

CSE 410: Systems Programming

Recitation 2: Introduction to Make

Introduction

This exercise will help you become familiar with Make. Make is a tool for managing the build of complex software projects, although it can be used for many other things as well. At its heart, it is a *dependency solver* that creates a *dependency graph* with software sources at the “leaves” and processed or built components at the nodes. Make then uses this information to determine whether the dependencies for a node in the graph have been met, and if so, executes user-defined instructions to create that node.

Make uses filesystem timestamps to determine that portions of the software are “stale” with respect to their sources, and then executes user-supplied commands to bring them up-to-date. In this manner a very large application can be re-compiled after minor changes very quickly, by re-compiling only those portions of the application that are affected by the changes.

As you will not be building particularly large applications in this course, the primary benefit that you will receive from using Make is automation. It should also reduce errors in builds by ensuring repeatability of the build process.

We will be using GNU Make, although most of the syntax and features that you will learn are the same as other popular Make implementations, such as the original AT&T Make and BSD Make.

1 Getting Started

Follow the GitHub Classroom link that has been provided to you in Piazza to create the starter repository for this assignment. Check it out following the same steps you used for the first lab of the semester. It will include a README, Makefile, and two C source files. You should not modify the C sources for this assignment, but you may wish to inspect them.

1.1 A Brief Introduction to Make

Your TAs will provide you with a more thorough introduction to Make. However, you may find this section useful for reference.

Make is invoked via the `make` command on most UNIX systems. It will read a file called `Makefile` from the current directory (unless otherwise specified), and then execute the first *rule* it finds in that file. Make rules are described below, and will be covered in more detail by your TAs. You can also provide a specific rule to be executed on the command line, by simply typing it after `make` on the command line; for example, `make hello` will instruct Make to consult the Makefile for a rule that will build the file `hello`, and then execute it.

Note that the order in which rules appear in the Makefile is not important *except for the first rule*, which is the rule that will be executed if no rule is specified on the command line. There is an idiomatic broad ordering of rules that many or most developers will follow, however, with the “biggest” and most important rules toward the top of the file, followed by more specific rules, followed by “phony” rules, or rules that don’t actually build anything. (There is an example phony rule named `clean` in the handout code that removes the built files. There are additional subtleties to phony rules, and this one is incompletely specified for simplicity, but it will work.)

For this lab you will need to use Make *variables*, create *dependency rules*, and assign *recipes* to those rules. The following Makefile demonstrates these things, and will be described below:

```
VARIABLE := value

var:
    echo VARIABLE is $(VARIABLE)

chain: var
    echo The var rule should have run before this.
```

The first line in this file sets the variable named `VARIABLE` to the value `value`. Note the `:=` syntax used for this assignment. Make will also accept `=`, although the meaning is slightly different. Once this variable is created, it can be referenced with `$(VARIABLE)`. You may also see `${VARIABLE}` used (with curly brackets instead of parenthesis), which has the same meaning.

The next pair of lines defines a *rule* for the *target* `var`, with a *recipe* to create it. If `make` is invoked with no arguments, this rule (being the first in the file) is the rule that will be run by default.

The syntax for a rule is as follows:

```
TARGET: DEPENDENCIES
      RECIPE
```

The syntax for Makefiles is picky; in particular, the whitespace before the recipe **must be** an ASCII Tab character. If you use spaces, Make will (if you are lucky) emit an error like one of these:

```
example-Makefile:4: *** missing separator (did you mean TAB instead of 8 spaces?).
      Stop.
```

```
example-Makefile:4: *** missing separator.  Stop.
```

It may, however, emit an even more cryptic message.

This example recipe declares that building the target named `TARGET` requires the dependencies `DEPENDENCIES`, and that producing `TARGET` from `DEPENDENCIES` is accomplished by running the UNIX commands described by `RECIPE`. The dependencies specified in `DEPENDENCIES` may be either other Make targets or filenames, separated by whitespace. Make targets with no associated file will always be considered out-of-date and run *once* each time Make is invoked. If they are filenames, `RECIPE` will be run only if the file `TARGET` is older than each of the files described in `DEPENDENCIES`. In the above example Makefile, note that neither `var` nor `chain` creates a file; therefore, both of these rules will be run any time they are a dependency of any requested target.

Each *line* of `RECIPE` is run as a UNIX shell script. Any UNIX command or combination of commands may be included. If any command in a recipe fails, the recipe will fail and Make will stop with an error. This will happen if, for example, your recipe is compiling C source code and you have an error in your sources. Multi-line UNIX scripts must be written as a single logical line by using semicolons to separate commands and escaping actual newlines with a backslash.

The example Makefile rule for `chain` therefore declares that the rule `var` must be run before `chain`, and that the target described by `chain` is satisfied by echoing a sentence to the terminal. Try creating a Makefile like the example above and then running `make`, `make var`, and `make chain` to see how this works.

1.2 The Make Manual

The manual for GNU Make is available online, but *does not use the `man` command that we discussed in class*. Instead, it uses GNU Info, which works somewhat differently. You can consult it by running `info make` at the prompt. You can press `H` while in the Info viewer for a summary of usage, or see `info info` for more detail. Some useful keys to know are `/`, which initiates a search, `u`, which moves to “up” to the next higher level of the document structure, `spacebar` to page down, and `q` to quit. The arrow keys will work as expected, and the enter key will jump to a hyperlink under the cursor.

2 Requirements

You must create a Makefile to compile the sources in the handout repository into an executable named `hello`, which is one of the more complicated “Hello, World” programs ever devised. The handout repository includes a starter Makefile, but it is missing some variable definitions and rules. You must:

- Define the variable `CFLAGS`.

It should hold the value `-g -Wall -Werror`. This variable is traditionally used to hold options passed to the C compiler.

- Fix the rule for the target `main.o`.

This rule must pass the extra option `-DMAIN` to the C compiler.

- Create a rule to compile `dohello.o` from `dohello.c` using `gcc`.

This rule must include `$(CFLAGS)` on the command line, and must define the preprocessor symbol `DOHELLO` to have the value `do_hello` using the option `-DDOHELLO=do_hello` passed to the C compiler.

Once these requirements are complete, typing `make` with no options must cause `hello` to be built, and running `hello` must produce the string "Hello, World!" at the terminal. *You should not modify any C source to accomplish this!*

3 Grading

Turn in the file `Makefile` (and only this file) to Autograder. Full credit on the Autograder tests will result in credit for this lab. Anything less than full credit on the Autograder will result in no credit given for this lab.