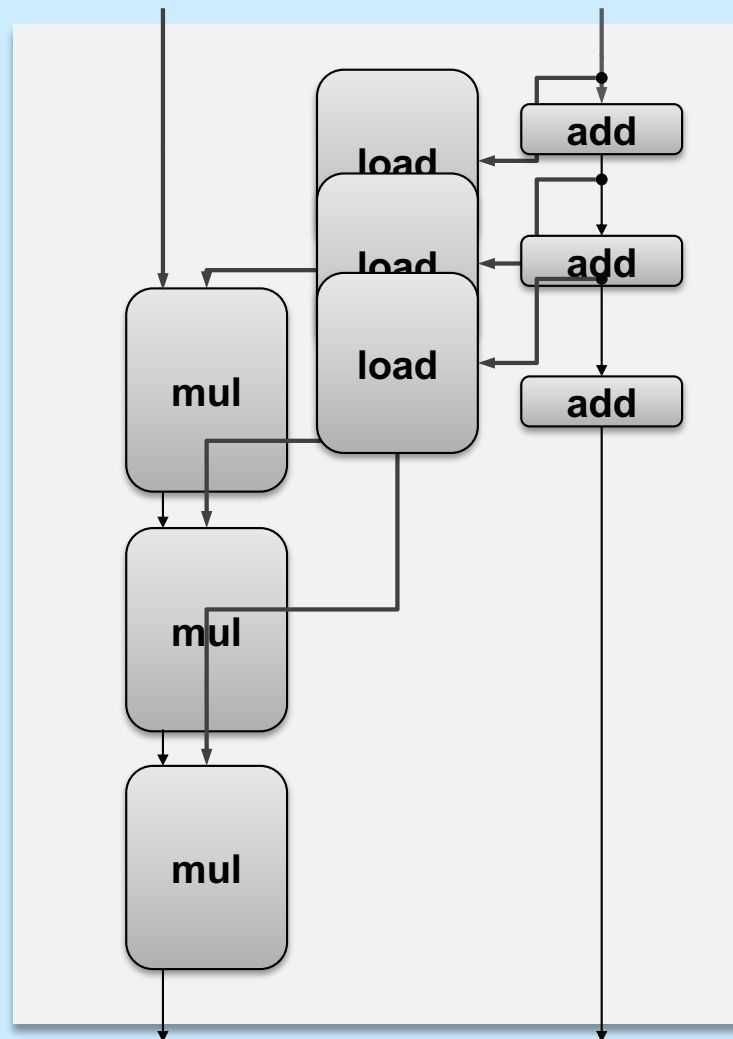# CS 33

Architecture and Optimization (3)

Eliminate Unneeded Memory References

```
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;
}
```

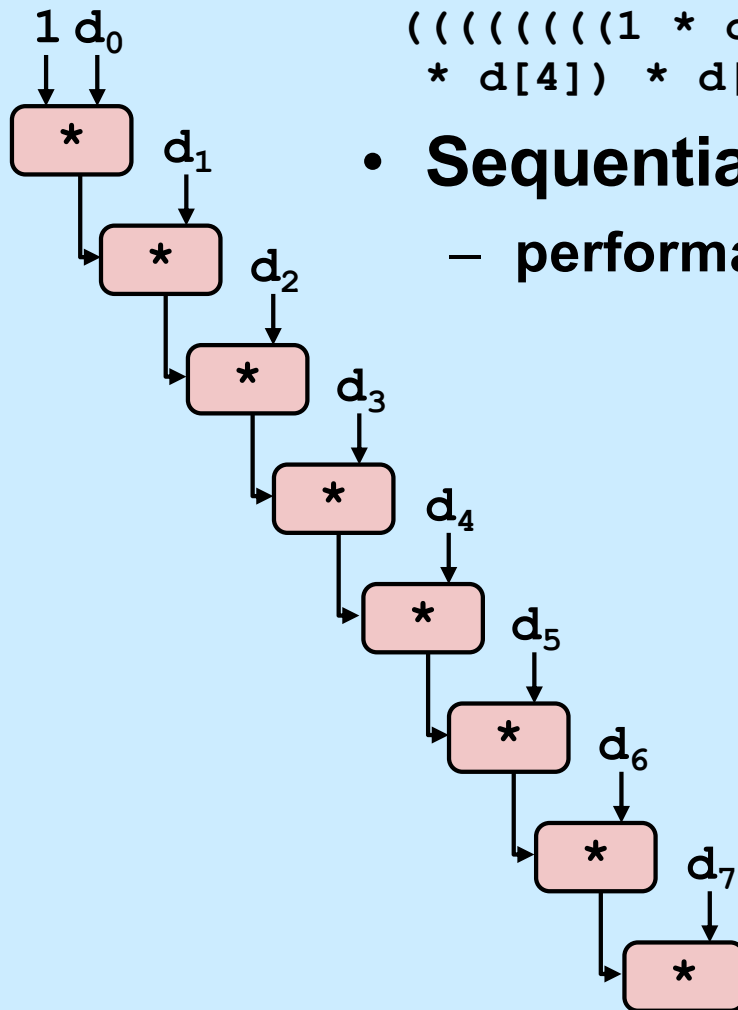| Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Combine1 –O1	12.0	12.0	12.0	13.0
Combine4	2.0	3.0	3.0	5.0

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

$1 \quad d_0$

$*$ $\quad d_1$

$*$ $\quad d_2$

$*$ $\quad d_3$

$*$ $\quad d_4$

$*$ $\quad d_5$

$*$ $\quad d_6$

$*$ $\quad d_7$

$*$

Loop Unrolling

```c
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**

XVI–5

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;
}
```
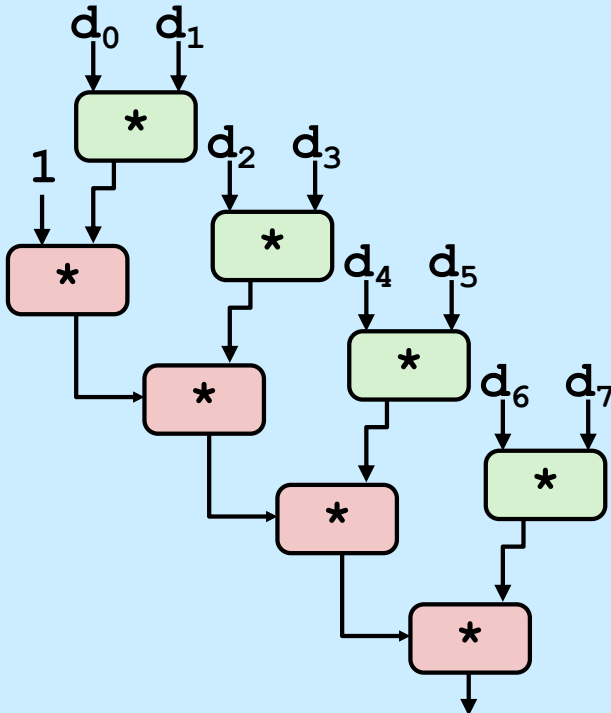
Compare to before

```
x = (x OP d[i]) 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	
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
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];
```
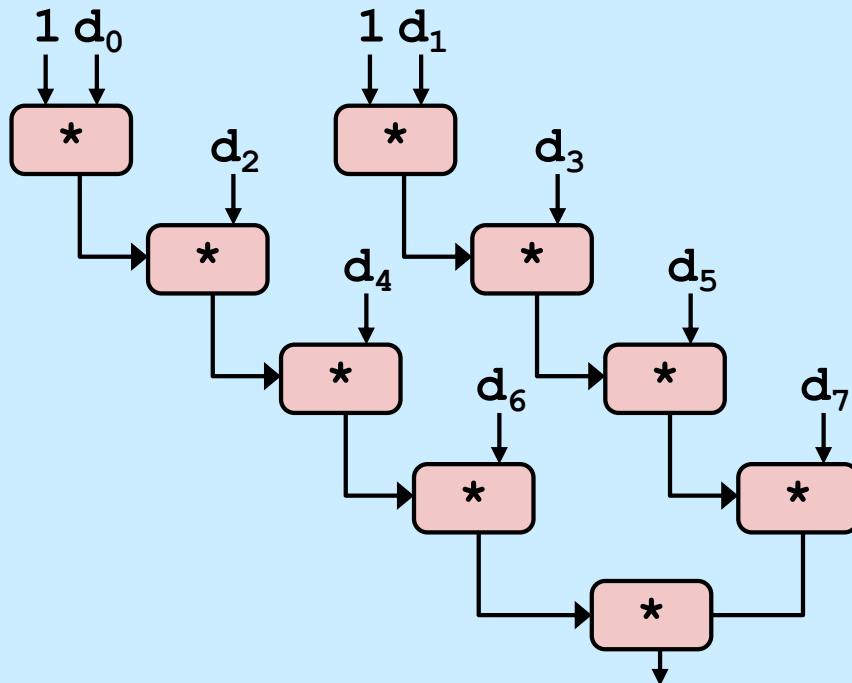
- **What changed:**
  - two independent "streams" of operations

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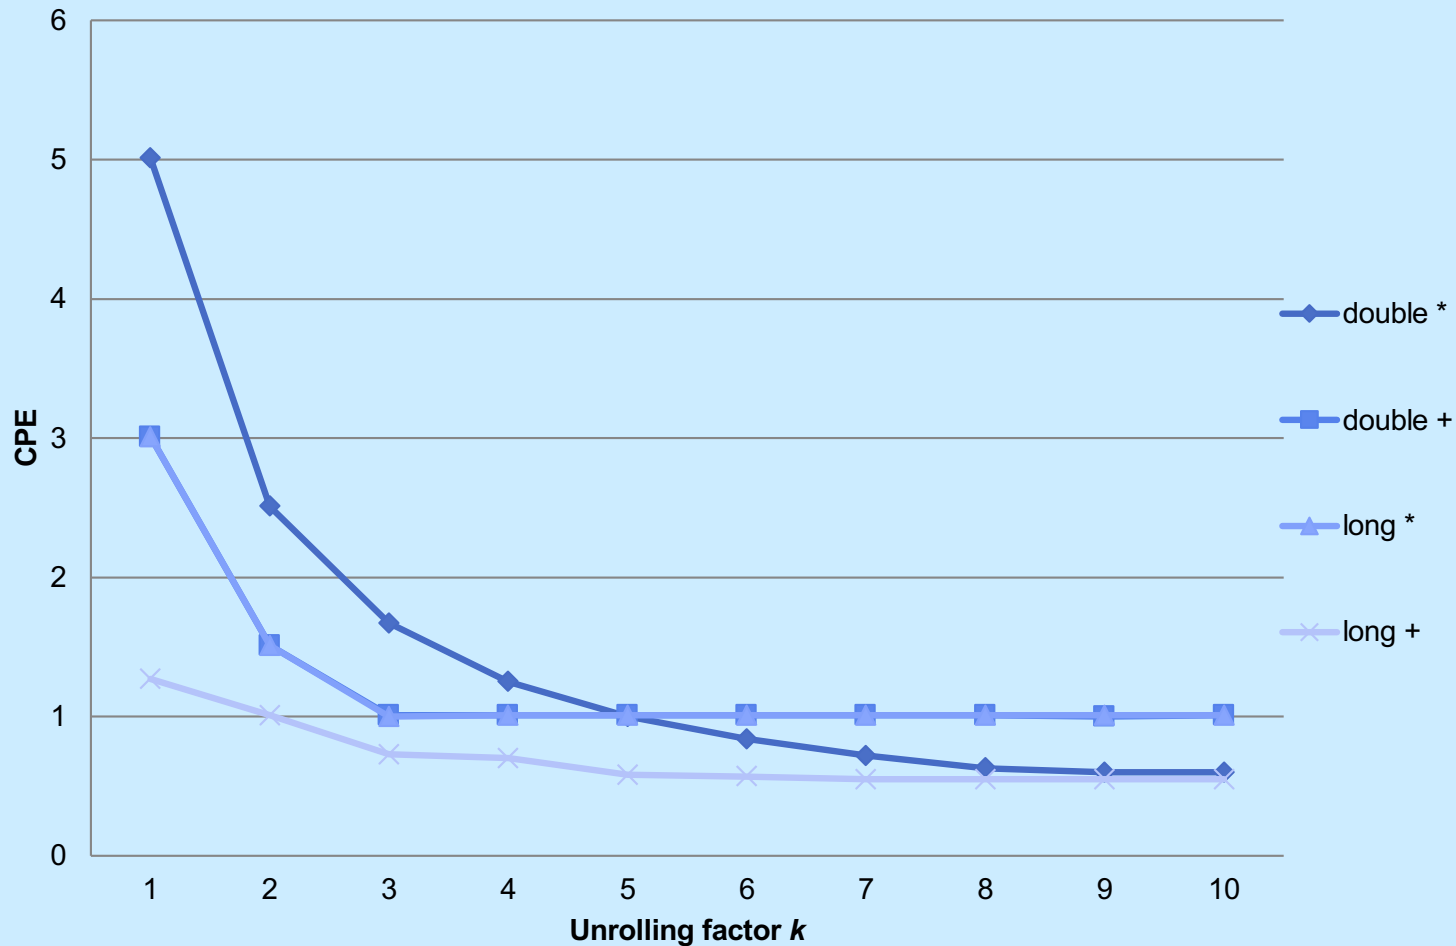
- **K-way loop unrolling with K accumulators**
 - **limited by number and throughput of functional units**

Achievable Performance

Method	Integer		Double FP	
Operation	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	
Operation	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

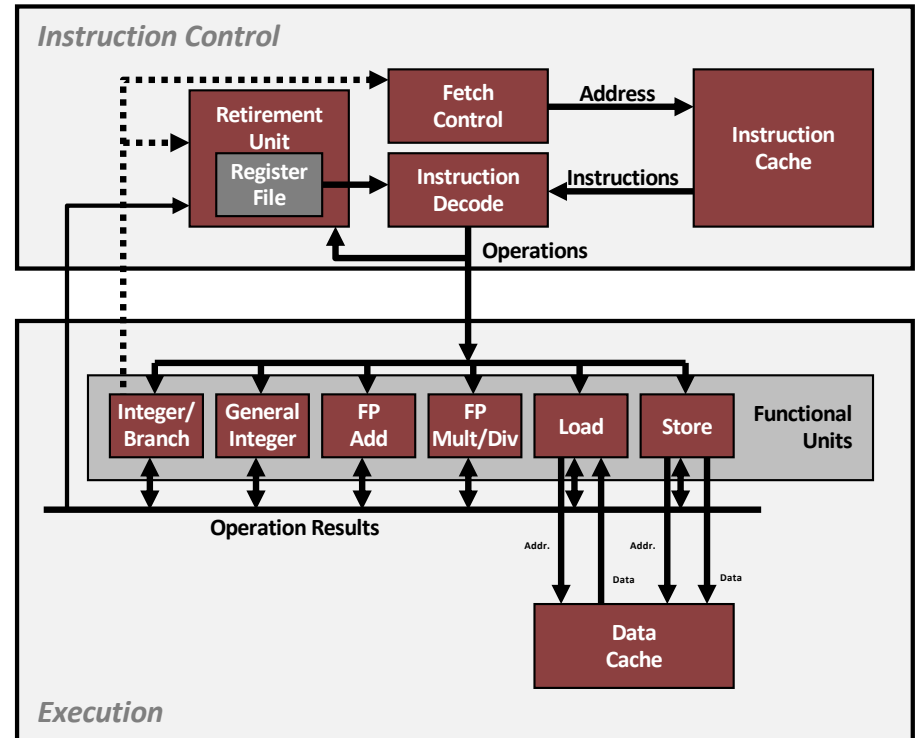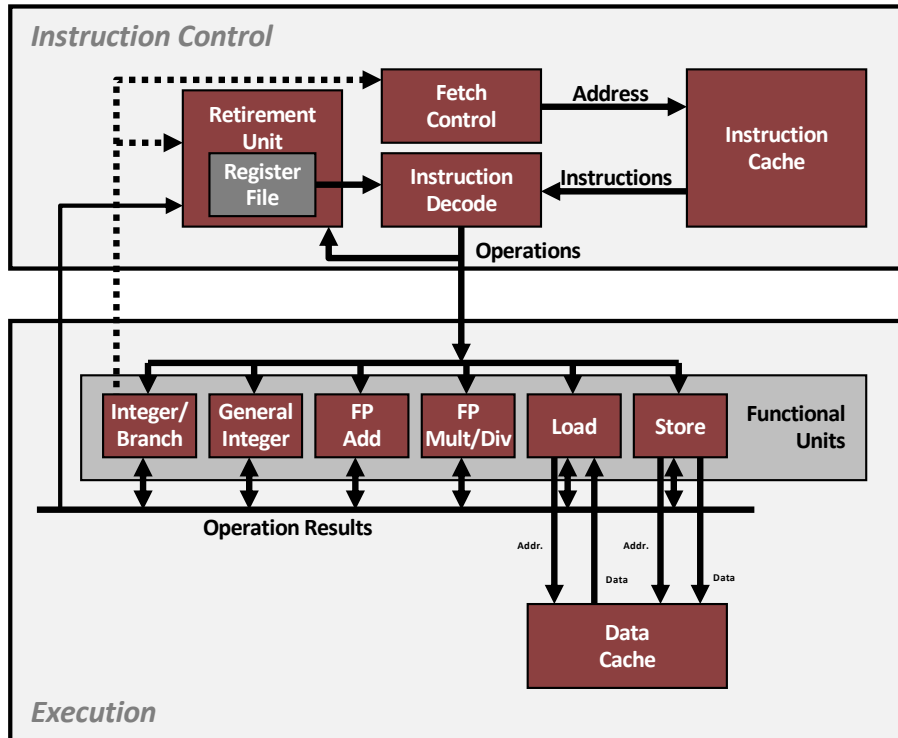- **Make use of SSE Instructions**
  - **parallel operations on multiple data elements**

Hyper Threading
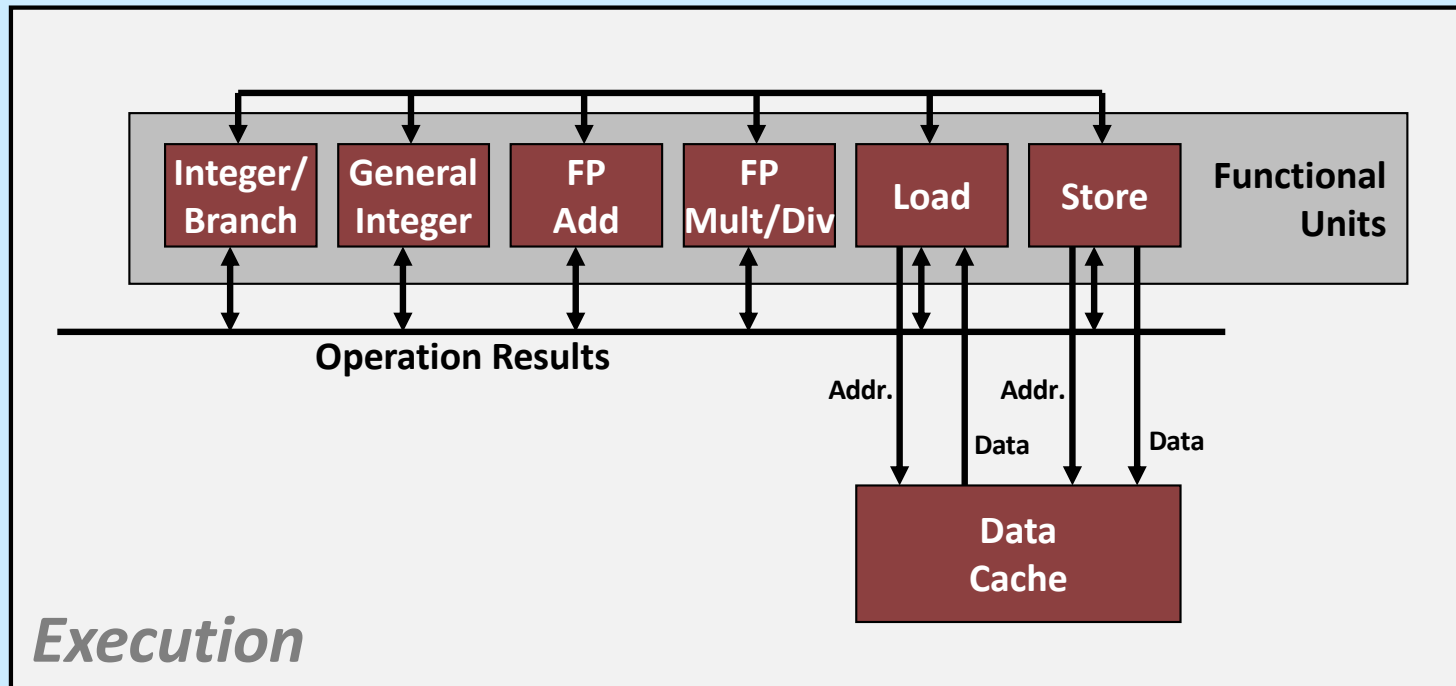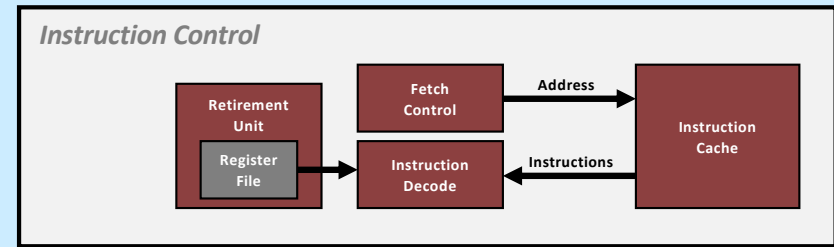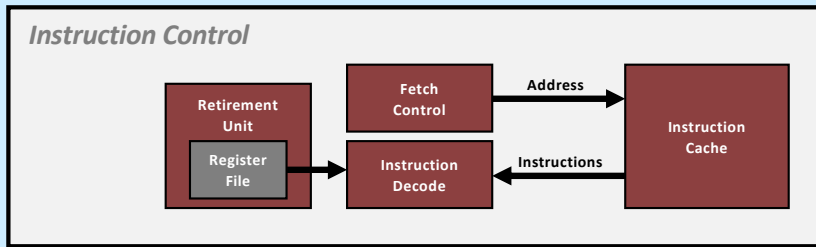
Multiple Cores

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 x w DRAM:
  - dw total bits organized as d **supercells** of size w bits

16 x 8 DRAM chip

(to/from CPU)

Memory controller

2 bits
/
addr

8 bits
/
data

cols

	0	1	2	3
0				
1				
2				
3				

rows

supercell (2,1)

Internal row buffer
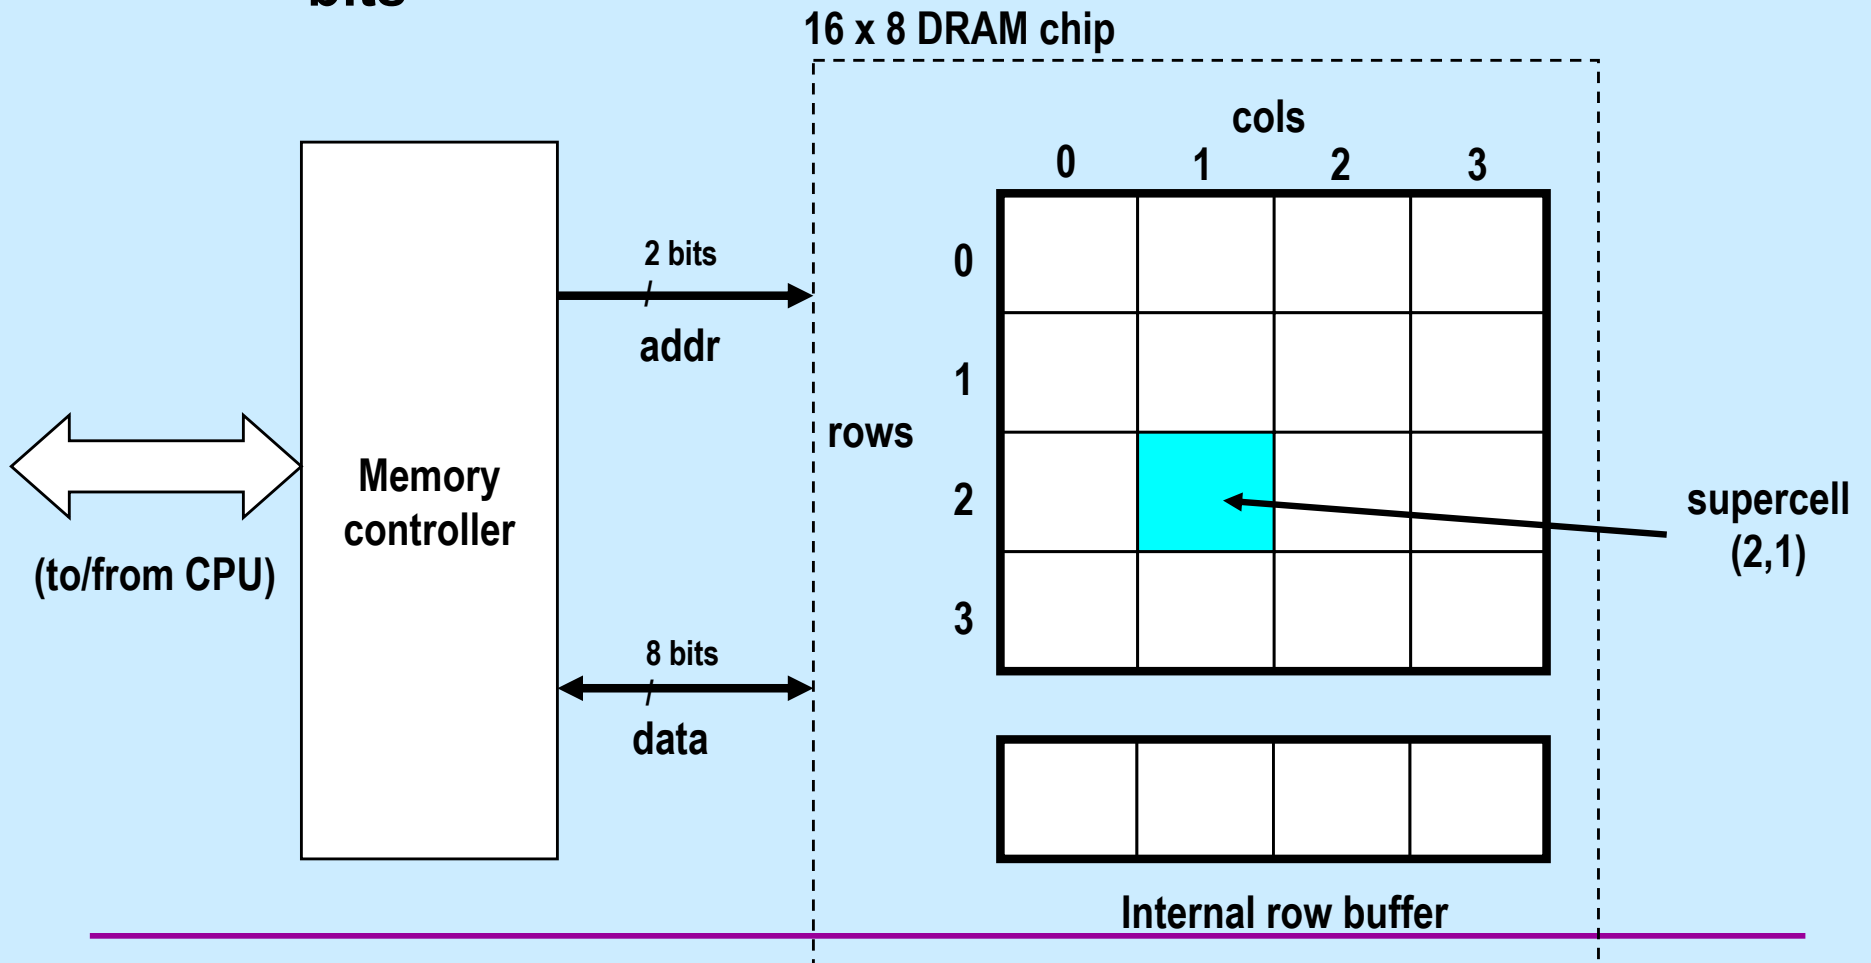
Reading DRAM Supercell (2,1)

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

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

16 x 8 DRAM chip

Cols

RAS = 2

Memory controller

addr

data

Rows

Internal row buffer
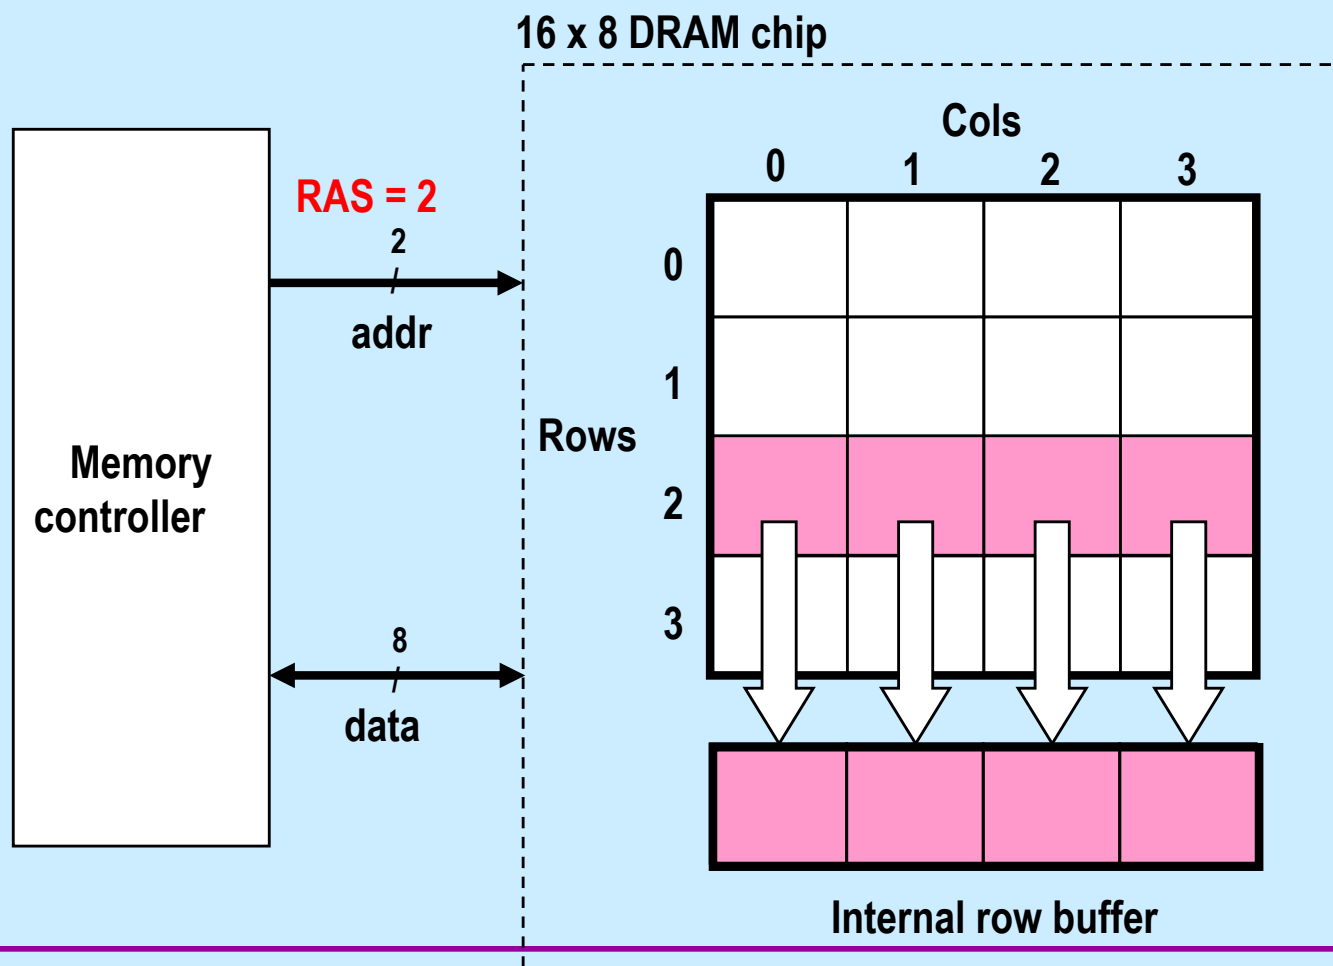
Reading DRAM Supercell (2,1)

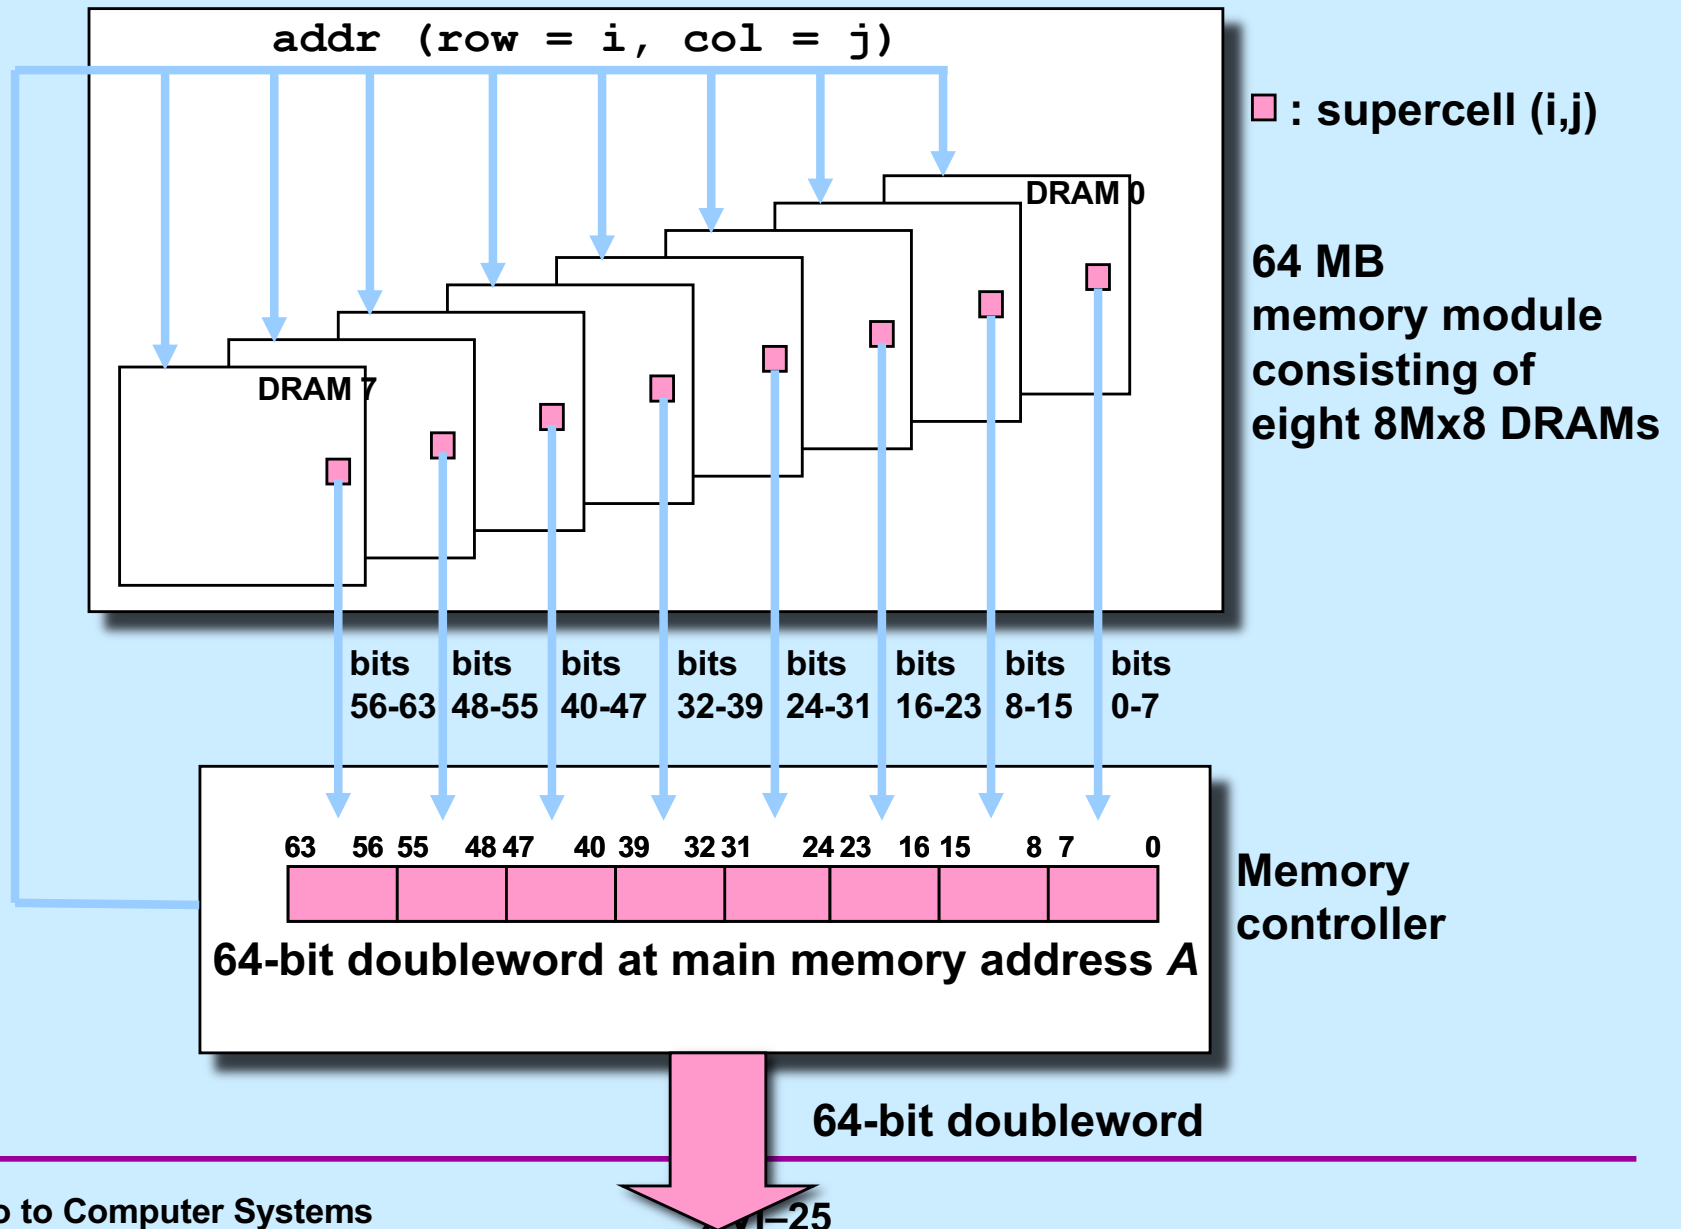**Step 2(a): column access strobe (CAS) selects column 1**

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

16 x 8 DRAM chip

To CPU

Memory controller

CAS = 1

2

addr

8

data

supercell (2,1)

supercell (2,1)

Cols

0   1   2   3

Rows

0

1

2

3

Internal row buffer

Memory Modules

addr (row = i, col = j)

■ : supercell (i,j)

64 MB
memory module
consisting of
eight 8Mx8 DRAMs

DRAM 0

DRAM 7

| bits 56-63 | bits 48-55 | bits 40-47 | bits 32-39 | bits 24-31 | bits 16-23 | bits 8-15 | bits 0-7 |

63   56 55   48 47   40 39   32 31   24 23   16 15   8 7   0

Memory controller

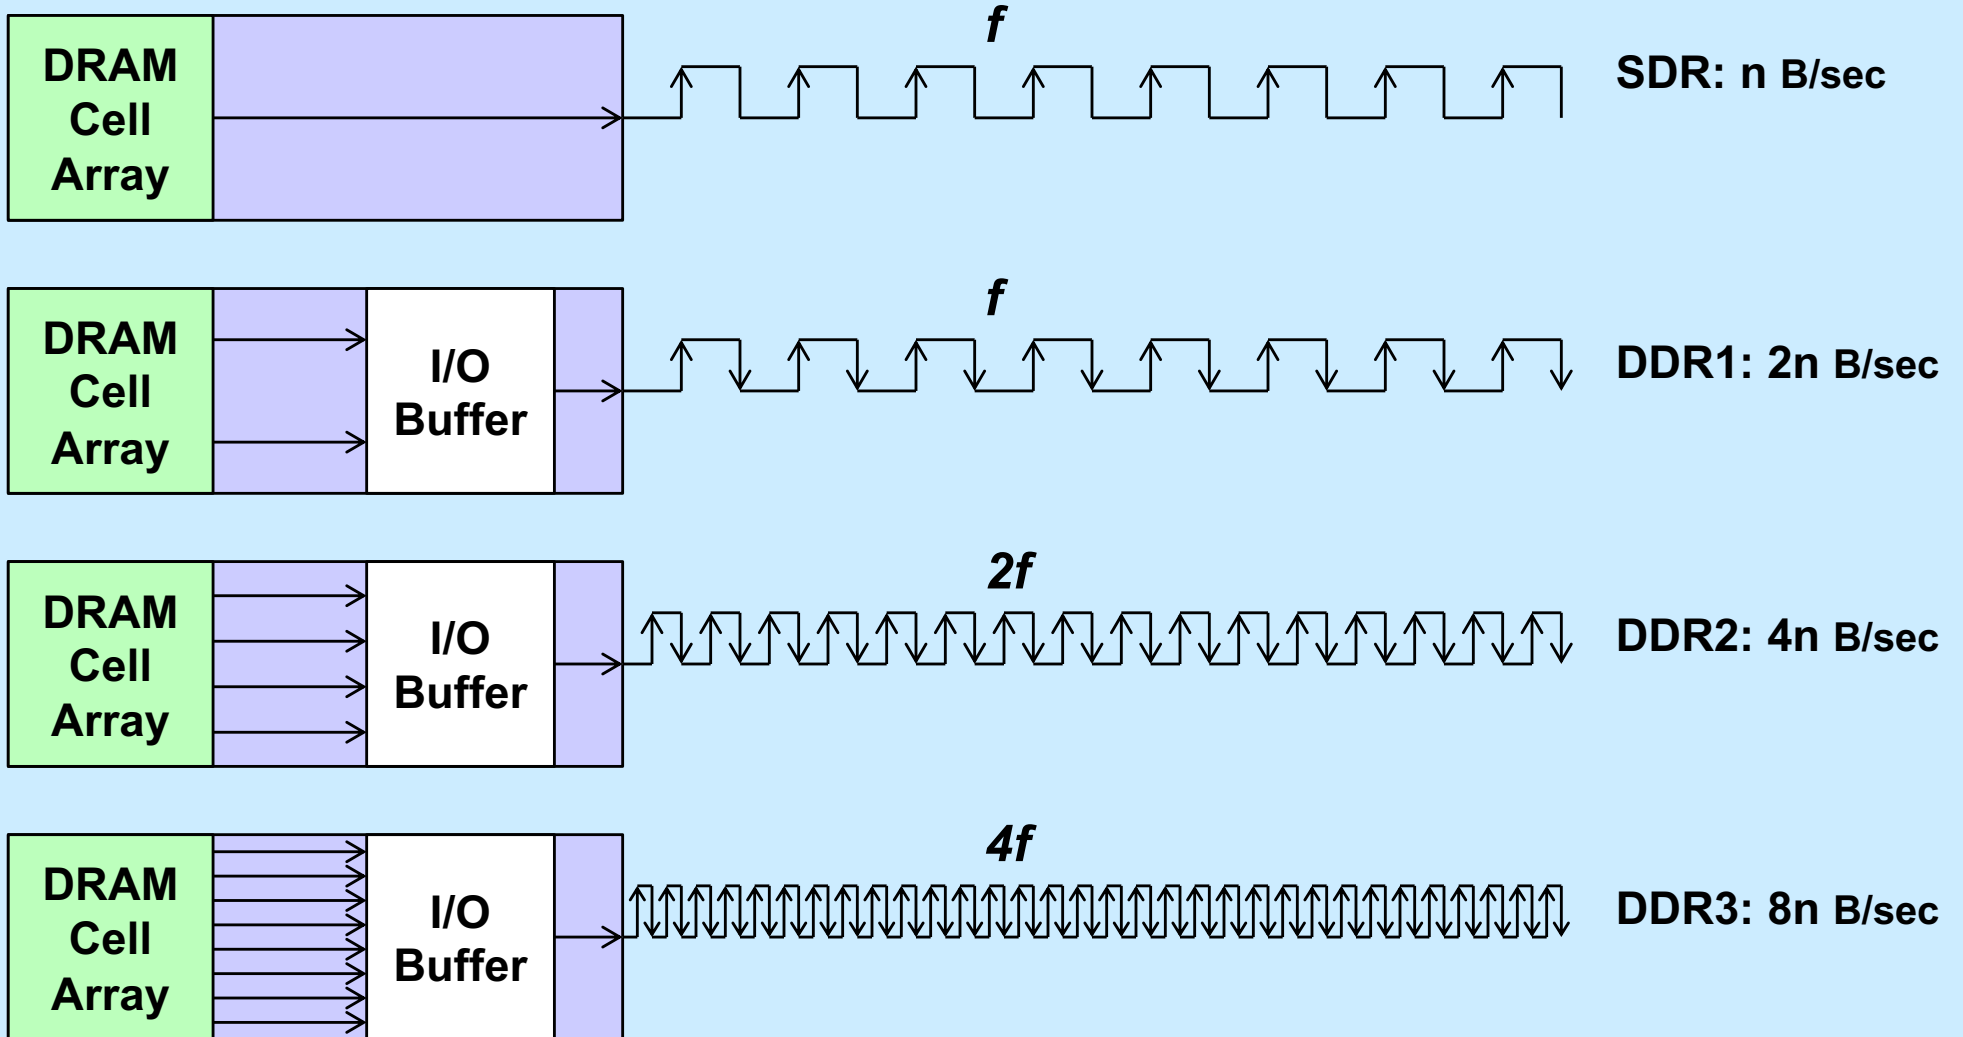**64-bit doubleword at main memory address *A***

64-bit doubleword

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

DRAM Cell Array		f — SDR: n B/sec
DRAM Cell Array	**I/O Buffer**	f — DDR1: 2n B/sec
DRAM Cell Array	**I/O Buffer**	$2f$ — DDR2: 4n B/sec
DRAM Cell Array	**I/O Buffer**	$4f$ — DDR3: 8n B/sec

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**
  - 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?**
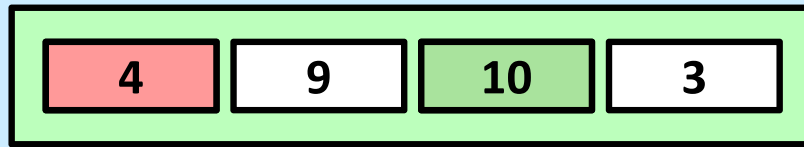
Caching to the Rescue

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
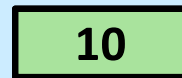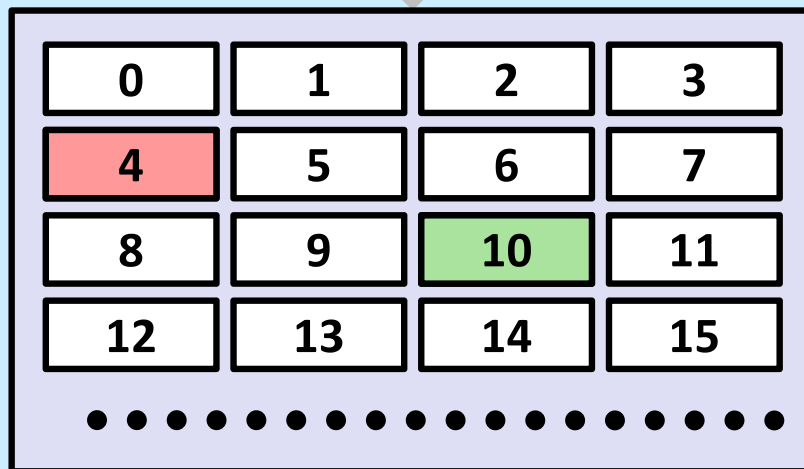- Typical system structure:

CPU chip

Register file

Cache memories

ALU

System bus   **Memory bus**

Bus interface

I/O bridge

Main memory
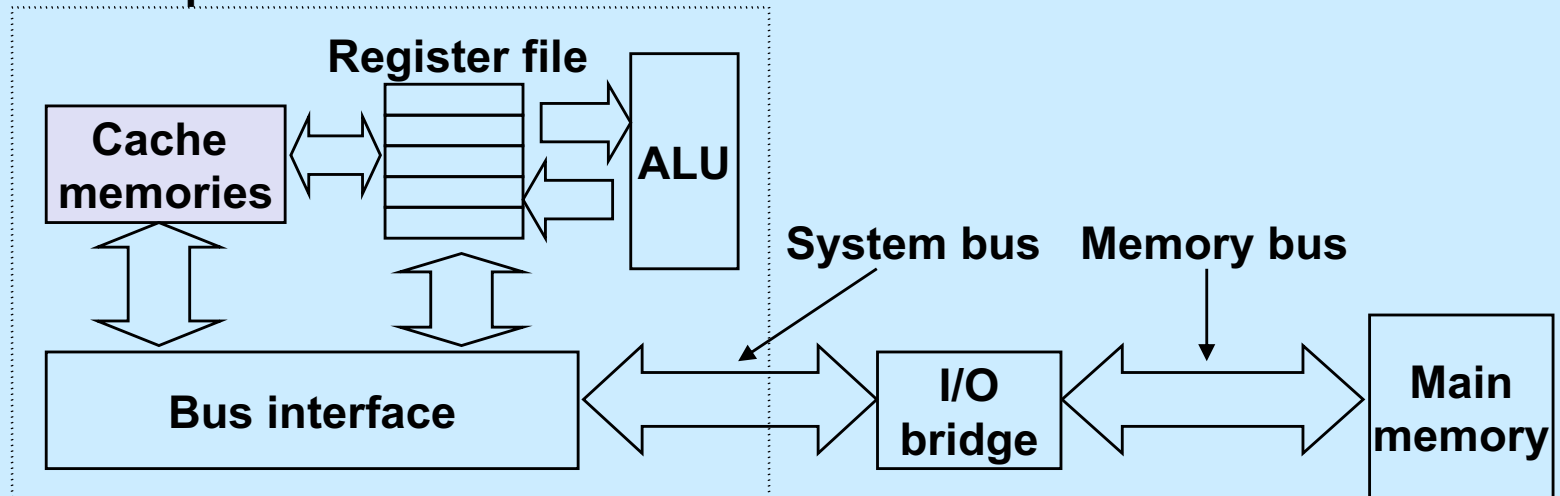
General Cache Concepts

Cache

4	9	10	3

Smaller, faster, more expensive memory caches a subset of the blocks

10

Data is copied in block-sized transfer units

Memory

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15

Larger, slower, cheaper memory viewed as partitioned into "blocks"

General Cache Concepts: Hit

Request: 14

Cache

| 8 | 9 | 14 | 3 |

Memory

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15

• • • • • • • • • • • • • • • • • •

Data in block b is needed

Block b is in cache:
Hit!

General Cache Concepts: Miss

Request: 12

Cache

8	12	14	3

12

Request: 12

Memory

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15

● ● ● ● ● ● ● ● ● ● ● ● ● ● ● ● ●

Data in block b is needed

Block b is not in cache:
Miss!

Block b is fetched from memory

Block b is stored in cache
- placement policy: determines where b goes
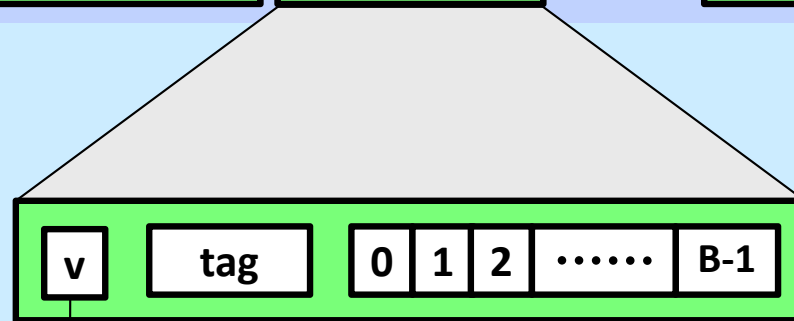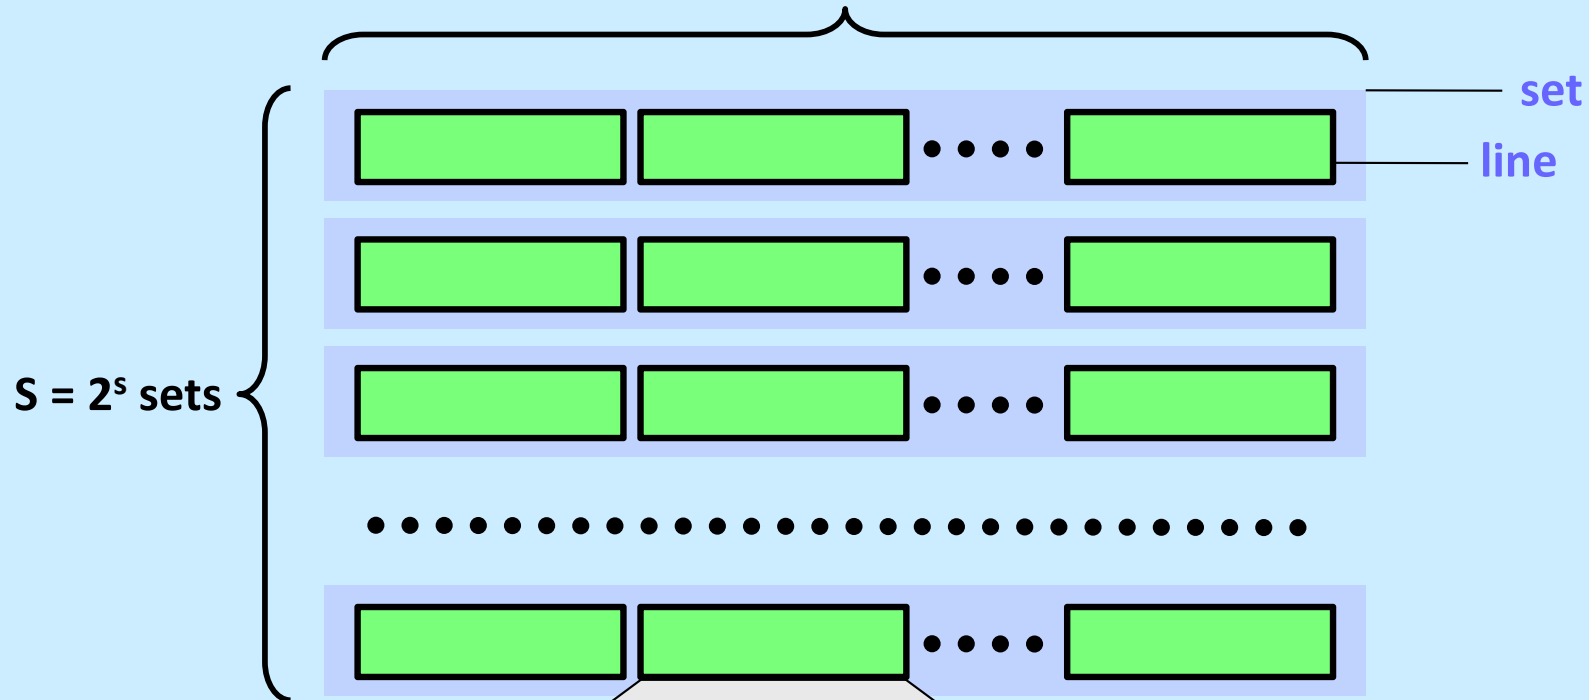- replacement policy: determines which block gets evicted (victim)
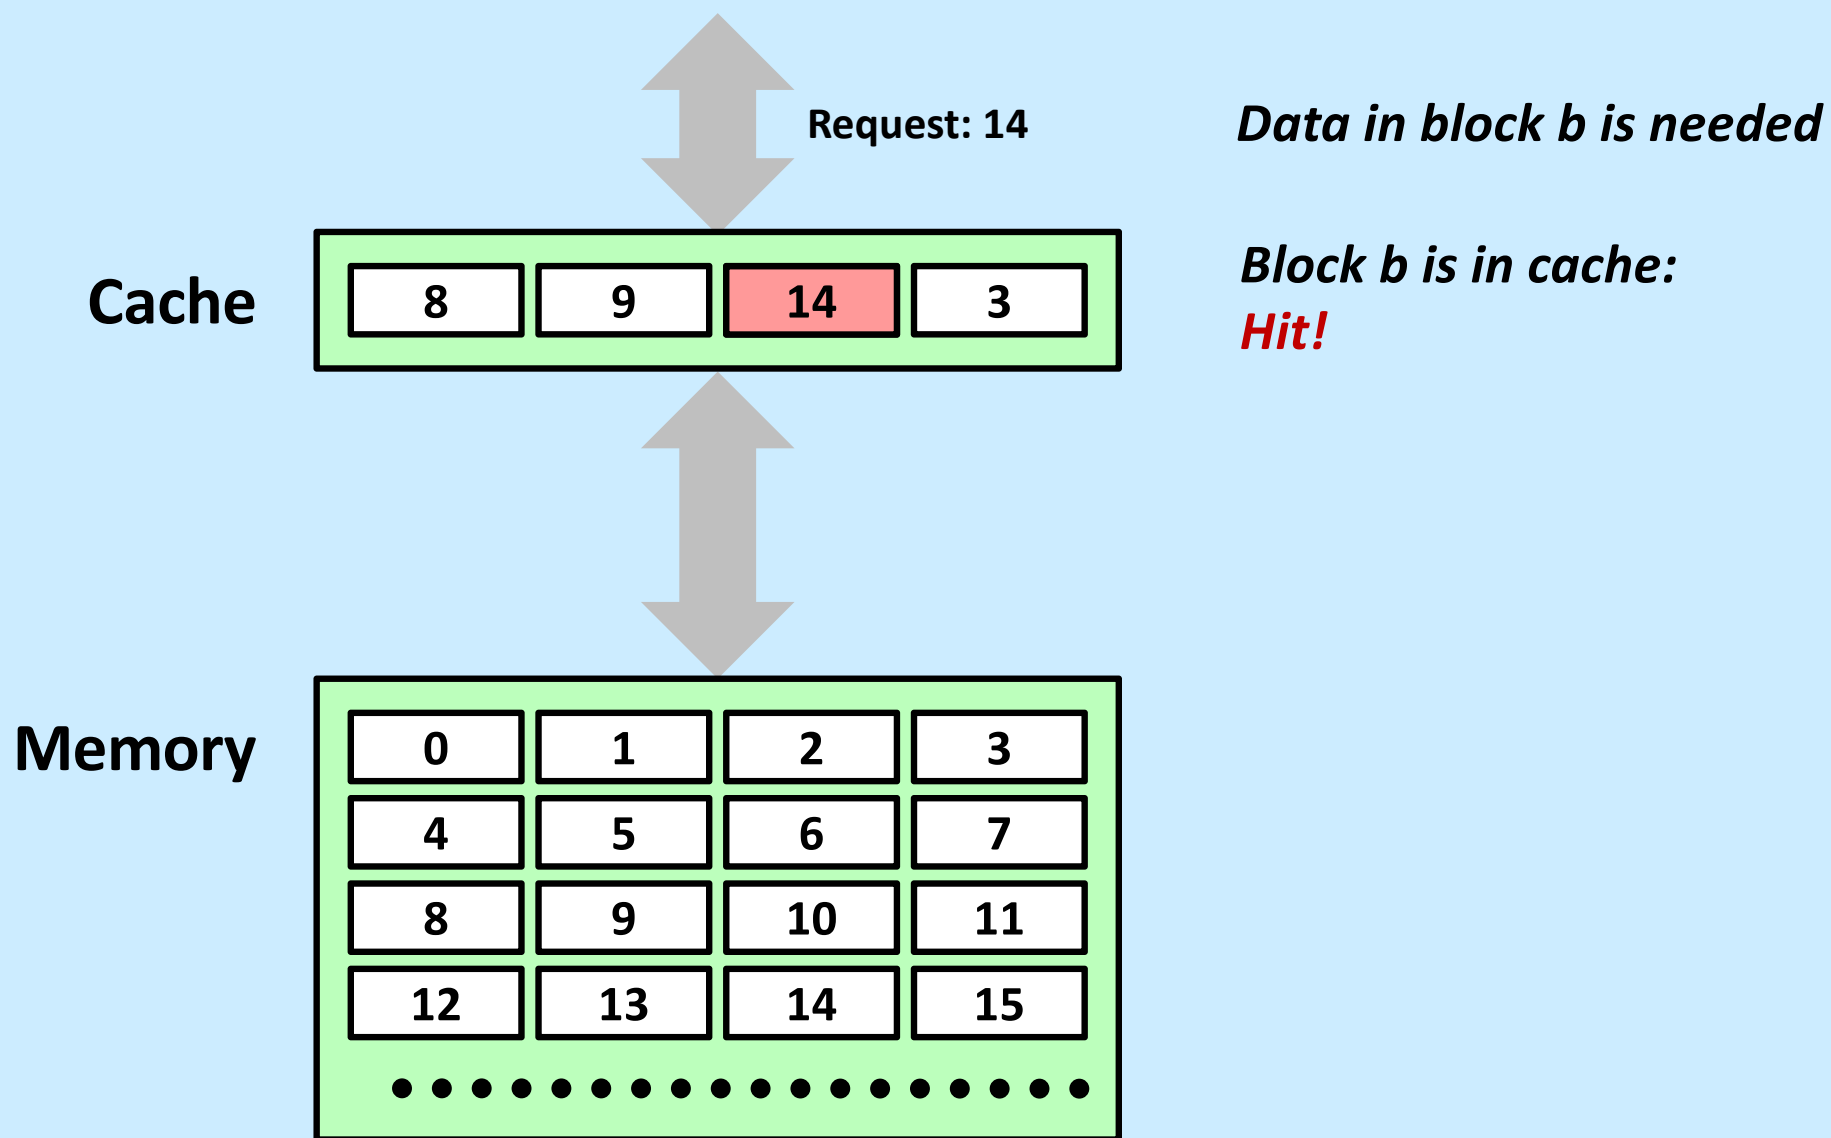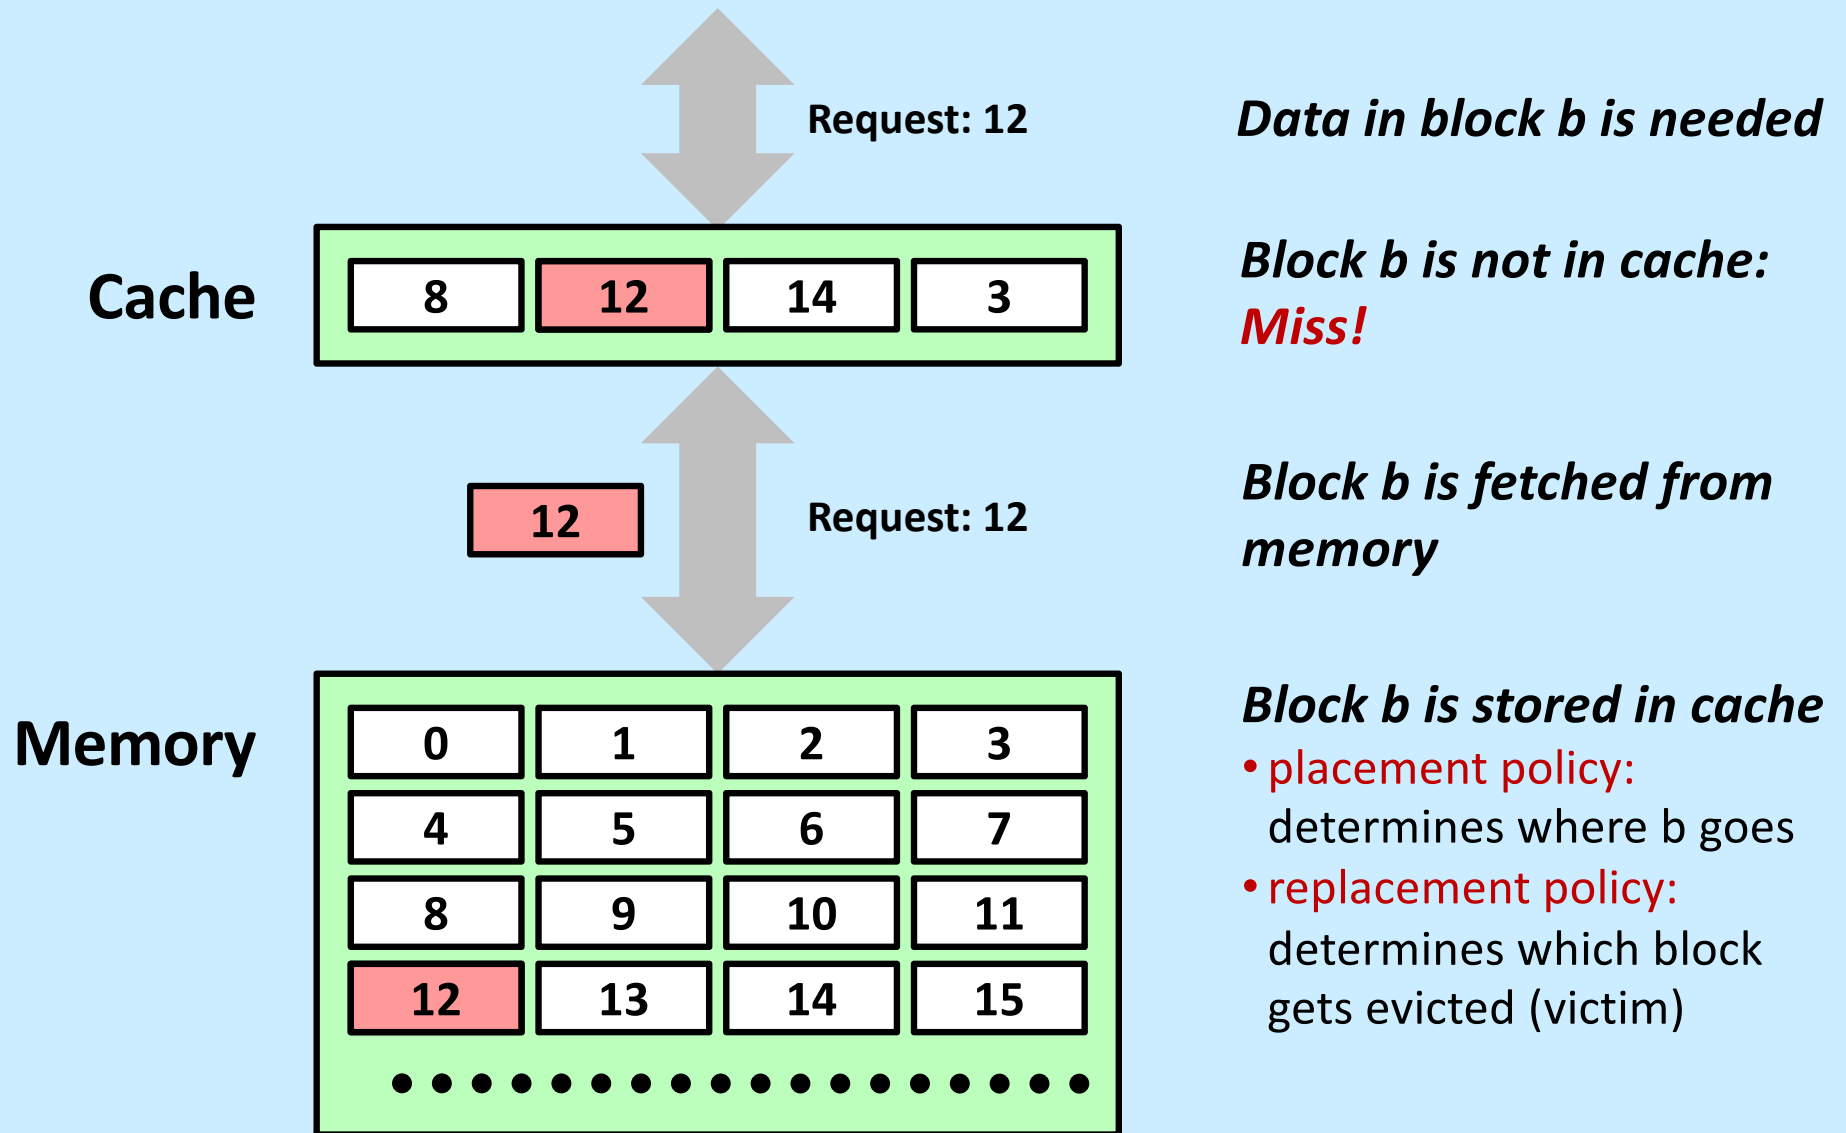
General Caching Concepts: Types of Cache Misses

- **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 mod 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

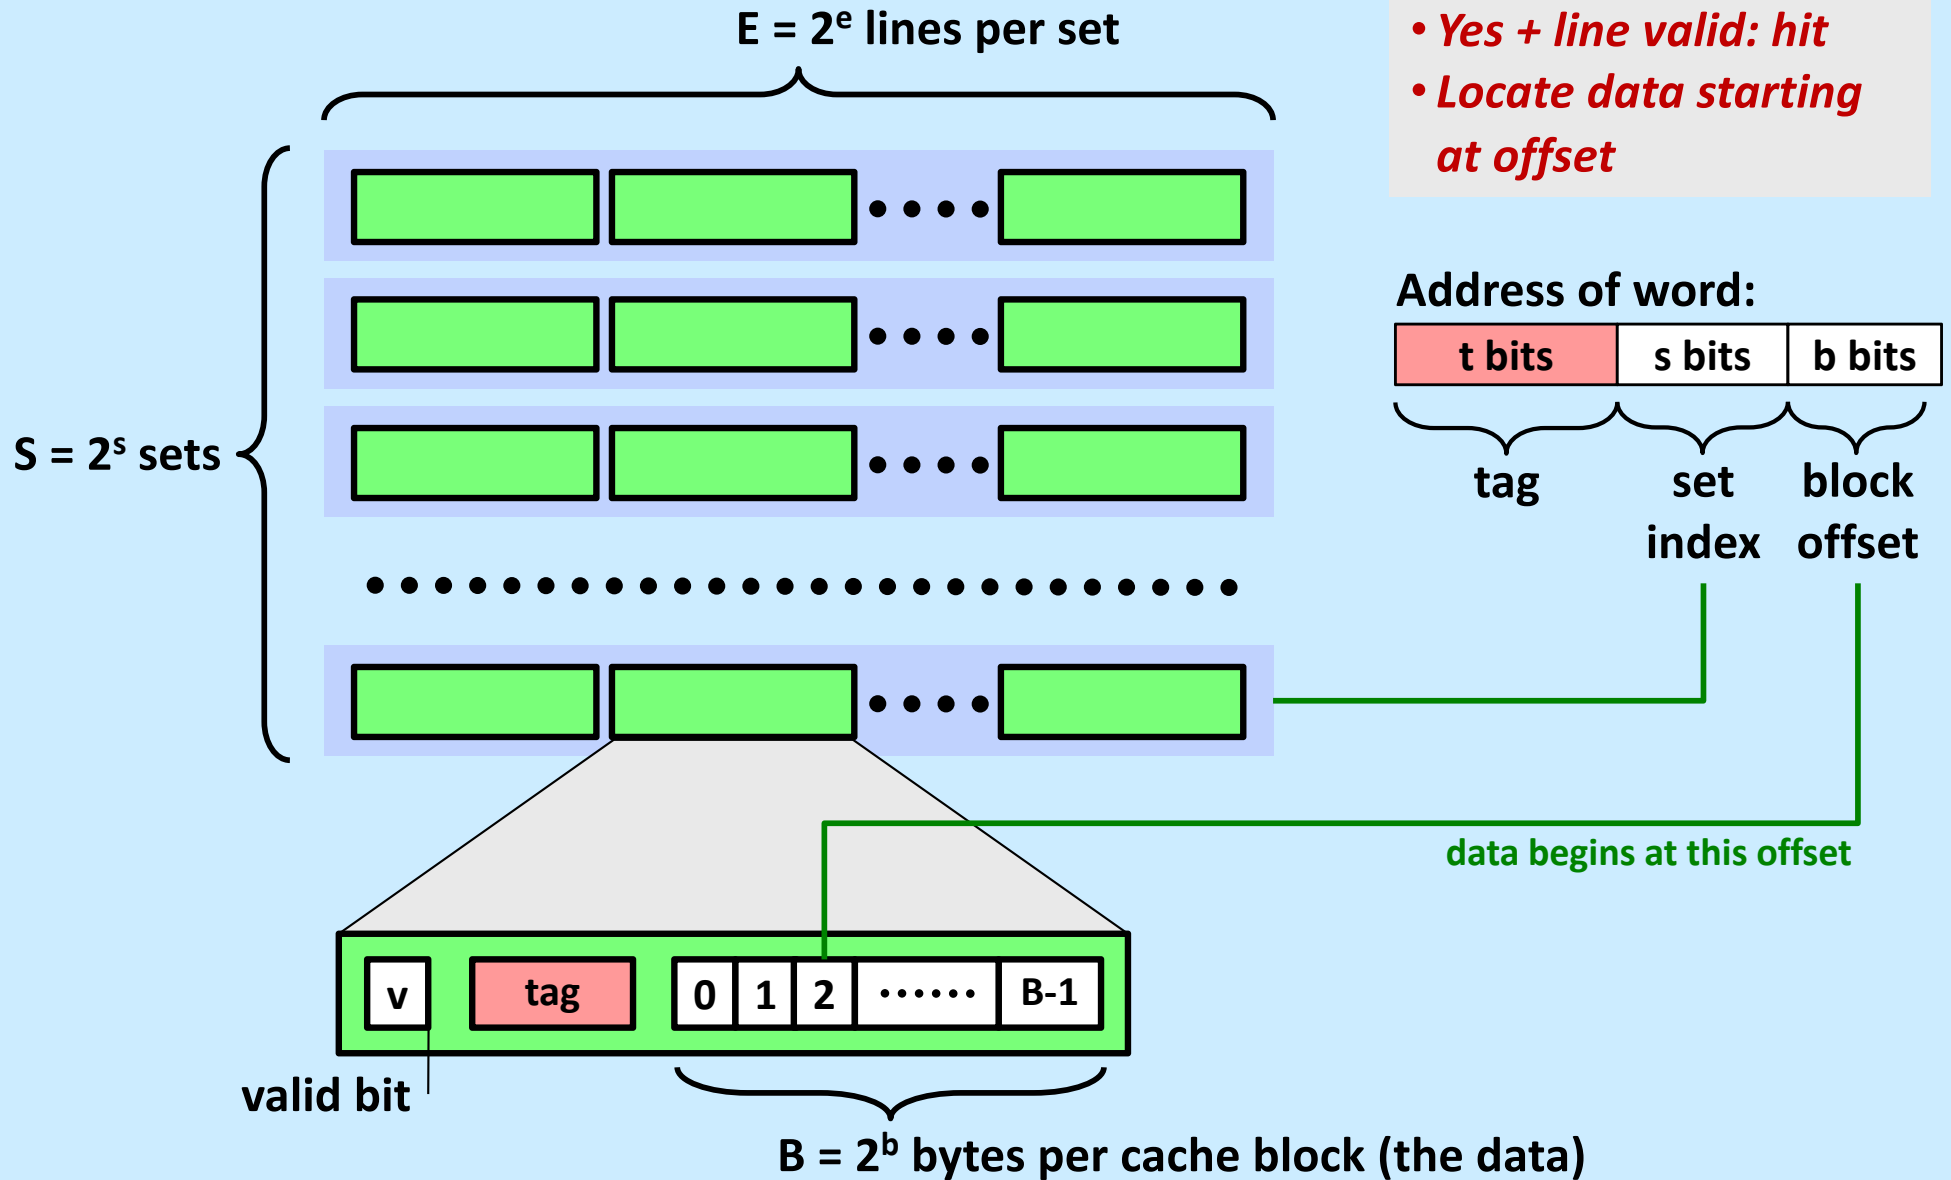# General Cache Organization (S, E, B)

$E = 2^e$ lines per set

set

line

$S = 2^s$ sets

valid bit

v | tag | 0 | 1 | 2 | ...... | B-1

Cache size:

C = S x E x B data bytes

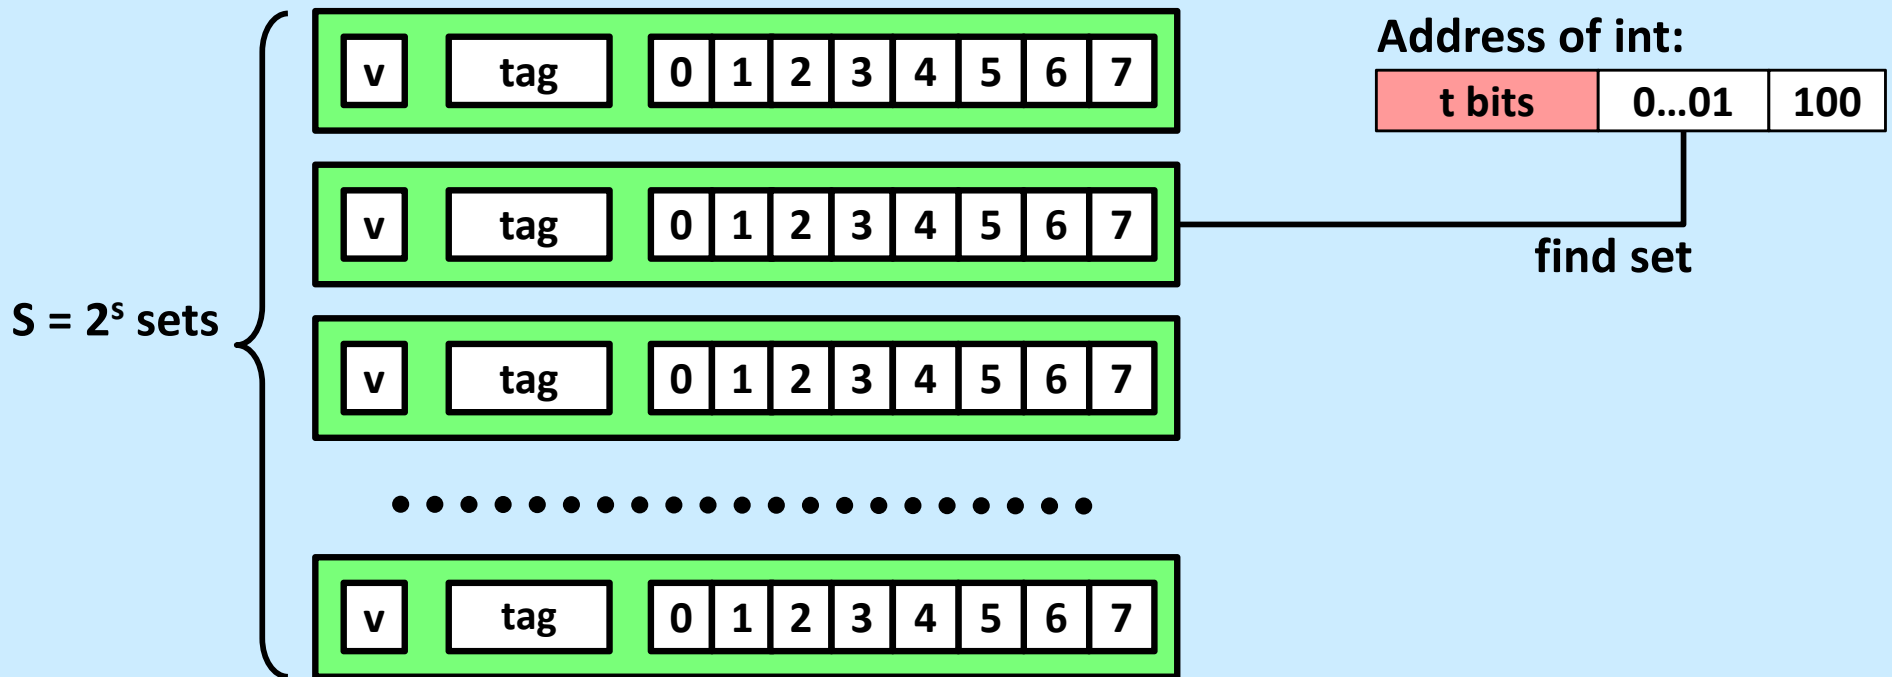$B = 2^b$ bytes per cache block (the data)

Cache Read

E = 2^e lines per set

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

S = 2^s sets
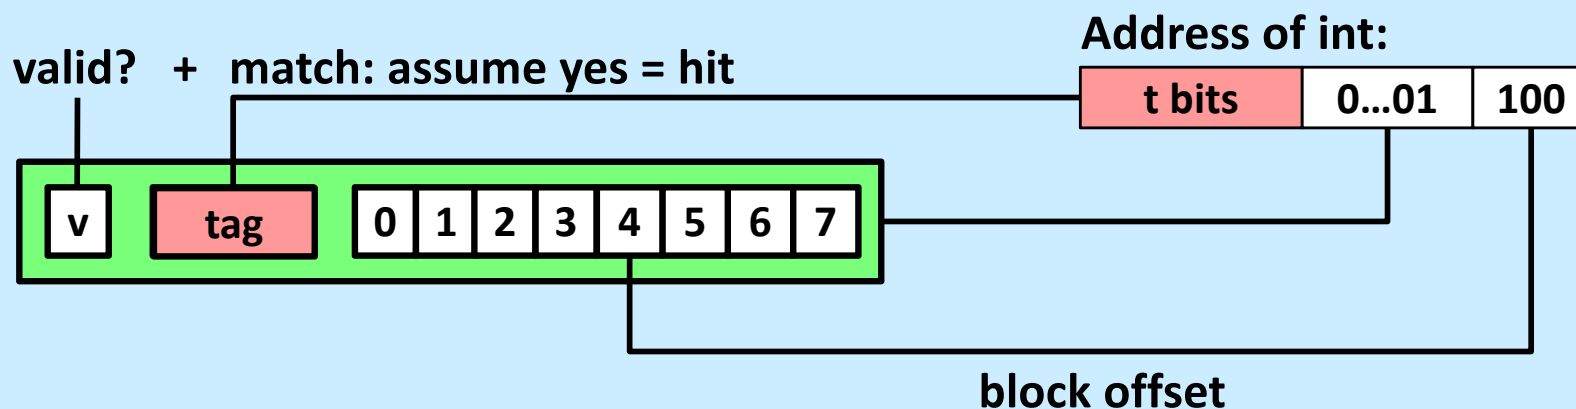
Address of word:

t bits	s bits	b bits

tag   set index   block offset

data begins at this offset

valid bit

v   tag   0   1   2   ......   B-1

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

Example: Direct Mapped Cache (E = 1)

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

$S = 2^s$ sets

Address of int:

| t bits | 0...01 | 100 |

find set

Example: Direct Mapped Cache (E = 1)

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

Address of int:

| t bits | 0...01 | 100 |

valid? + match: assume yes = hit

| v | tag | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |

block offset

Example: Direct Mapped Cache (E = 1)

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

Address of int:

valid? + match: assume yes = hit

| t bits | 0...01 | 100 |

| v | tag | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |

block offset

int (4 Bytes) is here
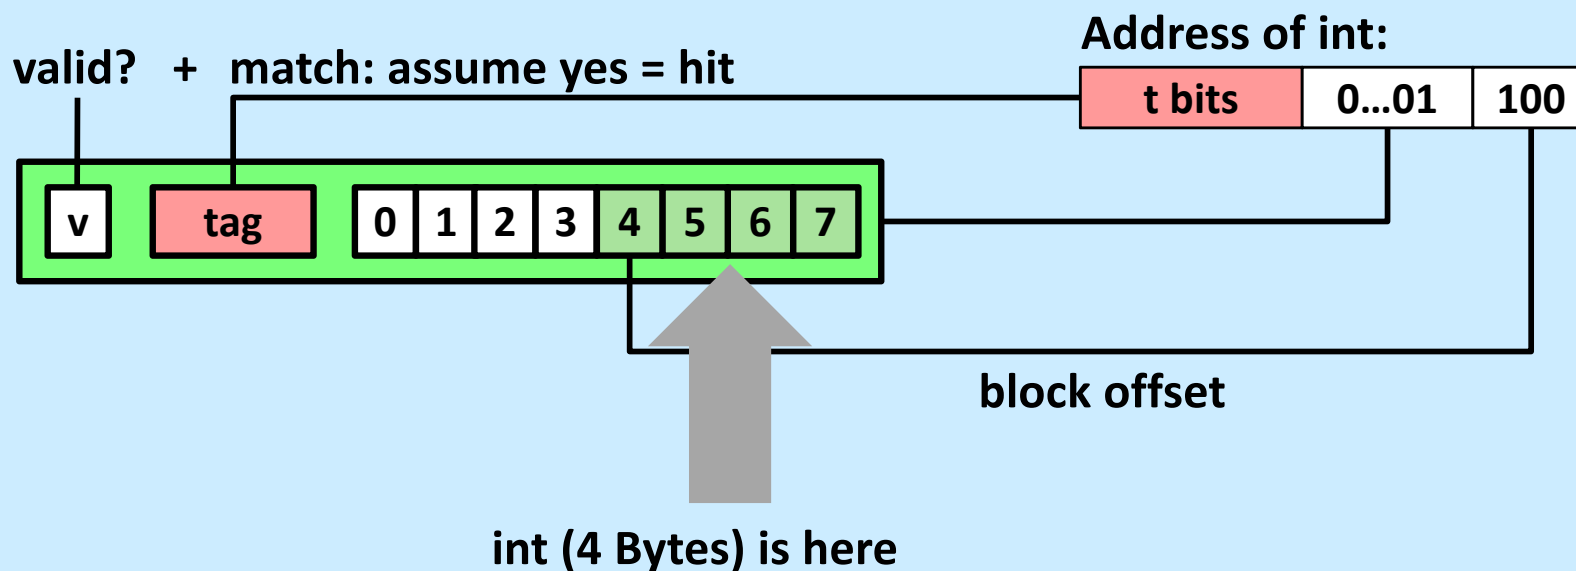
No match: old line is evicted and replaced

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	[0000_2],	miss
1	[0001_2],	hit
7	[0111_2],	miss
8	[1000_2],	miss
0	[0000_2]	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

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;

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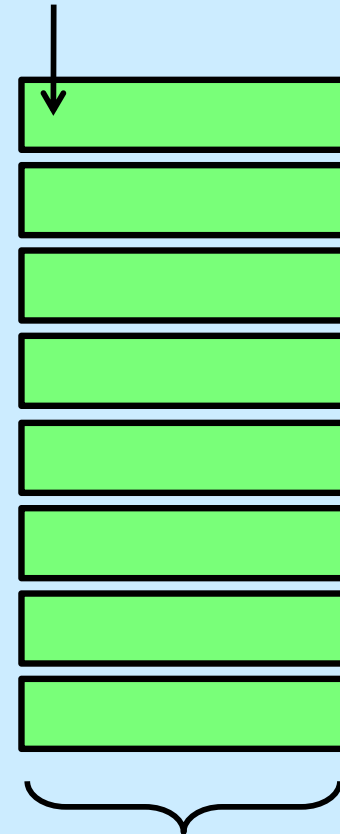
```
int sum_array_cols(double a[16][16])
{
    int i, j;
    double sum = 0;

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

}
```

32 B = 4 doubles
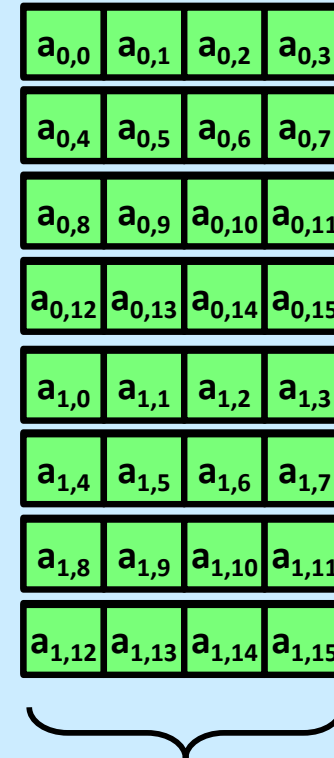
A Higher-Level Example

```
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;
}
```

$a_{0,0}$	$a_{0,1}$	$a_{0,2}$	$a_{0,3}$
$a_{0,4}$	$a_{0,5}$	$a_{0,6}$	$a_{0,7}$
$a_{0,8}$	$a_{0,9}$	$a_{0,10}$	$a_{0,11}$
$a_{0,12}$	$a_{0,13}$	$a_{0,14}$	$a_{0,15}$
$a_{1,0}$	$a_{1,1}$	$a_{1,2}$	$a_{1,3}$
$a_{1,4}$	$a_{1,5}$	$a_{1,6}$	$a_{1,7}$
$a_{1,8}$	$a_{1,9}$	$a_{1,10}$	$a_{1,11}$
$a_{1,12}$	$a_{1,13}$	$a_{1,14}$	$a_{1,15}$

32 B = 4 doubles
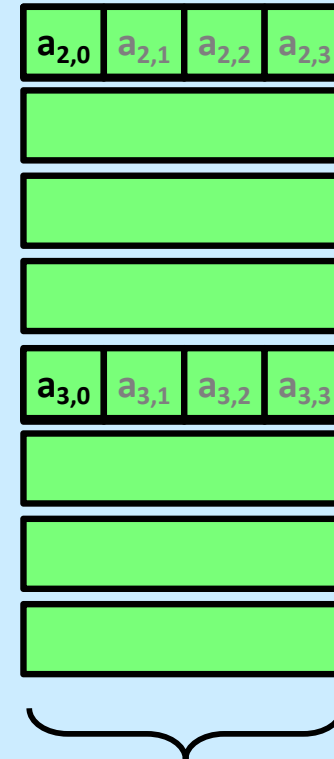
A Higher-Level Example

```
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; i++)
        for (i = 0; i < 16; j++)
            sum += a[i][j];
    return sum;
}
```
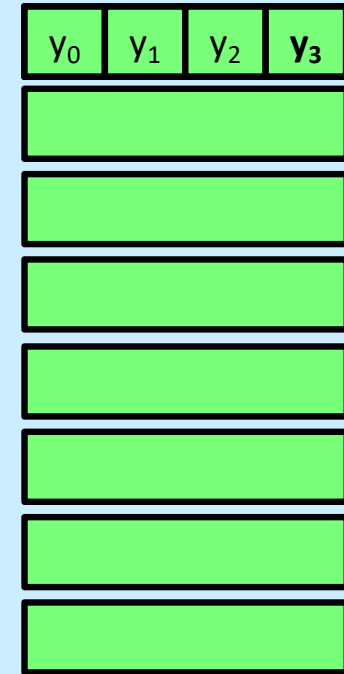
$a_{2,0}$	$a_{2,1}$	$a_{2,2}$	$a_{2,3}$

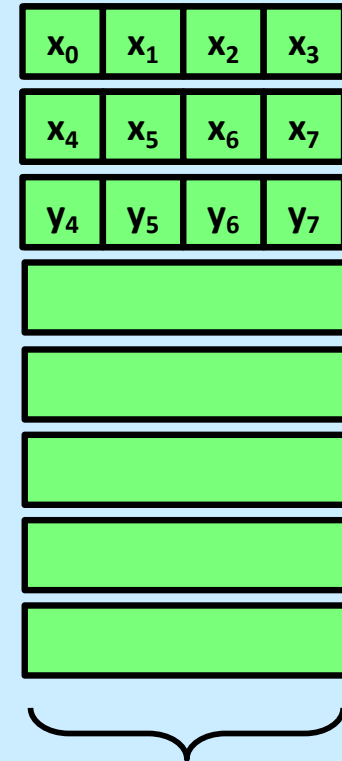| $a_{3,0}$	$a_{3,1}$	$a_{3,2}$	$a_{3,3}$

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;
}
```

y_0 y_1 y_2 **y_3**

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;
}
```

x_0	x_1	x_2	x_3
x_4	x_5	x_6	x_7
y_4	y_5	y_6	y_7

32 B = 4 doubles