

**Due:** Monday, 5 December at 2400

## 1 Introduction

For this final project, you will be writing a library package to provide facilities for manipulating graphs, plus two clients that use the package. The package must be general: no specifics of either client may find their way into the code for the graph package. We will check that by running our own clients against your package. That is, in contrast to past projects where we didn't care how you arranged the internals of your program as long as it functioned according to the specification, in this case, part of the project code—the API for the graph package—will be set in stone.

## 2 Client: Make

We've been using the MAKE program (specifically, GNU make: GMAKE) in most programming assignments this semester. For this project, we are going to implement a much stripped-down version. An input file (“makefile”) will consist of *rules* that indicate that one set of objects (called *targets*) depends on another set (called *prerequisites*). The syntax is

$$T_1 \quad T_2 \cdots T_n : P_1 \quad P_2 \cdots P_m$$

*command set*

where  $n \geq 1$ ,  $m \geq 0$ , and  $T_1$  starts in the first column. We'll call the first line of the rule its *header*. Each  $T_i$  or  $P_j$  is a string of non-whitespace characters other than ':', '=', '#', or '\'. The command set consists of lines that are either empty or else begin with a blank or tab. An *empty command set* consists only of whitespace (blanks, tabs, and newlines). The commands for one rule end at the next rule (or the end of file).

When  $n > 1$ , the rule is exactly equivalent to

$$\begin{array}{l} T_1 : P_1 \quad P_2 \cdots P_m \\ \quad \text{command set} \\ T_2 : P_1 \quad P_2 \cdots P_m \\ \quad \text{command set} \\ T_3 : P_1 \quad P_2 \cdots P_m \\ \quad \text{command set} \end{array}$$

Any line that starts with '#' is ignored. A header line may be continued by ending it with a backslash ('\'). The backslash and following newline get turned into one blank, so that the two lines become one.

## 2.1 Interpreting a Makefile

In general, the program is given a “makefile” in the format just described, plus a set of targets that it is to “rebuild,” and information about which objects currently exist and when they were last built. (In real life, MAKE finds out the latter for itself, by querying the file system. We will stop short of this level of realism). The program reads the Makefile sequentially.

After reading all the definitions, the program considers all the targets it has been asked to rebuild, and figures out which commands (if any) must be executed to rebuild each target, and in what order. For a target  $T$ ,

- If the makefile does not have a rule with  $T$  as the target, and  $T$  does not currently exist, there is an error. If  $T$  does exist, it is considered “built.”
- Otherwise, MAKE considers all the headers in the makefile whose target is  $T$  and collects the set of all prerequisites listed in those headers. It first builds all of these prerequisites, applying the procedure being described here. Next, if  $T$  does not exist, or if  $T$  is older than at least one of its prerequisites, MAKE executes the non-empty *commands* set from one of these rules, and then sets the creation time of  $T$  to be larger than that of any existing object (so that  $T$  is now the youngest object). It is an error if there is more than one non-empty set of commands for a target. MAKE never rebuilds something unnecessarily.

For this assignment, “executing” commands simply means printing them, after first performing the following substitutions:

- Replace all occurrences of  $\$^*$  with the set of all prerequisites.
- Replace all occurrences of  $\$?$  with the set of all prerequisites that are newer than the target.
- Remove leading whitespace.

In each case, to render a set of objects as a string, the objects are listed in alphabetical order, without duplicates, and separated by single blanks.

As usual, in response to input errors (in format, for example), you should output an error message on the standard error output and exit with a non-zero exit code. Do not output to the standard error output otherwise. Errors include missing prerequisites with no build instructions, attempting to build a target that has more than one non-empty commands set, syntax errors in the input files (see also §2.2), and circular dependencies.

## 2.2 Running Make

To run the MAKE client, a user will type

```
java make.Main [ -f MAKEFILE ] [ -D FILEINFO ] [ TARGET ... ]
```

where [...] indicates an optional argument. *MAKEFILE* (default **Makefile**) is the name of the makefile. *FILEINFO* (default **fileinfo**) describes the current existing objects and their ages (see below). Each *TARGET* is processed as a target in order (if there are no targets, use the target(s) of the first rule in the makefile as the default).

The *FILEINFO* argument is a file containing lines of the form

*NAME CHANGEDATE*

where *CHANGEDATE* is an integer indicating a time (the larger the time, the younger the named object). On a Unix system, this would be the number of seconds since the Epoch (the beginning of 1 January 1970), not counting leap seconds (making 2038 ( $2^{31}$  seconds past the Epoch) the Unix Y2K year). However, for us, the only thing that matters is that larger is later (so for the current time, you can choose any value larger than any of those in the *FILEINFO* file).

### 3 Client: Trip finder

I suspect you all have or have used some sort of GPS device by now, or gotten directions from the Google<sup>tm</sup> map service. Such systems “know” about a network of roads, and given two end points (and perhaps some waypoints in between) will pick out a shortest route. This next client will be a stripped-down version. Specifically, this program, given a map and a request—a sequence of two or more points on the map—will produce a shortest route from the first point to the second, third, and so forth.

The map file will be in free format (a sequence of “words” containing no whitespace and separated by whitespace—blanks, tabs, newlines, etc.). Information will have the following form:

$C_0$  *R* *L* *D*  $C_1$

where each  $C_i$  designates a place, *R* is the name of a road *L* is a numeric distance (a floating-point number), and *D* is one of the strings NS, EW, WE, SN. For example, the entry

Montara US\_1 56.0 NS Santa\_Cruz

would mean that *US\_1* run North to South from Montara to Santa Cruz, for a distance of 56.0 miles, and that there are no map points in between (I guess we have a rather idiosyncratically sparse map). Our road connections will all be two-way, so the sample entry also indicates a South to North-running route (SN) from Santa Cruz to Montara.

A request will consist of a file containing the names of two or more points on the map, separated by whitespace. For example,

Berkeley San\_Francisco Santa\_Cruz

In response, your program is supposed to print out a route in this format:

From Berkeley:

1. Take University Ave west for 0.1 miles.
2. Take Martin Luther King Jr Way south for 1.2 miles.
3. Take Ashby Ave west for 1.8 miles.
4. Take I-580 west for 1.0 miles.
5. Take I-80 west for 8.4 miles to San Francisco.
6. Take US-101 south for 34.5 miles.
7. Take CA-85 south for 13.3 miles.
8. Take CA-17 south for 22.2 miles.
9. Take CA-1 north for 1.0 miles.

To Santa Cruz.

When two adjacent segments use the same road with the same orientation (NS, EW), join them into one segment. Convert underscores to blanks when printing.

As usual, in response to input errors (in format, for example), you should output an error message on the standard error output and exit with a non-zero exit code. Do not output to the standard error output otherwise.

### 3.1 Running Trip Finder

To run the TRIPFINDER client, a user will type

```
java trip.Main [ -m MAP ] [ -o OUT ] [ REQUEST ]
```

*MAP* (default `Map`) is the name of the file containing the map data. *REQUEST* (default, the standard input) contains the request. *OUT* (default, the standard output) receives the solution.

## 4 The Graph Package

You should implement these two clients using a graph package whose interface we will supply. The package `graph` will export an abstract class `graph.Graph`, defining the API, and no other public classes (aside from the usual unit-testing class). We will test your graph package separately, exercising methods that might not be used by either of the clients you implement, so unit tests will be particularly useful. Likewise, we will test your clients using *our* graph package, so that your clients must rely only on the published spec, and not on any special behavior of your particular implementation.

## 5 Your Task

The staff directory will contain skeleton files for this project in `proj3`.

Please read *General Guidelines for Programming Projects*.

Be sure to include tests of your program (that is part of the grade). The makefile we provide has a convenient target for running such tests. Our skeleton directory contains a couple of trivial tests, but *these do not constitute an adequate set of tests!* Make up your tests ahead of time and update your makefile to run them. To help with testing and debugging, we will provide our own version of the program, so that you can test your program against ours (we'll be on the lookout for illegal moves). More details will follow.

The input to your program will come from fallible humans. Therefore, part of the problem is dealing gracefully with errors. When the user makes an illegal move, tell him and just give him another chance. We don't care about the message format you use to do this.

We will expect your finished programs to be workmanlike, and of course, will enforce the mechanical style standards. Make sure all methods are adequately commented—meaning that after reading the name, parameters, and comment on a method, you don't need to look at the code to figure out what a call will do. Don't leave debugging print statements lying around. In fact, don't use them; learn to use the debugger (either `gjdb` or that of **Eclipse** or your favorite Java-system vendor).

The nature of the graph package is such that you'll probably want to do extensive JUnit testing on it. Always feel free to add private setup procedures for testing purposes only that set up conditions (such as graphs) that are used by several different tests.

Our testing of your projects (but not our grading!) will be automated. The testing program will be finicky, so be sure that

```
gmake check
```

runs your tests.

## 6 Advice

Use `svn` regularly. It provides backup. It's the safest way to transport things from a copy of your project on a local computer to (`svn commit`) and from (`svn update`) the instructional machines. Frequent commits will limit the amount of work you have to do when a file gets messed up. Use `svn status` to check for files that should be added and to make sure all necessary commits have been done.

Use of `gmake style` not only encourages proper formatting, but also shows you missing documentation or left-over internal comments that might indicate work that still needs to be done, and can find certain error-prone constructs or bad practices.