

# ESE532: System-on-a-Chip Architecture

Day 11: February 20, 2017  
Real Time



## Today

- Real Time
  - Demands
  - Challenges
    - Algorithms
    - Architecture
  - Analysis
  - Scheduling

## Message

- Real-Time applications demand different discipline from best-effort tasks
- Look more like synchronous circuits
- Can sequentialize like processor
  - But must avoid/rethink typical processor common-case optimizations

## Real Time Tasks

- What applications demand real-time computing tasks?

## Real-Time Tasks

- Human consumed media:
  - video, audio, games, UI, graphics
- Control
  - Anti-lock brakes, cruise-control, auto-pilot, UAV, self-driving car, industrial automation
- Stock trading
- Network traffic handling
- Crypto (avoid information leak)

## Real Time Guarantees

- What guarantees might we want for real-time tasks?

## Real-Time Guarantees

- Attention/processing within fixed interval
  - Sample new value every XX ms
  - Produce new frame every 30 ms
- Bounded response time
  - Respond to keypress within 20 ms
  - Detect object within 100 ms
  - Return search results within 200 ms

Penn ESE532 Spring 2017 -- DeHon

7

## Synchronous Circuit Model

- A simple synchronous circuit is a good “model” for real-time task
  - Run at fixed clock rate
  - Take input every cycle
  - Produce output every cycle
  - Complete computation between input and output
  - Designed to run at fixed-frequency
    - Critical path meets frequency requirement

Penn ESE532 Spring 2017 -- DeHon

8

## Preclass 1

- How implement spatial pipeline?



```
float a[MAX], b[MAX], c[MAX], d;  
// stuff to load/define a, b, d ... maybe an outer loop  
for (i=0; i<MAX; i++) c[i]=(a[i]+b[i])*d;
```

Penn ESE532 Spring 2017 -- DeHon

9

## Historically

- Real-Time concerns grew up in EE
  - Because an analog circuit was the only way could meet frequency demands
  - ...later a dedicated digital circuit...
- Where worried about
  - Signal processing, video, control, ...

Penn ESE532 Spring 2017 -- DeHon

10

## Technological Change

- Why not be satisfied with this answer today?
  - For real-time task need dedicated synchronous circuit?

Penn ESE532 Spring 2017 -- DeHon

11

## Performance Scaling

- As circuit speeds increased
  - Can meet real-time performance demands with heavy sequentialization
- Circuit and processor clocks
  - from MHz to GHz
- Many real-time task rates unchanged
  - 44KHz audio, 33 frames/second video
- Even 100MHz processor
  - Can implement audio in a small fraction of its computational throughput capacity

Penn ESE532 Spring 2017 -- DeHon

12

## HW/SW Co-Design

- Computer Engineers – know can implement anything as hardware or software
- Want freedom to move between hardware and software to meet requirements
  - Performance, costs, energy

Penn ESE532 Spring 2017 -- DeHon

13

## Real-Time Challenge

- Meet real-time demands / guarantees
  - Economically using programmable architectures
- Sequentialize and share resources with deterministic, guaranteed timing

Penn ESE532 Spring 2017 -- DeHon

14

## Preclass 2

- Time for loop iteration case (a)?


```
top: bzero r4, exit
     ld f1, r1 // f1<-*r1 (read a[i])
     ld f2, r2 // f2<-*r2 (read b[i])
     fpadd f3,f1,f2 // f3<-f1+f2 (floating-point)
     fpmul f5,f3,f4 // f5<-f3*f4 (floating-point)
     st f5, r3 // *r3<-f5 (write c[i])
     subi r4,#1,r4 // r4<-r4-1 (decrement i)
     addi r1,#4,r1 // r1<-r1+4 (update a ptr)
     addi r2,#4,r2 // r2<-r2+4 (update b ptr)
     addi r3,#4,r3 // r3<-r1+4 (update c ptr)
     b top
```

exit:

Penn ESE532 Spring 2017 -- DeHon

15

## Preclass 2 Processor



Penn ESE532 Spring 2017 -- DeHon

16

## Preclass 2

- Time for loop iteration case (b)?

```
top: bzero r4, exit
     ld f1, r1 // f1<-*r1 (read a[i])
     ld f2, r2 // f2<-*r2 (read b[i])
     fpadd f6,f1,f2 // f6<-f1+f2 (floating-point) ** different **
     fpmul f5,f3,f4 // f5<-f3*f4 (floating-point)
     subi r4,#1,r4 // r4<-r4-1 (decrement i)
     addi r1,#4,r1 // r1<-r1+4 (update a ptr)
     addi r2,#4,r2 // r2<-r2+4 (update b ptr)
     addi r3,#4,r3 // r3<-r1+4 (update c ptr)
     mv f3, f6 //f3 <- f6 ** new **
     st f5, r3 // *r3<-f5 (write c[i]) ** different place **
     b top
```

exit:

Penn ESE532 Spring 2017 -- DeHon

17

## Data-dependent hazard

- Stalls instruction pipeline
  - Only when data needed before computed

Penn ESE532 Spring 2017 -- DeHon

18

## Observe

- Instructions on “General Purpose” processors take variable number of cycles

## Preclass 3

- How many cycles?

```
sum=0;
for (i=0;i<32;i++) {
    sum+=(0-(b%2)) & a;
    b=b>>1;
    a=a<<1;
}
```

## Preclass 3

- How many cycles?

```
sum=0;
for (;b!=0;b=b>>1) {
    if (b%2==1)
        sum+=a;
    a=a<<1;
}
```

## Observe

- Data-dependent branching, looping
  - Means variable time for operations

## Two Challenges

1. Architecture – Hardware have variable (data-dependent) delay
  - Esp. for General-Purpose processors
    - Instructions take different number of cycles
2. Algorithm – computational specification have variable (data-dependent) operations
  - Different number of instructions

## Algorithm

- What programming constructs are data-dependent (variable delay)?

## Programming Constructs

- Conditionals: if/then/else
- Loops without compile-time determined bounds
  - While with termination expressions
  - For with data-dependent bounds
- Recursion
- Hash tables, memoization
- Interrupts
  - I/O events, time-slice

Penn ESE532 Spring 2017 -- DeHon

25

## Programming Constructs

- Dynamic Dataflow
  - Variable rates
  - Switch/select operators

Penn ESE532 Spring 2017 -- DeHon

26

## ...like Hardware

- Many problematic constructs similar to C/ Programming-Language constructs need to avoid for hardware
  - Dynamic allocation (malloc)
  - Recursive functions
  - Loops without determined bounds
  - Mux-conversion/predications for if/then/else

Penn ESE532 Spring 2017 -- DeHon

27

## Architecture

- What processor constructs are variable delay?

Penn ESE532 Spring 2017 -- DeHon

28

## Processor Variable Delay

- Data hazards
- Caches
- Data-dependent branching / branch delays
- Speculative issue
  - Out-of-Order, branch prediction
- Dynamic arbitration for shared resources
  - Bus, I/O, Crossbar output, memory, ...

Penn ESE532 Spring 2017 -- DeHon

29

## Cache Predictable?

- Is an element in or out of cache?
  - Accessed before?
  - Had an address conflict?
  - Depend on access pattern
- If shared
  - Did someone else write it?
  - Depends on everything else sharing

Penn ESE532 Spring 2017 -- DeHon

30

## Hardware Architecture

- Same “optimizations” can cause variable delay even in dedicated hardware data path
  - Caches
  - Common-case optimizations
  - Pipeline stalls


## What can we do?

- Explicitly managed memory
  - Scheduled
  - Multi-threaded
- Deadlines
- Offline-scheduled resource sharing

## Explicitly Managed Memory

- Make memory hierarchy visible
  - Use Scratchpad memories instead of caches
  - E.g. local memory in HW5
- Explicitly move data between memories
  - E.g. DMA into OCM, movement into local memory

## Explicitly Managed Memory




## Fixed Delays (1)

- Drop dynamic data hazards, branch speculation
- Data becomes available after a predictable time
- Branches take effect at a fixed time
  - Likely delayed
- Schedule to delays to get correct data

## Fixed Delay Example

- Branches occur
  - 1 cycle later (uncond)
  - 3 cycles later
- Non-FP data
  - Available on 2<sup>nd</sup> instr
- FP data
  - Available on 6<sup>th</sup> instr



## Preclass 4

- Where 2(b) loop code not work?

```

top: bzero r4, exit
    ld f1, r1 // f1<-*r1 (read a[i])
    ld f2, r2 // f2<-*r2 (read b[i])
    fpadd f6,f1,f2 // f6<-f1+f2 (floating-point) ** different **
    fpmul f5,f3,f4 // f5<-f3*f4 (floating-point)
    subi r4,#1,r4 // r4<-r4-1 (decrement i)
    addi r1,#4,r1 // r1<-r1+4 (update a ptr)
    addi r2,#4,r2 // r2<-r2+4 (update b ptr)
    addi r3,#4,r3 // r3<-r1+4 (update c ptr)
    mv f3, f6 //f3 <- f6                ** new **
    st f5, r3 // *r3<-f5 (write c[i])    ** different place **
    b top
exit:
    
```

Penn ESE532 Spring 2017 -- DeHon

37

## Preclass 4

- How need to fix?

```

top: bzero r4, exit
    ld f1, r1 // f1<-*r1 (read a[i])
    ld f2, r2 // f2<-*r2 (read b[i])
    fpadd f6,f1,f2 // f6<-f1+f2 (floating-point) ** different **
    fpmul f5,f3,f4 // f5<-f3*f4 (floating-point)
    subi r4,#1,r4 // r4<-r4-1 (decrement i)
    addi r1,#4,r1 // r1<-r1+4 (update a ptr)
    addi r2,#4,r2 // r2<-r2+4 (update b ptr)
    addi r3,#4,r3 // r3<-r1+4 (update c ptr)
    mv f3, f6 //f3 <- f6                ** new **
    st f5, r3 // *r3<-f5 (write c[i])    ** different place **
    b top
exit:
    
```

Penn ESE532 Spring 2017 -- DeHon

38

## Fixed-Delay (2)

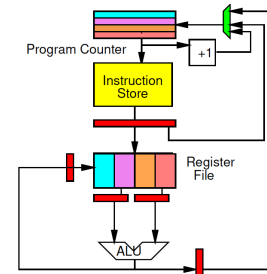
- Drop dynamic data hazards, branch speculation
- Pipeline processor
- But only feed one instruction per thread through processor at time
  - Each instruction completes before next issues (no dependencies)
- Use pipeline to issue from multiple threads
  - For throughput, not latency

Penn ESE532 Spring 2017 -- DeHon

39

## Multithreaded Pipeline

- Only one instruction per thread in pipeline
- Looks like PIPEDEPTH slower processors
- No interlock/bypass
  - Smaller control
  - Faster cycle?




Penn ESE532 Spring 2017 -- DeHon

40

## Multithreaded Pipeline

- Can run multiple threads
- Non-real-time threads can share
- Timing of threads not impact each other
- Non-real-time threads take variable time
  - Not interfere with real-time thread slots



Penn ESE532 Spring 2017 -- DeHon

41

## Deadline Instruction

- Set a hardware counter for thread
- Demand counter reach 0 before allowed to continue
- Orderly way to tolerate variable instructions in algorithm
- Model: fixed rate of attention
  - Stall if get there early
  - Similar to flip-flop on a logic path
    - Wait for clock edge to change value
- Model: fixed-time

Penn ESE532 Spring 2017 -- DeHon

42

## Offline Schedule Resource Sharing

- Don't arbitrate
- Decide up-front when each shared resource can be used by each thread or processor
  - Simple fixed schedule
  - Detailed Schedule
- What
  - Memory bank, bus, I/O, network link, ...



Penn ESE532 Spring 2017 -- DeHon

43

## Time-Multiplexed Bus

Fixed by hardware master

- 4 masters share a bus
- Each master gets to make a request on the bus every 4<sup>th</sup> cycle
  - If doesn't use it, goes idle



Penn ESE532 Spring 2017 -- DeHon

44

## Time Multiplexed Bus


- Regular schedule
- Fixed bus slot schedule of length  $N >$  masters
  - (probably a multiple)
- Assign owner for each slot
  - Can assign more slots to one
- E.g.  $N=8$ , for 4 masters
  - Schedule (1 2 1 3 1 2 1 4)

Penn ESE532 Spring 2017 -- DeHon

45

## Fully Scheduled

- At extreme, fully schedule which tasks gets resource on each cycle



Penn ESE532 Spring 2017 -- DeHon

46


## Fully Scheduled

- At extreme, fully schedule which tasks gets resource on each cycle
- Sensible if all master's sharing resource are also fully scheduled, running in lock-step
- Think of instruction field for bus

Penn ESE532 Spring 2017 -- DeHon

47

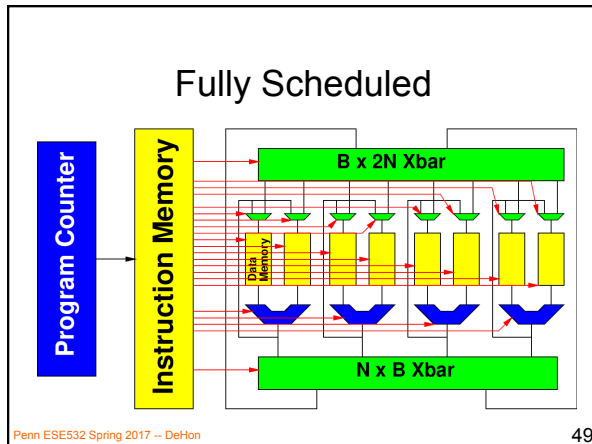
## Fully Scheduled (before instr)



Penn ESE532 Spring 2017 -- DeHon

48





- ### SoC Opportunity
- Can choose which resources are shared
  - Can dedicate resources to tasks
  - Isolate real-time tasks/portions of tasks from best-effort
    - Separate hardware/processors
    - Separate memories, network
- Penn ESE532 Spring 2017 -- DeHon 50

- ### Different Goals
- |  |  |
|--|--|
| <p><b>Real Time</b></p> <ul style="list-style-type: none"> <li>• Willing to recompile to new hardware</li> <li>• Want time on hardware predictable</li> <li>• Willing to schedule for delays in particular hardware</li> </ul> | <p><b>General Purpose</b></p> <ul style="list-style-type: none"> <li>• ISA fixed</li> <li>• Want to run same assembly on different implementations</li> <li>• Tolerate different delays for different hardware</li> <li>• Run faster on newer, larger implementations</li> </ul> |
|--|--|
- Penn ESE532 Spring 2017 -- DeHon 51

- ### WCET
- WCET – Worst-Case Execution Time
  - Analysis when working with algorithms and architectures with data-dependent delay
    - Need to meet real time
    - Calculate the worst-case runtime of a task
      - Like calculating the critical path (but harder)
      - Worst-case delay of instructions
      - Worst-case path through code
      - Worst-case # loop iterations
- Penn ESE532 Spring 2017 -- DeHon 52

- ### Big Ideas:
- Real-Time applications demand different discipline from best-effort tasks
  - Look more like synchronous circuits and hardware discipline
  - Can sequentialize like processor
    - But must avoid/rethink typical processor common-case optimizations
    - Offline calculate static schedule for computation and sharing
- Penn ESE532 Spring 2017 -- DeHon 53

- ### Admin
- Reading for Monday on Web
  - No lecture Wednesday
  - No lab/office hours this week
    - Use piazza
  - HW6 two week assignment (No HW7)
    - Recommend do base part (3,4) by Friday
    - Can finish (5) by following Friday
  - Exam next Wednesday
- Penn ESE532 Spring 2017 -- DeHon 54