# Arrays, Structs & Alignment

#### CSE 351 Summer 2024

**Instructor:** Ellis Haker

#### **Teaching Assistants:**

Naama Amiel Micah Chang Shananda Dokka Nikolas McNamee Jiawei Huang



#### Administrivia

- Today
  - HW10 due (11:59pm)
  - Mid-Quarter Survey out on Canvas
    - Part of your participation grade
- Wednesday, 7/17
  - RD13 due (1pm)
  - HW11 due (11:59pm)
  - Mid-Quarter Survey due (11:59pm)

- Friday, 7/19
  - RD14 due (1pm)
  - HW12 due (11:59pm)
  - Lab2 due (11:59pm)
- Quiz 2 starts next Monday!
  - $\circ$   $\;$  Due the following week

# **Layers of Computing Revisited**

- Back to Hardware for today
  - How are compound data types (arrays, structs) stored in memory?
  - How do we get individual elements/fields out of them?
- Why now?
  - Knowing a little about assembly will help you understand this
  - Will be helpful for future lectures

Software Applications (written in Java, Python, C, etc.)

Programming Languages & Libraries (e.g. Java Runtime Env, C Standard Lib)

> Operating System (e.g. MacOS, Windows, Linux)

Hardware (e.g. CPU, memory, disk, network, peripherals)

#### **Review Questions**



1. How much space (in bytes) does one instance of struct ll\_node take?

#### **Lecture Topics**

- Arrays
  - Array review
  - Arrays in C
  - Multidimensional (nested) arrays
  - Multilevel arrays
- Structs
  - $\circ \quad \text{Structs in C}$
  - Struct memory layout
  - Alignment

Sector Se	Arrays start at 0	
Lua Matlab	Arrays start at 1	
Perl	Arrays can start wherever 「\_(ツ)_/「	
	Arrays start at 4, stop at 6, restart at 1, stop again at 3, restart at 7 then continue on	

#### **Recap: Arrays**

- T A[N]  $\rightarrow$  array A of type T and length N
  - <u>Contiguously</u> allocated region of N\*sizeof(T) bytes
  - Identifier A evaluates to the address of the array (type T\*)



#### **Recap: Array Accesses**



*a* = starting address of x

	Expression	Туре	Value
	х	int*	a
2 Levers	x[4]	int	5
x[2]	x + 1	int*	a + 4
Ĺ	<b>~~~</b> &x[2]	int *	o-t 8
seve -	→*(x + 1)	int	F
+613	x[5]	int	??? wherever is art abbress and

#### **Arrays in Memory Example**



#### **C** Details: Arrays and Pointers

- Array variables are *almost* identical to pointers
  - char\* string and char[] string are nearly identical declarations
  - Subtle differences: initialization, sizeof(), etc.
- An array name is an *expression*, <u>not</u> a variable!
  - $\circ$  Evaluates to the address of the first (0<sup>th</sup>) element
  - Translates to a label in assembly
    - Doesn't store the address in a variable (unlike a pointer)
    - Read-only, can't re-assign array name

#### **C** Details: Arrays and Functions

- Allocated on the stack, only allocated while scope is valid
  - Ex: What's wrong with this code?

```
char* foo() {
    char string[32];
    . . .
    return string;
}
```

• An array is passed into a function as a pointer



#### **Lecture Topics**

#### • Arrays

- Array review
- $\circ \quad \text{Arrays in C}$
- Multidimensional (nested) arrays
- Multilevel arrays
- Structs
  - $\circ \quad \text{Structs in C}$
  - Struct memory layout
  - Alignment

#### **Nested Array Example**

- sea evaluates to an int\*\*
- What's the layout in memory?

#### **Nested Array Example (pt 2)**



Row-major order: each row stored contiguously
 Guaranteed (in C)



## **Multi-Dimensional (Nested) Arrays**

- Declaration: T A [R] [C];
  - $\circ$  2D array of type T
  - **R** rows, **C** columns
  - Each element requires sizeof(T) bytes
- How big is this array?
  - o R\*C\*sizeof(T) bytes
- Arrangement: row-major ordering

Conceptual view:





#### Nested Array Row Access

- Given T A[R][C]
  - A[i] is the array of elements in row i
  - Pointer arithmetic:
    - A is the address of the start of the array
    - Starting address of row i = A + i\*C\*sizeof(T) int A[R][C]



#### **Nested Array Element Access**



i. (. lize of (t) = row affect

#### **Lecture Topics**

#### • Arrays

- $\circ$  Array review
- $\circ \quad \text{Arrays in C}$
- Multidimensional (nested) arrays
- Multilevel arrays
- Structs
  - $\circ \quad \text{Structs in C}$
  - Struct memory layout
  - Alignment

## <u>Multilevel</u> Array

This is how Jove stores "20 comps"



// Multi-level array
int\* univ[3] = {uw,
columbia, princeton};



#### **Multilevel Array <u>Element</u> Access**



- <u>Ex</u>: univ[1][3]
  - Requires two memory reads. 1) to get pointer to row array. 2) to get element.
  - Mem[Mem[univ + 1\*8] + 3\*4]

#### **TLDR: Array Element Accesses**

• Syntax looks the same, but memory layout is different



A[i][j] = Mem[A+(i\*C+j)\*sizeof(T)]

#### Multilevel



A[i][j] = Mem[Mem[A+i\*ptr\_size]+j\*sizeof(T)]

#### **Lecture Topics**

- Arrays
  - Array review
  - $\circ$  Arrays in C
  - Multidimensional (nested) arrays
  - Multilevel arrays
- Structs
  - Structs in C
  - Struct memory layout
  - Alignment

Me, who mostly has experience with Java and OOP, trying out C and just using structs and functions:



#### Structs in C

- User-defined compound data type <u>structured</u> group of variables
  - Kinda like a Java object, but no methods or inheritance, just fields

#### Example:

```
struct album {
    char* title;
    char* artist;
    int year_released;
};
```

```
struct album album1;
album1.title = "Heavy Rocks";
album1.artist = "Boris";
album1.year_released = 2002;
struct album album2;
album2.title = "Heathen";
album2.artist = "Thou";
album2.year_released = 2014;
```

#### **Struct Definition**

- Does not declare a variable, just defines the type
- Variable type is "struct <name>"
  - have to write "struct" every time :(
- Variable declarations are like any other data type:

struct album album1; // instance
struct album\* a\_ptr; // pointer
struct album playlist[5]; // array



• Can also combine struct definitions and instance declarations:

struct album { ... } a, \*ptr = &a; sene cs sene cs struct album 2 sine cs struct album 2 struct album 3 struct album

# Typedef in C

- A way to create an *alias* for another data type:
  - typedef <data type> <alias>;
  - Alias can be used interchangeably with the original data type
- Example: typedef unsigned long ul; unsigned long x = 12131989; ul x = 12131989;
- Can combine definition and typedef don't have to type "struct" every time



#### **Scope of Struct Definition**

- Why is the placement of a struct definition important?
  - $\circ$  C compiler is dumb, reads through file top—bottom
  - Declaring a variable creates space for it somewhere
    - Without type definition, program doesn't know how much space allocate!
- Structs follow normal C scope rules
- In practice, almost always define at the top of a file

Side note: similar rules apply to functions in C. Must define *before* they're used!

# **Accessing Struct Fields**

- Given a struct <u>instance</u>, use the '.' operator struct album a1; a1.title = "Deathconsciousness";
- Given a <u>pointer</u> to a struct:
   struct album\* p1 = &a1;
  - Two equivalent options:
    - Dereference, then use '•'

```
(*al).artist = "Have a Nice Life";
```

■ Use '->'

```
p1->year_released = 2008;
```

```
struct album {
    char* title;
    char* artist;
    int year_released;
};
```

#### Aside: Java

- An instance of a Java class is like a pointer ("reference") to a struct containing the fields
  - $\circ$  So Java's x.f is like C's x->f
  - Ignoring methods and inheritance, that's a future lecture!
- Arrays are similar. Allocated elsewhere, variable is a pointer to the array

#### **Lecture Topics**

#### • Arrays

- Array review
- $\circ \quad \text{Arrays in C}$
- Multidimensional (nested) arrays
- Multilevel arrays

#### • Structs

- $\circ \quad \text{Structs in C}$
- Struct memory layout
- Alignment

#### **Struct Representation**

```
struct album {
    char* title;
    char* artist;
    int year_released;
};
```



- Contiguously allocated in memory
- Fields ordered according to declaration order
  - Even if another ordering would be more efficient!
  - Why? Easier to debug, programmer can predict what offsets fields are at
- Compiler determines size + position of fields
  - Machine-level program has no understanding of structs

#### **Accessing Struct Fields (pt 2)**

```
struct album {
    char* title;
    char* artist;
    int year_released;
} a, p = &a;
```



• Compiler knows the offset of each field



#### **Pointers to Struct Fields**





• We can get the addresses of fields within a struct!



# **Recap: Alignment**

- Primitive-type data is stored at a starting address that is a **multiple of its size** 
  - Required on some systems, advised in x86
- Why?... Performance
  - Allows hardware optimizations
    - Some (non-x86) hardware will not work without it
  - Speeds up memory accesses (future lecture!)



#### Alignments in x86-64

Useful Gor Lobs 5!

- Primitive of *K* bytes must have an address that is divisible by *K*
- **Useful Fact**: if a number is divisible by 2<sup>n</sup>, then its binary representation will

end in <i>n</i> 0's	Data Type	Size (bytes)	Address
ety	char	1	No restrictions
	short	2	Lowest bit is 0: 0b0
ab 5.	int, float	4	Lowest 2 bits are 0: 0b00
	long, double	8	Lowest 3 bits are 0: 0b000
	long double	16	Lowest 4 bits are 0: 0b0000

# **Alignment in Structs**

- Each field must be aligned to the size of its data type
  - Internal fragmentation: unused space between fields to that each field is aligned



#### Unaligned

с	i[0]	i[1]	v
0	1		9

#### Aligned

с	3 bytes	i[0]	i[1]	4 bytes	v
0		4			16

How can we minimize internal fragmentation?

# Alignment in Structs (pt 2)

- Internal fragmentation: unused space between fields to that each field is aligned
  - Changing ordering of fields can reduce the amount of padding needed



# Before c 3 bytes i[0] i[1] 4 bytes v 0 4 16

After

7.11.01			
v	i[0]	i[1]	с
0	8		16

# Alignment in Structs (pt 3)

- <u>Within struct</u>: fields must be aligned to the size of their types
- Overall struct
  - Must be aligned to  $K_{max}$ 
    - K<sub>max</sub> = largest alignment requirement of any field
  - Add external fragmentation at the end of the struct to maintain alignment

#### Before

v	i[0]	i[1]	с
0	8		16

#### After

v	i[0]	i[1]	с	7 bytes
0	8		16	17



# **Why External Fragmentation?**

- Arrays of structs
  - If overall structure is aligned to  $K_{max}$ , each field will be aligned for every element in the array





## Alignment of Structs (pt 4)

- Compiler will do the following:
  - Maintain the order of fields specified in the C code
  - Add internal fragmentation to preserve alignment of fields
  - Add external fragmentation to preserve alignment of the overall struct

#### **Programming With Structs Tips**

- Because of padding, size of a struct != the sum of the sizes of each field
  - $\circ$  Use sizeof() to get the true size
- To save space, declare larger fields first



#### **Practice Question**

1. Minimize the size of the struct by re-ordering the fields.



1. sizeof(struct old) == 32 B. What sizeof(struct new)?



## **Summary: Arrays**

- Contiguously allocated
- Array name evaluates to starting address
  - Not a variable! Becomes a label in assembly
- Multidimensional arrays stored in row-major order: T A[R][C]
  - A[i] = array of row i = A + i\*C\*sizeof(T)
  - o A[i][j] = element j of row i = Mem[A + (i\*C + j)\*sizeof(T)]
- Multilevel arrays are arrays of *pointers* to other arrays: **T**\* A[**R**] = {...}
  - o A[i] = Mem[A + i\*sizeof(pointer)]
  - o A[i][j] = Mem[Mem[A+i\*sizeof(pointer)] + j\*sizeof(T)]

### **Summary: Structs**

- User defined types containing fields
  - Fields allocated in order declared by programmer
- Accessing Fields
  - $\circ$   $\,$  For an instance of the struct, use  $\, {\scriptstyle \bullet } \,$
  - $\circ$  For a pointer, use ->
    - Equivalent to \* (dereference) and then .
- Compiler adds padding to maintain alignment
  - Each field must be aligned to its size: add internal fragmentation
  - Struct must be aligned to  $K_{max}$ : add external fragmentation
  - Minimize fragmentation by ordering fields from largest to smallest in C code