## Computer Architecture: SIMD/Vector/GPU

Prof. Onur Mutlu (edited by seth)
Carnegie Mellon University

## Data Parallelism

- Concurrency arises from performing the same operations on different pieces of data
- Single instruction multiple data (SIMD)
- E.g., dot product of two vectors
- Contrast with data flow
- Concurrency arises from executing different operations in parallel (in a data driven manner)
- Contrast with thread ("control") parallelism
- Concurrency arises from executing different threads of control in parallel
- SIMD exploits instruction-level parallelism
- Multiple instructions concurrent: instructions happen to be the same


## Vector Processing: Exploiting Regular (Data) Parallelism

## SIMD Processing

- Single instruction operates on multiple data elements
- In time or in space
- Multiple processing elements
- Time-space duality
- Array processor: Instruction operates on multiple data elements at the same time
- Vector processor: Instruction operates on multiple data elements in consecutive time steps


## Array vs. Vector Processors

## ARRAY PROCESSOR

VECTOR PROCESSOR


Instruction Stream
LD $\quad \mathrm{VR} \leftarrow \mathrm{A}[3: 0]$ ADD VR $\leftarrow \mathrm{VR}, 1$ MUL VR $\leftarrow \mathrm{VR}, 2$ ST $\quad \mathrm{A}[3: 0] \leftarrow \mathrm{VR}$

Same op @ same time


SIMD Array Processing vs. VLIW

- VLIW



## Vector Processors

- A vector is a one-dimensional array of numbers
- Many scientific/commercial programs use vectors
for ( $\mathrm{i}=0 ; \mathrm{i}<=49 ; \mathrm{i}++$ )
$C[i]=(A[i]+B[i]) / 2$
- A vector processor is one whose instructions operate on vectors rather than scalar (single data) values
- Basic requirements
- Need to load/store vectors $\rightarrow$ vector registers (contain vectors)
- Need to operate on vectors of different lengths $\rightarrow$ vector length register (VLEN)
- Elements of a vector might be stored apart from each other in memory $\rightarrow$ vector stride register (VSTR)
- Stride: distance between two elements of a vector


## Vector Processors (II)

- A vector instruction performs an operation on each element in consecutive cycles
- Vector functional units are pipelined
- Each pipeline stage operates on a different data element
- Vector instructions allow deeper pipelines
- No intra-vector dependencies $\rightarrow$ no hardware interlocking within a vector
- No control flow within a vector
- Known stride allows prefetching of vectors into cache/memory


## Vector Processor Disadvantages

-- Works (only) if parallelism is regular (data/SIMD parallelism) ++ Vector operations
-- Very inefficient if parallelism is irregular
-- How about searching for a key in a linked list?
To program a vector machine, the compiler or hand coder must make the data structures in the code fit nearly exactly the regular structure built into the hardware. That's hard to do in first place, and just as hard to change. One tweak, and the low-level code has to be rewritten by a very smart and dedicated programmer who knows the hardware and often the subtleties of the application area. Often the rewriting is

## Vector Processor Advantages

```
+ No dependencies within a vector
```

- Pipelining, parallelization work well
- Can have very deep pipelines, no dependencies!
+ Each instruction generates a lot of work
- Reduces instruction fetch bandwidth
+ Highly regular memory access pattern
- Interleaving multiple banks for higher memory bandwidth
- Prefetching
+ No need to explicitly code loops
- Fewer branches in the instruction sequence


## Vector Processor Limitations

-- Memory (bandwidth) can easily become a bottleneck, especially if

1. compute/memory operation balance is not maintained
2. data is not mapped appropriately to memory banks

## Vector Registers

- Each vector data register holds N M-bit values
- Vector control registers: VLEN, VSTR, VMASK
- Vector Mask Register (VMASK)
- Indicates which elements of vector to operate on
- Set by vector test instructions
- e.g., VMASK[i] $=\left(V_{k}[i]==0\right)$
- Maximum VLEN can be N
- Maximum number of elements stored in a vector register



## Vector Functional Units

- Use deep pipeline (=> fast clock) to execute element operations
- Simplifies control of deep pipeline because elements in vector are independent



## Memory Banking

- Example: 16 banks; can start one bank access per cycle
- Bank latency: 11 cycles
- Can sustain 16 parallel accesses if they go to different banks

| $\begin{array}{\|l} \text { Bank } \\ 0 \end{array}$ | Bank $1$ | Bank $2$ |  |
| :---: | :---: | :---: | :---: |



## Vector Memory System



## Scalar Code Execution Time

- Scalar execution time on an in-order processor with 1 bank
- First two loads in the loop cannot be pipelined: 2*11 cycles
- $4+50 * 40=2004$ cycles
- Scalar execution time on an in-order processor with 16 banks (word-interleaved)
- First two loads in the loop can be pipelined
- $4+50 * 30=1504$ cycles
- Why 16 banks?
- 11 cycle memory access latency
- Having 16 ( $>11$ ) banks ensures there are enough banks to overlap enough memory operations to cover memory latency


## Scalar Code Example

- For I $=0$ to 49
- $C[i]=(A[i]+B[i]) / 2$
- Scalar code

| MOVI RO $=50$ | 1 |
| :---: | :---: |
| MOVA R1 $=$ A | 1304 dynamic instructions |
| MOVA R2 $=$ B | 1 |
| MOVA R3 $=$ C | 1 |
| LD R4 $=$ MEM[R1++] | 11 ;autoincrement addressing |
| LD R5 = MEM[R2++] | 11 |
| ADD R6 = R4 + R5 | 4 |
| SHFR R7 = R6 >> 1 | 1 |
| ST MEM[R3++] = R7 | 11 |
| DECBNZ --RO, X | 2 ;decrement and branch if NZ |

## Vectorizable Loops

- A loop is vectorizable if each iteration is independent of any other
- For I $=0$ to 49

$$
\square \mathrm{C}[\mathrm{i}]=(\mathrm{A}[\mathrm{i}]+\mathrm{B}[\mathrm{i}]) / 2 \quad 7 \text { dynamic instructions }
$$

- Vectorized loop:

| MOVI VLEN $=50$ | 1 |
| :--- | :--- |
| MOVI VSTR $=1$ | 1 |
| VLD V0 $=\mathrm{A}$ | $11+$ VLN -1 |
| VLD V1 $=\mathrm{B}$ | $11+$ VLN -1 |
| VADD V2 $=\mathrm{V} 0+\mathrm{V} 1$ | $4+V L N-1$ |
| VSHFR V3 $=\mathrm{V} 2 \gg 1$ | $1+V L N-1$ |
| VST $C=V 3$ | $11+V L N-1$ |

## Vector Code Performance

## - No chaining

- i.e., output of a vector functional unit cannot be used as the input of another (i.e., no vector data forwarding)
- One memory port (one address generator)
- 16 memory banks (word-interleaved)

- 285 cycles


## Vector Code Performance - Chaining

- Vector chaining: Data forwarding from one vector functional unit to another



## Vector Chaining

- Vector chaining: Data forwarding from one vector functional unit to another

- Chaining and 2 load ports, 1 store port in each bank

- 79 cycles



## Questions (I)

- What if \# data elements > \# elements in a vector register?
- Need to break loops so that each iteration operates on \# elements in a vector register
- E.g., 527 data elements, 64-element VREGs
- 8 iterations where VLEN $=64$
- 1 iteration where VLEN $=15$ (need to change value of VLEN)
- Called vector strip-mining
- What if vector data is not stored in a strided fashion in memory? (irregular memory access to a vector)
- Use indirection to combine elements into vector registers
- Called scatter/gather operations


## Gather/Scatter Operations

Want to vectorize loops with indirect accesses:

```
for (i=0; i<N; i++)
    A[i] = B[i] + C[D[i]]
```

Indexed load instruction (Gather)
LV vD, rD \# Load indices in D vector
LVI vC, rC, vD \# Load indirect from rC base
LV vB, rB \# Load B vector
ADDV.D vA,vB,vC \# Do add
SV vA, rA
\# Store result

## Gather/Scatter Operations

- Gather/scatter operations often implemented in hardware to handle sparse matrices
- Vector loads and stores use an index vector which is added to the base register to generate the addresses

| Index Vector | Data Vector | Equivalent |
| :---: | :---: | :---: |
|  |  |  |
| 1 | 3.14 | 3.14 |
| 3 | 6.5 | 0.0 |
| 7 | 71.2 | 6.5 |
| 8 | 2.71 | 0.0 |
|  |  | 0.0 |
|  |  | 0.0 |
|  |  | 0.0 |
|  |  | 71.2 |
|  |  | 2.7 |

## Conditional Operations in a Loop

- What if some operations should not be executed on a vector (based on a dynamically-determined condition)?
loop: $\quad$ if $(a[i]!=0)$ then $b[i]=a[i] * b[i]$
goto loop
- Idea: Masked operations
- VMASK register is a bit mask determining which data element should not be acted upon

VLD VO = A
VLD V1 = B
VMASK = (VO $!=0)$
VMUL V1 $=$ V0 $*$ V1
VST B = V1

- Does this look familiar? This is essentially predicated execution.


## Another Example with Masking

$$
\begin{aligned}
& \text { for }(i=0 ; i<64 ;++i) \\
& \quad \text { if }(a[i]>=b[i]) \text { then } c[i]=a[i] \\
& \text { else } c[i]=b[i]
\end{aligned}
$$

Steps to execute loop

1. Compare $A, B$ to get VMASK
2. Masked store of $A$ into $C$
3. Complement VMASK
4. Masked store of B into C

## Some Issues

## - Stride and banking

- As long as they are relatively prime to each other and there are enough banks to cover bank access latency, consecutive accesses proceed in parallel
- Storage of a matrix
- Row major: Consecutive elements in a row are laid out consecutively in memory
- Column major: Consecutive elements in a column are laid out consecutively in memory
- You need to change the stride when accessing a row versus column


## Masked Vector Instructions

Simple Implementation

- execute all N operations, turn off result writeback according to mask

| $\mathrm{M}[7]=1$ | $\mathrm{~A}[7]$ | $\mathrm{B}[7]$ |
| :--- | :--- | :--- |
| $\mathrm{M}[6]=0$ | $\mathrm{~A}[6]$ | $\mathrm{B}[6]$ |
| $\mathrm{M}[5]=1$ | $\mathrm{~A}[5]$ | $\mathrm{B}[5]$ |
| $\mathrm{M}[4]=1$ | $\mathrm{~A}[4]$ | $\mathrm{B}[4]$ |
| $\mathrm{M}[3]=0$ | $\mathrm{~A}[3]$ | $\mathrm{B}[3]$ |
| $\mathrm{M}[2]=0$ |  |  |

Density-Time Implementation

- scan mask vector and only execute elements with non-zero masks



## Array vs. Vector Processors, Revisited

- Array vs. vector processor distinction is a "purist's"
- They exploit data parallelism in both time and space
distinction
- Most "modern" SIMD processors are a combination of both

Instruction Stream
LD $\quad \mathrm{VR} \leftarrow \mathrm{A}[3: 0]$ ADD VR $\leftarrow \mathrm{VR}, 1$ MUL VR $\leftarrow \mathrm{VR}$, 2 ST $\quad \mathrm{A}[3: 0] \leftarrow \mathrm{VR}$

ARRAY PROCESSOR
PE0 PE1 PE2 PE3 AD MUD

Same op @ same time


Vector Instruction Execution



## Vector Instruction Level Parallelism

## Can overlap execution of multiple vector instructions

- example machine has 32 elements per vector register and 8 lanes
- Complete 24 operations/cycle while issuing 1 short instruction/cycle



## Vector/SIMD Processing Summary

- Vector/SIMD machines good at exploiting regular data-level parallelism
- Same operation performed on many data elements
- Improve performance, simplify design (no intra-vector dependencies)
- Performance improvement limited by vectorizability of code
- Scalar operations limit vector machine performance
- Amdahl's Law
- CRAY-1 was the fastest SCALAR machine at its time!
- Many existing ISAs include (vector-like) SIMD operations - Intel MMX/SSEn/AVX, PowerPC AltiVec, ARM Advanced SIMD


## Automatic Code Vectorization



## Intel Pentium MMX Operations

- Idea: One instruction operates on multiple data elements simultaneously
- Ala array processing (yet much more limited)
- Designed with multimedia (graphics) operations in mind


No VLEN register Opcode determines data type:
8 8-bit bytes
4 16-bit words
2 32-bit doublewords
1 64-bit quadword
Stride always equal to 1 .

Peleg and Weiser, "MMX Technology Extension to the Intel Architecture," IEEE Micro, 1996

## MMX Example: Image Overlaying (I)


figure 8. Chroma keying: image overlay using a background color.

figure 9. Generating the selection bit mask.

MMX Example: Image Overlaying (II)
PAND MM4, MM1 PANDN MM1, MM3





Figure 10 . Using the mask with logical $M M X$ instructions to perform a conditional select.

| Mova | mm3, mem1 | $P$ Load eight pixels from woman's image |
| :---: | :---: | :---: |
| Movq | mm4, mem2 | / Load eight pixels from the blossom image |
| Pampeqb | mm1, mm3 |  |
| Pand | mm4, mm1 |  |
| Pandn | mm1, mm3 |  |
| Por | mm4, mm1 |  |

[^0]
## High-Level View of a GPU



Loop Iterations as Threads


## Concept of "Thread Warps" and SIMT

- Warp: A set of threads that execute the same instruction (on different data elements) $\rightarrow$ SIMT (Nvidia-speak)
- All threads run the same kernel
- Warp: The threads that run lengthwise in a woven fabric ...



## SIMT Memory Access

- Same instruction in different threads uses thread id to index and access different data elements



## Sample GPU SIMT Code (Simplified)



## Latency Hiding with "Thread Warps"

- Warp: A set of threads that execute the same instruction (on different data elements)
- Fine-grained multithreading
- One instruction per thread in pipeline at a time (No branch prediction)
- Interleave warp execution to hide latencies
- Register values of all threads stay in register file
- No OS context switching
- Memory latency hiding

- Graphics has millions of pixels


## CPU Program

## ( float *a, float* $b$, float * $c$, int $N$ ) int index; <br> for (int $\mathrm{i}=0 ; \mathrm{i}<\mathrm{N} ;+\mathrm{i}$ ) <br> for (int j $=0 ; \mathrm{j}<\mathrm{N} ;++\mathrm{j}$ ) $\{$ index $=i+j^{*} N$; <br> $c[$ index] $=a[$ index] $+b[$ index]; \} <br> \} <br> int main () \{ <br> add matrix (a, b, c, N)

## GPU Program

## global add matrix

( float *a, float *b, float *c, int N) \{
int $\mathrm{i}=$ blockldx. x * blockDim. $\mathrm{x}+$ threadld $\mathrm{x} . \mathrm{x}$; Int $j=$ blockldx. $y$ * blockDim. $y+$ threadIdx. $y$ int index $=i+j^{*} N$;
if $(\mathrm{i}<\mathrm{N} \& \& \mathrm{j}<\mathrm{N})$
$c[$ index] $=a[$ index] $+b[$ index]

Int main() \{
dim3 dimBlock( blocksize, blocksize) ; dim3 dimGrid (N/dimBlock.x, N/dimBlock.y); add_matrix<<<dimGrid, dimBlock>>>( $a, b, c, N$ ):

## Warp-based SIMD vs. Traditional SIMD

- Traditional SIMD contains a single thread
- Lock step
- Programming model is SIMD (no threads) $\rightarrow$ SW needs to know vector length
- ISA contains vector/SIMD instructions
- Warp-based SIMD consists of multiple scalar threads executing in a SIMD manner (i.e., same instruction executed by all threads)
- Does not have to be lock step
- Each thread can be treated individually (i.e., placed in a different warp) $\rightarrow$ programming model not SIMD
- SW does not need to know vector length
- Enables memory and branch latency tolerance
- ISA is scalar $\rightarrow$ vector instructions formed dynamically
- Essentially, it is SPMD programming model implemented on SIMD hardware


## SPMD

- Single procedure/program, multiple data
- This is a programming model rather than computer organization
- Each processing element executes the same procedure, except on different data elements
- Procedures can synchronize at certain points in program, e.g. barriers
- Essentially, multiple instruction streams execute the same program
- Each program/procedure can 1) execute a different control-flow path, 2) work on different data, at run-time
- Many scientific applications programmed this way and run on MI MD computers (multiprocessors)
- Modern GPUs programmed in a similar way on a SIMD computer


## Control Flow Problem in GPUs/SIMD

- GPU uses SIMD pipeline to save area on control logic.
- Group scalar threads into warps
- Branch divergence occurs when threads inside warps branch to different execution
 paths.



## Dynamic Warp Formation

- Idea: Dynamically merge threads executing the same instruction (after branch divergence)
- Form new warp at divergence
- Enough threads branching to each path to create full new warps



## What About Memory Divergence?

- Modern GPUs have caches
- Ideally: Want all threads in the warp to hit (without conflicting with each other)
- Problem: One thread in a warp can stall the entire warp if it misses in the cache.
- Need techniques to
- Tolerate memory divergence
- Integrate solutions to branch and memory divergence


## Dynamic Warp Formation/Merging

- Idea: Dynamically merge threads executing the same instruction (after branch divergence)

- Fung et al., "Dynamic Warp Formation and Scheduling for Efficient GPU Control Flow," MI CRO 2007.



## NVIDIA GeForce GTX 285

- NVIDIA-speak:
- 240 stream processors
- "SIMT execution"
- Generic speak:
- 30 cores
- 8 SIMD functional units per core


NVIDIA GeForce GTX 285 "core"


64 KB of storage for fragment contexts (registers)
= SIMD functional unit, control shared across 8 units
= instruction stream decode
= execution context storage

NVIDIA GeForce GTX 285


There are 30 of these things on the GTX 285: 30,720 threads

lide credit: Kayvon Fatahalian

64 KB of storage for thread contexts (registers)

- Groups of 32 threads share instruction stream (each group is a Warp)
- Up to 32 warps are simultaneously interleaved
- Up to 1024 thread contexts can be stored


[^0]:    tional select.

