Due: See website for due date.

What to submit: Upload a tar archive that contains a text file answers.txt with your answers for the questions not requiring code, as well as individual files for those that do, as listed below.

This exercise is intended to reinforce the content of the lectures related to linking using small examples.

As some answers are specific to our current environment, you must again do this exercise on our rlogin cluster.

Our verification system will reject your submission if any of the required files are not in your submission. If you want to submit for a partial credit, you still need to include all the above files.

1. Linking Cush

In this part of the exercise, you are asked to make small changes to the cush starter code to induce linker errors and changes to the executable. To that end, you should clone a fresh copy of the starter code with

```
git clone git@git.cs.vt.edu:cs3214-staff/cs3214-cush.git
cd cs3214-cush
(cd posix_spawn; make)
```

The 3 parts are independent and you should undo any changes you made for one part before continuing to the next.

1. After changing 2 lines ¹ in 2 files, cush rebuilds but the build fails with:²

```
cc -Wall -Werror -Wmissing-prototypes -I../posix_spawn -g -O2
    -fsanitize=undefined -o cush -L../posix_spawn
    cush.o shell-grammar.o list.o shell-ast.o termstate_management.o
    utils.o signal_support.o -lspawn -ll -lreadline
/usr/bin/ld:
    termstate_management.o:/....cush/src/termstate_management.h:6:
    multiple definition of `terminal_fd';
    cush.o:/...cush/src/termstate_management.h:6: first defined here
    collect2: error: ld returned 1 exit status
    make: *** [Makefile:27: cush] Error 1
```

What lines were changed? Provide the output as a patch (which you can produce via git diff.)

2. After a single line change from the original code, the project builds but the following nm command shows this:

```
$ nm cush | grep shell_pgrp
0000000000420080 B shell_pgrp
```

What line was changed and how? Provide the output as a patch.

3. After changing a single keyword on a single line, the project build fails with

```
cc -Wall -Werror -Wmissing-prototypes -I../posix_spawn -g -02
    -fsanitize=undefined -o cush -L../posix_spawn cush.o
    shell-grammar.o list.o shell-ast.o termstate_management.o utils.o
    signal_support.o -lspawn -ll -lreadline
/usr/bin/ld: termstate_management.o: in function `termstate_save':
/.../cs3214-cush/src/termstate_management.c:48:
```

¹This is not counting any empty/whitespace lines. Also, in your reproduction, the line numbers do not need to match.

²I introduced newlines for readability.

```
undefined reference to `saved_tty_state'
/usr/bin/ld: termstate_management.o: in function
    `termstate_give_terminal_back_to_shell':
/.../cs3214-cush/src/termstate_management.c:107:
    undefined reference to `saved_tty_state'
/usr/bin/ld: termstate_management.o: in function `termstate_save':
/.../cs3214-cush/src/termstate_management.c:48:
    undefined reference to `saved_tty_state'
collect2: error: ld returned 1 exit status
make: *** [Makefile:27: cush] Error 1
```

What line was changed and how? Provide the output as a patch.

Remember to revert to the original starter code for each part!

2. Baking Pie

From past courses (CS 2505, CS 2506) you are familiar with threats that can affect vulnerable applications that contain buffer overflows. Some of the exploits that targeted such applications made assumptions about the way in which they were built and run.

One particular assumption relates to the virtual addresses at which functions or data can be found, which traditionally has been static. Recently, some OS have instead adopted position-independent executables as a default where the locations of functions and global variables is randomized from run to run.

For instance, a hypothetical program pie.c could be built either as a regular executable like so:

gcc pie.c -o no.pie

Or as a position-independent example like so:

gcc -fPIE -pie pie.c -o pie

When running the ./no.pie version, the output will always be:

\$./no.pie 0x404000

But when running the ./pie version, the output will vary randomly, perhaps like so:

```
$ ./pie
0x5562f12a9000
$ ./pie
0x5564fe021000
$ ./pie
0x564deec67000
```

Write pie.c!

3. Link Time Optimization

Traditional separate compilation and linking has an important drawback: since the intermediate representation created by the compiler is no longer available at link time, potential interprocedural optimizations cannot be performed. For instance, the linker cannot inline functions or replace calls to functions that produce constant results with their values.

Link Time Optimization (LTO) overcomes this drawback by preserving the compiler's intermediate representation and passing it along to the linker which can then perform whole-program optimization across modules. Languages such as Rust use LTO to be able to perform optimizations across the different source files that are part of a crate.

In this part of the exercise, you will be looking at how LTO works in a current compiler (gcc 11.4.1).

Create or copy the following files lto.h, lto1.c and lto2.c:

```
typedef struct {
    double x, y;
} Solution;
Solution cramer(double a, double b, double c,
    double d, double e, double f);
```

```
#include <stdio.h>
#include "lto.h"
int
main()
{
    Solution s = cramer(1, 2, 4, -2, 1, -3);
    return (int)(s.x + s.y);
}
```

Compile and build the two files using the following commands:

gcc -03 -flto -c lto1.c lto2.c
gcc -03 -flto lto1.o lto2.o -o lto

Then answer the following questions:

- 1. Use objdump -d to find the code for the main() in the final lto executable. Copy and paste the body of main (the disassembled machine code)!
- 2. Now compile these programs without LTO like so:

gcc -03 lto1.c lto2.c -o nolto

Use objdump -d nolto to look at the main function, and reproduce the assembly code here.

Explain in your own words what the compiler and linker did when LTO was enabled and how this was possible using LTO but not when LTO was not being used.

3. What is the output of

```
./lto; echo $?
and why?
```

4. Building Redis

Redis is a popular in-memory data store that is widely used in industry as a cache or database. It was created by Salvatore Sanfilippo, better known as antirez.

In this part of the exercise you will look at how redis is compiled and built in order to observe how compilers and linkers are used in a larger software project.

Your answers will be specific to the version of the GCC tool chain installed on rlogin this semester.

- 1. Download and extract the source code of Redis 7.2.4.
- 2. Change into the directory into which you've extracted the source code and build it using the command

make -j V=1

The make program will run a number of commands to configure, compile, and link multiple executables that are part of Redis.

- 3. Which of the commands the make program invokes takes the longest amount of time (about 44 seconds currently) and why? Copy and paste the command here.
- 4. One of the executables the make program built is the redis-server executable in the src directory. When this executable is linked, 5 libraries are statically linked with it. List them.

- 5. How much space does the redis-server executable take up on disk? (Use ls -l to find out.)
- 6. The command size (without any arguments) gives you an estimate of how much memory is needed if the executable were fully loaded into memory, broken down by text/code, data, and bss. Run size on redis-server. Roughly what fraction of the executable is taken up by code?
- 7. The command strip strips an executable of those parts that are not necessary to run it. Run strip on the redis-server executable and measure its size with ls -l. Then run size on the stripped executable. Roughly what fraction of the stripped executable is taken up by code?
- 8. One of the best practices we had discussed was to ensure that functions that are used only within a single C file are declared with the static keyword. This way they will not leak into the global namespace and encapsulation is preserved.

List 3 functions in the Redis code that could have been made static but weren't (yet). **Hint:**

You can obtain .o files with symbol tables for analysis by recompiling without LTO. To that end, do

```
make clean
rm src/.make-settings
make -j OPTIMIZATION=-00
```

then you should be able to examine the symbol tables of the object modules in the source directory with nm.

List the names of these 3 functions as they appear in the symbol table. (Think about how you can test your answers!)

9. Last, but not least, do not forget to remove the source directory from your rlogin file space. It takes up about 200MB of space.