

# CS 33

## Architecture and Optimization (3)

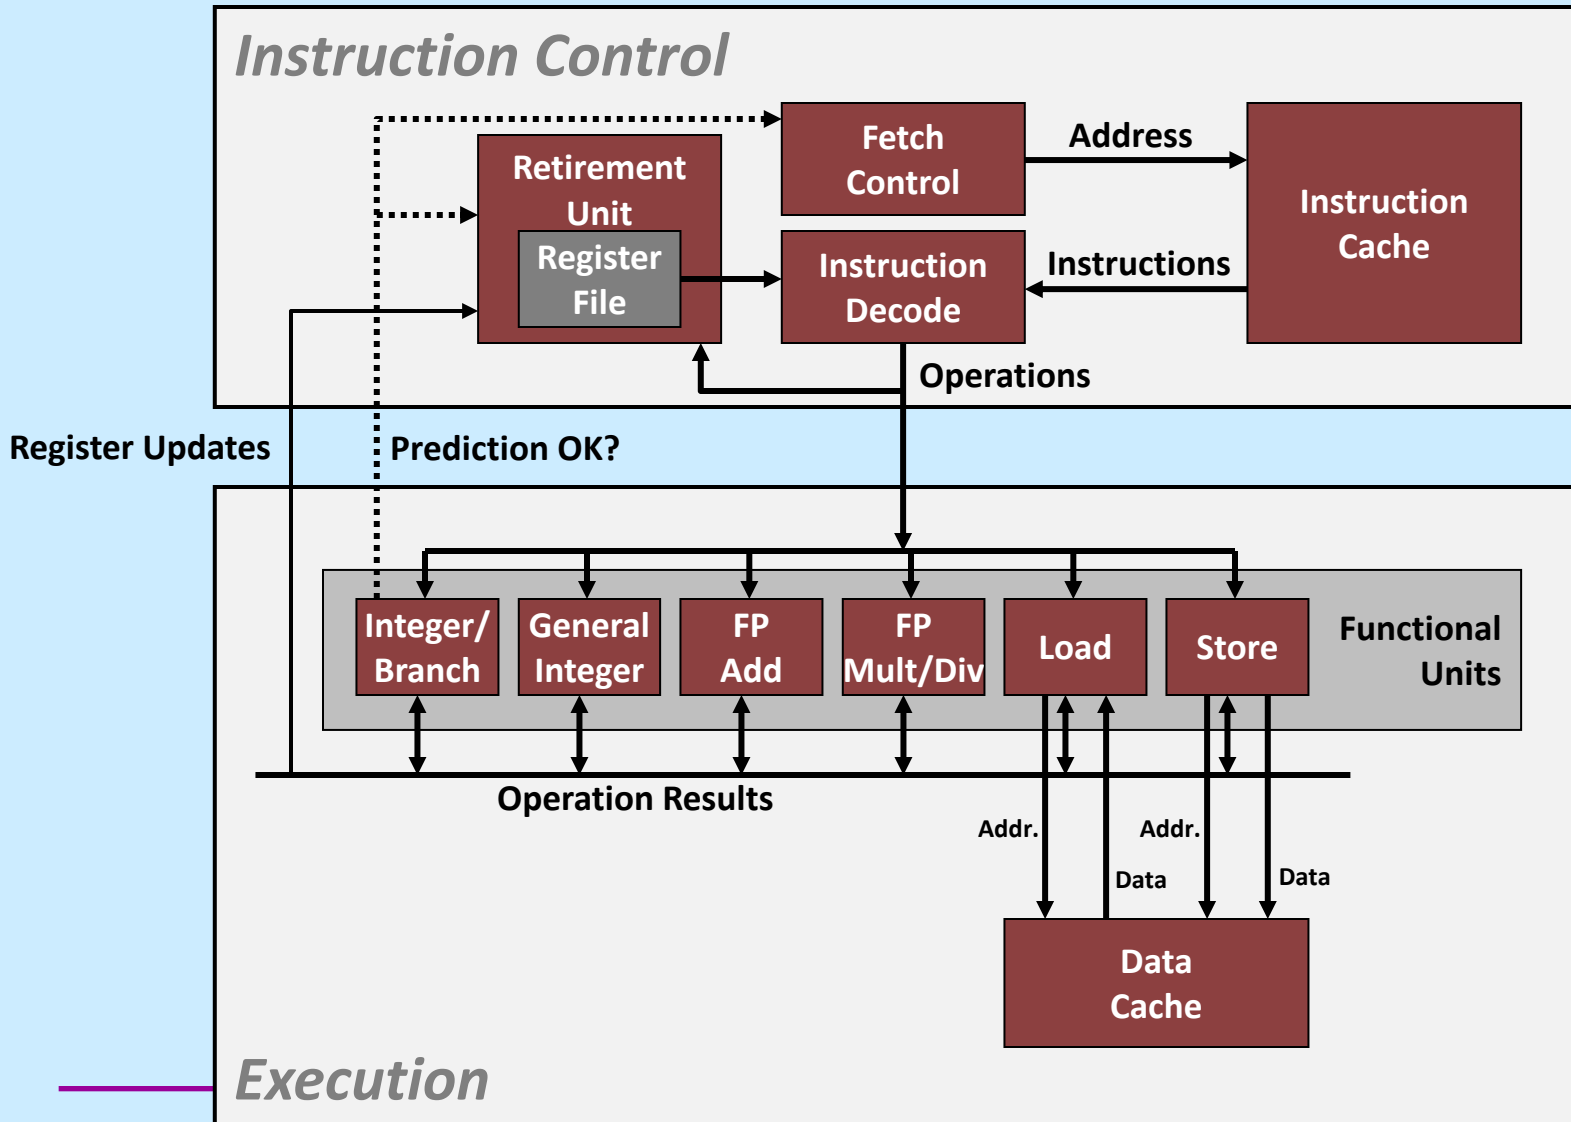
# The program so far

```
void combine4(vec_ptr_t v, data_t *dest) {
    int i;
    int length = vec_length(v);
    data_t *d = get_vec_start(v);
    data_t t = IDENT;
    for (i = 0; i < length; i++)
        t = t OP d[i];
    *dest = t;
}
```

Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine1 -O1	12.0	12.0	12.0	13.0
Combine4	2.0	3.0	3.0	5.0

*Can we do better?*

# Modern CPU Design



# Haswell CPU

- **Instruction characteristics**

<i><b>Instruction</b></i>	<i><b>Latency</b></i>	<i><b>Cycles/Issue</b></i>	<i><b>Capacity</b></i>
<b>Integer Add</b>	<b>1</b>	<b>1</b>	<b>4</b>
<b>Integer Multiply</b>	<b>3</b>	<b>1</b>	<b>1</b>
<b>Integer/Long Divide</b>	<b>3-30</b>	<b>3-30</b>	<b>1</b>
<b>Single/Double FP Add</b>	<b>3</b>	<b>1</b>	<b>1</b>
<b>Single/Double FP Multiply</b>	<b>5</b>	<b>1</b>	<b>2</b>
<b>Single/Double FP Divide</b>	<b>3-15</b>	<b>3-15</b>	<b>1</b>
<b>Load</b>	<b>4</b>	<b>1</b>	<b>2</b>
<b>Store</b>	<b>-</b>	<b>1</b>	<b>2</b>

# Haswell CPU Performance Bounds

	Integer		Floating Point	
	+	*	+	*
Latency	1.00	3.00	3.00	5.00
Throughput	4.00	1.00	1.00	2.00

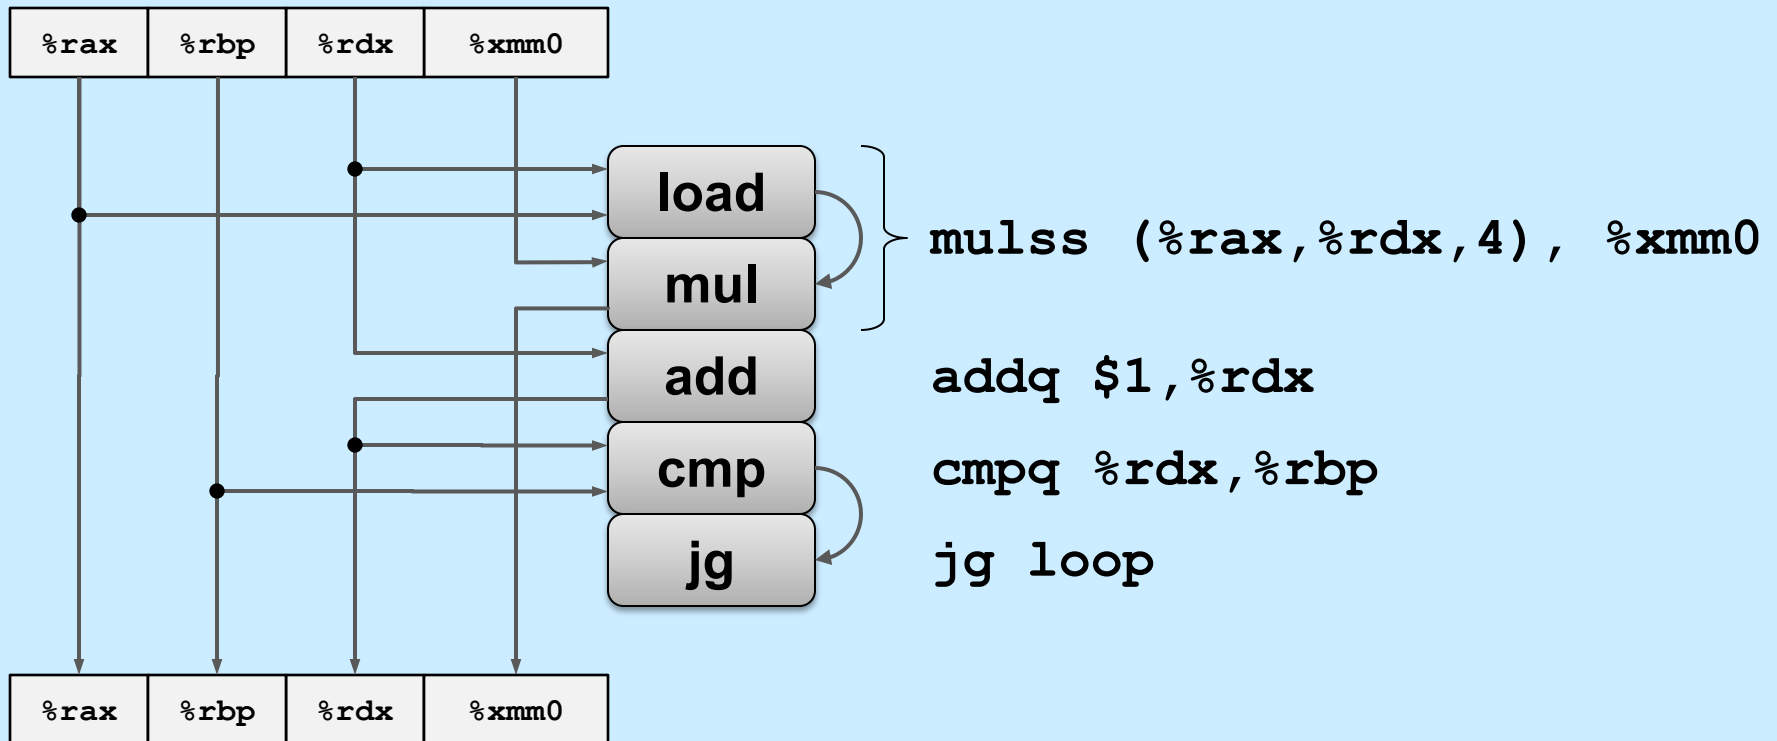
# x86-64 Compilation of Combine4

- Inner loop (case: SP floating-point multiply)

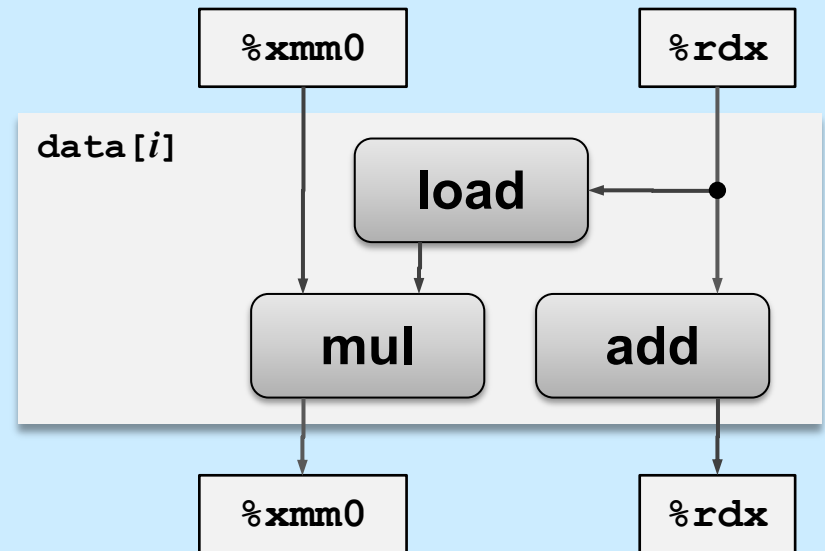

```
.L519:                # Loop:
  mullss (%rax,%rdx,4), %xmm0 # t = t * d[i]
  addq $1, %rdx             # i++
  cmpq %rdx, %rbp          # Compare length:i
  jg .L519                  # If >, goto Loop
```

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Combine4	1.27	3.00	3.00	5.00
Latency bound	1.00	3.00	3.00	5.0
Throughput bound	0.25	1.00	1.00	0.50

# Inner Loop




# Data-Flow Graphs of Inner Loop

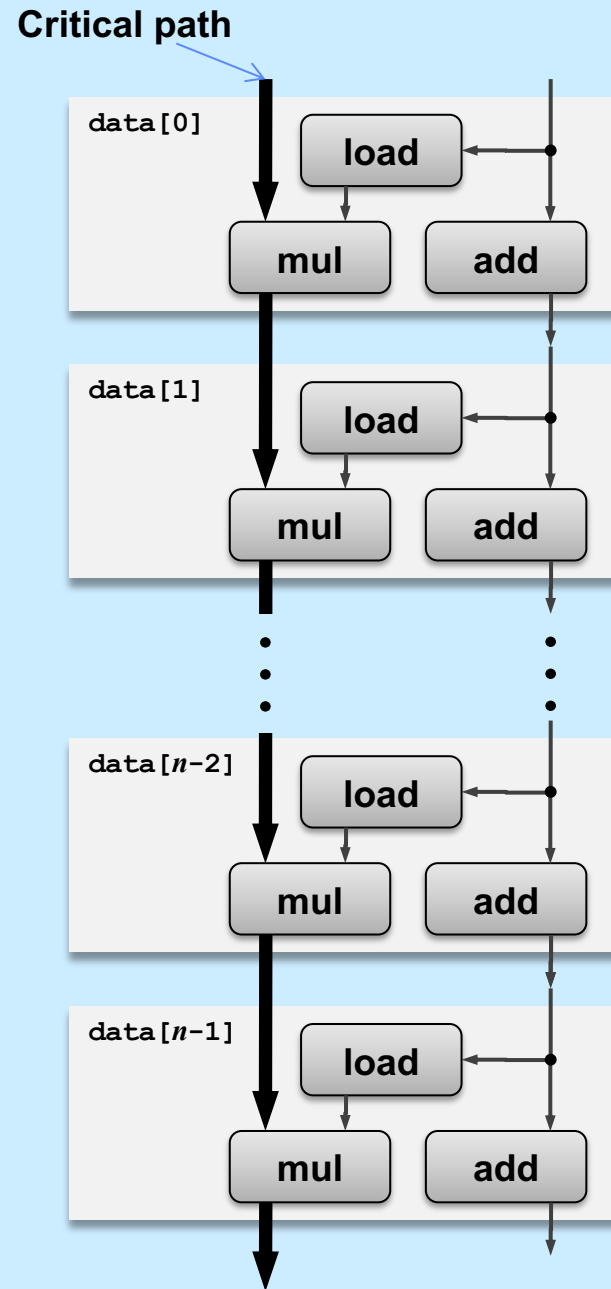




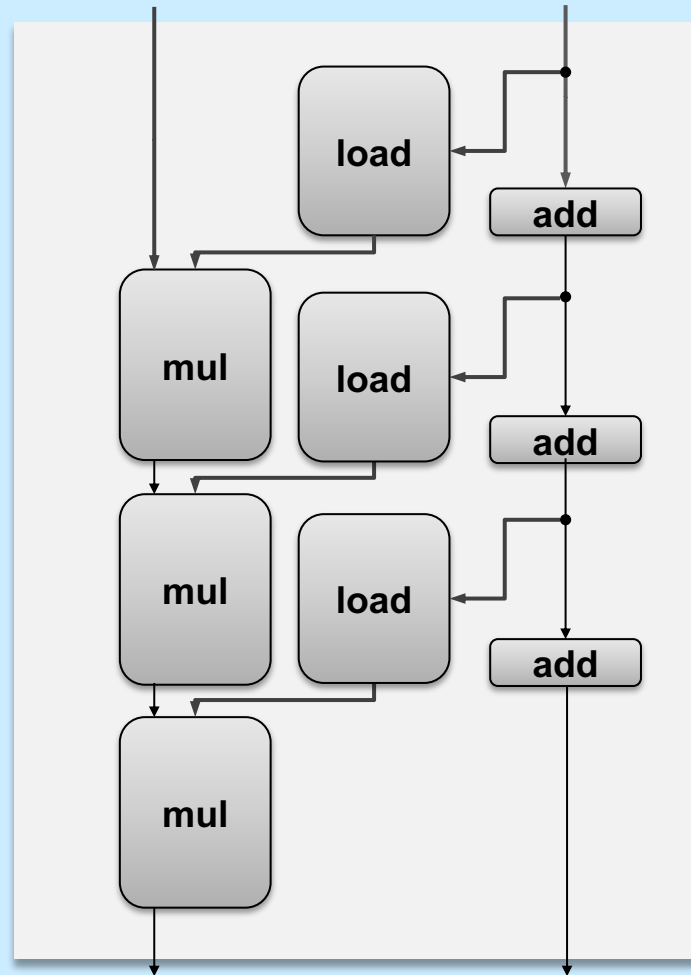
# Relative Execution Times



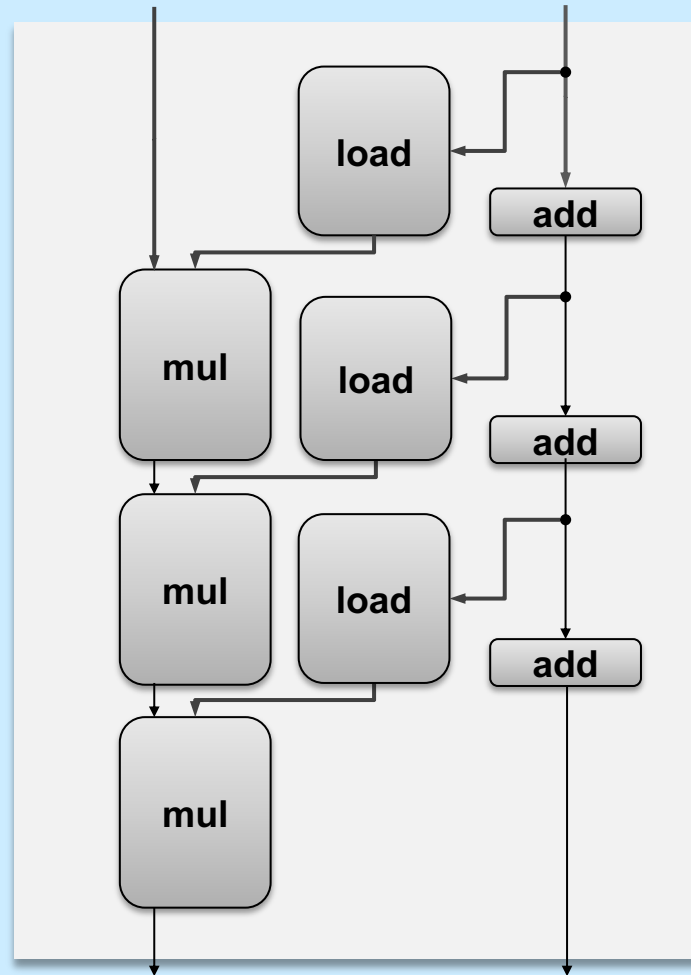
# Data Flow Over Multiple Iterations




# Pipelined Data-Flow Over Multiple Iterations



# Pipelined Data-Flow Over Multiple Iterations



# Pipelined Data-Flow Over Multiple Iterations



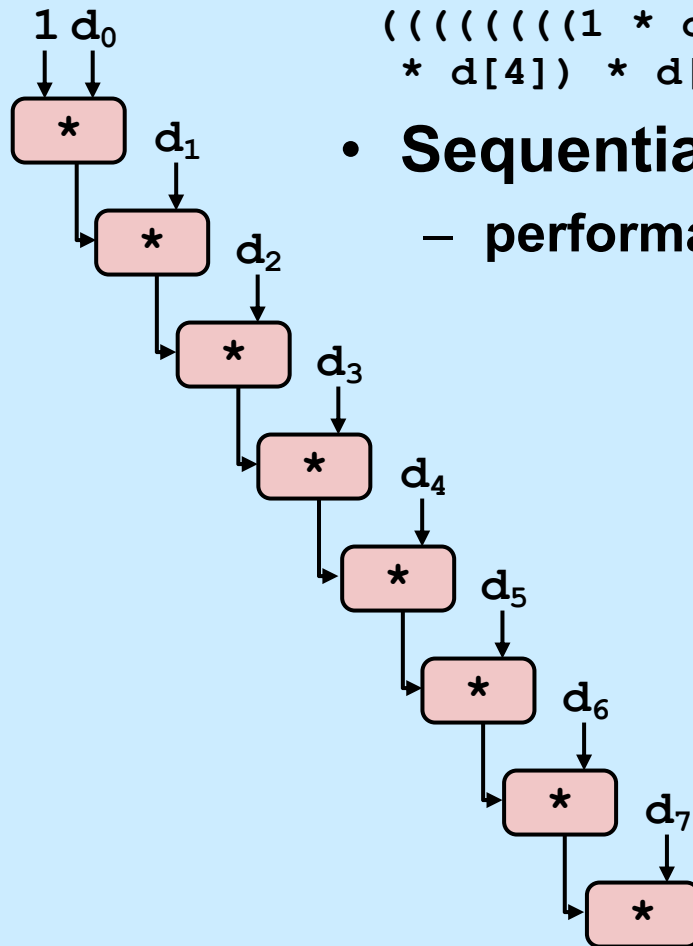
# Combine4 = Serial Computation (OP = \*)

- **Computation (length=8)**

$(((((1 * d[0]) * d[1]) * d[2]) * d[3]) * d[4]) * d[5]) * d[6]) * d[7])$

- **Sequential dependence**

– performance: determined by latency of OP



# Loop Unrolling

```
void unroll2x(vec_ptr_t v, data_t *dest)
{
    int length = vec_length(v);
    int limit = length-1;
    data_t *d = get_vec_start(v);
    data_t x = IDENT;
    int i;
    /* Combine 2 elements at a time */
    for (i = 0; i < limit; i+=2) {
        x = (x OP d[i]) OP d[i+1];
    }
    /* Finish any remaining elements */
    for (; i < length; i++) {
        x = x OP d[i];
    }
    *dest = x;
}
```

- **Perform 2x more useful work per iteration**

# Effect of Loop Unrolling

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Combine4	1.27	3.00	3.00	5.00
Unroll 2x	1.01	3.00	3.00	5.00
Latency bound	1.0	3.0	3.0	5.0
Throughput bound	0.25	1.0	1.0	0.5

- **Helps integer add**
  - reduces loop overhead
- **Others don't improve. *Why?***
  - still sequential dependency

```
x = (x OP d[i]) OP d[i+1];
```



# Loop Unrolling with Reassociation

```
void unroll2xra(vec_ptr_t v, data_t *dest)
{
    int length = vec_length(v);
    int limit = length-1;
    data_t *d = get_vec_start(v);
    data_t x = IDENT;
    int i;
    /* Combine 2 elements at a time */
    for (i = 0; i < limit; i+=2) {
        x = x OP (d[i] OP d[i+1]);
    }
    /* Finish any remaining elements */
    for (; i < length; i++) {
        x = x OP d[i];
    }
    *dest = x;
}
```


Compare to before

$x = (x \text{ OP } d[i]) \text{ OP } d[i+1];$

- Can this change the result of the computation?
- Yes, for FP. *Why?*

# Reassociated Computation

```
x = x OP (d[i] OP d[i+1]);
```



- **What changed:**

- ops in the next iteration can be started early (no dependency)

- **Overall Performance**

- N elements, D cycles  
latency/op
- should be  $(N/2+1)*D$  cycles:  
**CPE = D/2**
- measured CPE slightly worse for integer addition (there are other things going on)

# Effect of Reassociation

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Combine4	1.27	3.00	3.00	5.00
Unroll 2x	1.01	3.00	3.00	5.00
Unroll 2x, reassociate	1.01	1.51	1.51	2.51
Latency bound	1.0	3.0	3.0	5.0
Throughput bound	.25	1.0	1.0	.5

- Nearly 2x speedup for int \*, FP +, FP \*
  - reason: breaks sequential dependency

```
x = x OP (d[i] OP d[i+1]);
```

# Loop Unrolling with Separate Accumulators

```
void unroll2xp2x(vec_ptr_t v, data_t *dest)
{
    int length = vec_length(v);
    int limit = length-1;
    data_t *d = get_vec_start(v);
    data_t x0 = IDENT;
    data_t x1 = IDENT;
    int i;
    /* Combine 2 elements at a time */
    for (i = 0; i < limit; i+=2) {
        x0 = x0 OP d[i];
        x1 = x1 OP d[i+1];
    }
    /* Finish any remaining elements */
    for (; i < length; i++) {
        x0 = x0 OP d[i];
    }
    *dest = x0 OP x1;
}
```

- Different form of reassociation

# Effect of Separate Accumulators


Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine4	1.27	3.00	3.00	5.00
Unroll 2x	1.01	3.00	3.00	5.00
Unroll 2x, reassociate	1.01	1.51	1.51	2.51
Unroll 2x parallel 2x	.81	1.51	1.51	2.51
Latency bound	1.0	3.0	3.0	5.0
Throughput bound	.25	1.0	1.0	.5

- 2x speedup (over unroll 2x) for int \*, FP +, FP \*
  - breaks sequential dependency in a “cleaner,” more obvious way

```
x0 = x0 OP d[i];  
x1 = x1 OP d[i+1];
```

# Separate Accumulators

```
x0 = x0 OP d[i];  
x1 = x1 OP d[i+1];
```



- **What changed:**
  - two independent “streams” of operations
- **Overall Performance**
  - N elements, D cycles latency/op
  - should be  $(N/2+1)*D$  cycles:  
**CPE = D/2**
  - Integer addition improved, but not yet at predicted value


**What Now?**

# Quiz 1

**We're making progress. With two accumulators we get a two-fold speedup. With three accumulators, we can get a three-fold speedup. How much better performance can we expect if we add even more accumulators?**

- a) It keeps on getting better as we add more and more accumulators**
- b) It's limited by the latency bound**
- c) It's limited by the throughput bound**
- d) It's limited by something else**

# Performance



- **K-way loop unrolling with K accumulators**
  - **limited by number and throughput of functional units**



# Achievable Performance

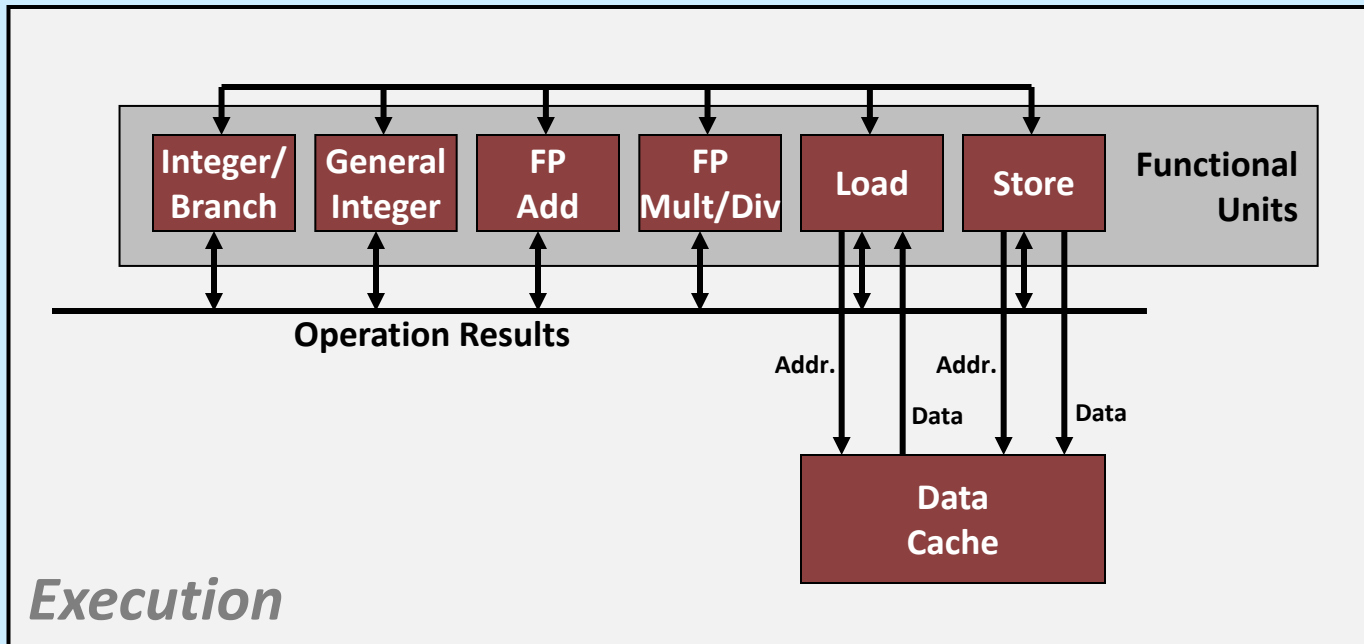


Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine4	1.27	3.0	3.0	5.0
Achievable scalar	.52	1.01	1.01	.54
Latency bound	1.00	3.00	3.00	5.00
Throughput bound	.25	1.00	1.00	.5

# Using Vector Instructions

Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine4	1.27	3.0	3.0	5.0
Achievable Scalar	.52	1.01	1.01	.54
Latency bound	1.00	3.00	3.00	5.00
Throughput bound	.25	1.00	1.00	.5
Achievable Vector	.05	.24	.25	.16
Vector throughput bound	.06	.12	.25	.12

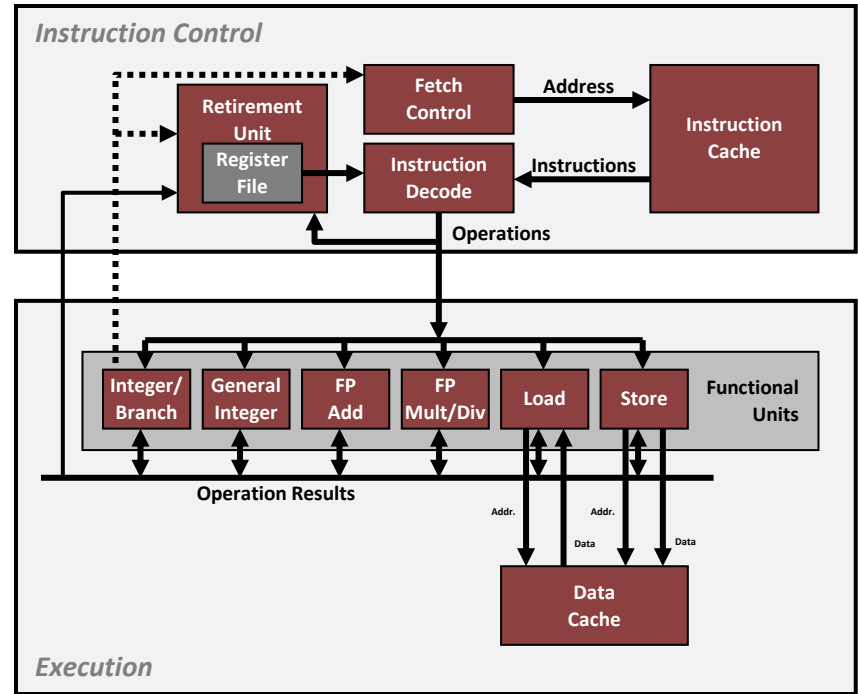
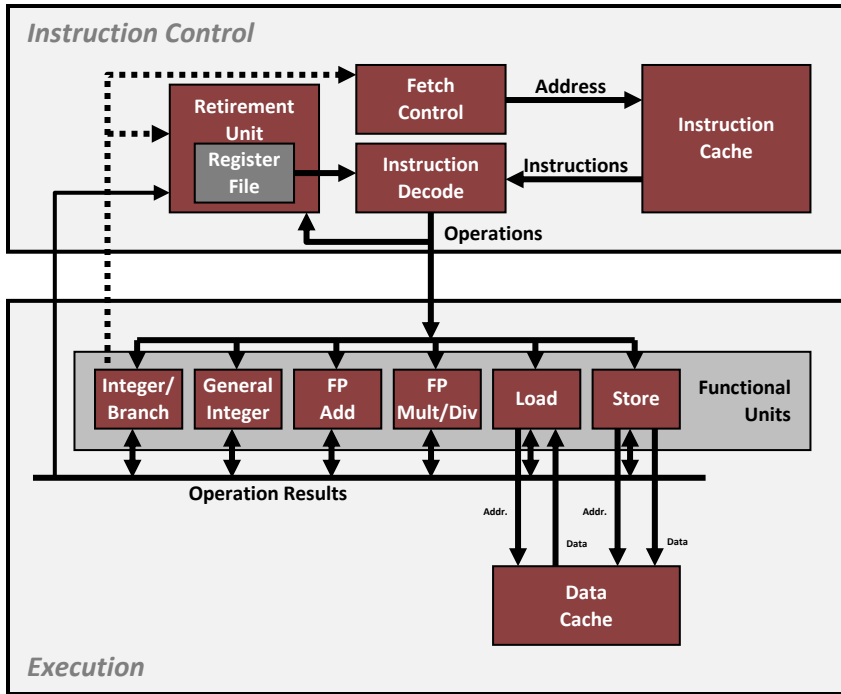
- **Make use of SSE Instructions**
  - parallel operations on multiple data elements

# Hyper Threading



# Multiple Cores

## Chip



Other Stuff

More Cache

Other Stuff

# CS 33

## Memory Hierarchy I

# Random-Access Memory (RAM)

- **Key features**
  - **RAM** is traditionally packaged as a chip
  - basic storage unit is normally a **cell** (one bit per cell)
  - multiple RAM chips form a memory
- **Static RAM (SRAM)**
  - each cell stores a bit with a four- or six-transistor circuit
  - retains value indefinitely, as long as it is kept powered
  - relatively insensitive to electrical noise (EMI), radiation, etc.
  - faster and more expensive than DRAM
- **Dynamic RAM (DRAM)**
  - each cell stores bit with a capacitor; transistor is used for access
  - value must be refreshed every 10-100 ms
  - more sensitive to disturbances (EMI, radiation,...) than SRAM
  - slower and cheaper than SRAM

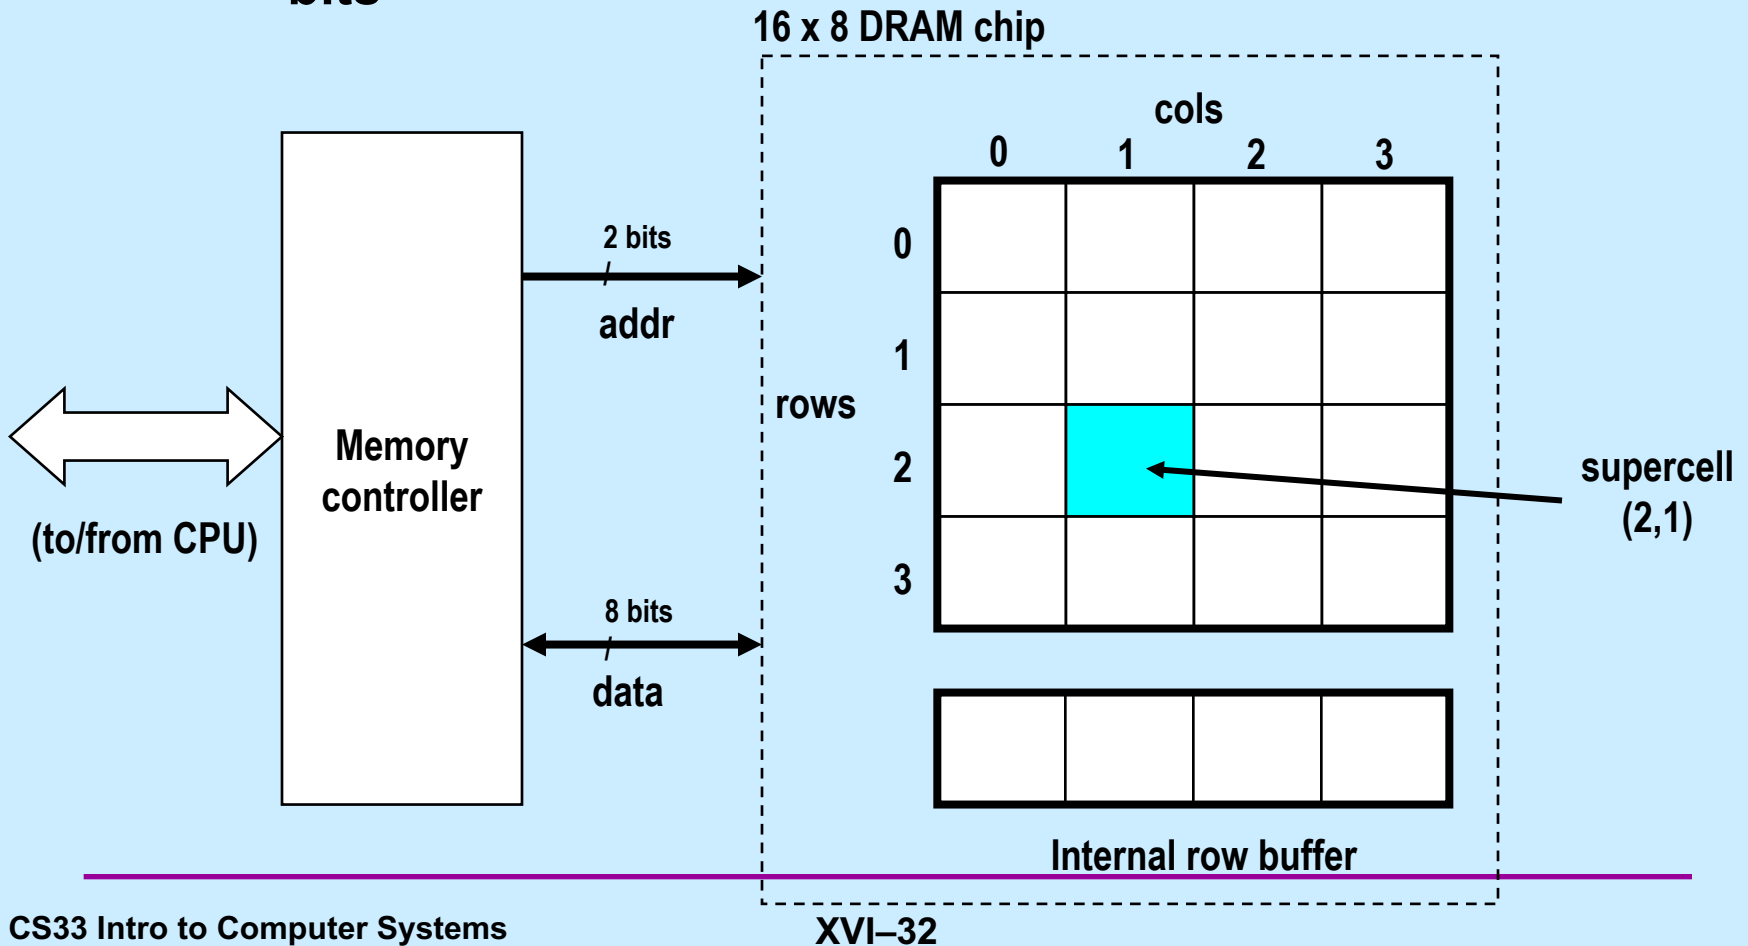
# SRAM vs DRAM Summary

	Trans. per bit	Access time	Needs refresh?	Needs EDC?	Cost	Applications
SRAM	4 or 6	1X	No	Maybe	100x	Cache memories
DRAM	1	10X	Yes	Yes	1X	Main memories, frame buffers

- **EDC = error detection and correction**
  - to cope with noise, etc.

# Conventional DRAM Organization

- $d \times w$  DRAM:
  - $dw$  total bits organized as  $d$  **supercells** of size  $w$  bits






# Reading DRAM Supercell (2,1)

Step 1(a): row access strobe (**RAS**) selects row 2


Step 1(b): row 2 copied from DRAM array to row buffer



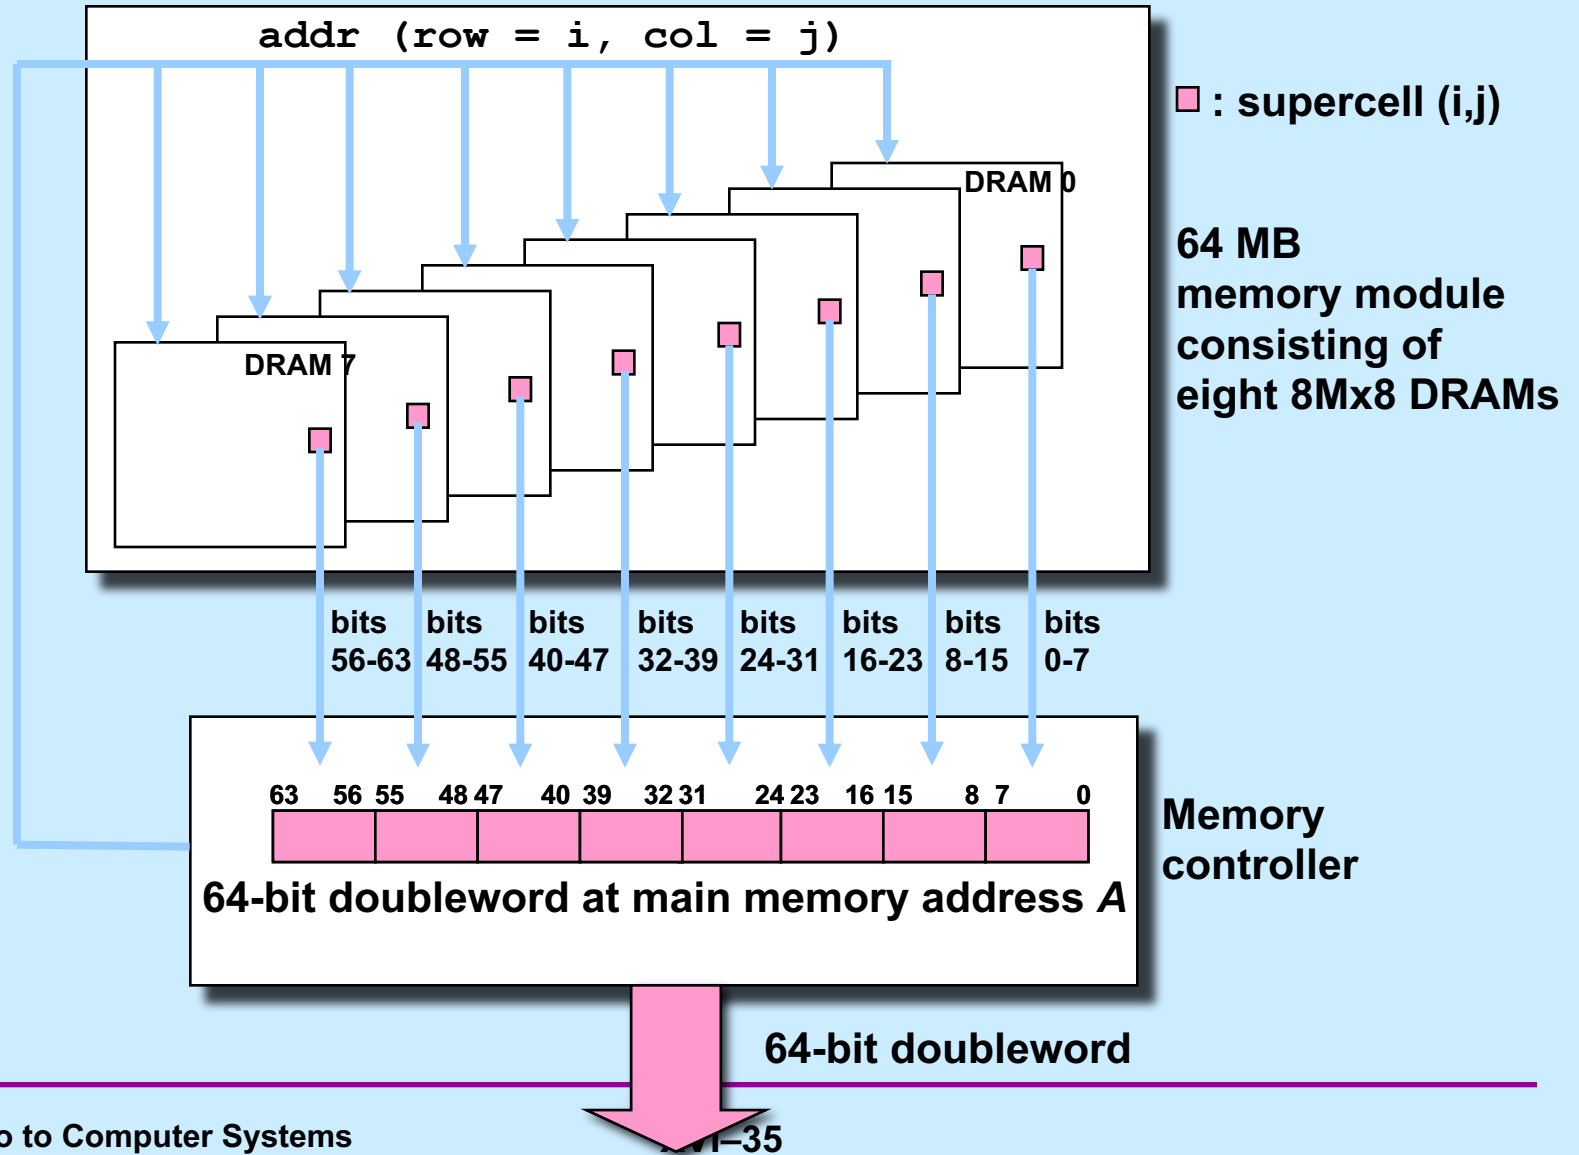
# Reading DRAM Supercell (2,1)

Step 2(a): column access strobe (**CAS**) selects column 1

Step 2(b): supercell (2,1) copied from buffer to data lines, and eventually back to the CPU




# Memory Modules



# Enhanced DRAMs

- **Basic DRAM cell has not changed since its invention in 1966**
  - commercialized by Intel in 1970
- **DRAMs with better interface logic and faster I/O:**
  - **synchronous DRAM (SDRAM or SDR)**
    - » uses a conventional clock signal instead of asynchronous control
    - » allows reuse of the row addresses (e.g., RAS, CAS, CAS, CAS)
  - **double data-rate synchronous DRAM (DDR SDRAM)**
    - » **DDR1**
      - twice as fast: 16 consecutive bytes xfr'd as fast as 8 in SDR
    - » **DDR2**
      - 4 times as fast: 32 consecutive bytes xfr'd as fast as 8 in SDR
    - » **DDR3**
      - 8 times as fast: 64 consecutive bytes xfr'd as fast as 8 in SDR

# Enhanced DRAMs



# DDR4

- **Memory transfer speed increased by a factor of 16 (twice as fast as DDR3)**
  - no increase in DRAM Cell Array speed (same as SDR)
  - 16 times more data transferred at once
    - » 64 adjacent bytes fetched from DRAM
      - just like DDR3

# Quiz 2

**A program is loading randomly selected bytes from memory. These bytes will be delivered to the processor on a DDR4 system at a speed that's  $n$  times that of an SDR system, where  $n$  is:**

- a) 8**
- b) 4**
- c) 2**
- d) 1**

# A Mismatch

- **A processor clock cycle is ~0.3 nsecs**
  - Older SunLab machines (Intel Core i5-4690) run at 3.5 GHz
- **Basic operations take 1 – 10 clock cycles**
  - .3 – 3 nsecs
- **Accessing memory takes 70-100 nsecs**
- **How is this made to work?**



# Caching to the Rescue

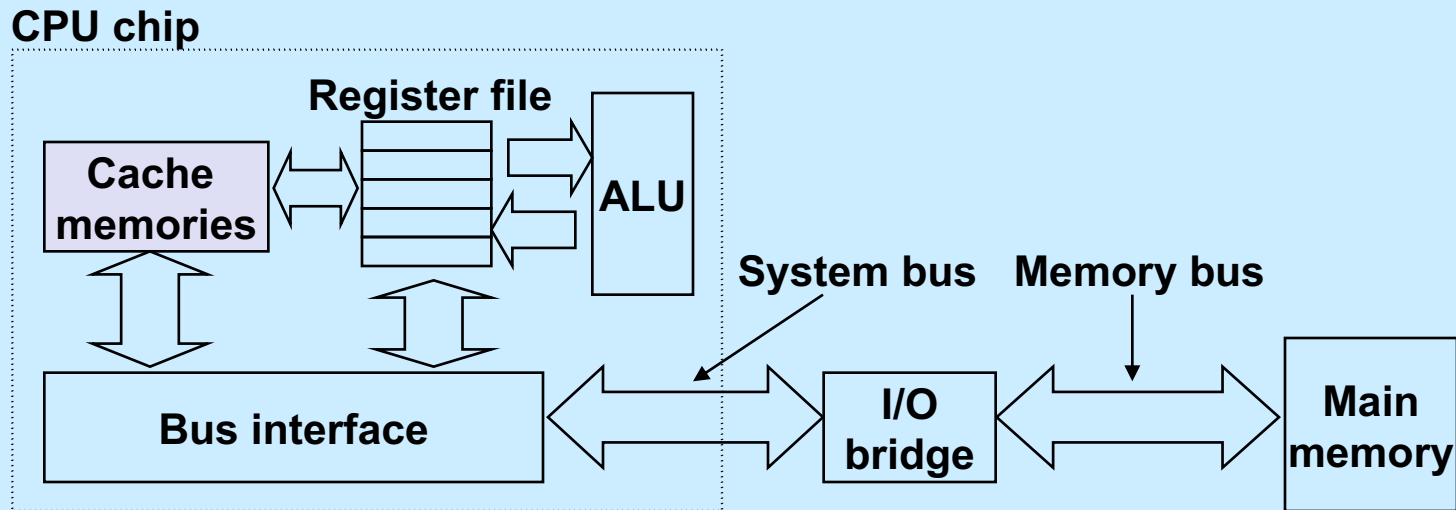
**CPU**

**Cache**

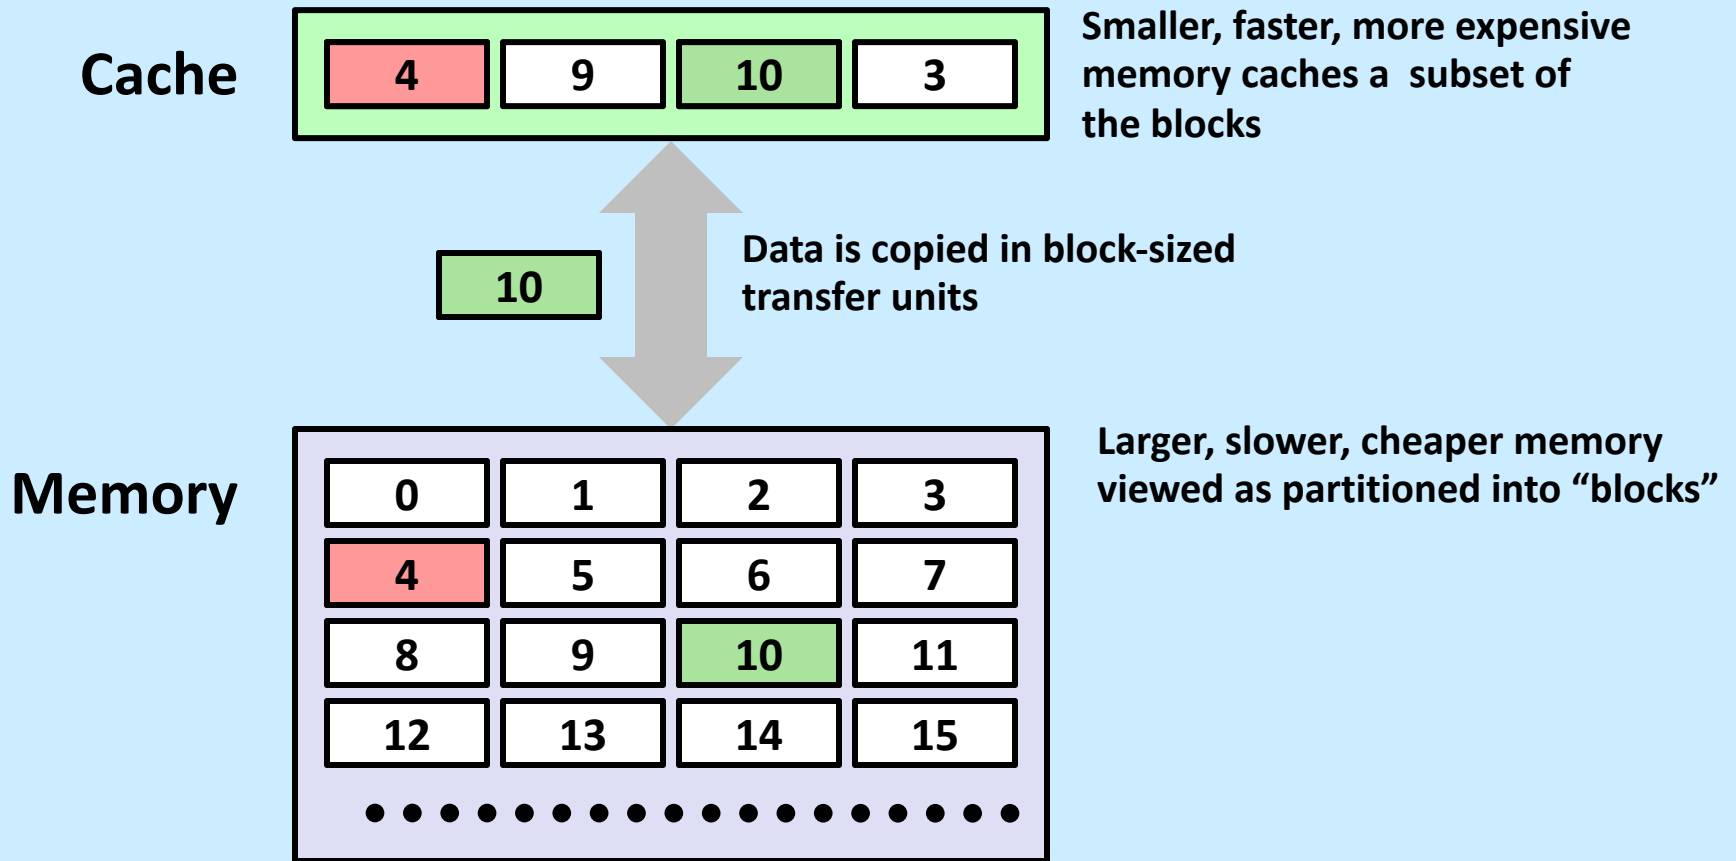


# Cache Memories

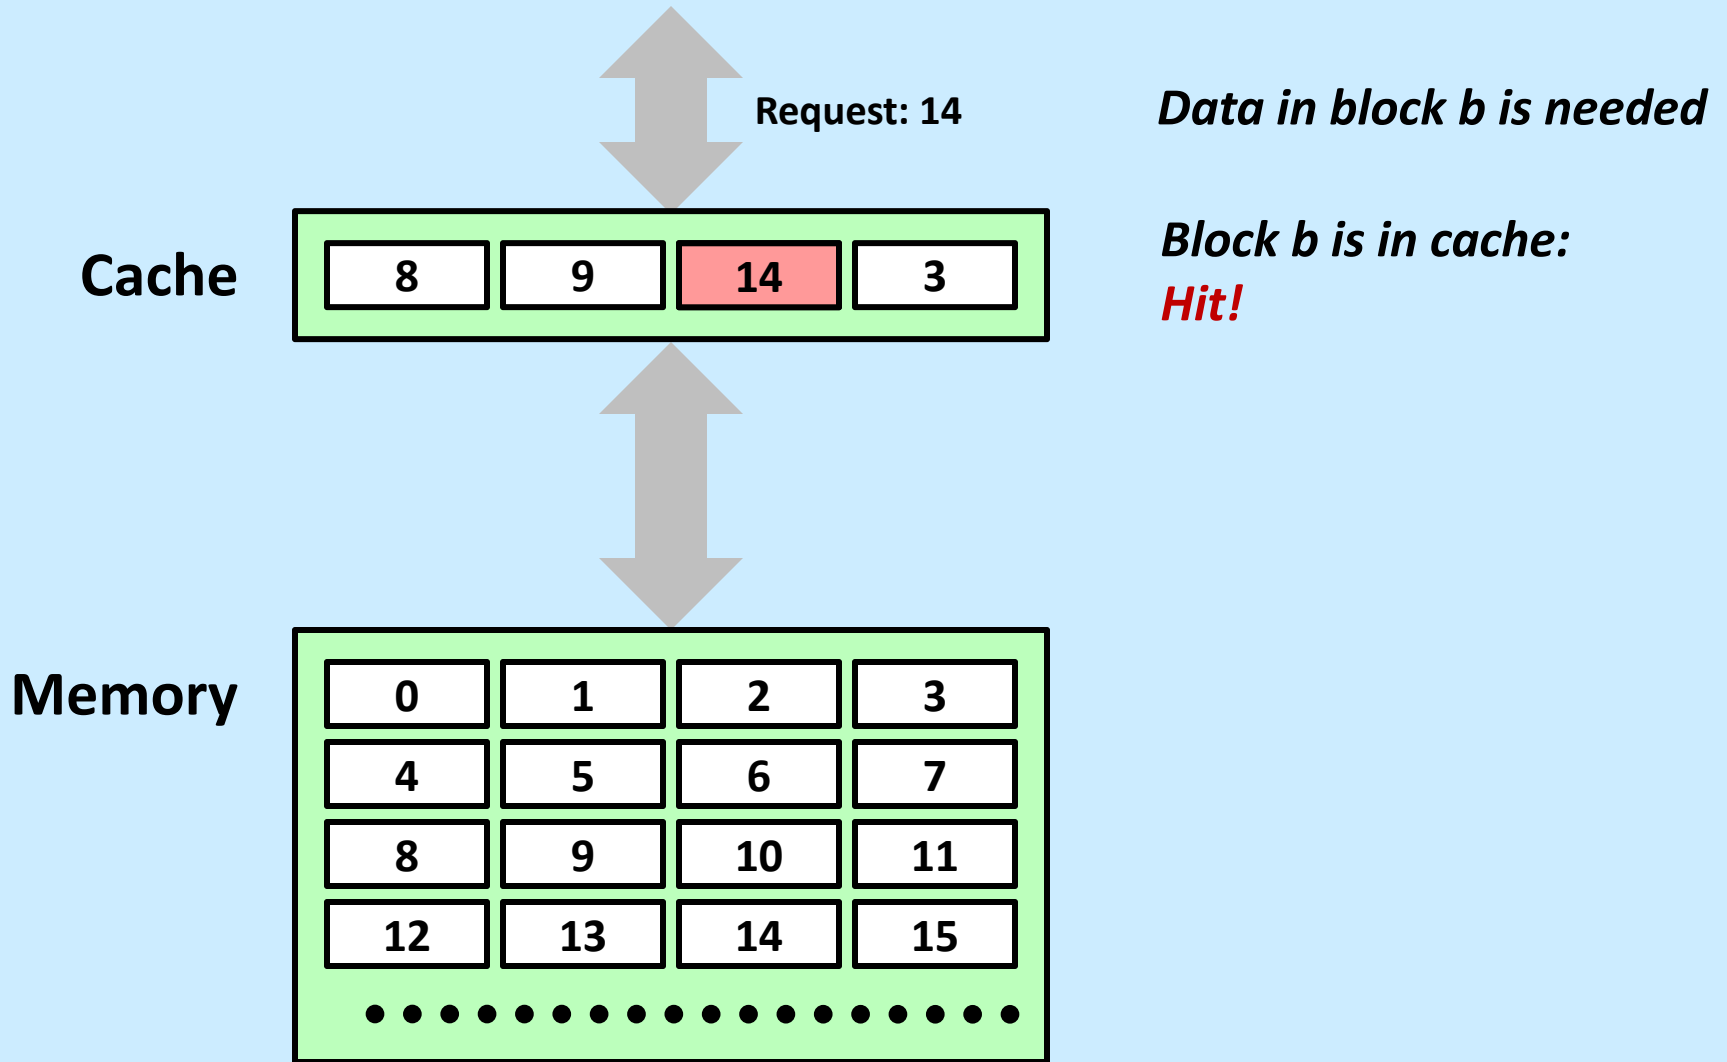
- **Cache memories** are small, fast SRAM-based memories managed automatically in hardware
  - hold frequently accessed blocks of main memory
- CPU looks first for data in caches (e.g., L1, L2, and L3), then in main memory
- Typical system structure:



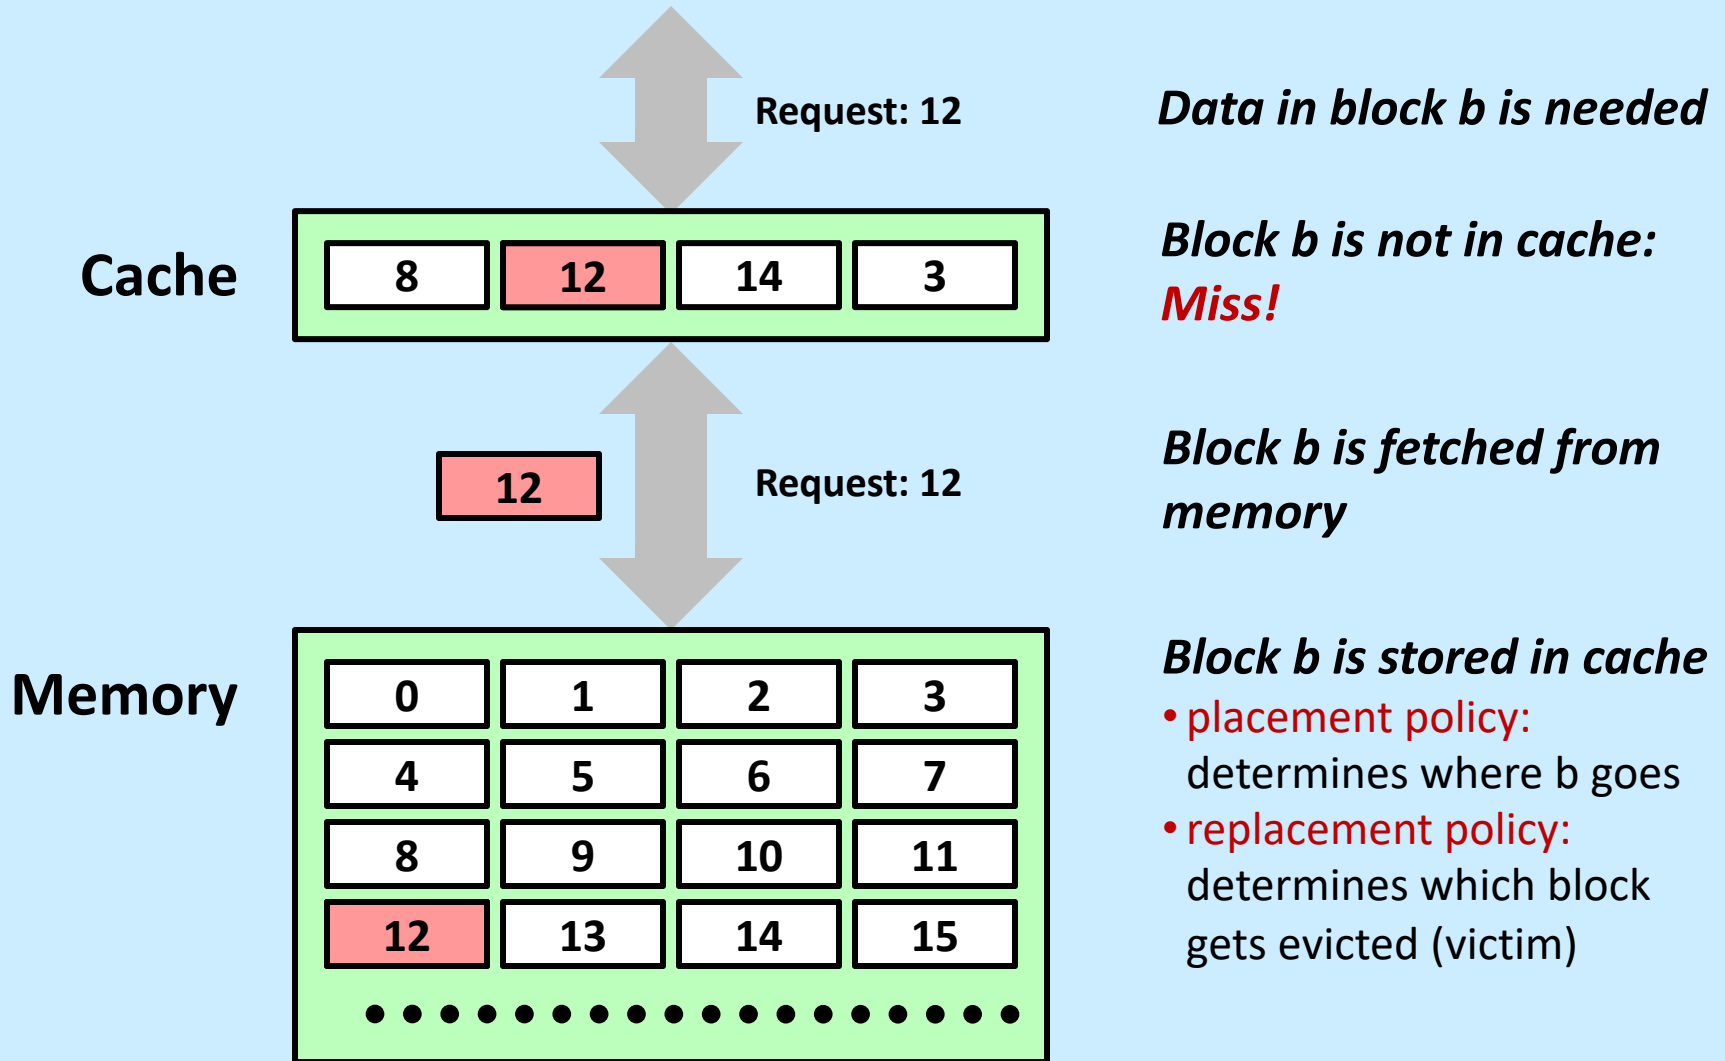
# General Cache Concepts



# General Cache Concepts: Hit



# General Cache Concepts: Miss




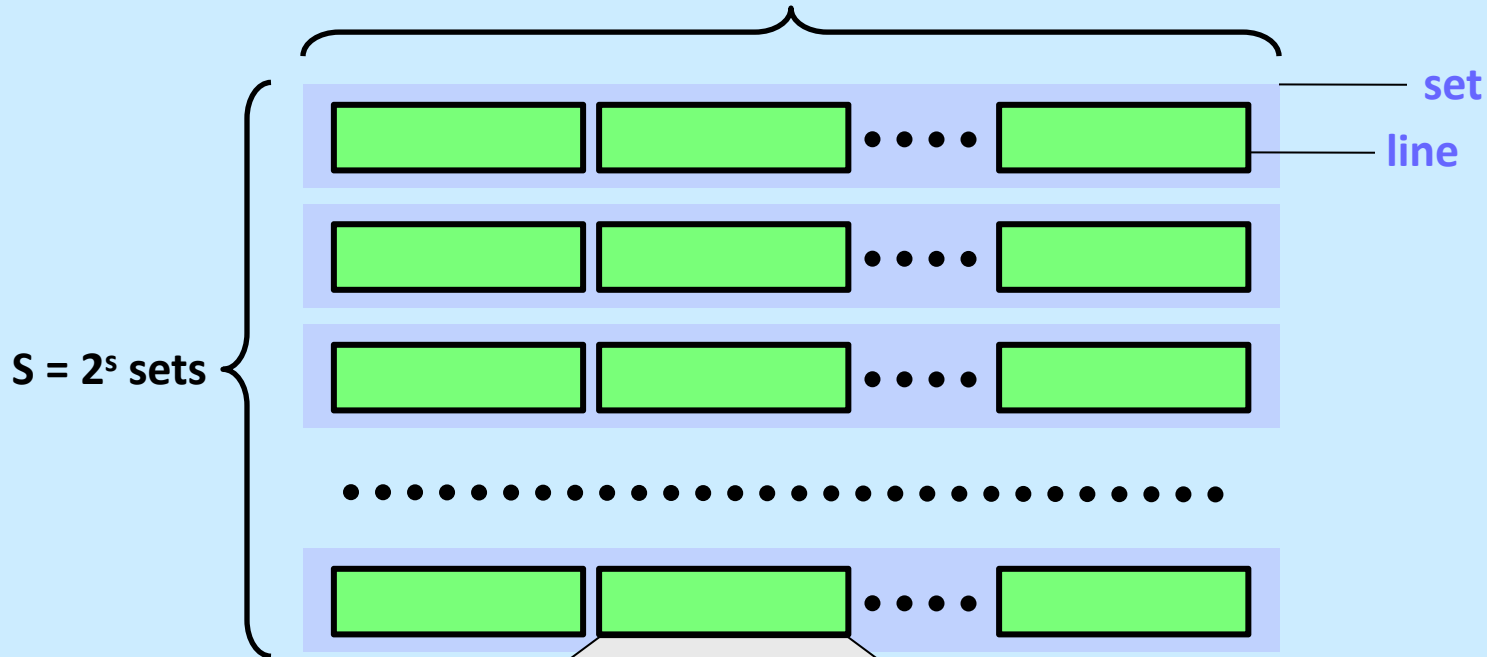
# General Caching Concepts:

## Types of Cache Misses

- **Cold (compulsory) miss**
  - cold misses occur because the cache is empty
- **Conflict miss**
  - most caches limit blocks to a small subset (sometimes a singleton) of the block positions in RAM
    - » e.g., block  $i$  in RAM must be placed in block  $(i \bmod 4)$  in the cache
  - conflict misses occur when the cache is large enough, but multiple data objects all map to the same cache block
    - » e.g., referencing blocks 0, 8, 0, 8, 0, 8, ... would miss every time
- **Capacity miss**
  - occurs when the set of active cache blocks (**working set**) is larger than the cache

# General Cache Organization (S, E, B)

$E = 2^e$  lines per set



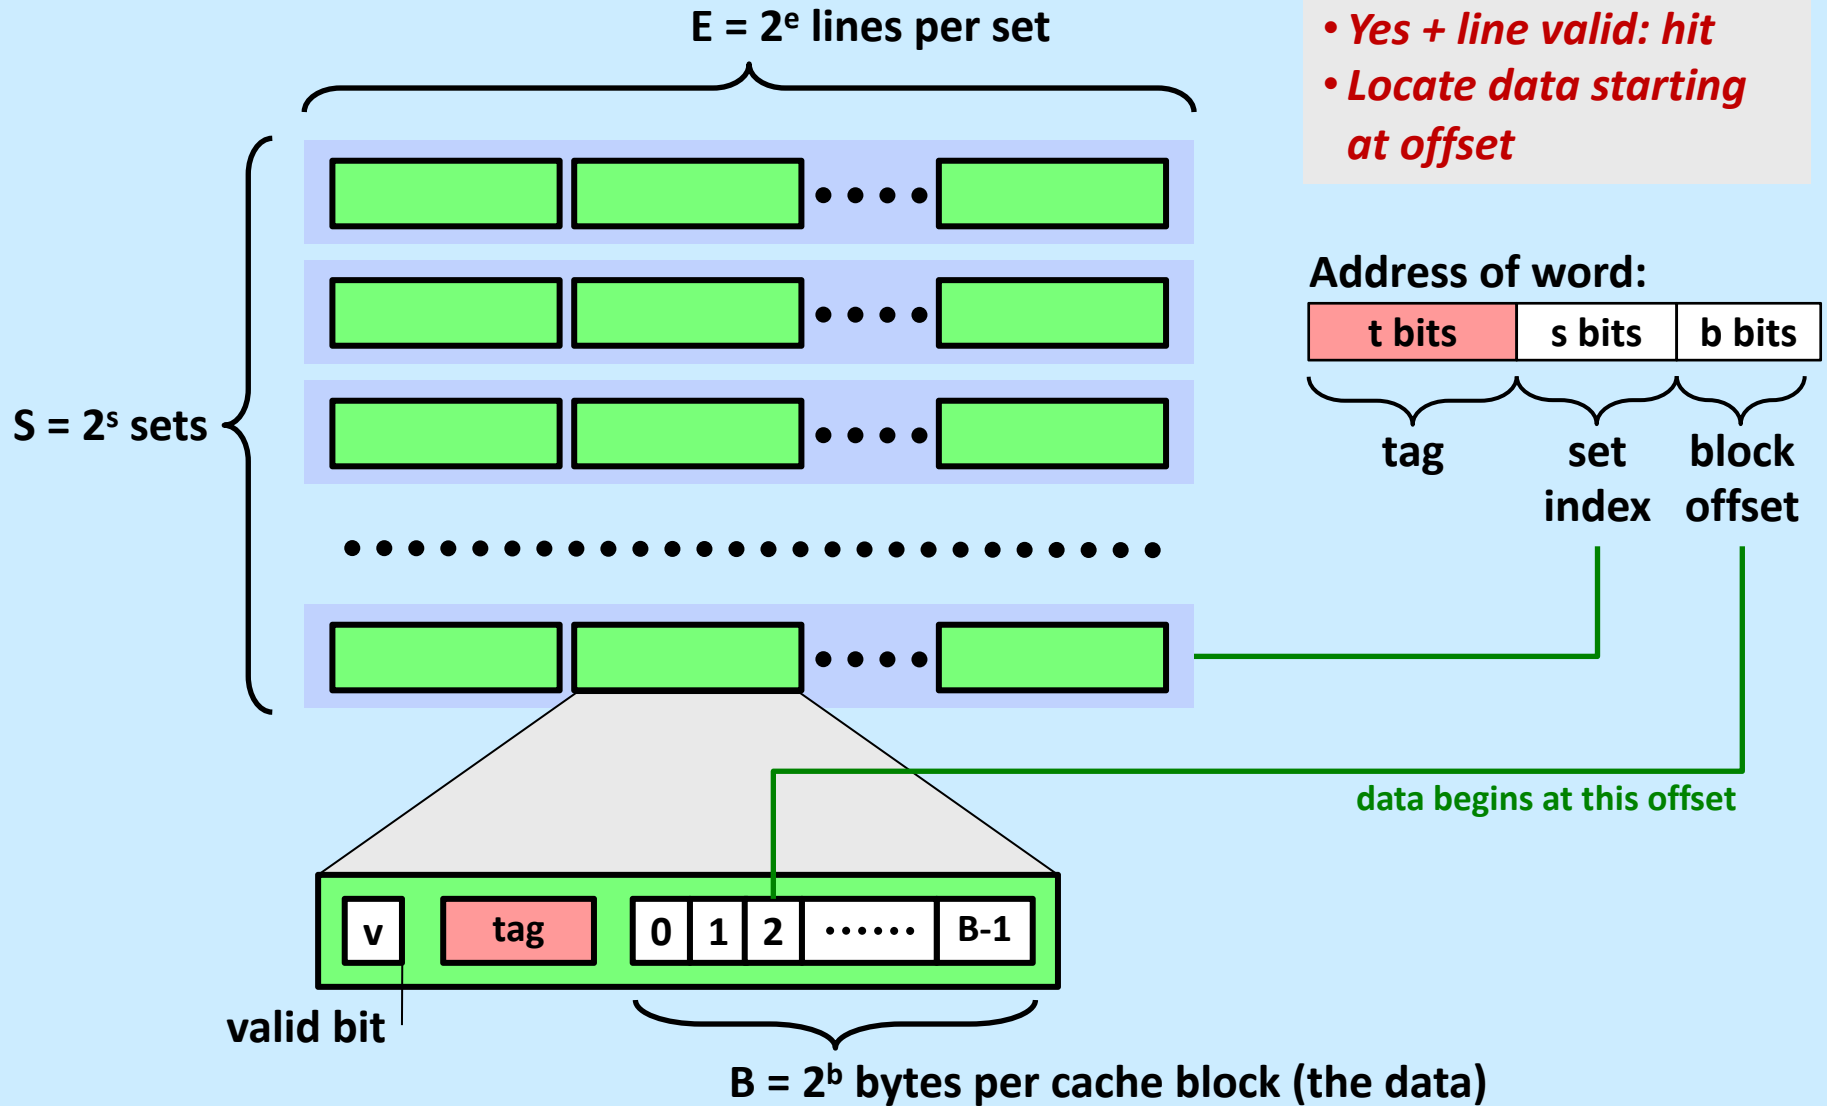
$B = 2^b$  bytes per cache block (the data)

**Cache size:**

$$C = S \times E \times B \text{ data bytes}$$

# Cache Read

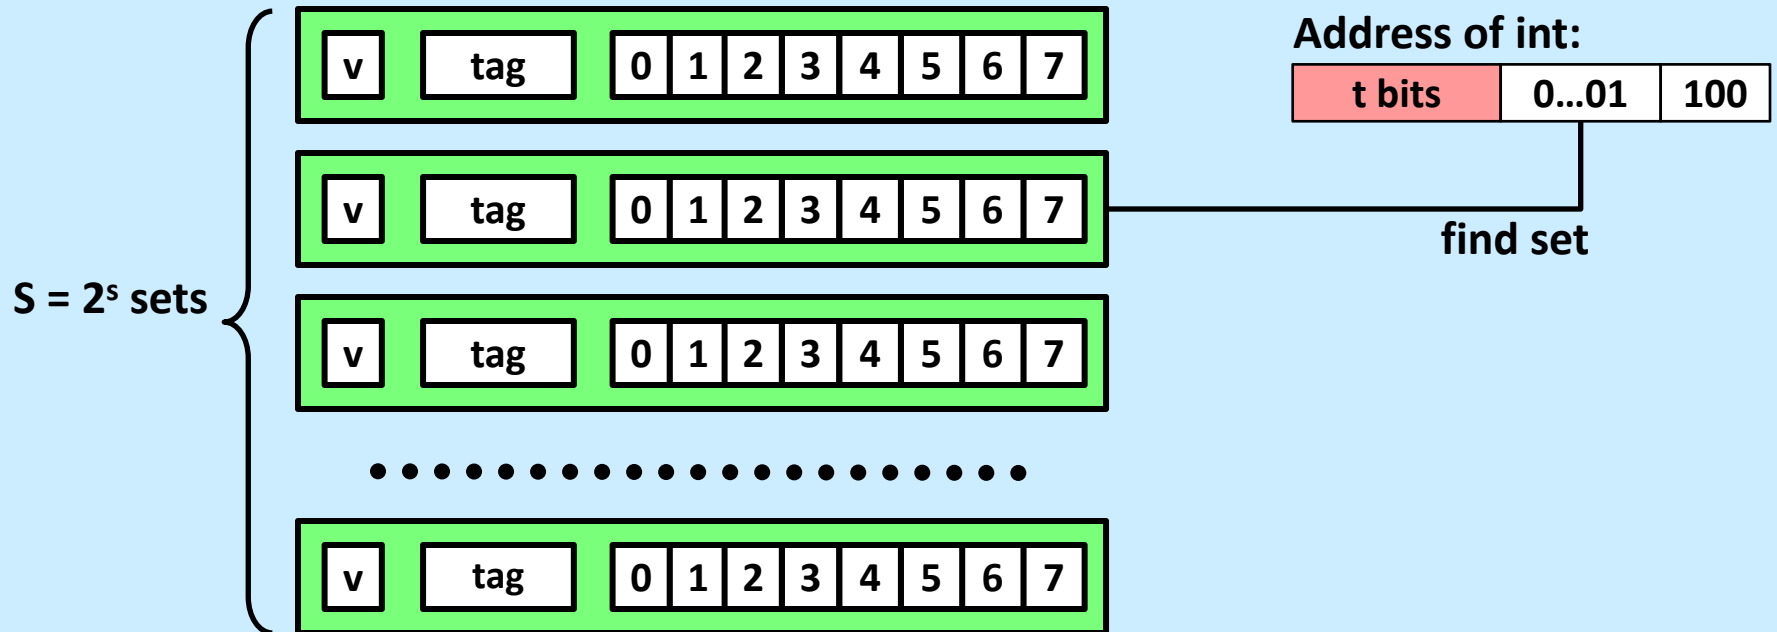
- *Locate set*
- *Check if any line in set has matching tag*
- *Yes + line valid: hit*
- *Locate data starting at offset*






# Example: Direct Mapped Cache (E = 1)

Direct mapped: one line per set  
Assume: cache block size 8 bytes



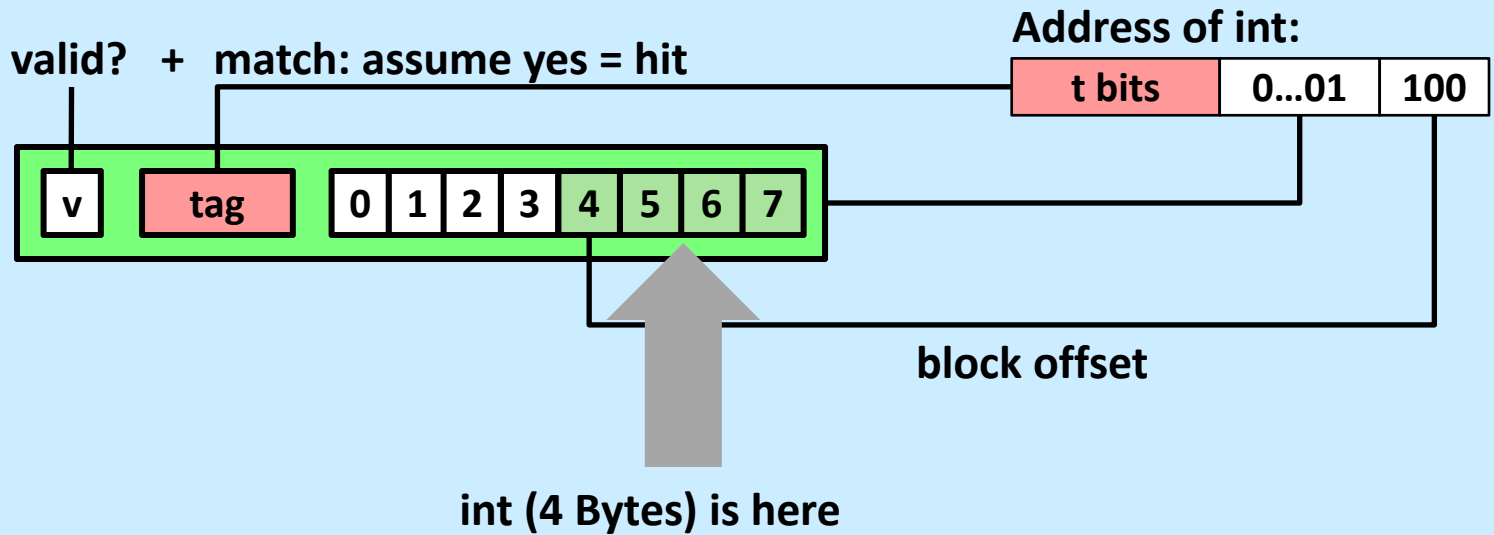
# Example: Direct Mapped Cache (E = 1)

Direct mapped: one line per set  
Assume: cache block size 8 bytes



# Example: Direct Mapped Cache (E = 1)

Direct mapped: one line per set  
Assume: cache block size 8 bytes



**No match:** old line is evicted and replaced

# Direct-Mapped Cache Simulation

t=1	s=2	b=1
x	xx	x

M=16 byte addresses, B=2 bytes/block,  
S=4 sets, E=1 Blocks/set

Address trace (reads, one byte per read):

0	[ <u>0000</u> <sub>2</sub> ],	miss
1	[ <u>0001</u> <sub>2</sub> ],	hit
7	[ <u>0111</u> <sub>2</sub> ],	miss
8	[ <u>1000</u> <sub>2</sub> ],	miss
0	[ <u>0000</u> <sub>2</sub> ]	miss

	v	Tag	Block
Set 0	1	0	M[0-1]
Set 1			
Set 2			
Set 3	1	0	M[6-7]

# A Higher-Level Example

*Ignore the variables sum, i, j*

assume: cold (empty) cache,  
a[0][0] goes here

```
int sum_array_rows(double a[16][16])
{
    int i, j;
    double sum = 0;

    for (i = 0; i < 16; i++)
        for (j = 0; j < 16; j++)
            sum += a[i][j];
    return sum;
}
```

```
int sum_array_cols(double a[16][16])
{
    int i, j;
    double sum = 0;

    for (j = 0; i < 16; i++)
        for (i = 0; j < 16; j++)
            sum += a[i][j];
    return sum;
}
```



32 B = 4 doubles


# A Higher-Level Example

```
int sum_array_rows(double a[16][16])
{
    int i, j;
    double sum = 0;

    for (i = 0; i < 16; i++)
        for (j = 0; j < 16; j++)
            sum += a[i][j];
    return sum;
}
```

```
int sum_array_cols(double a[16][16])
{
    int i, j;
    double sum = 0;

    for (j = 0; i < 16; i++)
        for (i = 0; j < 16; j++)
            sum += a[i][j];
    return sum;
}
```



32 B = 4 doubles


# A Higher-Level Example

```
int sum_array_rows(double a[16][16])
{
    int i, j;
    double sum = 0;

    for (i = 0; i < 16; i++)
        for (j = 0; j < 16; j++)
            sum += a[i][j];
    return sum;
}
```

```
int sum_array_cols(double a[16][16])
{
    int i, j;
    double sum = 0;

    for (j = 0; j < 16; j++)
        for (i = 0; i < 16; i++)
            sum += a[i][j];
    return sum;
}
```



32 B = 4 doubles

# Conflict Misses: Aligned

```
double dotprod(double x[8], double y[8]) {  
    double sum = 0.0;  
    int i;  
  
    for (i=0; i<8; i++)  
        sum += x[i] * y[i];  
  
    return sum;  
}
```




32 B = 4 doubles



# Different Alignments

```
double dotprod(double x[8], double y[8]) {  
    double sum = 0.0;  
    int i;  
  
    for (i=0; i<8; i++)  
        sum += x[i] * y[i];  
  
    return sum;  
}
```



32 B = 4 doubles