

ESE532: System-on-a-Chip Architecture


Day 11: October 7, 2019
Data Movement
(Interconnect, DMA)



Penn ESE532 Fall 2019 -- DeHon

Preclass 1

- N processors
- Each: 1 read, 10 cycle compute, 1 write
- Memory: 1 read or write per cycle
- How many processors can support before saturate memory capacity?



Penn ESE532 Fall 2019 -- DeHon

2

Schedule Memory Port


12	13	14	15	16	17	18	19	20	21	22	23	24	25
P1.1 write	P1.2 read	P2.1 write	P2.2 read	P3.1 write	P3.2 read	P4.1 write	P4.2 read	P5.1 write	P5.2 read	P6.1 write	P6.2 read	P1.2 write	P1.3 read
P1 compute f on 2 nd iteration													
P2 compute f on 2 nd iteration													

Penn ESE532 Fall 2019 -- DeHon

3

Bottleneck

- Sequential access to a common memory can become the bottleneck



Penn ESE532 Fall 2019 -- DeHon

4


Previously

- Want data in small memories
 - Low latency, high bandwidth
- FPGA has many memories all over fabric

Penn ESE532 Fall 2019 -- DeHon

5

Embedded Memory in FPGA



ZU3EG (Ultra96) has 216 36Kb BRAMs

Penn ESE532 Fall 2019 -- DeHon

6

Previously

- Want data in small memories
 - Low latency, high bandwidth
- FPGA has many memories all over fabric
- Want C arrays in small memories
 - Partitioned so can perform enough reads (writes) in a cycle to avoid memory bottleneck

Today

- Interconnect Infrastructure
- Data Movement Threads
- Peripherals
- DMA -- Direct Memory Access

Message

- Need to move data
- Shared interconnect to make physical connections
- Useful to move data as separate thread of control
 - Dedicated a processor is inefficient
 - Useful to have dedicated data-movement hardware: Direct Memory Access (DMA)

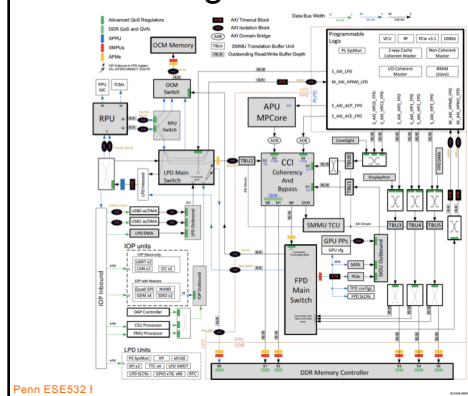
Memory and I/O Organization

- 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

Term: Peripheral

- “On the edge (or periphery) of something”
- Peripheral device – device used to put information onto or get information off of a computer
 - E.g.
 - Keyboard, mouse, modem, USB flash drive, ...

Programmable SoC



UG1085
Xilinx
UltraScale
Zynq
TRM
(p27)

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

Penn ESE532 Fall 2019 -- DeHon

13

How move data?

- Abstractly, using stream links.
- Connect stream between producer and consumer.
- Ideally: dedicated wires

Penn ESE532 Fall 2019 -- DeHon

14

Dedicated Wires?

- What might prevent us from having dedicated wires between all communicating units?

Penn ESE532 Fall 2019 -- DeHon

15

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
 - Source or destination may be sequentialized on hardware

Penn ESE532 Fall 2019 -- DeHon

16

Model

- Programmable, possibly shared interconnect




Penn ESE532 Fall 2019 -- DeHon

17

Simple Realization

Shared Bus


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



Penn ESE532 Fall 2019 -- DeHon

Alternate: Crossbar

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




19

Penn ESE532 Fall 2019 -- DeHon

Simplistic FPGA (illustrate possibility)

Day 8

- Every LUT input has a mux
- Every such mux has $m=(N+1)$ inputs
 - An input for each LUT output (N 2-LUTs)
 - An input for each Circuit Input (I Circuit inputs)
- Each Circuit Output has an m -input mux




20

Penn ESE532 Fall 2019 -- DeHon

Alternate: Crossbar

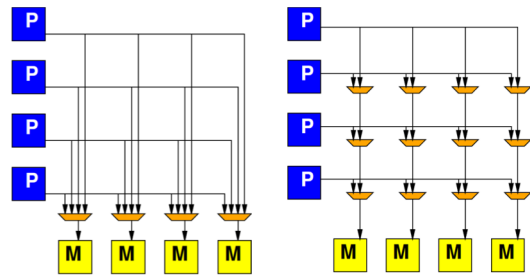
- Provide programmable connection between all sources and destinations
- Any destination can be connected to any single source



21

Penn ESE532 Fall 2019 -- DeHon

Crossbar




22

Penn ESE532 Fall 2019 -- DeHon

Preclass 2

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




23

Penn ESE532 Fall 2019 -- DeHon

Crossbar

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



24

Penn ESE532 Fall 2019 -- DeHon

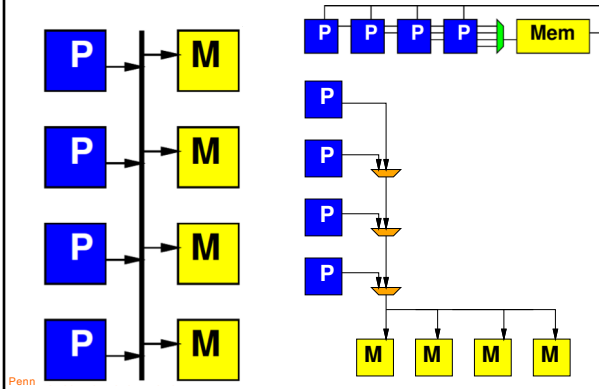
General Interconnect

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

Interconnect


- How might get design points between bus and crossbar?
- How could reduce number
 - Inputs to crossbar?
 - Outputs from crossbar?

Mux-based Bus




Multiple Busses

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




Share Crossbar Outputs

- Group set of outputs together on a bus




Share Crossbar Inputs

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



Delay

- Delay proportional to distance
- Pipeline bus to keep cycle time down
 - Take many cycles to travel long distance
 - ...but fewer cycles when distance small





31

Penn ESE532 Fall 2019 -- DeHon

Local Interconnect

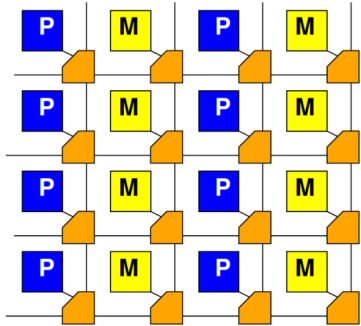
- How many cycles from:
 - PE3 to PE2
 - PE3 to PE1
 - PE3 to PE4



32

Penn ESE532 Fall 2019 -- DeHon


Mesh



33

Penn ESE532 Fall 2019 -- DeHon

Hierarchical Busses




34

Penn ESE532 Fall 2019 -- DeHon

Interconnect

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



35

Penn ESE532 Fall 2019 -- DeHon

Long Latency Memory Operations

36

Penn ESE532 Fall 2019 -- DeHon

Day 3

- 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 3, Preclass 2

- 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 3

- Throughput using 3 threads on 3 processors: P1, P2, P3?

```
P1: for(i=0;i<MAX;i++) Astream.write(a[i]);  
P2: while(1) {Astream.read(aval); Bstream.write(f(aval));}  
P3: for(i=0;i<MAX;i++) Bstream.read(b[i]);
```


Fetch (Write) Threads

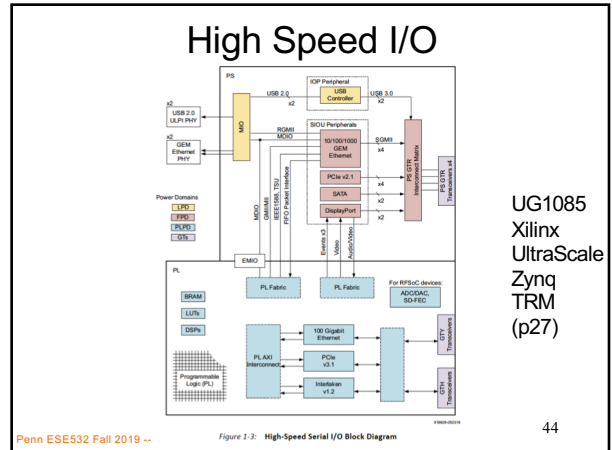
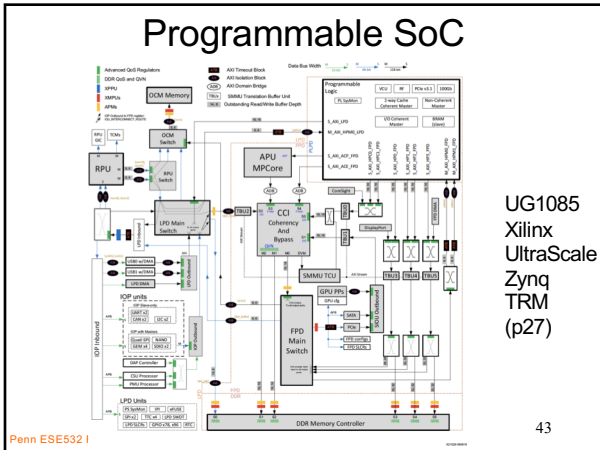
- Potentially useful to move data in separate thread
- Especially when
 - Long (potentially variable) latency to data source (memory)
- Useful to split request/response

Peripherals

Input and Output

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



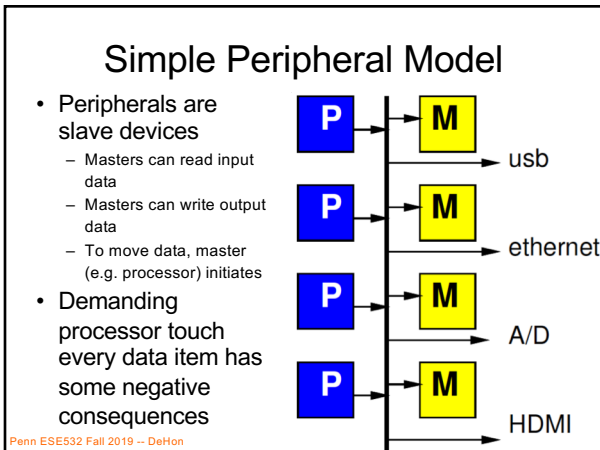
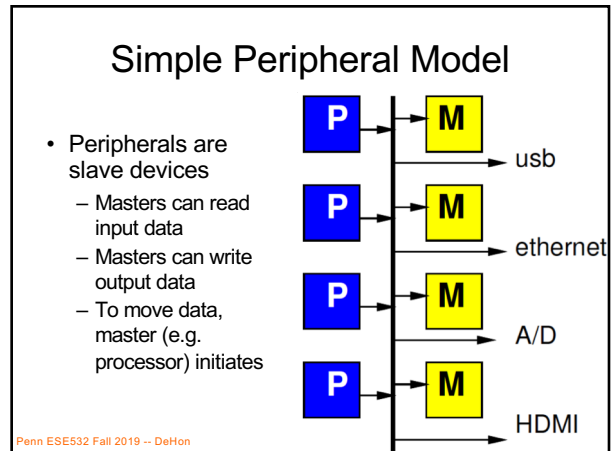


Masters and Slaves

- 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

45

Penn ESE532 Fall 2019 -- DeHon



Timing Demands


- Must read each input before overwritten
- Must write each output within real-time window
- Must guarantee processor scheduled to service each I/O at appropriate frequency
- How many cycles between 32b input words for 1Gb/s network and 32b, 1GHz processor?
 - Consider input data shifted into register 1b per ns
 - Must read out 32b register before overwritten

48

Penn ESE532 Fall 2019 -- DeHon

Refine Model

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



Penn ESE532 Fall 2019 -- DeHon

DMA


Direct Memory Access

50

Penn ESE532 Fall 2019 -- DeHon

Preclass 4a

```
P1: for(i=0;i<MAX;i++) Astream.write(a[i]);
```




```
int *p;
P1: for(p=&(a[0]);p<&(a[MAX]);i++) Astream.write(*p);
```

51

Penn ESE532 Fall 2019 -- DeHon

Preclass 4

- How much hardware?
 - Counter bits?
 - Registers?
 - Comparators?
 - Control Logic gates? (4cd)



- Compare to MicroBlaze
 - small RISC Processor optimized for Xilinx
 - minimum config 630 6-LUTs

52

Penn ESE532 Fall 2019 -- DeHon

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

53

Penn ESE532 Fall 2019 -- DeHon

DMA

- Direct Memory Access (DMA)
- “Direct” – inputs (and outputs) don’t have to be indirectly handled by the processor between memory and I/O
- I/O unit directly reads/writes memory



54

Penn ESE532 Fall 2019 -- DeHon

DMA

- Direct Memory Access (DMA)
- Peripheral as Master
 - Can write **directly** into (read from) memory
 - Saves processor from copying
 - Reduces demand to schedule processor to service

55

Penn ESE532 Fall 2019 -- 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

56

Penn ESE532 Fall 2019 -- DeHon

DMA Engine

57

Penn ESE532 Fall 2019 -- DeHon

Programmable DMA Engine

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

58

Penn ESE532 Fall 2019 -- 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 (or channel)
- Swap among threads
 - Simplest: round-robin:
 - 1, 2, 3, .. K, 1, 2, 3, .. K, 1,

59

Penn ESE532 Fall 2019 -- DeHon

Programmable SoC

60


Penn ESE532 I

UG1085
Xilinx
UltraScale
Zynq
TRM
(p27)

Hardwired and Programmable

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

UG1085
Xilinx
UltraScale
Zynq
TRM
Ch. 10
(p519)



61

Penn ESE532 Fall 2019 -- DeHon

Example

- Networking Application



- Header on processor
- Payload (encrypt, checksum) on FPGA
- DMA from ethernet → main memory
- DMA main memory → BRAM
- Stream between payload components
- DMA from chksum to ethernet out

62

Penn ESE532 Fall 2019 -- DeHon

Automation

- SDSoC will automatically take care of DMA of memory to and from accelerators in FPGA fabric
 - Inserting logic, programming DMA
- Mostly need to be aware is happening
 - Understand impacts on performance
- Have some options to control
 - With pragmas
 - With choice of data sizes
 - Explore HW6

63

Penn ESE532 Fall 2019 -- DeHon

Big Ideas

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

64

Penn ESE532 Fall 2019 -- DeHon

Admin

- Midterm Wednesday, here, classtime
 - Review Tuesday evening office hours
- Day 13 (Monday) reading
 - Chapter nine of *Parallel Programming for FPGAs* (available on web)
 - DRAM reading if not read on Day 3
- HW6 out soon??
 - We learned quite a bit putting it together
 - Hopefully, assignment passes on to you

65

Penn ESE532 Fall 2019 -- DeHon