

COMP 321: Introduction to Computer Systems

Assignment 3: Linking

Assigned: 1/31/19, Due: 2/14/19, 11:55 PM

Important: This assignment may be done individually or in pairs. Be sure to carefully read the course policies for assignments (including the honor code policy) on the assignments page of the course web site:

<http://www.clear.rice.edu/comp321/html/assignments.html>

Overview

You will answer three questions about x86-64 linking and complete the implementation of the `readjcf` program, which reads the linkage information from a Java Class File (JCF). The goals of this assignment are as follows:

- Improve your understanding of the x86-64 linking process.
- Gain familiarity with `objdump` and `readelf`.
- Gain an understanding of the similarities and differences between x86-64 ELF files and Java virtual machine class files.
- Improve your understanding of data representation and I/O in C.

Problems

Write solutions to each of the following problems. Each problem indicates its point value.

1. Textbook Problem 7.6¹ (8 points)

To generate identical object files `m.o` and `swap.o` to those assumed by the textbook, use the commands `gcc -g -Wall -Wextra -Werror -c m.c` and `gcc -g -Wall -Wextra -Werror -c swap.c`.

2. Textbook Problem 7.9 (8 points)

Although the textbook says that the program will print the value `0x48` when you run it on an x86-64 Linux system, when you compile and run the program on CLEAR, which is an x86-64 Linux system, it nonetheless prints a different value. Specifically, when you compile the program using the command `clang -Wall -Wextra -Werror bar6.c foo6.c -o program` and run `program`, it prints the value `0x55`. In short, the value printed by the program is also dependent on the compiler used to create the executable, and not just the architecture and operating system. This difference in the values printed should not affect your answer to the question.

¹The problems in the third edition of the textbook are different from those in the second edition.

3. Textbook Problem 7.12 (8 points)

For both parts of this question, you should also explain how you arrived at your answer.

To generate an identical object file `m.o` to that assumed by the textbook, use the command `gcc -g -Wall -Wextra -Werror -c m.c`.

4. Programming Problem, `readjcf` (76 points)

The compilation and execution of Java programs has both similarities and differences with respect to C programs. First, the `javac` compiler does not output x86-64 ELF files. It outputs files that are instead in the Java Class File (JCF) format. Second, the machine code within these files is not x86-64 machine code. It is machine code for a hypothetical computer called the Java Virtual Machine (JVM). Essentially, on CLEAR, this machine code is executed by a very fast JVM simulator². Nonetheless, there are many conceptual similarities between the JCF and ELF file formats. In other words, to link together a Java program consisting of many classes, the JCF format must support many of the same features as the ELF file format. We learn about these similarities and differences through the implementation of the `readjcf` program.

First, we will explain the entire JCF format, then we will give a specification of the `readjcf` program that you will write. Although we need to describe the entire JCF format, the `constant_pool`, `fields`, and `methods` arrays are of particular interest for implementing the `readjcf` program.

Java Class File (JCF) Format:

The Java Class File (JCF) format is defined by the following pseudo-C structure. Note that in this definition, `u1`, `u2`, and `u4` refer to an unsigned one-, two-, and four-byte quantity, respectively, and that all multibyte data is stored in big-endian format:³

```
ClassFile {
    u4 magic;
    u2 minor_version;
    u2 major_version;
    u2 constant_pool_count;
    cp_info constant_pool[constant_pool_count-1];
    u2 access_flags;
    u2 this_class;
    u2 super_class;
    u2 interfaces_count;
    u2 interfaces[interfaces_count];
    u2 fields_count;
    field_info fields[fields_count];
    u2 methods_count;
    method_info methods[methods_count];
    u2 attributes_count;
    attribute_info attributes[attributes_count];
}
```

²This simulator is itself written in C++ and compiled to an x86-64 ELF executable.

³The full specification for the JCF format is located at <https://docs.oracle.com/javase/specs/jvms/se7/html/jvms-4.html>. This document presents the subset of the specification that is relevant to the assignment.

This structure is pseudo-C because the `constant_pool`, `interfaces`, `fields`, `methods`, and `attributes` arrays are variable sized, which is not allowed in C structures.

The magic field of a valid JCF must equal `0xCAFEBABE`. The `minor_version` and `major_version` fields contain version information about the class.

The `constant_pool` section of the JCF is most similar to a combination of the `.rodata` section, symbol table, and relocation information of an ELF file. The `constant_pool` contains all of the class' constants, including strings, numbers, references, and the names and types of classes, interfaces, fields, and methods. The `constant_pool` array contains one less entry than the count because the zeroth entry of the array is omitted. In other words, the first constant in the pool has index one.

The `access_flags` field is a bit vector that specifies access permissions for the class. The `this_class` and `super_class` fields are indices into the constant pool that reference the names of the current class and the super class. The `interfaces` section contains an array of indices into the constant pool for all of the interfaces that the class implements. The `fields` and `methods` arrays contain information about every single field and method in the class. Lastly, the `attributes` array contains specific attributes about the class.

There are fourteen different types of constants in the constant pool: `CONSTANT_Class`, `CONSTANT_Fieldref`, `CONSTANT_Methodref`, `CONSTANT_InterfaceMethodref`, `CONSTANT_String`, `CONSTANT_Integer`, `CONSTANT_Float`, `CONSTANT_Long`, `CONSTANT_Double`, `CONSTANT_NameAndType`, `CONSTANT_Utf8`, `CONSTANT_MethodHandle`, `CONSTANT_MethodType`, and `CONSTANT_InvokeDynamic`. The `cp_info` structures for all of these different constants all share the same general form:

```
cp_info {
    u1 tag;
    u1 info[];
}
```

The `tag` field indicates which type of constant the structure represents and, subsequently, the exact size of the `info` array.

The `CONSTANT_Class` constant contains the name of a class and is represented with the following structure:

```
CONSTANT_Class_info {
    u1 tag;
    u2 name_index;
}
```

The `name_index` field is a valid index into the `constant_pool` array. The element at that index corresponds to a `CONSTANT_Utf8`.

The `CONSTANT_Fieldref`, `CONSTANT_Methodref`, and `CONSTANT_InterfaceMethodref` constants indicate references made in the class to fields, methods, and interface methods, respectively. The following structure is used to represent all of these constants:

```
CONSTANT_{Field, Method, InterfaceMethod}ref_info {
    u1 tag;
```

```
    u2 class_index;
    u2 name_and_type_index;
}
```

The `class_index` and `name_and_type_index` fields indicate the class, name, and type of the field, method, or interface method being referenced. The fields are indices into the constant pool that correspond to `CONSTANT_Class` and `CONSTANT_NameAndType` constants, respectively.

The `CONSTANT_String`, `CONSTANT_Integer`, `CONSTANT_Float`, `CONSTANT_Long`, `CONSTANT_Double`, `CONSTANT_MethodHandle`, `CONSTANT_MethodType`, and `CONSTANT_InvokeDynamic` constants have the following structure definitions:

```
CONSTANT_String_info {
    u1 tag;
    u2 string_index;
}

CONSTANT_{Integer, Float}_info {
    u1 tag;
    u4 bytes;
}

CONSTANT_{Long, Double}_info {
    u1 tag;
    u4 high_bytes;
    u4 low_bytes;
}

CONSTANT_MethodHandle_info {
    u1 tag;
    u1 reference_kind;
    u2 reference_index;
}

CONSTANT_MethodType_info {
    u1 tag;
    u2 descriptor_index;
}

CONSTANT_InvokeDynamic_info {
    u1 tag;
    u2 bootstrap_method_attr_info;
    u2 name_and_type_index;
}
```

Although you must necessarily pay attention to the sizes of the various data types in the above structures, the actual meaning of some of these structures' data members can remain opaque to you throughout the assignment. For instance, in the case of `CONSTANT_InvokeDynamic`, the

`bootstrap_method_attr_info` is an index into the attributes table of the JCF. However, since you are not required to do any special processing of JCF attributes in your program, you may ignore what this actually refers to.

The JCF format specifies that the `CONSTANT_Long` and `CONSTANT_Double` constants occupy two entries in the constant pool. The specification says that, “If a `CONSTANT_Long_info` or `CONSTANT_Double_info` structure is the item in the `constant_pool` table at index n , then the next usable item in the pool is located at index $n + 2$. The `constant_pool` index $n + 1$ must be valid but is considered unusable.”

The `CONSTANT_NameAndType` constant represents the name and type of a field or method. It is defined by the following structure:

```
CONSTANT_NameAndType_info {
    u1 tag;
    u2 name_index;
    u2 descriptor_index;
}
```

The `name_index` and `descriptor_index` fields indicate the name and description string of the field, method, or interface method being referenced. The fields are indices into the constant pool that correspond to `CONSTANT_Utf8` constants.

The `CONSTANT_Utf8` constant represents a UTF-8 string that is not NUL-terminated. UTF-8 is a variable-width character encoding that can represent any character in the Unicode character set, which includes over 110,000 characters from 100 scripts. UTF-8 is the most common character encoding used on the web, in part because ASCII encoding is a subset of UTF-8 character encoding. The `CONSTANT_Utf8` constant is defined by the following structure:

```
CONSTANT_Utf8_info {
    u1 tag;
    u2 length;
    u1 bytes[length];
}
```

The `field_info` and `method_info` structures both have the same definition:

```
{field, method}_info {
    u2 access_flags;
    u2 name_index;
    u2 descriptor_index;
    u2 attributes_count;
    attribute_info attributes[attributes_count];
}
```

The `access_flags` field is a bit vector that specifies access information about the field or method, including whether the method is public, protected, or private. The only flag that is necessary for this assignment is the `ACC_PUBLIC` flag. The `name_index` and `descriptor_index` are indices into the constant pool to a `CONSTANT_Utf8` that contains the name and description string of the field or method, respectively. Finally, the structure contains an array of attributes.

The `attribute_info` structure is defined by the following pseudo-C structure:

```

attribute_info {
    u2 attribute_name_index;
    u4 attribute_length;
    u1 info[attribute_length];
}

```

There are many more specific types of attributes defined by the Java Class File format, but they all follow the format of the previous structure. For this assignment, we will be ignoring all attribute information, so we do not need to delve into the specific types of attributes.

readjcf Specification:

The `readjcf` program will read a single Java Class File and print out the class' dependencies and exports, as requested by flags. Additionally, the program will perform some basic verification of the JCF. The program accepts three command line flags: “-d”, “-e”, and “-v”. The flags “-d” and “-e” indicate whether or not the program should print the class file's dependencies and exports, respectively. The flag “-v” enables extra print statements for debugging. If no flags are specified, then no dependencies or exports should be printed, but the class file should still be read and verified.

We provide a main routine that reads and parses the command line arguments appropriately. The usage of the program is as follows:

```
readjcf [-d] [-e] [-v] <input file>
```

We will not use the `-v` flag when testing your program. Moreover, we are not specifying what the output from the extra print statements should be. That is up to you.

For the purposes of the `readjcf` program, we define a dependency as any reference made in the class, including self-references. These references are made by the `CONSTANT_Fieldref`, `CONSTANT_Methodref`, and `CONSTANT_InterfaceMethodref` constants in the constant pool. We define an export as any field or method in the class that has the access flag `ACC_PUBLIC` set.

The dependencies and exports should be printed in the order that they are encountered in the JCF. Due to the format of the JCF, this means that all of the dependencies will be printed before the exports. The output format for printing the dependencies and exports in the class file is exactly as follows:

```

Dependency - <class name>.<name> <descriptor>
Export - <name> <descriptor>

```

For a concrete example, the following is output from the reference solution:

```

Dependency - java/lang/Object.equals (Ljava/lang/Object;)Z
Dependency - java/lang/Object.hashCode ()I
Export - <init> (D)V
Export - toString ()Ljava/lang/String;
Export - equals (Ljava/lang/Object;)Z
Export - hashCode ()I

```

The program must also perform a simple form of verification of the JCF. Your `readjcf` program must verify the following:

- The magic number is `0xCAFEBABE`.
- The class file is not truncated and does not contain extra bytes.
- All indices into the constant pool that are followed while printing the dependencies and exports are valid and the constant at that index is of the expected type (according to its tag).

The verification should be performed as the file is processed. If at any time the file fails verification, the program should use the `readjcf_error` function to print an error message to `stderr`, then immediately quit.

We are providing you with a partially implemented program for reading these files. The provided program includes a main function that opens the input file and calls the rest of the processing functions in the correct order for the JCF format. You need to implement the following functions:

```
static int      print_jcf_constant(struct jcf_state *jcf,
                                  uint16_t index, uint8_t expected_tag);
static int      process_jcf_header(struct jcf_state *jcf);
static int      process_jcf_constant_pool(struct jcf_state *jcf);
static void     destroy_jcf_constant_pool(struct jcf_constant_pool *pool);
static int      process_jcf_body(struct jcf_state *jcf);
static int      process_jcf_interfaces(struct jcf_state *jcf);
static int      process_jcf_attributes(struct jcf_state *jcf);
```

The `print_jcf_constant` function prints a constant from the constant pool to `stdout`. `process_jcf_header` reads the header (the `magic`, `minor_version`, and `major_version` fields) from the JCF. `process_jcf_constant_pool` reads and stores the constant pool, then prints out all of the class' dependencies, if requested. `destroy_jcf_constant_pool` deallocates all of the memory allocated to store the constant pool by `process_jcf_constant_pool`. `process_jcf_body` reads the body (the `access_flags`, `this_class`, and `super_class` fields) from the JCF. `process_jcf_interfaces` reads the interfaces count and the array of interfaces. `process_jcf_attributes` reads an attributes count, followed by an array of attributes.

Notes

- All of the multibyte data in the JCF is stored in big-endian format, which means that the data must be byteswapped before it may be accessed on CLEAR. To byteswap 16-bit and 32-bit integers use the functions `ntohs` and `ntohl`, respectively.
- The `jcf_cp_utf8_info` structure contains a zero byte array `bytes` as the last field of the structure. The compiler and the `sizeof()` function do not include space for this array in their representation of the structure because it is zero bytes long. This causes the `bytes` field to point to the byte immediately after the structure. If you only allocate `sizeof(struct jcf_cp_utf8_info)` bytes, then the `bytes` field points to invalid memory. If you instead allocate `(sizeof(struct jcf_cp_utf8_info) + numbytes)` bytes, you will have allocated memory for the `bytes` array. This means that the `bytes` array now points to an array that is `numbytes` long.

- The UTF-8 strings in the constant pool are not NUL-terminated. The standard I/O library functions will work properly on NUL-terminated UTF-8 strings. You will need to NUL-terminate the strings yourself.
- The UTF-8 strings in the constant pool can be zero bytes long. The `fread` function will always return 0, an indication of error, when it is asked to read an element whose `size` equals 0.
- The `JCF_CONSTANT_Long` and `JCF_CONSTANT_Double` constants contain 8 bytes worth of data after the tag and occupy two spots in the constant pool. This means that, for example, you encounter one of these constants at index 1, then the next constant is located at index 3.
- The `access_flags` fields are bit vectors. Specific flags may be tested for by performing the boolean and of the field and the mask. For example, `access_flags & ACC_PUBLIC` will be true iff the `ACC_PUBLIC` flag is set.
- We provide you with the implementation of `process_jcf_fields` and `process_jcf_methods`. These functions provide a good example as to how to implement the rest of the program.
- The `fread` function can read structures from a file. This is the preferred way to read data in your program.

Getting Started

If you choose to work in a pair, only one of you should follow the below instructions for creating a private repository. In a later section, we will explain how to share this repository with the partner who didn't create it.

Please visit the web page at

```
https://classroom.github.com/a/bUnSsMPn
```

This page should be titled "RICE-COMP321-SPRING19". Moreover, it should say, "Accept the Linking assignment", and have a button labeled "Accept the invitation"⁴. Please accept the invitation.

Upon accepting the invitation, you will be redirected to another web page. This page confirms that you have accepted the assignment and provides a link to your personal repository for the assignment. Click the link to go to your personal repository.

The web page for your personal repository has a button labeled "Clone or download". Click this button. You should now see text field with a URL. Copy or remember this URL.

Login to the CLEAR system if you have not already done so. Type the following:

```
git clone [Copy the URL for your repo here]
```

You will be prompted for your `github` username and password.

Once the clone operation is complete, you will have a directory named

```
linking-[YOUR github ID]
```

⁴You may have to login to `github` to see the invitation. If so, you will be prompted for your username and password at the top of the page.

Most of the files provided for this assignment are stored in subdirectories named `7.6`, `7.9`, and `readjcf`, the first two correspond to the textbook problem number. You may use the files in `7.6` and `7.9` in any way to answer the three textbook problems, including recompiling them, compiling them with other files, using `readelf` or `objdump`, or any other way you see fit. Carefully note that the compiler options may be different for each problem and are given above. The `readjcf` subdirectory contains an executable named `readjcf_ref`. This is the reference solution for `readjcf`. Run this program to resolve any questions you have about how your program should behave. *Your program should emit output that is byte-for-byte identical to the reference solution for any combination of the flags “-d” and “-e”.* You are not responsible for matching the exact output of the reference solution with the flag “-v”. That said, attempting to match the output of the reference solution with the flag “-v” is a good way to develop and test your program in small steps.

Turning in Your Assignment

If you choose to work in a pair, your code and writeup must include both partners’ names and NetIDs.

To turn in your code and the writeup of your problem solutions, you *must* use `git push` to copy your work to the `github` remote repository. We will *only* look at the last version that you pushed before the deadline, taking into account your use of slip days. If you are working with a partner, it doesn’t matter which partner performs the final push. As a precaution against accidental loss of your code or writeup, we encourage you to push periodically. Please note, the *only* files that you need to turn in are `readjcf.c` and `writeup.txt`. In other words, these are the only two files on which you should ever perform `git add`. The `writeup.txt` file should contain both the answers to problems `7.6`, `7.9`, and `7.12` as well as the writeup of the `readjcf` program, which includes the normal design and testing strategy sections.

For grading your submission, we will use the `Makefile` that was originally provided to compile your code. Therefore, your code should not rely on any modifications to the `Makefile` for correct compilation.

As a sanity check, you should use your web browser to visit your assignment repo. Make sure that what you see in the browser is consistent what you think you have pushed.

How to add your partner as a collaborator on your repository

To work with a partner, the assignment acceptor from the section “Getting Started” must add the partner as a collaborator by following these steps:

1. Navigate your browser to the assignment repository page, using the URL you received when you initially accepted the assignment.
2. Near the top of the `github` repository page, there is a tab labeled “Settings”. There is a small gear icon directly to the left of the “Settings” title. Click on the “Settings” tab.
3. Once on the “Settings” page, you should find a group of operations labeled “Options” (“Options” is in boldface.) The first such option is labeled “Collaborators and teams”. Click on the “Collaborators and teams” operation.
4. The “Options” pane will still be visible on the left. On the right of the “Options” pane, you will find two panes stacked vertically. The top right pane is labeled “Teams”. The lower right pane is labeled “Collaborators”. At the bottom of the “Collaborators” pane, there is a text box paired with a button labeled “Add collaborator”. The “Add collaborator” button will be slightly greyed out.
5. Click inside the “Collaborators” text box. You should see a text cursor appear inside the text box.

6. Enter your partner's github ID in the text box. Presumably, github will recognize the ID. When recognized, the github ID and user's actual name appear just below the text box. At this point, you should click the "Add collaborator" button.
7. You will be informed that you are "Awaiting a response" from your collaborator.
8. At this point, your partner should check email for the invitation. When your partner receives the invitation, please click the provided acceptance link. [The acceptance step requires that the partner be logged into github].
9. A short time later, the original repo acceptor will see confirmation that the invitation has been accepted.

Congratulations! You now have a shared repository for your assignment. Both partners will be able to clone, push, and pull from the assignment repository. (The last of which, i.e., pull, we will explain in lab.)