ECE 697J – Advanced Topics in Computer Networks

The Intel Internet Exchange Architecture

10/30/03

Overview

- We have looked at
  - Network processor concepts
  - Different NP architectures
  - Commercial examples
- Next few weeks
  - In-depth look at one particular architecture
  - Intel Internet Exchange Architecture
  - Intel IXP1200 network processors
- Looking at details of platform will bring up some issues that we have ignored so far
- Good opportunity to discuss some software concepts

Intel IXA

- Intel Internet eXchange Architecture (IXA)
  - Intel's offering of network processors
  - Combination of NP hardware and software environment
- IXA Hardware:
  - Intel offers several network processors
  - We'll discuss IXP1200 (Internet eXchange Processor)
  - IXP1200 is one of the most widely used NPs in research
- IXA Software:
  - Software Development Environment (SDK 2.0) for packet processors and control processor
  - Allows detailed simulation without hardware
  - Controls hardware
- We'll start with hardware

Intel IXP NPs

- Intel offers variety of network processors
- IXP1200
  - OC-3 to OC-12
  - Diverse applications, one of the first NPs
- IXP420, IXP421, IXP422, IXP425
  - Low-end, access points, home office
  - VPN, VoIP, firewall, wireless
- IXP2400
  - OC-12 to OC-48
  - Network access and edge
- IXP2800
  - OC-48 to OC-192
  - Network edge and core
- IXP2850
  - Added encryption coprocessors
  - VPNs, SANs

IXP1200

IXP1200 Features

- Single chip with
  - One embedded RISC processor
  - Six programmable packet processors
  - Multiple, independent onboard buses
  - Processor synchronization mechanisms
  - Small amount of onboard memory
  - One low-speed serial line interface
  - Multiple interfaces for external memories
  - Multiple interfaces for external I/O buses
  - A coprocessor for hash computation
  - Other functional units

System Architecture

- Two memory buses
- One PCI bus
- One IX bus for network interfaces
- Serial line interface
- SRAM bus also for FlashROM

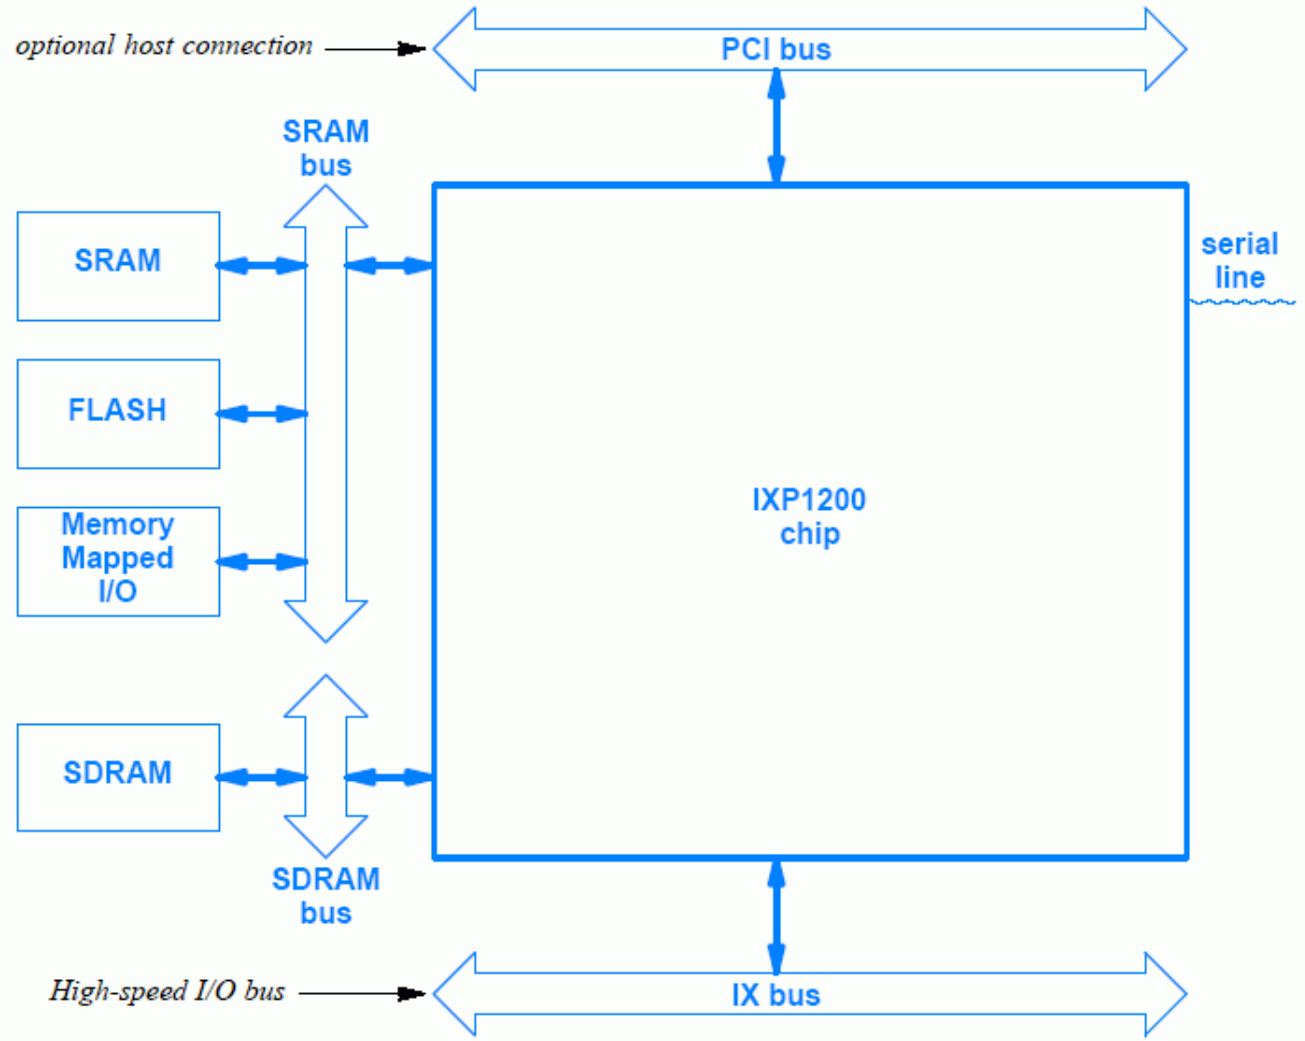

IXP1200 Bus Speeds

- Bandwidth for buses:

Type	Bus Width	Clock Rate	Data Rate
Serial line	(NA)	(NA)	38.4 Kbps
PCI bus	32 bits	33-66 MHz	2.2 Gbps
IX bus	64 bits	66-104 MHz	4.4 Gbps
SDRAM bus	64 bits	≤ 232 MHz	928.0 MBps†
SRAM bus	16 or 32 bits	≤ 232 MHz	464.0 MBps

- IX bus provides more bandwidth than PCI (as expected)
  - Connection to network interfaces
  - Connection to other IXPs for multi-chip architectures
- SDRAM provides more bandwidth than SRAM

Internal IXP1200 Components

Quantity	Component	Purpose
1	Embedded RISC processor	Control, higher layer protocols, and exceptions
6	Packet processing engines	I/O and basic packet processing
1	SRAM access unit	Coordinate access to the external SRAM bus
1	SDRAM access unit	Coordinate access to the external SDRAM bus
1	IX bus access unit	Coordinate access to the external IX bus
1	PCI bus access unit	Coordinate access to the external PCI bus
several	Onboard buses	Internal control and data transfer

IXP1200 Organization

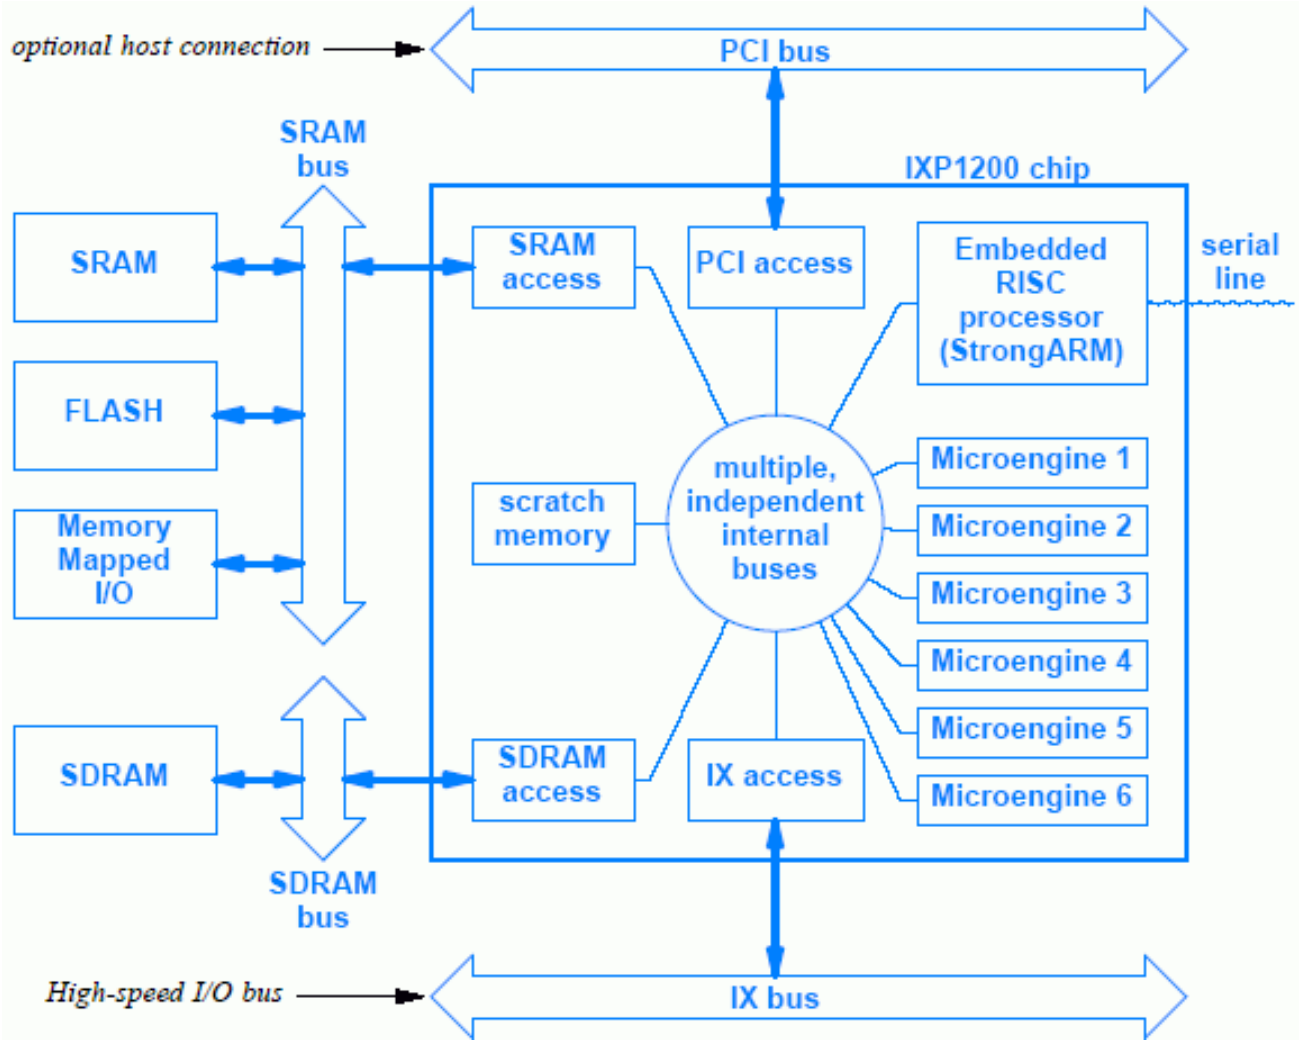

IXP1200 Processor Hierarchy

- Different processors:

Processor Type	Onboard?	Programmable?
General-Purpose Processor	no	yes
Embedded RISC Processor	yes	yes
I/O Processors	yes	yes
Coprocessors	yes	no
Physical Interfaces	no	no

- General-Purpose Processor
  - CPU of host system on which IXP resides
  - High-level functions, like routing
- Embedded RISC Processor (StrongARM)
  - Runs conventional operating system (e.g., Linux)
  - Manages Microengines, configures system, processes exceptions
- I/O Processors (Microengines)
  - Fast path, main packet processors

IXP1200 Coprocessors

- Hash Unit
  - Computes 48-bit or 64-bit hash in hardware
  - Adaptive polynomial hash function
  - Anybody have any idea why this could be useful?? ☺
- Four timers
- Real-time clock
- JTAG interface for testing
- IX bus controller
- FBI (FIFO Bus Interface) unit

IXP1200 Memory Hierarchy

- Different Memories:

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

IXP1200 Memories

- Caches are transparent to programmer
- Memories that programmer focuses on:
  - SRAM, SDRAM, Scratch (Stratchpad)
- Other memory features:

Memory Type	Addressable Data Unit (bytes)	Relative Access Time	Special Features
Scratch	4	12 - 14	synchronization via test-and-set and other bit manipulation, atomic increment
SRAM	4	16 - 20	stack manipulation, bit manipulation, read/write locks
SDRAM	8	32 - 40	direct transfer path to I/O devices

Synchronization

- How to implement synchronization between processors/threads?
- Example: shared packet counter
- Problem: increment requires read and write
  - Write after write (WAW) causes problems
- Solution?
  - Semaphore: atomic test and set operation
- Other problems
  - Deadlocks

Memory Addressing

- Addressable data units vary with memory type
  - Smallest addressable unit
- Typical units:
  - Word (16 bits)
  - Longword/long (32 bits)
  - Quadword (64 bits)
- Programmer needs to carefully plan memory layout
  - Address needs to be adapted to addressable units
  - Data structures that cross boundaries require multiple accesses

Complexity

- Each unit has many components
- Example: SRAM access unit

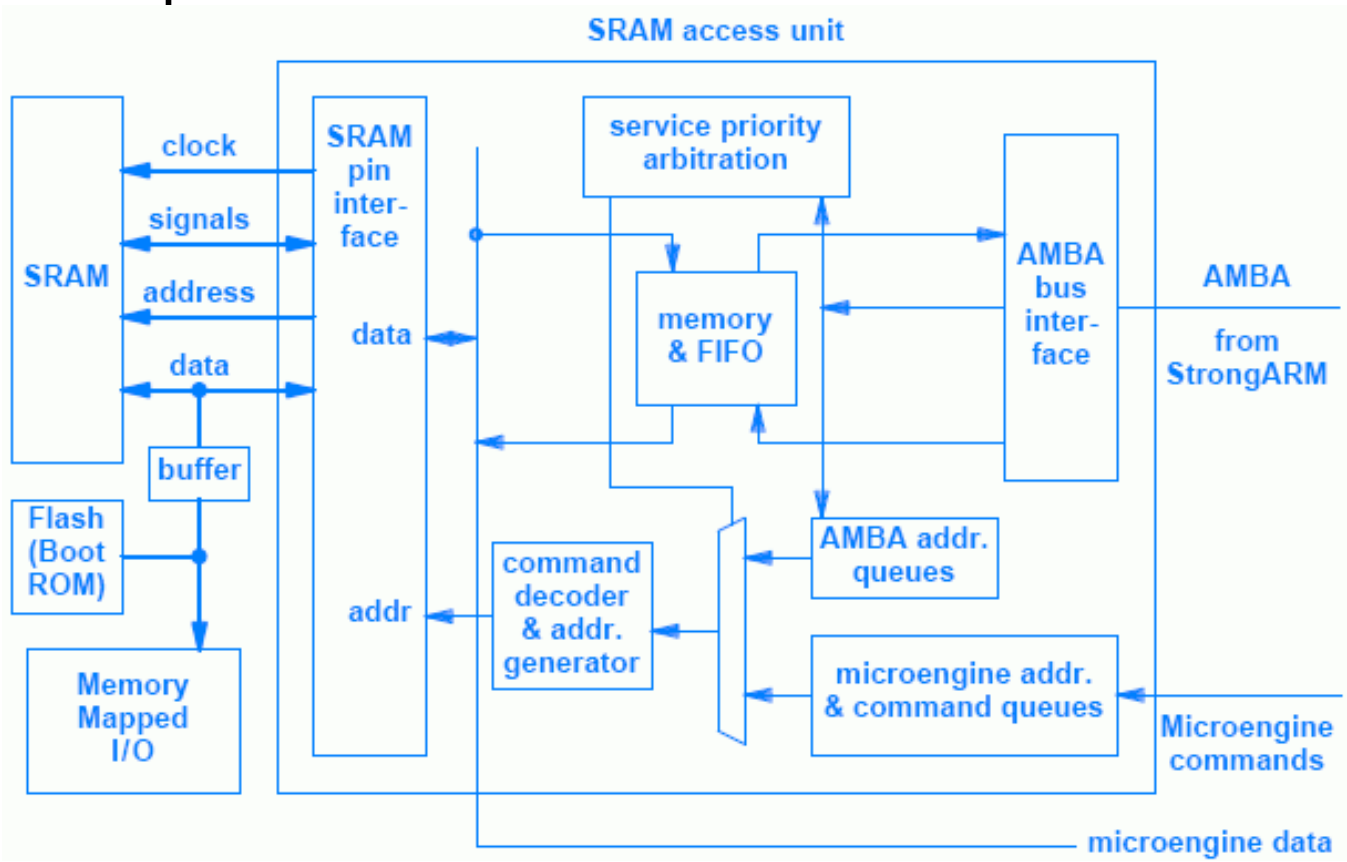

Next Class

- More details on IXP1200 will be covered
  - Embedded RISC Processor (StrongARM core) – Chapter 19
  - Packet Processing Hardware (Microengines, FBI) – Chapter 20
- Read Chapters 19 and 20