

Structs & Alignment

CSE 351 Winter 2019

Instructors:

Max Willsey

Luis Ceze

Teaching Assistants:

Britt Henderson

Lukas Joswiak

Josie Lee

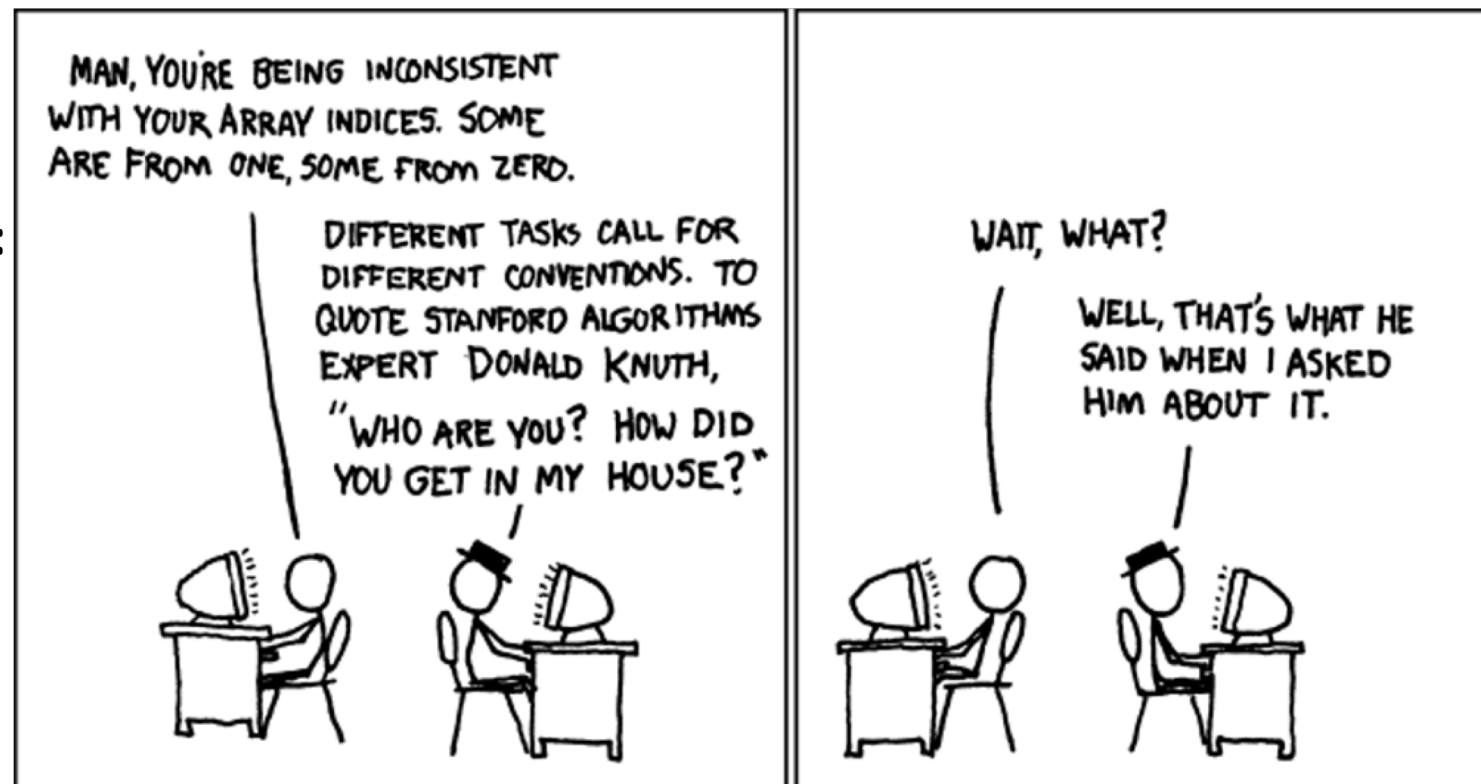
Wei Lin

Daniel Snitkovsky

Luis Vega

Kory Watson

Ivy Yu



Administrivia

- ❖ Snow Day! Online office hours
- ❖ Mid-survey due Thursday (2/14)
- ❖ Homework 3 due Friday (2/15)
- ❖ **Take Home Midterm** (Thursday 2/14)
 - Due that night!

Roadmap

C:

```
car *c = malloc(sizeof(car));
c->miles = 100;
c->gals = 17;
float mpg = get_mpg(c);
free(c);
```

Java:

```
Car c = new Car();
c.setMiles(100);
c.setGals(17);
float mpg =
    c.getMpg();
```


Assembly
language:

```
get_mpg:
    pushq    %rbp
    movq     %rsp, %rbp
    ...
    popq     %rbp
    ret
```

Machine
code:


```
0111010000011000
100011010000010000000010
1000100111000010
110000011111101000011111
```

Computer
system:



Memory & data
Integers & floats
x86 assembly
Procedures & stacks
Executables
Arrays & structs
Memory & caches
Processes
Virtual memory
Memory allocation
Java vs. C

OS:



Data Structures in Assembly

❖ Arrays

- One-dimensional
- Multi-dimensional (nested)
- Multi-level

❖ Structs

- Alignment

❖ ~~Unions~~

Multi-Level Array Example

Multi-Level Array Declaration(s):

```
int cmu[5] = { 1, 5, 2, 1, 3 };  
int uw[5] = { 9, 8, 1, 9, 5 };  
int ucb[5] = { 9, 4, 7, 2, 0 };
```

```
int* univ[3] = {uw, cmu, ucb};
```

2D Array Declaration:

```
zip_dig univ2D[3] = {  
    { 9, 8, 1, 9, 5 },  
    { 1, 5, 2, 1, 3 },  
    { 9, 4, 7, 2, 0 }  
};
```

Is a multi-level array the
same thing as a 2D array?

NO


One array declaration = one contiguous block of memory

Multi-Level Array Example

```
int cmu[5] = { 1, 5, 2, 1, 3 };  
int uw[5]  = { 9, 8, 1, 9, 5 };  
int ucb[5] = { 9, 4, 7, 2, 0 };
```

```
int* univ[3] = {uw, cmu, ucb};
```

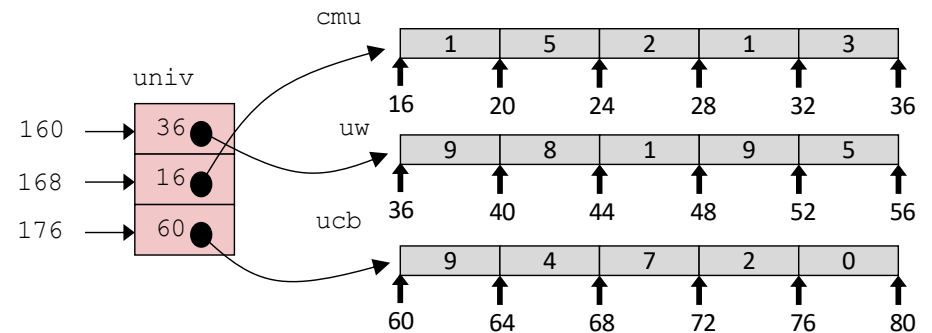
- ❖ Variable `univ` denotes array of 3 elements
- ❖ Each element is a pointer
 - 8 bytes each
- ❖ Each pointer points to array of `ints`



Note: this is how Java represents multi-dimensional arrays

Element Access in Multi-Level Array

```
int get_univ_digit
(int index, int digit)
{
    return univ[index][digit];
}
```



```
salq    $2, %rsi           # rsi = 4*digit
addq    univ(,%rdi,8), %rsi # p = univ[index] + 4*digit
movl    (%rsi), %eax        # return *p
ret
```


❖ Computation

- Element access $\text{Mem}[\text{Mem}[\text{univ} + 8 * \text{index}] + 4 * \text{digit}]$
- Must do **two memory reads**
 - First get pointer to row array
 - Then access element within array
- But allows inner arrays to be different lengths (not in this example)

Array Element Accesses


Nested array

```
int get_sea_digit
(int index, int digit)
{
    return sea[index][digit];
}
```



Multi-level array

```
int get_univ_digit
(int index, int digit)
{
    return univ[index][digit];
}
```




Access *looks* the same, but it isn't:

Mem[sea+20*index+4*digit]

Mem[**Mem**[univ+8*index]+4*digit]

Multi-Level Referencing Examples



<u>Reference</u>	<u>Address</u>	<u>Value</u>	<u>Guaranteed?</u>
<code>univ[2][3]</code>			
<code>univ[1][5]</code>			
<code>univ[2][-2]</code>			
<code>univ[3][-1]</code>			
<code>univ[1][12]</code>			

- C code does not do any bounds checking
- Location of each lower-level array in memory is *not* guaranteed

Summary

- ❖ Contiguous allocations of memory
- ❖ **No bounds checking** (and no default initialization)
- ❖ Can usually be treated like a pointer to first element
- ❖ **int** a[4][5]; → array of arrays
 - all levels in one contiguous block of memory
- ❖ **int*** b[4]; → array of pointers (to arrays)
 - First level in one contiguous block of memory
 - Each element in the first level points to another “sub” array
 - Parts anywhere in memory

Data Structures in Assembly

❖ Arrays

- One-dimensional
- Multi-dimensional (nested)
- Multi-level

❖ Structs

- Alignment

❖ ~~Unions~~

Structs in C

- ❖ Way of defining compound data types
- ❖ A structured group of variables, possibly including other structs

```
typedef struct {  
    int lengthInSeconds;  
    int yearRecorded;  
} Song;
```

```
Song song1;
```

```
song1.lengthInSeconds = 213;  
song1.yearRecorded    = 1994;
```

```
Song song2;
```

```
song2.lengthInSeconds = 248;  
song2.yearRecorded    = 1988;
```

```
typedef struct {  
    int lengthInSeconds;  
    int yearRecorded;  
} Song;
```

song1
lengthInSeconds: 213
yearRecorded: 1994

song2
lengthInSeconds: 248
yearRecorded: 1988

Struct Definitions

❖ Structure definition:

- Does NOT declare a variable
- Variable type is “**struct name**”

```
struct name {  
    /* fields */  
};
```

Easy to forget
semicolon!

```
struct name name1, *pn, name_ar[3];
```

pointer

array

❖ Joint struct definition and typedef

- Don't need to give struct a name in this case

```
struct nm {  
    /* fields */  
};  
typedef struct nm name;  
name n1;
```



```
typedef struct {  
    /* fields */  
} name;  
name n1;
```

Scope of Struct Definition

- ❖ Why is placement of struct definition important?
 - What actually happens when you declare a variable?
 - Creating space for it somewhere!
 - Without definition, program doesn't know how much space

```
struct data {  
    int ar[4];  
    long d;  
};
```

← Size = _____ bytes

Size = _____ bytes →

```
struct rec {  
    int a[4];  
    long i;  
    struct rec* next;  
};
```

- ❖ Almost always define structs in global scope near the top of your C file
 - Struct definitions follow normal rules of scope

Accessing Structure Members

- ❖ Given a struct instance, access member using the `.` operator:

```
struct rec r1;  
r1.i = val;
```

```
struct rec {  
    int a[4];  
    long i;  
    struct rec *next;  
};
```

- ❖ Given a *pointer* to a struct:

```
struct rec *r;  
r = &r1; // or malloc space for r to point to
```

We have two options:

- Use `*` and `.` operators: `(*r).i = val;`
- Use `->` operator for short: `r->i = val;`

- ❖ **In assembly:** register holds address of the first byte
 - Access members with offsets


Java side-note

```
class Record { ... }  
Record x = new Record();
```

- ❖ An instance of a class is like a *pointer to* a struct containing the fields
 - (Ignoring methods and subclassing for now)
 - So Java's $x.f$ is like C's $x \rightarrow f$ or $(*x).f$
- ❖ In Java, almost everything is a pointer ("*reference*") to an object
 - Cannot declare variables or fields that are structs or arrays
 - Always a *pointer* to a struct or array
 - So every Java variable or field is ≤ 8 bytes (but can point to lots of data)

Structure Representation

```
struct rec {  
    int a[4];  
    long i;  
    struct rec *next;  
};  
  
struct rec *r;
```




❖ Characteristics

- Contiguously-allocated region of memory
- Refer to members within structure by names
- Members may be of different types

Structure Representation


```
struct rec {  
    int a[4];  
    long i;  
    struct rec *next;  
};  
  
struct rec *r;
```



- ❖ Structure represented as block of memory
 - Big enough to hold all of the fields
- ❖ Fields ordered according to declaration order
 - Even if another ordering would be more compact
- ❖ Compiler determines overall size + positions of fields
 - Machine-level program has no understanding of the structures in the source code

Accessing a Structure Member

```
struct rec {  
    int a[4];  
    long i;  
    struct rec *next;  
};  
  
struct rec *r;
```



```
long get_i(struct rec *r)  
{  
    return r->i;  
}
```


❖ Compiler knows the *offset* of each member within a struct

- Compute as $(r + \text{offset})$
 - Referring to absolute offset, so no pointer arithmetic

```
# r in %rdi, index in %rsi  
movq 16(%rdi), %rax  
ret
```

Exercise: Pointer to Structure Member

```
struct rec {  
    int a[4];  
    long i;  
    struct rec *next;  
};  
  
struct rec *r;
```



```
long* addr_of_i(struct rec *r)  
{  
    return &(r->i);  
}
```


```
# r in %rdi  
_____, %rax  
ret
```

```
struct rec** addr_of_next(struct rec *r)  
{  
    return &(r->next);  
}
```

```
# r in %rdi  
_____, %rax  
ret
```

Generating Pointer to Array Element

```
struct rec {  
    int a[4];  
    long i;  
    struct rec *next;  
};  
  
struct rec *r;
```



❖ Generating Pointer to Array Element

- Offset of each structure member determined at compile time
- Compute as:
 $r + 4 * \text{index}$


```
int* find_addr_of_array_elem  
    (struct rec *r, long index)  
{  
    return &r->a[index];  
}
```

\searrow
 $\&(r \rightarrow a[\text{index}])$

```
# r in %rdi, index in %rsi  
leaq  (%rdi,%rsi,4), %rax  
ret
```


Nested Struct

```
struct foo {  
    long a;  
    long b;  
    struct bar my_bar;  
};  
  
struct bar {  
    long x;  
    long y;  
};  
  
struct foo *f;
```



Nested Struct

```
struct foo {  
    long a;  
    long b;  
    struct foo my_foo;  
};
```



Review: Memory Alignment in x86-64

- ❖ *Aligned* means that any primitive object of K bytes must have an address that is a multiple of K
- ❖ Aligned addresses for data types:

K	Type	Addresses
1	char	No restrictions
2	short	Lowest bit must be zero: $\dots 0_2$
4	int, float	Lowest 2 bits zero: $\dots 00_2$
8	long, double, *	Lowest 3 bits zero: $\dots 000_2$

Alignment Principles

❖ Aligned Data


- Primitive data type requires K bytes
- Address must be multiple of K
- Required on some machines; advised on x86-64

❖ Motivation for Aligning Data

- Memory accessed by (aligned) chunks of bytes (width is system dependent)
 - Inefficient to load or store value that spans quad word boundaries
 - Virtual memory trickier when value spans 2 pages (more on this later)
- Though x86-64 hardware will work regardless of alignment of data

Structures & Alignment

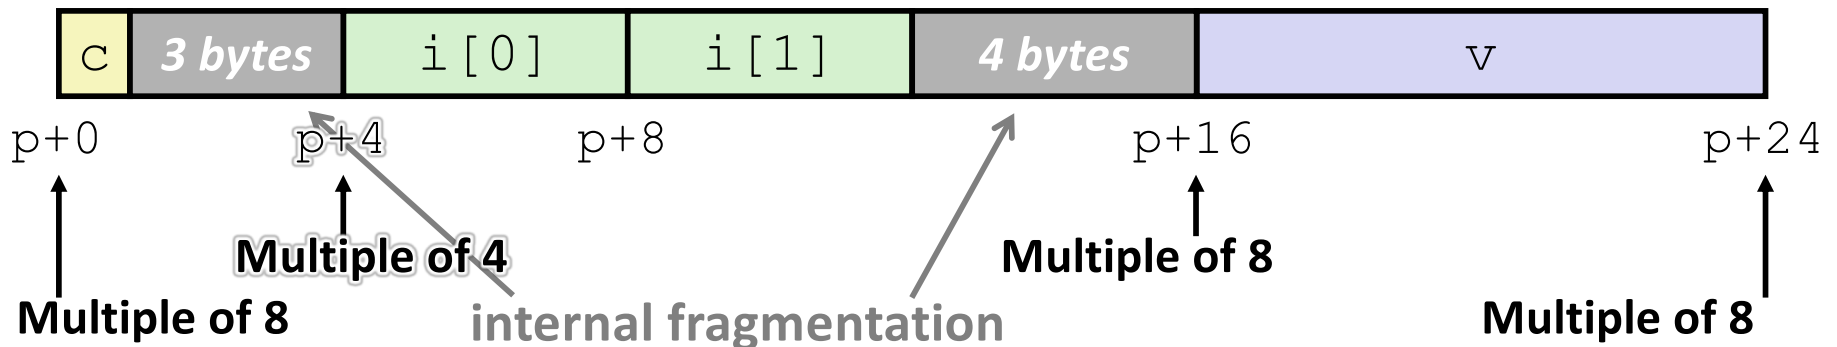
❖ Unaligned Data



```
struct S1 {  
    char c;  
    int i[2];  
    double v;  
} *p;
```

❖ Aligned Data

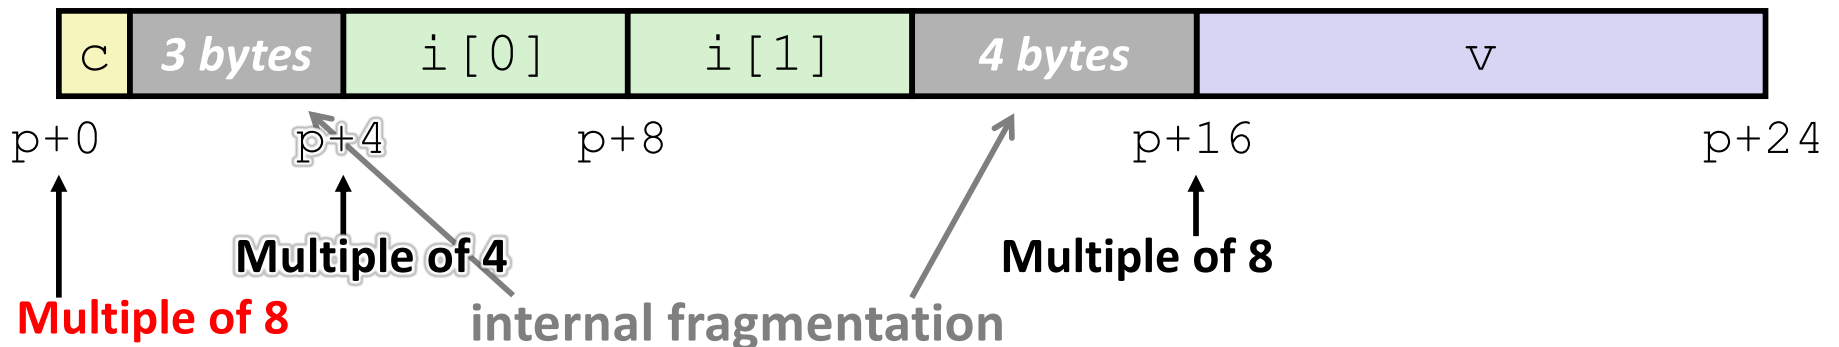
- Primitive data type requires K bytes
- Address must be multiple of K



Satisfying Alignment with Structures (1)

- ❖ Within structure:
 - Must satisfy each element's alignment requirement
- ❖ Overall structure placement
 - Each structure has alignment requirement K_{\max}
 - K_{\max} = Largest alignment of any element
 - Counts array elements individually as elements
 - Inner structs are aligned to *their* largest alignment
- ❖ Example:
 - $K_{\max} = 8$, due to `double` element

```
struct S1 {  
    char c;  
    int i[2];  
    double v;  
} *p;
```

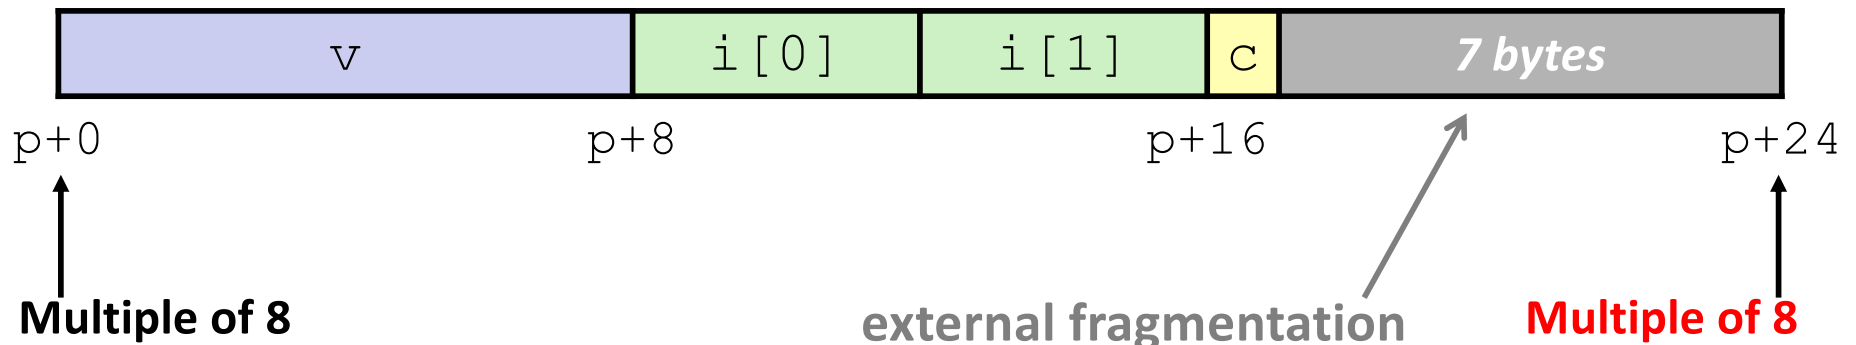


Satisfying Alignment with Structures (2)

- ❖ Can find offset of individual fields using `offsetof()`
 - Need to `#include <stddef.h>`
 - Example: `offsetof(struct S2, c)` returns 16

```
struct S2 {  
    double v;  
    int i[2];  
    char c;  
} *p;
```

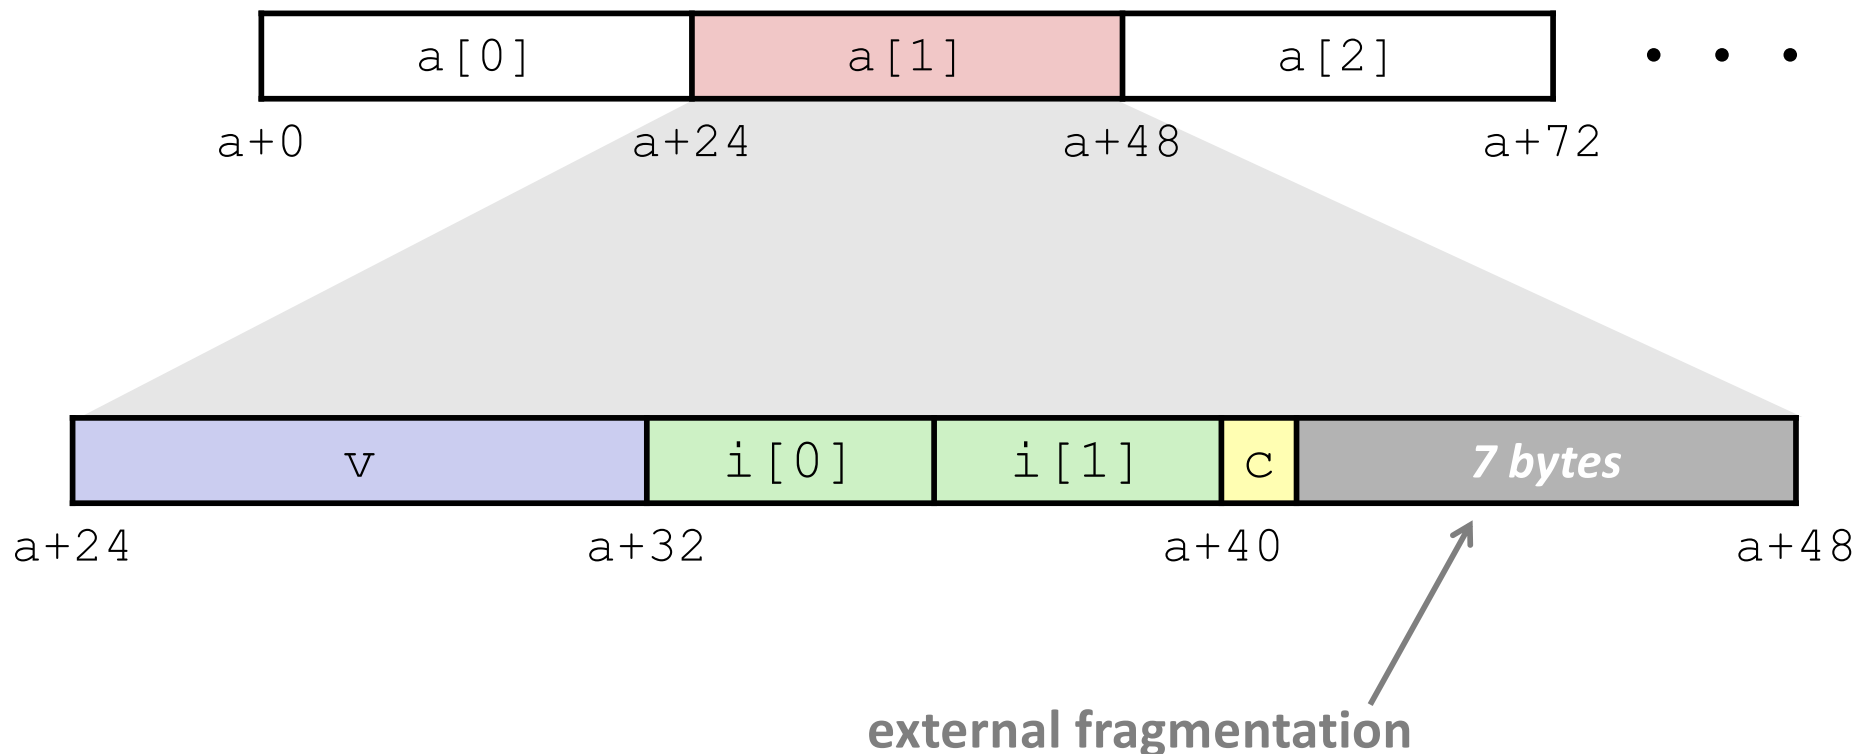
- ❖ For largest alignment requirement K_{\max} ,
overall structure size must be multiple of K_{\max}
 - Compiler will add padding **at end** of structure to meet overall structure alignment requirement



Arrays of Structures

- ❖ Overall structure length multiple of K_{max}
- ❖ Satisfy alignment requirement for every element in array

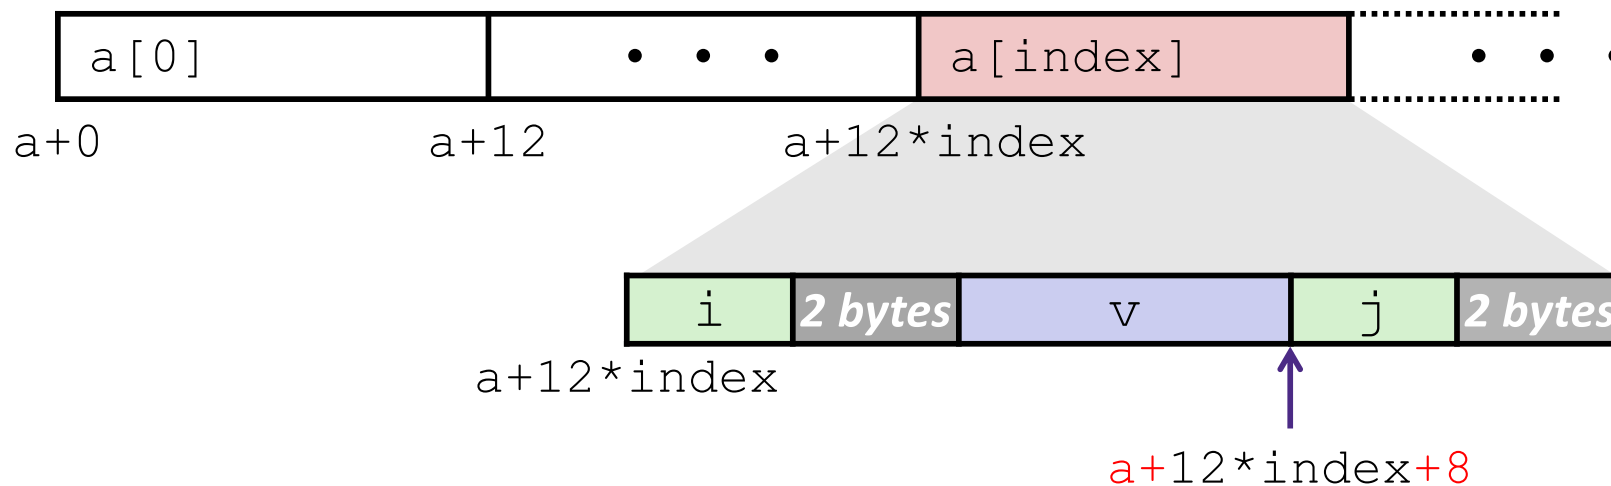
```
struct S2 {  
    double v;  
    int i[2];  
    char c;  
} a[10];
```



Accessing Array Elements

- ❖ Compute start of array element as: $12 * \text{index}$
 - `sizeof(S3) = 12`, including alignment padding
- ❖ Element `j` is at offset 8 within structure
- ❖ Assembler gives offset `a+8`

```
struct S3 {  
    short i;  
    float v;  
    short j;  
} a[10];
```



```
short get_j(int index)  
{  
    return a[index].j;  
}
```

```
# %rdi = index  
leaq (%rdi,%rdi,2),%rax # 3*index  
movzwl a+8(,%rax,4),%eax
```

Alignment of Structs

- ❖ Compiler will do the following:
 - Maintains declared *ordering* of fields in struct
 - Each **field** must be aligned *within* the struct (*may insert padding*)
 - `offsetof` can be used to get actual field offset
 - Overall struct must be **aligned** according to largest field
 - Total struct **size** must be multiple of its alignment (*may insert padding*)
 - `sizeof` should be used to get true size of structs



How the Programmer Can Save Space

- ❖ Compiler must respect order elements are declared in
 - Sometimes the programmer can save space by declaring large data types first

```
struct S4 {  
    char c;  
    int i;  
    char d;  
} *p;
```




```
struct S5 {  
    int i;  
    char c;  
    char d;  
} *p;
```



Peer Instruction Question

- ❖ Minimize the size of the struct by re-ordering the vars

```
struct old {  
    int i;  
  
    short s[3];  
  
    char *c;  
  
    float f;  
};
```



```
struct new {  
    int i;  
  
    _____;  
  
    _____;  
  
    _____;  
};
```

- ❖ What are the old and new sizes of the struct?

sizeof(struct old) = _____

sizeof(struct new) = _____

Summary

- ❖ Arrays in C
 - Aligned to satisfy every element's alignment requirement
- ❖ Structures
 - Allocate bytes in order declared
 - Pad in middle and at end to satisfy alignment