Master Informatics Eng.

2017/18 *A.J.Proença*

Data Parallelism 1 (vector & SIMD extensions) (most slides are borrowed)

AJProença, Advanced Architectures, MiEI, UMinho, 2017/18

Instruction and Data Streams

		Data Streams		
		Single	Multiple	
Instruction Single Streams		SISD: Intel Pentium 4	SIMD : SSE instructions of x86	
	Multiple	MISD : No examples today	MIMD: Intel Xeon e5345	

1



~~

Instruction and Data Streams

		Data Streams		
		Single	Multiple	
Instruction Streams	Single	SISD: Intel Pentium 4	SIMD : SSE instructions of x86	
	Multiple	MISD : No examples today	MIMD: Intel Xeon e5345	

- SPMD: Single Program Multiple Data
 - A parallel program on a MIMD computer
 - Conditional code for different processors

Chapter 7 — Multicores, Multiprocessors, and Clusters — 3

Introduction

Introduction

- SIMD architectures can exploit significant datalevel parallelism for:
 - matrix-oriented scientific computing
 - media-oriented <u>image</u> and <u>sound</u> processing
- SIMD is more energy efficient than MIMD
 - only needs to fetch one instruction per data operation
 - makes SIMD attractive for personal mobile devices
- SIMD allows programmers to continue to think sequentially





Vector Architectures

- Basic idea:
 - Read sets of data elements (<u>gather</u> from memory) into "vector registers"
 - Operate on those registers
 - Store/<u>scatter</u> the results back into memory
- Registers are controlled by the compiler
 - Used to hide memory latency
 - Leverage memory bandwidth



Vector Architectures





Can overlap execution of multiple vector instructions - Consider machine with 32 elements per vector register and 8 lanes:



Complete 24 operations/cycle while issuing 1 short instruction/cycle 8/19/2009 John Kubiatowicz Parallel Architecture: 35





AJProença, Advanced Architectures, MiEI, UMinho, 2017/18

VMIPS Instructions

- ADDVV.D: add two vectors
- ADDVS.D: add vector to a scalar
- LV/SV: vector load and vector store from address
- Example: DAXPY (Double-precision <u>A x X Plus Y</u>)

L.D	F0,a	;	load scalar a
LV	V1,Rx	;	load vector X
MULVS.D	V2,V1,F0	;	vector-scalar multiply
LV	V3,Ry	;	load vector Y
ADDVV	V4,V2,V3	;	add
SV	Ry,V4	;	store the result

 Requires the execution of 6 instructions versus almost 600 for MIPS (assuming DAXPY is operating on a vector with 64 elements)



9

Vector Architectures

11

Vector Architectures

Vector Execution Time

- Execution time depends on three factors:
 - Length of operand vectors
 - Structural hazards
 - Data dependencies
- VMIPS functional units consume one element per clock cycle
 - Execution time is approximately the vector length
- Convoy
 - Set of vector instructions that could potentially execute together in one unit of time, *chime*

Copyright © 2012, Elsevier Inc. All rights reserved.

Challenges

- Start up time
 - Latency of vector functional unit
 - Assume the same as Cray-1
 - Floating-point add => 6 clock cycles
 - Floating-point multiply => 7 clock cycles
 - Floating-point divide => 20 clock cycles
 - Vector load => 12 clock cycles
- Improvements:
 - > 1 element per clock cycle (1)
 - Non-64 wide vectors (2)
 - IF statements in vector code (3)
 - Memory system optimizations to support vector processors (4)
 - Multiple dimensional matrices (5)
 - Sparse matrices (6)
 - Programming a vector computer (7)

