

Problem Set 3: The Game of Fifteen

due by 7:00pm on Fri 10/1

Per the directions at this document's end, submitting this problem set involves submitting source code on `cloud.cs50.net` as well as filling out a Web-based form (the latter of which will be available after lecture on Wed 9/29), which may take a few minutes, so best not to wait until the very last minute, lest you spend a late day unnecessarily.

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.
- Introduce you to literature in computer science.
- Implement a party favor.

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 in writing by the course's instructor. Collaboration in the completion of problem sets is not permitted unless otherwise stated by some problem set's specification.

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 course's instructor.

You may 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. If the course refers some matter to the Administrative Board and the outcome for some student is *Warn*, *Admonish*, or *Disciplinary Probation*, the course reserves the right to impose local sanctions on top of that outcome for that student that may include, but not be limited to, a failing grade for work submitted or for the course itself.

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 `cloud.cs50.net` and execute the command below.¹

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

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 home directory. The `-r` switch triggers a “recursive” copy. Navigate your way to your copy by executing the command below.

```
cd ~/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/
```

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

- Don't forget about `help.cs50.net`! 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!²

And, of course, there are plenty of office hours: <http://www.cs50.net/ohs/>.

¹ Note that there's a space between `pset3` and that last tilde (`~`), the latter of which, recall, represents your home directory.

² Thus, only we will make fun.

Find.

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

```
cd ~/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 program that uses `rand` to generate a whole bunch of pseudorandom numbers, one per line. (Remember `rand` from Problem Set 1?) Go ahead and compile this program by executing the command below.

```
make generate
```

Now 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 to generate. The second, `s`, is optional, as the brackets are meant to imply; if supplied, it represents the value that the pseudorandom-number generator should use as its seed. (Remember from Problem Set 1 what a seed is?) 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 with a value for `s` too (*e.g.*, 0), as in the below.

```
./generate 10 0
```

Now run that same command again:

```
./generate 10 0
```

Bet you saw the same "random" sequence of ten numbers again? Yup, that's what happens if you don't vary a pseudorandom number generator's initial seed.

- 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 (the version of C that we've been using). If Problem Set 1 feels like a long time ago, you might want to read up on `rand` and `srand` again at the URLs below.

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

Or you can execute the commands below.

```
man rand
man srand
```

Once done commenting `generate.c`, re-compile the program to be sure you didn't break anything by re-executing the command below.

```
make generate
```

If `generate` no longer compiles properly, take a moment to fix what you broke!

Now, recall that `make` automates compilation of your code so that you don't have to execute `gcc` manually along with a whole bunch of switches. 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` won't be able to infer from context anymore how to compile your code; 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? It actually used a configuration file that we wrote. 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 -Wformat=0 -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, recall, 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 -Wformat=0 -o find find.c helpers.c -lcs50 -lm
```

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 -Wformat=0 -o find find.c helpers.c -lcs50 -lm
```

Per the dependencies implied above (after the colon), any changes to `find.c`, `helpers.c`, or `helpers.h` 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 the CS50 Library to return `INT_MAX`, a constant that, 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.

- 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 the CS50 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`:

```
if (!sort(haystack, size))
{
    printf("Could not sort haystack.\n");
    return 2;
}
```

(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`):

```
bool 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 Yes, this running time is possible, `hax0r`, because you may assume that each of the arrays' numbers will be non-negative and less than `LIMIT`, a constant defined in `generate.c`.⁴ However, realize that the array might contain duplicates. Take care, though, not to alter our declaration of `sort`. Its prototype must remain:

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

As this return type of `bool` implies, this function must not return a sorted array; it must instead "destructively" sort the actual array that it's passed by moving around the values therein, thereafter returning `true` if and only if sorting was successful. (Odds are it will always be successful, but, if you employ certain tricks, it's possible to, say, run out of memory.) As we'll discuss in Week 4, arrays are not passed "by value" but instead "by reference," which means that `sort` will not be passed a copy of an array but, rather, the original array itself.

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 they way they are!

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!

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 `cloud.cs50.net`, you may execute the below.

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

- Need help? Head to `help.cs50.net`!
- Now that `sort` (presumably) works, it's time to improve upon `search`, the other function that lives in `helpers.c`. Notice that our version implements "linear search," whereby `search` looks for `value` by iterating over the integers in `array` linearly, from left to right. Rip out the lines that we've written and re-implement `search` as binary search, that divide-and-conquer strategy that we employed in Week 0 in order to search through phone book.⁵ You are welcome to take an iterative or a recursive approach. If you pursue the 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.⁶



⁵ No need to tear anything in half.

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

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.⁷



- Navigate your way to `~/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.⁸ 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.

⁷ Figure adapted from http://en.wikipedia.org/wiki/Fifteen_puzzle.

⁸ 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.

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  _
```

Incidentally, recall that the positions of tiles numbered 1 and 2 should only be swapped (as they are in the 4×4 example above) if the board has an odd number of tiles (as does the 4×4 example above). If the board has an even number of tiles, those positions should not be swapped. Consider, for instance, the 3×3 example below:

```
8  7  6
5  4  3
2  1  _
```

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

You're welcome to write your own functions and even change the prototypes of functions we wrote. 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 `cloud.cs50.net`, 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:

<http://webdocs.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!

- [help.cs50.net!](http://help.cs50.net)

How to Submit.

In order to submit this problem set, you must first execute a command on `cloud.cs50.net` and then submit a (brief) form online; the latter will be posted after lecture on Wed 9/29.

- SSH to `cloud.cs50.net`, if not already there, and then execute:

```
cd ~/hacker3
```

If informed that there is “no such file or directory,” odds are you skipped some step(s) earlier! Not a problem, we’ll help you fix in a moment. If the command indeed works, proceed to execute this command next:

```
ls
```

At a minimum, you should see `fifteen` and `find`. If not, odds are you skipped some more steps earlier! If you do see those files, you’re ready to submit your source code to us:⁹

```
~cs50/pub/bin/submit hacker3
```

That command will essentially copy your entire `~/hacker2` directory into CS50’s own account, where your TF will be able to access it. You’ll know that the command worked if you are informed that your “work HAS been submitted.” If you instead encounter an error that doesn’t appear to be a mistake on your part, do try running the command one or more additional times. (The `submit` program’s a bit old and annoyingly flakey.) But if informed that your “work HAS been submitted,” it indeed has.

⁹ Note that there is no slash between this tilde (~) and `cs50`.

Now, if you instead see some message informing you that your code is not where it should be, you'll need to fix that before you submit. Review the last two pages of Problem Set 1's spec for some tips!

- Anytime after lecture on Wed 9/29 but before this problem set's deadline, head to the URL below where a short form awaits:

`http://www.cs50.net/psets/3/`

If not already logged in, you'll be prompted to log into the course's website.

Once you have submitted that form (as well as your source code), you are done!

This was Problem Set 3.

HACKER EDITION