

Introduction

1/9 - 2006

Overview

- ✓ Course topic and scope
- ✓ Examples of asymmetric processors
- ✓ (Very) short intro of the Intel IXP 2400 network processors

INF5062: The Course

People

- ✓ Morten Pedersen (TA)
email: mortp @ ifi
- ✓ Carsten Griwodz
email: griff @ ifi
- ✓ Pål Halvorsen
email: paalh @ ifi

About INF5062: Topic & Scope

- ✓ **Content:** The course gives ...
 - ... an overview of asymmetric multi-core processors in general and network processor cards in particular (architectures and use)
 - ... an introduction of how to program the Intel IXP 2400 network processors
 - ... some ideas of how to use/program asymmetric multi-core processors (guest lectures and paper presentations)

About INF5062: Topic & Scope

- ✓ **Lab-assignments:**
An important part of the course are lab-assignments where the students should make a program for the **Intel IXP2400** network processor
 1. **protocol statistics** – download and run *wwpingbump* and then extend it to give processor, interface and protocol statistics
 2. **packet bridge with ARP support** – forward packet to correct interface (of 3 available)
 3. **transparent load balancer** – balance load and forward packets to the right machine in a cluster of two with same IP address
 4. **HTTP protocol translator** – add support in the transparent load balancer for HTTP streaming having an RTSP/RTP server

About INF5062: Exam

- ✓ **Prerequisite – mandatory assignments:**
 - lab assignment 1: protocol statistics
 - presentation of a relevant paper
- ✓ **Graded assignments (counting 33% each):**
 - lab assignment 3: transparent load balancer
 - deliver code
 - short demo/explanation of code
 - lab assignment 4: HTTP protocol translator
 - deliver code and a short report
 - present and demonstrate
- ✓ **Final oral exam (counting 33%): early December 2006**
 - excerpts IXP documentation
 - lecture slides
 - presented papers
 - content of lab assignments

INF5062 – programming asymmetric multi-core processors 2006 Carsten Griwodt & Pål Halvorsen

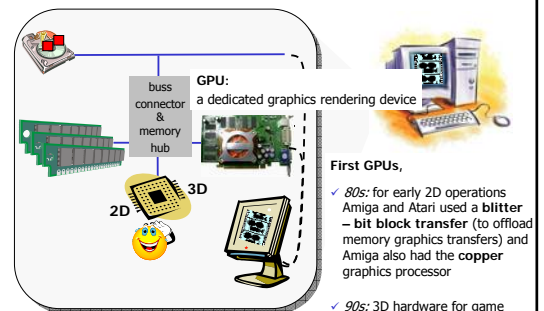
Available Resources

- ✓ Resources will be placed at
 - <http://www.ifi.uio.no/~{paalh | griff}/INF5062>
 - Login: *inf5062*
 - Password: *ixp*
 - Manuals, papers, code example, ...

INF5062 – programming asymmetric multi-core processors 2006 Carsten Griwodt & Pål Halvorsen

Background and Motivation 1: Graphics Processing Units

Graphics Processing Units (GPUs)



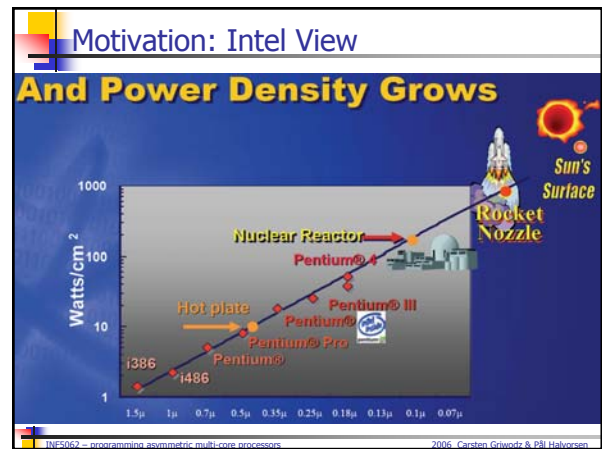
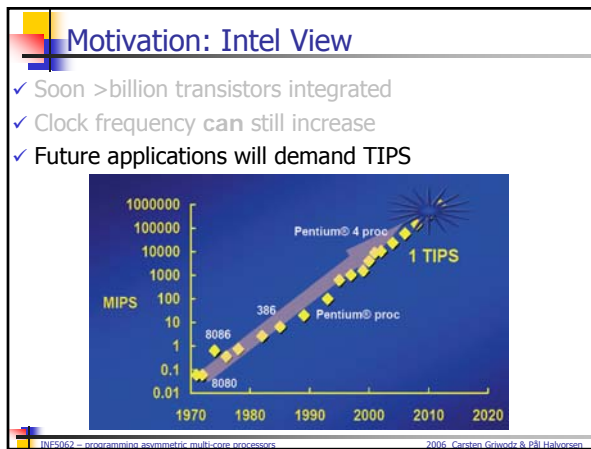
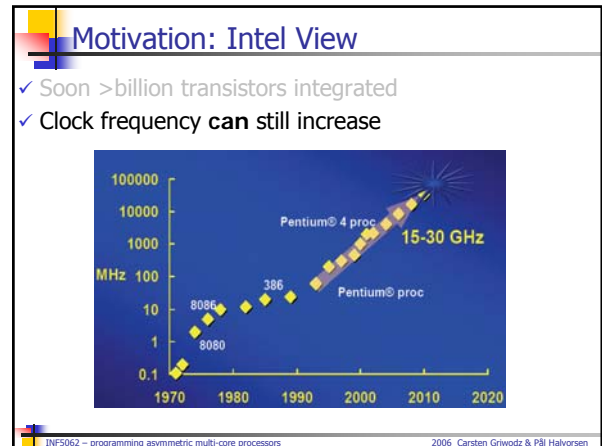
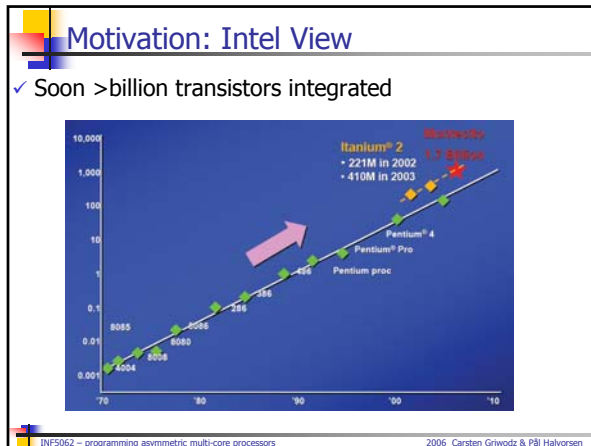
INF5062 – programming asymmetric multi-core processors 2006 Carsten Griwodt & Pål Halvorsen

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
- ✓ **BUT: how should it be programmed and which tasks should go where?**

INF5062 – programming asymmetric multi-core processors 2006 Carsten Griwodt & Pål Halvorsen

Background and Motivation 2: Moore's Law for single cores



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 consume less energy

↳ **asymmetric multi-core processors**

2006 Carsten Grawoetz & Pål Halvorsen

Co-Processors

- ✓ Commodore Amiga was one of the earlier machines that used multiple processors
 - blitter & copper (as we saw for the GPUs)
 - Motorola 680x0
 - IBM power
 - old Motorola kept for backwards compatibility and parts of the OS not ported
- ✓ 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, where the i487 actually was a full 486DX knocking the 486SX to sleep.

2006 Carsten Grawoetz & Pål Halvorsen

Intel Multi-Core Processors

- ✓ Single-die multi-processors is the new era and the next logical step in driving Moore's Law into another decade

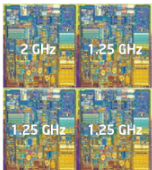


Figure 5. In an asymmetric multi-processor, the cores can run at different speeds to dynamically vary the CPL.




Figure 6. 25% relative performance boost for parallel programs

Figure 3. Run enables higher

Highly parallel Moderately parallel Sequential

INFS067 - programming asymmetric multi-core processors 2006, Carsten Grawodt & Pål Halvorsen

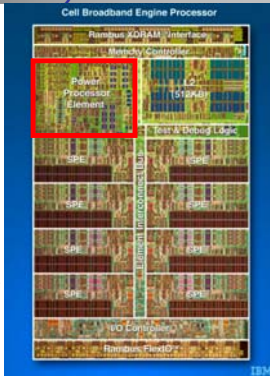
Intel Multi-Core Processors

- ✓ The operating systems can handle only a limited number of threads, e.g., 64 in Windows (2005)
- ✓ Where does the doubling stop? How many applications need more than 64 threads? Performance?
 - ➔ software limits is the issue
- ✓ Application specific engines?
 - ➔ Intel Academic Forum 2006: YES
 - ➔ **But: Programming model?????**

INFS067 - programming asymmetric multi-core processors 2006, Carsten Grawodt & Pål Halvorsen

STI (Sony, Toshiba, IBM) Cell

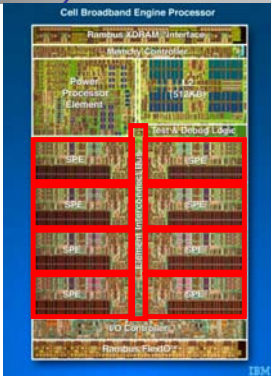
- ✓ Cell is a 9-core processor
 - combining a light-weight general-purpose processor with multiple co-processors into a coordinated whole
 - Power Processing Element (PPE)
 - conventional Power processor
 - not supposed to perform all operations itself, acting like a controller
 - running conventional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



INFS067 - programming asymmetric multi-core processors 2006, Carsten Grawodt & Pål Halvorsen

STI (Sony, Toshiba, IBM) Cell

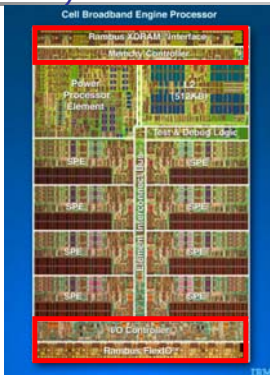
- Synergistic Processing Elements (SPE)
 - specialized co-processors for specific types of code, i.e., very high performance vector processors
 - local stores
 - can do general purpose operations
 - the PPE can start, stop, interrupt and schedule processes running on an SPE
- Element Interconnect Bus (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



INFS067 - programming asymmetric multi-core processors 2006, Carsten Grawodt & Pål Halvorsen

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



INFS067 - programming asymmetric multi-core processors 2006, Carsten Grawodt & Pål Halvorsen

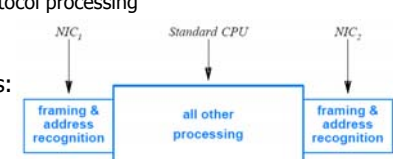
STI (Sony, Toshiba, IBM) Cell

- 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
 - 7 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

INFS067 - programming asymmetric multi-core processors 2006, Carsten Grawodt & Pål Halvorsen

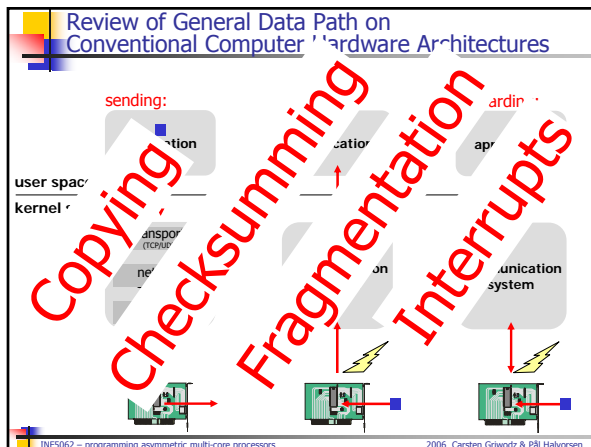
Background and Motivation 3: Network traffic increase

Software-Based Network System

- ✓ Uses conventional, shared hardware (e.g., a PC)
- ✓ Software
 - runs the entire system
 - allocates memory
 - controls I/O devices
 - performs all protocol processing
- ✓ First generation network systems:
 

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwaldt & Pål Halvorsen

Review of General Data Path on Conventional Computer Hardware Architectures

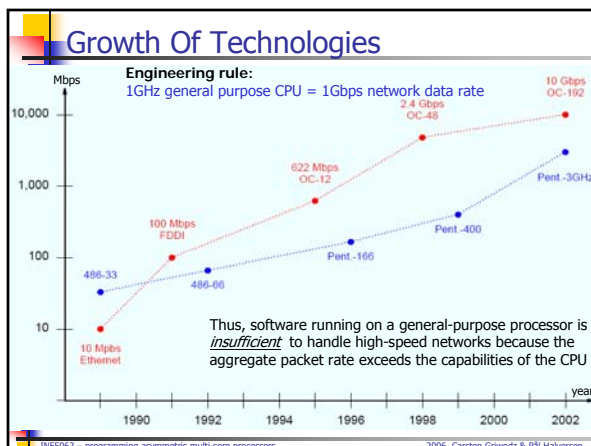


INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwaldt & Pål Halvorsen

Question:

- ✓ Which is growing faster?
 - network bandwidth
 - processing power
- ✓ Note: if network bandwidth is growing faster
 - CPU may be the bottleneck
 - need special-purpose hardware
 - conventional hardware will become irrelevant
- ✓ Note: if processing power is growing faster
 - no problems with processing
 - network/busses will be bottlenecks

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwaldt & Pål Halvorsen

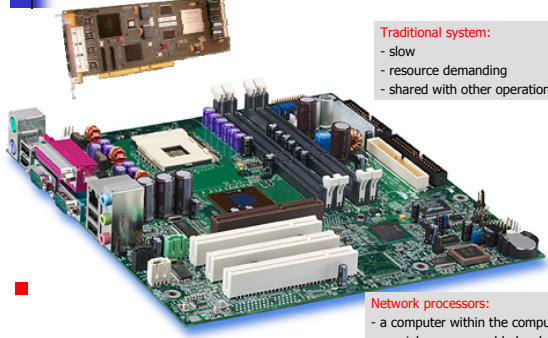


Network Processors: The Idea in a Nutshell

- ✓ Many designs through many generations (varying amount of HW & SW)
- ✓ Include support for protocol processing and I/O on one chip
 - General-purpose processor(s) for control tasks
 - Special-purpose processor(s) for packet processing and table lookup
 - Include functional units for tasks such as checksum computation, hashing, ...
- ⇒ Call the result a *network processor*
- ✓ We are here – they exist
- ⇒ **BUT: programming model??**

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwaldt & Pål Halvorsen

Network Processors: Main Idea



Traditional system:

- slow
- resource demanding
- shared with other operations

Network processors:

- a computer within the computer
- special, programmable hardware
- offloads host resources

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen

Explosion of Commercial Products

- ✓ 1990 → 2000: network processors transformed from interesting curiosity to mainstream product
 - reduction in both overall costs and time to market
- 2002: over 30 vendors with a wide range of architectures
- e.g.,
 - Multi-Chip Pipeline (Agere)
 - Augmented RISC Processor (Alchemy)
 - Embedded Processor Plus Coprocessors (Applied Micro Circuit Corporation)
 - Pipeline of Homogeneous Processors (Cisco)
 - Pipeline of Heterogeneous Processors (EZchip)
 - Configurable Instruction Set Processors (Cognigine)
 - Extensive And Diverse Processors (IBM)
 - Flexible RISC Plus Coprocessors (Motorola)
 - Internet Exchange Processor (Intel)
 - ...

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen

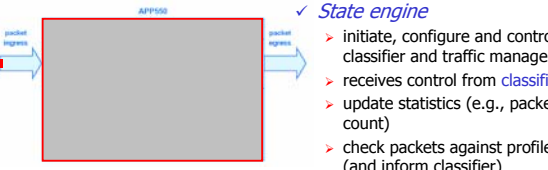
Agere PayloadPlus: A Short Overview

Agere PayloadPlus (APP)

- ✓ Agere PayloadPlus (APP)
 - consists of both programmable hardware and software
 - consists of both data and control planes (i.e., slow and fast plane)
- ✓ APP defines HW architectures, SW mechanisms, interconnection mechanisms and interfaces, **BUT** does not specify how to implement them
- ✓ Several versions of APP exist differing in the number and types of functional units, degree of parallelism and internal bandwidth (2. generation: 5 models)

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen

APP Conceptual Pipeline



✓ **Classifier**

- extract packets from ingress
- classify packet
- send statistics to **state engine**
- reassemble blocks
- pass packet to **forwarder** together with classification decision

✓ **State engine**

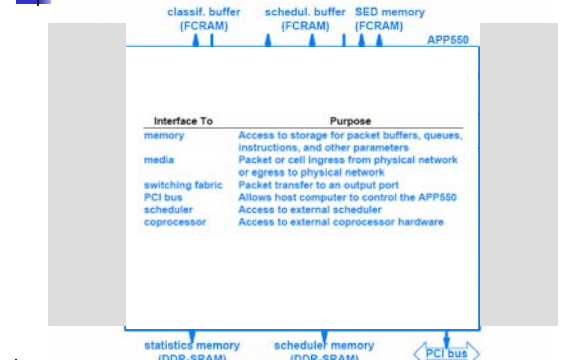
- initiate, configure and control classifier and traffic manager
- receives control from **classifier**
- update statistics (e.g., packet count)
- check packets against profiles (and inform classifier)

✓ **Forwarder**

- get packet from **classifier**
- perform traffic shaping and management
- fragment packet (if necessary)
- modify headers (if necessary)

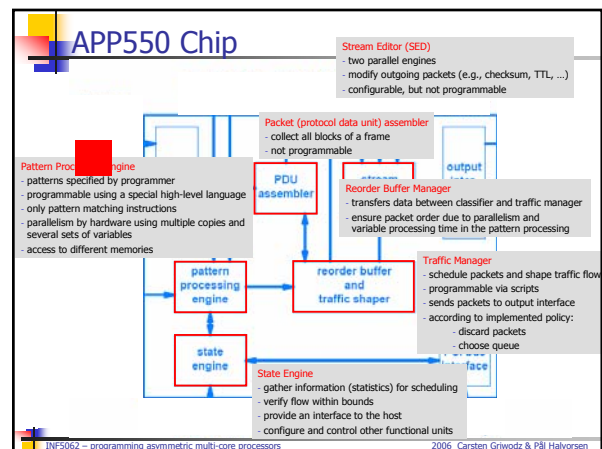
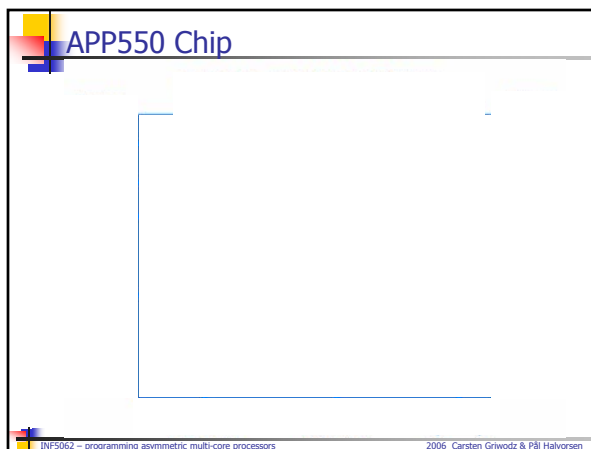
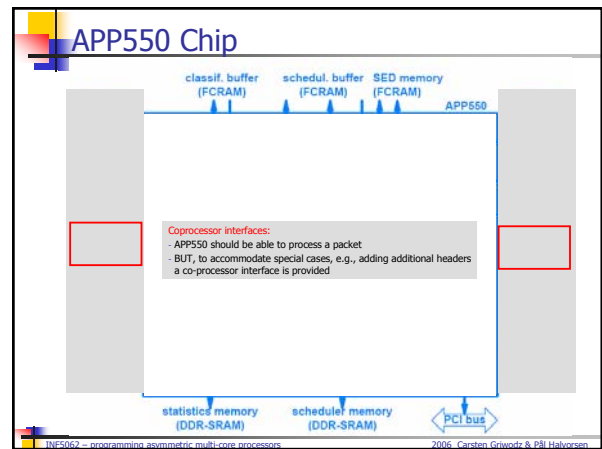
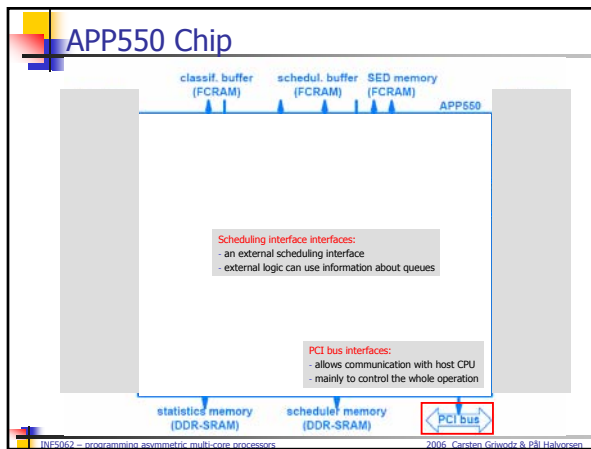
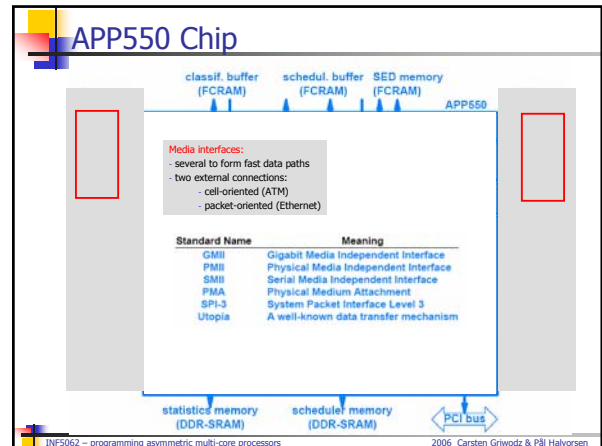
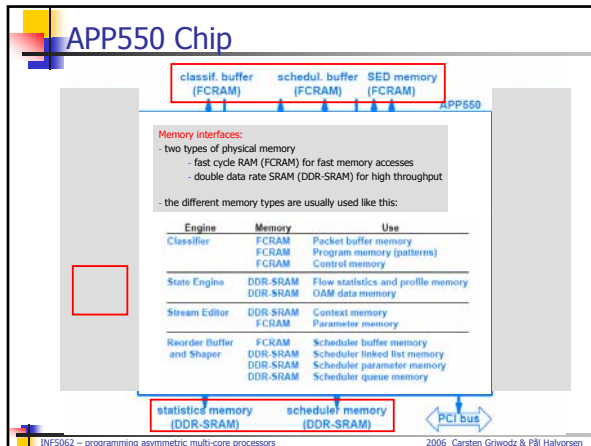
INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen

APP550 Chip



Interface To	Purpose
memory	Access to storage for packet buffers, queues, instructions, and other parameters
media	Packet or cell ingress from physical network or egress to physical network
switching fabric	Packet transfer to an output port
PCI bus	Allows host computer to control the APP550
scheduler coprocessor	Access to external scheduler hardware

INF5062 – programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen



APP550 Full Duplex

- ✓ Clock rate for APP550 is 233 MHz
- ✓ One chip cannot manage packet at wire speed in both directions – often two in parallel (one each direction)

- overhead: all features needed in both direction?
- classification only one direction
 - ➔ checks outgoing packets and enqueues using special queue

INF5067 – programming asymmetric multi-core processors 2006 Carsten Grunz & Pål Halvorsen

Intel IXP1200 / 2400: A Short Overview

IXA: Internet Exchange Architecture

- ✓ IXA is a broad term to describe the Intel network architecture (HW & SW, control- & data plane)
- ✓ IXP: Internet Exchange Processor
 - processor that implements IXA
 - IXP1200 is the first IXP chip (4 versions)
 - IXP2xxx has now replaced the first version
- ✓ IXP1200 basic features
 - 1 embedded 232 MHz StrongARM
 - 6 packet 232 MHz *µ*engines
 - onboard memory
 - 4 x 100 Mbps Ethernet ports
 - multiple, independent busses
 - low-speed serial interface
 - interfaces for external memory and I/O busses
 - ...
- ✓ IXP2400 basic features
 - 1 embedded 600 MHz XScale
 - 8 packet 600 MHz *µ*engines
 - 3 x 1 Gbps Ethernet ports
 - ...

INF5067 – programming asymmetric multi-core processors 2006 Carsten Grunz & Pål Halvorsen

IXP1200 Architecture

SRAM bus:

- shared bus (several external units)
- usually control rather than data
- rate 3.71 Gbps

PCI bus:

- allow IXP to connect to I/O devices
- enable use of host CPU
- rate 2.2 Gbps

Serial line:

- connects to the RISC
- intended for control and management
- rate 38 Kbps

SDRAM bus:

- provide access to external SDRAM memory used to store packets
- can also pass addresses, control/store operations, etc.
- rate 7.42 Gbps

IX (Intel eXchange) bus:

- enable higher rates compared to PCI form fast path (IXP and high-speed interfaces)
- interface to other DP cards
- 4.4 Gbps

INF5067 – programming asymmetric multi-core processors 2006 Carsten Grunz & Pål Halvorsen

IXP1200 Architecture

RISC processor:

- StrongARM running Linux
- control, higher layer protocols and exceptions
- 232 Mhz

Access units:

- coordinate access to external units

Scratchpad:

- on-chip memory
- used for IPC and synchronization

Microengines:

- low-level devices with limited set of instructions
- transfers between memory devices
- packet processing
- 232 Mhz

INF5067 – programming asymmetric multi-core processors 2006 Carsten Grunz & Pål Halvorsen

IXP1200 Processor Hierarchy

Processor Type	Onboard?	Programmable?

General-Purpose Processor:

- used for control and management
- running general applications

RISC processor:

- chip configuration interface (serial line)
- control, higher layer protocols and exceptions

I/O processors (microengines):

- transfers between memory devices
- packet processing

Coprocessors:

- real-time clock and timers
- IX bus controller
- hashing unit
- ...

Physical interface processors:

- implement layer 1 & 2 processing

INF5067 – programming asymmetric multi-core processors 2006 Carsten Grunz & Pål Halvorsen

IXP2400 Memory Hierarchy

Memory Type	Maximum Size	On Chip?	Typical Use
GP Registers	128 regs.	yes	Intermediate computation
Inst. Cache	16 Kbytes	yes	Recently used instructions
Data Cache	8 Kbytes	yes	Recently used data
Mini Cache	512 bytes	yes	Data that is reused once
Write buffer	unspecified	yes	Write operation buffer
Scratchpad	4 Kbytes	yes	IPC and synchronization
Inst. Store	64 Kbytes	yes	Microengine instructions
FlashROM	8 Mbytes	no	Bootstrap
SRAM	8 Mbytes	no	Tables or packet headers
SDRAM	256 Mbytes	no	Packet storage

INFS062 - programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen

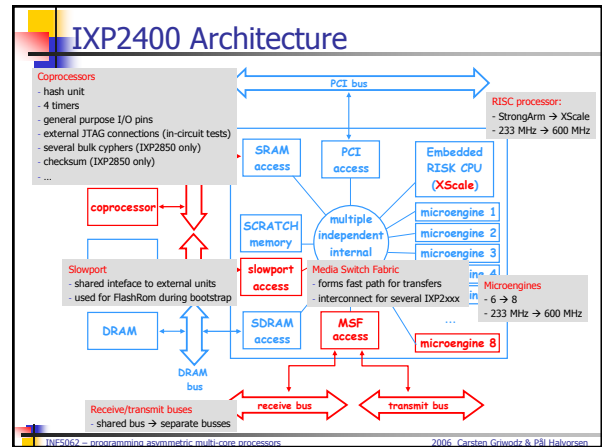
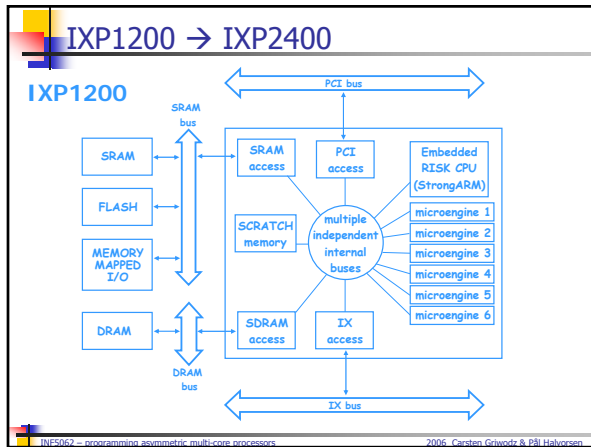
IXP2400 Memory Hierarchy

Memory Type	Addressable Data Unit (bytes)	Relative Access Time	Special Features

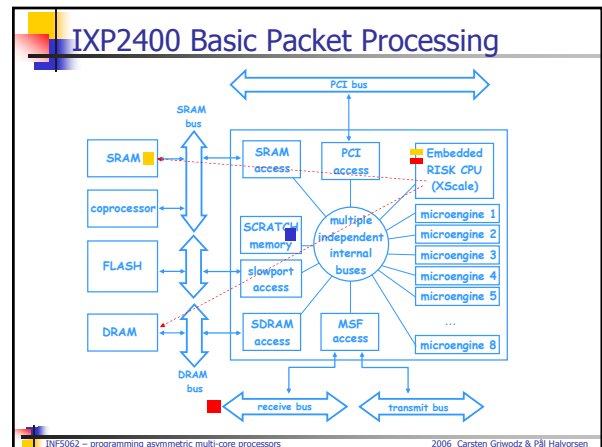
Different **memory** types...
 ✓ ...are organized into different addressable data units (words or longwords)
 ✓ ...have different access times
 ✓ ...connected to different buses

Therefore, to achieve optimal performance, programmers must **understand the organization** and **allocate** items from the **appropriate type**

INFS062 - programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen



- ### IXP2400 Architecture
- ✓ Memory
 - generally more of everything
 - generally larger gap between CPUs and memory access in terms of cycles
 - **local memory** on each microengine
 - saving temporary results
 - private per packet processor
 - small (2560 bytes)
 - low latency (one cycle)
 - accessed through special registers
- INFS062 - programming asymmetric multi-core processors 2006 Carsten Grunwald & Pål Halvorsen



The End: Summary

- ✓ Asymmetric multi-core processors are already
- ⇒ **Challenge: programming**
 - should know the capabilities of the system
 - identify which parts of a program that should run where
 - different methods to program the different components
- ✓ We will use **Intel IXP2400** as an example which offers...
 - ...embedded processor plus parallel packet processors
 - ...connections to external memories and buses
- ✓ Next time: how to start programming these monsters