

15-213 Recitation

Caches and C Review

Your TAs

Monday, Feb 20th, 2023

Attack Lab Conclusion

- If you enjoyed this lab:
 - Consider [15-330](#) Introduction to Computer Security
 - Consider joining the hacking team at CMU - [PPP](#)
 - Don't use functions vulnerable to buffer overflow (like gets)
 - Use functions that allow you to specify buffer lengths:
 - fgets instead of gets
 - strncpy instead of strcpy
 - strncat instead of strcat
 - snprintf instead of sprintf
 - Use sscanf and fscanf with input lengths (%213s)
 - Stack protection makes buffer overflow very hard...
 - But very hard ≠ impossible!

Agenda

- Logistics
- Cache Lab
- Cache Concepts
- Activity 1: Traces
- C Review
- Activity 2: Getopt()
- Appendix: Examples, Style, Git, fscanf

Logistics

- Cache Lab is due **Thursday, March 2nd** at 11:59pm
- NO Midterm!
- Drop date **Monday, Feb. 27th**
- C Bootcamp happened on Sunday, Feb 19th
- Make sure you have Github working so you can commit your code!


Cache Lab: Cache Simulator Hints

- Goal: Count hits, misses, evictions and # of dirty bytes
- Procedure
 - Least Recently Used (LRU) replacement policy
 - Structs are good for storing cache line parts (valid bit, tag, LRU counter, etc.)
 - A cache is like a 2D array of cache lines

```
struct cache_line cache[S][E];
```
- Your simulator needs to handle different values of S, E, and b (block size) given at run time
 - Dynamically allocate memory!
- Dirty bytes: any payload byte whose corresponding cache block's dirty bit is set (i.e. the payload of that block has been modified, but not yet written back to main memory)

Cache Concepts

Cache Organization



Cache Read

- Address of word: | t bits | s bits | b bits |
 - Tag: t bits
 - Set index: s bits
 - Block offset: b bits
- Steps:
 - Use set index to get appropriate set
 - Loop through lines in set to find matching tag
 - If found and valid bit is set: hit
 - Locate data starting at block offset

Tying it all together: Bomblab

```
(gdb) disas phase_1
Dump of assembler code for function phase_1:
   0x0000000000400e80 <+0>:      sub    $0x8,%rsp
   0x0000000000400e84 <+4>:      mov    $0x604420,%esi
   0x0000000000400e89 <+9>:      callq 0x401326 <strings_not_equal>
   0x0000000000400e8e <+14>:     test   %al,%al
   0x0000000000400e90 <+16>:     je     0x400e97 <phase_1+23>
   0x0000000000400e92 <+18>:     callq 0x401577 <explode_bomb>
   0x0000000000400e97 <+23>:     add    $0x8,%rsp
   0x0000000000400e9b <+27>:     retq
End of assembler dump.
```

Tying it all together: Bomblab

```
(gdb) disas phase_1
Dump of assembler code for function phase_1:
   0x0000000000400e80 <+0>:      sub    $0x8,%rsp
   0x0000000000400e84 <+4>:      mov    $0x604420,%esi
   0x0000000000400e89 <+9>:      callq 0x401326 <strings_not_equal>
   0x0000000000400e8e <+14>:     test  %al,%al
   0x0000000000400e90 <+16>:     je     0x400e97 <phase_1+23>
   0x0000000000400e92 <+18>:     callq 0x401577 <explode_bomb>
   0x0000000000400e97 <+23>:     add   $0x8,%rsp
   0x0000000000400e9b <+27>:     retq
End of assembler dump.
```

Tying it all together: Bomblab

```
tianxinx@bambooshark:~$ getconf -a | grep CACHE
LEVEL1_ICACHE_SIZE          32768
LEVEL1_ICACHE_ASSOC         4
LEVEL1_ICACHE_LINESIZE     32
LEVEL1_DCACHE_SIZE          32768
LEVEL1_DCACHE_ASSOC         8
LEVEL1_DCACHE_LINESIZE     64
LEVEL2_CACHE_SIZE           262144
LEVEL2_CACHE_ASSOC          8
LEVEL2_CACHE_LINESIZE      64
LEVEL3_CACHE_SIZE           8388608
LEVEL3_CACHE_ASSOC          16
LEVEL3_CACHE_LINESIZE      64
LEVEL4_CACHE_SIZE           0
LEVEL4_CACHE_ASSOC          0
LEVEL4_CACHE_LINESIZE      0
tianxinx@bambooshark:~$
```

For the L1 dCache (data)

$C = 32768$ (32 kb)

$E = 8$

$B = 64$

$S = 64$

How did we get S?

Tying it all together: Bomblab

- 64 bit address space: $m = 64$
- $b = 6$
- $s = 6$
- $t = 52$

Tying it all together: Bomblab

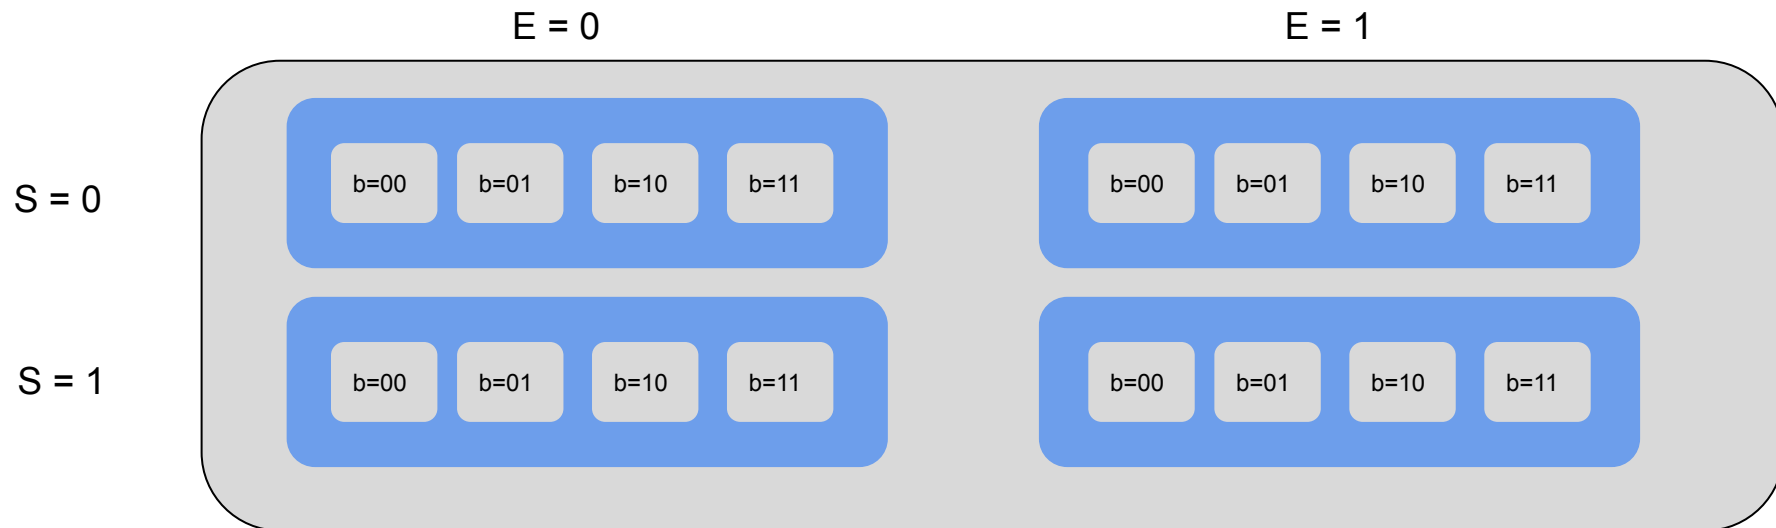
0x00604420 → 0b0000000011000000100010000100000

- tag bits: 0000000011000000100
- set index bits: 010000
- block offset bits: 100000

Activity 1: Traces

Tracing a Cache

Example Cache: `-s 1 -E 2 -b 2` ($S=2$ $B=4$)



Example Trace

L - Load

S - Store

Memory Location

Size

Jack.trace

L 0,4

S 0,4

L 0,1

L 6,1

L 5,1

L 6,1

L 7,1

Example Trace

Jack.trace

L 0,4

S 0,4


L 0,1

L 6,1

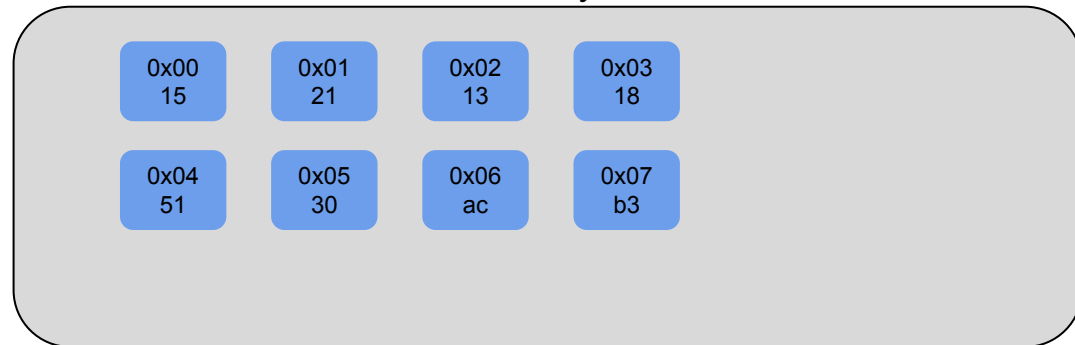
L 5,1

L 6,1

L 7,1



Memory



Example Trace

Jack.trace

L 0,4 M

S 0,4

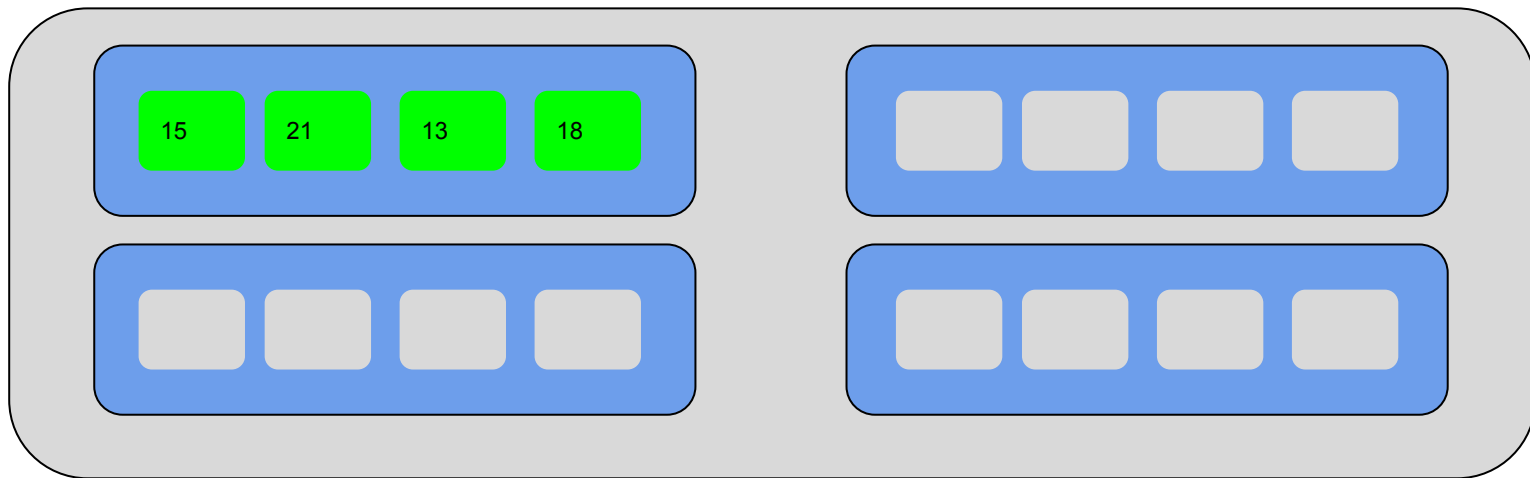
L 0,1

L 6,1

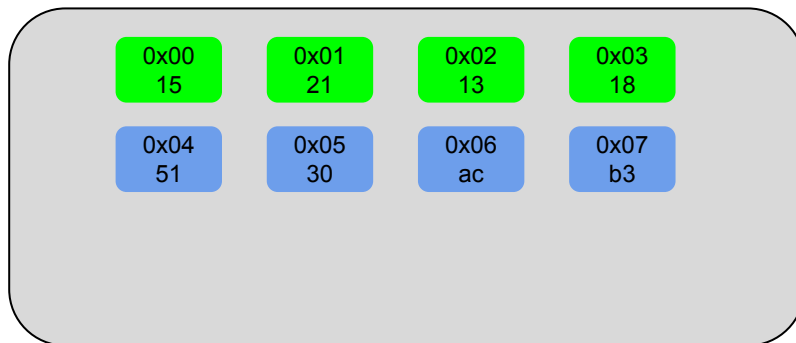
L 5,1

L 6,1

L 7,1



Memory



Why that line?

Where are those values from?

Example Trace

Jack.trace

L 0,4 M

S 0,4 H


L 0,1

L 6,1


L 5,1

L 6,1

L 7,1



Memory



What happens if values change?

Example Trace

Jack.trace

L 0,4 M

S 0,4 H

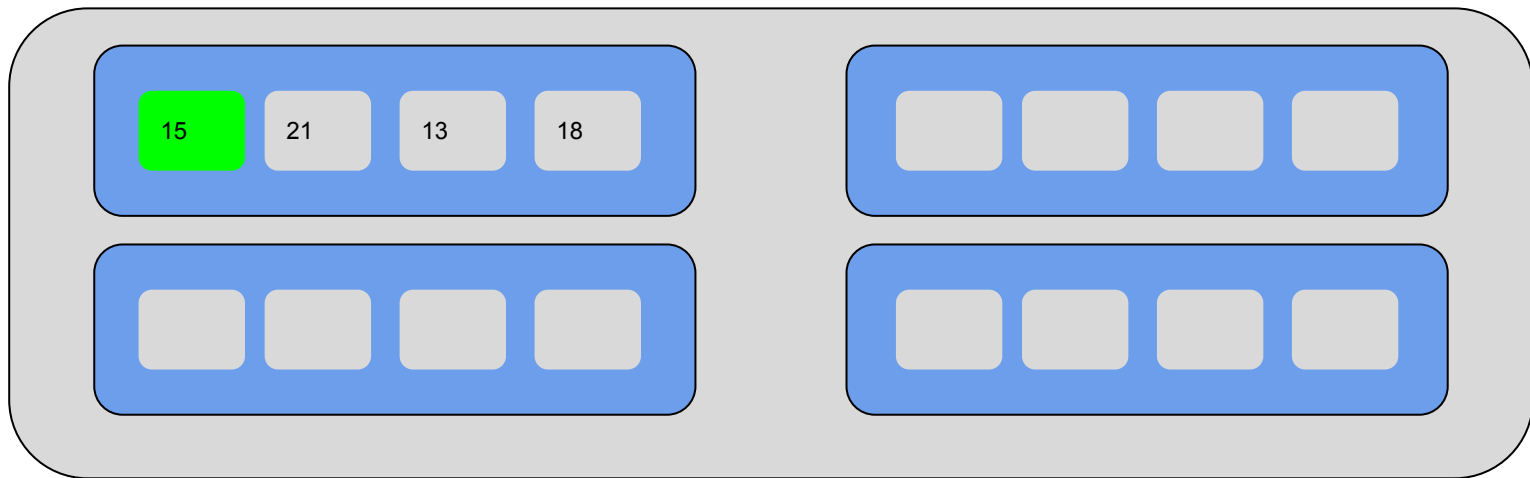
L 0,1 H

L 6,1

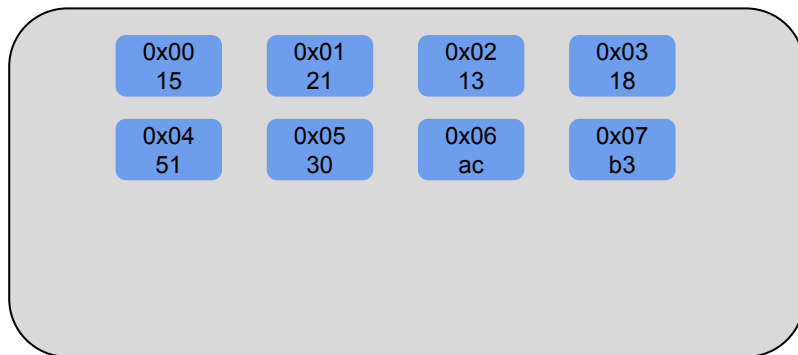
L 5,1

L 6,1

L 7,1



Memory



Why is this still a hit?

What would happen if we had not previously loaded all four bytes?

Example Trace

Jack.trace

L 0,4 M

S 0,4 H


L 0,1 H

L 6,1 M

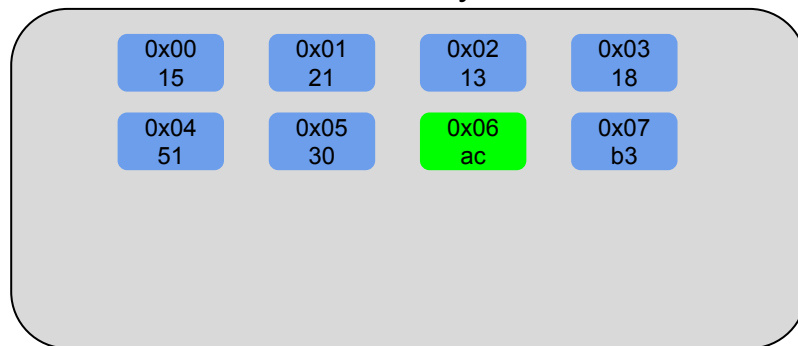
L 5,1

L 6,1

L 7,1

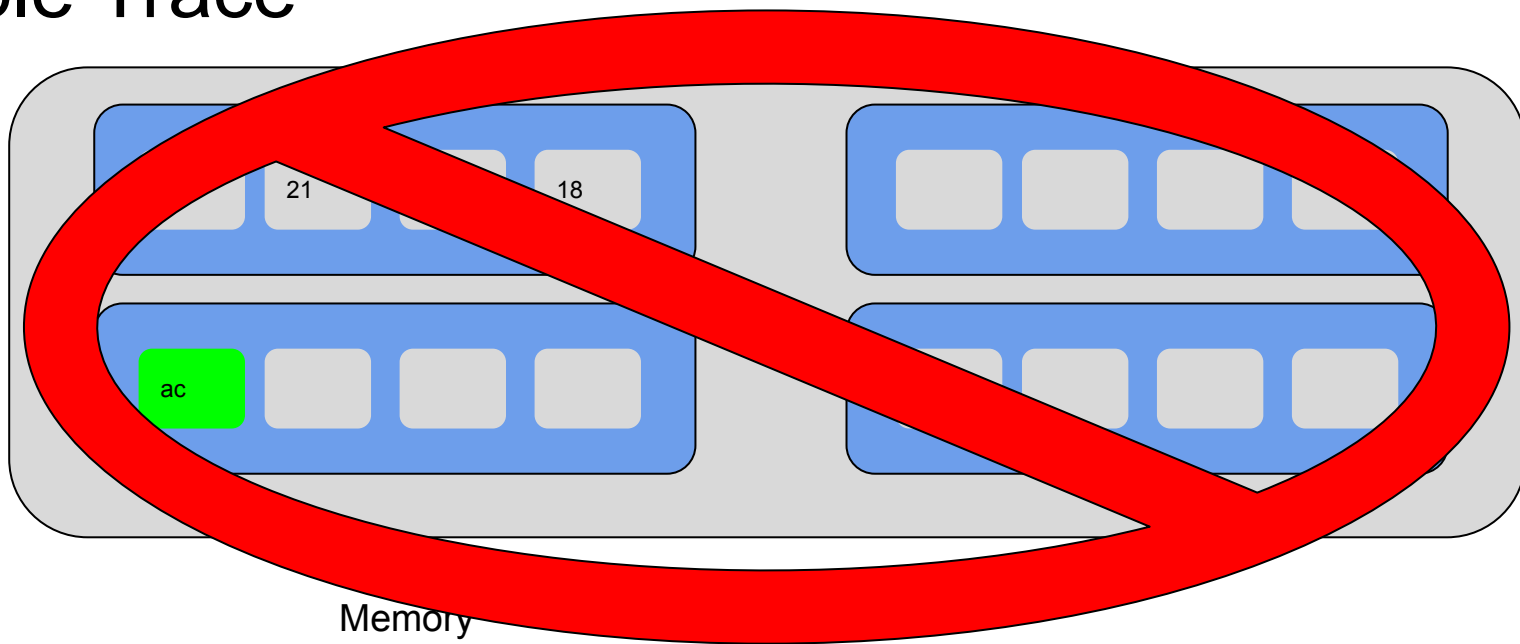


Memory



Just one Byte?

Example Trace



Jack.trace

L 0,4 M

S 0,4 H


L 0,1 H

L 6,1 M

L 5,1

L 6,1

L 7,1



Just one Byte?

NO!

Example Trace

Jack.trace

L 0,4 M

S 0,4 H

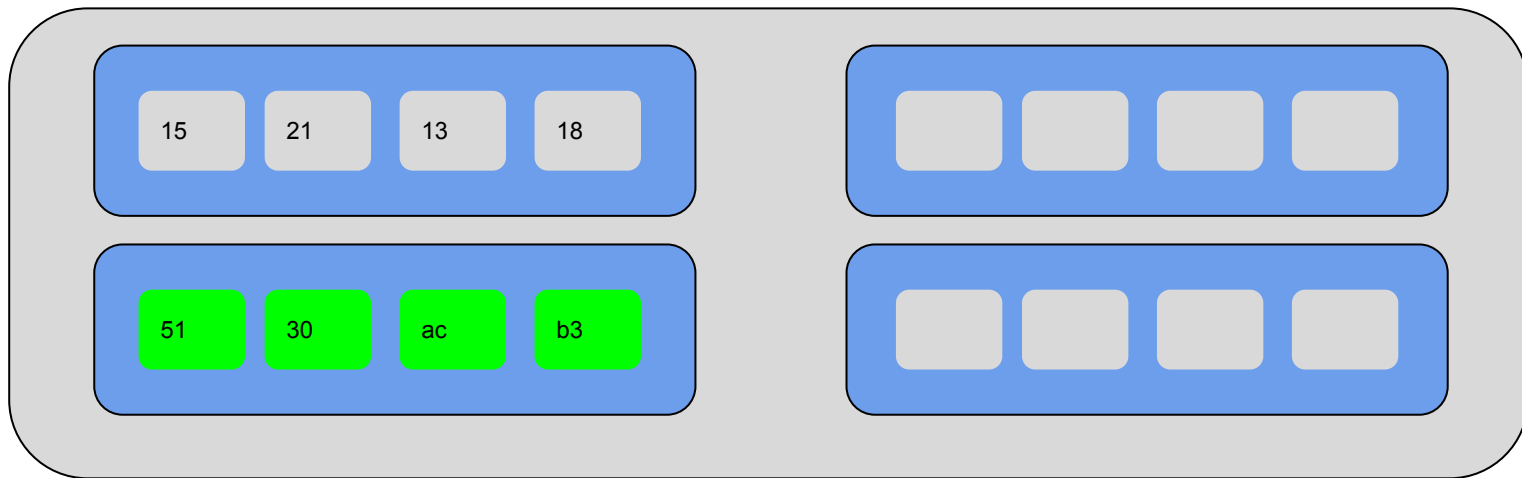
L 0,1 H

L 6,1 M

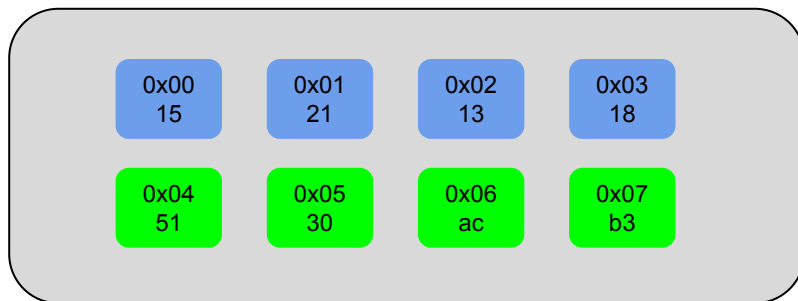
L 5,1

L 6,1

L 7,1



Memory



Why below and not above?

Why load all four bytes?

Example Trace

Jack.trace

L 0,4 M

S 0,4 H


L 0,1 H

L 6,1 M


L 5,1 H

L 6,1

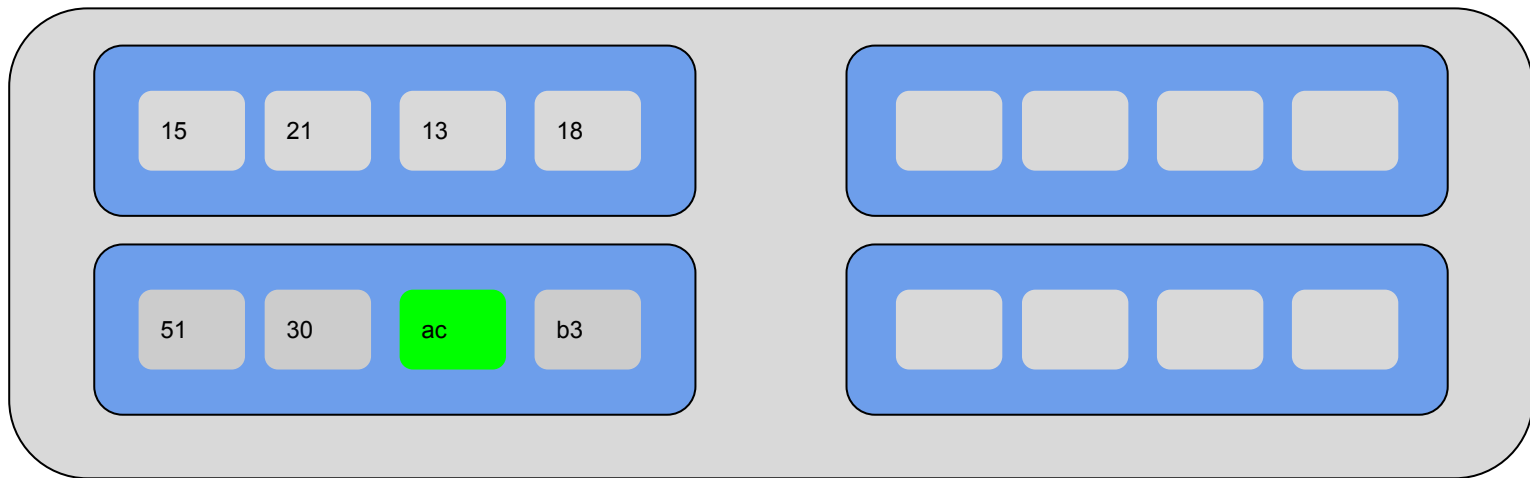
L 7,1



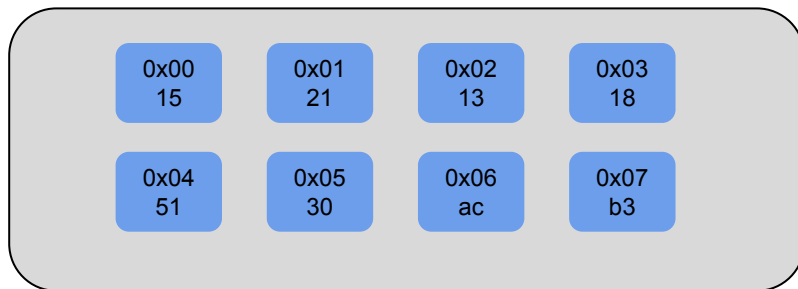
Memory



Example Trace



Memory



Jack.trace

L 0,4 M

S 0,4 H

L 0,1 H

L 6,1 M

L 5,1 H

L 6,1 H

L 7,1

Example Trace

Jack.trace

L 0,4 M

S 0,4 H

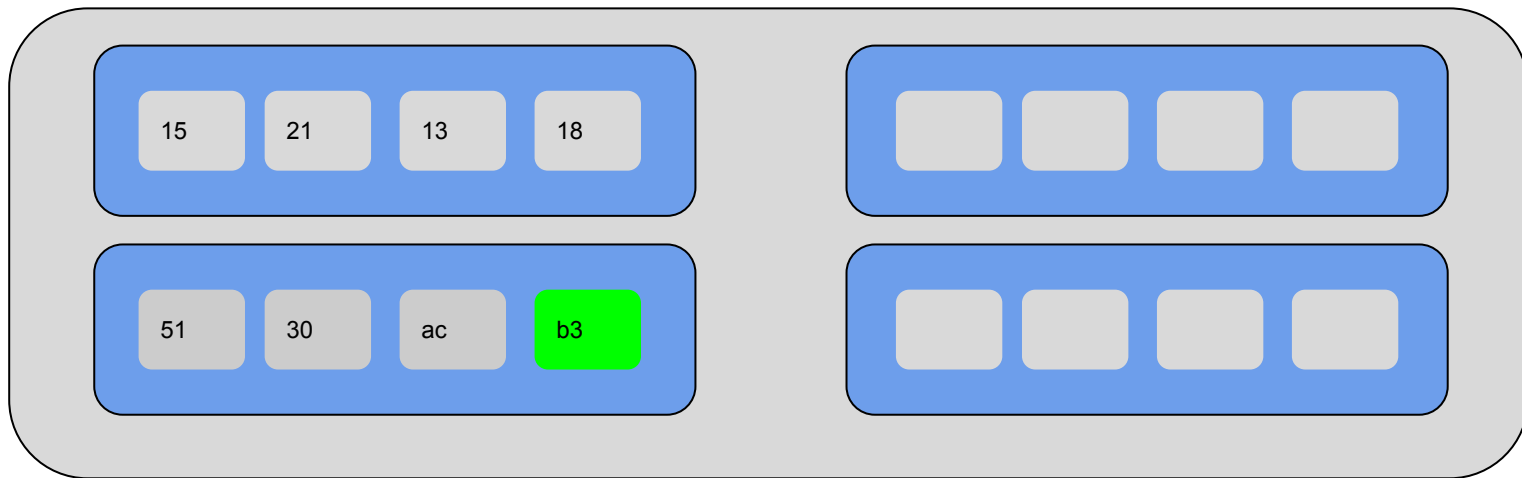
L 0,1 H

L 6,1 M


L 5,1 H

L 6,1 H

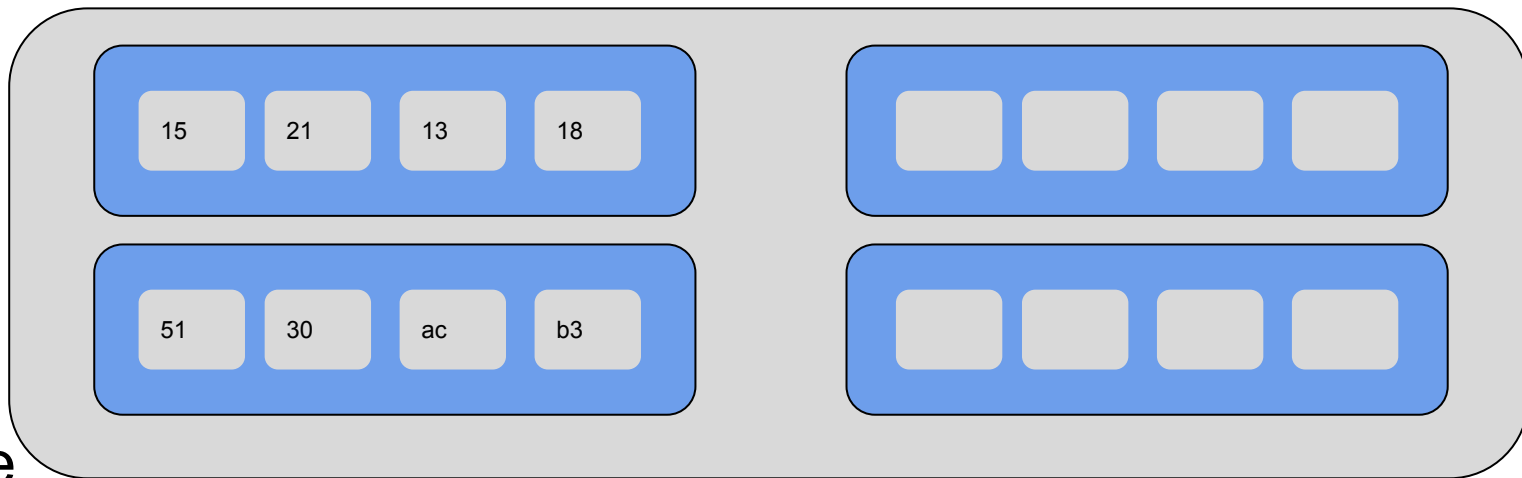
L 7,1 H



Memory

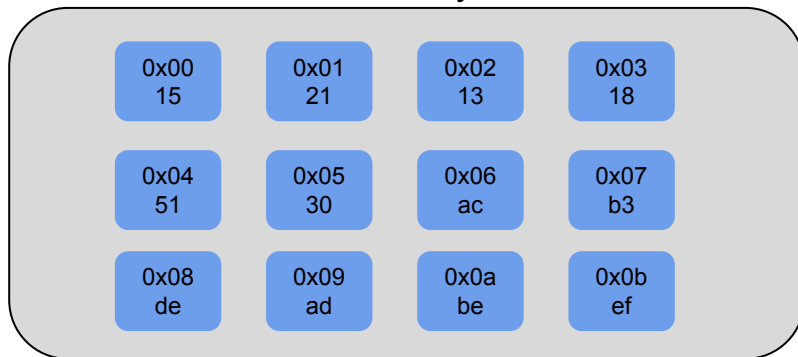


Example Trace



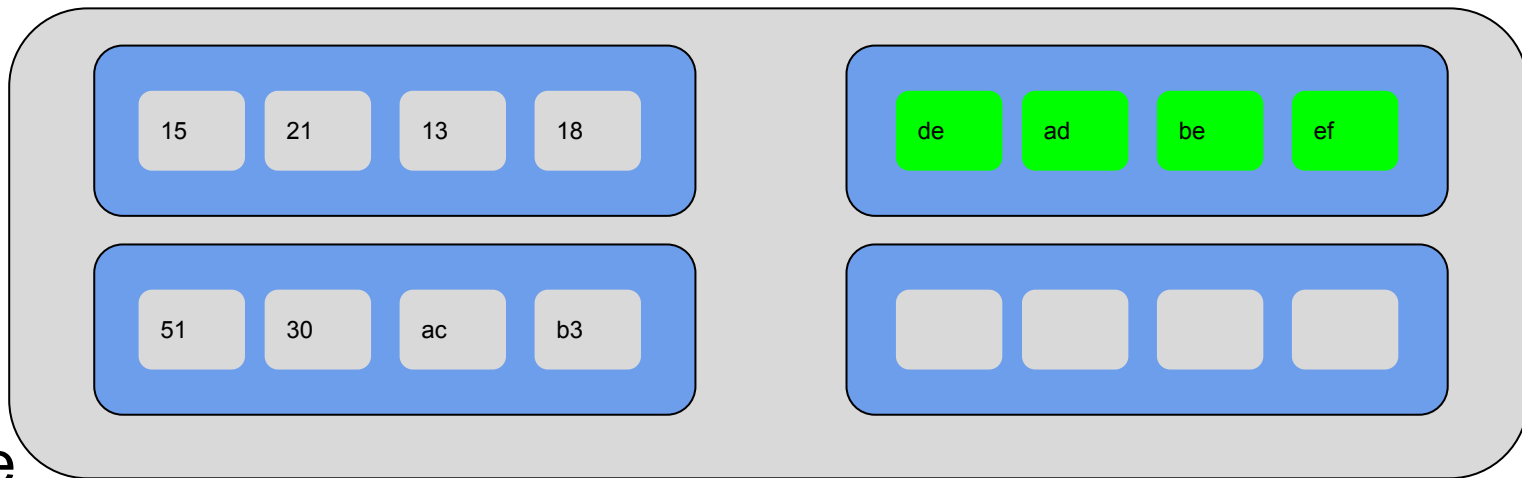
Jack2.trace
L 8,4 M

Memory



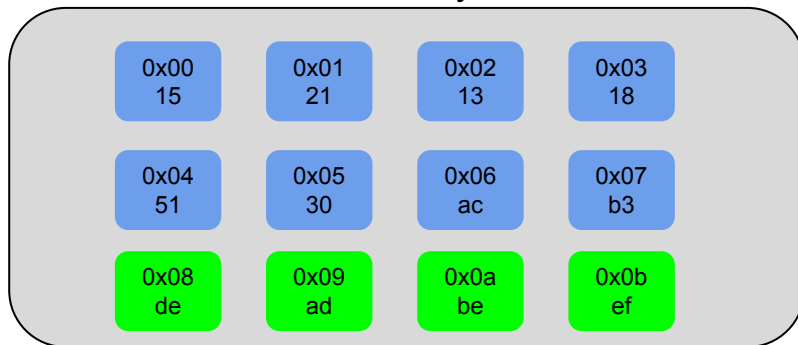
What would happen if we
loaded from memory address
0x08?

Example Trace



Jack2.trace
L 8,4 M

Memory



What would happen if we loaded from memory address 0x08?

C Review

C bootcamp is your go-to!

C Review: Pointers

- Pointer: stores address of some value in memory
- Dereferencing a NULL pointer causes segfault
- Dereferencing a pointer: `*p`
- Access address of a value: `p = &v`

C Review: Pointers

- What is wrong with this code?

```
1 int main(int argc, char** argv) {
2     int *a = (int*) malloc(213 * sizeof(int));
3     for (int i=0; i<213; i++) {
4         if (a[i] == 0) a[i]=i;
5         else a[i]=-i;
6     }
7     return 0;
8 }
```

C Review: Pointers

- malloc can fail!

```
1 int main(int argc, char** argv) {
2     int *a = (int*) malloc(213 * sizeof(int));
3     if (a == NULL) return 0;
4     for (int i=0; i<213; i++) {
5         if (a[i] == 0) a[i]=i;
6         else a[i]=-i;
7     }
8     return 0;
9 }
```


C Review: Pointers

- Allocated memory is not initialized!

```
1 int main(int argc, char** argv) {
2     int *a = (int*) calloc(213, sizeof(int));
3     if (a == NULL) return 0;
4     for (int i=0; i<213; i++) {
5         if (a[i] == 0) a[i]=i;
6         else a[i]=-i;
7     }
8     return 0;
9 }
```

C Review: Pointers

- All allocated memory must be freed!

```
1 int main(int argc, char** argv) {
2     int *a = (int*) calloc(213, sizeof(int));
3     if (a == NULL) return 0;
4     for (int i=0; i<213; i++) {
5         if (a[i] == 0) a[i]=i;
6         else a[i]=-i;
7     }
8     free(a);
9     return 0;
10 }
```

C Review: Arrays

- Initializing your array
 - `int *a = calloc(4, sizeof(int));`
 - Allocated on Heap
 - `int a[4];`
 - Allocated on stack
- Where does the following point to?

```
int a[4] = {1,2,3,4};
```

- `a[0]`
- `*(a + 3)`

```
char *listOfName[4] = {"Alice", "Bob", "Cherry"};
```


- `(listOfName + 1)`
- `*(listOfName + 1)`

C Review: Structs + Unions

Struct:

- Groups list of variables under one block in memory


```
struct temp {  
    int i;  
    char c;  
};
```



Union:

- Store different data types in same region of memory
- Many ways to refer to same memory location

```
union temp {  
    int i;  
    char c;  
};
```



C Review: Valgrind

- What is Valgrind?
 - Tool used for debugging memory use
- Valgrind may...
 - Find corrupted memory
 - Find potential memory leaks and double frees
 - Detects invalid memory reads and writes
- To learn more... `man valgrind` and check the appendix

C Review Conclusion

- Did you know each concept? If not...
 - Refer to the C Bootcamp slides
- Were the concepts so easy you were bored? If not...
 - Refer to the C Bootcamp slides
- When in doubt...
 - Refer to the C Bootcamp slides
- This will be *very* important for the rest of this class, so make sure you are comfortable with the material covered or check the C Bootcamp recording!

C Programming Style

- Write comments and then implement functionality
- Communicate meaning through naming choices
- Code should be testable. Modularity supports this
- Use consistent formatting
- Common bugs: memory and file descriptor leaks, check errors and failure conditions
- Warning: *Dr. Evil* has returned for code reviews on Cache Lab! 😊
 - Refer to 213 Style suggestions:
<http://cs.cmu.edu/~213/codeStyle.html>

Part 0: reading man pages!

- Reading man pages is important!
- To get started, either:
 - `$ man getopt` on Terminal
 - Google “man getopt”
- Overall, what does getopt do?
- What arguments does it take?
- How can you use it in a program?
- <https://linux.die.net/man/3/getopt>

Activity: getopt()

Part 1: Activity Setup

- Split up into groups of 2-3 people
- One person needs a laptop
- Log in to a Shark machine, and type:

```
$ wget https://www.cs.cmu.edu/~213/activities/rec6.tar  
$ tar -xvf rec6.tar  
$ cd rec6
```

Part 1: getopt_example.c

```
$ make getopt_example
```

```
$ ./getopt_example (ARGUMENTS)
```


- What does getopt_example.c do?
- How does the program process its arguments?
- i.e. formatting specifics?
- What does the -v argument do? The -n argument?
- Hint: try `$./getopt_example -v -n 5`

Part 1: getopt_example.c

- What does getopt_example.c do?
 - Takes in a number as input + “counts” to that number.
 - Verbose (-v) : prints all numbers counting up to that number)
- Formatting specifics
 - Use -(ARG) to get getopt to process the argument
 - -v: Enables verbose mode
 - -n:NUM with NUM as user input

```
while ((opt = getopt(argc, argv, "vn:")) != -1) {  
    switch (opt) {  
        case 'v':  
            verbose = 1;  
            break;  
        case 'n':  
            n = atoi(optarg);  
            break;  
        default:  
            fprintf(stderr, "usage: ...");  
            exit(1);  
    }  
}
```

Returns -1 when
done parsing



Parses value to
store in n b/c colon

If you get stuck...

- Reread the writeup
- Look at CS:APP Chapter 6
- Review lecture notes (<http://cs.cmu.edu/~213>)
- Come to Office Hours
- Post private question on Piazza
- `man malloc, man valgrind, man gdb`

Cache Lab Tips!

- Review cache and memory lectures
 - Ask if you don't understand something
- Start early, this can be a challenging lab!
- Don't get discouraged!
 - If you try something that doesn't work, take a well deserved break, and then try again
- Good luck!

Practice Problems

Class Question / Discussions

- We'll work through a series of questions
- Write down your answer for each question
- You can discuss with your classmates

What Type of Locality?

- The following function exhibits which type of locality? Consider *only* array accesses.

```
void who(int *arr, int size) {  
    for (int i = 0; i < size-1; ++i)  
        arr[i] = arr[i+1];  
}
```

- A.** Spatial
- B.** Temporal
- C.** Both A and B
- D.** Neither A nor B

What Type of Locality?

- The following function exhibits which type of locality? Consider *only* array accesses.

```
void who(int *arr, int size) {  
    for (int i = 0; i < size-1; ++i)  
        arr[i] = arr[i+1];  
}
```

- A. Spatial
- B. Temporal
- C.** Both A and B
- D. Neither A nor B

What Type of Locality?

- The following function exhibits which type of locality? Consider *only* array accesses.

```
void coo(int *arr, int size) {  
    for (int i = size-2; i >= 0; --i)  
        arr[i] = arr[i+1];  
}
```

- A.** Spatial
- B.** Temporal
- C.** Both A and B
- D.** Neither A nor B

What Type of Locality?

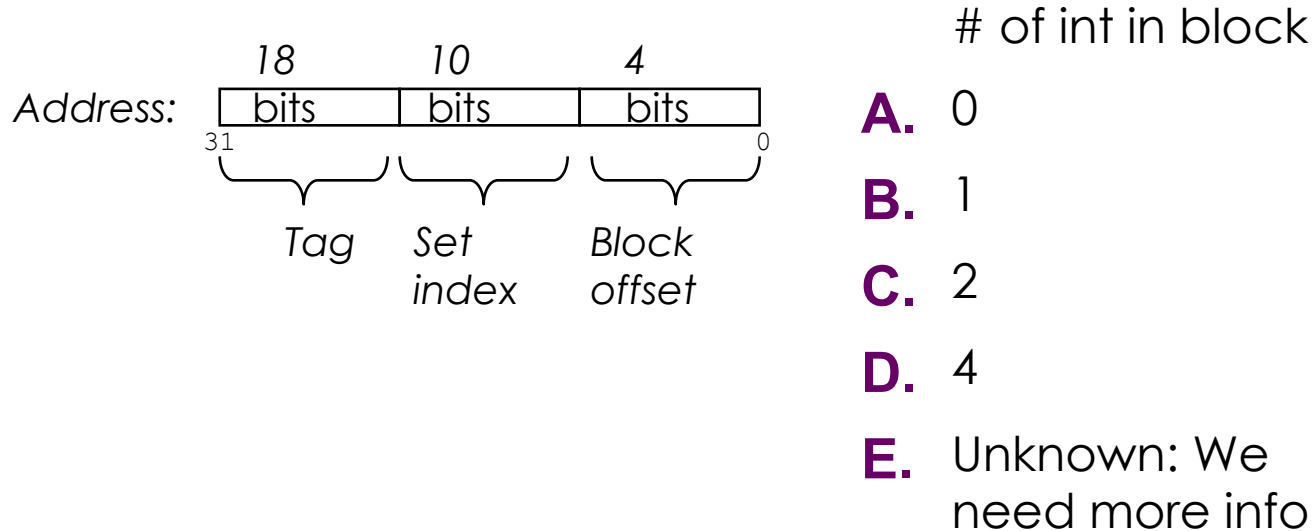
- The following function exhibits which type of locality? Consider *only* array accesses.

```
void coo(int *arr, int size) {  
    for (int i = size-2; i >= 0; --i)  
        arr[i] = arr[i+1];  
}
```

- A.** Spatial
- B.** Temporal
- C.** Both A and B
- D.** Neither A nor B

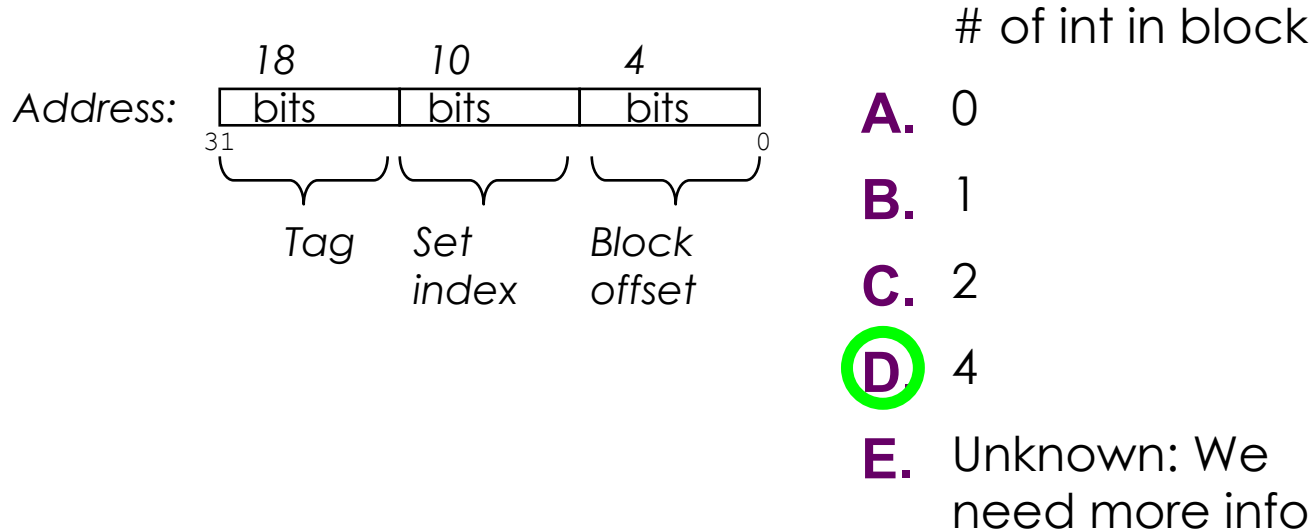
Calculating Cache Parameters

- Given the following address partition, how many `int` values will fit in a single data block?



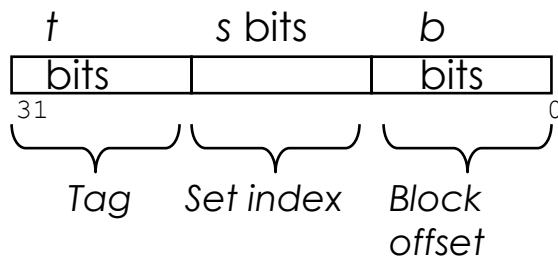
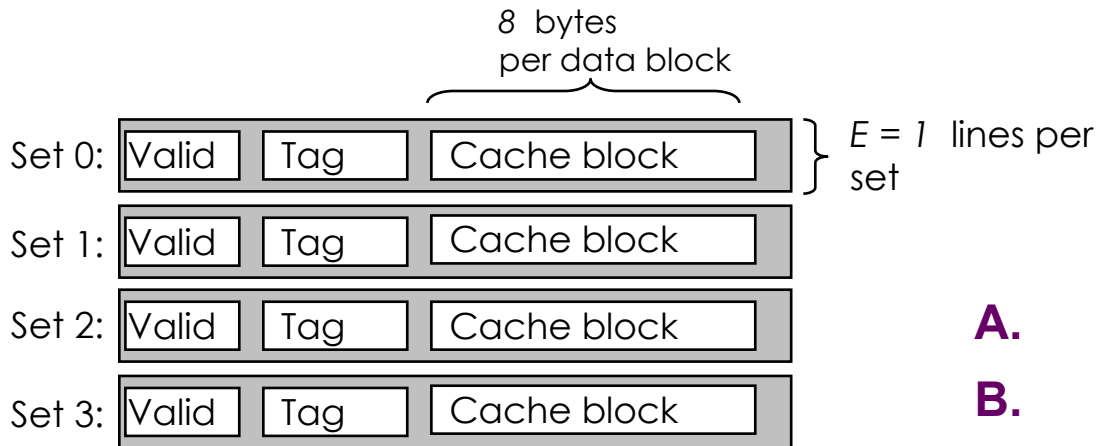
Calculating Cache Parameters

- Given the following address partition, how many `int` values will fit in a single data block?



Direct-Mapped Cache Example



- Assuming a 32-bit address (i.e. $m=32$), how many bits are used for tag (t), set index (s), and block offset (b).



	t	s	b
A.	1	2	3
B.	27	2	3
C.	25	4	3
D.	1	4	8
E.	20	4	8

Direct-Mapped Cache Example


- Assuming a 32-bit address (i.e. $m=32$), how many bits are used for tag (t), set index (s), and block offset (b).



	t	s	b
A.	1	2	3
B.	27	2	3
C.	25	4	3
D.	1	4	8
E.	20	4	8

Which Set Is it?

- Which set is the address **0xFA1C** located in?



Set # for 0xFA1C


A. 0

B. 1

C. 2



D. 3

E. More than one of the above



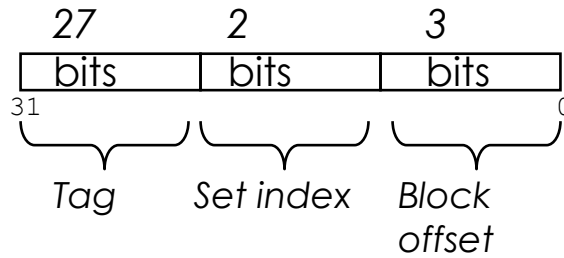
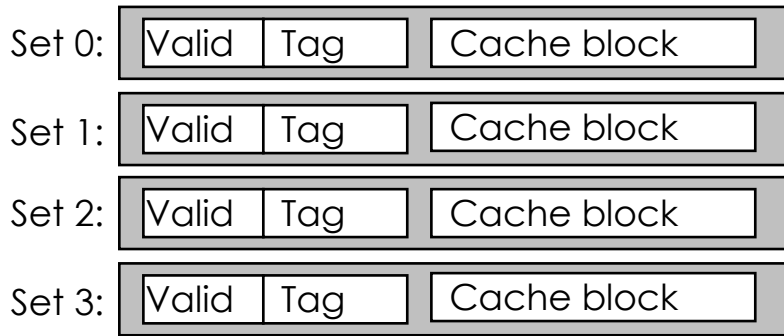
Which Set Is it?

- Which set is the address **0xFA1C** located in?



Cache Block Range

- What range of addresses will be in the same block as address **0xFA1C**? 8 bytes per data block

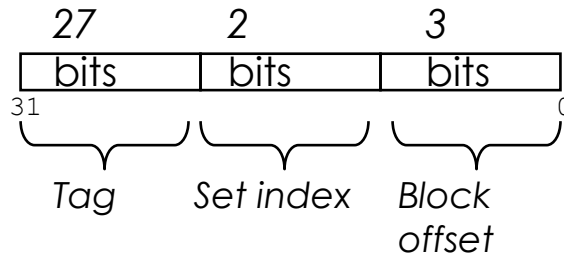
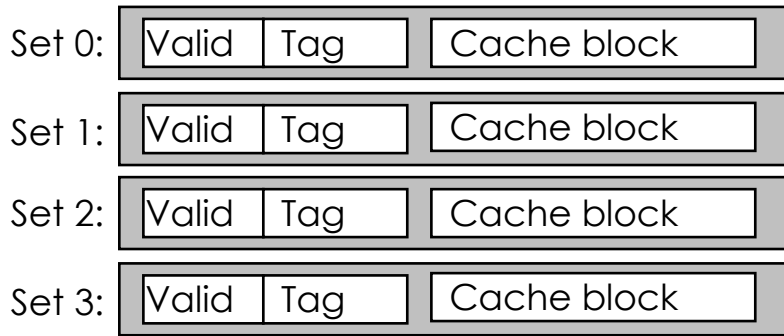


Addr. Range

- A.** 0xFA1C
- B.** 0xFA1C – 0xFA23
- C.** 0xFA1C – 0xFA1F
- D.** 0xFA18 – 0xFA1F
- E.** It depends on the access size (byte, word, etc)

Cache Block Range

- What range of addresses will be in the same block as address **0xFA1C**? 8 bytes per data block



Addr. Range

- A.** 0xFA1C
- B.** 0xFA1C – 0xFA23
- C.** 0xFA1C – 0xFA1F
- D.** 0xFA18 – 0xFA1F
- E.** It depends on the access size (byte, word, etc)

Cache Misses

If $N = 16$, how many bytes does the loop access of a ?

<pre>int foo(int* a, int N) { int i; int sum = 0; for(i = 0; i < N; i++) { sum += a[i]; } return sum; }</pre>	Accessed Bytes
A	4
B	16
C	64
D	256

Cache Misses

If $N = 16$, how many bytes does the loop access of a ?

```
int foo(int* a, int N)
{
    int i;
    int sum = 0;
    for(i = 0; i < N; i++)
    {
        sum += a[i];
    }
    return sum;
}
```

Accessed
Bytes

- | | |
|----------|-----|
| A | 4 |
| B | 16 |
| C | 64 |
| D | 256 |

Cache Misses

Consider a 32 KB cache in a 32 bit address space. The cache is 8-way associative and has 64 bytes per block. A LRU (Least Recently Used) replacement policy is used. What is the miss rate on 'pass 1'?

```
void muchAccessSoCacheWow(int *bigArr){
    // 48 KB array of ints
    int length = (48*1024)/sizeof(int);

    int access = 0;

    // traverse array with stride 8

    // pass 1
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }

    // pass 2
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }
}
```

Miss Rate

- A** 0 %
- B** 25 %
- C** 33 %
- D** 50 %
- E** 66 %

Cache Misses

Consider a 32 KB cache in a 32 bit address space. The cache is 8-way associative and has 64 bytes per block. A LRU (Least Recently Used) replacement policy is used. What is the miss rate on 'pass 1'?

```
void muchAccessSoCacheWow(int *bigArr){
    // 48 KB array of ints
    int length = (48*1024)/sizeof(int);

    int access = 0;

    // traverse array with stride 8

    // pass 1
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }

    // pass 2
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }
}
```

Miss Rate

- | | |
|----------|------|
| A | 0 % |
| B | 25 % |
| C | 33 % |
| D | 50 % |
| E | 66 % |

Cache Misses

Consider a 32 KB cache in a 32 bit address space. The cache is 8-way associative and has 64 bytes per block. A LRU (Least Recently Used) replacement policy is used. What is the miss rate on 'pass 2'?

```
void muchAccessSoCacheWow(int *bigArr){
    // 48 KB array of ints
    int length = (48*1024)/sizeof(int);

    int access = 0;

    // traverse array with stride 8

    // pass 1
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }

    // pass 2
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }
}
```

Miss Rate

- | | |
|----------|------|
| A | 0 % |
| B | 25 % |
| C | 33 % |
| D | 50 % |
| E | 66 % |

Cache Misses

Consider a 32 KB cache in a 32 bit address space. The cache is 8-way associative and has 64 bytes per block. A LRU (Least Recently Used) replacement policy is used. What is the miss rate on 'pass 2'?

```
void muchAccessSoCacheWow(int *bigArr){
    // 48 KB array of ints
    int length = (48*1024)/sizeof(int);

    int access = 0;

    // traverse array with stride 8

    // pass 1
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }

    // pass 2
    for(int i = 0; i < length; i+=8){
        access = bigArr[i];
    }
}
```

Miss Rate

- | | |
|----------|------|
| A | 0 % |
| B | 25 % |
| C | 33 % |
| D | 50 % |
| E | 66 % |

Detailed explanation in Appendix!


Appendix

Appendix: C Programming Style

- Properly document your code
 - Function + File header comments, overall operation of large blocks, any tricky bits
- Write robust code – check error and failure conditions
- Write modular code
 - Use interfaces for data structures, e.g. `create/insert/remove/free` functions for a linked list
 - No magic numbers – use `#define` or `static const`
- Formatting
 - 80 characters per line (use Autolab's highlight feature to double-check)
 - Consistent braces and whitespace
- No memory or file descriptor leaks

Appendix: Git: What is Git?

- Most widely used version control system out there
- Version control:
 - Help track changes to your source over time
 - Help teams manage changes on shared code



Appendix: Git Usage

- Commit early and often!
 - At minimum at every major milestone
 - Commits don't cost anything!
- Popular stylistic conventions
 - Branches: short, descriptive names
 - Commits: A single, logical change. Split large changes into multiple commits.
 - Messages:
 - Summary: Descriptive, yet succinct
 - Body: More detailed description on **what** you changed, **why** you changed it, and what **side effects** it may have

Git Commands

- Clone: `git clone <clone-repository-url>`
- Add: `git add .` or `git add <file-name>`
- Push / Pull: `git push` / `git pull`
- Commit: `git commit -m "your-commit-message"`
 - Good commit messages are key!
 - Bad: "commit", "change", "fixed"
 - Good: "Fixed buffer overflow potential in AttackLab"

Appendix: Parsing Input with fscanf

- `fscanf(FILE *stream, const char *format, ...)`
 - “scanf” but for files
- Arguments
 1. A stream pointer, e.g. from `fopen()`
 2. Format string for parsing, e.g. “%c %d,%d”
 - 3+. **Pointers** to variables for parsed data
 - Can be pointers to stack variables
- Return Value
 - Success: # of parsed vars
 - Failure: EOF
- `man fscanf`

Appendix: fscanf() Example

```
FILE *pFile;
pFile = fopen("trace.txt", "r"); // Open file for reading

// TODO: Error check sys call

char access_type;
unsigned long address;
int size;

// Line format is " S 2f,1" or " L 7d0,3"
//      - 1 character, 1 hex value, 1 decimal value
while (fscanf(pFile, " %c %lx, %d", &access_type, &address, &size) > 0)
{
    // TODO: Do stuff
}

fclose(pFile); // Clean up Resources
```

Appendix: Discussion Questions

- What did the optimal transversal orders have in common?
- How does the pattern generalize to `int[8][8] A` and a cache that holds 4 lines each of 4 `int`'s?

Appendix: Valgrind

- Finding memory leaks
 - `$ valgrind -leak-resolution=high
-leak-check=full -show-reachable=yes
-track-fds=yes ./myProgram arg1 arg`
- Remember that Valgrind can be used for other things, like finding invalid reads and writes!

Appendix: \$ man 3 getopt

- `int getopt(int argc, char * const argv[], const char *optstring);`
 - `int argc` → argument count passed to `main()`
 - Note: includes executable, so `./a.out 1 2` has `argc=3`
 - `char * const argv` is argument string array passed to `main`
 - `const char *optstring` → string with command line arguments
 - Characters followed by colon require arguments
 - Find argument text in `char *optarg`
 - `getopt` can't find argument or finds illegal argument sets `optarg` to “?”
 - Example: “`abc:d:`”
 - `a` and `b` are boolean arguments (not followed by text)
 - `c` and `d` are followed by text (found in `char *optarg`)
- Returns: `getopt` returns `-1` when done parsing

Appendix: Clang / LLVM

- Clang is a (gcc equivalent) C compiler
 - Support for code analyses and transformation
 - Compiler will check you variable usage and declarations
 - Compiler will create code recording all memory accesses to a file
 - Useful for Cache Lab Part B (Matrix Transpose)