

cache performance (finish) / optimizations 1

last time

looking at spatial/temporal locality

rough estimates of misses

- look for innermost long loop

- count how much needs to be loaded each loop

- spatial locality: divide misses by elements/cache block

changing miss rates by changing loop orders

- for (i...) for (j...) \rightarrow for (j...) for (i...)

cache blocking introduction:

- changing loop orders often compromise: better locality for A, worse for B

- split up outer loops, do “blocks” of K indices at a time

- goal: keep K items active in cache where 1 would be before

 - improve spatial+temporal locality with less compromises

- alternate view: choose what parts of data to keep in cache in inner loop

- today: generalize and finish

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 iterations, and (assuming N large):

about 1 misses from A per j-loop iteration

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

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

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

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

$N^3 \div (2 \cdot \text{block size})$ 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 iterations, and (assuming N large):

about 1 misses from A per j-loop iteration

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

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

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

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

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

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": */  
       $C_{i+0,j} += A_{i+0,k+0} * B_{k+0,j}$   
       $C_{i+0,j} += A_{i+0,k+1} * B_{k+1,j}$   
       $C_{i+1,j} += A_{i+1,k+0} * B_{k+0,j}$   
       $C_{i+1,j} += A_{i+1,k+1} * B_{k+1,j}$   
    }  
  }  
}
```

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 A_{kj} :

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

load $I \times J$ elements of B_{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

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  
    incq    %rdx              // i += 1  
    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)

times unrolled	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

instruction cache/etc. overhead

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+0] * B[(k+0)*N+j];
      C[i*N+j+1] += A[i*N+k+1] * B[(k+1)*N+j+1];
    }
```

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
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];
      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];
      }
}
```

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];
    }
```


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



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
addq
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 _{v1}	→	%r8 _{v2}
addq %r11, %r8		addq %r11, %r8 _{v2}	→	%r8 _{v3}
addq %r12, %r8		addq %r12, %r8 _{v3}	→	%r8 _{v4}

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

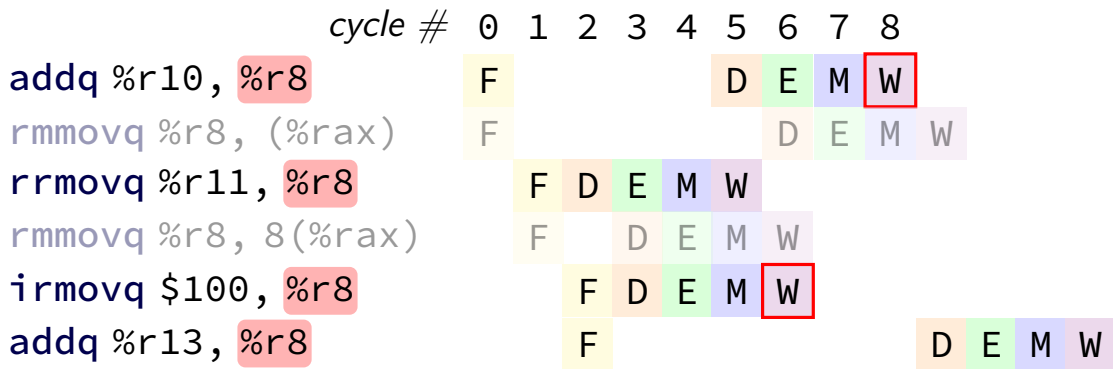
write-after-write example



out-of-order execution:

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

write-after-write example



out-of-order execution:

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

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

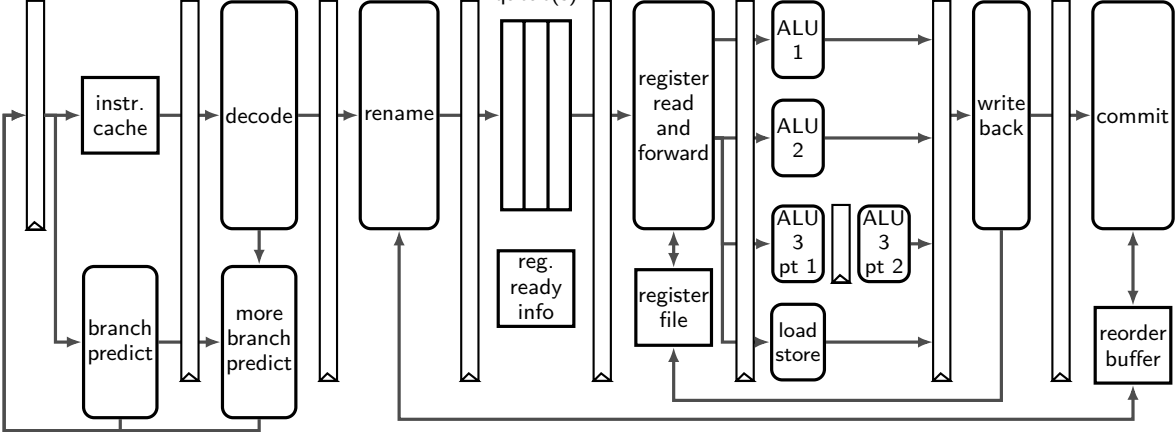
rename *architectural registers* to *physical registers*

different physical register for each version of architectural


track which physical registers are ready

compare physical register numbers to do forwarding

an OOO pipeline

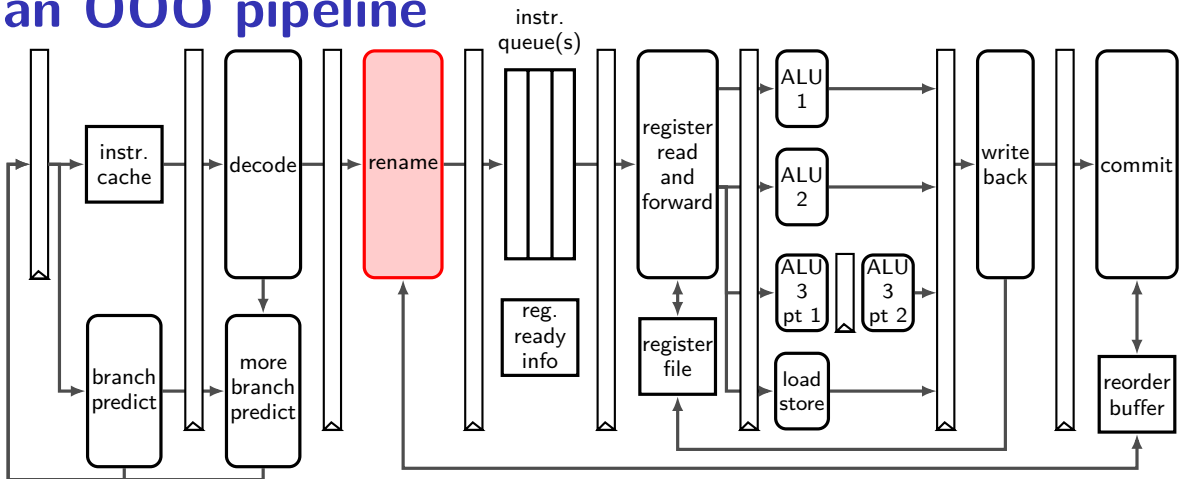


an OOO pipeline




branch prediction needs to happen before instructions decoded
done with cache-like tables of information about recent branches

an OOO pipeline



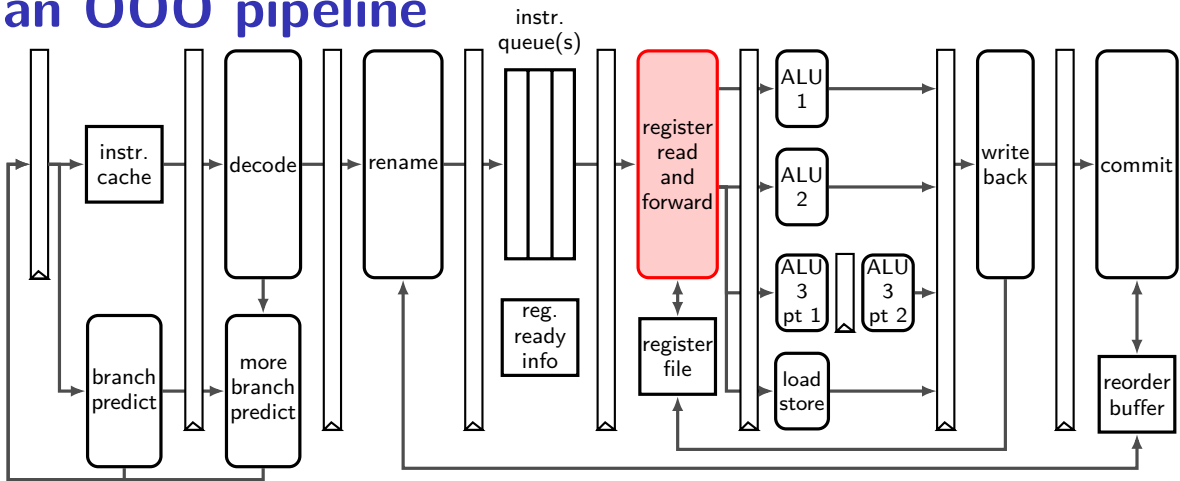
register renaming done here
stage needs to keep mapping from architectural to physical names

an OOO pipeline



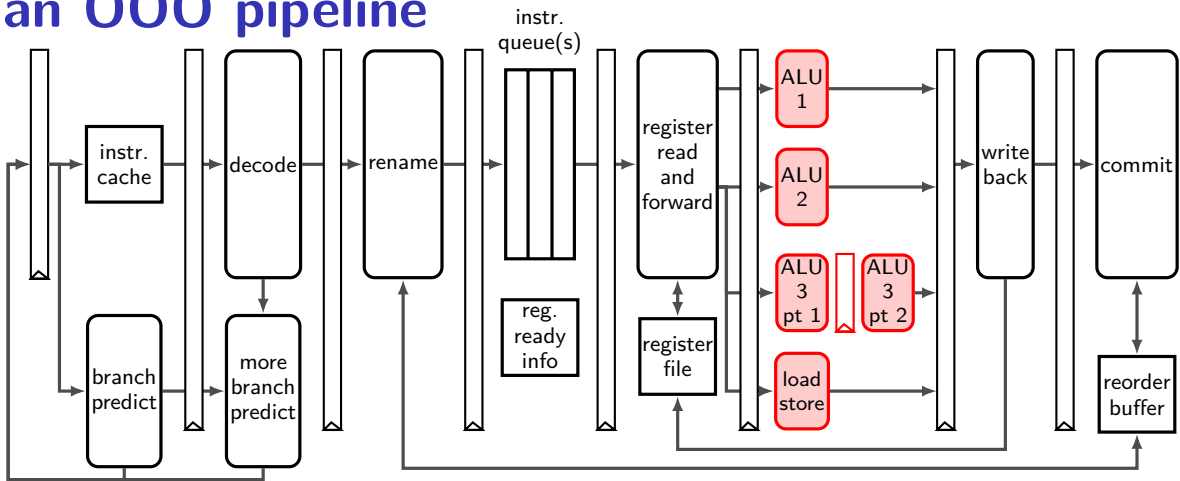
instruction queue holds pending renamed instructions combined with register-ready info to *issue* instructions (issue = start executing)

an OOO pipeline



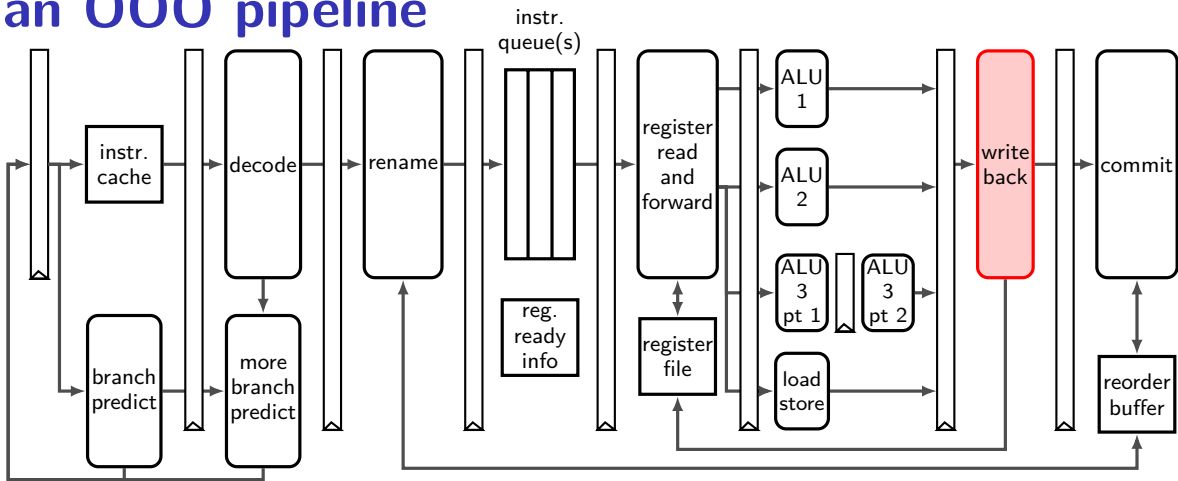
read from much larger register file and handle forwarding
register file: typically read 6+ registers at a time
(extra data paths wires for forwarding not shown)

an OOO pipeline



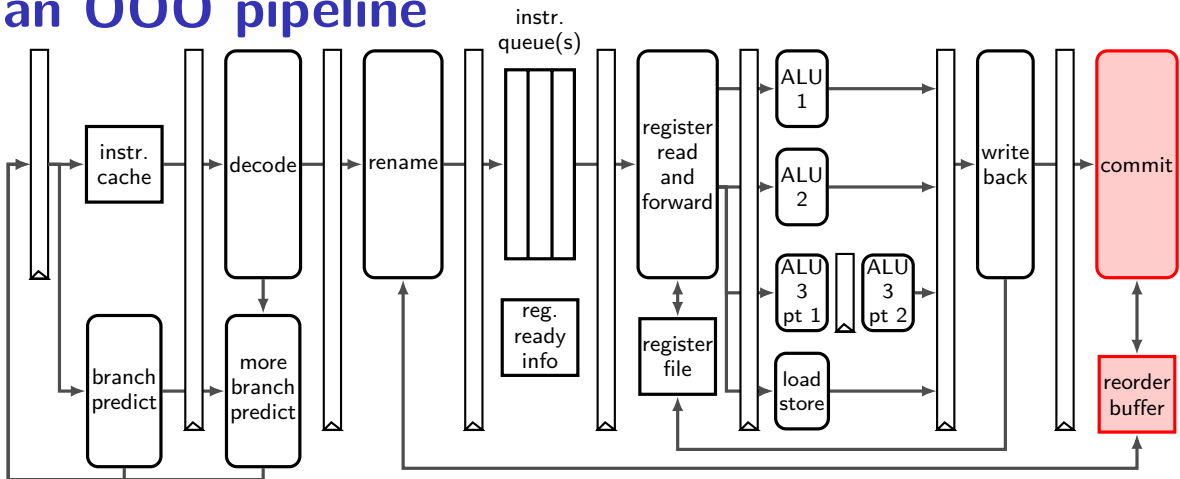
many *execution units* actually do math or memory load/store
some may have multiple pipeline stages
some may take variable time (data cache, integer divide, ...)

an OOO pipeline




writeback results to physical registers
register file: typically support writing 3+ registers at a time

an OOO pipeline



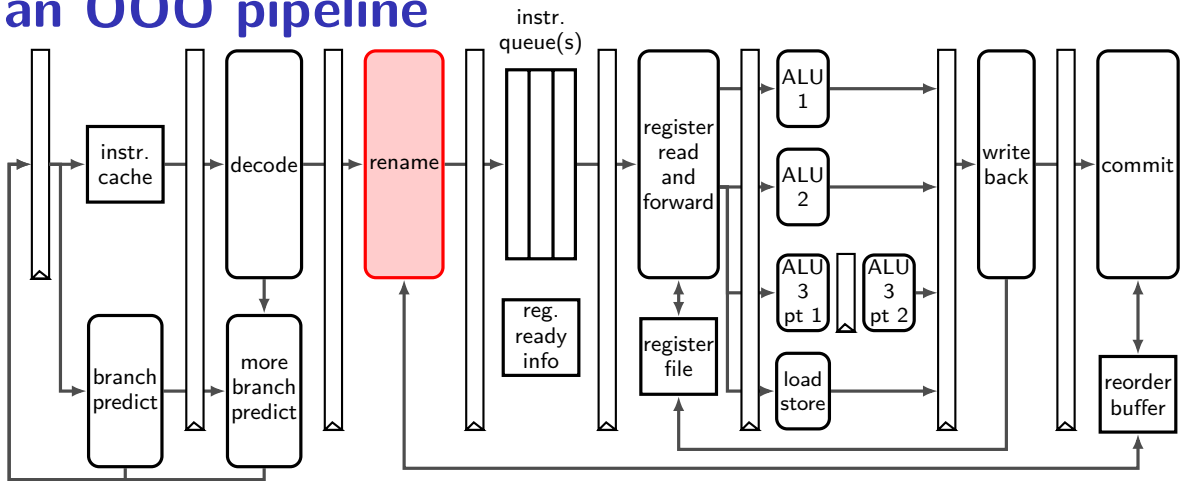
new commit (sometimes *retire*) stage finalizes instruction figures out when physical registers can be reused again

an OOO pipeline



commit stage also handles branch misprediction
reorder buffer tracks enough information to undo mispredicted instrs.

an OOO pipeline



register renaming

rename *architectural registers* to *physical registers*

architectural = part of instruction set architecture

different name for each version of architectural register

register renaming state

original
add %r10, %r8 ...
add %r11, %r8 ...
add %r12, %r8 ...

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming state

original
add %r10, %r8 ...
add %r11, %r8 ...
add %r12, %r8 ...

renamed

table for architectural (external)
and physical (internal) name
(for next instr. to process)

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming state

original
add %r10, %r8 ...
add %r11, %r8 ...
add %r12, %r8 ...

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

list of available physical registers
added to as instructions finish

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming state

original
add %r10, %r8 ...
add %r11, %r8 ...
add %r12, %r8 ...

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming example (1)

original

```
add %r10, %r8  
add %r11, %r8  
add %r12, %r8
```

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming example (1)

original	renamed
add %r10, %r8	add %x19, %x13 → %x18
add %r11, %r8	
add %r12, %r8	

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13 %x18
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming example (1)

original	renamed
add %r10, %r8	add %x19, %x13 → %x18
add %r11, %r8	add %x07, %x18 → %x20
add %r12, %r8	

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13 %x18 %x20
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming example (1)

original	renamed
add %r10, %r8	add %x19, %x13 → %x18
add %r11, %r8	add %x07, %x18 → %x20
add %r12, %r8	add %x05, %x20 → %x21

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18%x20
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming example (1)

original	renamed
add %r10, %r8	add %x19, %x13 → %x18
add %r11, %r8	add %x07, %x18 → %x20
add %r12, %r8	add %x05, %x20 → %x21

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18%x20
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
...	...

free reg list

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

```
addq %x19, %x13 → %x18
```

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13 %x18
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

```
addq %x19, %x13 → %x18
```

```
rmmovq %x18, (%x04) → (memory)
```

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original	renamed
addq %r10, %r8	addq %x19, %x13 → %x18
rmmovq %r8, (%rax)	rmmovq %x18, (%x04) → (memory)
subq %r8, %r11	
mrmovq 8(%r11), %r11	
irmovq \$100, %r8	
addq %r11, %r8	

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18
%r9	%x17
%r10	%x19
%r11	%x07
%r12	%x05
%r13	%x02

could be that %rax = 8+%r11
could load before value written!
possible data hazard!

not handled via register renaming

option 1: run load+stores in order

option 2: compare load/store addresses

%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

```
addq %x19, %x13 → %x18
rmmovq %x18, (%x04) → (memory)
subq %x18, %x07 → %x20
```

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18
%r9	%x17
%r10	%x19
%r11	%x07%x20
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

```
addq %x19, %x13 → %x18
rmmovq %x18, (%x04) → (memory)
subq %x18, %x07 → %x20
mrmovq 8(%x20), (memory) → %x21
```

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18
%r9	%x17
%r10	%x19
%r11	%x07%x20%x21
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

```
addq %x19, %x13 → %x18
rmmovq %x18, (%x04) → (memory)
subq %x18, %x07 → %x20
mrmovq 8(%x20), (memory) → %x21
irmovq $100 → %x23
```

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13 %x18 %x23
%r9	%x17
%r10	%x19
%r11	%x07 %x20 %x21
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming example (2)

original

```
addq %r10, %r8
rmmovq %r8, (%rax)
subq %r8, %r11
mrmovq 8(%r11), %r11
irmovq $100, %r8
addq %r11, %r8
```

renamed

```
addq %x19, %x13 → %x18
rmmovq %x18, (%x04) → (memory)
subq %x18, %x07 → %x20
mrmovq 8(%x20), (memory) → %x21
irmovq $100 → %x23
addq %x21, %x23 → %x24
```

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13%x18%x23%x24
%r9	%x17
%r10	%x19
%r11	%x07%x20%x21
%r12	%x05
%r13	%x02

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming exercise

original

```
addq %r8, %r9
movq $100, %r10
subq %r10, %r8
xorq %r8, %r9
andq %rax, %r9
```

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%r8	%x13
%r9	%x17
%r10	%x19
%r11	%x21
%r12	%x05
%r13	%x02
...	...

free
regs

%x18
%x20
%x21
%x23
%x24
...

register renaming: missing pieces

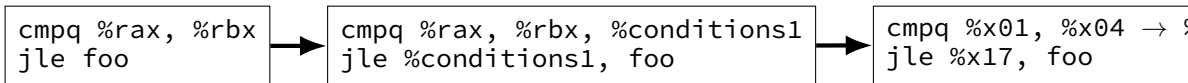

what about “hidden” inputs like `%rsp`, condition codes?

one solution: translate to instructions with additional register parameters

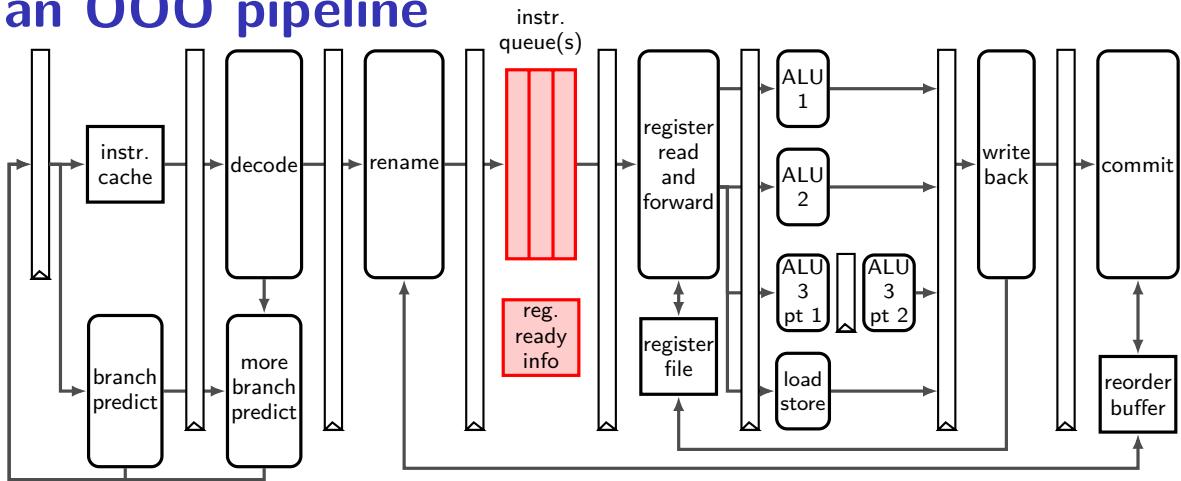
- making `%rsp` explicit parameter

- turning hidden condition codes into operands!

bonus: can also translate complex instructions to simpler ones



an OOO pipeline



instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending
%x07	pending
%x08	pending
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit

ALU 1

ALU 2

...

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending
%x07	pending
%x08	pending
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit cycle# 1
ALU 1
ALU 2 —

...

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending
%x07	pending
%x08	pending
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit cycle# 1
ALU 1
ALU 2 —

...

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending
%x08	pending
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit cycle# 1
ALU 1 1
ALU 2 —

...

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	...
ALU 1		1	2	
ALU 2		—	—	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	3	...
ALU 1		1	2	3	
ALU 2		—	—	—	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	3	...
ALU 1		1	2	3	
ALU 2		—	—	—	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending ready
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	3	4	...
ALU 1		1	2	3	4	
ALU 2		—	—	—	6	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending ready
%x10	pending
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	3	4	...
ALU 1		1	2	3	4	
ALU 2		—	—	—	6	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending ready
%x10	pending ready
%x11	pending
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	3	4	5	...
ALU 1		1	2	3	4	5	
ALU 2		—	—	—	6	7	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending ready
%x10	pending ready
%x11	pending ready
%x12	pending
%x13	pending
...	...

execution unit	cycle#	1	2	3	4	5	6	...
ALU 1		1	2	3	4	5	8	
ALU 2		—	—	—	6	7	—	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending ready
%x10	pending ready
%x11	pending ready
%x12	pending ready
%x13	pending
...	...

execution unit	cycle#	1	2	3	4	5	6	7	...
ALU 1		1	2	3	4	5	8	9	
ALU 2		—	—	—	6	7	—	...	

instruction queue and dispatch

instruction queue

#	instruction
1	addq %x01, %x05 → %x06
2	addq %x02, %x06 → %x07
3	addq %x03, %x07 → %x08
4	cmpq %x04, %x08 → %x09.cc
5	jne %x09.cc, ...
6	addq %x01, %x08 → %x10
7	addq %x02, %x09 → %x11
8	addq %x03, %x10 → %x12
9	cmpq %x04, %x11 → %x13.cc
...	...

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	ready
%x06	pending ready
%x07	pending ready
%x08	pending ready
%x09	pending ready
%x10	pending ready
%x11	pending ready
%x12	pending ready
%x13	pending ready
...	...

execution unit	cycle#	1	2	3	4	5	6	7	...
ALU 1		1	2	3	4	5	8	9	
ALU 2		—	—	—	6	7	—	...	

instruction queue and dispatch

instruction queue

#	instruction
1	<code>mrmovq (%x04) → %x06</code>
2	<code>mrmovq (%x05) → %x07</code>
3	<code>addq %x01, %x02 → %x08</code>
4	<code>addq %x01, %x06 → %x09</code>
5	<code>addq %x01, %x07 → %x10</code>

... ..

scoreboard

reg	status
%x01	ready
%x02	ready
%x03	ready
%x04	ready
%x05	
%x06	
%x07	
%x08	
%x09	
%x10	
...	...

execution unit cycle# 1 2 3 4 5 6 7 ...

ALU

data cache



assume 1 cycle/access

backup slides

a transformation

```
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 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 += 2) {  
        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 += 2) {  
        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 += 2) {
    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 += 2) {
        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}

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)

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

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

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

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

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

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

A_{ik} to $A_{i,k+1}$

—

■ B_{ki} to $B_{k+1,i}$

■ C_{ij}

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 kk :

for each i :

for each j :

for $k=kk, kk+1$:


$C_{ij} += A_{ik}$

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



have more than 4 cache blocks?

increasing kk increment would use more of them


cache blocking and miss rate



what about performance?



performance for big sizes



renaming with microcode

original
pushq %r8

renamed

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%rsp	%x11
...	...
%r8	%x13
%r9	%x17
%r10	%x19
...	...

free
regs

%x18
%x20
%x21
%x23
%x24
...

renaming with microcode

original
pushq %r8

renamed

iaddq \$8, %x11 → %x18

rmmovq %x13, 0(%x18) → (memory)

arch → phys
register map

%rax	%x04
%rcx	%x09
...	...
%rsp	%x11 %x18
...	...
%r8	%x13
%r9	%x17
%r10	%x19
...	...

pushq is really complicated
one implementation option:
split into simpler “microinstructions”
also exposes %rsp to register renaming

%x20
%x21
%x23
%x24
...

renaming and condition codes

original
`cmpq %r8, %r9`
`jle D`

renamed

arch \rightarrow phys
register map

<code>%rax</code>	<code>%x04</code>
<code>%rcx</code>	<code>%x09</code>
...	...
<code>%rsp</code>	<code>%x18</code>
...	...
<code>%r8</code>	<code>%x13</code>
<code>%r9</code>	<code>%x17</code>
<code>%r10</code>	<code>%x19</code>
...	...
<code>CC</code>	<code>%x04.cc</code>

free
regs

<code>%x18</code>
<code>%x20</code>
<code>%x21</code>
<code>%x23</code>
<code>%x24</code>
...

renaming and condition codes

original
`cmpq %r8, %r9`
`jle D`

renamed
`cmpq %x13, %x17 → %x18.cc`
`jle %x18.cc, D`

arch → phys
register map

<code>%rax</code>	<code>%x04</code>
<code>%rcx</code>	<code>%x09</code>
...	...
<code>%rsp</code>	<code>%x18</code>
...	...
<code>%r8</code>	<code>%x13</code>
<code>%r9</code>	<code>%x17</code>
<code>%r10</code>	<code>%x19</code>
...	...
<code>CC</code>	<code>%x04.cc%x18.cc</code>

one option for condition codes:
map to physical registers and track with renaming
not sure if real CPUs do this option
(complicates the commit stage?
more area for regs than alternative?)
alternative 1: entirely separate cond. code registers
(with separate free register list)
alternative 2: handle in 'in-order' part of pipeline?

<code>%x20</code>
<code>%x21</code>
<code>%x23</code>
<code>%x24</code>
...

renaming and condition codes

original
`cmpq %r8, %r9`
`jle D`

renamed
`cmpq %x13, %x17 → %x18.cc`
`jle %x18.cc, D`

arch → phys
register map

<code>%rax</code>	<code>%x04</code>
<code>%rcx</code>	<code>%x09</code>
...	...
<code>%rsp</code>	<code>%x18</code>
...	...
<code>%r8</code>	<code>%x13</code>
<code>%r9</code>	<code>%x17</code>
<code>%r10</code>	<code>%x19</code>
...	...
<code>CC</code>	<code>%x04.cc%x18.cc</code>

free
regs

<code>%x18</code>
<code>%x20</code>
<code>%x21</code>
<code>%x23</code>
<code>%x24</code>
...

renaming trickiness (1)

need to expose input + outputs

hidden dependencies on stack pointer, condition codes, memory, etc.

stack pointer + condition codes

- turn into visible register somehow

- alternative: force to execute in-order

memory: complex techniques we won't discuss

renaming trickiness (2)

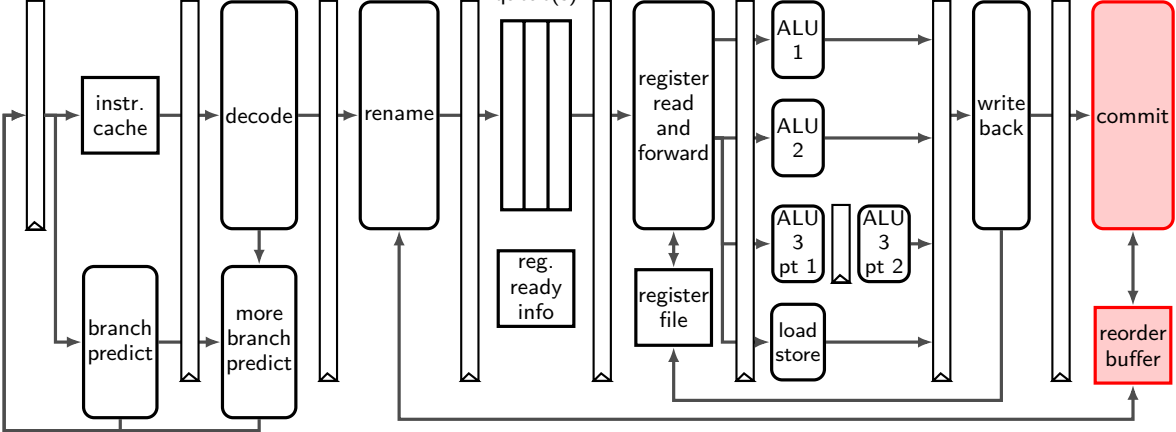
opportunity to translate complex instructions into simpler
commonly used for memory operands, probably some stack
instructions

`popq %rcx` → `addq, store`

`addq %rax, (%rbx)` → `load, addq, store`

...

an OOO pipeline



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: 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	%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



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
...
...

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		
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?
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: 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?
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		

free list

%x19
%x13
...
...

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?
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		

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?
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 ?