



Problem Set 3: The Game of Fifteen

due by 7:00pm on Fri 10/2

Be sure that your code is thoroughly commented to such an extent that lines' functionality is apparent from comments alone.

Goals.

- Introduce you to larger programs and programs with multiple source files.
- Empower you with Makefiles and RCS.
- Acquaint you with pseudorandom numbers.
- Introduce you to literature in computer science.
- Play (but call it work).

Recommended Reading.

- Section 17 of <http://www.howstuffworks.com/c.htm>.
- Chapters 20 and 23 of *Absolute Beginner's Guide to C*.
- Chapters 13, 15, and 18 of *Programming in C*.

diff pset3.pdf hacker3.pdf.

- Hacker Edition dares you to implement `sort` in $O(n)$ instead of $O(n^2)$.
- Hacker Edition asks you to play God.

Academic Honesty.

All work that you do toward fulfillment of this course's expectations must be your own unless collaboration is explicitly allowed (*e.g.*, by some problem set or the final project). Viewing or copying another individual's work (even if left by a printer, stored in an executable directory, or accidentally shared in the course's virtual terminal room) or lifting material from a book, website, or other source—even in part—and presenting it as your own constitutes academic dishonesty, as does showing or giving your work, even in part, to another student.

Similarly is dual submission academic dishonesty: you may not submit the same or similar work to this course that you have submitted or will submit to another. Nor may you provide or make available solutions to problem sets to individuals who take or may take this course in the future. Moreover, submission of any work that you intend to use outside of the course (*e.g.*, for a job) must be approved by the staff.

You are welcome to discuss the course's material with others in order to better understand it. You may even discuss problem sets with classmates, but you may not share code. In other words, you may communicate with classmates in English, but you may not communicate in, say, C. If in doubt as to the appropriateness of some discussion, contact the staff.

You may even turn to the Web for instruction beyond the course's lectures and sections, for references, and for solutions to technical difficulties, but not for outright solutions to problems on problem sets or your own final project. However, failure to cite (as with comments) the origin of any code or technique that you do discover outside of the course's lectures and sections (even while respecting these constraints) and then integrate into your own work may be considered academic dishonesty.

All forms of academic dishonesty are dealt with harshly.

Grades.

Your work on this problem set will be evaluated along three primary axes.

Correctness. To what extent is your code consistent with our specifications and free of bugs?

Design. To what extent is your code written well (*i.e.*, clearly, efficiently, elegantly, and/or logically)?

Style. To what extent is your code readable (*i.e.*, commented and indented with variables aptly named)?

Getting Started.

- SSH to `nice.fas.harvard.edu` and execute the command below.

```
cp -r ~cs50/pub/src/psets/hacker3/ ~/cs50/
```

That command will copy the staff's `hacker3/` directory, inside of which is some “distribution code,” files (and subdirectories) that you'll need for this problem set, into your own `~/cs50/` directory. The `-r` switch triggers a “recursive” copy. Navigate your way to your copy by executing the command below.

```
cd ~/cs50/hacker3/
```

If you list the contents of your current working directory (remember how?), you should see the below. If you don't, don't hesitate to ask the staff for assistance.

```
fifteen/ find/ questions.txt
```

As this output implies, most of your work for this problem set will be organized within two subdirectories.

- Don't forget that the course's website has a bulletin board! Not only can you post questions of your own, you can also search for or browse answers to questions already asked by others. And never fear asking “dumb questions.” Students' posts to the course's bulletin board are anonymized. Only the staff, not fellow students, will know who you are!¹

¹ Thus, only we will make fun.

Find.

- Now let's dive into the first of those subdirectories. Execute the command below.

```
cd ~/cs50/hacker3/find
```

If you list the contents of this directory, you should see the below.

```
helpers.c helpers.h Makefile find.c generate.c
```

Wow, that's a lot of files, eh? Not to worry, we'll walk you through them.

- Implemented in `generate.c` is a "pseudorandom-number generator" (PRNG), a program that outputs a whole bunch of random numbers, one per line. Actually, these numbers are generated not so much randomly as they are "pseudorandomly." Because a computer is a deterministic device (*i.e.*, it can only do what it's told to do), it can't just pick a number off the top of its head. However, algorithms exist that enable a computer to generate sequences of numbers that appear to be random in the sense that there's no obvious pattern to them. C provides a function called `rand()` for exactly this purpose. The language also provides a function called `srand()` that is used to "seed" the PRNG. To "seed" a PRNG means to feed an initial value, s , to its generating algorithm, g . Typically, the first number returned by a PRNG is $g(s)$; the second is $g(g(s))$; the third is $g(g(g(s)))$; and so forth. Hence, you can generate the same sequence of "random numbers" simply by seeding the PRNG with the same initial value. The current time, often measured in seconds since some particular moment in the past, is typically used as the seed to a PRNG so that the seed is not hard-coded into a program but instead dynamic.

Anyhow, go ahead and compile this program by executing the command below.

```
gcc -ggdb -std=c99 -Wall -o generate generate.c
```

Wow, that's quite the command, eh? It turns out you've been executing commands like that one all along. Prior to this problem set, anytime you typed

```
gcc
```

it was as though you were typing

```
gcc -ggdb -std=c99 -Wall
```

because we had "aliased" the former command to the latter to save you keystrokes and avoid confusion. (Remember `cs50setup`? That's one of the things it did for you.) If curious, the `-ggdb` switch tells `gcc` to include "debugging symbols" in your binaries to facilitate debugging with GDB, per Week 4's first lecture. The `-std=c99` switch tells GCC that your code might include C99

syntax.² And the `-Wall` switch tells GCC to report all possible warnings anytime it detects possible problems with your code.

It's time for some training wheels to come off, though! We've thus deactivated that alias. But we don't expect you to start typing ridiculously long commands. We'll just start using `make` even more. For now, though, go ahead and run the program you just compiled by executing the command below.

```
generate
```

You should be informed of the program's proper usage, per the below.

```
Usage: generate n [s]
```

As this output suggests, this program expects one or two command-line arguments. The first, `n`, is required; it indicates how many pseudorandom numbers you'd like the generate. The second, `s`, is optional, as implied by the brackets; if supplied, it represents the value that the pseudorandom-number generator should use as its seed. Go ahead and run this program again, this time with a value of, say, 10 for `n`, as in the below; you should see a list of 10 pseudorandom numbers.

```
generate 10
```

Run the program a third time using that same value for `n`; you should see a different list of 10 numbers. Now try running the program twice more, still using that same value for `n`, but this time also providing some value for `s` both times, as in the below; the output of both executions should be identical.

```
generate 10 0
```

Think of this last command, with its seed of 0, as having generated the PRNG's 0th possible sequence of 10 pseudorandom numbers.

- Now take a look at `generate.c` itself with Nano. (Remember how?) Comments atop that file explain the program's overall functionality. But it looks like we forgot to comment the code itself. Read over the code carefully until you understand each line and then comment our code for us, replacing each `TODO` with a phrase that describes the purpose or functionality of the corresponding line(s) of code. Realize that a comment flanked with `/*` and `*/` can span lines whereas a comment preceded by `//` can only extend to the end of a line; the latter is a feature of C99. If curious about `rand` and `srand`, pull up the URLs below.

<http://www.cs50.net/resources/cppreference.com/stdother/rand.html>
<http://www.cs50.net/resources/cppreference.com/stdother/srand.html>

² C99 is the dialect (*i.e.*, version) of C that we've been using. For the curious:
<http://en.wikipedia.org/wiki/C99>.

Or execute the commands below.

```
man 3 rand
man 3 srand
```

Note that if you instead execute the command below, you'll pull up the man page for a program (not function) named `rand` in section 1 of the Linux Programmer's Manual.

```
man rand
```

Functions, by contrast, tend to be documented in sections 2 and 3. To avoid any ambiguity, then, you sometimes need to tell `man` the section you want. If curious as to what's where, execute, believe it or not, the command below.

```
man man
```

Once done commenting `generate.c`, re-compile the program to be sure you didn't break anything. Rather than execute that long command from earlier, though, simply execute the command below.

```
make generate
```

As you may recall from Problem Set 1, `make` automates compilation of your code. Notice, in fact, how `make` just executed a pretty long command for you, per the tool's output. However, as your programs grow in size, `make` can't do absolutely everything for you; you'll need to start telling `make` how to compile your program, particularly when they involve multiple source (*i.e.*, `.c`) files. And so we'll start relying on "Makefiles," configuration files that tell `make` exactly what to do.

How did `make` know how to compile `generate` in this case? Using Nano, go ahead and look at the file called `Makefile` that's in the same directory as `generate.c`. This `Makefile` is essentially a list of rules that we wrote for you that tells `make` how to build `generate` from `generate.c` for you. The relevant lines appear below.

```
generate: generate.c
    gcc -ggdb -std=c99 -Wall -Werror -o generate generate.c
```

The first line tells `make` that the "target" called `generate` should be built by invoking the second line's command. Moreover, that first line tells `make` that `generate` is dependent on `generate.c`, the implication of which is that `make` will only re-build `generate` on subsequent runs if that file was modified since `make` last built `generate`. Neat time-saving trick, eh? In fact, go ahead and execute the command below again, assuming you haven't modified `generate.c`.

```
make generate
```

You should be informed that `generate` is already up to date. Incidentally, know that the leading whitespace on that second line is not a sequence of spaces but, rather, a tab. Unfortunately, `make` requires that commands be preceded by tabs, so be careful not to change them to spaces with Nano (which automatically converts tabs to four spaces), else you may encounter strange

errors! The `-Werror` flag, incidentally, tells `gcc` to treat warnings (bad) as though they're errors (worse) so that you're forced (in a good, instructive way!) to fix them.

- Now take a look at `find.c` with Nano. Notice that this program expects a single command-line argument: a “needle” to search for in a “haystack” of values. Once done looking over the code, go ahead and compile the program by executing the command below.

```
make find
```

Notice, per that command's output, that Make actually executed the below for you.

```
gcc -ggdb -std=c99 -Wall -Werror -o find find.c helpers.c -lcs50
```

Notice further that you just compiled a program comprising not one but two `.c` files: `helpers.c` and `find.c`. How did `make` know what to do? Well, again, open up `Makefile` to see the man behind the curtain. The relevant lines appear below.

```
find: find.c helpers.c helpers.h  
      gcc -ggdb -std=c99 -Wall -Werror -o find find.c helpers.c -lcs50
```

Per the dependencies implied above (after the colon), any changes to `helpers.c`, `helpers.h`, or `find.c` will compel `make` to rebuild `find` the next time it's invoked for this target.

Go ahead and run this program by executing, say, the below.

```
find 13
```

You'll be prompted to provide some hay (*i.e.*, some integers), one “straw” at a time. As soon as you tire of providing integers, hit `ctrl-d` to send the program an EOF (end-of-file) character. That character will compel `GetInt` from CS 50's library to return `INT_MAX`, which, per `find.c`, will compel `find` to stop prompting for hay. The program will then look for that needle in the hay you provided, ultimately reporting whether the former was found in the latter. In short, this program searches an array for some value.

It turns out you can automate this process of providing hay, though, by “piping” the output of `generate` into `find` as input. For instance, the command below passes 1,024 pseudorandom numbers to `find`, which then searches those values for 13.

```
generate 1024 | find 13
```

Alternatively, you can “redirect” `generate`'s output to a file with a command like the below.

```
generate 1024 > numbers.txt
```

You can then redirect that file's contents as input to `find` with the command below.

```
find 13 < numbers.txt
```

Let's finish looking at that `Makefile`. Notice the line below.

```
all: find generate
```

This target implies that you can build both `generate` and `find` simply by executing the below.

```
make all
```

Even better, the below is equivalent (because `make` builds a `Makefile`'s first target by default).

```
make
```

If only you could whittle this whole problem set down to a single command! Finally, notice these last lines in `Makefile`:

```
clean:  
    rm -f *.o a.out core find generate
```

This target allows you to delete all files ending in `.o` or called `a.out`, `core` (`tsk, tsk`), `find`, or `generate` simply by executing the command below.

```
make clean
```

Be careful not to add, say, `*.c` to that last line in `Makefile`! (Why?) Any line, incidentally, that begins with `#` is just a comment.

- Phew, lots of good stuff so far, and it's almost time to start coding. But one last lesson for you. From personal (traumatic) experience, you probably already know that backups are a good thing. What you might not know is that a number of Linux tools exist to facilitate the process of backing up source code. Starting with this problem set, you'll want to use a utility called RCS (Revision Control System) to make regular, incremental backups of your source code. Not only will RCS enable you to restore your most recently backed-up copy of a file in the event of trauma, it will also enable you to restore different versions of your source code, in the event you realize that the code you wrote a few days ago was much better than what you've been producing since.

Go ahead and "check in" (*i.e.*, backup) your initial version of `find.c` by executing the command below.

```
ci find.c
```

You'll be prompted for a description for this file. Go ahead and describe the purpose of this file in a few words, then enter `.` or hit `ctrl-d` on a line of its own to save the description. You should be informed that version 1.1 of this file, your "initial revision," has been checked in. If you list the contents of your current working directory, you'll notice that you now have a directory called `RCS` therein, inside of which is `find.c,v`, which is where RCS (*i.e.*, `ci`) records changes to your file.

Henceforth, anytime you want to check in your latest version of `find.c`, simply execute the same command as before, per the below.

```
ci find.c
```

No longer will you be prompted for a description but, rather, a “log message,” which is even more important than the file’s initial description. Log messages are supposed to help you remember what’s different between this version and your last (e.g., “I changed my `while` loop to a `do-while` loop”). Entering that message might be tedious, but, trust us, you’re not going to remember what was special about version 1.9 at 3:00am. without a little help. As before, enter `.` or hit `ctrl-d` on a line of its own to save your message. Each time you check in a newer version of your file, RCS will assign an appropriate version number.

Suppose, for future reference, that you want to restore, say, version 1.1 of `find.c`. If you don’t want to clobber (i.e., overwrite) your current version, be sure to check it in first! Then proceed to execute the command below.

```
co -r1.1 find.c
```

You should find that version 1.1 of `find.c` has been restored to your current working directory. So that you know which version of `find.c` is which, execute the command below to see your own log messages.

```
rlog find.c
```

To save time, know that you can check in multiple files at once, as with the command below.

```
ci *.c *.h
```

For disk space’s sake, RCS will allow you to check in source files on `nice.fas.harvard.edu` but not binaries.

- And now the fun begins! Notice that `find.c` calls `sort`, a function declared in `helpers.h`. Unfortunately, we forgot to implement that function fully in `helpers.c`! Take a peek at `helpers.c` with Nano, and you’ll see that `sort` returns immediately, even though `find`’s main function does pass it an actual array. To be sure, we could have put the contents of `helpers.h` and `helpers.c` in `find.c` itself. But it’s sometimes better to organize programs into multiple files, especially when some functions (e.g., `sort`) are essentially utility functions that might later prove useful to other programs as well, much like those in CS 50’s own library.

Incidentally, recall the syntax for declaring an array. Not only do you specify the array’s type, you also specify its size between brackets, just as we do for `haystack` in `find.c`:

```
int haystack[HAY_MAX];
```

But when passing an array, you only specify its name, just as we do when passing `haystack` to `sort` in `find.c`:

```
sort(haystack, size);
```

(Why do we also pass in the size of that array separately?)

When declaring a function that takes a one-dimensional array as an argument, though, you don't need to specify the array's size, just as we don't when declaring `sort` in `helpers.h` (and `helpers.c`):

```
void sort(int values[], int n);
```

Go ahead and implement `sort` so that the function actually sorts, from smallest to largest, the array of numbers that it's passed, in such a way that its running time is in $O(n)$, where n is the array's size.³ You may not alter the function's declaration. In particular, its return type must remain `void`. Rather than return a new, sorted array, then, the function must instead "destructively" sort the actual array that it's passed. However, you may assume that each of the arrays' numbers will be non-negative and less than `LIMIT`, a constant defined in `generate.c`.⁴ But the array might contain duplicates.

Don't forget to check in `helpers.c` before making your changes! (Remember how?)

We leave it to you to determine how to test your implementation of `sort`. But don't forget that `printf` and, now, `gdb` are your friends. And don't forget that you can generate the same sequence of pseudorandom numbers again and again by explicitly specifying `generate`'s seed. Before you ultimately submit, though, be sure to remove any such calls to `printf`, as we like our programs' outputs just the way they are!

Incidentally, check out **Resources** on the course's website for a great little quick-reference guide for `gdb`!

If you'd like to play with the staff's own implementation of `find` on `nice.fas.harvard.edu`, you may execute the below.

```
~cs50/pub/solutions/hacker3/find
```

- Bulletin board!
- Now that `sort` (presumably) works, you can improve upon `search`. Notice that our version implements linear search. Rip out the lines that we've written and re-implement `search` as binary search! You are welcome to take an iterative or a recursive approach. If you pursue the

³ Technically, because we've bounded with a constant the amount of hay that `find` will accept (and because the value of `sort`'s second parameter is bounded by an `int`'s 32 bits) the running time of `sort`, however implemented, is arguably $O(1)$. Even so, for the sake of this asymptotic challenge, think of the size of `sort`'s input as n .

⁴ Leverage that assumption!

latter, though, know that you may not change our declaration of `search`, but you may write a new, recursive function (that perhaps takes different parameters) that `search` itself calls.

The Game Begins.

- And now it's time to play. The Game of Fifteen is a puzzle played on a square, two-dimensional board with numbered tiles that slide. The goal of this puzzle is to arrange the board's tiles from smallest to largest, left to right, top to bottom, with an empty space in board's bottom-right corner, as in the below.⁵



Sliding any tile that borders the board's empty space into that space constitutes a "move." Although the configuration above depicts a game already won, notice how the tile numbered 12 or the tile numbered 15 could be slid into the empty space. Tiles may not be moved diagonally, though, or forcibly removed from the board.

Depicted below is just one configuration from which the puzzle is solvable.⁶



⁵ Figure from http://en.wikipedia.org/wiki/Fifteen_puzzle.

⁶ Figure adapted from http://en.wikipedia.org/wiki/Fifteen_puzzle.

- Navigate your way to `~/cs50/hacker3/fifteen/`, and take a look at `fifteen.c` with Nano. Within this file is the entire framework for The Game of Fifteen (and variants thereof).

Implement God Mode for this game.

First implement `init` in such a way that the board is initialized to a pseudorandom but solvable configuration.⁷ Then complete the implementation of `draw`, `move`, and `won` so that a human can actually play the game. But embed in the game a cheat, whereby, rather than typing an integer between 1 and $d^2 - 1$, where d is the board's height and width, the human can also type

GOD

to compel "the computer" to take control of the game and solve it (using any strategy, optimal or non-optimal), making, say, only four moves per second so that the human can actually watch. Presumably, you'll need to swap out `GetInt` for something more versatile. It's fine if your implementation of God Mode only works (bearably fast) for $d \leq 4$; you need not worry about testing God Mode for $d > 4$. Oh and you can't implement God Mode by remembering how `init` initialized the board (as by remembering the sequence of moves that got your program to some pseudorandom but solvable state). That'd be, um, cheating. At cheating.

To test your implementation, you can certainly try playing it yourself, with or without God Mode enabled. (Know that you can quit your program by hitting `ctrl-c`.) Be sure that you (and we) cannot crash your program, as by providing bogus tile numbers. And know that, much like you automated input into `find`, so can you automate execution of this game via input redirection if you store in some file a winning sequence of moves for some configuration.

Any design decisions not explicitly prescribed herein (*e.g.*, how much space you should leave between numbers when printing the board) are intentionally left to you. Presumably the board, when printed, should look something like the below (albeit pseudorandom), but we leave it to you to implement your own vision.

```
15 14 13 12
11 10  9  8
 7  6  5  4
 3  1  2
```

Feel free to tweak the appropriate argument to `usleep` to speed up animation. In fact, you're welcome to alter the aesthetics of the game. For (optional) fun with "ANSI escape sequences," including color, take a look at our implementation of `clear` and check out the URL below for more tricks.

http://isthe.com/chongo/tech/comp/ansi_escapes.html

⁷ To be clear, whereas the standard edition of this problem set requires that the board be initialized to a specific configuration, this Hacker Edition requires that it be initialized to a pseudorandom but still solvable configuration.

But we ask that you not alter the flow of logic in `main` so that we can automate some tests of your program. If in doubt as to whether some design decision of yours might run counter to the staff's wishes, simply contact your teaching fellow.

If you'd like to play with the staff's own implementation of `fifteen` on `nice.fas.harvard.edu`, including God Mode, you may execute the below.

```
~cs50/pub/solutions/hacker3/fifteen
```

Speaking of God Mode, where to begin? Well, first read up on this Game of Fifteen. Wikipedia is probably a good starting point:

```
http://en.wikipedia.org/wiki/N-puzzle
```

Then dive a bit deeper, perhaps reading up on an algorithm called A*.

```
http://en.wikipedia.org/wiki/A*_search_algorithm
```

Consider using "Manhattan distance" (aka "city-block distance") as your implementation's heuristic. If you find that A* takes up too much memory (particularly for $d \geq 4$), though, you might want to take a look at iterative deepening A* (IDA*) instead:

```
https://eprints.kfupm.edu.sa/34045/1/34045.pdf
```

Another paper on point is that here:

```
http://www.cs.ualberta.ca/~tony/RecentPapers/pami94.pdf
```

The staff's own implementation, meanwhile, utilizes an algorithm like that in this paper:

```
http://www.eng.unt.edu/ian/pubs/saml.pdf
```

You're welcome to expand your search for ideas beyond those in these papers, but take care that your research does not lead you to actual code. Curling up with others' pseudocode is fine, but please click away if you stumble upon actual implementations (whether in C or other languages).

Alright, get to it, implement this game. But first check in `fifteen.c`!

Bulletin board!!

Submitting Your Work.

- Ensure that your work is in `~/cs50/hacker3/`. Submit your work by executing the command below.

```
cs50submit hacker3
```

Thereafter, follow any on-screen instructions until you receive visual confirmation of your work's successful submission. You will also receive a "receipt" via email to your FAS account, which you should retain until term's end. You may re-submit as many times as you'd like; each resubmission will overwrite any previous submission. But take care not to re-submit after the problem set's deadline, as only your latest submission's timestamp is retained.

HACKER EDITION