

ESE532: System-on-a-Chip Architecture

Day 4: January 25, 2017
Data-Level Parallelism



Penn ESE532 Spring 2017 -- DeHon

Today

- Data-level Parallelism
 - For Parallel Decomposition
 - Architectures
 - Concepts
 - NEON

Penn ESE532 Spring 2017 -- DeHon

2

Message

- Data Parallelism easy basis for decomposition
- Data Parallel architectures can be compact – pack more computations onto a die

Penn ESE532 Spring 2017 -- DeHon

3

Preclass 1

- 300 news articles
- Count total occurrences of a string
 - How can we exploit data-level parallelism on task?
 - How much parallelism can we exploit?

Penn ESE532 Spring 2017 -- DeHon

4

Parallel Decomposition

Penn ESE532 Spring 2017 -- DeHon

5

Data Parallel

- Data-level parallelism can serve as an organizing principle for parallel task decomposition
- Run computation on independent data in parallel

Penn ESE532 Spring 2017 -- DeHon

6

Exploit

- Can exploit with
 - Threads
 - Instruction-level Parallelism
 - Fine-grained Data-Level Parallelism

Thread Exploit DP


- How exploit threads for data-parallel text search?

SPMD

- Single Program Multiple Data
- Only need to write code once
 - Get to use many times

ILP Exploit DP

- How exploit ILP for DP?



Pipeline Exploit

- How exploit hardware pipeline for text search?

Common Examples

- What are common examples of DLP?
 - Signal Processing
 - Simulation
 - Numerical Linear Algebra
 - Graphics
 - Image Processing
 - Optimization
 - Other?

Hardware Architectures


Idea

- If we're going to perform the same operations on different data, exploit that to reduce area, energy
- Reduced area means can have more computation on a fixed-size die.

SIMD


- Single Instruction Multiple Data

Shared Instruction



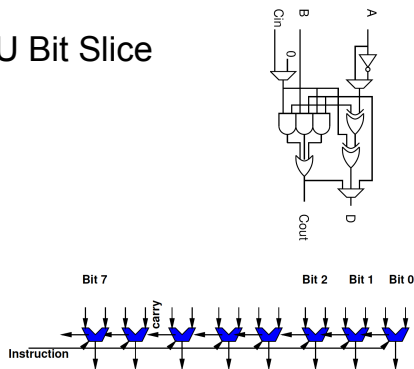
W-bit ALU as SIMD

- Familiar idea
- A W-bit ALU (W=8, 16, 32, 64, ...) is SIMD
- Each bit of ALU works on separate bits
 - Performing the same operation on it
 - Trivial to see bitwise AND, OR, XOR
 - Also true for ADD (each bit performing Full Adder)
- Share one instruction across all ALU bits




ALU Bit Slice

-




Register File

- Small Memory
- Usually with multiple ports
 - Ability to perform multiple reads and writes simultaneously
- Small
 - To make it fast (small memories fast)
 - Multiple ports are expensive



Preclass 2

- Area $W=16$?
- Area $W=128$?
- Number in 10^8
 - $W=16$
 - $W=128$
- Perfect Pack Ratio?




Penn ESL

19

Preclass 2

- W for single datapath in 10^8 ?
- Perfect 16b pack ratio?
- Compare $W=128$ perfect pack ratio?



Penn ESL

20

ALU vs. SIMD ?


- What's different between
 - 128b wide ALU
 - SIMD datapath supporting eight 16b ALU operations

Penn ESE532 Spring 2017 -- DeHon

21

Segmented Datapath

- Relatively easy (few additional gates) to convert a wide datapath into one supporting a set of smaller operations
 - Just need to squash the carry at points



Penn ESE532 Spring 2017 -- DeHon

22

Segmented Datapath


- Relatively easy (few additional gates) to convert a wide datapath into one supporting a set of smaller operations
 - Just need to squash the carry at points
- But need to keep instructions (description) small
 - So typically have limited, homogeneous widths supported

Penn ESE532 Spring 2017 -- DeHon

23

Segmented 128b Datapath

- 1x128b, 2x64b, 4x32b, 8x16b




Penn ESE532 Spring 2017 -- DeHon

24

Terminology: Vector Lane

- Each of the separate segments called a **Vector Lane**
- For 16b data, this provides 8 vector lanes



Penn ESE532 Spring 2017 -- DeHon

25

Opportunity

- Don't need 64b variables for lots of things
- Natural data sizes?
 - Audio samples?
 - Input from A/D?
 - Video Pixels?
 - X, Y coordinates for 4K x 4K image?

Penn ESE532 Spring 2017 -- DeHon

26

Vector Computation

- Easy to map to SIMD flow if can express computation as operation on vectors
 - Vector Add
 - Vector Multiply
 - Dot Product

Penn ESE532 Spring 2017 -- DeHon

27


Concepts

Penn ESE532 Spring 2017 -- DeHon

28

Vector Register File

- Need to be able to feed the SIMD compute units
 - Not be bottlenecked on data movement to the SIMD ALU
- Wide RF to supply
- With wide path to memory




Penn ESE532 Spring 2017 -- DeHon

29

Point-wise Vector Operations

- Easy – just like wide-word operations (now with segmentation)




Penn ESE532 Spring 2017 -- DeHon

30

Point-wise Vector Operations

- ...but alignment matters.
- If not aligned, need to perform data movement operations to get aligned



Penn ESE532 Spring 2017 -- DeHon

31

Vector Length

- May not match physical hardware length
- What happens when
 - Vector length > hardware SIMD operators?
 - Vector length < hardware SIMD operators?
 - Vector length % hwd operators != 0
 - E.g. vector length 20, for 8 hwd operators

Penn ESE532 Spring 2017 -- DeHon

32

Skipping Elements?

- How does this work with datapath?
- for (i=0; i<64; i=i+2)
 - $c[i]=a[i]+b[i]$

Penn ESE532 Spring 2017 -- DeHon

33

Stride

- Stride: the distance between vector elements used
- for (i=0; i<64; i=i+2)
 - $c[i]=a[i]+b[i]$
- Accessing data with stride=2

Penn ESE532 Spring 2017 -- DeHon

34

Load/Store

- Strided load/stores
 - Some architectures will provide strided memory access that compact when read into register file
- Scatter/gather
 - Some architectures will provide memory operations to grab data from different places to construct a dense vector

Penn ESE532 Spring 2017 -- DeHon

35

Conditionals?


- What happens if want to do something different?
- For (i=0; i<8; i++)
 - if (a[i]<b[i])
 - $d[i]=a[i]+c[i]$
 - else
 - $d[i]=b[i]+c[i]$

Penn ESE532 Spring 2017 -- DeHon

36

Conditionals

- Only have one Program Counter
 - Cannot implement conditional via branching




Penn ESE532 Spring 2017 -- DeHon

37

Conditionals

- Only have one instruction
 - Cannot perform separate operations on each ALU in datapath



Penn ESE532 Spring 2017 -- DeHon

38

Conditionals

- Only have one Program Counter
 - Cannot implement conditional via branching
- Only have one instruction
 - Cannot perform separate operations on each ALU in datapath
- Must perform an invariant operation sequence

Penn ESE532 Spring 2017 -- DeHon

39

Invariant Operation


- If $(a[i] < b[i])$
 - then $d[i] = a[i] + c[i]$
 - else $d[i] = b[i] + c[i]$
 - What's in each register as go through sequence?
1. $T1[i] = a[i] < b[i]$
 2. $T2[i] = \sim T1[i]$
 3. $T3[i] = \sim (T2[i])$
 4. $T2[i] = a[i] \& T2[i]$
 5. $T3[i] = b[i] \& T3[i]$
 6. $d[i] = c[i] + T2[i]$
 7. $d[i] = c[i] + T3[i]$

Penn ESE532 Spring 2017 -- DeHon

40

If → Mux Conversion

- Can always transform into a data independent sequence
- Often convenient to think of IF's as Multiplexers
- If (cond)
 - o = a
- else
 - o = b



Penn ESE534 Spring 2016 -- DeHon

41

Predicated Operation

- Many architectures will provide a predicated operation
 - $p[i]: d[i] = c[i] + a[i]$
 - $\sim p[i]: d[i] = c[i] + b[i]$
- Only perform operation when predicate matches instruction

Penn ESE532 Spring 2017 -- DeHon

42

Predicated Operation

- What does this do to instructions must be issued?
 - What does this do to efficiency?
 - Useful operations performed per cycle
- $p[i]=a[i]<b[i]$
 - $p[i]: d[i]=c[i] + a[i]$
 - $\sim p[i]: d[i]=c[i] + b[i]$

Penn ESE532 Spring 2017 -- DeHon

43

Nested Conditionals

- What happens with nested conditionals?

Penn ESE532 Spring 2017 -- DeHon

44

Dot Product

- What happens when need a dot product?
- $res=0;$
- for ($i=0;i<N;i++$)
 - $res+=a[i]*b[i]$

Penn ESE532 Spring 2017 -- DeHon

45

Reduction


- Common operations where want to perform a combining operation to reduce a vector to a scalar
 - Sum values in vector
 - AND, OR
- Reduce Operation

Penn ESE532 Spring 2017 -- DeHon

46

Reduce Tree

- Efficiently handled with reduce tree




Penn ESE532 Spring 2017 -- DeHon

47

Reduce in Pipeline

- Comes almost for free in pipeline



Penn ESE532 Spring 2017 -- DeHon


48

Vector Reduce Instruction

- Usually include support for vector reduce operation
 - Doesn't need to add much to delay
 - Maybe even faster than performing larger operation
 - 8 16x16 multiplies with sum reduce less complex than one 128x128 multiply
 - ...can exploit datapath of larger operation


Dot Product Revisited

- for (i=0;i<N;i++)
 - res+=a[i]*b[i]
- With 3 cycle pipelined multiply
- a in R0—R4
- b in R4—R7
- What happens if try to implement dot product as:
 - MPY R0, R4, R14
 - ADD R14, R15, R15
 - MPY R1, R5, R14
 - ADD R14, R15, R15
 - ...




Dot Product Revisited

- How should order (reformulate) instructions exploiting data-level parallelism?
- for (i=0;i<N;i++)
 - res+=a[i]*b[i]
- a in R0—R4
- b in R4—R7

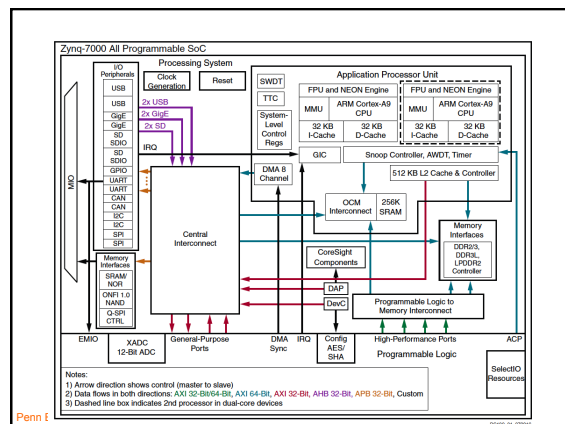


Pipelined Vector Units

- Will get both pipelining and parallel vector lanes
- Exploit data-level parallelism for both



Neon



Neon Vector

- 128b wide register file, 16 registers
- Support
 - 2x64b
 - 4x32b (also Single-Precision Float)
 - 8x16b
 - 16x8b

Penn ESE532 Spring 2017 -- DeHon

55

Sample Instructions

- VADD – basic vector
- VCEQ – compare equal
 - Sets to all 0s or 1s, useful for masking
- VMIN – avoid using if's
- VMLA – accumulating multiply
- VPADAL – maybe useful for reduce
- VEXT – for “shifting” vector alignment
- VLDn – deinterleaving load

Penn ESE532 Spring 2017 -- DeHon

56

Neon Notes

- Didn't see
 - Vector-wide reduce operation
 - Conditionals within vector lanes
- Do need to think about operations being pipelined within lanes

Penn ESE532 Spring 2017 -- DeHon

57

Big Ideas

- Data Parallelism easy basis for decomposition
- Data Parallel architectures can be compact – pack more computations onto a chip
 - SIMD, Pipelined
 - Benefit by sharing (instructions)
 - Performance can be brittle
 - Drop from peak as mismatch

Penn ESE532 Spring 2017 -- DeHon

58

Admin

- Reading for Day 5 on web
- Talk on Thursday by Ed Lee (UCB)
 - 3pm in Wu and Chen
- HW2 due Friday
- HW3 out (soon...)
 - Different partners

Penn ESE532 Spring 2017 -- DeHon

59