

CS 451 / 551 / ECE 541

ADVANCED
COMPUTER ARCHITECTURE

SESSION no. 20


University of Idaho

VECTOR PROCESSORS

• SIMD

"STREAM", MATRIX-VECTOR OPS
PROBLEM - EXPRESS IN TERMS OF VECTOR
OPERATIONS - "VECTORIZE"

• N-DIM MATRIX OPS - DECOMPOSE
INTO ~~2~~ SERIES OF 1-D OPS.



VECTOR REGISTER : HOLD VECTORS

University of Idaho

MIPS - 64 ELEMENTS x 64 BITS

. DOUBLE PRECISION F.P.

. LONG INT

. SMALLER FIELDS

. 2 x 32 BIT


. 4 x 16 BIT SHORT

. 8 x 8 BIT BYTES

VECTOR OPS

ONE INSTR \Rightarrow 64 COPIES OF THE OP

ADDV ADDR



University of Idaho

ADVANTAGES OF VECTOR OVER SCALAR?

1. FASTER

- MINIMIZE FETCH & DECODE
- OPTIMIZE HARDWARE FOR SPEED (RESTRICTED INSTRUCTION SET)

2. POWER SAVINGS

CONVY - STACK INSTRUCTIONS, INPUT TO PIPELINE,
 • INDEPENDENT (NO DEPENDENCIES)



ENHANCEMENTS - PERFORMANCE & FLEXIBILITY

1. MULTIPLE ELEMENTS / CLOCK CYCLE


• MULTIPLE EXECUTION UNITS

* 2. MULTIPLE VECTOR LENGTHS
ADD LENGTH REGISTER


3. IF STATEMENTS

for (x vector \rightarrow vector \times 64
vectors $R \times x$) f_i

$z = x + 3;$



mask 0 0 1 0 1 0 0 0



- 4. MEMORY BUW
- 5. MULTIPLE DIMENSIONS (MATRICES)
- 6. SPARSE MATRICES - INDEX VECTOR
- 7. PROGRAMMING

MULTIPLE LENGTH ALIGNED VECTORS

DATA VECTOR LENGTH REGISTER ALIGN REGISTER LENGTH

MVL = max vector length \forall $l \in \text{align}$


$\forall l = \text{vector length (odd size piece)}$

$n = \text{size of input data}$

ex $n = 16$

$n = 87$

$\forall l = 16 \cdot 5 + 7$




TO REPRESENT 87 WORDS

5 16-BIT VECTORS

1 7-BIT VECTOR

```
low = 0
VL = n % MVL - START WITH ODD (START)
SEGMENT = 7
for (j=0; j <= (n / MVL); j++)
  for (i=low; i < low + VL; i++) {
    Y[i] = a * X[i] + Y[i]
    low = low + VL;
    VL = MVL - USE LONGER SEGMENT
  }
  = 16 FOR REMAINDER
(5 TIMES)
```

IN HARDWARE

. VECTOR LENGTH REGISTER

. INITIALLY - SET TO 7

. ONLY DO 7 OPS.

. THEN SET TO 16


~~5. MULTIDIMENSIONAL ARRAYS~~

5. MULTIDIMENSIONAL ARRAYS

e.g. MATRIX MULTIPLY

```
for (i=0; i<100; i++)  
  for (j=0; j<100; j++)  
    A[i][j] = 0.0  
  for (k=0; k<100; k++)
```

$$A[i][j] = A[i][j] + B[i][k] * D[k][j]$$



"2D → 1D SCAN WITH STRIDE"

	col0	col1	col2	col3
row0	a ₀₀	a ₀₁	a ₀₂	a ₀₃
row1	a ₁₀	a ₁₁	a ₁₂	a ₁₃
row2	a ₂₀	a ₂₁	a ₂₂	a ₂₃
row3	a ₃₀	a ₃₁	a ₃₂	a ₃₃



ROW MAJOR ORDER

a ₀₀	a ₀₁	a ₀₂	a ₀₃
a ₁₀	a ₁₁	a ₁₂	a ₁₃
a ₂₀	a ₂₁	a ₂₂	a ₂₃
a ₃₀	a ₃₁	a ₃₂	a ₃₃

~~adjacent~~

row scan - adjacent
in memory

col scan
· multiply index by 4
4 = stride

NEED ABILITY TO LOAD VECTORS

- ONE WITH ADJACENT MEMORY LOCATIONS (ROW SCAN)
- ONE WITH "STRIDE" (COL SCAN)

LYWS $V_1, (R1, R2)$ STRIDE
↓ ↓ ↓
DEST. BASE ADDR
VECTOR

$$V_1 = \text{MEM}[R1], \text{MEM}[R1 + R2], \\ \text{MEM}[R1 + 2 \cdot R2]$$


GATHER NON CONTIGUOUS MEMORY
INTO CONTIGUOUS VECTOR
PUSH VECTORS THRU PIPELINE

University of Idaho

SYWS - INVERSE

"STORE VECTOR WITH STRIDE"

4. MAKE MEMORY BW MATCH PROCESSOR SPEED




University of Idaho

MEMORY BANKS

• CONTIGUOUS ADDRESS SPACE

• SEPARATE PHYSICAL MEMORIES



- READ OR WRITE VECTORS - EACH VECTOR ELEMENT FROM A DIFFERENT "MEMORY BANK." IN PARALLEL.
- AMORTIZE ACCESS LATENCY

EXAMPLE P. 277

Q: How many "banks"

Do I need to keep
processor busy?

CRAY T932

32 processors

4 loads, 2 stores / clock cycle

2.167 ns / clock cycle

15 ns mem cycle time

between memory accesses

32 proc. \times 6 $\frac{\text{accesses}}{\text{proc.}}$ = 192 accesses/
cycle

$$\frac{15}{2.167} = 7 \text{ cycles/access}$$

192 \times 7 = 1344 memory banks;
has 1024 memory banks