

# vector instructions

# last time

## compiler limitations

- mostly one function at a time

- guessing at code size v. speed tradeoffs

- guessing at typically number of iterations/sizes

## aliasing (two names for one value)

- typically two pointers that may point to same variable/array

- prevents optimizations that reorder calculations

## inlining

- copy body of function in place of calls

- remove function call/argument passing/etc. overhead

- often allows additional optimizations

## removing redundant operations in loops

# quiz Q1 (1)

satSum:

```
movq %rdi, %rax
addq %rsi, %rax
cmpq %rsi, %rax
jb retMax
cmpq %rdi, %rax
jb retMax
ret
```

retMax:

```
movq $-1, %rax
ret
```

(avoid inlined) first movq:

compiler instead changes registers used in calling code

(keep inlined) jb

need something to use result of comparison

yes, could replace with cmov, but that's pretty similar

## quiz Q1 (2)

satSum:

```
    movq %rdi, %rax
    addq %rsi, %rax
    cmpq %rsi, %rax
    jb  retMax
    cmpq %rdi, %rax
    jb  retMax
    ret
```

retMax:

```
    movq $-1, %rax
    ret
```

(avoid inlined) first ret (or second ret)

(keep inlined) second movq

need something to store -1 somewhere

## quiz Q2 (aliasing)

```
for (int i = 0; i < N; i += 1) {  
    for (int j = 0; j < N; j += 1) {  
        A[i] += (A[i] - B[j]) * C[j];  
    }  
}
```

loop unrolling: okay, N, i, j can't have aliases

storing A[i] once per i-loop: not okay, updating A[i] could update B[j], C[j]

swapping i, j: not okay, updating A[i] later make different overlapping B[j], C[j] values be read

cache blocking: not okay, same reason

## quiz Q2 (aliasing, swapping)

suppose  $N = 2$ ,  $A = \{1, 1\}$ ,  $B = \{0, 0\}$ ,  $A = C$

```
A[i] += (A[i] - B[j]) * C[j];
```

```
/* for (i) for (j) order */
```

```
/* i = 0, j = 0 */ A[0] = 1 + (1 - 0) * 1 = 2
```

```
/* i = 0, j = 1 */ A[0] = 2 + (2 - 0) * 1 = 4
```

```
/* i = 0, j = 0 */ A[1] = 1 + (1 - 0) * 4 = 5
```

```
/* i = 1, j = 1 */ A[1] = 5 + (5 - 0) * 5 = 25
```

```
/* versus */
```

```
/* i = 0, j = 0 */ A[0] = 1 + (1 - 0) * 1 = 2
```

```
/* i = 1, j = 0 */ A[1] = 1 + (1 - 0) * 2 = 3
```

```
/* i = 0, j = 1 */ A[0] = 2 + (2 - 0) * 3 = 8
```

```
/* i = 1, j = 1 */ A[1] = 3 + (3 - 0) * 3 = 9
```

## quiz Q3 (small iter count)

[bad for perf] unrolling the inner most loop

special case handling or non-multiples of unrollig count

[not bad] loading and storing  $A[i]$  once per i-loop iteration

[not bad] swapping loops

[bad for perf] cache blocking to improve cache locality

extra loops; possibly special case for non-mutliples of block size

## quiz Q4 (machine code)

[more code] unrolling the inner most loop  
more copies of loop body + special cases

[not more code] loading and storing  $A[i]$  once per i-loop iteration  
probbaly same # of instructions  
maybe one or tow more instructions?  
(add X, (Y)  $\rightarrow$  mov (Y), Z; ... add X, Z ...; mov Z, (Y))

[not more code] swapping loops

[more code] cache blocking to improve cache locality  
extra loops; possibly special case for non-mutliples of block size



# looplab speeds on my desktop

original assembly: 2.0 cycles/element

unrolled x2: 1.0 cycles element

unrolled x4: 1.0 cycles element

unrolled x8: 1.0 cycles element

unrolled x8, 4 accumulators: 0.5 cycles element

# looplab speeds on my desktop

original assembly: 2.0 cycles/element

unrolled x2: 1.0 cycles element

unrolled x4: 1.0 cycles element

unrolled x8: 1.0 cycles element

unrolled x8, 4 accumulators: 0.5 cycles element

Clang 6 optimized code: 0.13 cycles/element

GCC optimized code: 0.14 cycles/element

# looplab speeds on my desktop

original assembly: 2.0 cycles/element

unrolled x2: 1.0 cycles element

unrolled x4: 1.0 cycles element

unrolled x8: 1.0 cycles element

unrolled x8, 4 accumulators: 0.5 cycles element

Clang 6 optimized code: 0.13 cycles/element

GCC optimized code: 0.14 cycles/element

how? instructions that add *16 pairs of shorts* at once!

“vector” or “SIMD” (single instruction multiple data) instruction

## unvectorized add (original)

```
unsigned int A[512], B[512];  
...  
for (int i = 0; i < N; i += 1) {  
    A[i] = A[i] + B[i];  
}
```

# unvectorized add (unrolled)

```
unsigned int A[512], B[512];  
...  
for (int i = 0; i < 512; i += 8) {  
    A[i+0] = A[i+0] + B[i+0];  
    A[i+1] = A[i+1] + B[i+1];  
    A[i+2] = A[i+2] + B[i+2];  
    A[i+3] = A[i+3] + B[i+3];  
    A[i+4] = A[i+4] + B[i+4];  
    A[i+5] = A[i+5] + B[i+5];  
    A[i+6] = A[i+6] + B[i+6];  
    A[i+7] = A[i+7] + B[i+7];  
}
```

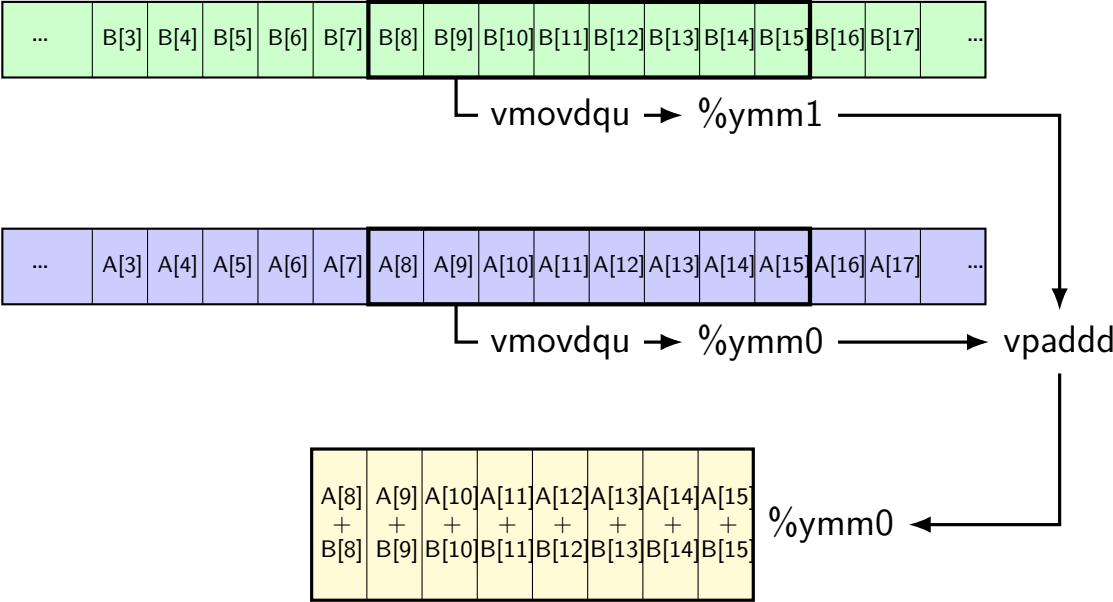
goal: use SIMD add instruction to do all 8 adds above  
SIMD = single instruction, multiple data

# desired assembly

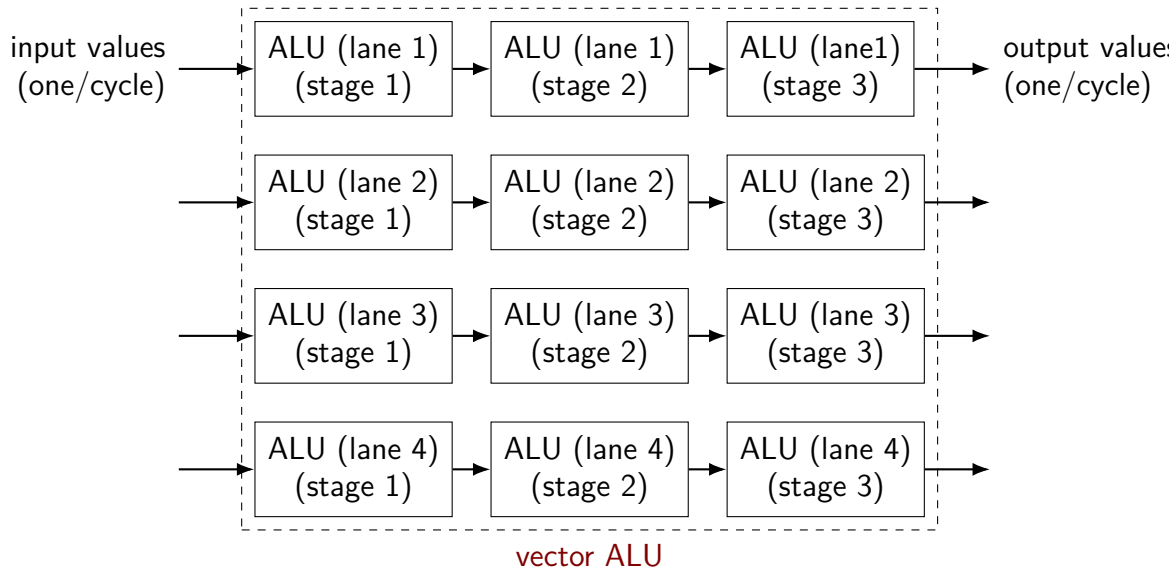
```
xor %rax, %rax
the_loop:
vmovdqu A(%rax), %ymm0
vmovdqu B(%rax), %ymm1
vpadd %ymm1, %ymm0, %ymm0
vmovdqu %ymm0, A(%rax)
addq $32, %rax
cmpq $2048, %rax
jne the_loop
```

```
/* load 256 bits of A into ymm0 */
/* load 256 bits of B into ymm1 */
/* ymm1 + ymm0 -> ymm0 */
/* store ymm0 into A */
/* increment index by 32 bytes */
/* offset < 2048 (= 512 * 4) bytes */
```

# vector add picture



# one view of vector functional units





# why vector instructions?

lots of logic not dedicated to computation

- instruction queue

- reorder buffer

- instruction fetch

- branch prediction

- ...

adding vector instructions — little extra control logic

...but a lot more computational capacity

# vector instructions and compilers

compilers can sometimes figure out how to use vector instructions  
(and have gotten much, much better at it over the past decade)

but easily messed up:

- by aliasing

- by conditionals

- by some operation with no vector instruction

- ...

very non-intuitive for me when compiler will/will not use vector instructions

## vector intrinsics

if compiler doesn't work...

could write vector instruction assembly by hand

second option: “intrinsic functions”

C functions that compile to particular instructions

## vector intrinsics: add example

```
int A[512], B[512];
```

```
for (int i = 0; i < 512; i += 8) {
```

```
    // "si256" --> 256 bit integer
```

```
    // a_values = {A[i], A[i+1], ..., A[i+7]} (8 x 32 bits)
```

```
    __m256i a_values = _mm256_loadu_si256((__m256i*) &A[i]);
```

```
    // b_values = {B[i], B[i+1] ..., A[i+7]} (8 x 32 bits)
```

```
    __m256i b_values = _mm256_loadu_si256((__m256i*) &B[i]);
```

```
    // add eight 32-bit integers
```

```
    // sums = {A[i] + B[i], A[i+1] + B[i+1], ....., A[i+7] + B[i+7]}
```

```
    __m256i sums = _mm256_add_epi32(a_values, b_values);
```

```
    // {A[i], A[i+1], A[i+2], A[i+3], ..., A[i+7]} = sums
```

```
    _mm256_storeu_si256((__m256i*) &A[i], sums);
```

```
}
```

## vector intrinsics: add example

```
int A[512]
```

special type `__m256i` — “256 bits of integers”  
other types: `__m256` (floats), `__m128d` (doubles)

```
for (int i = 0; i < 512; i += 8) {  
    // "si256" --> 256 bit integer  
    // a_values = {A[i], A[i+1], ..., A[i+7]} (8 x 32 bits)  
    __m256i a_values = _mm256_loadu_si256((__m256i*) &A[i]);  
    // b_values = {B[i], B[i+1] ..., A[i+7]} (8 x 32 bits)  
    __m256i b_values = _mm256_loadu_si256((__m256i*) &B[i]);  
  
    // add eight 32-bit integers  
    // sums = {A[i] + B[i], A[i+1] + B[i+1], ....., A[i+7] + B[i+7]}  
    __m256i sums = _mm256_add_epi32(a_values, b_values);  
  
    // {A[i], A[i+1], A[i+2], A[i+3], ..., A[i+7]} = sums  
    _mm256_storeu_si256((__m256i*) &A[i], sums);  
}
```

## vector intrinsics: add example

functions to store/load

si256 means "256-bit integer value"

u for "unaligned" (otherwise, pointer address must be multiple of 32)

```
// "si256" --> 256 bit integer
```

```
// a_values = {A[i], A[i+1], ..., A[i+7]} (8 x 32 bits)
```

```
__m256i a_values = _mm256_loadu_si256((__m256i*) &A[i]);
```

```
// b_values = {B[i], B[i+1] ..., A[i+7]} (8 x 32 bits)
```

```
__m256i b_values = _mm256_loadu_si256((__m256i*) &B[i]);
```

```
// add eight 32-bit integers
```

```
// sums = {A[i] + B[i], A[i+1] + B[i+1], ....., A[i+7] + B[i+7]}
```

```
__m256i sums = _mm256_add_epi32(a_values, b_values);
```

```
// {A[i], A[i+1], A[i+2], A[i+3], ..., A[i+7]} = sums
```

```
_mm256_storeu_si256((__m256i*) &A[i], sums);
```

```
}
```

## vector intrinsics: add example

```
int A[512], B[512];
```

```
for (int i = 0; i < 512; i += 8) {
```

```
    // "si256" -- function to add (8 x 32 bits)  
    // a_values = epi32 means "8 32-bit integers" (8 x 32 bits)  
    __m256i a_values = _mm256_loadu_si256((__m256i*) &A[i]);
```

```
    // b_values = {B[i], B[i+1] ..., A[i+7]} (8 x 32 bits)  
    __m256i b_values = _mm256_loadu_si256((__m256i*) &B[i]);
```

```
    // add eight 32-bit integers
```

```
    // sums = {A[i] + B[i], A[i+1] + B[i+1], ..., A[i+7] + B[i+7]}
```

```
    __m256i sums = _mm256_add_epi32(a_values, b_values);
```

```
    // {A[i], A[i+1], A[i+2], A[i+3], ..., A[i+7]} = sums
```

```
    _mm256_storeu_si256((__m256i*) &A[i], sums);
```

```
}
```

## vector intrinsics: different size

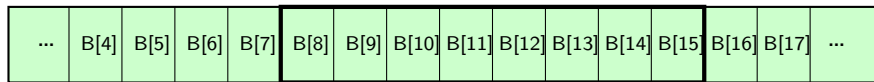
```
long A[512], B[512]; /* instead of int */
...
for (int i = 0; i < 512; i += 4) {
    // a_values = {A[i], A[i+1], A[i+2], A[i+3]} (4 x 64 bits)
    __m256i a_values = _mm256_loadu_si256((__m256i*) &A[i]);
    // b_values = {B[i], B[i+1], B[i+2], B[i+3]} (4 x 64 bits)
    __m256i b_values = _mm256_loadu_si256((__m256i*) &B[i]);
    // add four 64-bit integers: vpaddq %ymm0, %ymm1
    // sums = {A[i] + B[i], A[i+1] + B[i+1], ...}
    __m256i sums = _mm256_add_epi64(a_values, b_values);
    // {A[i], A[i+1], A[i+2], A[i+3]} = sums
    _mm256_storeu_si256((__m256i*) &A[i], sums);
}
```



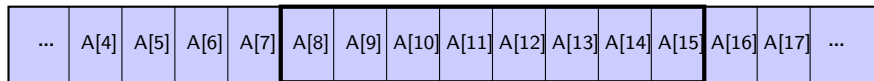
## vector intrinsics: different size

```
long A[512], B[512]; /* instead of int */
...
for (int i = 0; i < 512; i += 4) {
    // a_values = {A[i], A[i+1], A[i+2], A[i+3]} (4 x 64 bits)
    __m256i a_values = _mm256_loadu_si256((__m256i*) &A[i]);
    // b_values = {B[i], B[i+1], B[i+2], B[i+3]} (4 x 64 bits)
    __m256i b_values = _mm256_loadu_si256((__m256i*) &B[i]);
    // add four 64-bit integers: vpaddq %ymm0, %ymm1
    // sums = {A[i] + B[i], A[i+1] + B[i+1], ...}
    __m256i sums = _mm256_add_epi64(a_values, b_values);
    // {A[i], A[i+1], A[i+2], A[i+3]} = sums
    _mm256_storeu_si256((__m256i*) &A[i], sums);
}
```

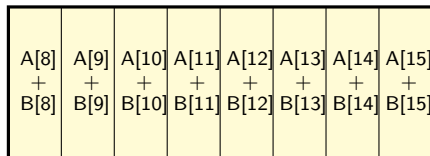
# vector add picture (intrinsics)



`_mm256_loadu_si256`  
(asm: vmovdqu) → `b_values`  
(%ymm1?)

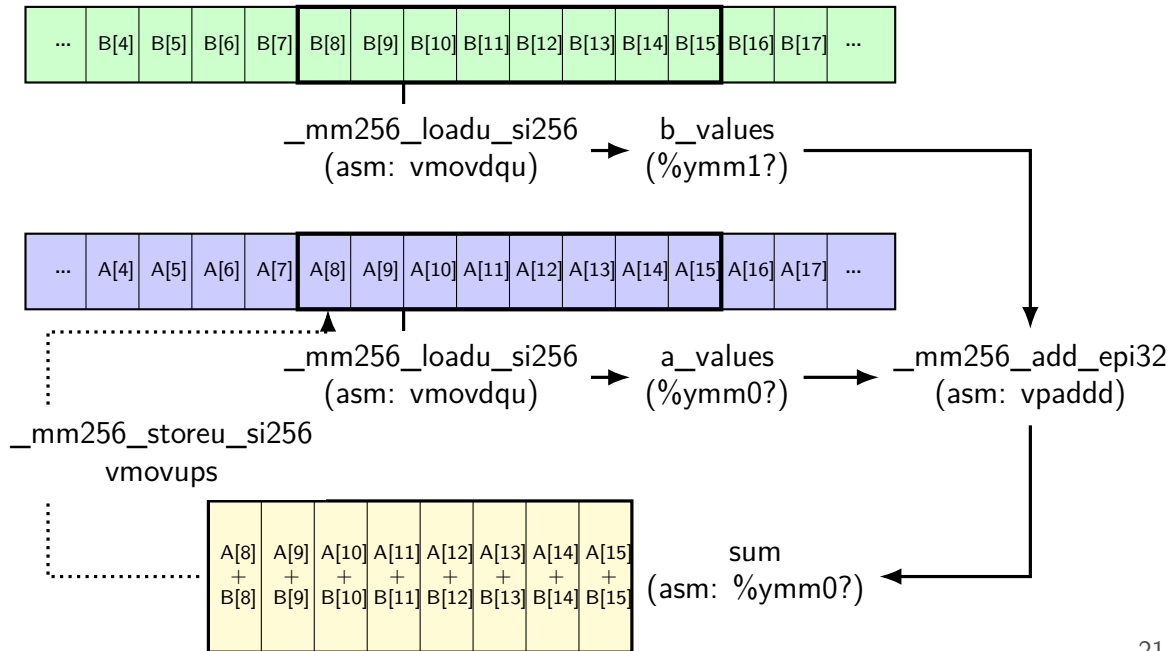


`_mm256_loadu_si256`  
(asm: vmovdqu) → `a_values`  
(%ymm0?) → `_mm256_add_epi32`  
(asm: vpaddd)



`sum`  
(asm: %ymm0?)

# vector add picture (intrinsics)



## exercise

```
long foo[8] = {1,1,2,2,3,3,4,4};
long bar[8] = {2,2,2,3,3,3,4,4};
__mm256i foo0_as_vector = _mm256_loadu_si256((__m256i*)&foo[0])
__mm256i foo4_as_vector = _mm256_loadu_si256((__m256i*)&foo[4])
__mm256i bar0_as_vector = _mm256_loadu_si256((__m256i*)&bar[0])

__mm256i result = _mm256_add_epi64(foo0_as_vector, foo4_as_vector);
result = _mm256_mullo_epi64(result, bar0_as_vector);
_mm256_storeu_si256((__m256i*) &bar[4], result);
```

Final value of bar array?

- A. {2,2,2,3,12,12,24,24}
- B. {2,2,2,3,15,15,28,28}
- C. {2,2,2,3,10,10,20,20}
- D. {12,12,24,24,3,3,4,4}
- E. {14,14,26,27,3,3,4,4}
- F. {14,14,26,27,12,12,24,24}
- G. something else

# matrix multiply

```
void matmul(unsigned int *A, unsigned int *B, unsigned int *C)
{
    for (int k = 0; k < N; ++k)
        for (int i = 0; i < N; ++i)
            for (int j = 0; j < N; ++j)
                C[i * N + j] += A[i * N + k] * B[k * N + j];
}
```

(simple version, no cache blocking, no avoiding aliasing between C, B, A,...)

# matmul unrolled

```
void matmul(unsigned int *A, unsigned int *B, unsigned int *C) {
    for (int k = 0; k < N; ++k) {
        for (int i = 0; i < N; ++i)
            for (int j = 0; j < N; j += 8) {
                /* goal: vectorize this */
                C[i * N + j + 0] += A[i * N + k] * B[k * N + j + 0];
                C[i * N + j + 1] += A[i * N + k] * B[k * N + j + 1];
                C[i * N + j + 2] += A[i * N + k] * B[k * N + j + 2];
                C[i * N + j + 3] += A[i * N + k] * B[k * N + j + 3];
                C[i * N + j + 4] += A[i * N + k] * B[k * N + j + 4];
                C[i * N + j + 5] += A[i * N + k] * B[k * N + j + 5];
                C[i * N + j + 6] += A[i * N + k] * B[k * N + j + 6];
                C[i * N + j + 7] += A[i * N + k] * B[k * N + j + 7];
            }
    }
}
```

(NB: would probably also want to do cache blocking...)

## handy intrinsic functions for matmul

`_mm256_set1_epi32` — load eight copies of a 32-bit value into a 256-bit value

instructions generated vary; one example: `vmovd + vpbroadcastd`

`_mm256_mullo_epi32` — multiply eight pairs of 32-bit values, give lowest 32-bits of results

generates `vpmulld`

# vectorizing matmul

*/\* goal: vectorize this \*/*

`C[i * N + j + 0] += A[i * N + k] * B[k * N + j + 0];`

`C[i * N + j + 1] += A[i * N + k] * B[k * N + j + 1];`

`...`

`C[i * N + j + 6] += A[i * N + k] * B[k * N + j + 6];`

`C[i * N + j + 7] += A[i * N + k] * B[k * N + j + 7];`

---



# vectorizing matmul

```
/* goal: vectorize this */
```

```
C[i * N + j + 0] += A[i * N + k] * B[k * N + j + 0];
```

```
C[i * N + j + 1] += A[i * N + k] * B[k * N + j + 1];
```

```
...
```

```
C[i * N + j + 6] += A[i * N + k] * B[k * N + j + 6];
```

```
C[i * N + j + 7] += A[i * N + k] * B[k * N + j + 7];
```

---

```
// load eight elements from C
```

```
Cij = _mm256_loadu_si256((__m256i*) &C[i * N + j + 0]);
```

```
... // manipulate vector here
```

```
// store eight elements into C
```

```
_mm_storeu_si256((__m256i*) &C[i * N + j + 0], Cij);
```

# vectorizing matmul

```
/* goal: vectorize this */  
C[i * N + j + 0] += A[i * N + k] * B[k * N + j + 0];  
C[i * N + j + 1] += A[i * N + k] * B[k * N + j + 1];  
...  
C[i * N + j + 6] += A[i * N + k] * B[k * N + j + 6];  
C[i * N + j + 7] += A[i * N + k] * B[k * N + j + 7];
```

---

```
// load eight elements from B  
Bkj = _mm256_loadu_si256((__m256i*) &B[k * N + j + 0]);  
... // multiply each by B[i * N + k] here
```

# vectorizing matmul

*/\* goal: vectorize this \*/*

$C[i * N + j + 0] += A[i * N + k] * B[k * N + j + 0];$

$C[i * N + j + 1] += A[i * N + k] * B[k * N + j + 1];$

...

$C[i * N + j + 6] += A[i * N + k] * B[k * N + j + 6];$

$C[i * N + j + 7] += A[i * N + k] * B[k * N + j + 7];$

---

*// load eight elements starting with B[k \* n + j]*

$Bkj = \_mm256\_loadu\_si256((\_m256i*) \&B[k * N + j + 0]);$

*// load eight copies of A[i \* N + k]*

$Aik = \_mm256\_set1\_epi32(A[i * N + k]);$

*// multiply each pair*

$multiply\_results = \_mm256\_mullo\_epi32(Aik, Bkj);$

# vectorizing matmul

```
/* goal: vectorize this */  
C[i * N + j + 0] += A[i * N + k] * B[k * N + j + 0];  
C[i * N + j + 1] += A[i * N + k] * B[k * N + j + 1];  
...  
C[i * N + j + 6] += A[i * N + k] * B[k * N + j + 6];  
C[i * N + j + 7] += A[i * N + k] * B[k * N + j + 7];
```

---

```
Cij = _mm256_add_epi32(Cij, multiply_results);  
// store back results  
_mm256_storeu_si256(..., Cij);
```

# matmul vectorized

```
__m256i Cij, Bkj, Aik, multiply_results;
```

```
// Cij = {Ci,j, Ci,j+1, Ci,j+2, ..., Ci,j+7}
```

```
Cij = _mm256_loadu_si256((__m256i*) &C[i * N + j]);
```

```
// Bkj = {Bk,j, Bk,j+1, Bk,j+2, ..., Bk,j+7}
```

```
Bkj = _mm256_loadu_si256((__m256i*) &B[k * N + j]);
```

```
// Aik = {Ai,k, Ai,k, ..., Ai,k}
```

```
Aik = _mm256_set1_epi32(A[i * N + k]);
```

```
// Aik_times_Bkj = {Ai,k × Bk,j, Ai,k × Bk,j+1, Ai,k × Bk,j+2, ..., Ai,k × Bk,j+7}
```

```
multiply_results = _mm256_mullo_epi32(Aij, Bkj);
```

```
// Cij = {Ci,j + Ai,k × Bk,j, Ci,j+1 + Ai,k × Bk,j+1, ...}
```

```
Cij = _mm256_add_epi32(Cij, multiply_results);
```

```
// store Cij into C
```

```
_mm256_storeu_si256((__m256i*) &C[i * N + j], Cij);
```

## vector exercise (2)

```
long A[1024], B[1024];  
...  
for (int i = 0; i < 1024; i += 1)  
    for (int j = 0; j < 1024; j += 1)  
        A[i] += B[i] * B[j];
```

(casts omitted below to reduce clutter:)

```
for (int i = 0; i < 1024; i += 4) {  
    A_part = _mm256_loadu_si256(&A[i]);  
    Bi_part = _mm256_loadu_si256(&B[i]);  
    for (int j = 0; j < 1024; /* BLANK 1 */) {  
        Bj_part = _mm256_/* BLANK 2 */;  
        A_part = _mm256_add_epi64(A_part,  
            _mm256_mullo_epi64(Bi_part, Bj_part));  
    }  
    _mm256_storeu_si256(&A[i], A_part);  
}
```

What goes in BLANK 1 and BLANK 2?

- A. `j += 1, loadu_si256(&B[j])`    B. `j += 4, loadu_si256(&B[j])`  
C. `j += 1, set1_epi64(B[j])`        D. `j += 4, set1_epi64(B[j])`

## vector exercise 2 explanation

```
for (int i = 0; i < 1024; i += 1)
    for (int j = 0; j < 1024; j += 1)
        A[i] += B[i] * B[j];
```

*/\* -- transformed into -- \*/*

```
for (int i = 0; i < 1024; i += 4)
    for (int j = 0; j < 1024; j += 1) {
        A[i+0] += B[i+0] * B[j];
        A[i+1] += B[i+1] * B[j];
        A[i+2] += B[i+2] * B[j];
        A[i+3] += B[i+3] * B[j];
    }
```

*/\* not the much harder to vectorize: \*/*

```
for (int i = 0; i < 1024; i += 1)
    for (int j = 0; j < 1024; j += 4) {
        A[i] += B[i] * B[j+0];
        A[i] += B[i] * B[j+1];
        A[i] += B[i] * B[j+2];
        A[i] += B[i] * B[j+3];
    }
```

# moving values in vectors?

sometimes values aren't in the right place in vector

example:

have: [1, 2, 3, 4]

want: [3, 4, 1, 2]

there are instructions/intrinsics for doing this  
called shuffling/swizzling/permute/...

sometimes might need combination of them

worst-case: could rearrange on stack..., I guess



# example shuffling operation (1)

goal: [1, 2, 3, 4] to [3, 4, 1, 2] (64-bit values)

```
/* x = {1, 2, 3, 4} */  
__m256i x = _mm256_setr_epi64x(1, 2, 3, 4);  
__m256i result = _mm256_permute4x64_epi64(  
    x,  
    /* index 2, then 3, then 0, then 1 */  
    2 | (3 << 2) | (0 << 4) | (1 << 6)  
    /* could also write _MM_SHUFFLE(1, 0, 3, 2) */  
);  
/* result = {3, 4, 1, 2} */
```

# other vector instructions

multiple extensions to the X86 instruction set for vector instructions

early versions (128-bit vectors): SSE, SSE2, SSE3, SSSE3, SSE4.1, SSE4.2

128-bit vectors

this class (256-bit): AVX, AVX2

not this class (512+-bit): AVX-512

512-bit vectors

also other ISAs have these: e.g. NEON on ARM, MSA on MIPS, AltiVec/VMX on POWER, ...

GPUs are essentially vector-instruction-specialized CPUs

## other vector interfaces

intrinsics (our assignments) one way

some alternate programming interfaces

have compiler do more work than intrinsics

e.g. CUDA, OpenCL, GCC's vector instructions

# other vector instructions features

more flexible vector instruction features:

- invented in the 1990s

- often present in GPUs and being rediscovered by modern ISAs

reasonable conditional handling

better variable-length vectors

ability to load/store non-contiguous values

some of these features in AVX2/AVX512

# alternate vector interfaces

intrinsics functions/assembly aren't the only way to write vector code

e.g. GCC vector extensions: more like normal C code

types for each kind of vector

write + instead of `_mm_add_epi32`

e.g. CUDA (GPUs): looks like writing multithreaded code, but each thread is vector “lane”

# optimizing real programs

ask your compiler to try first

spend effort where **it matters**

e.g. 90% of program time spent reading files, but optimize computation?

e.g. 90% of program time spent in routine A, but optimize B?

# profilers

first step — tool to determine where you spend time

tools exist to do this for programs

example on Linux: `perf`

# example

```
Samples: 37K of event 'cycles', Event count (approx.): 3736755513
```

Children	Self	Command	Shared Object	Symbol
+ 100.00%	0.00%	hclrs-with-debu	hclrs-with-debuginfo	[.] _start
+ 100.00%	0.00%	hclrs-with-debu	libc-2.31.so	[.] __libc_start_main
+ 100.00%	0.00%	hclrs-with-debu	hclrs-with-debuginfo	[.] main
+ 100.00%	0.00%	hclrs-with-debu	hclrs-with-debuginfo	[.] std::sys_common::backtrace::__rust_begin_short_backt
+ 100.00%	0.00%	hclrs-with-debu	hclrs-with-debuginfo	[.] hclrs::main
+ 99.99%	9.75%	hclrs-with-debu	hclrs-with-debuginfo	[.] hclrs::program::RunningProgram::run
+ 60.37%	31.67%	hclrs-with-debu	hclrs-with-debuginfo	[.] hclrs::ast::SpannedExpr::evaluate
+ 41.34%	23.29%	hclrs-with-debu	hclrs-with-debuginfo	[.] hashbrown::map::make_hash
+ 18.08%	18.07%	hclrs-with-debu	hclrs-with-debuginfo	[.] <std::collections::hash::map::DefaultHasher as core:
+ 16.33%	0.68%	hclrs-with-debu	hclrs-with-debuginfo	[.] hclrs::program::Program::process_register_banks
+ 9.54%	3.15%	hclrs-with-debu	hclrs-with-debuginfo	[.] std::collections::hash::map::HashMap<K,V,S>::get
+ 9.10%	9.09%	hclrs-with-debu	libc-2.31.so	[.] __memcmp_avx2_movbe
+ 6.11%	2.10%	hclrs-with-debu	hclrs-with-debuginfo	[.] hashbrown::map::HashMap<K,V,S>::get_mut
+ 2.32%	0.88%	hclrs-with-debu	hclrs-with-debuginfo	[.] std::collections::hash::map::HashMap<K,V,S>::get
+ 1.45%	0.52%	hclrs-with-debu	hclrs-with-debuginfo	[.] hashbrown::map::HashMap<K,V,S>::insert
0.37%	0.11%	hclrs-with-debu	hclrs-with-debuginfo	[.] <alloc::string::String as core::clone::Clone>::clone
0.19%	0.19%	hclrs-with-debu	libc-2.31.so	[.] malloc



## an infinite loop

```
int main(void) {  
    while (1) {  
        /* waste CPU time */  
    }  
}
```

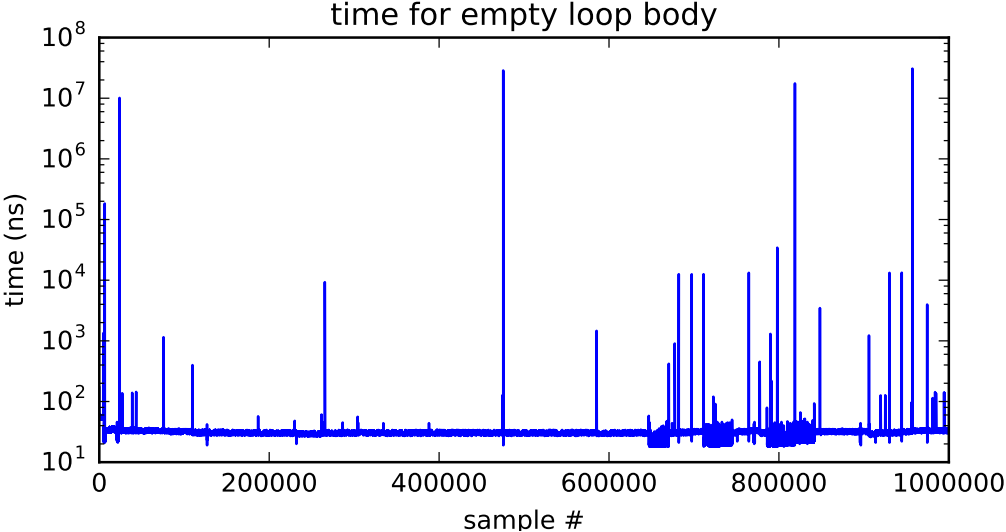
If I run this on a shared department machine, can you still use it?  
...if the machine only has one core?

## timing nothing

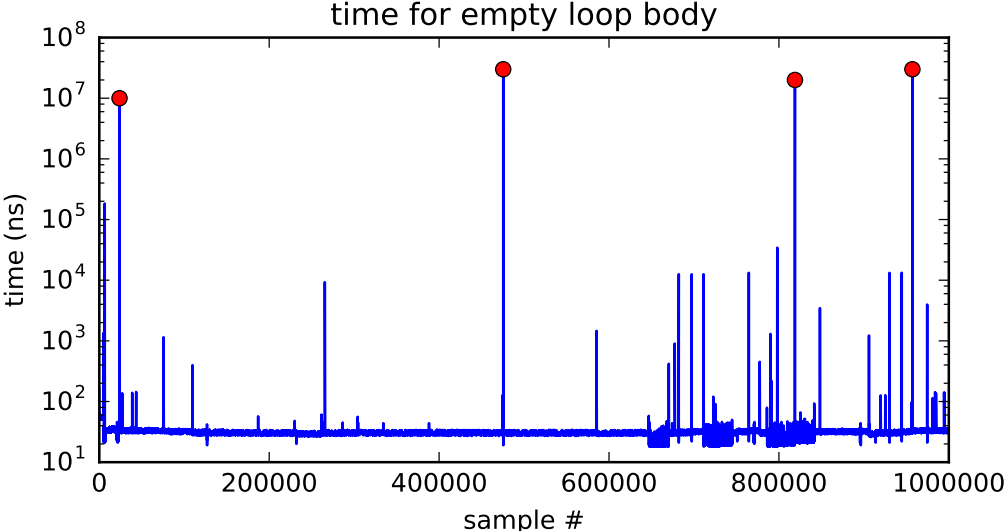
```
long times[NUM_TIMINGS];
int main(void) {
    for (int i = 0; i < N; ++i) {
        long start, end;
        start = get_time();
        /* do nothing */
        end = get_time();
        times[i] = end - start;
    }
    output_timings(times);
}
```

same instructions — **same difference** each time?

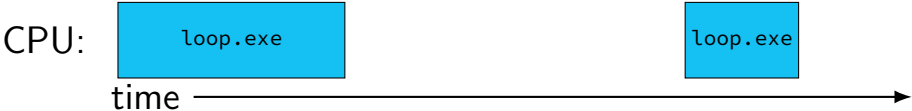
# doing nothing on a busy system



# doing nothing on a busy system



# time multiplexing



# time multiplexing



...

```
call get_time
```

```
// whatever get_time does
```

```
movq %rax, %rbp
```

———— million cycle delay ————

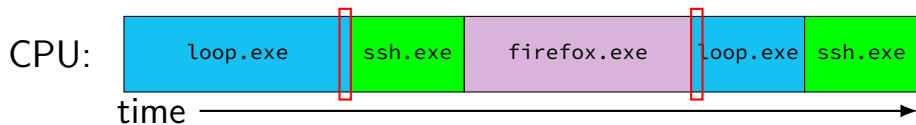
```
call get_time
```

```
// whatever get_time does
```

```
subq %rbp, %rax
```

...

# time multiplexing



...

```
call get_time
```

```
    // whatever get_time does
```

```
movq %rax, %rbp
```

———— million cycle delay ————

```
call get_time
```


```
    // whatever get_time does
```

```
subq %rbp, %rax
```

...

# time multiplexing really




 = operating system



# time multiplexing really



 = operating system

exception happens

return from exception

# OS and time multiplexing

starts running instead of normal program

mechanism for this: **exceptions** (later)

saves old program counter, registers somewhere

sets new registers, jumps to new program counter

called **context switch**

saved information called **context**

# context

all registers values

`%rax %rbx, ..., %rsp, ...`

condition codes

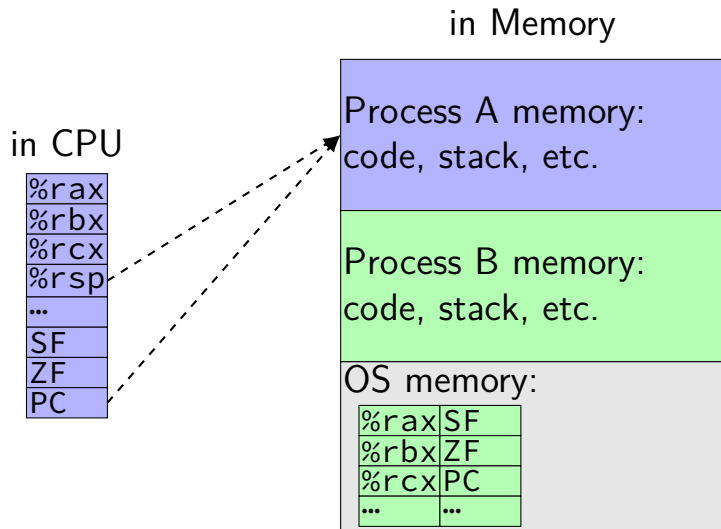
program counter

i.e. all visible state in your CPU except memory

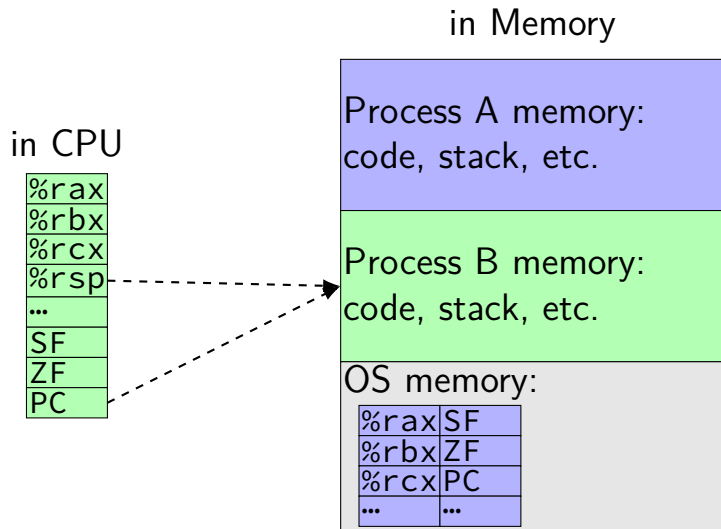
# context switch pseudocode

```
context_switch(last, next):  
    copy_preexception_pc last->pc  
    mov rax, last->rax  
    mov rcx, last->rcx  
    mov rdx, last->rdx  
    ...  
    mov next->rdx, rdx  
    mov next->rcx, rcx  
    mov next->rax, rax  
    jmp next->pc
```

# contexts (A running)



# contexts (B running)



# memory protection

reading from another program's memory?

Program A

```
0x10000: .word 42
// ...
// do work
// ...
movq 0x10000, %rax
```

Program B

```
// while A is working:
movq $99, %rax
movq %rax, 0x10000
...
```

# memory protection

reading from another program's memory?

Program A

```
0x10000: .word 42
        // ...
        // do work
        // ...
        movq 0x10000, %rax
```

Program B

```
// while A is working:
movq $99, %rax
movq %rax, 0x10000
...
```

result: %rax is ...

- A. 42
- B. 99
- C. 0x10000
- D. 42 or 99 (depending on timing/program layout/etc)
- E. 42 or program might crash (depending on ...)
- F. 99 or program might crash (depending on ...)
- G. 42 or 99 or program might crash (depending on ...)
- H. something else



# memory protection

reading from another program's memory?

Program A

```
0x10000: .word 42
// ...
// do work
// ...
movq 0x10000, %rax
```

Program B

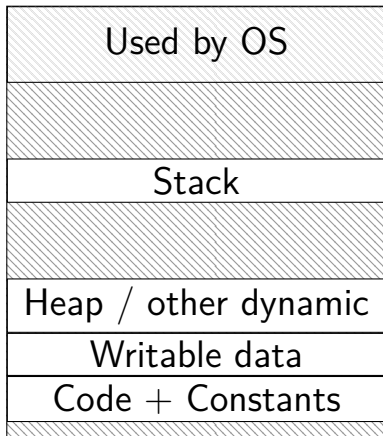
```
// while A is working:
movq $99, %rax
movq %rax, 0x10000
...
```

result: %rax is 42 (always)

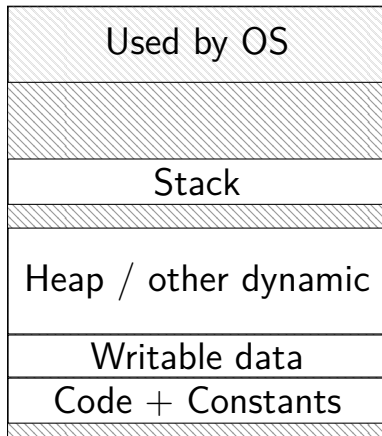
result: **might crash**

# program memory (two programs)

Program A



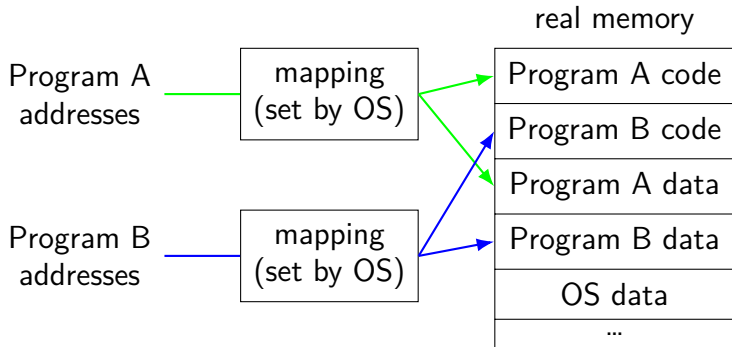
Program B



# address space

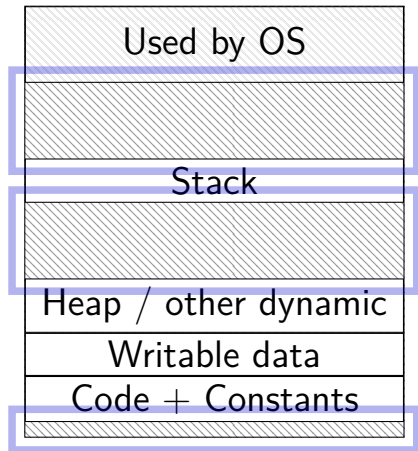
programs have **illusion of own memory**

called a program's **address space**

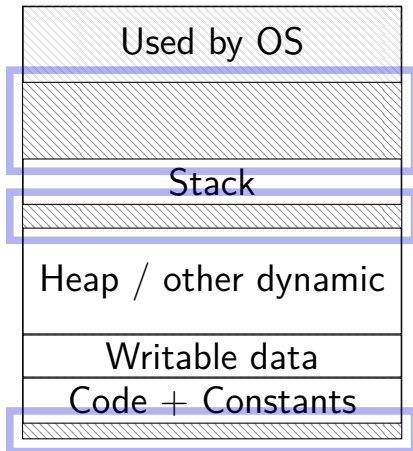


# program memory (two programs)

Program A



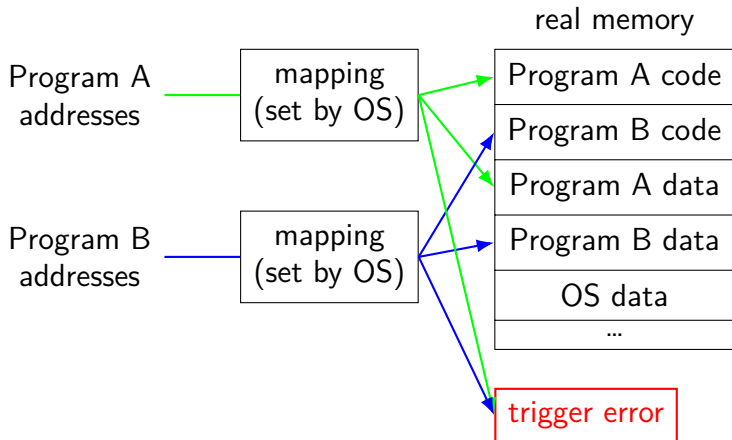
Program B



# address space

programs have **illusion of own memory**

called a program's **address space**



# address space mechanisms

topic after exceptions

called **virtual memory**

mapping called **page tables**

mapping part of what is changed in context switch

# context

all registers values

`%rax %rbx, ..., %rsp, ...`

condition codes

program counter

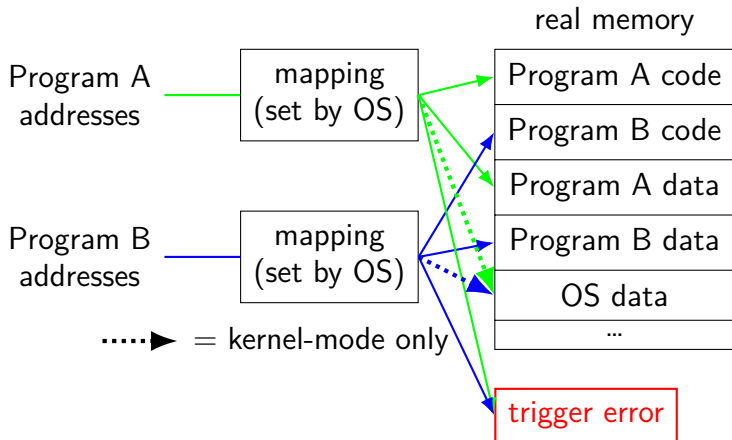
~~i.e. all visible state in your CPU except memory~~

**address space**: map from program to real addresses

# address space

programs have **illusion of own memory**

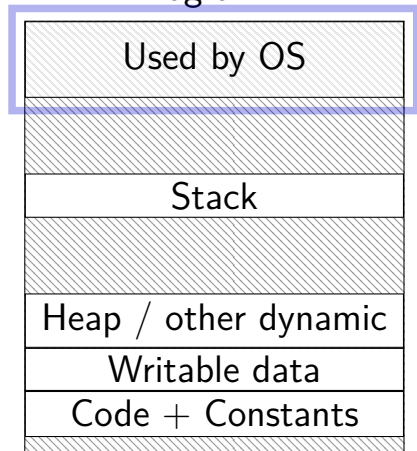
called a program's **address space**



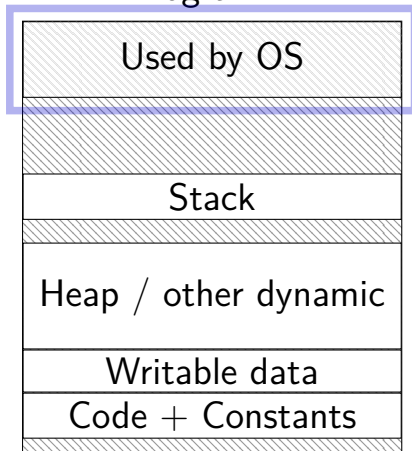


# program memory (two programs)

Program A



Program B



# backup slides

# 128-bit version, too

history: 256-bit vectors added in extension called AVX (c. 2011)

before: 128-bit vectors added in extension called SSE (c. 1999)

128-bit intrinsics exist, too:

`__m256i` becomes `__m128i`

`_mm256_add_epi32` becomes `_mm_add_epi32`

`_mm256_loadu_si256` becomes `_mm_loadu_si128`

# fickle compiler vectorization (1)

GCC 8.2 and Clang 7.0 generate vector instructions for this:

```
#define N 1024
void foo(unsigned int *A, unsigned int *B) {
    for (int k = 0; k < N; ++k)
        for (int i = 0; i < N; ++i)
            for (int j = 0; j < N; ++j)
                B[i * N + j] += A[i * N + k] * A[k * N + j];
}
```

but not:

```
#define N 1024
void foo(unsigned int *A, unsigned int *B) {
    for (int i = 0; i < N; ++i)
        for (int j = 0; j < N; ++j)
            for (int k = 0; k < N; ++k)
                B[i * N + j] += A[i * N + k] * A[j * N + k];
}
```

## fickle compiler vectorization (2)

Clang 5.0.0 generates vector instructions for this:

```
void foo(int N, unsigned int *A, unsigned int *B) {  
    for (int k = 0; k < N; ++k)  
        for (int i = 0; i < N; ++i)  
            for (int j = 0; j < N; ++j)  
                B[i * N + j] += A[i * N + k] * A[k * N + j];  
}
```

but not: (fixed in later versions)

```
void foo(long N, unsigned int *A, unsigned int *B) {  
    for (long k = 0; k < N; ++k)  
        for (long i = 0; i < N; ++i)  
            for (long j = 0; j < N; ++j)  
                B[i * N + j] += A[i * N + k] * A[k * N + j];  
}
```

## exercise: when optimizations backfire...

Which of these optimizations are likely to **increase** machine code size? (**Select all that apply.**)

Which of these optimizations are likely to **increase** number of instructions executed? (**Select all that apply.**)

- A. cache blocking
- B. function inlining
- C. loop unrolling
- D. moving a calculation outside a loop
- E. multiple accumulators (after loop unrolling)

# loop unrolling performance

on my laptop with 992 elements (fits in L1 cache)

work/loop iteration	cycles/element	instructions/element
1	1.33	4.02
2	1.03	2.52
4	1.02	1.77
8	1.01	1.39
16	1.01	1.21
32	1.01	1.15

1.01 cycles/element — latency bound

# multiple accumulators

```
int i;
long sum1 = 0, sum2 = 0;
for (i = 0; i + 1 < N; i += 2) {
    sum1 += A[i];
    sum2 += A[i+1];
}
// handle leftover, if needed
if (i < N)
    sum1 += A[i];
sum = sum1 + sum2;
```



# multiple accumulators performance

on my laptop with 992 elements (fits in L1 cache)

16x unrolling, variable number of accumulators

accumulators	cycles/element	instructions/element
1	1.01	1.21
2	0.57	1.21
4	0.57	1.23
8	0.59	1.24
16	0.76	1.57

starts hurting after too many accumulators

why?

# multiple accumulators performance

on my laptop with 992 elements (fits in L1 cache)

16x unrolling, variable number of accumulators

accumulators	cycles/element	instructions/element
1	1.01	1.21
2	0.57	1.21
4	0.57	1.23
8	0.59	1.24
16	0.76	1.57

starts hurting after too many accumulators

why?

## 8 accumulator assembly

```
sum1 += A[i + 0];  
sum2 += A[i + 1];  
...  
...
```

---

```
addq    (%rdx), %rax    // sum1 +=  
addq    8(%rdx), %rcx   // sum2 +=  
subq    $-128, %rdx     // i +=  
addq    -112(%rdx), %rbx // sum3 +=  
addq    -104(%rdx), %r11 // sum4 +=  
...  
.....  
cmpq    %r14, %rdx
```

register for each of the sum1, sum2, ...variables:

# 16 accumulator assembly

compiler runs out of registers

starts to use the stack instead:

```
movq    32(%rdx), %rax // get A[i+13]
addq    %rax, -48(%rsp) // add to sum13 on stack
```

code does **extra cache accesses**

also — already using all the adders available all the time

so performance increase not possible

# multiple accumulators performance

on my laptop with 992 elements (fits in L1 cache)

16x unrolling, variable number of accumulators

accumulators	cycles/element	instructions/element
1	1.01	1.21
2	0.57	1.21
4	0.57	1.23
8	0.59	1.24
16	0.76	1.57

starts hurting after too many accumulators

why?

# maximum performance

2 additions per element:

- one to add to sum

- one to compute address (part of mov)

3/16 add/sub/cmp + 1/16 branch per element:

- over 16 because loop unrolled 16 times

- loop overhead

- compiler not as efficient as it could have been

$2 + 3/16 + 1/16 = 2 + 1/4$  instructions per element

# hardware limits on my machine

4(?) register renamings per cycle

(Intel doesn't really publish exact numbers here...)

4-6 instructions decoded/cycle

(depending on instructions)

4(?) microinstructions committed/cycle

4 (add or cmp+branch executed)/cycle

# hardware limits on my machine

4(?) register renamings per cycle

(Intel doesn't really publish exact numbers here...)

4-6 instructions decoded/cycle

(depending on instructions)

4(?) microinstructions committed/cycle

4 (add or cmp+branch executed)/cycle

$(2 + 1/4) \div 4 \approx 0.57$  cycles/element



## getting over this limit

the  $+1/4$  was from loop overhead

solution: more loop unrolling!

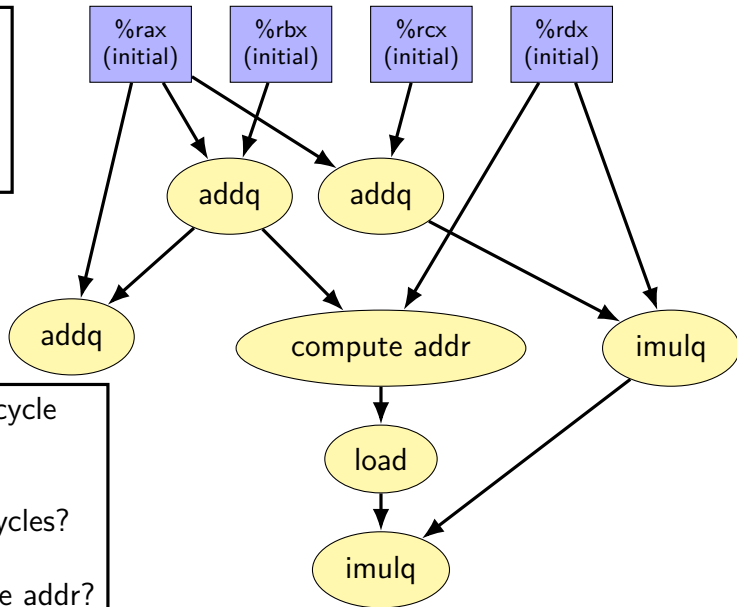
common theme with optimization:

fix one bottleneck (need to do adds one after the other)

find another bottleneck

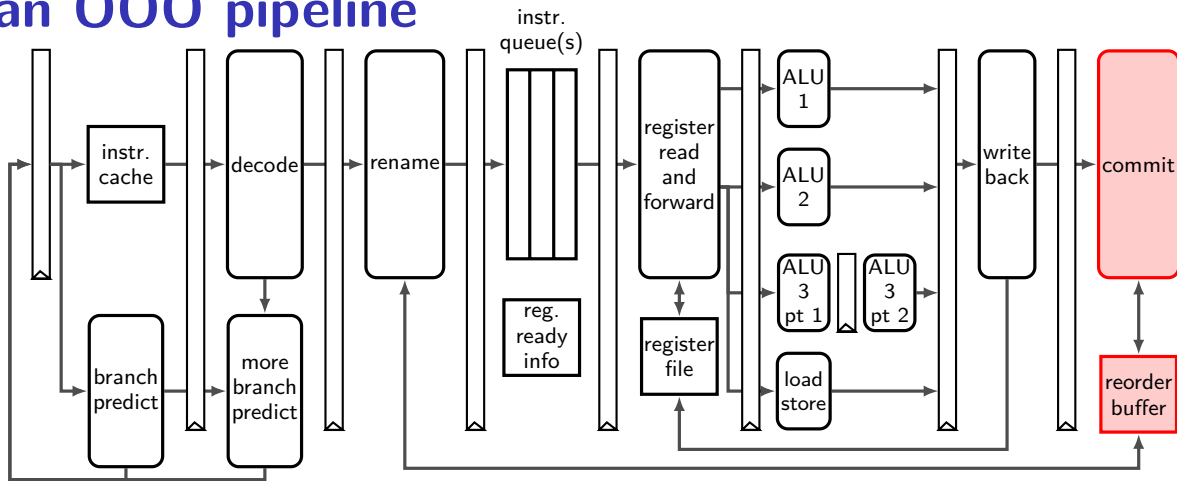
# a data flow example

```
addq %rax, %rbx
addq %rax, %rcx
imulq %rdx, %rcx
movq (%rbx, %rdx), %r8
imulq %r8, %rcx
addq %rax, %rbx
```



addq, compute addr: 1 cycle  
imulq: 3 cycle latency  
load: 3 cycle latency  
Q1: latency bound on cycles?  
Q2: what can be done  
at same time as compute addr?

# an OOO pipeline



# reorder buffer: on rename

phys → arch. reg  
for new instrs

<b>arch. reg</b>	<b>phys. reg</b>
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07
...	...

free list

%x19
%x23
...
...

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07
...	...

free list

%x19
%x23
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		

reorder buffer contains instructions started,  
but not fully finished new entries created on rename  
(not enough space? stall rename stage)

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07
...	...

free list

%x19
%x23
...
...

remove here  
when committed →

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		

add here  
on rename →

place newly started instruction at end of buffer  
remember at least its destination register  
(both architectural and physical versions)

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	<del>%x07</del> %x19
...	...

free list

%x19
%x23
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		
32	0x1230	%rdx / %x19		

add here  
on rename



next renamed instruction goes in next slot, etc.

# reorder buffer: on rename

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x23
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x23		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		
32	0x1230	%rdx / %x19		

add here  
on rename





# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x13
...
...

remove here  
when committed



reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x24		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31		
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34		
19	0x1249	%rax / %x38		
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x13
...
...

remove here  
when committed →

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x24		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓

instructions marked done in reorder buffer  
when result is computed  
but not removed from reorder buffer ('committed') yet

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

phys → arch. reg when committed  
for committed

arch. reg	phys. reg
%rax	%x30
%rcx	%x28
%rbx	%x23
%rdx	%x21
...	...

free list

%x19
%x13
...
...

remove here  
→

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x24		
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓

commit stage tracks architectural to physical register map for committed instructions

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

free list

%x19
%x13
...
%x23

phys → arch. reg when committed  
for committed

arch. reg	phys. reg
%rax	%x30
%rcx	%x28
%rbx	%x23 %x24
%rdx	%x21
...	...

remove here  
→

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x24	✓	
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓
32	0x1230	%rdx / %x19		

when next-to-commit instruction is done  
update this register map and free register list  
and remove instr. from reorder buffer

# reorder buffer: on commit

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x07 %x19
...	...

phys → arch. reg remove here  
for committed when committed

arch. reg	phys. reg
%rax	%x30
%rcx	%x28
%rbx	%x23 %x24
%rdx	%x21
...	...

free list

%x19
%x13
...
%x23

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	<del></del>
15	0x1239	%rax / %x30		
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32		
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC		
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12		✓
32	0x1230	%rdx / %x19		

when next-to-commit instruction is done  
update this register map and free register list  
and remove instr. from reorder buffer

# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x19
...	...

free list

%x19
%x13
...
...

phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	
<del>15</del>	<del>0x1239</del>	<del>%rax / %x30</del>	<del>✓</del>	
<del>16</del>	<del>0x1242</del>	<del>%rcx / %x31</del>	<del>✓</del>	
<del>17</del>	<del>0x1244</del>	<del>%rcx / %x32</del>	<del>✓</del>	
<del>18</del>	<del>0x1248</del>	<del>%rdx / %x34</del>	<del>✓</del>	
<del>19</del>	<del>0x1249</del>	<del>%rax / %x38</del>	<del>✓</del>	
→ 20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		

# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x12
%rcx	%x17
%rbx	%x13
%rdx	%x19
...	...

free list

%x19
%x13
...
...

phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	
<del>15</del>	<del>0x1239</del>	<del>%rax / %x30</del>	<del>✓</del>	
<del>16</del>	<del>0x1242</del>	<del>%rcx / %x31</del>	<del>✓</del>	
<del>17</del>	<del>0x1244</del>	<del>%rcx / %x32</del>	<del>✓</del>	
<del>18</del>	<del>0x1248</del>	<del>%rdx / %x34</del>	<del>✓</del>	
<del>19</del>	<del>0x1249</del>	<del>%rax / %x38</del>	<del>✓</del>	
→ 20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		

when committing a mispredicted instruction...  
this is where we undo mispredicted instructions

# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x38
%rcx	%x32
%rbx	%x24
%rdx	%x34
...	...

phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...



free list

%x19
%x13
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
<del>14</del>	<del>0x1233</del>	<del>%rbx / %x24</del>	<del>✓</del>	
<del>15</del>	<del>0x1239</del>	<del>%rax / %x30</del>	<del>✓</del>	
<del>16</del>	<del>0x1242</del>	<del>%rcx / %x31</del>	<del>✓</del>	
<del>17</del>	<del>0x1244</del>	<del>%rcx / %x32</del>	<del>✓</del>	
<del>18</del>	<del>0x1248</del>	<del>%rdx / %x34</del>	<del>✓</del>	
<del>19</del>	<del>0x1249</del>	<del>%rax / %x38</del>	<del>✓</del>	
20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		



copy commit register map into rename register map  
so we can start fetching from the correct PC



# reorder buffer: commit mispredict (one way)

phys → arch. reg  
for new instrs

arch. reg	phys. reg
%rax	%x38
%rcx	%x32
%rbx	%x24
%rdx	%x34
...	...

phys → arch. reg  
for committed

arch. reg	phys. reg
%rax	%x30 %x38
%rcx	%x31 %x32
%rbx	%x23 %x24
%rdx	%x21 %x34
...	...



free list

%x19
%x13
...
...

reorder buffer (ROB)

instr num.	PC	dest. reg	done?	mispred? / except?
14	0x1233	%rbx / %x24	✓	
15	0x1239	%rax / %x30	✓	
16	0x1242	%rcx / %x31	✓	
17	0x1244	%rcx / %x32	✓	
18	0x1248	%rdx / %x34	✓	
19	0x1249	%rax / %x38	✓	
20	0x1254	PC	✓	✓
21	0x1260	%rcx / %x17		
...	...	...	...	...
31	0x129f	%rax / %x12	✓	
32	0x1230	%rdx / %x19		



...and discard all the mispredicted instructions  
(without committing them)

## better? alternatives

- can take snapshots of register map on each branch

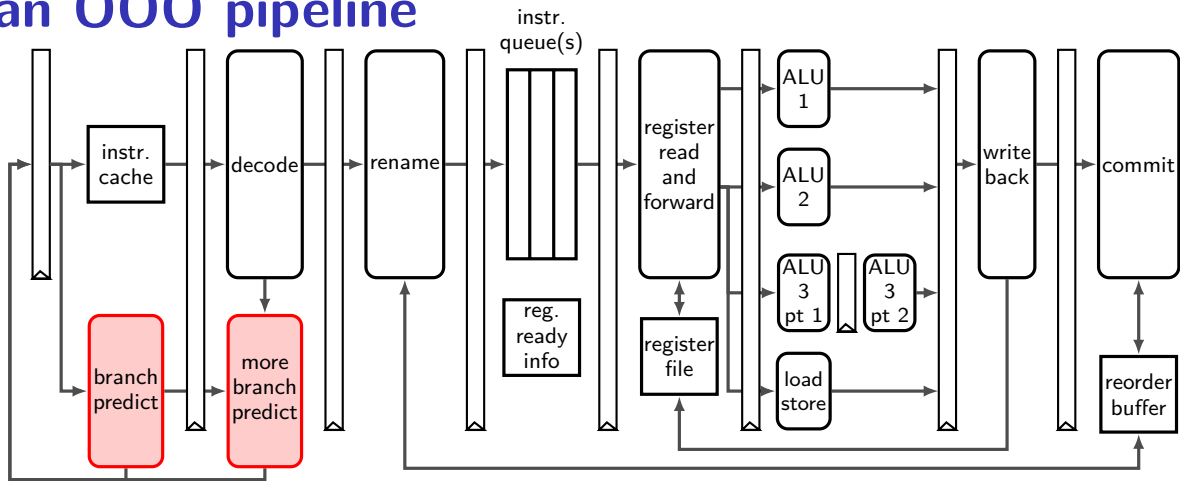
  - don't need to reconstruct the table  
(but how to efficiently store them)

- can reconstruct register map before we commit the branch instruction

  - need to let reorder buffer be accessed even more?

- can track more/different information in reorder buffer

# an OOO pipeline



# branch target buffer

can take several cycles to fetch+decode jumps, calls, returns

still want 1-cycle prediction of next thing to fetch

# BTB: cache for branches

idx	valid	tag	ofst	type	target	(more info?)	valid	...
0x00	1	0x400	5	Jxx	0x3FFFF3	...	1	...
0x01	1	0x401	C	JMP	0x401035	---	0	...
0x02	0	---	---	---	---	---	0	...
0x03	1	0x400	9	RET	---	...	0	...
...	...	...	...	...	...	...	...	...
0xFF	1	0x3FF	8	CALL	0x404033	...	0	...

```
0x3FFFF3:  movq %rax, %rsi
0x3FFFF7:  pushq %rbx
0x3FFFF8:  call 0x404033
0x400001:  popq %rbx
0x400003:  cmpq %rbx, %rax
0x400005:  jle 0x3FFFF3
...
0x400031:  ret
...
```

# BTB: cache for branches

idx	valid	tag	ofst	type	target	(more info?)	valid	...
0x00	1	0x400	5	Jxx	0x3FFFF3	...	1	...
0x01	1	0x401	C	JMP	0x401035	----	0	...
0x02	0	---	---	---	---	----	0	...
0x03	1	0x400	9	RET	----	...	0	...
...	...	...	...	...	...	...	...	...
0xFF	1	0x3FF	8	CALL	0x404033	...	0	...

```
0x3FFFF3:  movq %rax, %rsi
0x3FFFF7:  pushq %rbx
0x3FFFF8:  call 0x404033
0x400001:  popq %rbx
0x400003:  cmpq %rbx, %rax
0x400005:  jle 0x3FFFF3
...
0x400031:  ret
...
```

# BTB: cache for branches

idx	valid	tag	ofst	type	target	(more info?)	valid	...
0x00	1	0x400	5	Jxx	0x3FFFF3	...	1	...
0x01	1	0x401	C	JMP	0x401035	----	0	...
0x02	0	---	---	---	---	----	0	...
0x03	1	0x400	9	RET	----	...	0	...
...	...	...	...	...	...	...	...	...
0xFF	1	0x3FF	8	CALL	0x404033	...	0	...

```
0x3FFFF3:  movq %rax, %rsi
0x3FFFF7:  pushq %rbx
0x3FFFF8:  call 0x404033
0x400001:  popq %rbx
0x400003:  cmpq %rbx, %rax
0x400005:  jle 0x3FFFF3
...
0x400031:  ret
...
```

## aside on branch pred. and performance

modern branch predictors are very good

we might explore how later in semester (if time)

...usually can assume most branches will be predicted

but could be a problem if really no pattern

e.g. branch based on random number?

generally: measure and see



# if branch prediction is bad...

avoiding branches — conditional move, etc.

replace multiple branches with single lookup?

one misprediction better than  $K$ ?

## recall: shifts

we mentioned that compilers compile  $x/4$  into a shift instruction

they are really good at these types of transformation...

“strength reduction”: replacing complicated op with simpler one

but can't do without seeing special case (e.g. divide by constant)

# Intel Skylake OOO design

2015 Intel design — codename 'Skylake'

94-entry instruction queue-equivalent

168 physical integer registers

168 physical floating point registers

4 ALU functional units

but some can handle more/different types of operations than others

2 load functional units

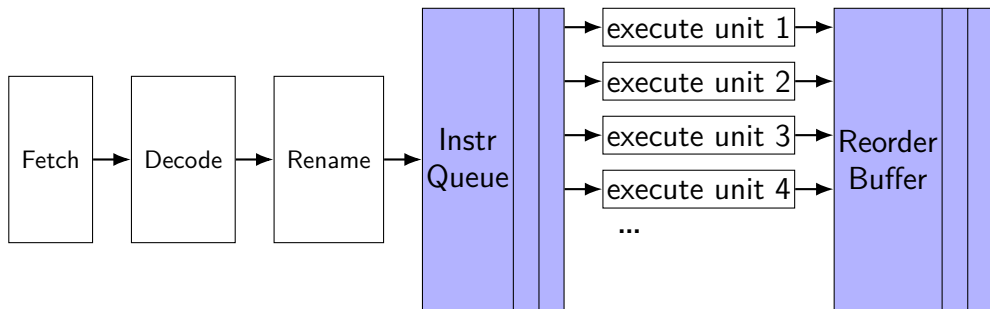
but pipelined: supports multiple pending cache misses in parallel

1 store functional unit

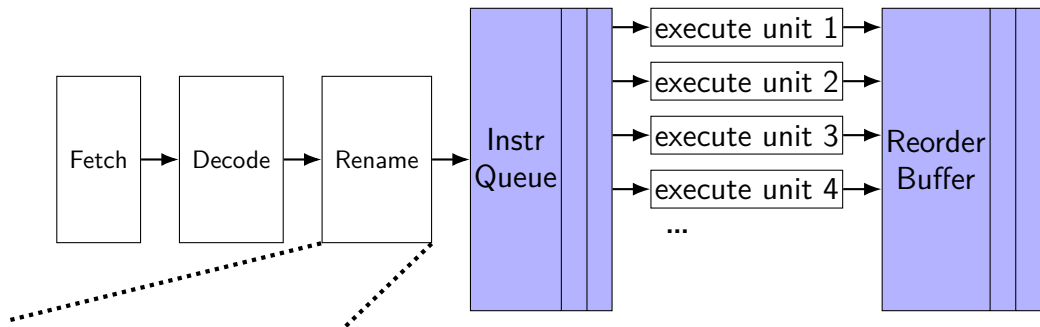
224-entry reorder buffer

determines how far ahead branch mispredictions, etc. can happen

# exceptions and OOO (one strategy)



# exceptions and OOO (one strategy)

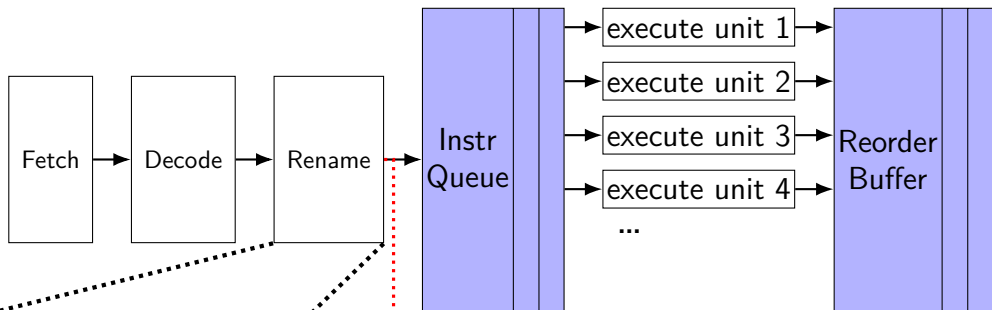


free regs for new instrs

X19
X23
...

arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

# exceptions and OOO (one strategy)



free regs for new instrs

X19
X23
...

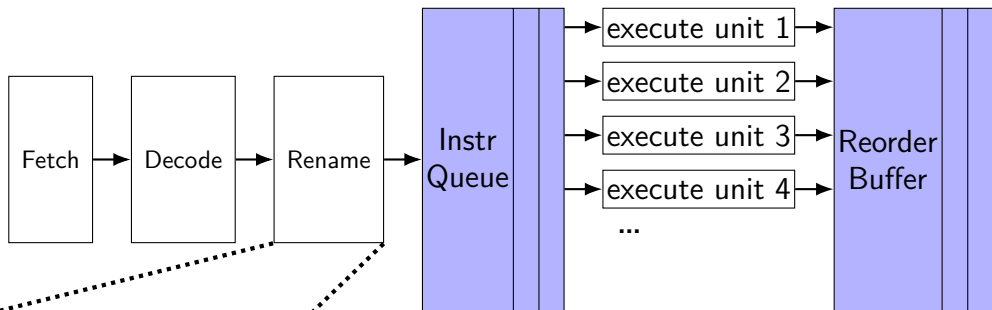
arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

done instrs  
committed in order

new instrs added

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32		
18	0x1248	RDX / X34		
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05		
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



free regs for new instrs      for complete instrs

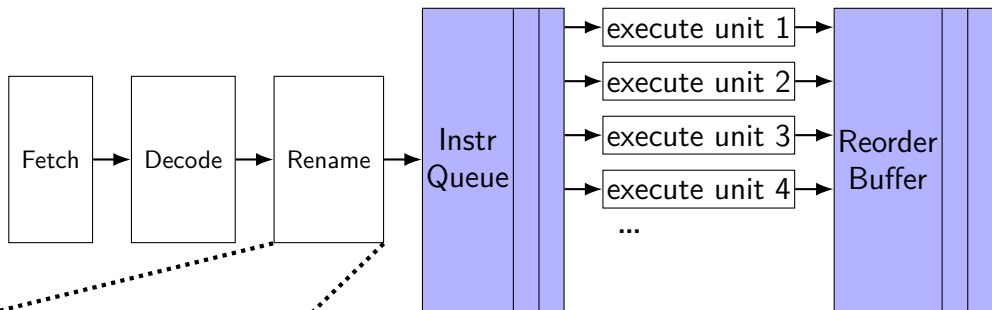
X19
X23
...

arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

arch. reg	phys. reg
RAX	X21
RCX	X2 X32
RBX	X48
RDX	X37
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34		
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05		
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



free regs for new instrs      for complete instrs

X19
X23
...

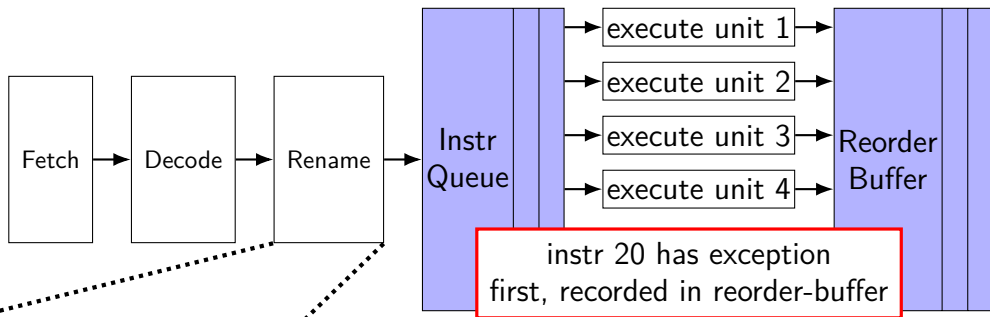
arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

arch. reg	phys. reg
RAX	X21
RCX	X2 X32
RBX	X48
RDX	X37
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	<del>RCX / X32</del>	✓	
18	0x1248	RDX / X34		
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05		
21	0x1260	R8 / X06		
...	...	...	...	...



# exceptions and OOO (one strategy)



free regs for new instrs      for complete instrs

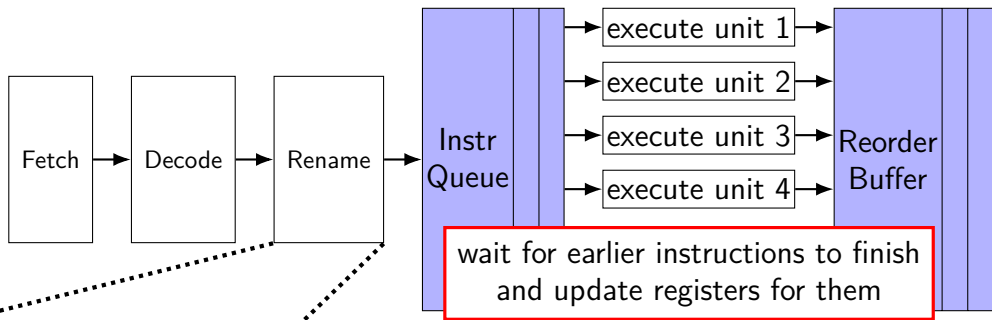
X19
X23
...

arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

arch. reg	phys. reg
RAX	X21
RCX	X2 X32
RBX	X48
RDX	X37
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34		
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05	✓	✓
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



free regs for new instrs    for complete instrs

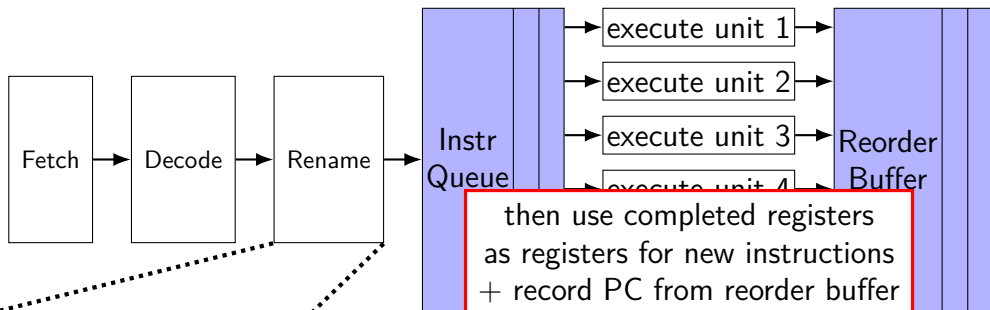
X19
X23
...

arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

arch. reg	phys. reg
RAX	X21 X38
RCX	X2 X32
RBX	X48
RDX	X37 X34
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34	✓	
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05	✓	✓
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



then use completed registers  
 as registers for new instructions  
 + record PC from reorder buffer  
 + jump to exception handler

free regs for new instrs for complet

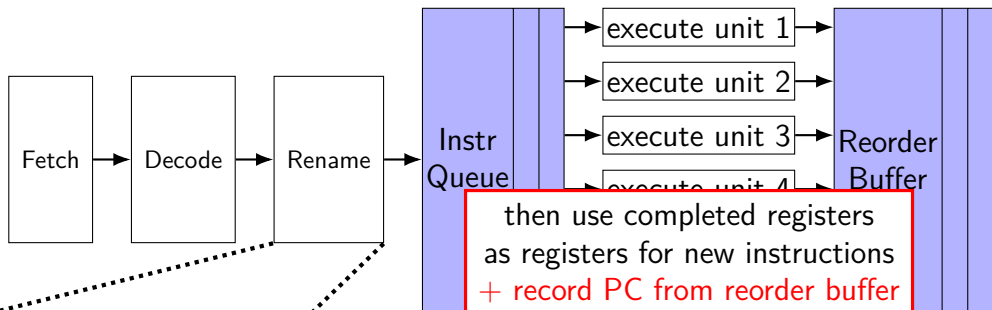
X19
X23
...

arch. reg	phys. reg
RAX	X38
RCX	X32
RBX	X48
RBX	X34
...	...

arch. reg	phys. reg
RAX	X21 X38
RCX	X2 X32
RBX	X48
RDX	X37 X34
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34	✓	
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05	✓	✓
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



free regs for new instrs for complet

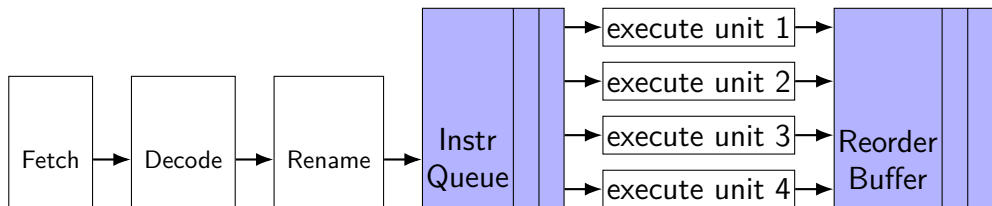
X19
X23
...

arch. reg	phys. reg
RAX	X38
RCX	X32
RBX	X48
RBX	X34
...	...

arch. reg	phys. reg
RAX	X21 X38
RCX	X2 X32
RBX	X48
RDX	X37 X34
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34	✓	
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05	✓	✓
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



variation: could store architectural reg. values instead of mapping for completed instrs.  
(and copy values instead of mapping on exception)

free regs for new instrs      for complete instrs

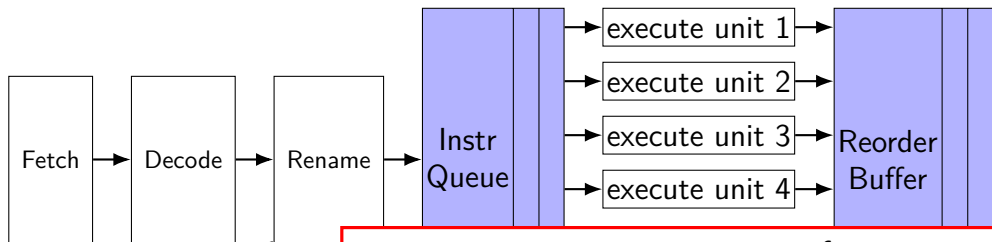
X19
X23
...

arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

arch. reg	value
RAX	0x12343
RCX	0x234543
RBX	0x56782
RDX	0xF83A4
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34	✓	
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05	✓	✓
21	0x1260	R8 / X06		
...	...	...	...	...

# exceptions and OOO (one strategy)



stopping instructions in progress for exception  
similar to how 'squashing' mispredicted instructions

free regs for new instrs for complete instrs

X19
X23
...

arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

arch. reg	phys. reg
RAX	X21 X38
RCX	X2 X32
RBX	X48
RDX	X37 X34
...	...

instr num.	PC	dest. reg	done?	except?
...	...	...	...	...
17	0x1244	RCX / X32	✓	
18	0x1248	RDX / X34	✓	
19	0x1249	RAX / X38	✓	
20	0x1254	R8 / X05	✓	✓
21	0x1260	R8 / X06		
...	...	...	...	...

# addressing efficiency

```
for (int kk = 0; kk < N; kk += 2) {  
    for (int i = 0; i < N; ++i) {  
        for (int j = 0; j < N; ++j) {  
            float Cij = C[i * N + j];  
            for (int k = kk; k < kk + 2; ++k) {  
                Cij += A[i * N + k] * B[k * N + j];  
            }  
            C[i * N + j] = Cij;  
        }  
    }  
}
```

tons of multiplies by N??

isn't that slow?

# addressing transformation

```
for (int kk = 0; k < N; kk += 2)
  for (int i = 0; i < N; ++i) {
    for (int j = 0; j < N; ++j) {
      float Cij = C[i * N + j];
      float *Bkj_pointer = &B[kk * N + j];
      for (int k = kk; k < kk + 2; ++k) {
        // Bij += A[i * N + k] * A[k * N + j~];
        Bij += A[i * N + k] * Bkj_pointer;
        Bkj_pointer += N;
      }
      C[i * N + j] = Bij;
    }
  }
```

transforms loop to **iterate with pointer**

**compiler** will often do this

increment/decrement by N ( $\times$  sizeof(float))



# addressing transformation

```
for (int kk = 0; k < N; kk += 2)
  for (int i = 0; i < N; ++i) {
    for (int j = 0; j < N; ++j) {
      float Cij = C[i * N + j];
      float *Bkj_pointer = &B[kk * N + j];
      for (int k = kk; k < kk + 2; ++k) {
        // Bij += A[i * N + k] * A[k * N + j~];
        Bij += A[i * N + k] * Bkj_pointer;
        Bkj_pointer += N;
      }
      C[i * N + j] = Bij;
    }
  }
```

transforms loop to **iterate with pointer**

**compiler** will often do this

increment/decrement by N ( $\times$  sizeof(float))

# addressing efficiency

compiler will **usually** eliminate slow multiplies  
doing transformation yourself often slower if so

```
i * N; ++i into i_times_N; i_times_N += N
```

way to check: see if assembly uses lots multiplies in loop

if it doesn't — do it yourself

## another addressing transformation

```
for (int i = 0; i < n; i += 4) {  
    C[(i+0) * n + j] += A[(i+0) * n + k] * B[k * n + j];  
    C[(i+1) * n + j] += A[(i+1) * n + k] * B[k * n + j];  
    // ...  
}
```

---

```
int offset = 0;  
float *Ai0_base = &A[k];  
float *Ai1_base = Ai0_base + n;  
float *Ai2_base = Ai1_base + n;  
// ...  
for (int i = 0; i < n; i += 4) {  
    C[(i+0) * n + j] += Ai0_base[offset] * B[k * n + j];  
    C[(i+1) * n + j] += Ai1_base[offset] * B[k * n + j];  
    // ...  
    offset += n;  
}
```

compiler will sometimes do this, too

## another addressing transformation

```
for (int i = 0; i < n; i += 4) {  
    C[(i+0) * n + j] += A[(i+0) * n + k] * B[k * n + j];  
    C[(i+1) * n + j] += A[(i+1) * n + k] * B[k * n + j];  
    // ...  
}
```

---

```
int offset = 0;  
float *Ai0_base = &A[k];  
float *Ai1_base = Ai0_base + n;  
float *Ai2_base = Ai1_base + n;  
// ...  
for (int i = 0; i < n; i += 4) {  
    C[(i+0) * n + j] += Ai0_base[offset] * B[k * n + j];  
    C[(i+1) * n + j] += Ai1_base[offset] * B[k * n + j];  
    // ...  
    offset += n;  
}
```

compiler will sometimes do this, too

## another addressing transformation

```
for (int i = 0; i < n; i += 20) {  
    C[(i+0) * n + j] += A[(i+0) * n + k] * B[k * n + j];  
    C[(i+1) * n + j] += A[(i+1) * n + k] * B[k * n + j];  
    // ...
```

---

```
int offset = 0;  
float *Ai0_base = &A[0*n+k];  
float *Ai1_base = Ai0_base + n;  
float *Ai2_base = Ai1_base + n;  
// ...  
for (int i = 0; i < n; i += 20) {  
    C[(i+0) * n + j] += Ai0_base[i*n] * B[k * n + j];  
    C[(i+1) * n + j] += Ai1_base[i*n] * B[k * n + j];  
    // ...  
    offset += n;
```

storing 20  $A_{iX\_base}$ ? — need the stack

maybe faster (quicker address computation)

maybe slower (can't do enough loads)

## another addressing transformation

```
for (int i = 0; i < n; i += 20) {  
    C[(i+0) * n + j] += A[(i+0) * n + k] * B[k * n + j];  
    C[(i+1) * n + j] += A[(i+1) * n + k] * B[k * n + j];  
    // ...
```

---

```
int offset = 0;  
float *Ai0_base = &A[0*n+k];  
float *Ai1_base = Ai0_base + n;  
float *Ai2_base = Ai1_base + n;  
// ...  
for (int i = 0; i < n; i += 20) {  
    C[(i+0) * n + j] += Ai0_base[i*n] * B[k * n + j];  
    C[(i+1) * n + j] += Ai1_base[i*n] * B[k * n + j];  
    // ...  
    offset += n;
```

storing 20  $A_{iX\_base}$ ? — need the stack

maybe faster (quicker address computation)

maybe slower (can't do enough loads)

# alternative addressing transformation

instead of:

```
float *Ai0_base = &A[0*n+k];
float *Ai1_base = Ai0_base + n;
// ...
for (int i = 0; i < n; i += 20) {
    C[(i+0) * n + j] += Ai0_base[i*n] * B[k * n + j];
    C[(i+1) * n + j] += Ai1_base[i*n] * B[k * n + j];
    // ...
}
```

---

could do:

```
float *Ai0_base = &A[k];
for (int i = 0; i < n; i += 20) {
    float *A_ptr = &Ai0_base[i*n];
    C[(i+0) * n + j] += *A_ptr * A[k * n + j];
    A_ptr += n;
    C[(i+1) * n + j] += *A_ptr * B[k * n + j];
    // ...
}
```

avoids spilling on the stack, but more dependencies

# alternative addressing transformation

instead of:

```
float *Ai0_base = &A[0*n+k];
float *Ai1_base = Ai0_base + n;
// ...
for (int i = 0; i < n; i += 20) {
    C[(i+0) * n + j] += Ai0_base[i*n] * B[k * n + j];
    C[(i+1) * n + j] += Ai1_base[i*n] * B[k * n + j];
    // ...
}
```

---

could do:

```
float *Ai0_base = &A[k];
for (int i = 0; i < n; i += 20) {
    float *A_ptr = &Ai0_base[i*n];
    C[(i+0) * n + j] += *A_ptr * A[k * n + j];
    A_ptr += n;
    C[(i+1) * n + j] += *A_ptr * B[k * n + j];
    // ...
}
```

avoids spilling on the stack, but more dependencies



# addressing efficiency generally

mostly: compiler does very good job itself

- eliminates multiplications, use pointer arithmetic

- often will do better job than if how typically programming would do it manually

sometimes compiler won't take the best option

- if spilling to the stack: can cause weird performance anomalies

- if indexing gets too complicated — might not remove multiply

if compiler doesn't, you can always make addressing simple yourself

- convert to pointer arith. without multiplies