

Linking

Kai Shen

1

Example C Program

main.c

```
int buf[2] = {1, 2};

int main()
{
    swap();
    return 0;
}
```

swap.c

```
extern int buf[];

int *bufp0 = &buf[0];
static int *bufp1;

void swap()
{
    int temp;

    bufp1 = &buf[1];
    temp = *bufp0;
    *bufp0 = *bufp1;
    *bufp1 = temp;
}
```

2

Static Linking

- Programs are translated and linked using a *compiler driver*:
 - `unix> gcc -O2 -g -o p main.c swap.c`
 - `unix> ./p`

main.c swap.c

Source files

Translators
(cpp, cc1, as)

Translators
(cpp, cc1, as)

Separately compiled relocatable object files

main.o

swap.o

Linker (ld)

*Fully linked executable object file
(contains code and data for all functions defined in main.c and swap.c)*

p

3

Why Linkers?

- Reason 1: Modularity
 - Program can be written as a collection of smaller source files, rather than one monolithic mass.
 - Can build libraries of common functions (more on this later)
 - e.g., Math library, standard C library
- Reason 2: Efficiency
 - Time: Separate compilation
 - Change one source file, compile, and then relink.
 - No need to recompile other source files.
 - Space: Libraries
 - Common functions can be aggregated into a single file...
 - Yet executable files and running memory images contain only code for the functions they actually use.

4

Linker Does Symbol Resolution

- Programs define and reference *symbols* (variables and functions):


```
void swap() {...} /* define symbol swap */
swap();          /* reference symbol a */
int *xp = &x;    /* define symbol xp, reference x */
```
- Symbol definitions are stored (by compiler) in *symbol table*.
 - Symbol table is an array of symbol entries;
 - Each entry includes name, size, and location of symbol.
- Linker associates each symbol reference with exactly one symbol definition.

5

Linker Does Relocation

- Merges separate code sections from multiple object files into single code section
 - same for data sections
- Relocates symbols from their relative locations in the object files to their final absolute memory locations in the executable.
- Updates all references to these symbols to reflect their new positions.

6

Three Kinds of Object Files (Modules)

- Relocatable object file (.o file or .a file)
 - Contains code and data in a form that can be combined with other relocatable object files to form executable object file.
- Executable object file (a.out file)
 - Contains code and data in a form that can be copied directly into memory and then executed.
- Dynamic link object file (.so file)
 - Special type of relocatable object file that can be loaded into memory and linked dynamically, at either load time or run-time.
 - Called *Dynamic Link Libraries* (DLLs) by Windows

7

Executable and Linkable Format (ELF)

- Standard binary format for object files
- Originally proposed by AT&T System V Unix
 - Later adopted by BSD Unix variants and Linux
- One unified format for
 - Relocatable object files (.o),
 - Executable object files (a.out)
 - Dynamic link object files (.so)
- Generic name: ELF binaries

8

ELF Object File Format

- ELF header
 - Word size, byte ordering, file type (.o, exec, .so), machine type, etc.
- Segment header table
 - Page size, virtual addresses memory segments (sections), segment sizes.
- .text section
 - Code
- .rodata section
 - Read only data: jump tables, ...
- .data section
 - Initialized global variables
- .bss section
 - Uninitialized global variables
 - "Block Started by Symbol"
 - "Better Save Space"
 - Has section header but occupies no space

ELF header
Segment header table (required for executables)
.text section
.rodata section
.data section
.bss section
.symtab section
.rel.txt section
.rel.data section
.debug section
Section header table

ELF Object File Format (cont.)

- .symtab section
 - Symbol table
 - Procedure and static variable names
 - Section names and locations
- .rel.text section
 - Relocation info for .text section
 - Addresses of instructions that will need to be modified in the executable
 - Instructions for modifying.
- .rel.data section
 - Relocation info for .data section
 - Addresses of pointer data that will need to be modified in the merged executable
- .debug section
 - Info for symbolic debugging (gcc -g)
- Section header table
 - Offsets and sizes of each section

ELF header
Segment header table (required for executables)
.text section
.rodata section
.data section
.bss section
.symtab section
.rel.txt section
.rel.data section
.debug section
Section header table

Linker Symbols

- Global symbols
 - Symbols defined by module *m* that can be referenced by other modules.
 - E.g.: non-**static** C functions and non-**static** global variables.
- External symbols
 - Global symbols that are referenced by module *m* but defined by some other module.
- Local symbols
 - Symbols that are defined and referenced exclusively by module *m*.
 - E.g.: C functions and variables defined with the **static** attribute.

Resolving Symbols

```

int buf[2] = {1, 2};

int main()
{
    swap();
    return 0;
}
main.c
            
```

```

extern int buf[];
int *bufp0 = &buf[0];
static int *bufp1;

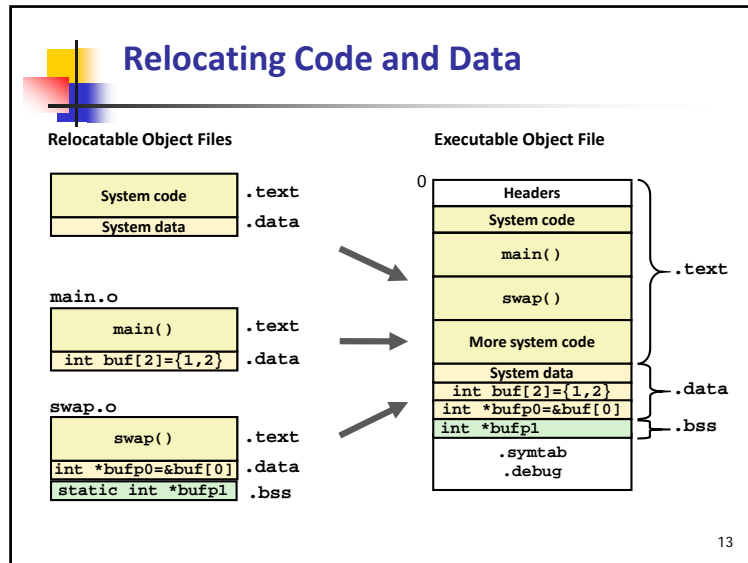
void swap()
{
    int temp;
    bufp1 = &buf[1];
    temp = *bufp0;
    *bufp0 = *bufp1;
    *bufp1 = temp;
}
swap.c
            
```

Global symbols: `buf` in `main.c`, `swap` in `swap.c`.

External symbols: `buf` in `swap.c` (referenced by `main.c`), `swap` in `main.c` (referenced by `swap.c`).

Local symbols: `temp` in `swap.c`.

Linker knows nothing of temp



Relocation

- Code section change due to relocation at link time

```
main.c
int buf[2] = {1,2};

int main()
{
    swap();
    return 0;
}
```

14

Relocation

- Data section change due to relocation at link time

```
swap.c
extern int buf[];

int *bufp0 = &buf[0];
static int *bufp1;

void swap() {
    int temp;

    bufp1 = &buf[1];
    temp = *bufp0;
    *bufp0 = *bufp1;
    *bufp1 = temp;
}
```

15

Packaging Commonly Used Functions

- How to package functions commonly used by programmers?
 - Math, I/O, memory management, string manipulation, etc.
- Awkward, given the linker framework so far:
 - Option 1:** Put all functions into a single source file
 - Programmers link big object file into their programs
 - Space and time inefficient
 - Option 2:** Put each function in a separate source file
 - Programmers explicitly link appropriate binaries into their programs
 - More efficient, but burdensome on the programmer

16

Solution: Libraries

- Concatenate related relocatable object files into a single file with an index (called an *archive*).
- Enhance linker so that it tries to resolve unresolved external references by looking for the symbols in one or more archives.
- If an archive member file resolves reference, link it into the executable.

- Static libraries

17

Creating Static Libraries

```

    graph TD
      atoi_c[atoi.c] --> T1[Translator]
      printf_c[printf.c] --> T2[Translator]
      random_c[random.c] --> T3[Translator]
      T1 --> atoi_o[atoi.o]
      T2 --> printf_o[printf.o]
      T3 --> random_o[random.o]
      atoi_o --> ar[Archiver ar]
      printf_o --> ar
      random_o --> ar
      ar --> libc_a[libc.a C standard library]
  
```

```

    unix> ar rs libc.a \
    atoi.o printf.o ... random.o
  
```

- Archiver allows incremental updates
- Recompile function that changes and replace .o file in archive.

18

Commonly Used Libraries

libc.a (the C standard library)

- 8 MB archive of 1392 object files.
- I/O, memory allocation, signal handling, string handling, data and time, random numbers, integer math

libm.a (the C math library)

- 1 MB archive of 401 object files.
- floating point math (sin, cos, tan, log, exp, sqrt, ...)

```

% ar -t /usr/lib/libc.a | sort
...
fork.o
...
fprintf.o
fpu_control.o
fputc.o
freopen.o
fscanf.o
fseek.o
fstab.o
...

% ar -t /usr/lib/libm.a | sort
...
e_acos.o
e_acosf.o
e_acosh.o
e_acoshf.o
e_acoshl.o
e_acosl.o
e_asin.o
e_asinf.o
e_asinl.o
...
  
```

19

Linking with Static Libraries

```

    graph TD
      main2_c[main2.c] --> T[Translators cpp, cc1, as]
      vector_h[vector.h] --> T
      T --> main2_o[main2.o Relocatable object files]
      addvec_o[addvec.o] --> ar[Archiver ar]
      multivec_o[multivec.o] --> ar
      ar --> libvector_a[libvector.a]
      libvector_a --> ld[Linker ld]
      main2_o --> ld
      libc_a[libc.a Static libraries] --> ld
      ld --> p2[p2 Fully linked executable object file]
  
```

20

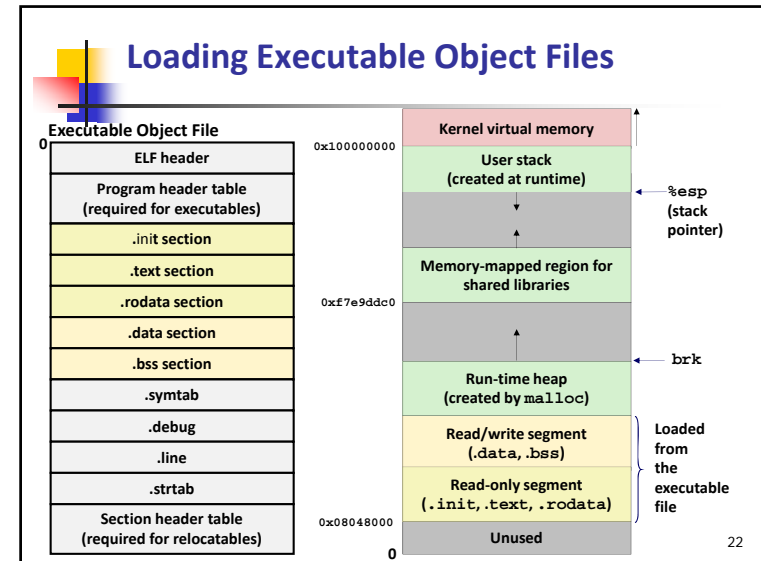
Using Static Libraries

- Linker's algorithm for resolving external references:
 - Scan `.o` files and `.a` files in the command line order.
 - During the scan, keep a list of the current unresolved references.
 - As each new `.o` or `.a` file, *obj*, is encountered, try to resolve each unresolved reference in the list against the symbols defined in *obj*.
 - If any entries in the unresolved list at end of scan, then error.
- Problem:
 - Command line order matters!
 - Moral: put libraries at the end of the command line.

```

unix> gcc -L. libtest.o -lmine
unix> gcc -L. -lmine libtest.o
libtest.o: In function `main':
libtest.o(.text+0x4): undefined reference to `libfun'
    
```

21



Dynamic Link Libraries

- Static libraries have the following disadvantages:
 - Duplication (every function need std libc)
 - In stored executables
 - In the running executables
 - Minor bug fixes of system libraries require each application to explicitly relink
- Modern solution: Dynamic link libraries
 - Object files that contain code and data that are loaded and linked into an application *dynamically*, at either *load-time* or *run-time*

23

Dynamic Link Libraries (cont.)

- Dynamic linking can occur when executable is first loaded and run (load-time linking).
 - Common case for Linux, handled automatically by the dynamic linker (`ld-linux.so`).
 - Standard C library (`libc.so`) usually dynamically linked.
- Dynamic linking can also occur after program has begun (run-time linking).
 - In Linux, this is done by calls to the `dlopen()` interface.
 - Distributing software.
 - High-performance web servers.
 - Runtime library interpositioning.
- Dynamic link library routines can be shared by multiple processes.
 - More on this when we learn about virtual memory

24

Position-Independent Code

- Dynamic link libraries need position-independent code

25

Case Study: Library Interpositioning

- Library interpositioning : powerful linking technique that allows programmers to intercept calls to arbitrary functions
- Interpositioning can occur at:
 - Compile time: When the source code is compiled
 - Link time: When the relocatable object files are statically linked to form an executable object file
 - Load/run time: When an executable object file is loaded into memory, dynamically linked, and then executed.

26

Some Interpositioning Applications

- Security
 - Confinement (sandboxing)
 - Interpose calls to libc functions.
 - Behind the scenes encryption
 - Automatically encrypt otherwise unencrypted network connections.
- Monitoring and Profiling
 - Count number of calls to functions
 - Characterize call sites and arguments to functions
 - Malloc tracing
 - Detecting memory leaks

27

Example program

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

int main()
{
    free(malloc(10));
    printf("hello, world\n");
    exit(0);
}
```

hello.c

- Goal: trace the addresses and sizes of the allocated and freed blocks, without modifying the source code.
- Three solutions: interpose on the lib malloc and free functions at compile time, link time, and load/run time.

28

Compile-time Interpositioning

malloc.h

```
#define malloc(size) mymalloc(size, __FILE__, __LINE__ )
#define free(ptr) myfree(ptr, __FILE__, __LINE__ )

void *mymalloc(size_t size, char *file, int line);
void myfree(void *ptr, char *file, int line);
```

- In mymalloc()/myfree(), do whatever interpositioning work before calling malloc()/free()

29

Link-time Interpositioning

```
gcc -O2 -Wall -Wl,--wrap,malloc -Wl,--wrap,free \
-o hello1 hello.c mymalloc.o
```

- The “-Wl” flag passes argument to linker
- Telling linker “--wrap,malloc” tells it to resolve references in a special way:
 - Refs to malloc should be resolved as __wrap_malloc
 - Refs to __real_malloc should be resolved as malloc

30

Load/Run-time Interpositioning

```
linux> make hellor
gcc -O2 -Wall -DRUNTIME -shared -fPIC -o mymalloc.so mymalloc.c
gcc -O2 -Wall -o hellor hello.c
linux> make runr
(LD_PRELOAD="/usr/lib64/libdl.so ./mymalloc.so" ./hellor)
```

- The LD_PRELOAD environment variable tells the dynamic linker to resolve unresolved refs (e.g., to malloc) by looking in libdl.so and mymalloc.so first.
 - libdl.so necessary to resolve references to the dlopen functions.

31

Disclaimer

These slides were adapted from the CMU course slides provided along with the textbook of “Computer Systems: A programmer’s Perspective” by Bryant and O’Hallaron. The slides are intended for the sole purpose of teaching the computer organization course at the University of Rochester.

32