


Table of Contents

Load Instructions	1
Vector Arithmetic	2
Reorder Instructions	3
Reorder Instructions: Complex Numbers	4
Reorder Instructions: 4x4 Transpose	5
Load 4 reals	6
Store 4 reals	7
Load 4 complex	8
Matrix Multiplication, vector loads	9, 10
Matrix Multiplication, scalar loads	11, 12
Matrix Multiplication, Analysis	13
Formal Vectorization, Vector	14
$A \cdot x = b \rightarrow$ code	15
parallel	16, 17
WHT	18
The Bar Operator	19, 20
Short Vector FFT	21


Instructions: 1-7
 Basic building blocks:
 - loads & stores 8-12
 - shuffles 13-15
 - arithmetic 16-18

vector indexing:




in most instructions the order of operands matters $\frac{D}{0}$

Load Instructions (SSE and later)




$a = _mm_load_ps(m);$
 $a = p[0];$
 $a = _mm_loadu_ps(m);$

aligned, explicit
 - ~~4~~, implicit
 unaligned



$a = _mm_loadl_pi(a, m)$
 $b = _mm_loadh_pi(b, m)$







$a = _mm_load_ss(m)$


stores are analogous

Constants (Special load instructions) (SSE and later)

$c = _mm_set_ps(1.0, 2.0, 3.0, 4.0);$
 $d = _mm_set1_ps(1.0);$
 $e = _mm_set_ss(1.0);$
 $f = _mm_set_zero_ps();$



Vector arithmetic (SSE and later)



$$c = _mm_addps(a, b)$$


Same:

$$c = _mm_subps(a, b)$$

$$c = _mm_mulps(a, b)$$


⋮

Scalar arithmetic (SSE and later)



$$c = _mm_addss(a, b)$$

Add Sub (SSE3 and later)




$$c = _mm_addsubps(a, b)$$

(alternate add & sub of vector elements)

Reorders Instructions (SSE and later)

Unpack lo


$$c = _mm_unpacklo_ps(a, b)$$




$$\begin{aligned} c_0 &= b_0 \\ c_1 &= b_1 \\ c_2 &= a_0 \\ c_3 &= a_1 \end{aligned}$$

Unpack hi

$$c = _mm_unpackhi_ps(a, b)$$




$$\begin{aligned} c_0 &= a_2 \\ c_1 &= a_3 \\ c_2 &= b_2 \\ c_3 &= b_3 \end{aligned}$$



Shuffle

$$c = _mm_shuffle_ps(a, b, _MM_SHUFFLE(r, k, j, i))$$



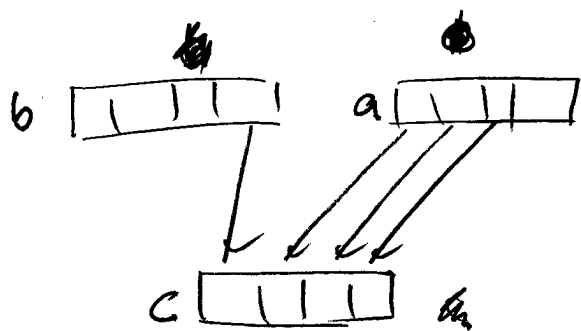
$$\begin{aligned} c_0 &= a_i \\ c_1 &= a_j \\ c_2 &= b_k \\ c_3 &= b_l \end{aligned}$$

any element of b any element of a

$i, j, k, l \in \{0, \dots, 3\}$
immediate

align (SSE3 and later)


5



"any consecutive 4 elements of the concatenation of a and b goes into c"

```
c = _mm_castps_si128 ( _mm_alignr_epi8 (
    _mm_castsi128_ps(a),
    _mm_castsi128_ps(b) ) );
```

shuffle (SSE3 and later)



can be given at runtime

"c is filled with any element of b or 0" (specified by a)


```
_mm_shuffle_epi8 ( )
```

shuffle (SSSE3 and later)

6

- wrong last time

- correct



- not the same as Altivec

Blend instruction

a | b | c | d

e | f | g | h

0

a | b | 0 | h

mask (int or vector)


"c is filled with elements of a or b at the same position on 0"

Matrix-Vector Product (SSE4 and later)

(7)


Dot-product instruction

- mm_dp-ps(a, b, mask) magic



"computes the pointwise product of a and b and writes an arbitrary sum of the resulting numbers into ~~all~~ selected elements of c (the others are set to 0)"

Example: load 4 real numbers from arbitrary memory locations (SSE)



```
# define SCALAR_LOAD (out, m0, m1, m2, m3)
{
  a = _mm_load_ss(m0);
  b = _mm_load_ss(m1);
  c = _mm_shuffle_ps(a, b, _MM_SHUFFLE(1, 0, 1, 0));
  d = _mm_load_ss(m2);
  e = _mm_load_ss(m3);
  f = _mm_shuffle_ps(d, e, _MM_SHUFFLE(1, 0, 1, 0));
  out = _mm_shuffle_ps(c, f, _MM_SHUFFLE(1, 0, 1, 0));
}
```

7 instructions

Note: - Whenever possible avoid this by ~~aligning~~ restructuring the algorithm or data format to have aligned vector loads (see page 1)

- ~~can~~ equivalent to macro on page 11 (but the above is "safer")

Other gather/scatter implementation (BAD) Don't do it like this (10)

float f[20] = { ... };

--declspec(align(16)) g[4];

--m128 vf;

g[0] = f[3];

g[1] = f[5];

g[2] = f[12];

g[3] = f[17];

vf = mm_loadps(g);
// operations


--mm_storeps(g, vf);

f[7] = g[0];

f[13] = g[1];

f[11] = g[2];

f[17] = g[3];



loads innocent, but is really bad
most people try that at some point
look at the assembly to see for yourself

Set instruction ~~type and usage~~ (SSE and later)

(1)

You can do: (see page 1)

```
_mm128 vf = _mm_set_ps(0.0, 3.0, 2.0, 1.0);
```

→ 1 vector load of 128 bit constant

Compiler lets you do this type of use

```
float f[20] = { ... };
```

```
_mm128 vf = _mm_set_ps(f[3], f[5], f[12], f[17]);
```

however internally: 4 loads, 3 shuffles


equivalent to page 8

Do not use `_mm_set_ps()` on variables if you can avoid it! (see page 8)

Example: load 4 complex numbers (load 4 pairs of numbers)


(SSE and /etc)

Strided access



6 instructions

Unit stride Access




4 instructions
(Benefit of consecutive data)

Same for store ops


Reorder Instructions - Examples (SSE and later)

Interleaved Complex \rightarrow split complex: L_2^8




```
c = _mm_shuffle_ps(a, b, _MM_SHUFFLE(2, 0, 2, 0));
d = _mm_shuffle_ps(a, b, _MM_SHUFFLE(3, 1, 3, 1));
```

Split Complex \rightarrow interleaved complex: L_4^8



```
c = _mm_unpacklo_ps(a, b);
d = _mm_unpackhi_ps(a, b);
```


Reverse Vector: J_4



```
b = _mm_shuffle_ps(a, a, _MM_SHUFFLE(0, 1, 2, 3));
```

~~Shift by 1~~

Shift by 1 (SSE)



2 instructions

$$t = \text{mm_shuffle_ps}(b, a, \text{MM_SHUFFLE}(3, 2, 1, 0));$$


$$c = \text{mm_shuffle_ps}(b, t, \text{MM_SHUFFLE}(1, 3, 2, 1));$$

SSE3: $\text{palign} \quad \text{mm_alignr_epi8}()$

1 instruction

~~$\text{psrnb} \quad \text{mm_shuffle_epi8}()$~~

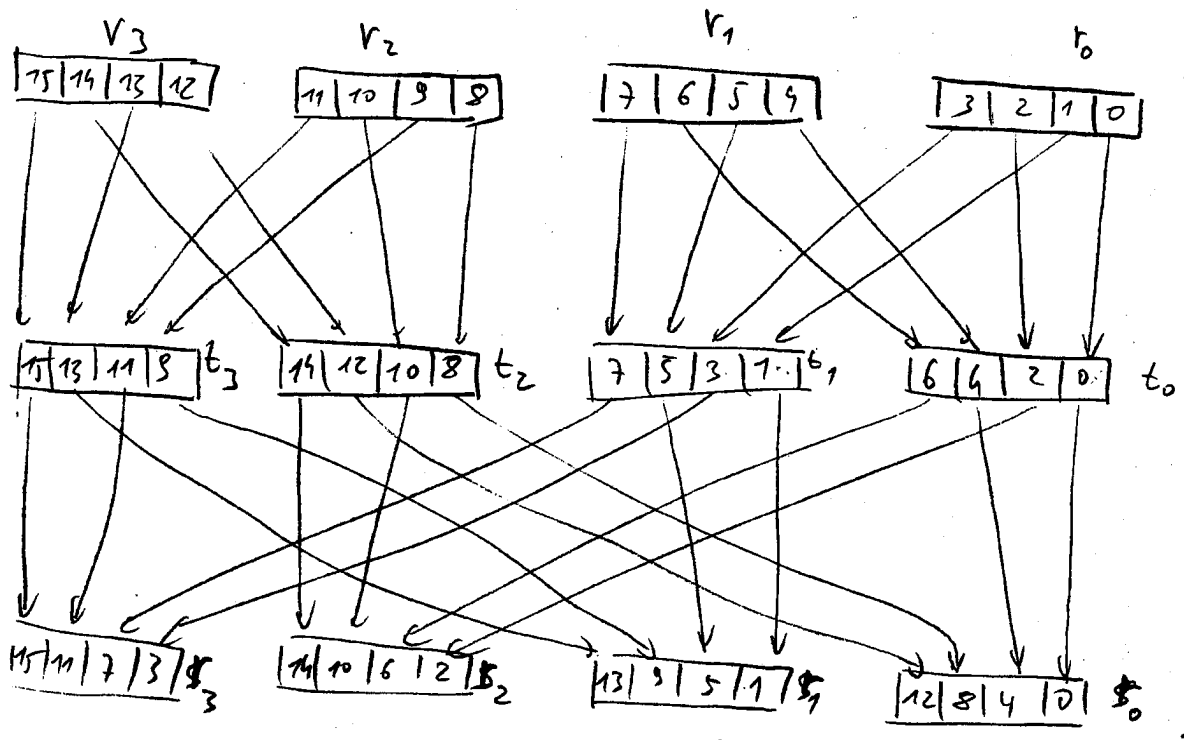
Reorder Instructions - Examples



may need fixity

vectors are rows

Transpose 4x4: $\frac{16}{4} = \frac{4 \oplus 4}{2} \oplus \frac{4 \oplus 4}{2} \oplus \frac{4 \oplus 4}{2} \oplus \frac{4 \oplus 4}{2}$




defined macro (SSE and later)

```
#define MM_TRANSPOSE4PS(v0, v1, v2, v3) {
    __m128 t0, t1, t2, t3;
    t0 = _mm_shuffle_ps(v0, v1, _MM_SHUFFLE(2, 0, 2, 0));
    t1 = _mm_shuffle_ps(v0, v1, _MM_SHUFFLE(3, 1, 3, 1));
    t2 = _mm_shuffle_ps(v2, v3, _MM_SHUFFLE(2, 0, 2, 0));
    t3 = _mm_shuffle_ps(v2, v3, _MM_SHUFFLE(3, 1, 3, 1));
    v0 = _mm_shuffle_ps(t0, t2, _MM_SHUFFLE(2, 0, 2, 0));
    v1 = _mm_shuffle_ps(t1, t3, _MM_SHUFFLE(2, 0, 2, 0));
    v2 = _mm_shuffle_ps(t0, t2, _MM_SHUFFLE(3, 1, 3, 1));
    v3 = _mm_shuffle_ps(t1, t3, _MM_SHUFFLE(3, 1, 3, 1));
}
```

3


MM_TRANSPOSE2_PD also exists (for doubles)

Matrix-Vector Product (SSE and later)



needs redrawing


↓



5 vector loads
4 vector multiplies
but how to do the sums?


transpose

MV



3 adds, 8 shuffles

other solutions (SSE3 and higher)



Matrix-Vector Product (SSE4)

(17)

needs
rednary

$$\begin{array}{l}
 \boxed{0|0|0|x} = \sum \boxed{\quad|\quad|\quad|\quad} \cdot \begin{array}{c} \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \end{array} \\
 \boxed{0|0|x|0} = \sum \boxed{\quad|\quad|\quad|\quad} \cdot \begin{array}{c} \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \end{array} \\
 \boxed{0|x|0|0} = \sum \boxed{\quad|\quad|\quad|\quad} \cdot \begin{array}{c} \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \end{array} \\
 \boxed{x|0|0|0} = \sum \boxed{\quad|\quad|\quad|\quad} \cdot \begin{array}{c} \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \end{array}
 \end{array}$$

$$\sum \boxed{x|x|x|x}$$

4 dot product instructions
3 add instructions

or 4 dot products

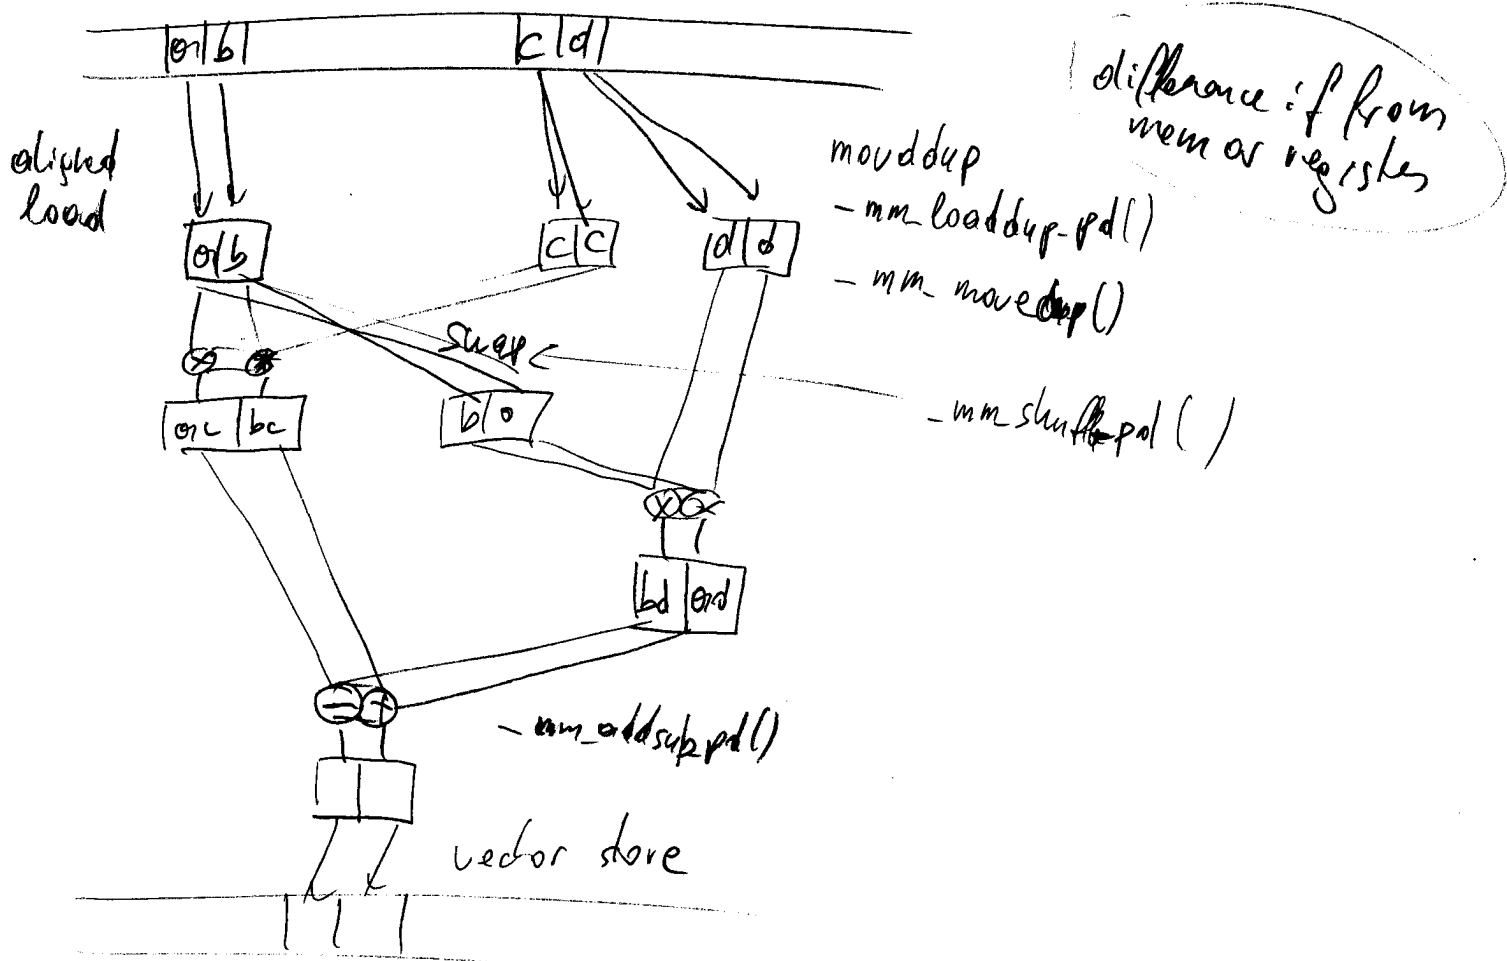
+ scalar store ops :

$$\boxed{\quad|\quad|\quad|x} = \sum \boxed{\quad|\quad|\quad|\quad} \cdot \begin{array}{c} \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \\ \boxed{\quad} \end{array}$$

performance implications?

Complex Multiplication double SSE3

$$(a+ib)(c+id) = (ac-bd) + i(ad+bc)$$



SSE, SSE2: see slides