




Fundamentals

Please follow these instructions to get set up:

github.com/rtfeldman/c-workshop-v1



roc-lang.org



zed.dev



Introduction

What is C?

Why is C still popular?

Why no Windows for this course?

Course overview

1. What is C?



PostgreSQL



redis



ANDROID




1972: initial release of C



Stevie Wonder - Superstition
(October 1972)

Stevie Wonder Superstition 1972



The Story of ABBA 1972 Creating History

1972: initial release of C



Dennis Ritchie, creator of C

Ken Thompson, UNIX designer


1972: initial release of C



PDP-11

1972: initial release of C





PDP-11

256KB memory

1-core CPU

125 KHz

1972: initial release of C



redis



PostgreSQL

MySQL



PDP-11

256KB memory
1-core CPU
0.001 GHz

UNIX®



MICROSOFT
WINDOWS



ANDROID



Linux™

mac
OS



iOS



Ruby



python

1972: initial release of C

What is C?

“Portable **assembly**”

Zero-overhead programming

Almost **zero-safety** programming

The language other languages use to
talk to each other

Why is C still popular?

Maximum possible performance

thanks to **Zero-overhead** programming

Much more ergonomic than Assembly

Much simpler than C++

Why is C still popular?

Local Static HTTP Server Perf Olympics

Time

9 ms

5 ms

3 ms

4 ms

5 ms

3 ms

5 ms

2 ms



http-server

↓ Weekly Downloads

2,360,975

Time

6 ms

1 ms

4 ms

3 ms

2 ms

2 ms

2 ms

2 ms



simple-http-server



52,089

Downloads all time

Time

2 ms

3 ms

2 ms

2 ms

1 ms

3 ms

2 ms

1 ms



the ~300 lines of C code we write in this workshop using only the stdlib

Low-overhead C alternatives

C++ (very complex, huge in game dev)

Rust (Zed, Biome, Ripgrep, ...) - my course!

Zig (Ghostty, Bun, TigerBeetle, Roc, ...)

many others (D, Odin, JAI, Carbon, ...)

Low-overhead C alternatives

all support C interop

because C is **the** language that

other languages use to talk to each other

C is the most **universal** language!

Why no Windows for this course?

We're using OS-specific APIs

macOS, Linux, and BSD all use POSIX APIs

WSL also supports POSIX APIs

Windows-specific APIs are often similar

Course Overview

Strings & Numbers

File I/O

Static Webserver

Node.js C Addon

Part 1: Hello, Metal!

"Hello, World!" in C

What the hardware sees

Performance implications

Improving ergonomics

HELLO WORLD

```
#include <unistd.h> imports write()  
main() is a function that returns nothing ("void")  
void main() {  
    write(1, "Hello, World!", 13);  
}
```

1 means write to stdout

HELLO WORLD

```
#include <unistd.h>
```

```
void main() {  
    write(1, "Hello, World!", 13);  
}
```

13 is the length of the string

WHAT THE HARDWARE SEES

```
.LC0:  
    .string "Hello, World!"
```

```
main:    "LC0" is "local constant number zero"  
    mov  edx, 13  
    mov  esi, OFFSET FLAT:.LC0  
    mov  edi, 1  
    jmp  write
```

```
.LC0:  
  .string "Hello, World!"
```

```
main:
```

```
mov edx, 13  
mov esi, OFFSET FLAT:.LC0  
mov edi, 1  
jmp write
```

```
write(1, "Hello, World!", 13);
```

```
.LC0:  
    .string "Hello, World!"
```

```
main: "set these CPU memory slots"
```

```
    mov edx, 13  
    mov esi, OFFSET FLAT:.LC0  
    mov edi, 1  
    jmp write
```

```
write(1, "Hello, World!", 13)
```

```
.LC0:  
    .string "Hello, World!"
```

```
main:
```

```
    mov edx, 13
```

```
    mov esi, OFFSET FLAT:.LC0
```

```
    mov edi, 1
```

```
    jmp write    "jump to the OS-provided write function"
```

(the OS will read from those 3 CPU memory slots)

Q: How fast does this program run?

A: As fast as the hardware can possibly run it!

.LC0:

.string "Hello, World!"

main:

mov edx, 13

mov esi, OFFSET FLAT:.LC0

mov edi, 1

jmp write only the OS can make write faster

This is the minimal set of instructions. ***Unbeatably fast!***



```
1 .LC0:
2 .string "Hello, World!"
3 main:
4 mov edx, 13
5 mov esi, OFFSET FLAT:.LC0
6 mov edi, 1
7 jmp write
```



```
1 .LC0:
2 .string "Hello World!"
3 main:
4 push rbp
5 mov rbp, rsp
6 mov esi, OFFSET FLAT:.LC0
7 mov edi, OFFSET FLAT:std::cout
8 call std::basic_ostream<char, std::char_traits<char>>
9 mov eax, 0
10 pop rbp
11 ret
```



```
1 main_main_pc0:
2 TEXT main.main(SB), ABIInternal, $64-0
3 CMPQ SF, 16(R14)
4 PCDATA $0, $-2
5 JLS main_main_pc76
6 PCDATA $0, $-1
7 PUSHQ BP
8 MOVQ SF, BP
9 SUBQ $56, SP
10 FUNCDATA $0, gcllocals.g2BeySu+wFnoycgXfElmcg==(SB)
11 FUNCDATA $1, gcllocals.EaPwxs275yYlhHMVZLak6g==(SB)
12 FUNCDATA $2, main.main.stkobj(SB)
13 LEAQ type:string(SB), DX
14 MOVQ DX, main..autotmp_8+40(SP)
15 LEAQ main..stmp_0(SB), DX
16 MOVQ DX, main..autotmp_8+48(SP)
17 MOVQ os.Stdout(SB), BX
18 NOP
19 LEAQ go:itab.*os.File,io.Writer(SB), AX
20 LEAQ main..autotmp_8+40(SP), CX
21 MOVL $1, DI
22 MOVQ DI, SI
23 PCDATA $1, $0
24 CALL fmt.Fprintln(SB)
25 ADDQ $56, SP
26 POPQ BP
27 RET
28 main_main_pc76:
29 NOP
30 PCDATA $1, $-1
31 PCDATA $0, $-2
32 CALL runtime.morestack_noctxt(SB)
33 PCDATA $0, $-1
34 JMP main_main_pc0
35 TEXT type:.eq.sync/atomic.Pointer[os.dirInfo](SB), DUPOK|NOSPLIT|NOFRAME|ABIInternal, $0-16
36 FUNCDATA $0, gcllocals.TjPuuCwd1CpTaRQGRKTrYw==(SB)
37 FUNCDATA $1, gcllocals.J5F+7Qw7O7ve2QcWC7DpeQ==(SB)
38 FUNCDATA $5, type:.eq.sync/atomic.Pointer[os.dirInfo].arginfo1(SB)
39 FUNCDATA $6, type:.eq.sync/atomic.Pointer[os.dirInfo].argliveinfo(SB)
40 PCDATA $3, $1
41 MOVQ (AX), CX
42 CMPQ (BX), CX
43 SETEQ AL
44 RET
```



```
1 .LC0:
2 .string "Hello, World!"
3 main:
4 mov edx, 13
5 mov esi, OFFSET FLAT:._LC0
6 mov edi, 1
7 jmp write
```

```
1 .LC0:
2 .string "Hello World!"
3 main:
4 push rbp
5 mov rbp, rsp
6 mov esi, OFFSET FLAT:._LC0
7 mov edi, OFFSET FLAT:std::cout
8 call std::basic_ostream<char, std::char_traits<char>>
9 mov eax, 0
10 pop rbp
11 ret
```

Rust, Zig, JAI, Odin, Carbon...

compared to C:

more ergonomic features

minimal extra overhead

HELLO WORLD

```
#include <unistd.h>
```

```
void main() {  
    write(1, "Hello, World!", 13);  
}
```

If anything benchmarks faster than this...

...then the benchmark must be wrong!

HELLO WORLD

```
#include <unistd.h>
```

```
int main() {  
    write(1, "Hello, World!", 13);  
  
    return 0;  
}
```

HELLO WORLD

```
#include <stdio.h>
(write is in unistd.h, printf is in stdio.h)
int main() {
    printf("Hello, World!");
    wrapper around write()
    return 0;
}
```

HELLO WORLD

```
#include <stdio.h>
```

```
int main() {  
    printf("Hello, World!");  
    "formatted"  
    return 0;  
}
```

HELLO WORLD

```
#include <stdio.h>
```

```
int main() {  
    int num = 42;  
    printf("The number is: %d", num);  
  
    return 0;  
}
```

HELLO WORLD

```
#include <stdio.h>
```

```
int main() {  
    int num = 42;  
    printf("The number is: %d", num);  
    "formatted"  
    return 0;  
}
```

HELLO WORLD

```
#include <stdio.h>
```

```
int main() {  
    int num = 42;  
    printf("The number is: %d", num);  
    Why "print"? And why no underscore before f?  
    return 0;  
}
```



PDP-11

1972: initial release of C

Summary of Part 1

"Hello, World!" in C

What the hardware sees

Performance implications

Improving ergonomics

Part 1 Exercise

Open `exercises/1.c`

See comments with 🙌 in them

Try doing what those comments say!

Part 2: Building HTTP Responses

Numbers in memory

Byte arrays

Null-terminated strings

Getting a string's length

00000000 00000001 00000010 00000011

0

$0+2^0$

$0+2^1$

$0+2^1+2^0$

0

1

2

3

4

00000100

5

00000101

6

00000110

7

00000111

8

00001000

HELLO WORLD

```
#include <unistd.h>
```

```
int main() {  
    write(1, "Hello, World!", 13);  
  
    return 0;  
}
```


01001000 01100101 01101100 01101100 01101111

48 65 108 108 111
'H' 'e' 'l' 'l' 'o'

ASCII → UTF-8

```
write(1, "Hello, World!", 13);
```

we decide how to interpret memory

write actually takes three integers!

```
write(1, "Hello, World!", 13);
```

01001000	01100101	01101100	01101100	01101111
----------	----------	----------	----------	----------

48	65	108	108	111
'H'	'e'	'l'	'l'	'o'

```
write(1, 547362942671452, 13);
```

the **memory address** of the
first byte in the string

01001000	01100101	01101100	01101100	01101111
----------	----------	----------	----------	----------

547362942671452

547362942671453

547362942671454
...

memory is basically a **gigantic array of bytes**

```
write(1, start index length  
547362942671452, 13);
```

the **memory address** of the
first byte in the string

memory is basically a **gigantic array of bytes**

start index

length

```
write(1, "Hello, World!", 13);
```

```
547362942671452
```

“Pointer” means “Memory Address”



Tsoding
@tsoding

Why people constantly complain that they don't understand pointers?
Pointer is just an index in a global array of bytes that we call "memory".
What's the problem?

7:20 AM · Feb 18, 2024 · **246.9K** Views

(“Monads are just monoids in the category of endofunctors, what's the problem?”)

HTTP Responses

<https://frontendmasters.com>

HTTP/1.1 200 OK

```
<!doctype html><html
```

72 84 84 80 47 49 46 49 32 50 48 48 32 79 75
H T T P / 1 . 1 2 0 0 0 K

"HTTP/1.1 200 OK"

72 84 84 80 47 49 46 49 32 50 48 48 32 79 75 0
H T T P / 1 . 1 2 0 0 0 K

"HTTP/1.1 200 OK"

72	84	84	80	47	49	46	49	32	50	48	48	32	79	75	0
H	T	T	P	/	1	.	1		2	0	0		0	K	

```
char *header = "HTTP/1.1 200 OK";
```

"address of a byte" (**char** means "byte")

aka "pointer to a byte" ("pointer" means "mem address")

in memory, **header** is an integer

(namely, the memory address of the first byte in the array)

72 84 84 80 47 49 46 49 32 50 48 48 32 79 75 0
H T T P / 1 . 1 2 0 0 0 K

```
char *header = "HTTP/1.1 200 OK";
```

```
write(1, header, 15);
```

72	84	84	80	47	49	46	49	32	50	48	48	32	79	75	0
H	T	T	P	/	1	.	1		2	0	0		0	K	

```
char *header = "HTTP/1.1 200 OK";
```

```
write(1, header, strlen(header));
```

```
72 84 84 80 47 49 46 49 32 50 48 48 32 79 75 0  
H T T P / 1 . 1 2 0 0 0 K
```

```
char *header = 415325;
```

```
write(1, header, strlen(header));
```

72	84	84	80	47	49	46	49	32	50	48	48	32	79	75	0
H	T	T	P	/	1	.	1		2	0	0		0	K	

```
char header[] = 415325;
```

```
write(1, 415325, strlen(415325));
```

72 84 84 80 47 49 46 49 32 50 48 48 32 79 75 0
H T T P / 1 . 1 / 2 0 0 0 K

415325
415326
415327

```
char *header = 415325;
```

```
write(1, 415325, strlen(415325));
```

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
72	84	84	80	47	49	46	49	32	50	48	48	32	79	75	0
H	T	T	P	/	1	.	1		2	0	0		0	K	

NULL-terminated strings


415325

```
char *header = 415325;
```

```
write(1, 415325, strlen(415325));
```

Isn't `strlen` incredibly slow?!

Why not store the length up front?



PDP-11

256KB memory

1-core CPU

125 KHz

1972: initial release of C

```
char *header = "HTTP/1.1 200 OK";  
write(1, header, strlen(header));  
15
```

Isn't `strlen` incredibly slow?

Yes, but this gets optimized away

```
#include <unistd.h>
```

```
char *header = "HTTP/1.1 200 OK";  
write(1, header, strlen(header));
```

```
#include <unistd.h>  
#include <string.h>
```

```
char *header = "HTTP/1.1 200 OK";  
write(1, header, strlen(header));
```

```
#include <unistd.h>
#include <string.h>
```

```
int main() {
    char header[] = "HTTP/1.1 200 OK";
    write(1, header, strlen(header));
    return 0;
}
```

```
#include <stdio.h>
#include <string.h>
```

```
int main() {
    char *header = "HTTP/1.1 200 OK";
    printf("Header: %s", header);
    return 0;
}
```

prints it as a string
(assumes null-terminated!)

```
#include <stdio.h>
#include <string.h>
```

```
int main() {
    char *header = "HTTP/1.1 200 OK";
    printf("Header: %d", header);
    return 0;
}
```

prints it as a 32-bit integer

```
#include <stdio.h>
#include <string.h>

int main() {
    char *header = "HTTP/1.1 200 OK";
    printf("Header: %zud", header);
    return 0;
}
```

prints it as a 64-bit integer

memory is basically a **gigantic array of bytes**

Summary of Part 2

Numbers in memory

Byte arrays

Null-terminated strings

Getting a string's length

Part 2 Exercise

Open `exercises/2.c`

See comments with  in them

Try doing what those comments say!

Part 3: Parsing HTTP Requests

Defining new functions

Iteration

Copying memory

Readonly memory

Incoming request:

GET /blog HTTP/1.1

Host: www.example.com

User-Agent: ...

File to open:

blog/index.html

Incoming request:

GET /blog

GET /blog/

File to open:

blog/index.html

blog/index.html

```
#include <stdio.h>
#include <string.h>

int main() {
    char *req = "GET /blog HTTP/1.1...";

    return 0;
}
```

```
#include <stdio.h>
#include <string.h>

int main() {
    char *req = "GET /blog HTTP/1.1...";
    char *path = to_path(req);

    return 0;
}
```

```
#include <stdio.h>
#include <string.h>

int main() {
    char *req = "GET /blog HTTP/1.1...";
    char *path = to_path(req);

    printf("Path: %s", path);

    return 0;
}
```

```
#include <stdio.h>
#include <string.h>
```

```
int main() {
    char *req = "GET /blog HTTP/1.1...";
    char *path = to_path(req);
    // ... ERROR: called before defined!
}
```

```
char *to_path(char *req) {
    // ...
```

```
char *to_path(char *req);  
// Forward-declare to_path
```

```
int main() {  
    char *req = "GET /blog HTTP/1.1...";  
    char *path = to_path(req);  
    // ...  
} ERROR: called before defined!
```

```
char *to_path(char *req) {  
    // ...
```

```
char *to_path(char *req) {
```

```
}
```

```
int main() {
```

```
integer integer
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:  "blog/index.html"
```

```
}
```

```
int main() {
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:            "blog/index.html"  
    char *start = req;
```

```
}
```

```
int main() {
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "blog/index.html"  
    char *start = req;  
    while (start[0] != ' ') {  
  
    }  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:        "blog/index.html"  
    char *start = req;  
    while (start[0] != ' ') {  
  
        start = start + 1;  
    }  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:        "blog/index.html"  
    char *start = req;  
    while (start[0] != ' ') {  
  
        start += 1;  
    }  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:            "blog/index.html"  
    char *start = req;  
    while (start[0] != ' ') {
```



```
        start++;
```

```
    }
```

```
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:        "blog/index.html"
    char *start = req;
    while (start[0] != ' ') {
        if (start[0] == '\\0')

            start++;
    }
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:        "blog/index.html"
    char *start = req;
    while (start[0] != ' ') {
        if (start[0] == 0)

            start++;
    }
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:        "blog/index.html"
    char *start = req;
    while (start[0] != ' ') {
        if (!start[0])

            start++;
    }
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:            "blog/index.html"
    char *start = req;
    while (start[0] != ' ') {
        if (!start[0])
            return NULL;

        start++;
    }
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:        "blog/index.html"  
    char *start = req;  
    while (start[0] != ' ') {  
        if (!start[0])  
            return NULL;  
        start++;  
    }  
}
```



```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:        "blog/index.html"
    char *start = req;
    while (start[0] != ' ') {
        if (!start[0]) {
            return NULL;
        }
        start++;
    }
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:            "blog/index.html"  
    char *start;  
    for (  
        start = req;  
        start[0] != ' ';  
        start++  
    ) {  
        if (!start[0]) { return NULL; }  
    }  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "blog/index.html"  
  
    for (    start is only in scope inside this loop  
           char *start = req;  
           start[0] != ' ';  
           start++)  
    {  
        if (!start[0]) { return NULL; }  
    }  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:            "blog/index.html"  
    char *start;  
    for (start = req; ... ) { ... }
```

```
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "/blog/index.html"  
    char *start;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "/blog/index.html"  
    char *start;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
  
    printf("Start: %s", start);  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "/blog/index.html"  
    char *start;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
  
    printf("Start: %s", start);  
    //      "Start: /blog HTTP/1.1..."  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET "  
    // Goal: ""  
    char *start;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
  
    printf("Start: %s", start);  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET "  
    // Goal: ""  
    char *start;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
    // start[0] == 0  
    printf("Start: %s", start);  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET "  
    // Goal: ""  
    char *start;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
    // start[0] == 0  
    printf("Start: %s", start);  
    //      "Start: "  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "/blog/index.html"  
    char *start, *end;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "/blog/index.html"  
    char *start, *end;  
    for (start = req; ...) { ... }  
    start++; // skip over the space  
    for (end = start; ...) { ... }  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "/blog/index.html"  
    char *start, *end; ... // set these
```

```
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "blog/index.html"  
    char *start, *end; ... // set these  
  
    // Ensure there's a "/" there  
  
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:  "blog/index.html"
    char *start, *end; ... // set these

    if (end[-1] == '/') {
        // e.g. if it's "GET /blog/..."
    }

}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog HTTP/1.1..."  
    // Goal:  "blog/index.html"  
    char *start, *end; ... // set these  
  
    if (end[-1] == '/') {  
        end--;  
    }  
  
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog HTTP/1.1..."
    // Goal:  "blog/index.html"
    char *start, *end; ... // set these

    if (end[-1] == '/') {
        end--;
    } else {
        end[0] = '/';
    }
    mutates original request string!
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog/ HTTP/1.1..."  
    // Goal:  "blog/index.html"  
    char *start, *end; ... // set these  
  
    // write "index.html\0" after '/'  
  
}
```

```
char *to_path(char *req) {  
    // Input: "GET /blog/ HTTP/1.1..."  
    // Goal:  "blog/index.html"  
    char *start, *end; ... // set these  
  
    // write "index.html\0" after '/'  
    memcpy(end + 1, "index.html", 11);  
  
}
```

```
char *to_path(char *req) {
    // Input: "GET /blog/ HTTP/1.1..."
    // Goal:  "blog/index.html"
    char *start, *end; ... // set these

    char *default_file = "index.html";
    memcpy(
        end + 1,
        default_file,
        strlen(default_file) + 1);
}
```

```
const char *DEFAULT_FILE = "index...";
char *to_path(char *req) {
    // Input: "GET /blog/ HTTP/1.1..."
    // Goal:  "blog/index.html"
    char *start, *end; ... // set these
```

```
    memcpy(
        end + 1,
        DEFAULT_FILE,
        strlen(DEFAULT_FILE) + 1);
```

```
}
```

```
const char *DEFAULT_FILE = "index...";
char *to_path(char *req) {
    // Input: "GET /blog/ HTTP/1.1..."
    // Goal:  "blog/index.html"
    char *start, *end; ... // set these

    memcpy(end + 1, DEFAULT_FILE, ...);

    return start;
}
```

```
const char *DEFAULT_FILE = "index...";
char *to_path(char *req) {
    // Input: "GET /blog/ HTTP/1.1..."
    // Goal:  "blog/index.html"
    char *start, *end; ... // set these

    memcpy(end + 1, DEFAULT_FILE, ...);
    // ⚠️ What if req is too short?
    // (This comes up in the exercise)
    return start;
}
```

```
const char *DEFAULT_FILE = "index...";

char *to_path(char *req) {
    ...
}

int main() {
    char *req = "GET /blog HTTP/1.1...";

}
```

```
const char *DEFAULT_FILE = "index...";
```

```
char *to_path(char *req) {  
    ... mutates original request string!  
}
```

```
int main() {  
    char *req = "GET /blog HTTP/1.1...";  
    char *path = to_path(req);  
                bus error  
    printf("Path: %s", path);  
}
```

this string is a **readonly** constant!

"GET /blog HTTP/1.1...";

```
const char *DEFAULT_FILE = "index...";

char *to_path(char *req) {
    ...
}

int main() {
    char req[] = "GET /blog HTTP/1.1...";
    char *path = to_path(req);

    printf("Path: %s", path);
}
```

Summary of Part 3

Defining new functions

Iteration

Copying memory

Readonly memory

Part 3 Exercise

Open `exercises/3.c`

See comments with 🙌 in them

Try doing what those comments say!

Part 4: File I/O

Opening a file

Getting its size

Reading its contents

Stack vs. heap

Incoming request:

GET /blog

GET /blog/

File to open:

blog/index.html

blog/index.html

```
char *path = "example.txt";
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
  
if (fd == -1) {  
    // Handle errors  
}
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);
```

```
if (fd == -1) {  
    // Handle errors  
}
```

```
char buffer[100];  
ssize_t length = read(fd, buffer, 100);
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);
```

```
if (fd == -1) {  
    // Handle errors  
}
```

```
char buffer[100];  
ssize_t length = read(fd, buffer, 100);  
if (length == -1) { ... } // Handle errors
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);
```

```
if (fd == -1) {  
    // Handle errors  
}
```

what if I want to read *all* the bytes?

```
char buffer[100];  
ssize_t length = read(fd, buffer, 100);  
if (length == -1) { ... } // Handle errors  
printf("Bytes read: %zd\n", length);
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
if (fd == -1) { ... } // Handle errors  
  
// 1. Get the size of the file on disk
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
if (fd == -1) { ... } // Handle errors  
  
// 1. Get the size of the file on disk  
// 2. Make a buffer of that many bytes
```

```
close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors

// 1. Get the size of the file on disk
// 2. Make a buffer of that many bytes
// 3. Read that many bytes into the buffer

close(fd);
```

```
struct stat {
    dev_t st_dev; // ID of device w/ file on it
    ino_t st_ino; // inode number
    mode_t st_mode; // protection
    link_t st_nlink; // number of hard links
    uid_t st_uid; // user ID of owner
    gid_t st_gid; // group ID of owner
    dev_t st_rdev; // device ID if special file
    off_t st_size; // total size, in bytes
    blksize_t st_blksize; // block size for I/O
    blkcnt_t st_blocks; // # of blocks allocated
    time_t st_atime; // last access time
    time_t st_mtime; // last modification time
    time_t st_ctime; // last status change time
}; // 144 bytes total
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
if (fd == -1) { ... } // Handle errors
```

```
struct stat metadata;
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
if (fd == -1) { ... } // Handle errors  
  
struct stat metadata; // 144 bytes of memory  
  
close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors
```

```
struct stat metadata; // 144 bytes of memory
```

```
if (fstat(fd, &metadata) == -1) {
    // Error
}
```

address of the first byte of those 144 bytes

```
close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors

struct stat metadata;

if (fstat(fd, &metadata) == -1) {
    // Error
}
// Now metadata has had its 144B populated

close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors
```

```
struct stat metadata;
```

Why doesn't fstat just return the stat struct?

```
if (fstat(fd, &metadata) == -1) {
```

```
    // Error
```

```
}
```

```
// Now metadata has had its 144B populated
printf("%ld bytes\n", metadata.st_size);
```

```
close(fd);
```

Why doesn't `fstat` just return the `stat` struct?

Returning the `stat` struct means returning all 144B

If a caller wants to return it, they copy all 144B again

If another caller wants to return it...another 144B copy

This would get expensive!

Why doesn't `fstat` just return a stat address?

When functions return, memory for their variables gets freed. (Otherwise, calling them would leak memory!)

This means they can't safely return addresses to variables they defined locally.

```
char *to_path(char *req) { ... }
```

address into an *argument's* memory, not a local variable's memory!

Why doesn't `fstat` just return a stat address?

Functions can reserve *long-lived* memory chunks using `malloc()`, and return the address of that memory.

However, that memory will leak unless some other function later calls `free()` on the address.

Calling `free()` at the wrong time can cause bugs!

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
if (fd == -1) { ... } // Handle errors
```

```
struct stat metadata;
```

```
if (fstat(fd, &metadata) == -1) {
```

```
    // Error
```

```
}
```

```
// Now metadata has had its 144B populated  
printf("%ld bytes\n", metadata.st_size);
```

```
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);
```

```
if (fd == -1) {  
    // Handle errors  
}
```

what if I want to read *all* the bytes?

```
char buffer[100];  
ssize_t length = read(fd, buffer, 100);  
if (length == -1) { ... } // Handle errors  
printf("Bytes read: %zd\n", length);  
  
close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors

struct stat metadata;
if (fstat(fd, &metadata) == -1) { ... }

char buffer[100];
ssize_t length = read(fd, buffer, 100);
if (length == -1) { ... } // Handle errors
printf("Bytes read: %zd\n", length);

close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors

struct stat metadata;
if (fstat(fd, &metadata) == -1) { ... }
    segmentation fault
char buffer[metadata.st_size];
ssize_t length = read(fd, buffer, 100);
if (length == -1) { ... } // Handle errors
printf("Bytes read: %zd\n", length);

close(fd);
```

```
char *path = "example.txt";
int fd = open(path, O_RDONLY);
if (fd == -1) { ... } // Handle errors
```

```
struct stat metadata; stack is fixed-size by default
if (fstat(fd, &metadata) == -1) { ... }
// Reserve `st_size` bytes on the stack
char buffer[metadata.st_size];
ssize_t length = read(fd, buffer, 100);
if (length == -1) { ... } // Handle errors
printf("Bytes read: %zd\n", length);

close(fd);
```

0101110101011101011011010110101010101010100100101000101100101111100101000101010101010101010101010

executable global **stack** heeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeap
+constants vars

```
struct stat metadata;  
if (fstat(fd, &metadata) == -1) { ... }  
// Reserve `st_size` bytes on the stack  
char buffer[metadata.st_size];  
ssize_t length = read(fd, buffer, 100);  
if (length == -1) { ... } // Handle errors  
printf("Bytes read: %zd\n", length);  
  
close(fd);
```

0101110101011101011011010110101010101010100100101000101100101111100101000101010101010101010101010

executable global **stack** heeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeap
+constants vars

```
struct stat metadata;  
if (fstat(fd, &metadata) == -1) { ... }  
// Reserve `st_size` bytes on the heap  
char *buffer = malloc(metadata.st_size);  
ssize_t length = read(fd, buffer, 100);  
if (length == -1) { ... } // Handle errors  
printf("Bytes read: %zd\n", length);  
  
close(fd);
```

0101110101011101011011010110101010101010100100101000101100101111100101000101010101010101010101010

executable global **stack** heeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeap
+constants vars **on macOS 15.3, reads over 2GB failed**

```
struct stat metadata; fix: read multiple chunks  
if (fstat(fd, &metadata) == -1) { ... }  
char *buffer = malloc(metadata.st_size);  
ssize_t length = fails if too many bytes read at once  
    read(fd, buffer, metadata.st_size);  
if (length == -1) { ... } // Handle errors  
printf("Bytes read: %zd\n", length);  
free(buffer);  
close(fd);
```

```
char *path = "example.txt";  
int fd = open(path, O_RDONLY);  
if (fd == -1) { ... } // Handle errors
```

```
struct stat metadata; fix: read multiple chunks  
if (fstat(fd, &metadata) == -1) { ... }
```

```
char buffer[100]; stack is good for reading chunks!  
ssize_t length = read(fd, buffer, 100);  
if (length == -1) { ... } // Handle errors  
printf("Bytes read: %zd\n", length);  
no need to free(), and better perf than malloc()  
close(fd);
```

Part 4 Summary

Opening a file

Getting its size

Reading its contents

Stack vs. heap

Part 5: Network I/O

Opening a socket

Listening for connections

Reading from a socket

Writing to a socket

`int` socket_fd = **LIBC** ❤️ **FILE DESCRIPTORS**

```
socket(AF_INET, SOCK_STREAM, 0);
```

IPv4

TCP

protocol number

(use AF_INET6 for IPv6)

(use SOCK_DGRAM for UDP)

```
int socket_fd = socket( ... );
```

```
int opt = 1;
```

```
setsockopt(  
    socket_fd,  
    SOL_SOCKET,  
    SO_REUSEADDR,  
    &opt,  
    sizeof(opt)
```

Without this, sometimes when you're stopping and restarting the server a lot, you can get "address in use" even when the server isn't running.

```
);
```

```
// Handle errors
```

```
int socket_fd = socket( ... );
```

```
int opt = 1;
```

Make a local variable

```
setsockopt(  
    socket_fd,  
    SOL_SOCKET,  
    SO_REUSEADDR,  
    &opt,  
    sizeof(opt)
```

(a "boolean" in this case)

and pass in its address.

```
);
```

```
// Handle errors
```

```
int socket_fd = socket( ... );  
int opt = 1;  
setsockopt(  
    socket_fd,  
    SOL_SOCKET,  
    SO_REUSEADDR,  
    &opt,  
    sizeof(opt)  
);  
// Handle errors
```

`sizeof` returns the number of bytes the argument takes up in memory (4 bytes for `int`)

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );
```

```
struct sockaddr_in address; // IPv4
```

```
address.sin_family = AF_INET;  
address.sin_addr.s_addr = INADDR_ANY;  
address.sin_port = htons(8080);
```

localhost:8080

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );
```

```
struct sockaddr_in address;  
bind(  
    socket_fd,  
    (struct sockaddr *)&address,  
    sizeof(address)  
);  
// Handle errors
```

```
struct sockaddr_in {  
    sa_family_t sin_family;  
    uint16_t sin_port;  
    ...  
};
```

```
struct sockaddr_in6 {  
    sa_family_t sin6_family;  
    uint16_t sin6_port;  
    ...  
};
```

```
struct sockaddr {  
    sa_family_t sa_family;  
    char sa_data[];  
};
```

```
struct sockaddr_in {
    sa_family_t sin_family;
    uint16_t sin_port;
    ...
};
```

```
struct sockaddr_in6 {
    sa_family_t sin6_family;
    uint16_t sin6_port;
    ...
};
```

2 bytes

```
struct sockaddr {
    sa_family_t sa_family;
    char sa_data[];
};
```

covers bytes of multiple fields!

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );
```

```
struct sockaddr_in address;
```

```
bind(  
    socket_fd, confirming that this memory address  
can be treated as a sockaddr struct  
    (struct sockaddr *)&address,  
    sizeof(address)  
);
```

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );
```

```
struct sockaddr_in address;
```

```
bind(  
    socket_fd,  
    (struct sockaddr *)&address,  
    sizeof(address)
```

```
);
```

Tells bind how many **actual** bytes are in the second argument, since it may be more than a sockaddr struct

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );  
struct sockaddr_in address; bind( ... );  
  
listen(socket_fd, 4);
```

number of pending connections that can
stack up before we start refusing new ones

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );  
struct sockaddr_in address; bind( ... );  
listen(socket_fd, 4);  
printf("Listening on localhost:8080");
```

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );  
struct sockaddr_in address; bind( ... );  
listen(socket_fd, 4);  
printf("Listening on localhost:8080");  
char req[MAX_REQUEST_BYTES + 1];
```

```
int socket_fd = socket( ... );  
int opt = 1; setsockopt( ... );  
struct sockaddr_in address; bind( ... );  
listen(socket_fd, 4);  
printf("Listening on localhost:8080");  
char req[MAX_REQUEST_BYTES + 1];  
int addrLen = sizeof(address);
```

```
int socket_fd = socket( ... );
int opt = 1; setsockopt( ... );
struct sockaddr_in address; bind( ... );
listen(socket_fd, 4);
printf("Listening on localhost:8080");
char req[MAX_REQUEST_BYTES + 1];
int addrlen = sizeof(address);
while (1) {
    // Wait for requests to come in
}
```

```
int addrLen = sizeof(address);
while (1) {    wait for the queue to get a connection,
              then dequeue it as a new socket
    int req_socket_fd = accept(
        socket_fd,
        (struct sockaddr *)&address,
        (socklen_t *)&addrLen
    );
    // Get data from req_socket_fd
    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read(
        req_socket_fd,
        req,
        MAX_REQUEST_BYTES
    );
    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read( ... );
    char *path = to_path(req);

    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read( ... );
    char *path = to_path(req);
    int fd = open(path, O_RDONLY);

    close(fd);
    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read( ... );
    int fd = ... // read file contents

    close(fd);
    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read( ... );
    int fd = ... // read file contents

    write(1, contents, length);
    close(fd);
    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read( ... );
    int fd = ... // read file contents

    write(req_socket_fd, contents, ...);
    close(fd);
    close(req_socket_fd);
}
```

```
char req[MAX_REQUEST_BYTES + 1];
while (1) {
    int req_socket_fd = accept( ... );
    ssize_t bytes_read = read( ... );
    int fd = ... // read file contents
    write(req_socket_fd, "200 OK\n\n" ...);
    write(req_socket_fd, contents, ...);
    close(fd);
    close(req_socket_fd);
}
```

Recap of `while (1)` Loop

`accept()` waits for a new connection, then gives us a fd

`read()` gets bytes from that connection (from browser)

`to_path()` parses those req bytes into a filesystem path

`open()`, `fstat()`, `read()` to get the file's full contents

`write()` the full contents to the socket (after a header)

Part 5 Summary

Opening a socket

Listening for connections

Reading from a socket

Writing to a socket

Part 6: Preprocessor Macros

Bit Shifts

Switches

Preprocessor Macros

Platform-specific code

Incoming request:

GET /blog

GET /blog/

File to open:

blog/index.html

blog/index.html

Incoming request:

GET /blog

GET /blog/

GET /blog.png

GET /blog/logo.png

File to open:

blog/index.html

blog/index.html

blog.png

blog/logo.png

number number

```
if (extension == "png") {  
    // When will this code run?  
}
```

```
    strcmp(extension, "png")
```

```
switch (extension) {  
    case "png":  
        content_type = "image/png";  
        break;  
}
```

```
if (!strcmp(extension, "png")) {  
    content_type = "image/png";  
}
```

```
if (!strcmp(extension, "png")) {  
    content_type = "image/png";  
} else if (!strcmp(extension, "css")) {  
    content_type = "text/css";  
}
```

```
if (!strcmp(extension, "png")) {
    content_type = "image/png";
} else if (!strcmp(extension, "css")) {
    content_type = "text/css";
} else if (!strcmp(extension, "js")) {
    content_type = "application/javascript";
}
```

```
if (!strcmp(extension, "png")) {
    content_type = "image/png";
} else if (!strcmp(extension, "css")) {
    content_type = "text/css";
} else if (!strcmp(extension, "js")) {
    content_type = "application/javascript";
} else if (!strcmp(extension, "html")) {
    content_type = "text/html";
}
```

01110000 01101110 01100111

'p' 112 'n' 110 'g' 103

"png"

"css"

"js"

"html"

1,752,460,652

01101000 01110100 01101101 01101100

01101000 01110100 01101101 01101100

'h' 104 't' 116 'm' 109 'l' 108

01110000 01101110 01100111 00000000

'p' 112 'n' 110 'g' 103 '\0' 0

"png\0"

1,886,283,520

"css"

"js"

"html"

1,752,460,652

01101000 01101000 01101001 01101100

01101000 01110100 01101101 01101100

'h' 104 't' 116 'm' 109 'l' 108

01110000 01101110 01100111 00000000

'p' 112 'n' 110 'g' 103 '\0' 0

"png\0"

1,886,283,520

"css\0"

"js"

"html"

1,752,460,652

01101000 01110100 01101101 01101100

01101000 01110100 01101101 01101100

'h' 104 't' 116 'm' 109 'l' 108

01110000 01101110 01100111 00000000

'p' 112 'n' 110 'g' 103 '\0' 0

"png\0"

1,886,283,520

"css\0"

"js\0\0"

"html"

1,752,460,652

01101000 01101000 01101001 01101100

01101000 01110100 01101101 01101100

'h' 104 't' 116 'm' 109 'l' 108

number

```
switch (extension) {  
  case "png":  
    type = "image/png";  
    break;  
}
```

```
const int PNG = 1886283520;
```

number

```
switch (extension) {
```

```
    case PNG:
```

```
        type = "image/png";
```

```
        break;
```

```
}
```

```
const int PNG = 1886283520;
```

“**cast**—that is, treat extension as an int's address”

```
switch (*(int *)extension) {  
    case PNG:  
        type = "image/png";  
        break;  
}
```

```
const int PNG = 1886283520;
```

“get whatever 4-byte int is at that address”

```
switch (*(int *)extension) {  
    case PNG:  
        type = "image/png";  
        break;  
}
```

```
const int PNG = 1886283520;
const int HTML = 1752460652;


switch (*(int *)extension) {
    case PNG:
        type = "image/png";
        break;
    case HTML:
        type = "text/html";
        break;
}
```

```
const int PNG = 1886283520;
const int HTML = 1752460652;

switch (*(int *)extension) {
    case PNG:
        type = "image/png";
        break;
    case HTML:
        type = "text/html";
        break;
    default:
        type = "application/octet-stream";
}
```

varies by **endianness**
(for this workshop,
don't worry about it)

Bit Shifts



Bit Shifts

```
'h' | ('t' << 8) | ('m' << 16) | ('l' << 24)
```

"bitwise OR"

Bit Shifts

```
'h' | ('t' << 8) | ('m' << 16) | ('l' << 24)
```

```
to_int('h', 't', 'm', 'l')
```

```
to_int('c', 's', 's', '\0')
```

```
to_int('j', 's', '\0', '\0')
```

```
const int PNG = 1886283520;
```

```
const int HTML = 1752460652;
```

Bit Shifts

```
'h' | ('t' << 8) | ('m' << 16) | ('l' << 24)
```

```
to_int('h', 't', 'm', 'l')
```

```
to_int('c', 's', 's', '\0')
```

```
to_int('j', 's', '\0', '\0')
```

```
const int PNG = to_int('p', 'n', 'g', '\0');
```

```
const int HTML = to_int('h', 't', 'm', 'l');
```

compile-time constant

```
const int PORT 8080;  
printf("Listening on localhost:%d\n", PORT);
```

```
#define PORT 8080;  
printf("Listening on localhost:%d\n", PORT);
```

```
#define PORT 8080;
printf("Listening on localhost:%d\n", PORT);

#define FOURCHAR(a, b, c, d) \
    a | (b << 8) | (c << 16) | (d << 24)
```

```
#define PORT 8080;
printf("Listening on localhost:%d\n", PORT);

#define FOURCHAR(a, b, c, d) \
    a | (b << 8) | (c << 16) | (d << 24)

const int HTML = FOURCHAR('h', 't', 'm', 'l');
const int CSS = FOURCHAR('c', 's', 's', '\0');
const int JS = FOURCHAR('j', 's', '\0', '\0');
```

01101000	01110100	01101101	01101100
----------	----------	----------	----------

'h' 104 't' 116 'm' 109 'l' 108

```
const int PNG = FOURCHAR('p', 'n', 'g', '\0');
const int HTML = FOURCHAR('h', 't', 'm', 'l');

switch (*(int *)extension) {
    case PNG:
        type = "image/png";
        break;
    case HTML:
        type = "text/html";
        break;
    default:
        type = "application/octet-stream";
}
```



this is what most languages compile to!

```
if (!strcmp(extension, "png")) {  
    content_type = "image/png";  
} else if (!strcmp(extension, "css")) {  
    content_type = "text/css";  
} else if (!strcmp(extension, "js")) {  
    content_type = "application/javascript";  
} else if (!strcmp(extension, "html")) {  
    content_type = "text/html";  
}
```

```
const int PNG = FOURCHAR('p', 'n', 'g', '\0');
const int HTML = FOURCHAR('h', 't', 'm', 'l');

switch (*(int *)extension) {
    case PNG:
        type = "image/png";
        break;
    case HTML:
        type = "text/html";
        break;
    default:
        type = "application/octet-stream";
}
```



This optimization relies on:

1. Knowing extensions fit in 4 bytes
2. Knowing memory is just 1s and 0s
3. Using a language that supports this

The `sendfile()` function

We `read()` from a file descriptor...

...then `write()` to another fd...

...can we avoid this handoff work?

`sendfile()` takes a src and dest fd



```
ssize_t sendfile(  
    int out_fd,  
    int in_fd,  
    off_t *offset,  
    size_t count  
)
```



```
int sendfile(  
    int fd,  
    int s,  
    off_t offset,  
    off_t *len,  
    sf_hdr *hdr,  
    int flags  
)
```

sendfile(2) — Linux manual page

[NAME](#) | [LIBRARY](#) | [SYNOPSIS](#) | [DESCRIPTION](#) | [RETURN VALUE](#) | [ERRORS](#) | [VERSIONS](#) | [STANDARDS](#) | [HISTORY](#) | [NOTES](#) | [SEE ALSO](#) | [COLOPHON](#)

sendfile(2)

System Calls Manual

sendfile(2)

NAME [top](#)

sendfile – transfer data between file descriptors

LIBRARY [top](#)

Standard C library (*libc*, *-lc*)

SYNOPSIS [top](#)

```
#include <sys/sendfile.h>
```

```
ssize_t sendfile(int out_fd, int in_fd, off_t *_Nullable offset,  
                 size_t count);
```

NAME

sendfile -- send a file to a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/uio.h>
```

```
int
sendfile(int fd, int s, off_t offset, off_t *len, struct sf_hdr *hdr,
         int flags);
```

DESCRIPTION

The **sendfile()** system call sends a regular file specified by descriptor fd out a stream socket specified by descriptor s.



```
ssize_t sendfile(  
    int out_fd,  
    int in_fd,  
    off_t *offset,  
    size_t count  
)
```



```
int sendfile(  
    int fd,  
    int s,  
    off_t offset,  
    off_t *len,  
    sf_hdr *hdr,  
    int flags  
)
```

What happened to “**portable assembly?!**”

Target-Specific Preprocessor Macros

```
#ifdef __linux__
```

```
#include <sys/sendfile.h>
```

```
#endif
```

Target-Specific Preprocessor Macros

```
#ifdef __linux__  
    // Linux-specific code  
#elif defined(__APPLE__)  
    // macOS-specific code  
#else  
    #error "Unsupported OS"  
#endif
```

Target-Specific Preprocessor Macros

```
#ifdef __linux__
    sendfile(/* args Linux expects */);
#elif defined(__APPLE__)
    sendfile(/* args macOS expects */);
#else
    #error "Unsupported OS"
#endif
```

Summary of Part 6

Bit Shifts

Switches

Preprocessor Macros

Platform-specific code

Wrap-Up

Course Summary

Further Resources

memory is basically a **gigantic array of bytes**

```
write(1, "Hello, World!", 13);  
547362942671452
```

start index **length**

“Pointer” means “Memory Address”



Tsoding
@tsoding

Why people constantly complain that they don't understand pointers?
Pointer is just an index in a global array of bytes that we call "memory".
What's the problem?

7:20 AM · Feb 18, 2024 · **246.9K** Views

(“Monads are just monoids in the category of endofunctors, what's the problem?”)

What is C?



PostgreSQL




redis



ANDROID





PDP-11

256KB memory

1-core CPU

125 KHz

1972: initial release of C

Why is C still popular?

Maximum possible performance

thanks to **Zero-overhead** programming

Much more ergonomic than Assembly

Much simpler than C++

Why is C still popular?

Local Static HTTP Server Perf Olympics

Time

9 ms

5 ms

3 ms

4 ms

5 ms

3 ms

5 ms

2 ms



http-server

↓ Weekly Downloads

2,360,975

Time

6 ms

1 ms

4 ms

3 ms

2 ms

2 ms

2 ms

2 ms



simple-http-server



52,089

Downloads all time

Time

2 ms

3 ms

2 ms

2 ms

1 ms

3 ms

2 ms

1 ms



the ~300 lines of C code we write in this workshop using only the stdlib

What is C?

“Portable **assembly**”

Zero-overhead programming

Almost **zero-safety** programming

The language other languages use to
talk to each other

Node.js

[About this documentation](#)

[Usage and example](#)

[Assertion testing](#)

[Asynchronous context tracking](#)

[Async hooks](#)

[Buffer](#)

[C++ addons](#)

[C/C++ addons with Node-API](#)

[C++ embedder API](#)

[Child processes](#)

[Cluster](#)

[Command-line options](#)

[Console](#)

Node.js v23.10.0 | [▶ Table of contents](#) | [▶ Index](#) | [▶ Other versions](#) | [▶ Options](#)

C++ addons

#

Addons are dynamically-linked shared objects written in C++. The `require()` function can load addons as ordinary Node.js modules. Addons provide an interface between JavaScript and C/C++ libraries.

There are three options for implementing addons:

- Node-API
- `nan` ([Native Abstractions for Node.js](#))
- direct use of internal V8, libuv, and Node.js libraries

Unless there is a need for direct access to functionality which is not exposed by Node-API, use Node-API. Refer to [C/C++ addons with Node-API](#) for more information on Node-API.

When not using Node-API, implementing addons becomes more complex, requiring knowledge of multiple components and APIs:

- **V8**: the C++ library Node.js uses to provide the JavaScript implementation. It provides the mechanisms for creating objects, calling functions, etc. The V8's API is documented mostly in the `v8.h` header file (`deps/v8/include/v8.h` in the Node.js source tree), and is also available [online](#).
- **libuv**: The C library that implements the Node.js event loop, its worker threads and all of the


Building Large C(++) Projects

“How to Build Software From Source”

Andrew Kelley (Zig language creator)

2023 Software You Can Love

youtu.be/CwXixVcliP0



Brian **K**ernighan

Dennis **R**itchie

“**K&R**”

Software optimization resources

[See also my blog](#)

Contents

agner.org/optimize

- [Optimization manuals](#)
 - [Vector class library](#)
 - [Object file converter and disassembler](#)
 - [Subroutine library](#)
 - [ForwardCom: An open standard instruction set for high performance microprocessors](#)
 - [Test programs for measuring clock cycles in C++ and assembly code](#)
 - [Floating point exception tracking and NaN propagation](#)
 - [CPUID manipulation program](#)
 - [Links](#)
-

Optimization manuals

This series of five manuals describes everything you need to know about optimizing code for x86 and x86-64 family microprocessors, including optimization advices for C++ and assembly language, details about the microarchitecture and instruction timings of most Intel, AMD and VIA processors, and details about different compilers and calling conventions.

Operating systems covered: DOS, Windows, Linux, BSD, Mac OS X Intel based, 32 and 64 bits.

Note that these manuals are not for beginners.

1. Optimizing software in C++: An optimization guide for Windows, Linux and Mac platforms

This is an optimization manual for advanced C++ programmers. Topics include: The choice of platform and operating system. Choice of compiler and framework. Finding performance bottlenecks. The efficiency of different C++ constructs. Multi-core systems. Parallelization with vector operations. CPU dispatching. Efficient container class templates. Etc.

File name: optimizing_cpp.pdf, size: 1798079, last modified: 2024-Mar-15.

[Download](#).



COMPUTER,
ENHANCE!

Performance-Aware Programming

computerenhance.com

Low-overhead C alternatives



JAI


Low-overhead C alternatives

all support C interop

because C is **the** language that

other languages use to talk to each other

C is the most **universal** language!



roc-lang.org

I host a podcast!



software-unscripted.com