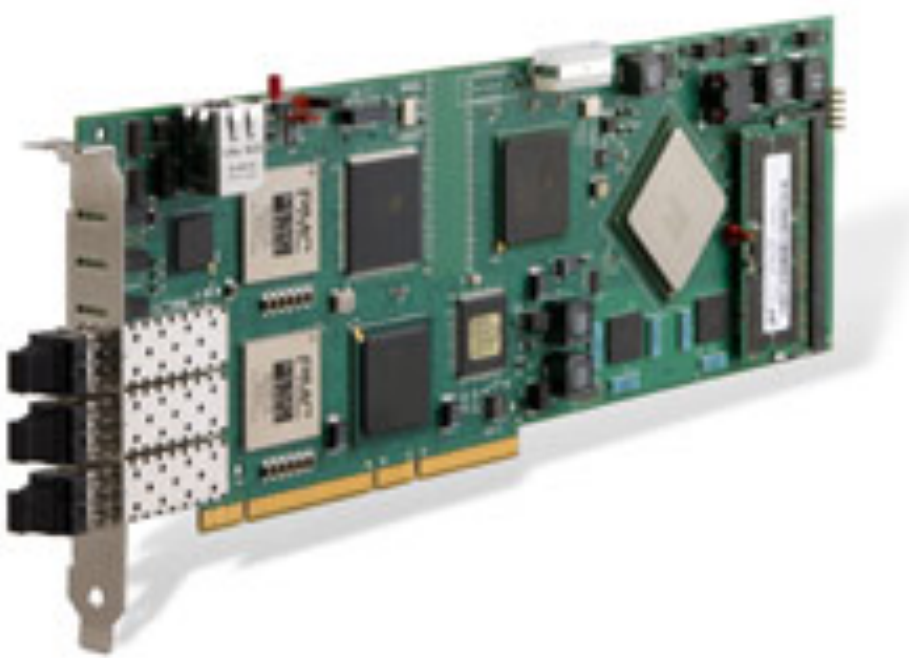

IXP2400 Architecture

IXP2400:

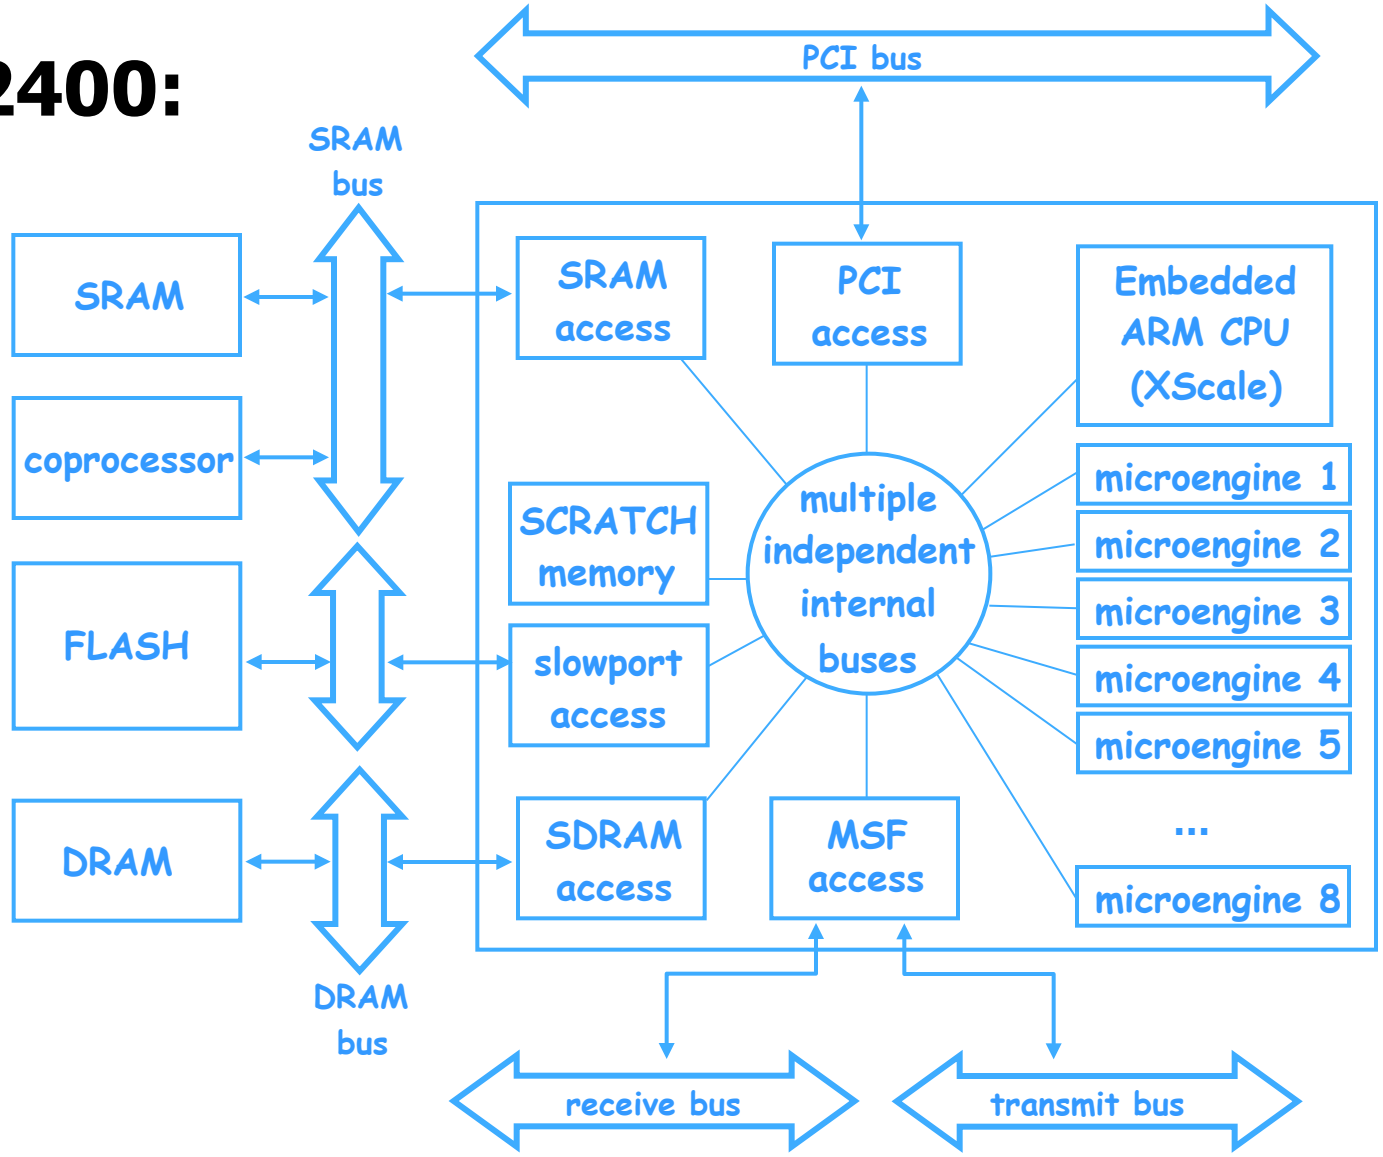

Intel IXP Network Processor

Overview

- Course topic and scope
- Background for the use and parallel processing using heterogeneous multi-core processors
- Examples of heterogeneous architectures

INF5063: The Course

People

- Håvard Espeland
  email: haavares @ ifi
- Håkon Kvale Stensland
  email: haakonks @ ifi
- Preben Nenseth Olsen
  email: prebenno @ ifi
- Carsten Griwodz
  email: griff @ ifi
- Professor Pål Halvorsen
  email: paalh @ ifi

Time and place

- **Lectures**:
  Tuesdays 14:15 - 16:00
  Prolog (Room 2465)

- **Group exercises**:
  Thursdays 09:15 - 12:00
  Pascal (Room 2452)

- There will be some weeks without exercises and lectures, see the detailed plan for each week!

About INF5063: Topic & Scope

- **Content:** The course gives …
  - … an overview of heterogeneous multi-core processors in general and three variants in particular and a modern general-purpose core (architectures and use)
  - … an *introduction* to working with heterogeneous multi-core processors
    - SSE*x* / AVX for x86
    - The Cell Broadband Engine Architecture
    - nVIDIA's family of GPUs and the CUDA/OpenCL programming framework
  - … some ideas of how to use/program heterogeneous multi-core processors

About INF5063: Topic & Scope

- **Tasks:**
  The important part of the course is lab-assignments where you program each of the three examples of heterogeneous multi-core processors

- 3 graded home exams (counting 33% each):
  - Deliver code
  - Make a demonstration and explain your design and code to the class

  1. On the **x86**
     - Video encoding – Improve the performance of *video compression* by using SSE instructions.
  2. On the **Cell processor**
     - Video encoding – Improve the performance of *video compression using the* Cell processor and its SPEs
  3. On the **nVIDIA graphics cards**
     - Video encoding – The same as above, but exploit the parallelism by using the GPU architecture

- Competition at the end of the course! Have the fastest Cell and/or GPU implementation of the code!

Available Resources

- Resources will be placed at
  - http://www.ifi.uio.no/~griff/INF5063
  - Login: *inf5063*
  - Password: *ixp*
  - Manuals, papers, code example, …

Background and Motivation:

Moore's Law

Motivation: Intel View

- >billion transistors integrated

2012:
- 2,6 billion - Intel 10-Core Xeon Westmere-EX
- 7,1 billion - nVIDIA GK110 (Kepler)

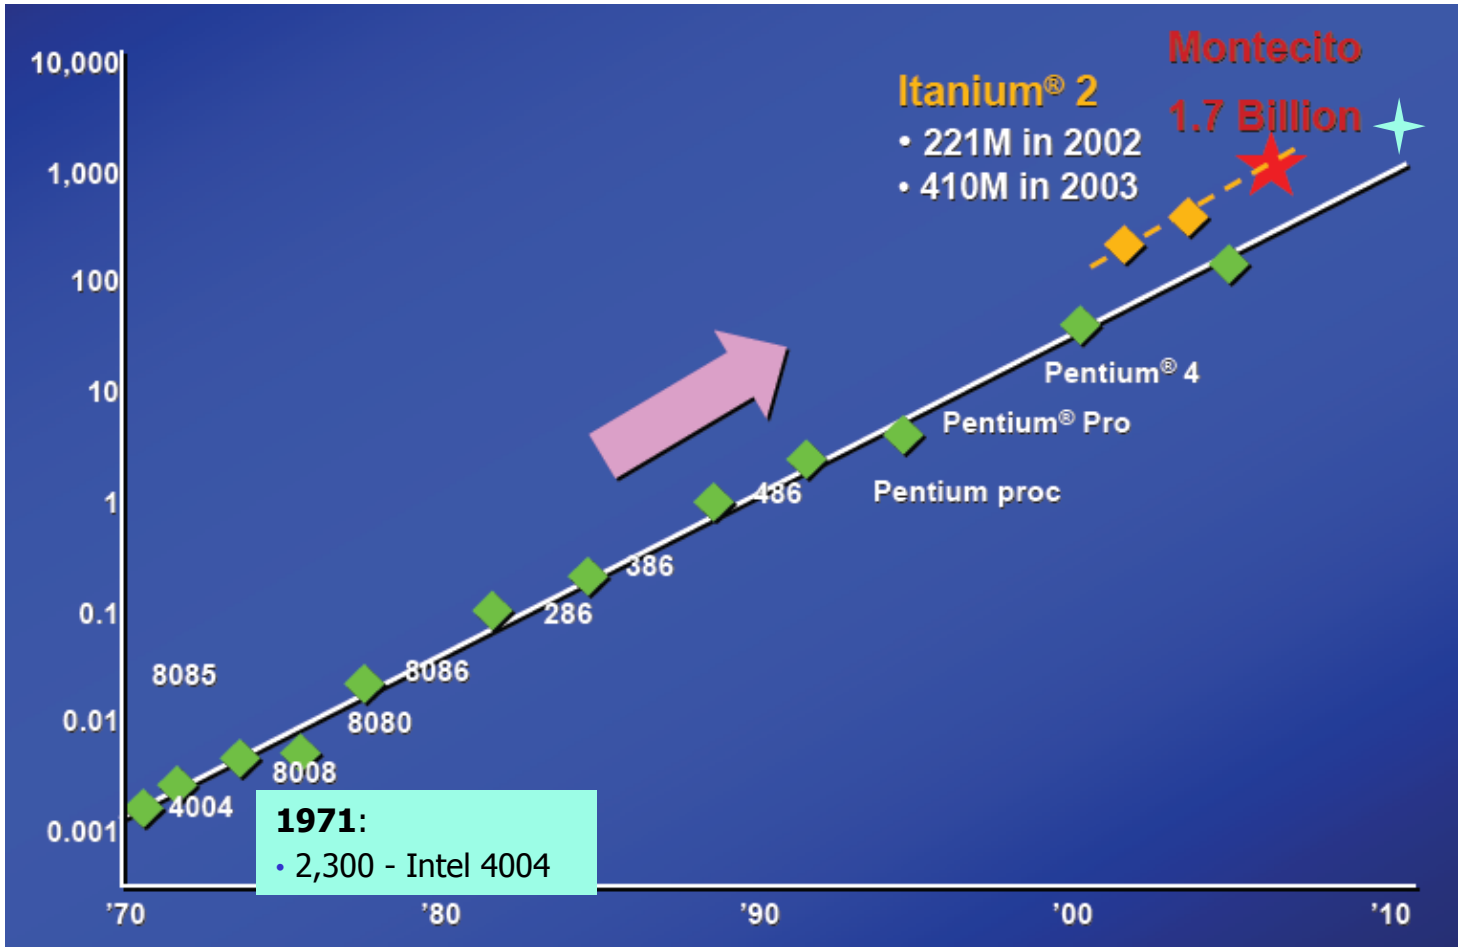

Motivation: Intel View

- >billion transistors integrated
- Clock frequency **can** still increase

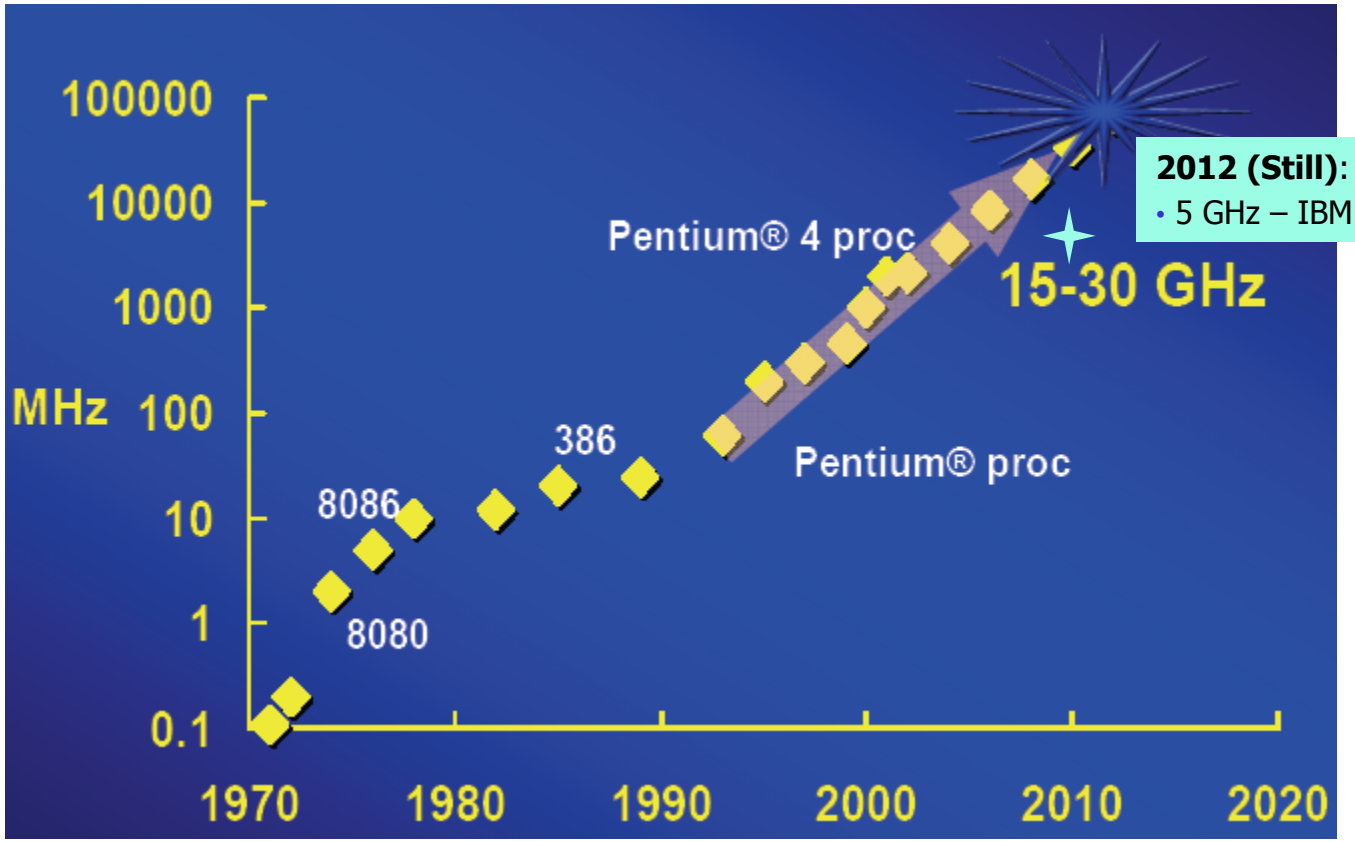

Motivation: Intel View

- >billion transistors integrated
- Clock frequency **can** still increase
- Future applications will demand TIPS

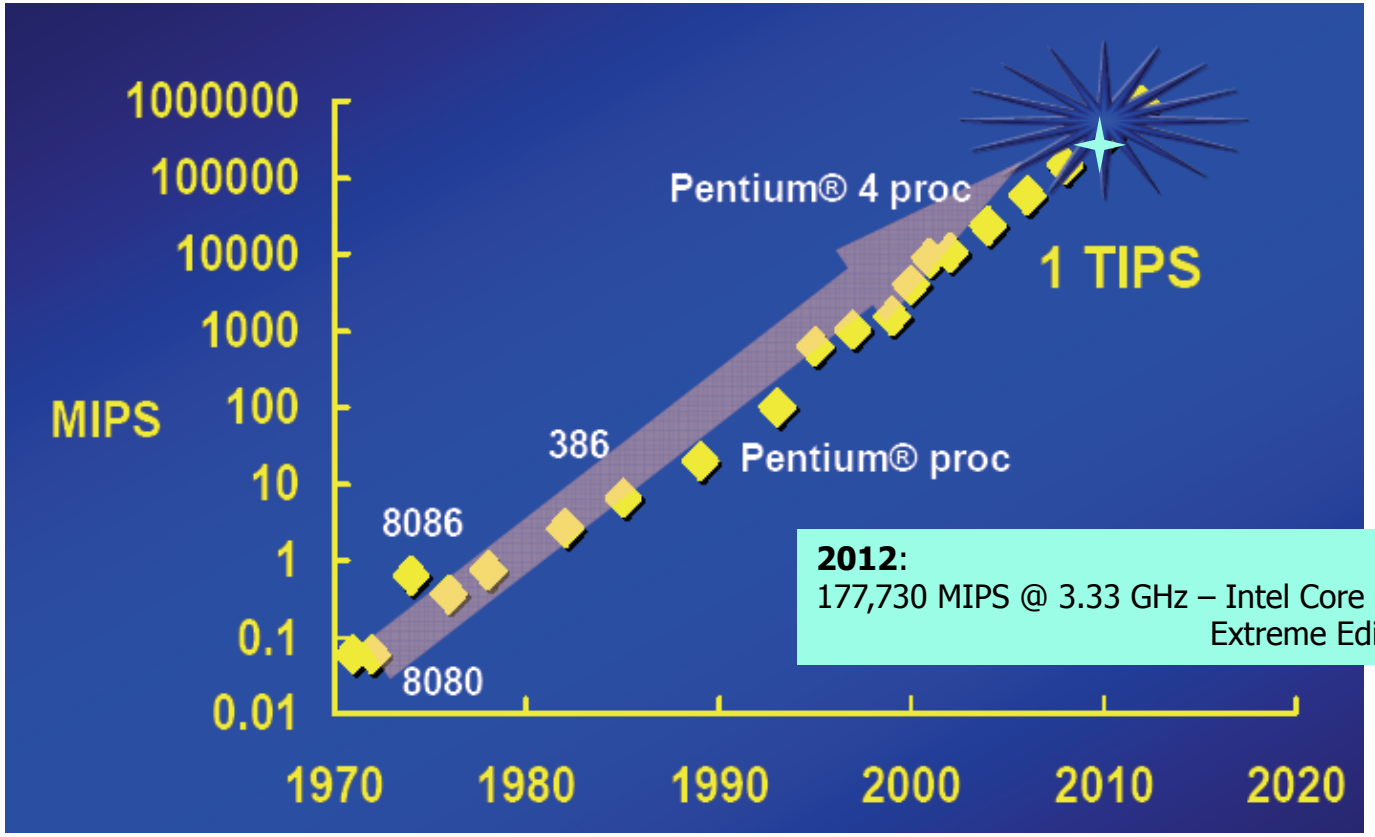

2012:
177,730 MIPS @ 3.33 GHz – Intel Core i7 Extreme Edition 3960X (6 cores)

Motivation: Intel View

- >billion transistors integrated
- Clock frequency **can** still increase
- Future applications will demand TIPS
- Power? Heat?

Motivation: Intel View

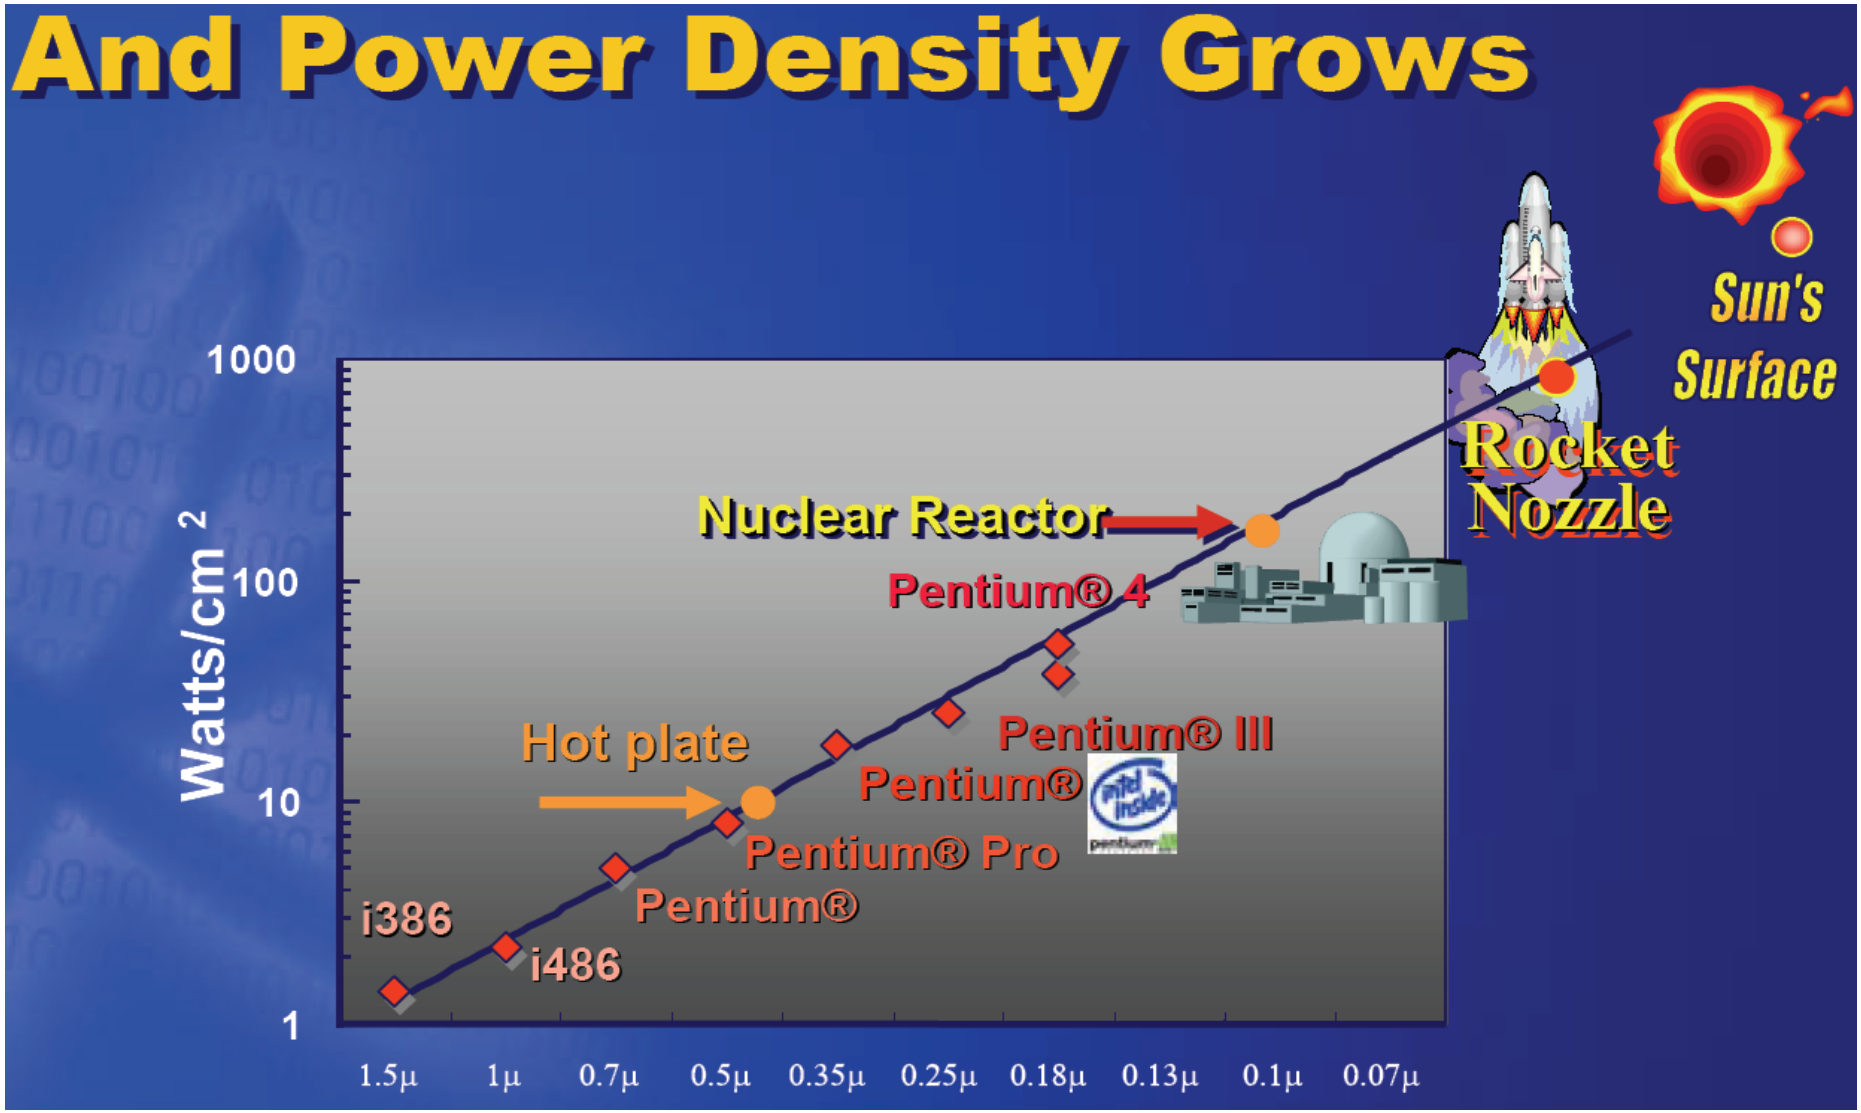

Motivation

"Future applications will demand TIPS"

"Think platform beyond a single processor"

"Exploit concurrency at multiple levels"

"Power will be the limiter due to complexity and leakage"

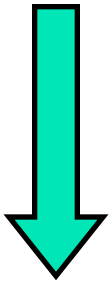

Distribute workload on multiple cores

Symmetric Multi-Core Processors

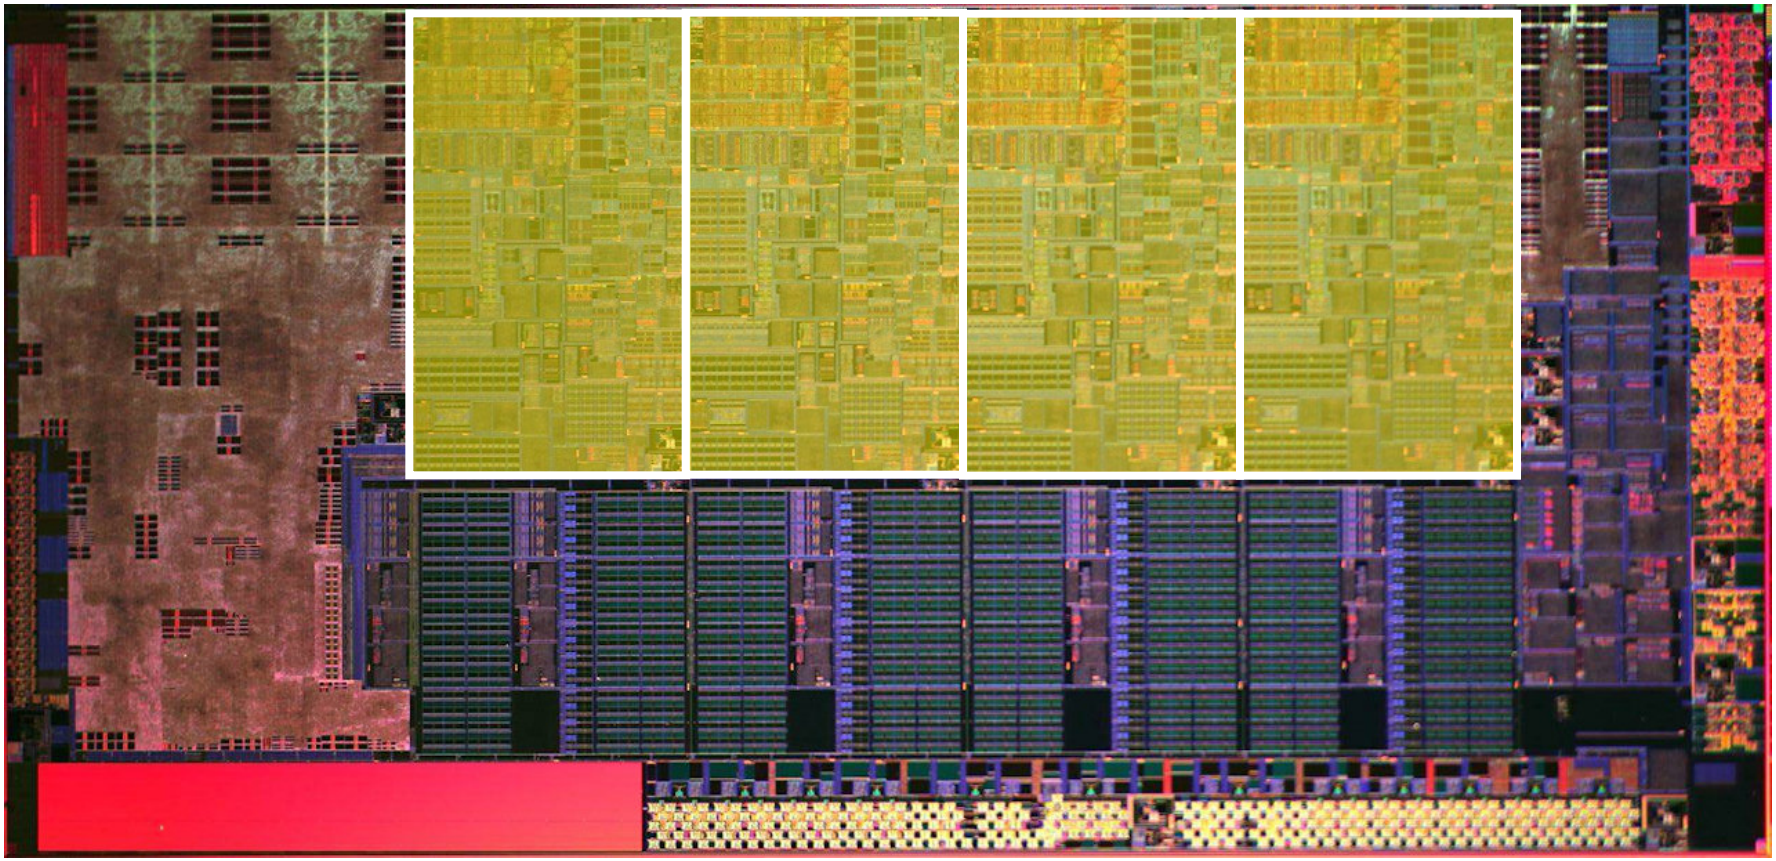

Intel Sandy Bridge

Symmetric Multi-Core Processors

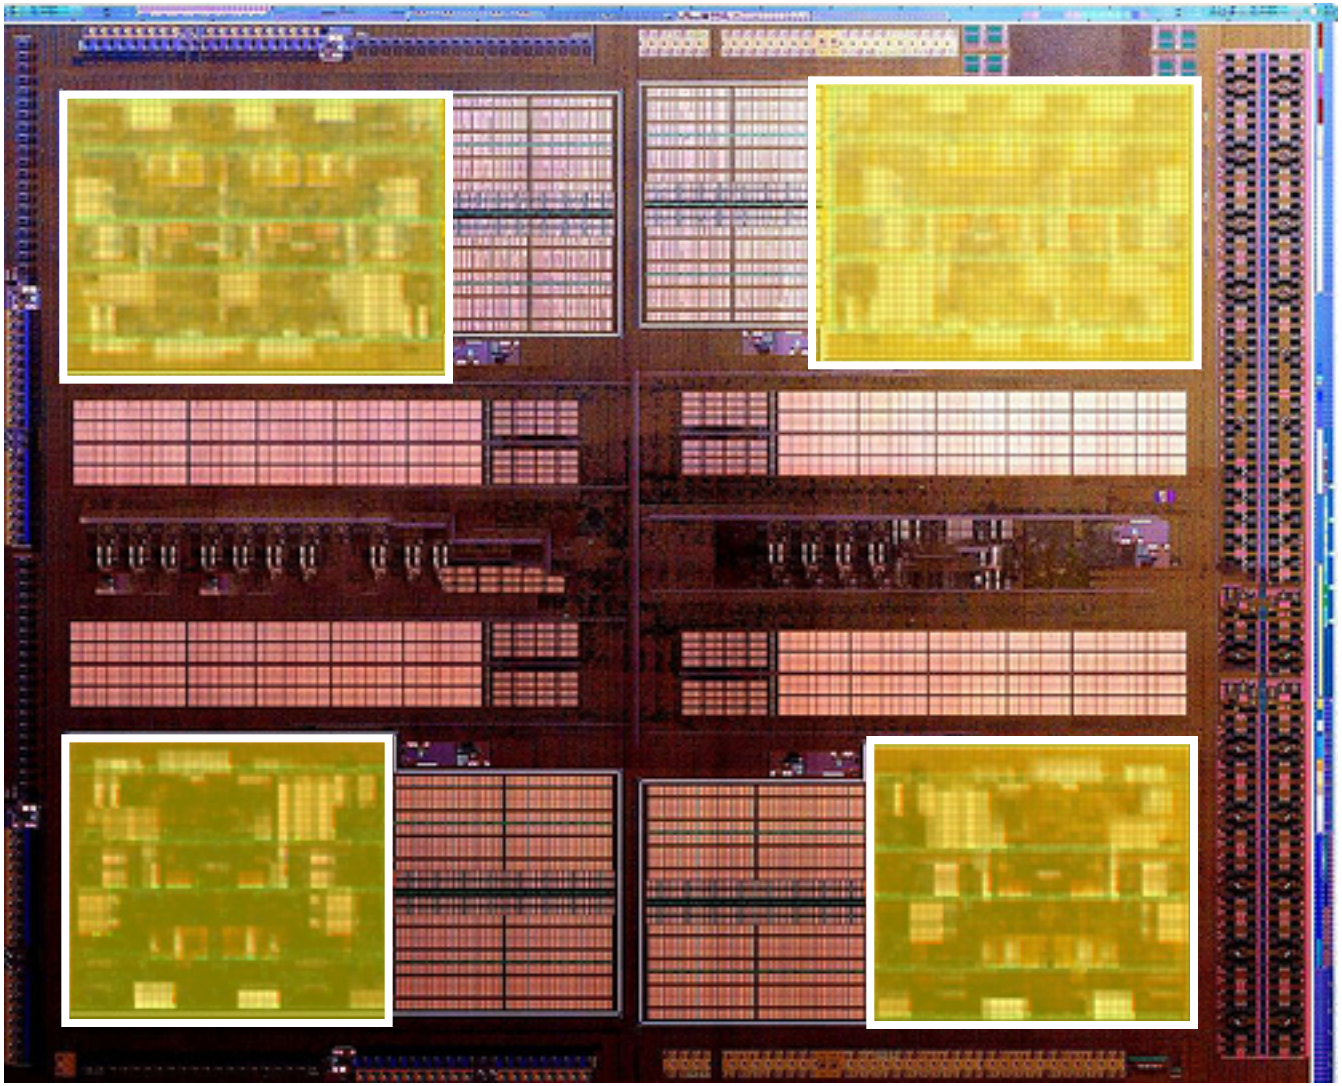

AMD “Bulldozer”

Intel Multi-Core Processors

- Symmetric multi-processors allow multi-threaded applications to achieve higher performance at less die area and power consumption than single-core processors

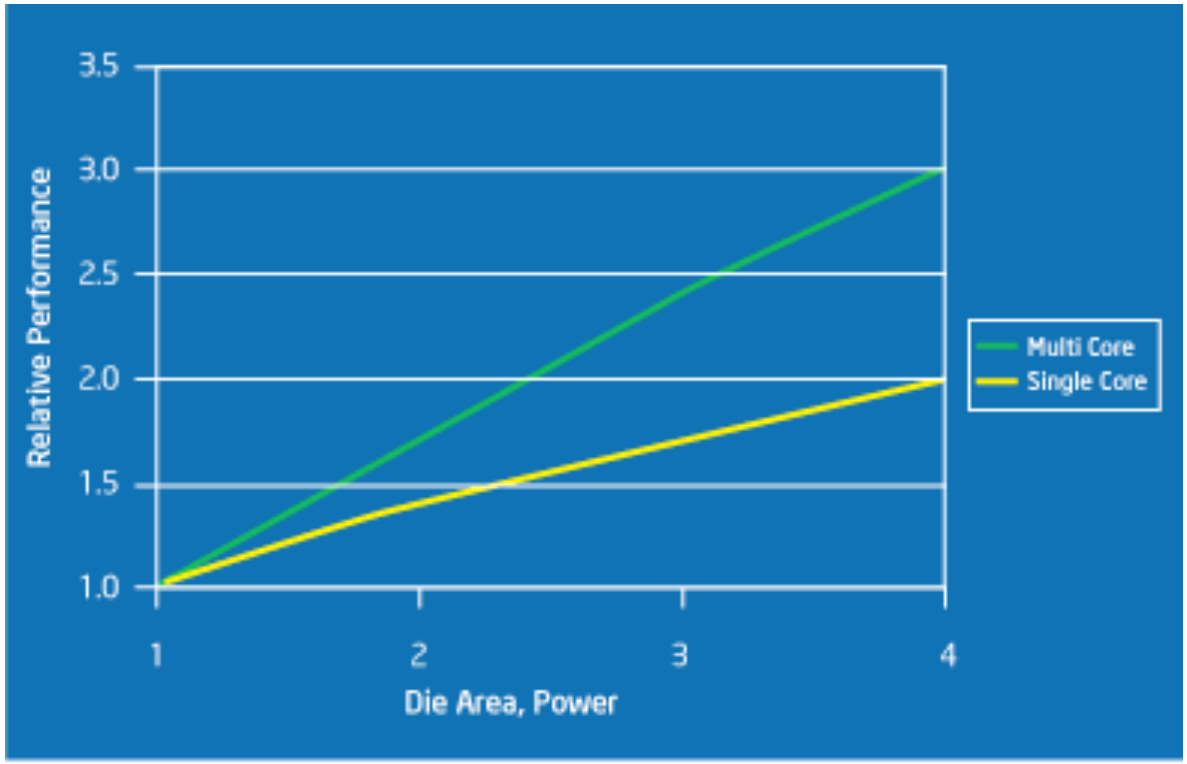

Symmetric Multi-Core Processors

- Good
  - Growing computational power
- Problematic
  - Growing die sizes
  - Unused resources
    - Some cores used much more than others
    - Many core parts frequently unused
- Why not spread the load better?
  - Functions exist only once per core
  - Parallel programming is hard

⇒ Asymmetric multi-core processors

Asymmetric Multi-Core Processors

- Asymmetric multi-processors consume power and provide increased computational power only on demand
- This is now done in all modern CPUs!

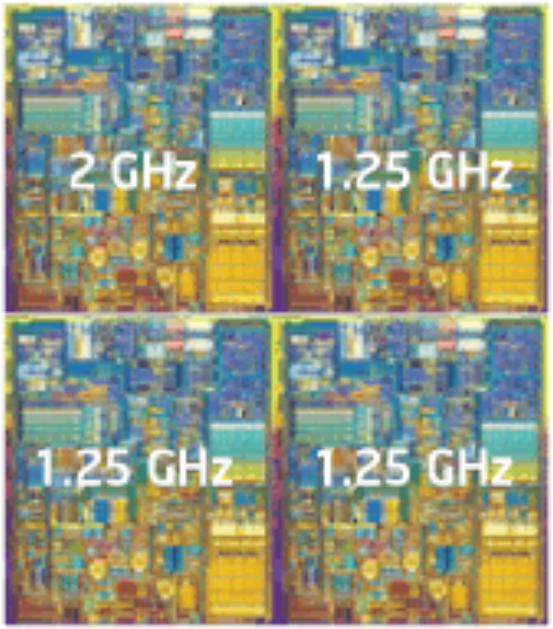

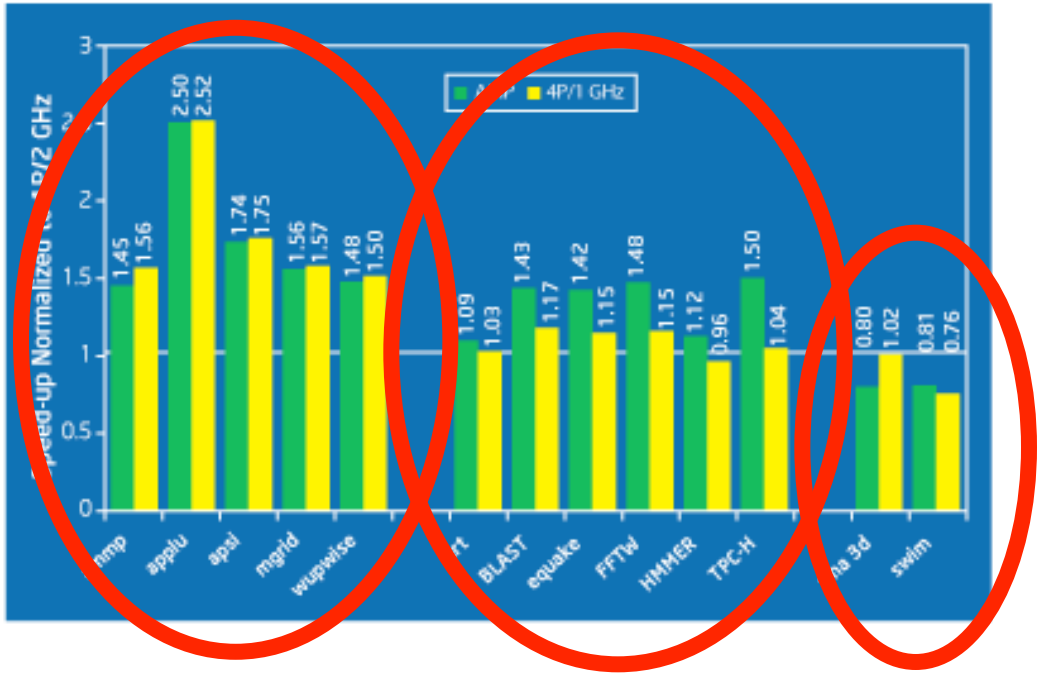

Highly parallel **Moderately parallel** **Sequential**

Motivation

"Future applications will demand TIPS"

"Think platform beyond a single processor"

"Exploit concurrency at multiple levels"

"Power will be the limiter due to complexity and leakage"

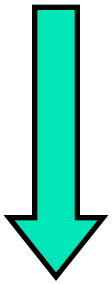

Distributed workload on multiple cores
+ simple processors are "easier" to program
+consume less energy

↳ **heterogeneous multi-core processors**

What now – are today's cores really "Symmetric"?

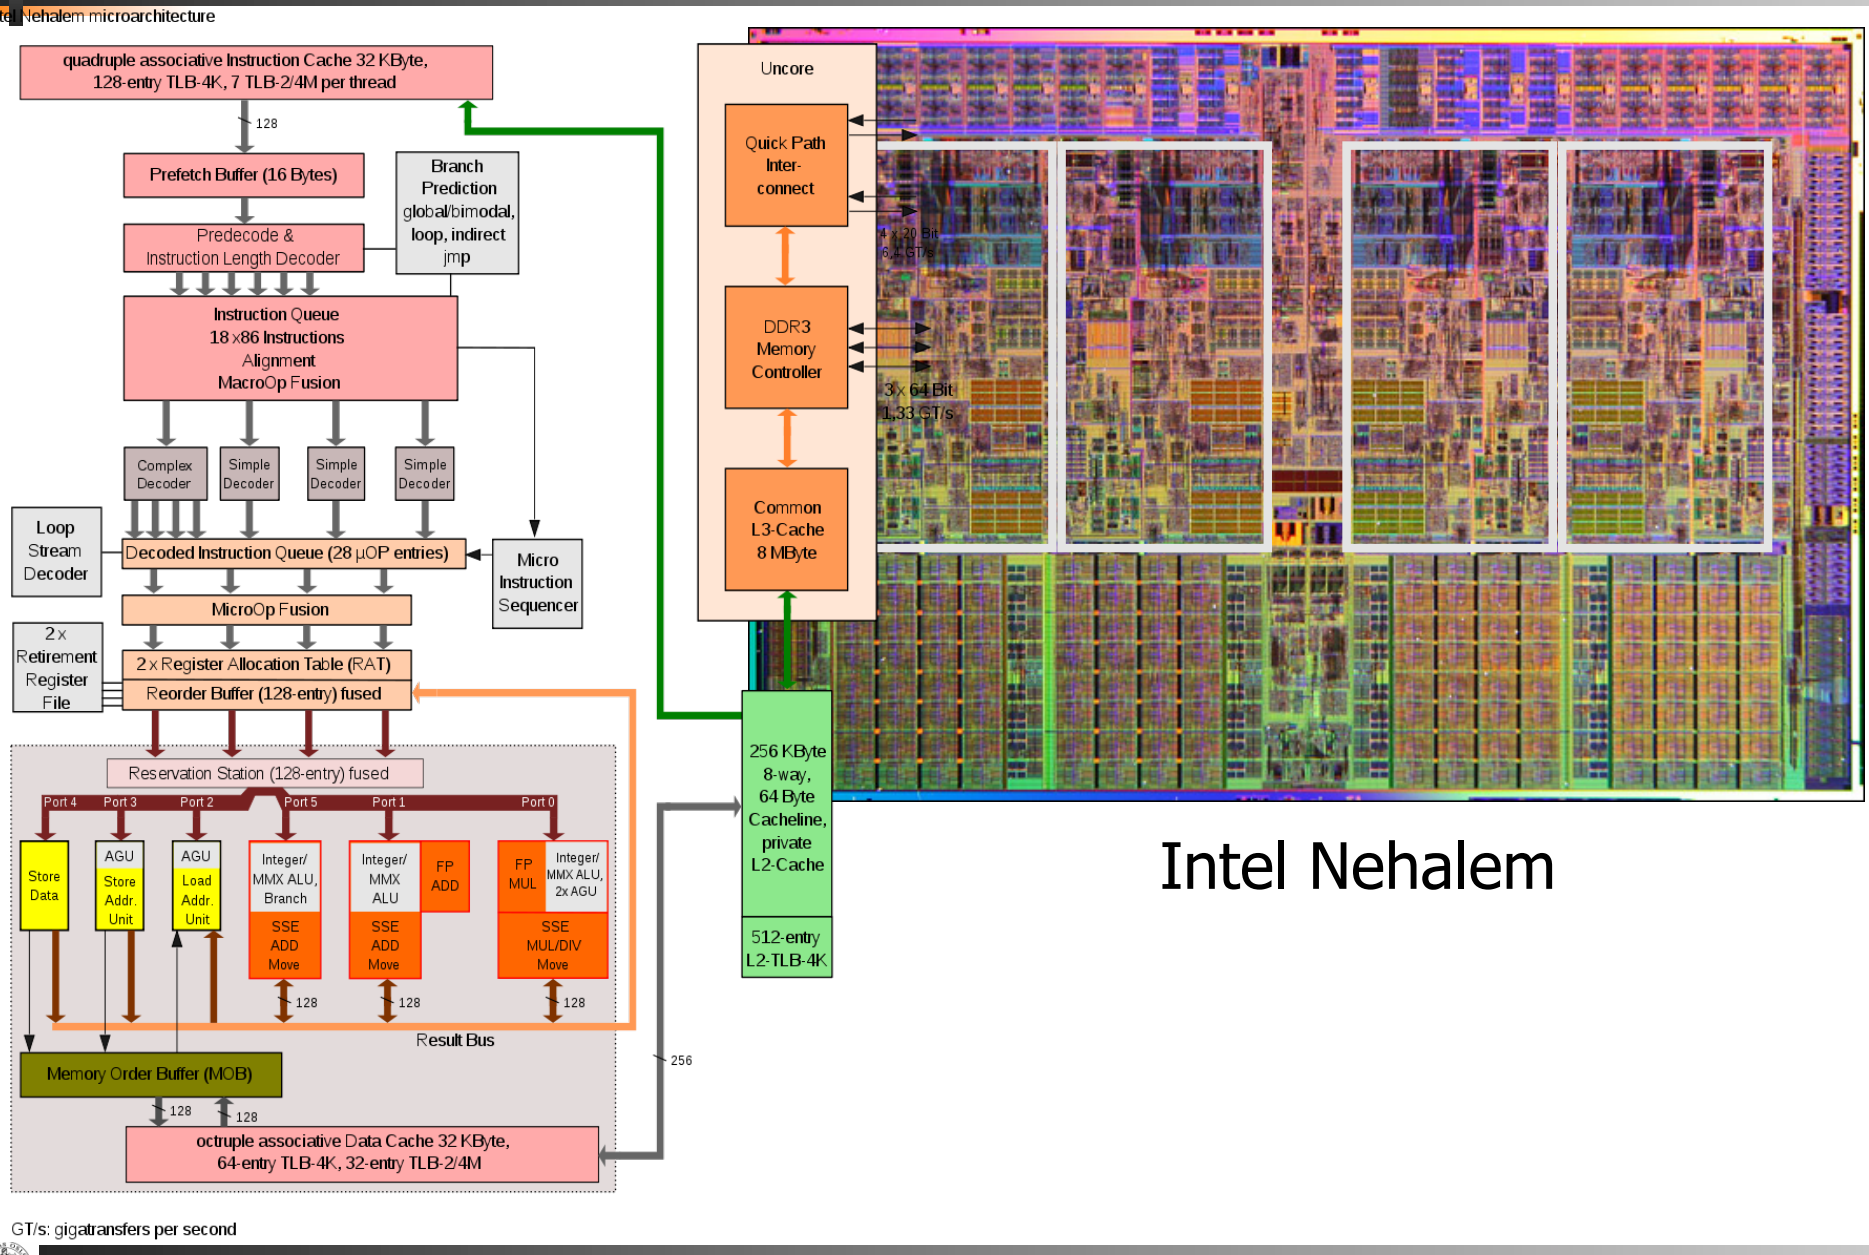

GT/s: gigatransfers per second

Co-Processors

- The original IBM PC included a socket for an Intel 8087 floating point co-processor (FPU)
  - 50-fold speed up of floating point operations
- Intel kept the co-processor up to i486
  - 486DX contained an optimized i487 block
  - Still separate pipeline (pipeline flush when starting and ending use)
  - Communication over an internal bus
- Commodore Amiga was one of the earlier machines that used multiple processors
  - Motorola 680x0 main processor
  - Blitter (block image transferrer - moving data, fill operations, line drawing, performing boolean operations)
  - Copper (Co-Processor - change address for video RAM on the fly)

nVIDIA Tegra 3 ARM SoC

- One of many multi-core processors for handheld devices
- 4 (5?) ARM Cortex-A9 processors
  - Out-of-order design
  - 32-bit ARMv7 instruction set
  - 4 cache-coherent cores
  - 1,6 GHz
- Several “dedicated” processors:
  - HD Video Decoder
  - HD Video Encoder
  - Audio Processor
  - Image Processor
  - Graphics Processor (fixed function)
- Next generation will bring faster cores and a programmable GPUs

Intel SCC (Single Chip Cloud Computer)

- Prototype processor from Intel's TerraScale project
- 24 «tiles» connected in a mesh
- 1.3 billion transistors
- Intel SCC Tile:
  - Two Intel P54C cores Pentium
  - In-order-execution design
  - IA32 architecture
  - NO SIMD units(**!!**)
- No cache coherency
- Only made a limited number of chips for research.

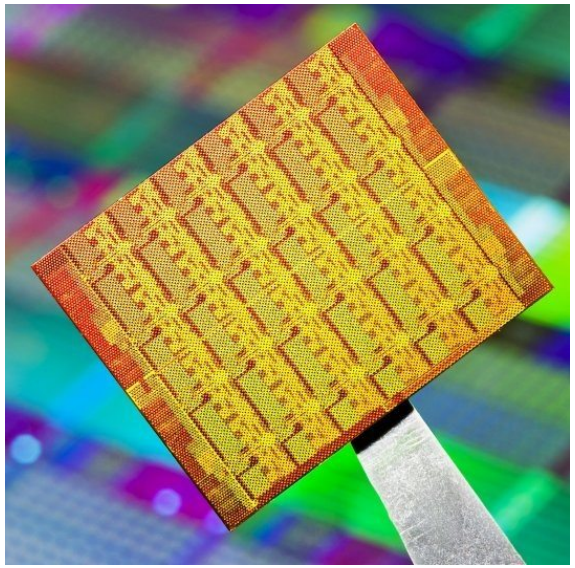

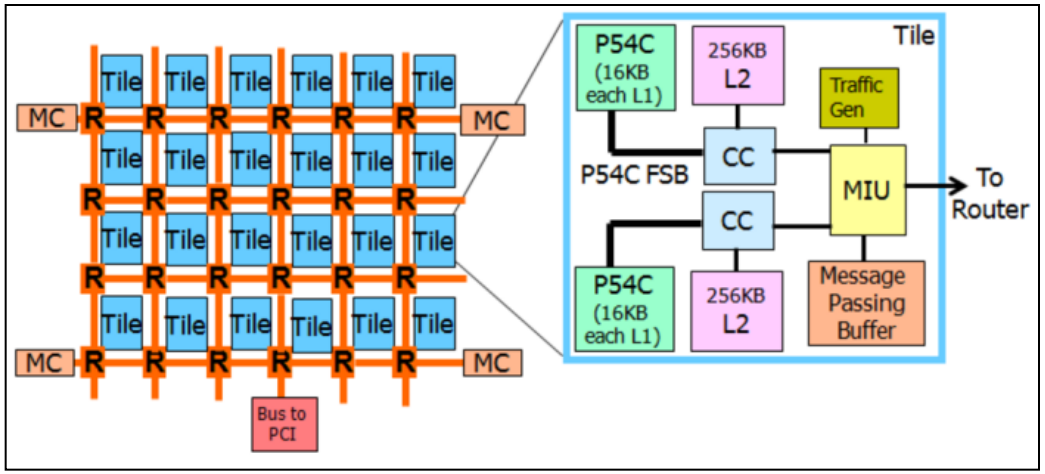

Intel «Larabee» / «Knights Ferry» / «Knights Corner»

- Launched in 2009 as a consumer GPU from Intel
- Canceled in May 2010 due to "disappointing performance"
- Relaunch in 2012 as a dedicated "accelerator" card (Intel Xeon Phi)

Designing GPUs this big is "fucking hard"!
Jonah Alben, NVIDIA's VP of GPU Engineering

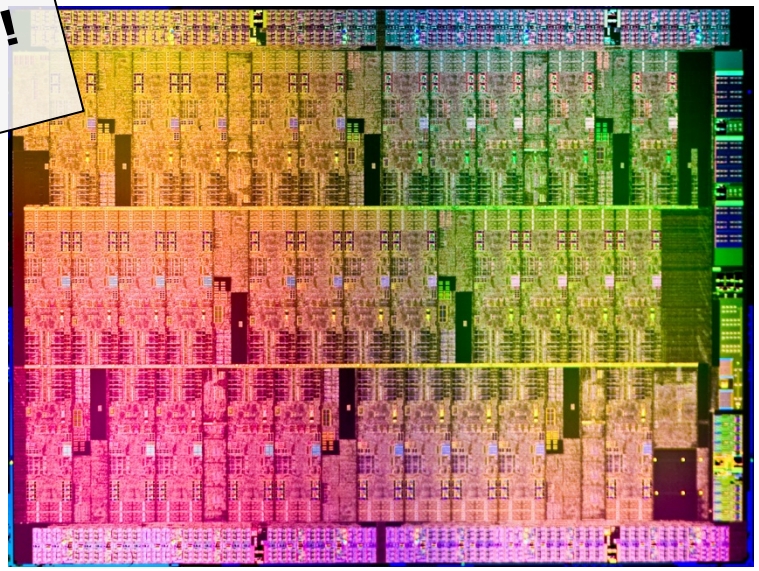

- 50+ cores connected with a ring bus
- Intel "Larabee" core:
  - Based on the Intel P54C Pentium core
  - In-order-execution design
  - Multi-threaded (4-way)
  - 64-bit architecture
  - 512-bit SIMD unit
  - 256 kB L2 Cache
- Cache coherency
- Software rendering

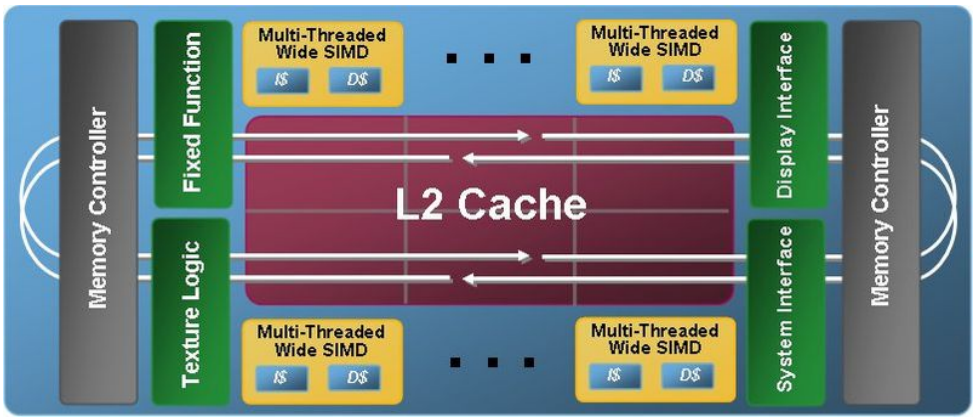

Graphics Processing Units (GPUs)

GPU:
a dedicated graphics rendering device

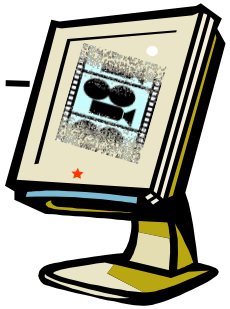

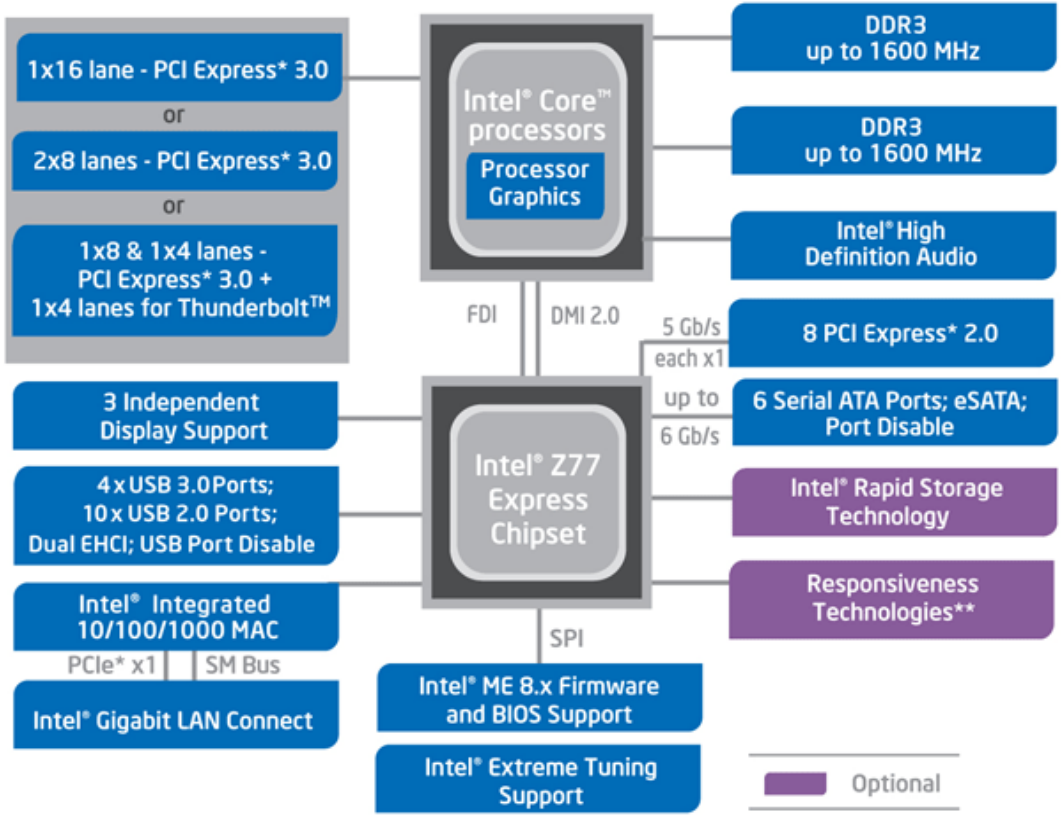

New powerful GPUs, e.g.,:

- ✓ *Nvidia GeForce GTX 680*
  - ✓ 1536 – 1058 MHz core
  - ✓ 2 GB memory
  - ✓ memory BW: 192,26 GB/sec
  - ✓ PCI Express 3.0
- ✓ similar to other manufacturers …

General Purpose Computing on GPU

- The
  - high arithmetic precision
  - extreme parallel nature
  - optimized, special-purpose instructions
  - available resources
  - …

  … of the GPU allows for general, non-graphics related operations to be performed on the GPU

- Generic computing workload is off-loaded from CPU and to GPU

⇒ **More generically:**
Heterogeneous multi-core processing

nVIDIA Kepler Architecture – GK104

- 3,54 billion transistors
- 1536 cores
- 256 bit memory bus (GDDR5)
- 192,26 GB/sec memory bandwidth
- 3090 Gflops single precision
- PCI Express 3.0

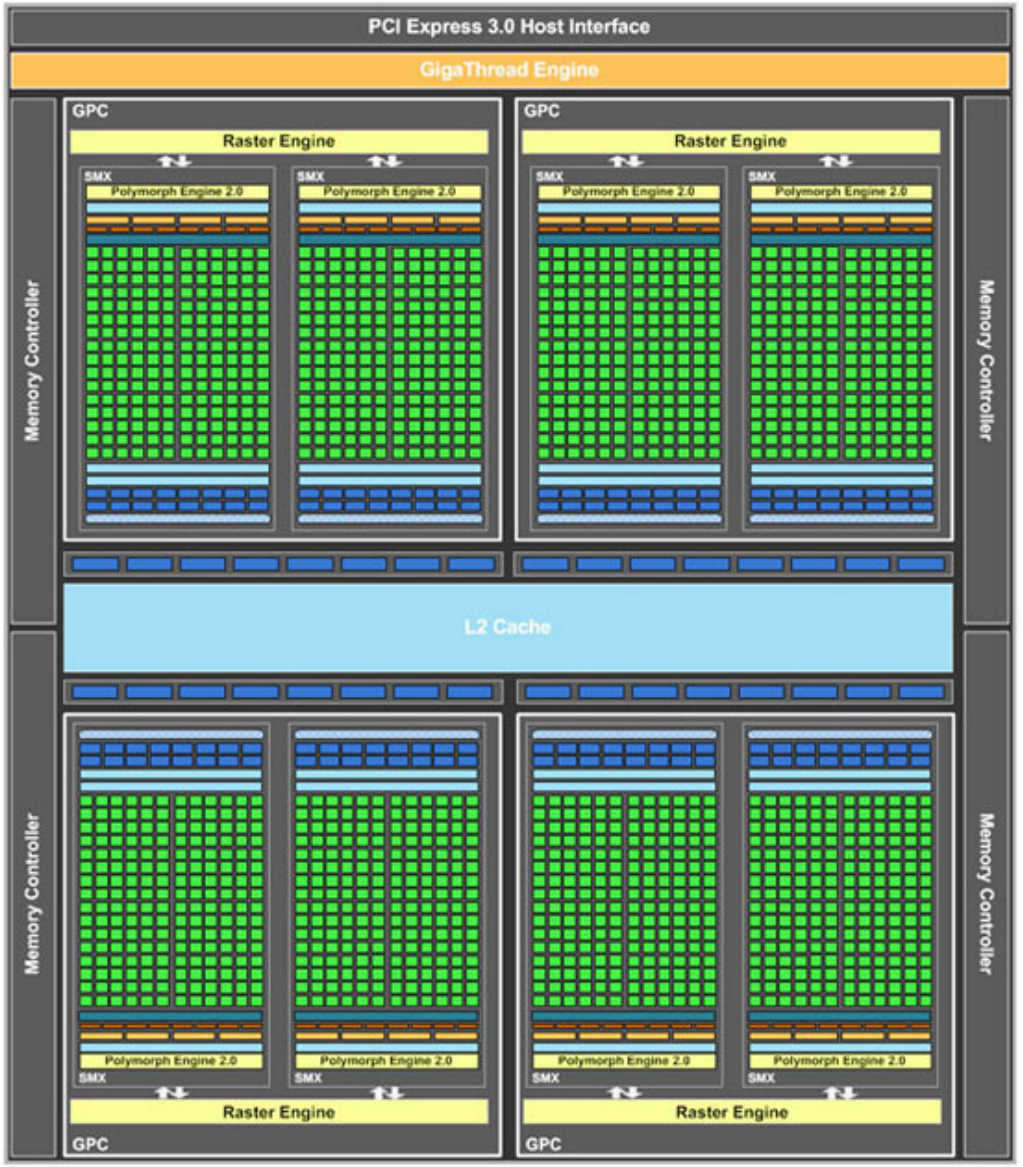

nVIDIA GK104 / GK110

- GK10x Streaming Multiprocessor (SMX)
  - Fundamental thread block unit
  - 192 stream processors (SPs) (scalar ALU for threads)
  - 32 super function units (SFUs) (cos, sin, log, ...)
  - 256 kB local register files (RFs)
  - 16 / 48 kB level 1 cache
  - 16 / 48 kB shared memory
  - 768 kB global level 2 cache

- GK110 Streaming Multiprocessor (SMX)
  - *In addition:*
  - *64 double-precision ALUs*
  - *48 kB Read-Only Data Cache*
  - 1536 kB global level 2 cache

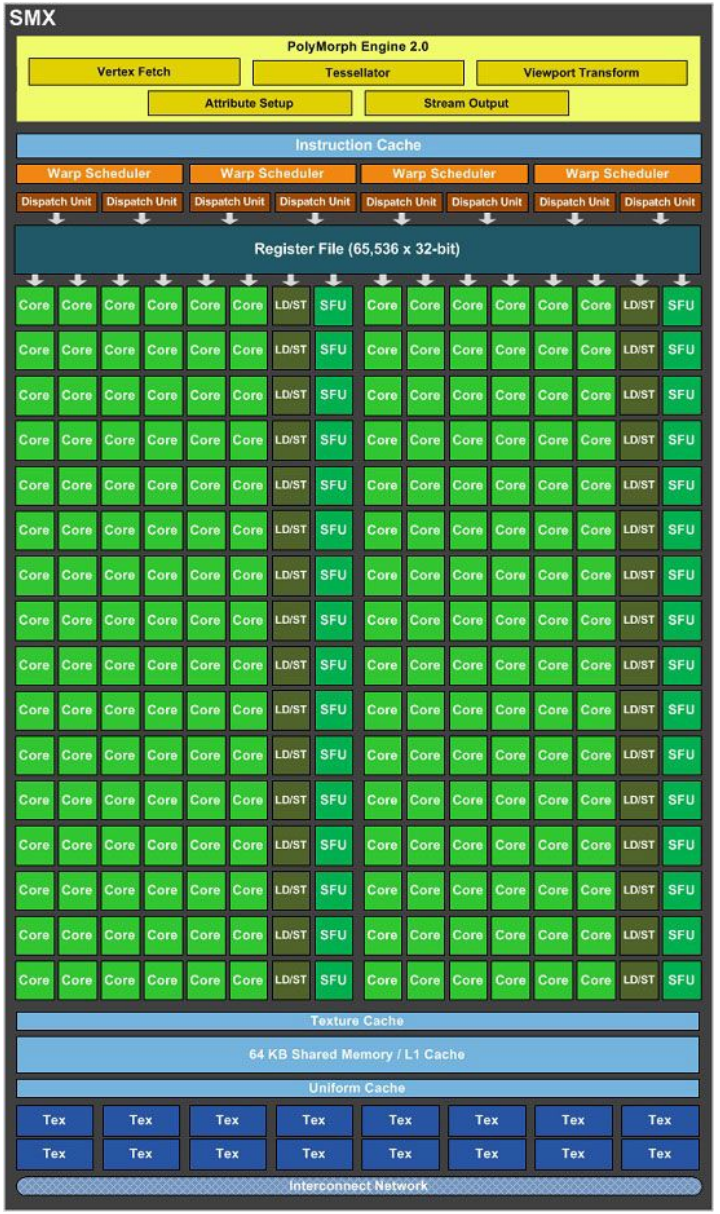

STI (Sony, Toshiba, IBM) Cell

- Motivation for the Cell
  - Cheap processor
  - Energy efficient
  - For games and media processing
  - Short time-to-market
- Conclusion
  - Use a multi-core chip
  - Design around an existing, power-efficient design
  - Add simple cores specific for game and media processing requirements

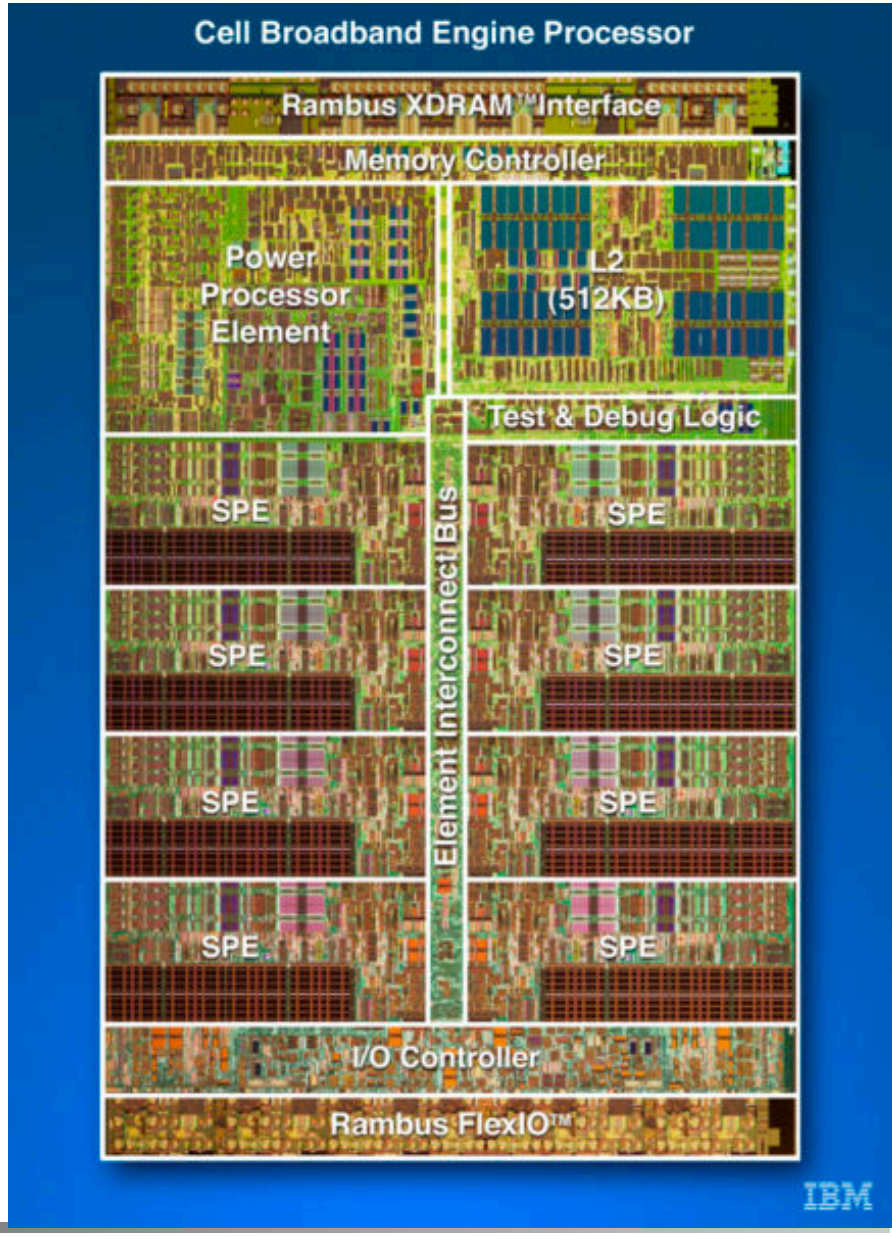

STI (Sony, Toshiba, IBM) Cell

- Cell is a 9-core processor
  - combining a light-weight general-purpose processor with multiple co-processors into a coordinated whole
  - *Power Processing Element* (PPE)
    - conventional Power processor
    - not supposed to perform all operations itself, acting like a controller
    - running conventional OSes
    - 16 KB instruction/data level 1 cache
    - 512 KB level 2 cache

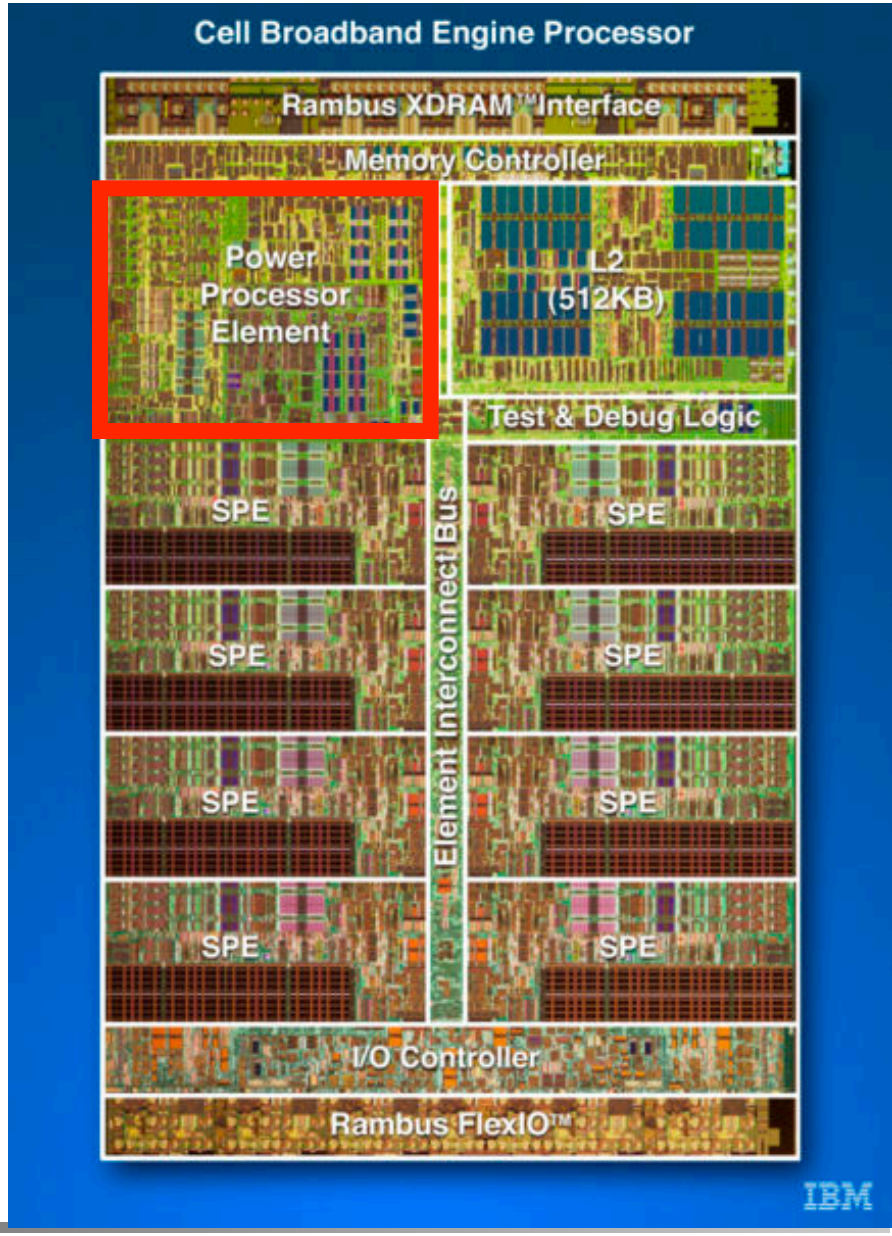

STI (Sony, Toshiba, IBM) Cell

- *Synergistic Processing Elements* (SPE)
  - specialized co-processors for specific types of code, i.e., very high performance vector processors
  - local stores
  - can do general purpose operations
  - the PPE can start, stop, interrupt and schedule processes running on an SPE
- Element Interconnect Bus (EIB)
  - internal communication bus
  - connects on-chip system elements:
    - PPE & SPEs
    - the memory controller (MIC)
    - two off-chip I/O interfaces
  - 25.6 GBps each way

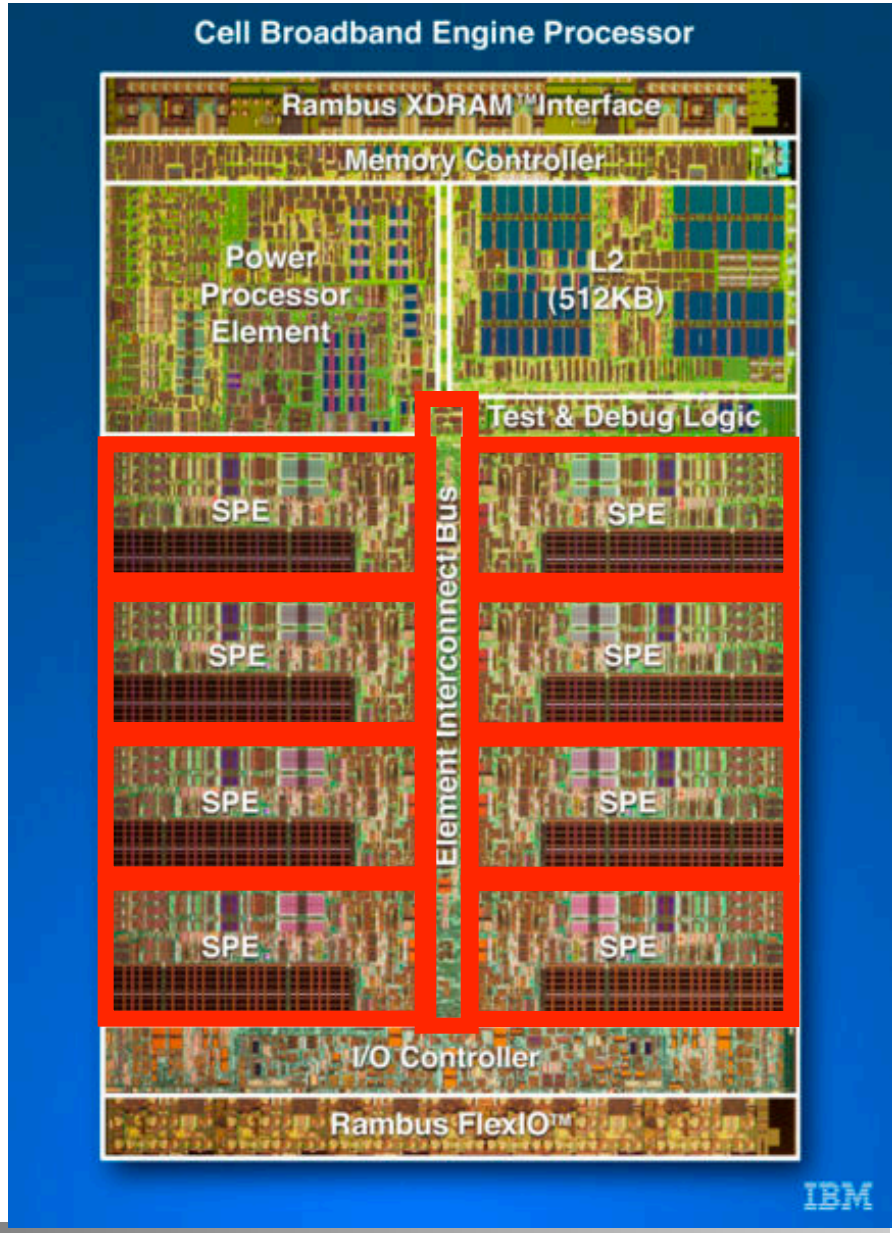

STI (Sony, Toshiba, IBM) Cell

- memory controller
  - Rambus XDRAM interface to Rambus XDR memory
  - dual channels at 12.8 GBps → 25.6 GBps
- I/O controller
  - Rambus FlexIO interface which can be clocked independently
  - dual configurable channels
  - maximum ~ 76.8 GBps

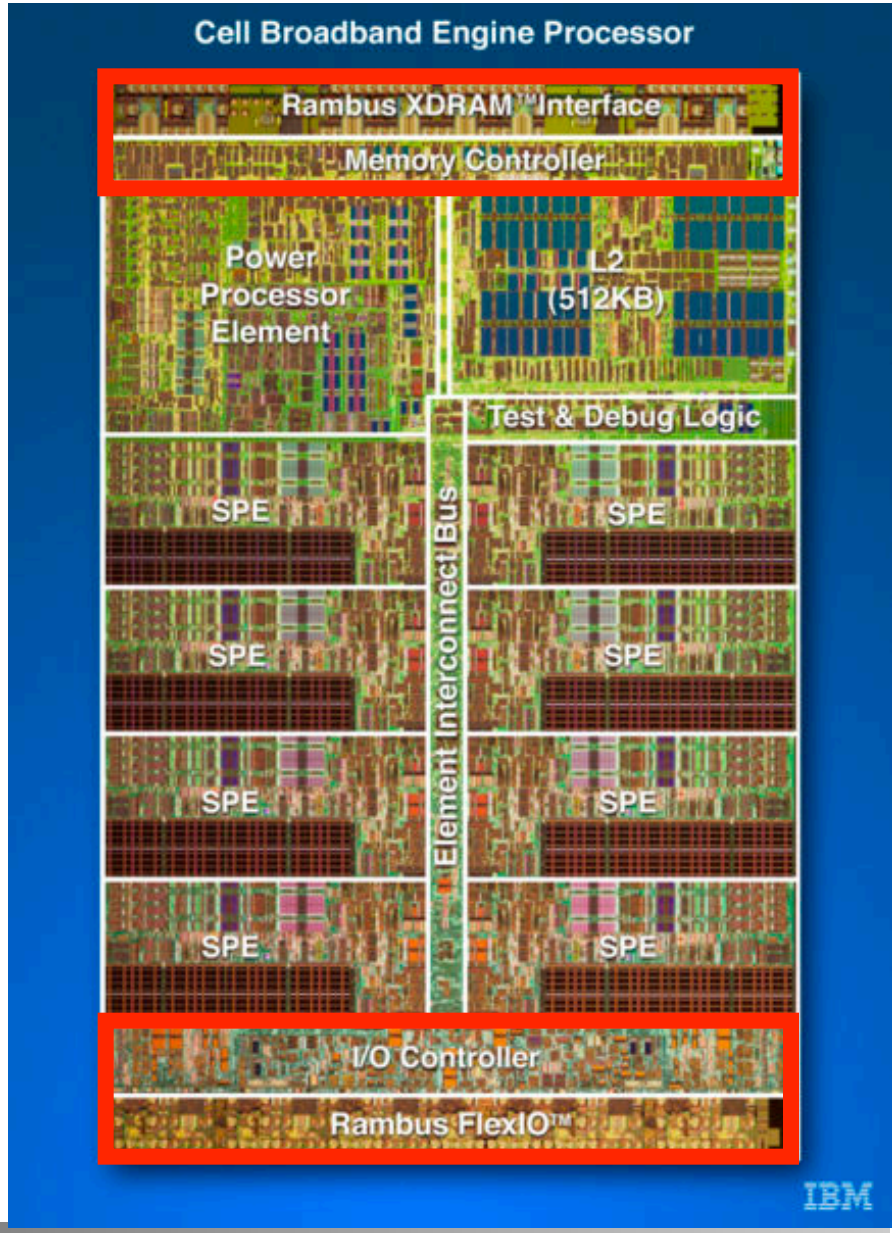

STI (Sony, Toshiba, IBM) Cell

- Cell has in essence traded running everything at moderate speed for the ability to run certain types of code at high speed

- used for example in
  - Sony PlayStation 3:
    - 3.2 GHz clock
    - 6 SPEs for general operations
    - 1 SPE for security for the OS
  - Toshiba home cinema:
    - decoding of 48 HDTV MPEG streams
      → dozens of thumbnail videos simultaneously on screen
  - IBM blade centers:
    - 3.2 GHz clock
    - Linux ≥ 2.6.11

The End: Summary

- Heterogeneous multi-core processors are already everywhere

⇨**Challenge: programming**

- **Need to know the capabilities of the system**
- **Different abilities in different cores**
- **Memory bandwidth**
- **Memory sharing efficiency**
- **Need new methods to program the different components**