



# Changelog

2022-11-02: simple blocking – expanded: correct typo of premature 'i += 2' for 'i += 1'

2022-11-05: simple blocking – counting loads: also correct typo of premature i += 2 for i += 1

# last time

counting cache misses in C code

- mapping of arrays to sets

- alignment (or not?)

- analyzing sets separately

approximate miss counting

- assessing locality (spatial/temporal)

- look at innermost loop, accesses to array

- same as previous access: 0% chance of miss

- adjacent:  $1/(\text{elems per block})$  chance of miss

- non-adjacent, not accessed recently: 100% chance of miss

## quiz Q2

normal version:

hit detection + data extraction take 5 cycles

+100 cycles on miss = miss penalty

105 cycles miss time

optimized version

hit detection takes 2 cycles

data extraction takes 5 cycles (done in parallel with hit detection)

+100 cycles after hit detection on miss

102 cycle miss time

$$90\% \cdot c + 10\% \cdot 102 = 14.7 \text{ cycle AMAT}$$

## quiz Q3 (1)

```
for (int j = 0; j < 2; j += 1) {  
    for (int i = 0; i < 4; i += 1) {  
        if (sum > array[i * 3 + j]) {  
            sum += array[i * 3 + j];  
        }  
    }  
}
```

accesses index 0, 3, 6, 9, 1, 4, 7, 10

## quiz Q3 (2)

0 miss: set 0 {0+1,--}; set 1 {--,--}  
3 miss: set 0 {0+1,--}; set 1 {2+3,--}  
6 miss: set 1 {0+1,--}; set 1 {2+3,6+7}  
9 miss: set 0 {0+1,8+9}; set 1 {2+3,6+7}  
1 hit  
4 miss: set 0 {0+1,4+5}; set 1 {2+3,6+7}  
7 hit  
10 miss

## quiz Q5

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

when  $i = 1$ : B index 1

when  $i = 2$ : B index 2, 2

when  $i = 3$ : B index 3, 3, 3

...

only first access to B for each  $i$  should be miss

first access only miss if not in same block as element of B from prior  $i$

$M$  possible  $i$

## quiz Q5 part 2

when  $i = 1$ : B index 1

when  $i = 2$ : B index 2, 2

when  $i = 3$ : B index 3, 3, 3

...

only first access to B for each  $i$  should be miss

first access only miss if not in same block as element of B from prior  $i$

$N$  possible  $i$

$1/4$  chance of  $i$  and  $i + 1$  being in different blocks

total misses  $N/4$



## quiz Q6

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

when  $i = 1$ : D index 1

when  $i = 2$ : D index 2,  $N+2$

when  $i = 3$ : D index 3,  $N+3$ ,  $2N+3$

...

when  $i = K$ : D index  $K$ ,  $N+K$ ,  $2N+K$ , ... $(K-1)N+K$

## quiz Q6 part 2

when  $i = 1$ : D index 1

when  $i = 2$ : D index 2,  $N+2$

when  $i = 3$ : D index 3,  $N+3$ ,  $2N+3$

...

when  $i = K$ : D index  $K$ ,  $N+K$ ,  $2N+K$ , ... $(K-1)N+K$

...

once  $i$  gets big enough, accessing lots of elements in inner loop  
once  $i$  gets big enough, not access same block without lots of accesses  
in between

so, except every access to D to be miss once once  $i$  big enough

total number of accesses to D is about  $N(N-1)/2 \approx N^2/2$

since  $N$  is large compared to cache, except most of them to be

## locality exercise (2)

```
/* version 2 */  
for (int i = 0; i < N; ++i)  
    for (int j = 0; j < N; ++j)  
        A[i] += B[j] * C[i * N + j]
```

```
/* version 3 */  
for (int ii = 0; ii < N; ii += 32)  
    for (int jj = 0; jj < N; jj += 32)  
        for (int i = ii; i < ii + 32; ++i)  
            for (int j = jj; j < jj + 32; ++j)  
                A[i] += B[j] * C[i * N + j];
```

exercise: which has better temporal locality in A? in B? in C?  
how about spatial locality?

## a transformation

```
for (int k = 0; k < N; k += 1)
    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];
```

---

```
for (int kk = 0; kk < N; kk += 2)
    for (int k = kk; k < kk + 2; ++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];
```

split the loop over  $k$  — should be exactly the same  
(assuming even  $N$ )

## a transformation

```
for (int k = 0; k < N; k += 1)
    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];
```

---

```
for (int kk = 0; kk < N; kk += 2)
    for (int k = kk; k < kk + 2; ++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];
```

split the loop over  $k$  — should be exactly the same  
(assuming even  $N$ )

## simple blocking

```
for (int kk = 0; kk < N; kk += 2)
  /* was here: for (int k = kk; k < kk + 2; ++k) */
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      /* load Aik, Aik+1 into cache and process: */
      for (int k = kk; k < kk + 2; ++k)
        C[i*N+j] += A[i*N+k] * B[k*N+j];
```

now **reorder** split loop — same calculations

## simple blocking

```
for (int kk = 0; kk < N; kk += 2)
    /* was here: for (int k = kk; k < kk + 2; ++k) */
    for (int i = 0; i < N; ++i)
        for (int j = 0; j < N; ++j)
            /* load Aik, Aik+1 into cache and process: */
            for (int k = kk; k < kk + 2; ++k)
                C[i*N+j] += A[i*N+k] * B[k*N+j];
```

now **reorder** split loop — same calculations

now handle  $B_{ij}$  for  $k + 1$  right after  $B_{ij}$  for  $k$

(previously:  $B_{i,j+1}$  for  $k$  right after  $B_{ij}$  for  $k$ )

## simple blocking

```
for (int kk = 0; kk < N; kk += 2)
  /* was here: for (int k = kk; k < kk + 2; ++k) */
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      /* load Aik, Aik+1 into cache and process: */
      for (int k = kk; k < kk + 2; ++k)
        C[i*N+j] += A[i*N+k] * B[k*N+j];
```

now **reorder** split loop — same calculations

now handle  $B_{ij}$  for  $k + 1$  right after  $B_{ij}$  for  $k$

(previously:  $B_{i,j+1}$  for  $k$  right after  $B_{ij}$  for  $k$ )



## simple blocking – expanded

```
for (int kk = 0; kk < N; kk += 2) {  
    for (int i = 0; i < N; ++i) {  
        for (int j = 0; j < N; ++j) {  
            /* process a "block" of 2 k values: */  
            C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];  
            C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];  
        }  
    }  
}
```

## simple blocking – expanded

```
for (int kk = 0; kk < N; kk += 2) {  
    for (int i = 0; i < N; ++i) {  
        for (int j = 0; j < N; ++j) {  
            /* process a "block" of 2 k values: */  
            C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];  
            C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];  
        }  
    }  
}
```

Temporal locality in  $C_{ij}$ s

## simple blocking – expanded

```
for (int kk = 0; kk < N; kk += 2) {  
    for (int i = 0; i < N; ++i) {  
        for (int j = 0; j < N; ++j) {  
            /* process a "block" of 2 k values: */  
            C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];  
            C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];  
        }  
    }  
}
```

More spatial locality in  $A_{ik}$

## simple blocking – expanded

```
for (int kk = 0; kk < N; kk += 2) {  
    for (int i = 0; i < N; ++i) {  
        for (int j = 0; j < N; ++j) {  
            /* process a "block" of 2 k values: */  
            C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];  
            C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];  
        }  
    }  
}
```

Still have good spatial locality in  $B_{kj}$ ,  $C_{ij}$

## recall: counting misses (kij-order)

```
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];
```

for A: about 1 misses per j-loop

total misses:  $N^2$

for B: about  $N \div \text{block size}$  miss per j-loop

total misses:  $N^3 \div \text{block size}$

for C: about  $N \div \text{block size}$  miss per j-loop

total misses:  $N^3 \div \text{block size}$

# counting misses for A (1)

```
for (int kk = 0; kk < N; kk += 2)
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];
      C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];
    }
```

access pattern for A:

$A[0*N+0]$ ,  $A[0*N+1]$ ,  $A[0*N+0]$ ,  $A[0*N+1]$  ... (repeats N times)

$A[1*N+0]$ ,  $A[1*N+1]$ ,  $A[1*N+0]$ ,  $A[1*N+1]$  ... (repeats N times)

...

...

# counting misses for A (1)

```
for (int kk = 0; kk < N; kk += 2)
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];
      C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];
    }
```

access pattern for A:

$A[0*N+0]$ ,  $A[0*N+1]$ ,  $A[0*N+0]$ ,  $A[0*N+1]$  ... (repeats N times)

$A[1*N+0]$ ,  $A[1*N+1]$ ,  $A[1*N+0]$ ,  $A[1*N+1]$  ... (repeats N times)

...

$A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$ ,  $A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$  ...

$A[0*N+2]$ ,  $A[0*N+3]$ ,  $A[0*N+2]$ ,  $A[0*N+3]$  ...

...

# counting misses for A (1)

```
for (int kk = 0; kk < N; kk += 2)
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];
      C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];
    }
```

access pattern for A:

$A[0*N+0]$ ,  $A[0*N+1]$ ,  $A[0*N+0]$ ,  $A[0*N+1]$  ... (repeats N times)

$A[1*N+0]$ ,  $A[1*N+1]$ ,  $A[1*N+0]$ ,  $A[1*N+1]$  ... (repeats N times)

...

$A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$ ,  $A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$  ...

$A[0*N+2]$ ,  $A[0*N+3]$ ,  $A[0*N+2]$ ,  $A[0*N+3]$  ...

...



## counting misses for A (2)

$A[0*N+0]$ ,  $A[0*N+1]$ ,  $A[0*N+0]$ ,  $A[0*N+1]$  ... (repeats N times)

$A[1*N+0]$ ,  $A[1*N+1]$ ,  $A[1*N+0]$ ,  $A[1*N+1]$  ... (repeats N times)

...

...

## counting misses for A (2)

$A[0*N+0]$ ,  $A[0*N+1]$ ,  $A[0*N+0]$ ,  $A[0*N+1]$  ... (repeats N times)

$A[1*N+0]$ ,  $A[1*N+1]$ ,  $A[1*N+0]$ ,  $A[1*N+1]$  ... (repeats N times)

...

$A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$ ,  $A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$  ...

$A[0*N+2]$ ,  $A[0*N+3]$ ,  $A[0*N+2]$ ,  $A[0*N+3]$  ...

...

likely cache misses: only first iterations of  $j$  loop

how many cache misses per iteration? usually one

$A[0*N+0]$  and  $A[0*N+1]$  usually in same cache block

## counting misses for A (2)

$A[0*N+0]$ ,  $A[0*N+1]$ ,  $A[0*N+0]$ ,  $A[0*N+1]$  ... (repeats  $N$  times)

$A[1*N+0]$ ,  $A[1*N+1]$ ,  $A[1*N+0]$ ,  $A[1*N+1]$  ... (repeats  $N$  times)

...

$A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$ ,  $A[(N-1)*N+0]$ ,  $A[(N-1)*N+1]$  ...

$A[0*N+2]$ ,  $A[0*N+3]$ ,  $A[0*N+2]$ ,  $A[0*N+3]$  ...

...

likely cache misses: only first iterations of  $j$  loop

how many cache misses per iteration? usually one

$A[0*N+0]$  and  $A[0*N+1]$  usually in same cache block

about  $\frac{N}{2} \cdot N$  misses total

## counting misses for B (1)

```
for (int kk = 0; kk < N; kk += 2)
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];
      C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];
    }
```

access pattern for B:

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

$B[2*N+0]$ ,  $B[3*N+0]$ , ...  $B[2*N+(N-1)]$ ,  $B[3*N+(N-1)]$

$B[4*N+0]$ ,  $B[5*N+0]$ , ...  $B[4*N+(N-1)]$ ,  $B[5*N+(N-1)]$

...

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

...

## counting misses for B (2)

access pattern for B:

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

$B[2*N+0]$ ,  $B[3*N+0]$ , ...  $B[2*N+(N-1)]$ ,  $B[3*N+(N-1)]$

$B[4*N+0]$ ,  $B[5*N+0]$ , ...  $B[4*N+(N-1)]$ ,  $B[5*N+(N-1)]$

...

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

...

## counting misses for B (2)

access pattern for B:

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

$B[2*N+0]$ ,  $B[3*N+0]$ , ...  $B[2*N+(N-1)]$ ,  $B[3*N+(N-1)]$

$B[4*N+0]$ ,  $B[5*N+0]$ , ...  $B[4*N+(N-1)]$ ,  $B[5*N+(N-1)]$

...

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

...

likely cache misses: any access, each time

## counting misses for B (2)

access pattern for B:

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

$B[2*N+0]$ ,  $B[3*N+0]$ , ...  $B[2*N+(N-1)]$ ,  $B[3*N+(N-1)]$

$B[4*N+0]$ ,  $B[5*N+0]$ , ...  $B[4*N+(N-1)]$ ,  $B[5*N+(N-1)]$

...

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

...

likely cache misses: any access, each time

how many cache misses per iteration? equal to # cache blocks in 2 rows

## counting misses for B (2)

access pattern for B:

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

$B[2*N+0]$ ,  $B[3*N+0]$ , ...  $B[2*N+(N-1)]$ ,  $B[3*N+(N-1)]$

$B[4*N+0]$ ,  $B[5*N+0]$ , ...  $B[4*N+(N-1)]$ ,  $B[5*N+(N-1)]$

...

$B[0*N+0]$ ,  $B[1*N+0]$ , ...  $B[0*N+(N-1)]$ ,  $B[1*N+(N-1)]$

...

likely cache misses: any access, each time

how many cache misses per iteration? equal to  $\#$  cache blocks in 2 rows

about  $\frac{N}{2} \cdot N \cdot \frac{2N}{\text{block size}} = N^3 \div \text{block size}$  misses



## simple blocking – counting misses

```
for (int kk = 0; kk < N; kk += 2)
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];
      C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];
    }
```

$\frac{N}{2} \cdot N$  j-loop executions and (assuming  $N$  large):

about 1 misses from  $A$  per j-loop

$N^2/2$  total misses (before blocking:  $N^2$ )

about  $2N \div$  block size misses from  $B$  per j-loop

$N^3 \div$  block size total misses (same as before blocking)

about  $N \div$  block size misses from  $C$  per j-loop

$N^3 \div (2 \cdot \text{block size})$  total misses (before:  $N^3 \div$  block size)

# simple blocking – counting misses

```
for (int kk = 0; kk < N; kk += 2)
  for (int i = 0; i < N; i += 1)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+kk+0] * B[(kk+0)*N+j];
      C[i*N+j] += A[i*N+kk+1] * B[(kk+1)*N+j];
    }
```

$\frac{N}{2} \cdot N$  j-loop executions and (assuming  $N$  large):

about 1 misses from  $A$  per j-loop

$N^2/2$  total misses (before blocking:  $N^2$ )

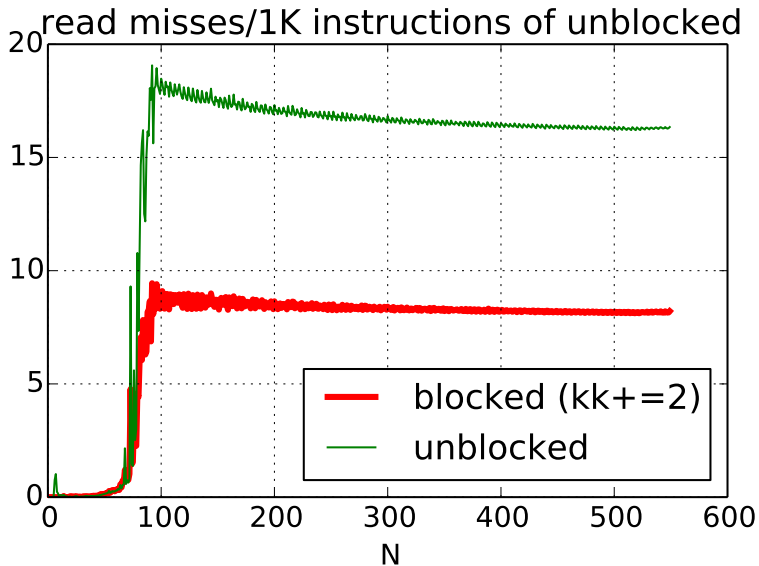
about  $2N \div$  block size misses from  $B$  per j-loop

$N^3 \div$  block size total misses (same as before blocking)

about  $N \div$  block size misses from  $C$  per j-loop

$N^3 \div (2 \cdot \text{block size})$  total misses (before:  $N^3 \div$  block size)

# improvement in read misses



## simple blocking (2)

same thing for  $i$  in addition to  $k$ ?

```
for (int kk = 0; kk < N; kk += 2) {
    for (int ii = 0; ii < N; ii += 2) {
        for (int j = 0; j < N; ++j) {
            /* process a "block": */
            for (int k = kk; k < kk + 2; ++k)
                for (int i = 0; i < ii + 2; ++i)
                    C[i*N+j] += A[i*N+k] * B[k*N+j];
        }
    }
}
```

## simple blocking — locality

```
for (int k = 0; k < N; k += 2) {  
  for (int i = 0; i < N; i += 2) {  
    /* load a block around Aik */  
    for (int j = 0; j < N; ++j) {  
      /* process a "block": */  
       $C_{i+0,j} += A_{i+0,k+0} * B_{k+0,j}$   
       $C_{i+0,j} += A_{i+0,k+1} * B_{k+1,j}$   
       $C_{i+1,j} += A_{i+1,k+0} * B_{k+0,j}$   
       $C_{i+1,j} += A_{i+1,k+1} * B_{k+1,j}$   
    }  
  }  
}
```

## simple blocking — locality

```
for (int k = 0; k < N; k += 2) {  
  for (int i = 0; i < N; i += 2) {  
    /* load a block around Aik */  
    for (int j = 0; j < N; ++j) {  
      /* process a "block": */  
      Ci+0,j += Ai+0,k+0 * Bk+0,j  
      Ci+0,j += Ai+0,k+1 * Bk+1,j  
      Ci+1,j += Ai+1,k+0 * Bk+0,j  
      Ci+1,j += Ai+1,k+1 * Bk+1,j  
    }  
  }  
}
```

now: more temporal locality in  $B$

previously: access  $B_{kj}$ , then don't use it again for a long time

## simple blocking — counting misses for A

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; i += 2)
    for (int j = 0; j < N; ++j) {
      Ci+0,j += Ai+0,k+0 * Bk+0,j
      Ci+0,j += Ai+0,k+1 * Bk+1,j
      Ci+1,j += Ai+1,k+0 * Bk+0,j
      Ci+1,j += Ai+1,k+1 * Bk+1,j
    }
```

$\frac{N}{2} \cdot \frac{N}{2}$  iterations of  $j$  loop

likely 2 misses per loop with  $A$  (2 cache blocks)

total misses:  $\frac{N^2}{2}$  (same as only blocking in  $K$ )

## simple blocking — counting misses for B

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; i += 2)
    for (int j = 0; j < N; ++j) {
      Ci+0,j += Ai+0,k+0 * Bk+0,j
      Ci+0,j += Ai+0,k+1 * Bk+1,j
      Ci+1,j += Ai+1,k+0 * Bk+0,j
      Ci+1,j += Ai+1,k+1 * Bk+1,j
    }
```

$\frac{N}{2} \cdot \frac{N}{2}$  iterations of  $j$  loop

likely  $2 \div$  block size misses per iteration with  $B$

total misses:  $\frac{N^3}{2 \cdot \text{block size}}$  (before:  $\frac{N^3}{\text{block size}}$ )



# simple blocking — counting misses for C

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; i += 2)
    for (int j = 0; j < N; ++j) {
      Ci+0,j += Ai+0,k+0 * Bk+0,j
      Ci+0,j += Ai+0,k+1 * Bk+1,j
      Ci+1,j += Ai+1,k+0 * Bk+0,j
      Ci+1,j += Ai+1,k+1 * Bk+1,j
    }
```

$\frac{N}{2} \cdot \frac{N}{2}$  iterations of  $j$  loop

likely  $\frac{2}{\text{block size}}$  misses per iteration with  $C$

total misses:  $\frac{N^3}{2 \cdot \text{block size}}$  (same as blocking only in K)

# simple blocking — counting misses (total)

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; i += 2)
    for (int j = 0; j < N; ++j) {
      Ci+0,j += Ai+0,k+0 * Bk+0,j
      Ci+0,j += Ai+0,k+1 * Bk+1,j
      Ci+1,j += Ai+1,k+0 * Bk+0,j
      Ci+1,j += Ai+1,k+1 * Bk+1,j
    }
```

before:

$$A: \frac{N^2}{2}; \quad B: \frac{N^3}{1 \cdot \text{block size}}; \quad C: \frac{N^3}{1 \cdot \text{block size}}$$

after:

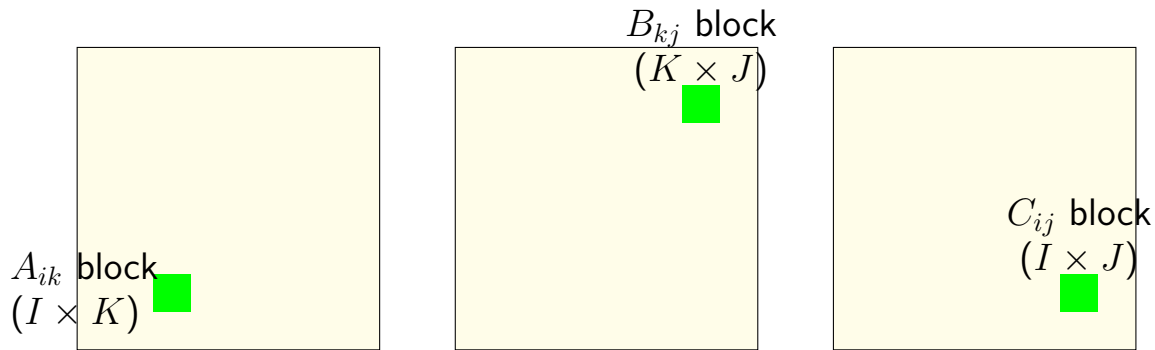
$$A: \frac{N^2}{2}; \quad B: \frac{N^3}{2 \cdot \text{block size}}; \quad C: \frac{N^3}{2 \cdot \text{block size}}$$

## generalizing: divide and conquer

```
partial_matrixmultiply(float *A, float *B, float *C
                        int startI, int endI, ...) {
    for (int i = startI; i < endI; ++i) {
        for (int j = startJ; j < endJ; ++j) {
            for (int k = startK; k < endK; ++k) {
                ...
            }
        }
    }
}

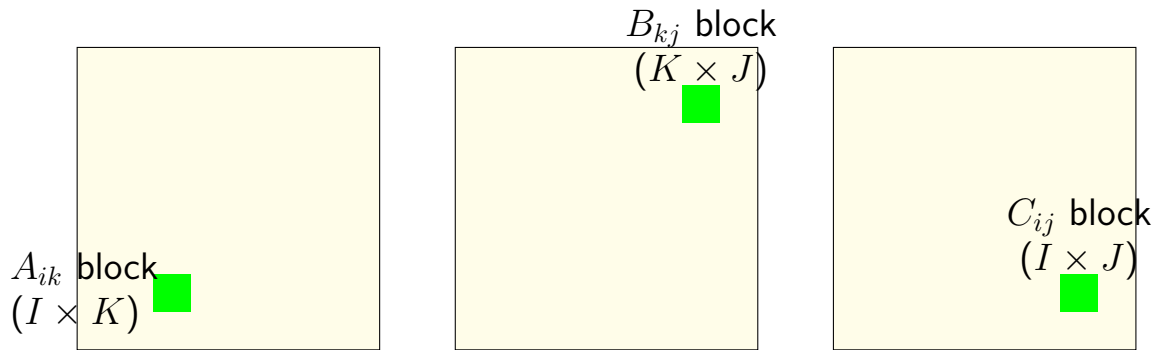
matrix_multiply(float *A, float *B, float *C, int N) {
    for (int ii = 0; ii < N; ii += BLOCK_I)
        for (int jj = 0; jj < N; jj += BLOCK_J)
            for (int kk = 0; kk < N; kk += BLOCK_K)
                ...
                /* do everything for segment of A, B, C
                   that fits in cache! */
                partial_matmul(A, B, C,
                               ii, ii + BLOCK_I, jj, jj + BLOCK_J,
                               kk, kk + BLOCK_K)
}
```

# array usage: matrix block $C_{ij} += A_{ik} \cdot B_{kj}$



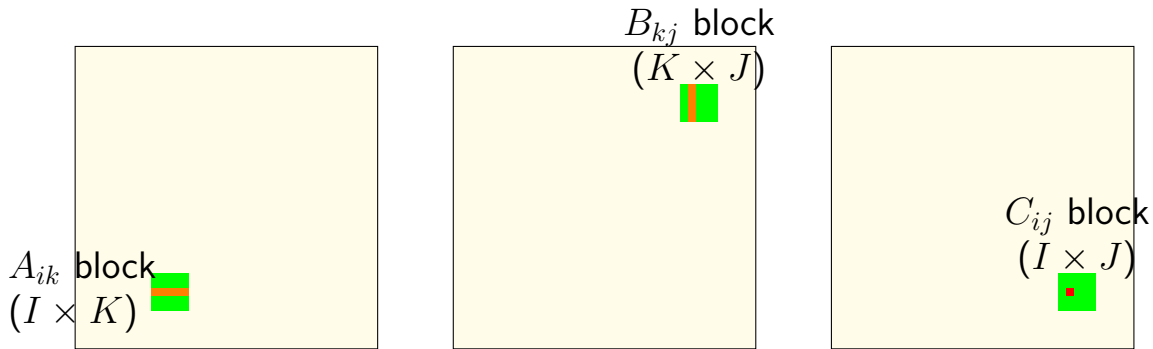
inner loops work on “matrix block” of A, B, C  
rather than rows of some, little blocks of others  
blocks fit into cache (b/c we choose  $I, K, J$ )  
where previous rows might not

# array usage: matrix block $C_{ij} += A_{ik} \cdot B_{kj}$



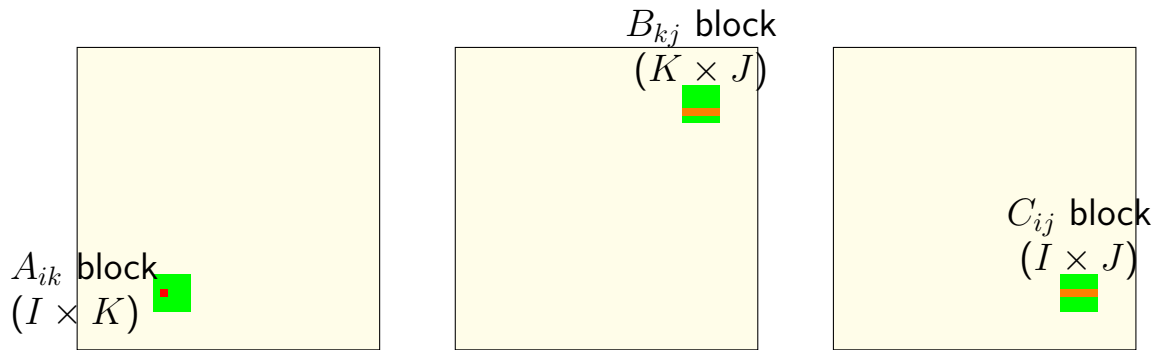
now (versus loop ordering example)  
some spatial locality in A, B, and C  
some temporal locality in A, B, and C

# array usage: matrix block $C_{ij} += A_{ik} \cdot B_{kj}$



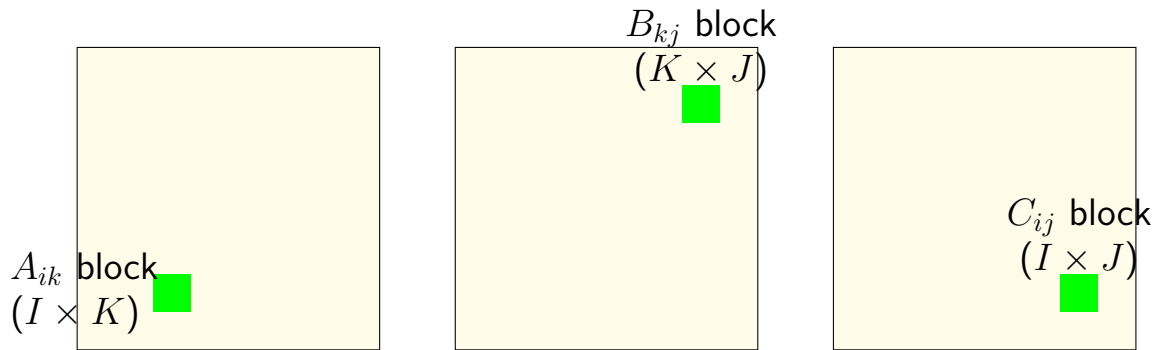
$C_{ij}$  calculation uses strips from  $A, B$   
 $K$  calculations for one cache miss  
good temporal locality!

# array usage: matrix block $C_{ij} += A_{ik} \cdot B_{kj}$



$A_{ik}$  used with entire strip of  $B$   $J$  calculations for one cache miss  
good temporal locality!

# array usage: matrix block $C_{ij} += A_{ik} \cdot B_{kj}$



(approx.)  $KIJ$  fully cached calculations  
for  $KI + IJ + KJ$  values need to be loaded per “matrix block”  
(assuming everything stays in cache)



# cache blocking efficiency

for each of  $N^3/IJK$  matrix blocks:

load  $I \times K$  elements of  $A_{ik}$ :

$\approx IK \div$  block size misses per matrix block

$\approx N^3/(J \cdot \text{blocksize})$  misses total

load  $K \times J$  elements of  $B_{kj}$ :

$\approx N^3/(I \cdot \text{blocksize})$  misses total

load  $I \times J$  elements of  $C_{ij}$ :

$\approx N^3/(K \cdot \text{blocksize})$  misses total

bigger blocks — more work per load!

catch:  $IK + KJ + IJ$  elements must fit in cache  
otherwise estimates above don't work

# cache blocking rule of thumb

fill the **most of the cache with useful data**

and do as much work as possible from that

example: my desktop 32KB L1 cache

$I = J = K = 48$  uses  $48^2 \times 3$  elements, or 27KB.

assumption: conflict misses aren't important

## exercise: miss estimating (3)

```
for (int kk = 0; kk < 1000; kk += 10)
  for (int jj = 0; jj < 1000; jj += 10)
    for (int i = 0; i < 1000; i += 1)
      for (int j = jj; j < jj+10; j += 1)
        for (int k = kk; k < kk + 10; k += 1)
          A[k*N+j] += B[i*N+j];
```

assuming: 4 elements per block

assuming: cache not close to big enough to hold 1K elements, but big enough to hold 500 or so

estimate: *approximately* how many misses for A, B?

hint 1: part of A, B loaded in two inner-most loops only needs to be loaded once

# loop ordering compromises

loop ordering forces compromises:

```
for k: for i: for j: c[i,j] += a[i,k] * b[j,k]
```

perfect temporal locality in  $a[i,k]$

bad temporal locality for  $c[i,j]$ ,  $b[j,k]$

perfect spatial locality in  $c[i,j]$

bad spatial locality in  $b[j,k]$ ,  $a[i,k]$

# loop ordering compromises

loop ordering forces compromises:

```
for k: for i: for j: c[i,j] += a[i,k] * b[j,k]
```

perfect temporal locality in  $a[i,k]$

bad temporal locality for  $c[i,j]$ ,  $b[j,k]$

perfect spatial locality in  $c[i,j]$

bad spatial locality in  $b[j,k]$ ,  $a[i,k]$

cache blocking: work on blocks rather than rows/columns  
have some temporal, spatial locality in everything

# cache blocking pattern

no perfect loop order? work on rectangular matrix blocks

size amount used in inner loops based on cache size

in practice:

- test performance to determine 'size' of blocks

# sum array ASM (gcc 8.3 -Os)

```
long sum_array(long *values, int size) {  
    long sum = 0;  
    for (int i = 0; i < size; ++i) {  
        sum += values[i];  
    }  
    return sum;  
}
```

sum\_array:

```
    xorl    %edx, %edx           // i = 0  
    xorl    %eax, %eax         // sum = 0
```

loop:

```
    cmpq    %edx, %esi         // if (i < size) break  
    jle    endOfLoop          // sum += values[i]  
    addq    (%rsi,%rdx,8), %rax // i += 1  
    incq    %rdx  
    jmp     loop
```

endOfLoop:

```
    ret
```

# loop unrolling (ASM)

```
loop:
    cmpl    %edx, %esi
    jle     endOfLoop           // if (i < size) break
    addq   (%rdi,%rdx,8), %rax  // sum += values[i]
    incq   %rdx                 // i += 1
    jmp    loop
endOfLoop:
```

---

```
loop:
    cmpl    %edx, %esi
    jle     endOfLoop           // if (i < size) break
    addq   (%rdi,%rdx,8), %rax  // sum += values[i]
    addq   8(%rdi,%rdx,8), %rax // sum += values[i+1]
    addq   $2, %rdx             // i += 2
    jmp    loop
    // plus handle leftover?
endOfLoop:
```



# loop unrolling (ASM)

loop:

```
    cmpl    %edx, %esi
    jle     endOfLoop           // if (i < size) break
    addq   (%rdi,%rdx,8), %rax  // sum += values[i]
    incq   %rdx                 // i += 1
    jmp    loop
```

endOfLoop:

size iterations  $\times$  5 instructions

---

loop:

```
    cmpl    %edx, %esi
    jle     endOfLoop           // if (i < size) break
    addq   (%rdi,%rdx,8), %rax  // sum += values[i]
    addq   8(%rdi,%rdx,8), %rax // sum += values[i+1]
    addq   $2, %rdx             // i += 2
    jmp    loop
    // plus handle leftover?
```

endOfLoop:

size  $\div$  2 iterations  $\times$  6 instructions

# loop unrolling (C)

```
for (int i = 0; i < N; ++i)
    sum += A[i];
```

---

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

## more loop unrolling (C)

```
int i;
for (i = 0; i + 4 <= N; i += 4) {
    sum += A[i];
    sum += A[i+1];
    sum += A[i+2];
    sum += A[i+3];
}
// handle leftover, if needed
for (; i < N; i += 1)
    sum += A[i];
```

# 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

# loop unrolling on MM

original code:

```
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];
    }
```

---

loop unrolling in  $j$  loop (not cache blocking)

```
for (int k = 0; k < N; ++k)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; j += 2) {
      C[i*N+j] += A[i*N+k] * B[k*N+j];
      C[i*N+j+1] += A[i*N+k] * B[k*N+j+1];
    }
```

# loop unrolling on MM

original code:

```
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];
    }
```

access order:

```
k=i=j=0: C[0*N+0], A[0*N+0], B[0*N+0]
k=i=0, j=1: C[0*N+1], A[0*N+0], B[0*N+1]
k=i=0, j=2: C[0*N+2], A[0*N+0], B[0*N+2]
k=i=0, j=3: C[0*N+3], A[0*N+0], B[0*N+3]
...
```

---

loop unrolling in  $j$  loop (not cache blocking)

```
for (int k = 0; k < N; ++k)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; j += 2) {
      C[i*N+j] += A[i*N+k] * B[k*N+j];
      C[i*N+j+1] += A[i*N+k] * B[k*N+j+1];
    }
```

access order:

```
k=i=j=0: C[0*N+0], A[0*N+0], B[0*N+0]
          C[0*N+1], A[0*N+0], B[0*N+1]
k=i=0, j=2: C[0*N+2], A[0*N+0], B[0*N+2]
          C[0*N+3], A[0*N+0], B[0*N+3]
...
```

# loop unrolling on MM

original code:

```
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];
    }
```

access order:

```
k=i=j=0: C[0*N+0], A[0*N+0], B[0*N+0]
k=i=0, j=1: C[0*N+1], A[0*N+0], B[0*N+1]
k=i=0, j=2: C[0*N+2], A[0*N+0], B[0*N+2]
k=i=0, j=3: C[0*N+3], A[0*N+0], B[0*N+3]
...
```

---

loop unrolling in  $j$  loop (not cache blocking)

```
for (int k = 0; k < N; ++k)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; j += 2) {
      C[i*N+j] += A[i*N+k] * B[k*N+j];
      C[i*N+j+1] += A[i*N+k] * B[k*N+j+1];
    }
```

access order:

```
k=i=j=0: C[0*N+0], A[0*N+0], B[0*N+0]
          C[0*N+1], A[0*N+0], B[0*N+1]
k=i=0, j=2: C[0*N+2], A[0*N+0], B[0*N+2]
          C[0*N+3], A[0*N+0], B[0*N+3]
...
```

# loop unrolling on MM

original code:

```
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];
    }
```

access order:

```
k=i=j=0: C[0*N+0], A[0*N+0], B[0*N+0]
k=i=0, j=1: C[0*N+1], A[0*N+0], B[0*N+1]
k=i=0, j=2: C[0*N+2], A[0*N+0], B[0*N+2]
k=i=0, j=3: C[0*N+3], A[0*N+0], B[0*N+3]
...
```

---

loop unrolling in  $j$  loop (not cache blocking)

```
for (int k = 0; k < N; ++k)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; j += 2) {
      C[i*N+j] += A[i*N+k] * B[k*N+j];
      C[i*N+j+1] += A[i*N+k] * B[k*N+j+1];
    }
```

access order:

```
k=i=j=0: C[0*N+0], A[0*N+0], B[0*N+0]
          C[0*N+1], A[0*N+0], B[0*N+1]
k=i=0, j=2: C[0*N+2], A[0*N+0], B[0*N+2]
            C[0*N+3], A[0*N+0], B[0*N+3]
...
```



# partial cache blocking in MM

original code:

```
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];
    }
```

---

(incomplete) cache blocking with only  $k$ :

**changes locality v. original (order of A, B, C accesses)**

```
for (int kk = 0; kk < N; kk += BLOCK_SIZE)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      for (int k = kk; k < kk + BLOCK_SIZE; ++k)
        C[i*N+j] += A[i*N+k+0] * B[k*N+j];
```

# loop unrolling v cache blocking (0)

cache blocking for  $k$  only: (with teeny 1 by 1 by 2 matrix blocks)

**changes locality v. original (order of A, B, C accesses)**

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

---

with loop unrolling added afterwards:

**same order of A, B, C accesses as above**

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+k+0] * B[(k+0)*N+j];
      C[i*N+j] += A[i*N+k+1] * B[(k+1)*N+j];
    }
```

# loop unrolling v cache blocking (0)

cache blocking for  $k$  only: (with teeny 1 by 1 by 2 matrix blocks)

**changes locality v. original (order of A, B, C accesses)**

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

---

with loop unrolling added afterwards:

**same order of A, B, C accesses as above**

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+k+0] * B[(k+0)*N+j];
      C[i*N+j] += A[i*N+k+1] * B[(k+1)*N+j];
    }
```

# loop unrolling v cache blocking

cache blocking for  $k$  only (1x1x2 blocks) *and* then loop unrolling

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j) {
      C[i*N+j] += A[i*N+k+0] * B[(k+0)*N+j];
      C[i*N+j] += A[i*N+k+1] * B[(k+1)*N+j];
    }
```

---

versus pretty useless loop unrolling in  $k$ -loop

**same order of A, B, C accesses as original**

```
for (int k = 0; k < N; k += 2) {
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      C[i*N+j] += A[i*N+k+0] * B[(k+0)*N+j];
  for (int i = 0; i < N; ++i)
    for (int j = 0; j < N; ++j)
      C[i*N+j] += A[i*N+k+1] * B[(k+1)*N+j];
}
```

# loop unrolling v cache blocking (1)

cache blocking for  $k,i$  only: (1 by 2 by 2 matrix blocks)

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

---

cache blocking for  $k,i$  and loop unrolling for  $i$ :

```
for (int k = 0; k < N; k += 2)
  for (int i = 0; i < N; i += 2)
    for (int j = 0; j < N; ++j)
      for(int kk = k; kk < k + 2; ++kk) {
        C[(i+0)*N+j] += A[(i+0)*N+kk] * B[(kk)*N+j];
        C[(i+1)*N+j] += A[(i+1)*N+kk] * B[(kk)*N+j];
      }
}
```

# interlude: real CPUs

modern CPUs:

execute **multiple instructions at once**

execute instructions **out of order** — whenever **values available**

# beyond pipelining: multiple issue

start **more than one instruction/cycle**

multiple parallel pipelines; many-input/output register file

**hazard handling much more complex**

	<i>cycle #</i>	0	1	2	3	4	5	6	7	8
addq %r8, %r9		F	D	E	M	W				
subq %r10, %r11		F	D	E	M	W				
xorq %r9, %r11			F	D	E	M	W			
subq %r10, %rbx			F	D	E	M	W			
...										

# beyond pipelining: out-of-order

find **later instructions to do** instead of stalling

lists of available instructions in pipeline registers  
take any instruction with available values

provide **illusion that work is still done in order**  
much more complicated hazard handling logic

	<i>cycle #</i>	0	1	2	3	4	5	6	7	8	9	10	11
<code>mrmovq 0(%rbx), %r8</code>		F	D	E	M	M	M	W	C				
<code>subq %r8, %r9</code>			F					D	E	W	C		
<code>addq %r10, %r11</code>				F	D	E	W					C	
<code>xorq %r12, %r13</code>					F	D	E	W					C
...													



# out-of-order and hazards

out-of-order execution makes hazards harder to handle

problems for forwarding:

- value in last stage may not be most up-to-date

- older value may be written back before newer value?

problems for branch prediction:

- mispredicted instructions may complete execution before squashing

which instructions to dispatch?

- how to quickly find instructions that are ready?

# out-of-order and hazards

out-of-order execution makes hazards harder to handle

problems for forwarding:

- value in last stage may not be most up-to-date

- older value may be written back before newer value?

problems for branch prediction:

- mispredicted instructions may complete execution before squashing

which instructions to dispatch?

- how to quickly find instructions that are ready?

# read-after-write examples (1)

	cycle #	0	1	2	3	4	5	6	7	8
<code>addq %r10, %r8</code>		F	D	E	M	W				
<code>addq %r11, %r8</code>			F	D	E	M	W			
<code>addq %r12, %r8</code>				F	D	E	M	W		

normal pipeline: two options for `%r8`?

choose the one from *earliest stage*

because it's from the most recent instruction

# read-after-write examples (1)

out-of-order execution:

%r8 from earliest stage might be from *delayed instruction*  
can't use same forwarding logic

addq %r11, %r8  
addq %r12, %r8

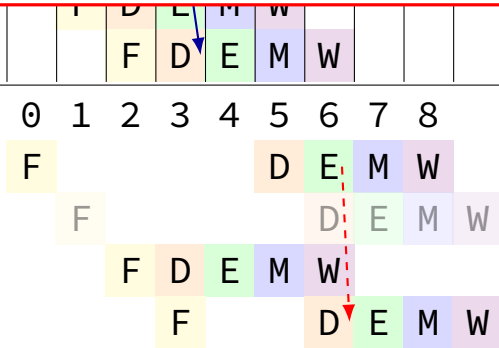
addq %r10, %r8

rmmovq %r8, (%rax)

irmovq \$100, %r8

addq %r13, %r8

cycle # 0 1 2 3 4 5 6 7 8



# register version tracking

goal: track **different versions of registers**

out-of-order execution: may compute versions at different times

only forward the **correct version**

strategy for doing this: preprocess instructions represent version info

makes forwarding, etc. lookup easier

## rewriting hazard examples (1)

addq %r10, %r8		addq %r10, %r8 <sub>v1</sub>	→	%r8 <sub>v2</sub>
addq %r11, %r8		addq %r11, %r8 <sub>v2</sub>	→	%r8 <sub>v3</sub>
addq %r12, %r8		addq %r12, %r8 <sub>v3</sub>	→	%r8 <sub>v4</sub>

---

read different version than the one written

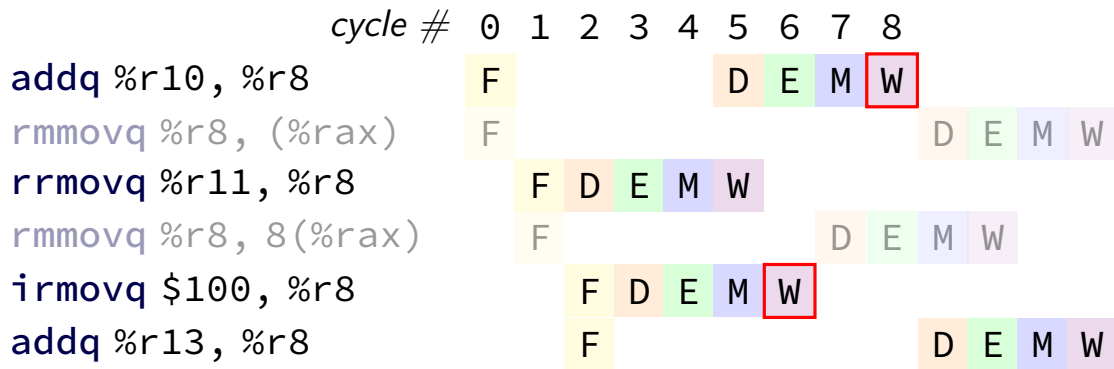
represent with three argument psuedo-instructions

forwarding a value? must match version *exactly*

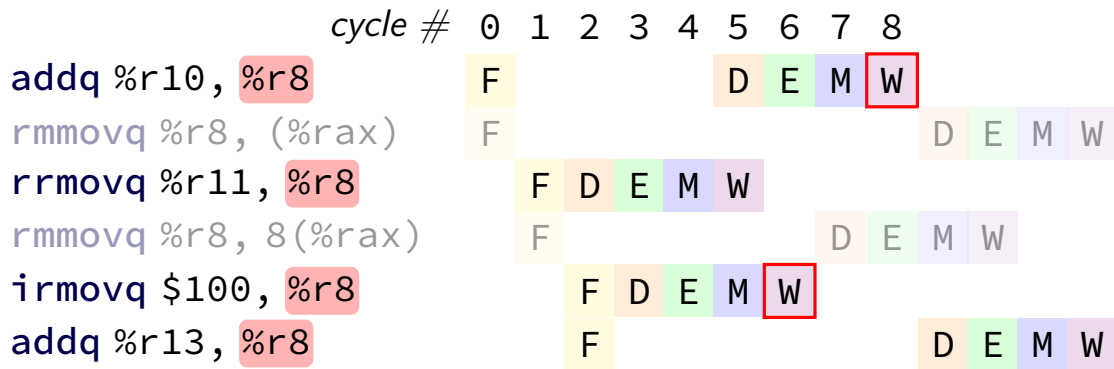
for now: version numbers

later: something simpler to implement

# write-after-write example



# write-after-write example

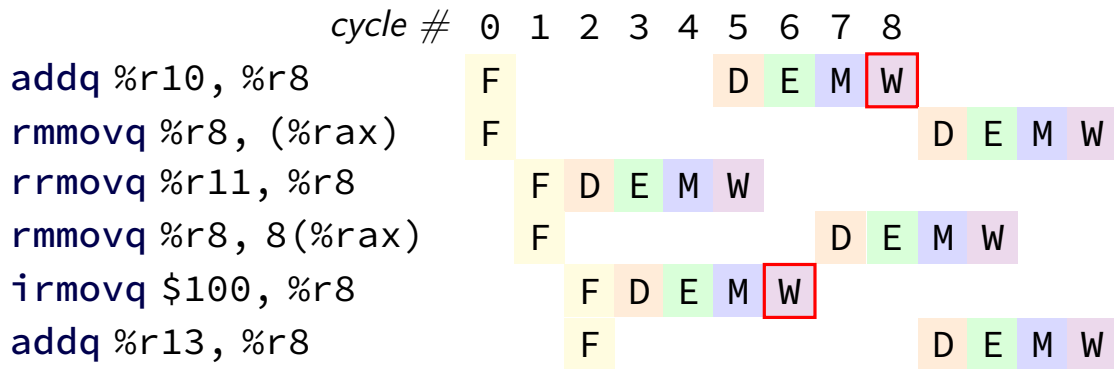


out-of-order execution:

if we don't do something, newest value could be overwritten!

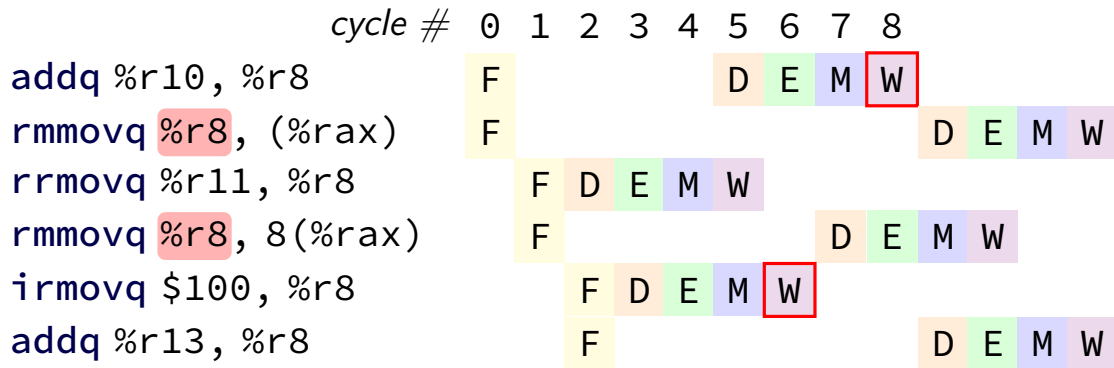


# write-after-write example



two instructions that haven't been started  
could need *different versions* of %r8!

# write-after-write example



# keeping multiple versions

for write-after-write problem: need to keep copies of multiple versions

both the new version and the old version needed by delayed instructions

for read-after-write problem: need to distinguish different versions

solution: have lots of extra registers

...and assign each version a new 'real' register

called register renaming

# register renaming

rename *architectural registers* to *physical registers*

different physical register for each version of architectural

track which physical registers are ready

compare physical register numbers to do forwarding

**backup slides**

## exercise: miss estimating (2)

```
for (int k = 0; k < 1000; k += 1)
  for (int i = 0; i < 1000; i += 1)
    for (int j = 0; j < 1000; j += 1)
      A[k*N+j] += B[i*N+j];
```

assuming: 4 elements per block

assuming: cache not close to big enough to hold 1K elements

estimate: *approximately* how many misses for  $A$ ,  $B$ ?

## simple blocking – with 3?

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

$\frac{N}{3} \cdot N$  j-loop iterations, and (assuming  $N$  large):

about 1 misses from  $A$  per j-loop iteration

$N^2/3$  total misses (before blocking:  $N^2$ )

about  $3N \div$  block size misses from  $B$  per j-loop iteration

$N^3 \div$  block size total misses (same as before)

about  $3N \div$  block size misses from  $C$  per j-loop iteration

$N^3 \div$  block size total misses (same as before)

## simple blocking – with 3?

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

$\frac{N}{3} \cdot N$  j-loop iterations, and (assuming  $N$  large):

about 1 misses from  $A$  per j-loop iteration

$N^2/3$  total misses (before blocking:  $N^2$ )

about  $3N \div$  block size misses from  $B$  per j-loop iteration

$N^3 \div$  block size total misses (same as before)

about  $3N \div$  block size misses from  $C$  per j-loop iteration

$N^3 \div$  block size total misses (same as before)



## more than 3?

can we just keep doing this increase from 3 to some large  $X$ ? ...

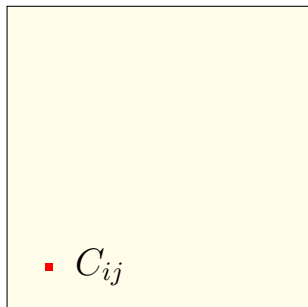
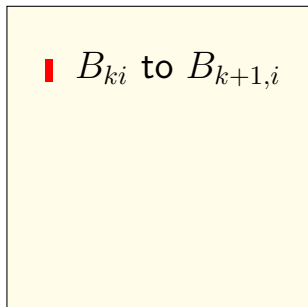
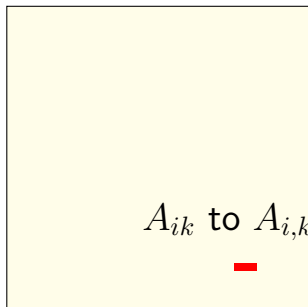
assumption:  $X$  values from A would stay in cache

$X$  too large — cache not big enough

assumption:  $X$  blocks from B would help with spatial locality

$X$  too large — evicted from cache before next iteration

## array usage (2 $k$ at a time)



for each  $kk$ :

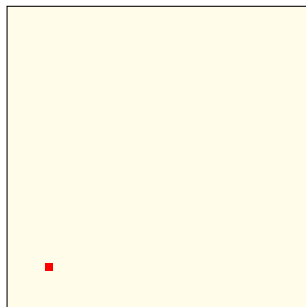
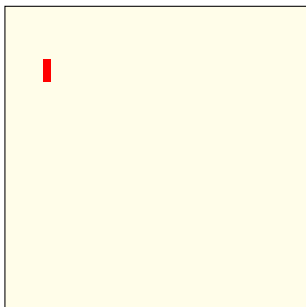
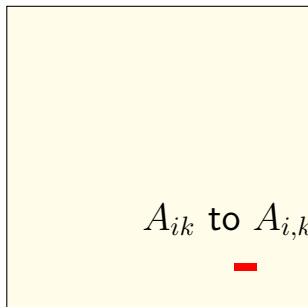
for each  $i$ :

for each  $j$ :

for  $k=kk, kk+1$ :

$$C_{ij} += A_{ik} \cdot B_{kj}$$

## array usage (2 $k$ at a time)



for each  $k$ :

for each  $i$ :

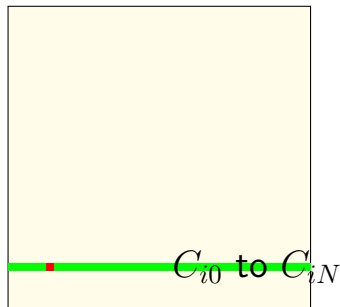
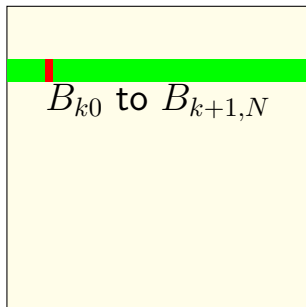
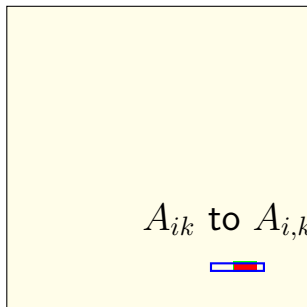
for each  $j$ :

for  $k=k, k+1$ :

$$C_{ij} += A_{ik} \cdot B_{kj}$$

within innermost loop  
good spatial locality in  $A$   
bad locality in  $B$   
good temporal locality in  $C$

## array usage (2 $k$ at a time)



for each  $kk$ :

for each  $i$ :

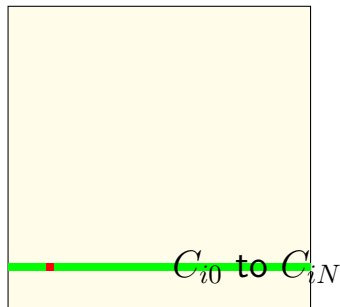
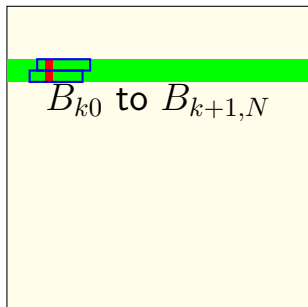
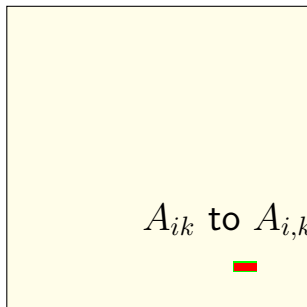
for each  $j$ :

for  $k=kk, kk+1$ :

$$C_{ij+} = A_{ik} \cdot B_{kj}$$

loop over  $j$ : better spatial locality  
over  $A$  than before;  
still good temporal locality for  $A$

## array usage (2 $k$ at a time)



for each  $k$ :

for each  $i$ :

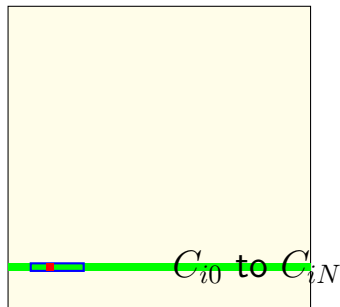
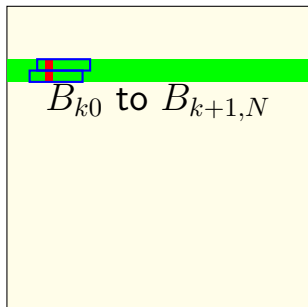
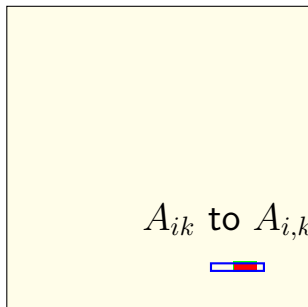
for each  $j$ :

for  $k=k, k+1$ :

$$C_{ij} += A_{ik} \cdot B_{kj}$$

loop over  $j$ : spatial locality over  $B$  is worse  
but probably not more misses  
cache needs to keep two cache blocks  
for next iter instead of one  
(probably has the space left over!)

# array usage (2 $k$ at a time)



for each  $k$ :

for each  $i$ :

for each  $j$ :

for  $k=k, k+1$ :

$$C_{ij} = A_{ik} \cdot$$

right now: only really care about  
keeping 4 cache blocks in  $j$  loop

have more than 4 cache blocks?

increasing  $k$  increment would use more of them

## exercise

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

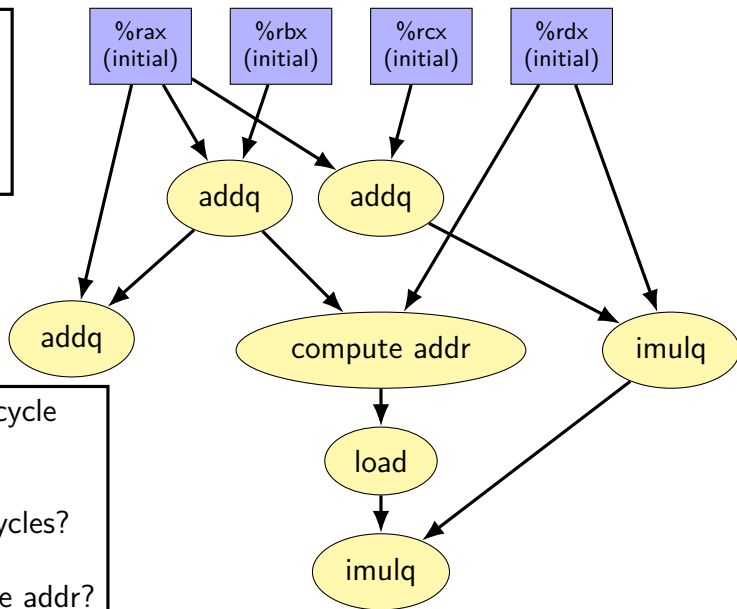
Which of the following suggests changing order of memory accesses?

```
/* version A */
for (int i = 0; i < N; ++i)
  for (int j = 0; j < N; j += 2) {
    A[i*N+j] += B[i] + C[j]
    A[i*N+j+1] += B[i] + C[j+1]
  }
```

```
/* version B */
for (int i = 0; i < N; i += 2)
  for (int j = 0; j < N; j += 2) {
    A[i*N+j] += B[i] + C[j];
    A[i*N+j+1] += B[i] + C[j+1];
    A[(i+1)*N+j] += B[i+1] + C[j];
    A[(i+1)*N+j+1] += B[i+1] + C[j+1];
  }
```

# a data flow example

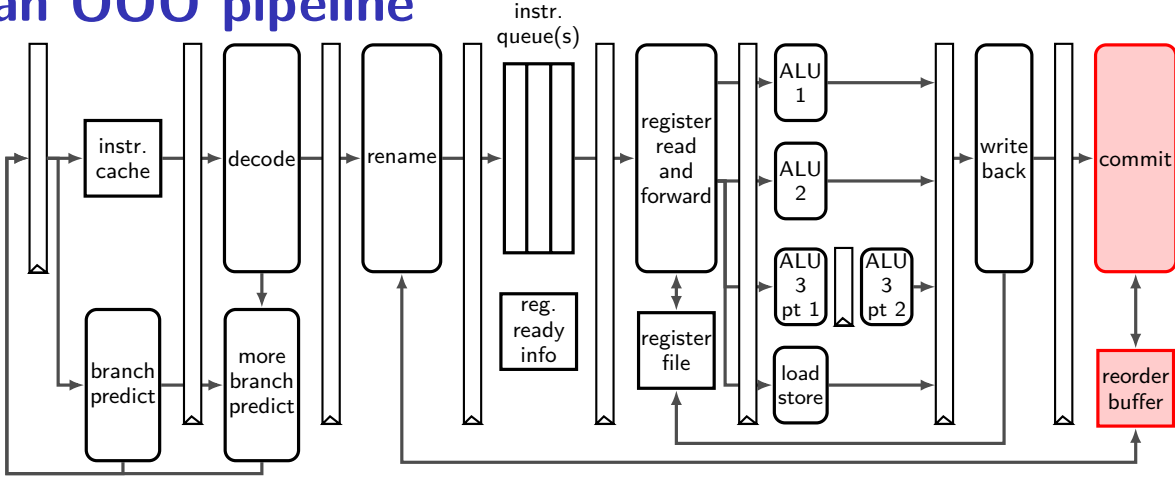
```
addq %rax, %rbx
addq %rax, %rcx
imulq %rdx, %rcx
movq (%rbx, %rdx), %r8
imulq %r8, %rcx
addq %rax, %rbx
```



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
...
...

reorder buffer (ROB)

remove here  
when committed →

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
...	...

phys → arch. reg when committed  
for committed

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

free list

%x19
%x13
...
%x23

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
...	...

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		

free list

%x19
%x13
...
...

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?
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		



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



...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)

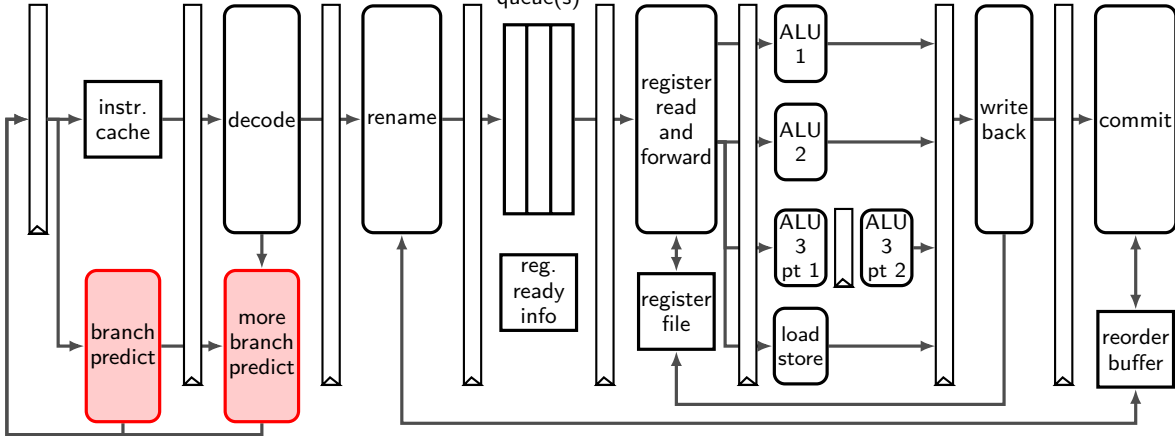
- can reconstruct register map before we commit the branch instruction

  - 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

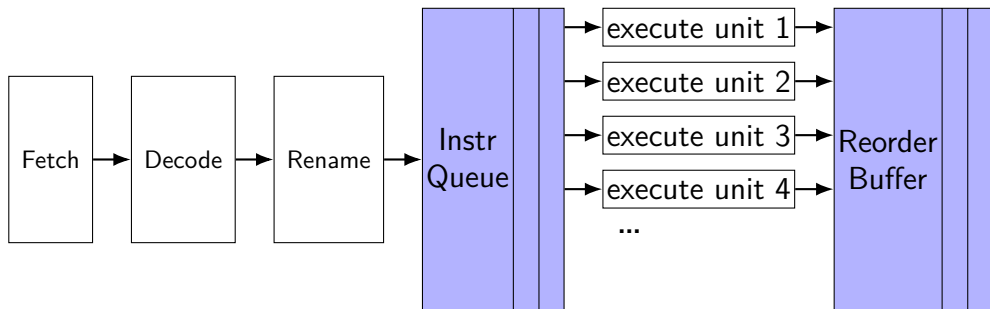
but pipelined: supports multiple pending cache misses in parallel

1 store functional unit

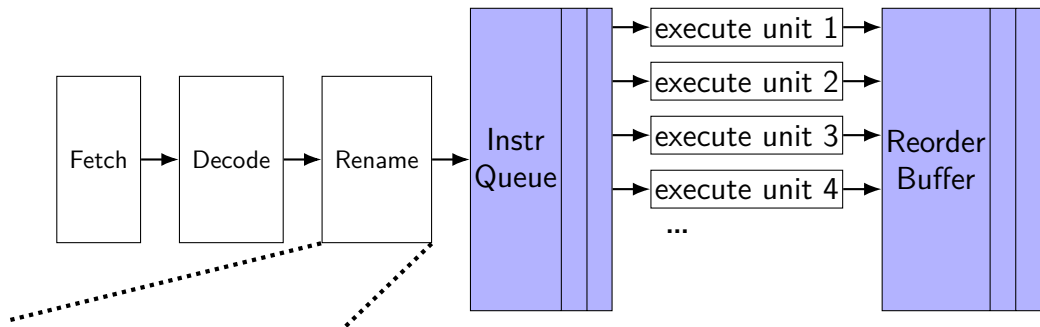
224-entry reorder buffer

determines how far ahead branch mispredictions, etc. can happen

# exceptions and OOO (one strategy)



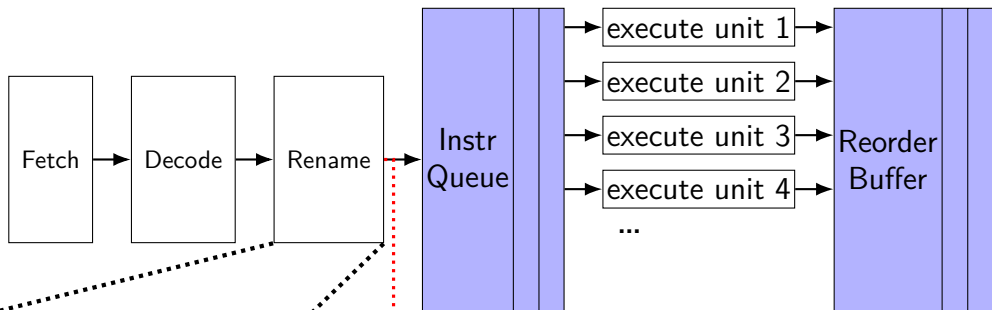
# exceptions and OOO (one strategy)



free regs for new instrs

X19	<b>arch. reg</b>	<b>phys. reg</b>
X23		
...		
	RAX	X15
	RCX	X17
	RBX	X13
	RBX	X07
	...	...

# exceptions and OOO (one strategy)



free regs for new instrs

X19
X23
...

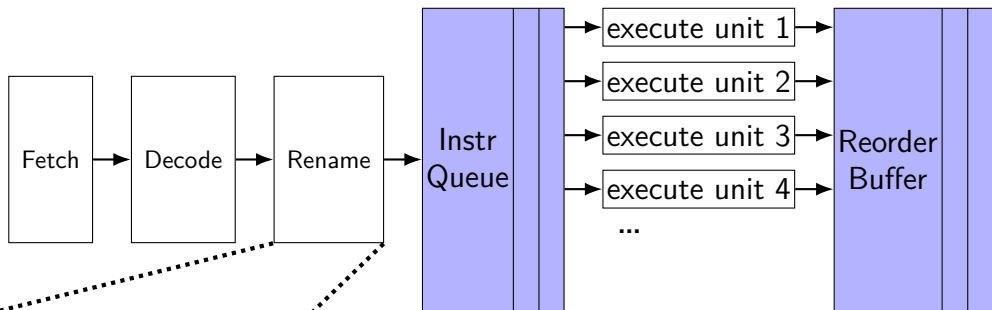
arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

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

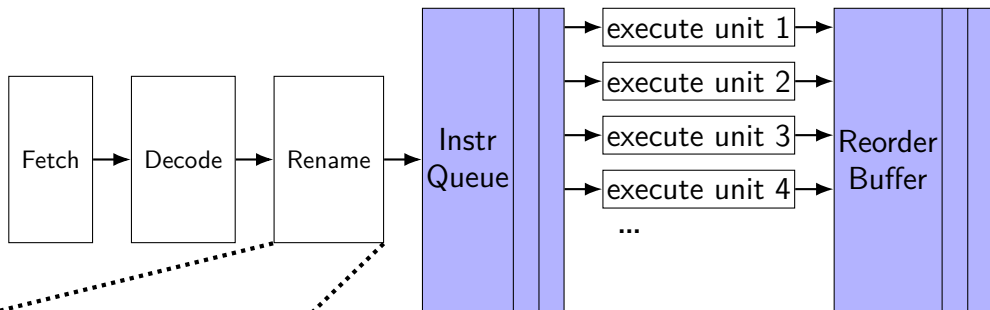
X19
X23
...

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

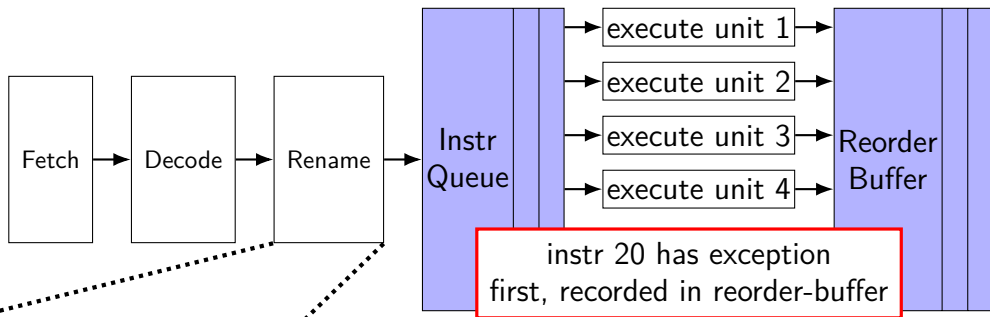
X19
X23
...

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

arch. reg	phys. reg
RAX	X21
RCX	X2
RCX	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

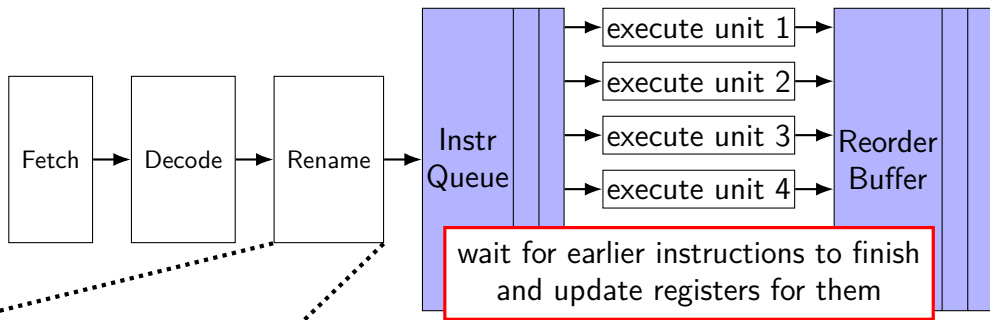
X19
X23
...

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
...

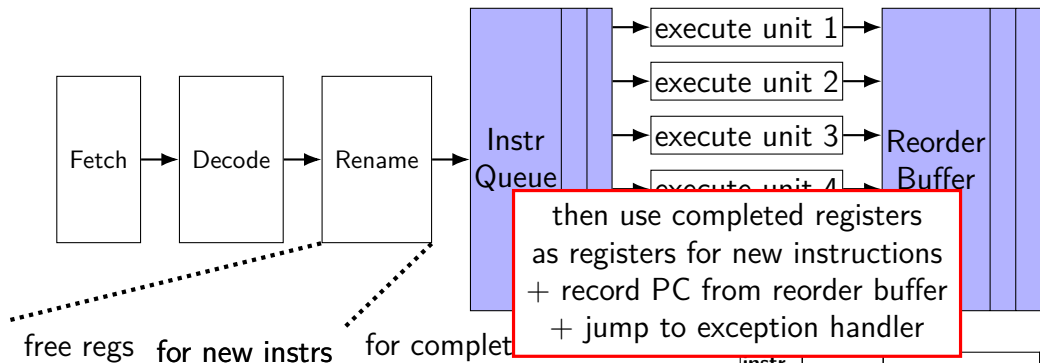
arch. reg	phys. reg
RAX	X15
RCX	X17
RBX	X13
RBX	X07
...	...

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)



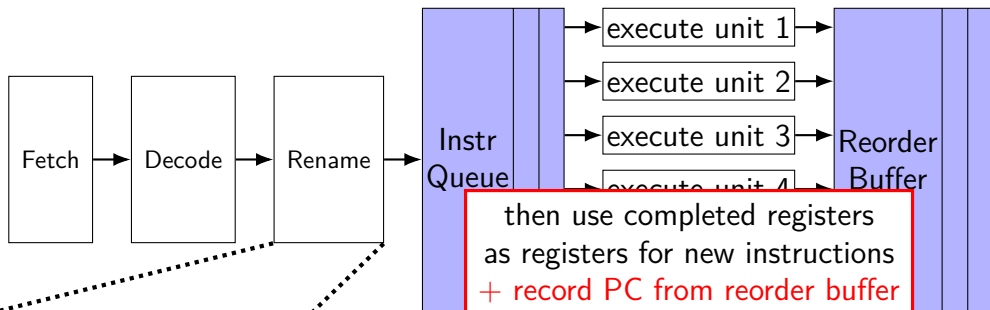
X19
X23
...

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

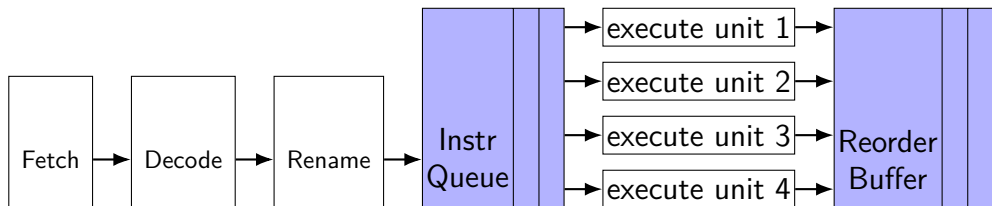
X19
X23
...

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

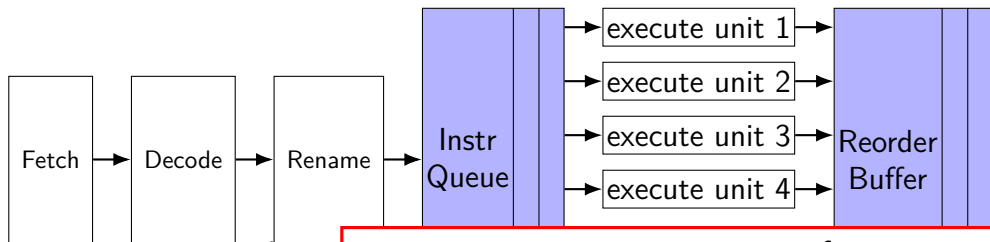
X19
X23
...

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