

ESE532: System-on-a-Chip Architecture

Day 8: February 8, 2017
Data Movement
(Interconnect, DMA)



Today

- Interconnect Infrastructure
- Data Movement Threads
- Peripherals
- DMA

Message

- Need to move data
- Shared Interconnect to make physical connections
- Useful to move data as separate thread of control
- Dedicating a processor to move data is inefficient
- Useful to have dedicated data-movement hardware: DMA

Memory and I/O Organization

- Architecture contains
 - Large memories
 - For density, necessary sharing
 - Small memories local to compute
 - For high bandwidth, low latency, low energy
 - Peripherals for I/O
- Need to move data
 - Among memories and I/O
 - Large to small and back
 - Among small
 - From Inputs, To Outputs

How move data?

- Abstractly, using stream links.
- Connect stream between producer and consumer.

- Ideally: dedicated wires

Dedicated Wires?

- Why might we not be able to have dedicated wires?

Making Connections


- Cannot always be dedicated wires
 - Programmable
 - Wires take up area
 - Don't always have enough traffic to consume the bandwidth of point-to-point wire
 - May need to serialize use of resource
 - E.g. one memory read per cycle

Penn ESE532 Spring 2017 -- DeHon

7

Model

- Programmable, possibly shared interconnect




Penn ESE532 Spring 2017 -- DeHon

8

Simple Realization

Shared Bus


- Write to bus with address of destination
- When address match, take value off bus
- Pros?
- Cons?



Penn ESE532 Spring 2017 -- DeHon

Alternate: Crossbar


- Provide programmable connection between all sources and destinations
- Any destination can be connected to any single source



Penn ESE532 Spring 2017 -- DeHon

10

Crossbar




Penn ESE532 Spring 2017 -- DeHon

11

Preclass 1

- K-input, O-output, W-bit wide Crossbar
- How many 2-input muxes?




Penn ESE532 Spring 2017 -- DeHon

12

Crossbar

- Provides high bandwidth
 - Minimal blocking
- Costs large amounts of area
 - Grows fast with inputs, outputs



13

Penn ESE532 Spring 2017 -- DeHon

General Interconnect

- Generally, want to be able to parameterize designs
- Here: tune area-bandwidth
 - Control how much bandwidth provide

14

Penn ESE532 Spring 2017 -- DeHon

Interconnect


- How might get design points between bus and crossbar?

15

Penn ESE532 Spring 2017 -- DeHon

Multiple Busses


- Think of crossbar as one bus per output
- Simple bus is one bus total
- In between,
 - How many simultaneous busses support?



Penn ESE532 Spring 2017 -- DeHon

Share Crossbar Outputs

- Group set of outputs together on a bus




17

Penn ESE532 Spring 2017 -- DeHon

Share Crossbar Inputs

- Group number of inputs together on an input port to crossbar



18

Penn ESE532 Spring 2017 -- DeHon


Locality in Interconnect

- How allow physically local items to be closer?

Penn ESE532 Spring 2017 -- DeHon

19


Hierarchical Busses



Penn ESE532 Spring 2017 -- DeHon

20

Mesh




Penn ESE532 Spring 2017 -- DeHon

21

Interconnect

- Will need an infrastructure for programmable connections
- Rich design space to tune area-bandwidth-locality
 - Will explore more later in course



Penn ESE532 Spring 2017 -- DeHon

22

Masters and Slaves

- Regardless of form, potentially have two kinds of entities on interconnect
- Master – can initiate requests
 - E.g. processor that can perform a read or write
- Slaves – can only respond to requests
 - E.g. memory that can return the read data from a read request

Penn ESE532 Spring 2017 -- DeHon

23

Long Latency Memory Operations

Penn ESE532 Spring 2017 -- DeHon

24

Last Time

- Large memories are slow
 - Latency increases with memory size
- Distant memories are high latency
 - Multiple clock-cycles to cross chip
 - Off-chip memories even higher latency

Day 7, Preclass 4

- 10 cycle latency to memory
- If must wait for data return, latency can degrade throughput
- 10 cycle latency + 10 op + (assorted)
 - More than 20 cycles / result

```
for(i=0;i<MAX;i++) {  
    in=a[i]; // memory read  
    out=f(in); // 10 cycle compute  
    b[i]=out;  
}
```

Preclass 2

- Throughput using 3 threads?

```
P1: for(i=0;i<MAX;i++) write_fifoA(a[i]);  
P2: while(1) write_fifoB(f(read_fifoA()));  
P3: for(i=0;i<MAX;i++) b[i]=read_fifoB();
```


Fetch (Write) Threads

- Potentially useful to move data in separate thread
- Especially when
 - Long (potentially variable) latency to data source (memory)
- Useful to split request/response

Peripherals


Input and Output

- Typical SoC has I/O with external world
 - Sensors
 - Actuators
 - Keyboard/mouse, display
 - Communications
- Also accessible from interconnect



Simple Peripheral Model


- Peripherals are slave devices
 - Masters can read input data
 - Masters can write output data
 - To move data, master (e.g. processor) initiates



Penn ESE532 Spring 2017 -- DeHon

Simple Model Implications

- What implication to processor grabbing/ moving each input (output) value?




Penn ESE532 Spring 2017 -- DeHon

32

Refine Model

- Give each peripheral local FIFO
- Processor must still move
- How does this change requirements and impact?



Penn ESE532 Spring 2017 -- DeHon

DMA

Penn ESE532 Spring 2017 -- DeHon

34

Preclass 3

- How much hardware to support fetch thread:
 - Counter bits?
 - Registers?
 - Comparators?
 - Other gates?
- Compare to MicroBlaze
 - (minimum config 630 6-LUTs)

Penn ESE532 Spring 2017 -- DeHon

35

Observe

- Modest hardware can serve as data movement thread
 - Much less hardware than a processor
 - Offload work from processors
- Small hardware allow peripherals to be Master devices on interconnect

Penn ESE532 Spring 2017 -- DeHon

36

DMA

- Direct Memory Access (DMA)
- Peripheral as Master
 - Can write directly into (read from) memory
 - Saves processor from copying

37

Penn ESE532 Spring 2017 -- DeHon

DMA Engine

- Data Movement Thread
 - Specialized Processor that moves data
- Act independently
- Implement data movement
- Can build to move data between memories (Slave devices)
- E.g., Implement P1, P3 in Preclass 3

38

Penn ESE532 Spring 2017 -- DeHon

DMA Engine

39

Penn ESE532 Spring 2017 -- DeHon

Programmable DMA Engine

- What copy from?
- Where copy to?
- Stride?
- How much?
- What size data?
- Loop?
- Transfer Rate?

40

Penn ESE532 Spring 2017 -- DeHon

Multithreaded DMA Engine

- One copy task not necessarily saturate bandwidth of DMA Engine
- Share engine performing many transfers (channels)
- Separate transfer state for each
 - Hence thread
- Swap among threads
 - E.g., round-robin

41

Penn ESE532 Spring 2017 -- DeHon

Zynq-7000 All Programmable SoC

Processing System: USB, 2x USB, 2x GigE, 2x SD, SD, SPI0, SPI1, SPI2, UART, CAN, CAN, CAN, I2C, I2C, I2C, SPI, SPI

Application Processor Unit: FPU and NEON Engine, ARM Cortex-A9 CPU, MMU, 32 KB I-Cache, 32 KB D-Cache, Snoop Controller, AWDT, Timer, DMA 8 Channel, OGM Interconnect, 256K SRAM, Memory Interconnect, Memory Interfaces, DDR2, DDR3, LPDDR2, LPDDR2 Controller, Consight Components, DAP, DAP, Programmable Logic to Memory Interconnect, High-Performance Ports, Programmable Logic, ACP, SelectIO Resources

Notes:
 1) Arrow direction shows control (master to slave)
 2) Data flows in both directions: AXI 32-Bit/64-Bit, AXI 64-Bit, AXI 32-Bit, AHB 32-Bit, APB 32-Bit, Custom
 3) Dashed line box indicates 2nd processor in dual-core devices

©2014 Xilinx, Inc.

Hardwired and Programmable

- Zynq has hardwired DMA engine
- Can also add data movement engines (Data Movers) in FPGA fabric

Penn ESE532 Spring 2017 -- DeHon

43

Big Ideas

- Need to move data
- Shared Interconnect to make physical connections – can tune area/bw
- Useful to
 - move data as separate thread of control
 - Have dedicated data-movement hardware: DMA

Penn ESE532 Spring 2017 -- DeHon

44

Admin

- Reading for Day 9 on web
- HW4 due Friday

Penn ESE532 Spring 2017 -- DeHon

45