

CSE 333

Lecture 8 - final C details, low-level I/O

Steve Gribble

Department of Computer Science & Engineering

University of Washington



Today's topics:

- a few final C details
 - ▶ header guards and other preprocessor tricks
 - ▶ extern, static and visibility of symbols
 - ▶ some topics for you to research on your own
- lower-level file access
 - ▶ open / read / write / close
 - instead of fopen / fread / fwrite / fclose

an #include problem

What happens when we compile foo.c?

```
typedef void *LinkedList;  
  
// more definitions below
```

ll.h

```
#include "ll.h"
```

ht.h

```
typedef void *HashTable;  
  
// A hypothetical function  
LinkedList HTKeyList(HashTable t);
```

```
#include "ll.h"  
#include "ht.h"
```

```
int main(int argc,  
          char **argv) {  
    // ... do stuff here ...  
    return 0;  
}
```

foo.c

an #include problem

What happens when we compile foo.c?

```
bash$ gcc -Wall -g -o foo foo.c
```


```
In file included from ht.h:1,  
                 from foo.c:2:
```

```
ll.h:1: error: redefinition of typedef 'LinkedList'
```

```
ll.h:1: note: previous declaration of 'LinkedList' was here
```

foo.c includes **ll.h** twice!

- 2nd time is indirectly via **ht.h**
- so, typedef shows up twice!
- *try using cpp to see this*



header guards

A commonly used C preprocessor trick to deal with this

- uses macro definition (`#define`)
- uses conditional compilation (`#ifndef` and `#endif`)

```
#ifndef _LL_H_                                     // .h
#define _LL_H_ 1

typedef void *LinkedList;

// more definitions below

#endif // _LL_H_
```

```
#ifndef _HT_H_                                     ht.h
#define _HT_H_

#include "ll.h"

typedef void *HashTable;

// A hypothetical function
LinkedList HTKeyList(HashTable t);

#endif // _HT_H_
```

Other preprocessor tricks

A way to deal with “magic constants”

```
int globalbuffer[1000];

void circalc(float rad,
             float *circumf,
             float *area) {
    *circumf = rad * 2.0 * 3.1415;
    *area = rad * 3.1415 * 3.1415;
}
```

bad code
(littered with magic constants)

```
#define BUFSIZE 1000
#define PI 3.14159265359

int globalbuffer[BUFSIZE];

void circalc(float rad,
             float *circumf,
             float *area) {
    *circumf = rad * 2.0 * PI;
    *area = rad * 3.1415 * PI;
}
```

better code

Macros

You can pass arguments to macros

```
#define ODD(x) ((x) % 2 != 0)

void foo(void) {
    if ( ODD(5) )
        printf("5 is odd!\n");
}
```

cpp →

```
void foo(void) {
    if ( ((5) % 2 != 0) )
        printf("5 is odd!\n");
}
```

Be careful of precedence issues; use parenthesis:

```
#define ODD(x) ((x) % 2 != 0)
#define BAD(x) x % 2 != 0

ODD(5 + 1);

BAD(5 + 1);
```

cpp →

```
((5 + 1) % 2 != 0);

5 + 1 % 2 != 0;
```

Namespace problem

If I define a global variable named “counter” in foo.c, is it visible in bar.c?

- if you use **external linkage**: yes
 - ▶ the name “**counter**” refers to the same variable in both files
 - ▶ the variable is defined in one file, declared in the other(s)
 - ▶ when the program is linked, the symbol resolves to one location
- if you use **internal linkage**: no
 - ▶ the name “**counter**” refers to different variables in each file
 - ▶ the variable must be defined in each file
 - ▶ when the program is linked, the symbols resolve to two locations

External linkage

```
#include <stdio.h>

// A global variable, defined and
// initialized here in foo.c.
// It has external linkage by
// default.
int counter = 1;

int main(int argc, char **argv) {
    printf("%d\n", counter);
    bar();
    printf("%d\n", counter);
    return 0;
}
```

foo.c

```
#include <stdio.h>

// "counter" is defined and
// initialized in foo.c.
// Here, we declare it, and
// specify external linkage
// by using the extern
// specifier.
extern int counter;

void bar(void) {
    counter++;
    printf("(b): counter %d\n",
           counter);
}
```

bar.c

Internal linkage

```
#include <stdio.h>

// A global variable, defined and
// initialized here in foo.c.
// We force internal linkage by
// using the static specifier.
static int counter = 1;

int main(int argc, char **argv) {
    printf("%d\n", counter);
    bar();
    printf("%d\n", counter);
    return 0;
}
```

foo.c

```
#include <stdio.h>

// A global variable, defined and
// initialized here in bar.c.
// We force internal linkage by
// using the static specifier.
static int counter = 100;

void bar(void) {
    counter++;
    printf("(b): counter %d\n",
           counter);
}
```

bar.c

Some gotchas

Every global (variables and functions) is extern by default.

- unless you specify the static specifier, if some other module uses the same name, you'll end up with a collision!
 - ▶ best case: compiler error
 - ▶ worst case: stomp all over each other
- it's good practice to:
 - ▶ use static to defend your globals (hide your private stuff!)
 - ▶ place external (i.e., global) declarations in a module's header file

Extern, static functions

```
// By using the static specifier, we are indicating
// that foo() should have internal linkage. Other
// .c files cannot see or invoke foo().
```

```
static int foo(int x) {
    return x*3 + 1;
}
```

```
// Bar is "extern" by default. Thus, other .c files
// could declare our bar() and invoke it.
```

```
int bar(int x) {
    return 2*foo(x);
}
```

bar.c

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

```
extern int bar(int);
```

```
int main(int argc, char **argv) {
    printf("%d\n", bar(5));
    return 0;
}
```

main.c

Somebody should get fired



C has a second, different use for the word “static”

- to declare the extent of a local variable
- if you declare a static local variable, then:
 - ▶ the storage for that variable is allocated when the program loads, in either the program’s .data or .bss segment
 - ▶ the variable retains its value across multiple function invocations

(see static_extent.c for an example)

Additional C topics

Homework: teach yourself the following topics

- bit-level manipulation in C: `~` `|` `&` `<<` `>>`
- string library functions provided by the C standard library
 - ▶ `#include <string.h>`
 - `strlen()`, `strcpy()`, `strdup()`, `strcat()`, `strcmp()`, `strchr()`, `strstr()`, ...
 - learn why **strncat** is safer (in the security sense) than **strcat**, etc.
 - ▶ `#include <stdlib.h>` or `#include <stdio.h>`
 - `atoi()`, `atof()`, `sprintf()`, `sscanf()`
- **man** pages are your friend!

Additional C topics

Teach yourself:

- the syntax for function pointers, including passing as args
- how to declare, define, and use a function that accepts a variable-lengthed number of arguments (varargs)
- unions and what they are good for
- what argc and argv are for in main

```
#include <stdio.h>

int main(int argc, char **argv) {
    int i;

    for (i = 0; i < argc, i++) {
        printf("%d: %s\n", i, argv[i]);
    }
    return 0;
}
```

argv.c

```
bash$ gcc -o argv argv.c
bash$ ./argv
0: ./argv
bash$ ./argv foo bar
0: ./argv
1: foo
2: bar
bash$
```

Additional C topics

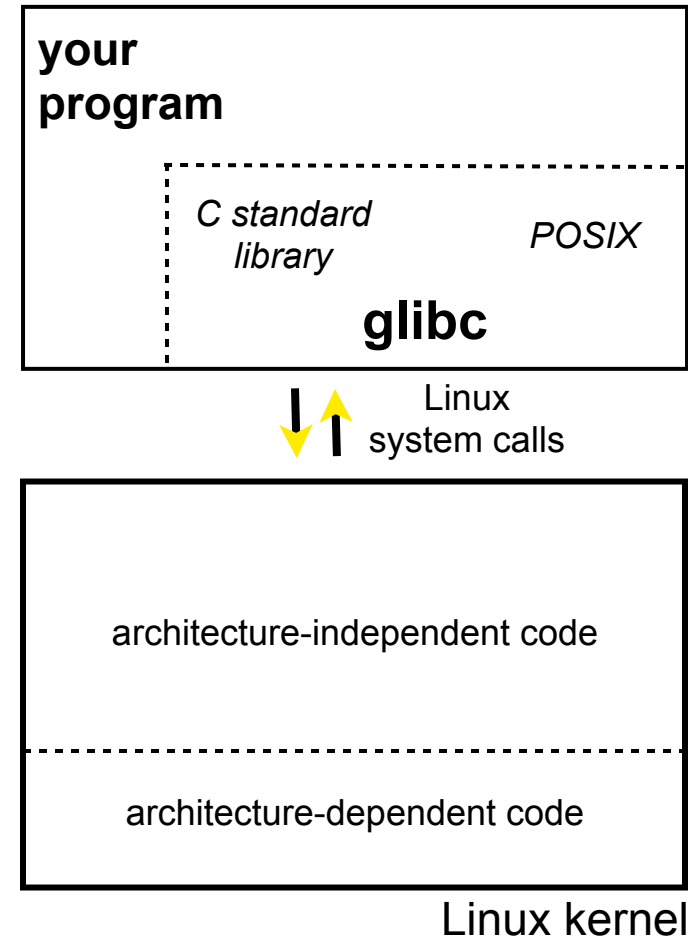
Teach yourself:

- the difference between pre-increment ($++v$) and post-increment ($v++$)
- the meaning of the “register” storage class
- harder: the meaning of the “volatile” storage class
 - pages 91, 92 of CARM

Lower-level file access

Remember this picture?

- your program can access many layers of APIs
 - ▶ C standard library
 - ▶ POSIX compatibility API
 - ▶ underlying OS system calls



So far...

You've used the C standard library to access files

- specifically, `fopen`, `fread`, `fwrite`, `fclose`, `fseek`
 - these provide a (FILE *) stream abstraction

These are convenient and portable...

- but, they are *buffered*
- and, they are implemented by using lower-level OS calls

Lower-level file access

Most UNIX-en support a common set of lower-level file access APIs

- open, read, write, close, fseek
 - ▶ similar in spirit to their fopen (etc.) counterparts
 - ▶ but, lower-level and unbuffered
 - ▶ and, less convenient
- you will have to use these for network I/O, so we might as well learn them now

open / close

To open a file...

- pass in the filename and access mode, similar to fopen
- get back a “file descriptor”
 - ▶ similar to a (FILE *) from fopen, but is just an int

```
#include <fcntl.h>

...

int fd = open("foo.txt",
              O_RDONLY);
if (fd == -1) {
    perror("open failed");
    exit(EXIT_FAILURE);
}

...

close(fd);
```

Reading from a file

```
ssize_t read(int fd, void *buf, size_t count);
```

- returns the # of bytes read
 - ▶ might be fewer bytes than you requested (!!!)
 - ▶ returns 0 if you're at end-of-file
 - ▶ return -1 on error
- warning: read has some very surprising error modes!

read() error modes

On error, the “errno” global variable is set

- you need to check it to see what kind of error happened

What errors might read() encounter?

- EBADF -- bad file descriptor
- EFAULT -- output buffer is not a valid address
- EINTR -- read was interrupted, please try again
 - ▶ argh!!!
- and many others

How to read() n bytes

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

...

char *buf = ...;
int bytes_left = n;
int result = 0;

while (bytes_left > 0) {
    result = read(fd, buf + (n-bytes_left), bytes_left);
    if (result == -1) {
        if (errno != EINTR) {
            // a real error happened, return an error result
        }
        // EINTR happened, do nothing and loop back around
        continue;
    }
    bytes_left -= result;
}
```

Other low-level functions

Read the man pages to learn about:

- **write()** -- write data
- **fsync()** -- flush data to the underlying device
- **opendir()**, **readdir()**, **closedir()** -- get a directory listing
 - ▶ make sure you read the section 3 version, e.g.:
 - man 3 opendir
 - ▶ kind of painful to use

A useful cheat-sheet

From a CMU systems programming course:

<http://www.cs.cmu.edu/~guna/15-123S11/Lectures/Lecture24.pdf>

Exercise 1

Write a program that:

- prompts the user to input a string (use `fgets()`)
 - ▶ assume the string is a sequence of whitespace-separated integers
 - ▶ e.g., “5555 1234 4 5543”
- converts the string into an array of integers
- converts an array of integers into an array of strings
 - ▶ where each element of the string array is the binary representation of the associated integer
- prints out the array of strings

Exercise 2

Modify the linked list code from last lecture / exercise 1

- add static declarations to any internal functions you implemented in `linkedlist.h`
- add a header guard to the header file
- write a Makefile
 - ▶ use Google to figure out how to add rules to the Makefile to produce a library (`liblinkedlist.a`) that contains the linked list code

See you on Wednesday!