

CSc 372, Fall 2006
Assignment 3
Due: Thursday, October 5 at 23:59:59

Problem 1. (7 points) longest.rb

Write a Ruby program that reads lines from standard input and upon end of file writes the longest line to standard output. If there are ties for the longest line, `longest` writes out all the lines that tie. If there is no input, `longest` produces no output.

Don't overlook the restrictions below.

Examples:

```
% cat lg.1
a test
for
the program
here
% ruby longest.rb < lg.1
the program
%
% cat lg.2
xx
a
YY
b
zz
% ruby longest.rb < lg.2
xx
YY
zz
% grep ^o /usr/share/dict/words | ruby longest.rb
overindividually
overintellectualization
% ruby longest.rb < /dev/null
% ruby longest.rb < /dev/null | wc -c
0
%
```

Restrictions:

- No comparisons, such as `<`, `==`, `!=`, `<=>`, and `String#casecmp` may be used.
- No arithmetic operations, such as addition and subtraction, may be used.
- The only datatypes you may use are integers and strings. In particular, you may not use arrays.

You are permitted to employ the comparison that's implicit in control structures. For example, statements like

```
while x do ...           # OK
if f(x) then ...        # OK
```

are permitted.

However, a statement such as

```
while x > 1 do          # NOT PERMITTED!
```

is not permitted—it contains a comparison ($x > 1$).

Problem 2. (7 points) `seqwords.rb`

For this problem you are to write a Ruby program that reads a series of words from standard input, one per line, and then prints lines with the words sequenced according to a series of specifications, also one per line and read from standard input.

Don't overlook the restrictions below.

Here is an example with four words and three specifications:

```
% cat sw.1
one
two
three
four
.
1
2
3
.
3
2
1
1
2
3
.
4
1
% ruby seqwords.rb < sw.1
one two three
three two one one two three
four one
%
```

Note that lines containing only a period (.) end the word list and also separate specifications. For output, words are separated with a single blank. Here's another example:

```
% cat sw.2
tick
.
1
.
1
1
% ruby seqwords.rb < sw.2
tick
tick tick
%
```

Assume that there is at least one word and at least one sequencing specification. Assume that each sequencing specification has at least one number. Assume that all entries in the sequencing specifications are integers and in range for the list of words. Assume that words are between 1 and 1000 characters in length, inclusive. Because periods separate specifications, the input will never end with a period.

Restriction: The only types you may use are integer and string. In particular, you may not use arrays.

Problem 3. (12 points) `minmax.rb`

Write a Ruby program that reads lines from standard input and determines which line(s) are the longest and shortest lines in the file. The minimum and maximum lengths are output along with the line numbers of the line(s) having that length.

```
% cat mm.1
just a
test
right here
x
% ruby minmax.rb < mm.1
Min length: 1 (4)
Max length: 10 (3)
%
```

The output indicates that the shortest line is line 4; it is one character in length. The longest line is line 3; it is ten characters in length.

Another example:

```
% cat mm.2
xxx
XX
YYY
YY
zzz
qqq
% ruby minmax.rb < mm.2
Min length: 2 (2, 4)
Max length: 3 (1, 3, 5, 6)
%
```

In this case, lines 2 and 4 are tied for being the shortest line. Likewise, several lines are tied for maximum length.

If the input file is empty, the program should output a single line that states "Empty file":

```
% ruby minmax.rb < /dev/null
Empty file
%
```

A final example, run on lectura:

```
% ruby minmax.rb < /usr/share/dict/words
Min length: 1 (9, 38, 40, 31876, 31878, 57049, 57051, 97101, 97103,
120349, 120351, 137529, 137530, 153807, 153809, 168908, 168909,
188220, 188221, 203658, 203660, 208077, 208079, 214569, 214571,
229273, 229277, 255473, 255475, 272147, 272151, 287461, 287463,
329584, 329586, 332857, 332859, 354559, 354561, 406354, 406355,
432266, 432268, 456211, 456213, 463230, 463232, 475253, 475254,
475885, 475887, 477707, 477708)
Max length: 45 (308762)
% ruby minmax.rb < /usr/share/dict/words | wc -l
2
%
```

Although output is shown wrapped across several lines, the first invocation above produces only two lines of output, as shown when the output is piped into `wc -l`.

IMPORTANT: DO NOT assume any maximum length for input lines.

Extra credit: (5 points) You might be inclined to have some repetitious code in your `minmax.rb`, such as an if-then-else that handles minimum lengths and a nearly identical if-then-else that handles maximum lengths. If your solution contains no repetitious code, you'll earn a bonus of five points.

To request inspection for this bonus, include this comment:

```
# Look! No repetition!
```

Even something like the following, albeit short, will be considered repetitious:

```
mins = [1]
maxs = [1]
```

(No, don't say `mins = maxs = [1]`. Both variables will reference the same list—a change in one will be seen in the other.)

Problem 4. (18 points) `xfield.rb`

For this problem you are to write a Ruby program that extracts columns of data from standard input. Command line arguments specify characters that delimit columns, the columns to extract, and strings to separate extracted columns in the final output.

Here is an input file:

```
% cat xf.1
one      1      1.0
two      2      2.0
three    3      3.0
four     4      4.0
twenty   20     20.0
%
```

By default, fields are considered to be delimited by one or more spaces. The English text and the real numbers could be extracted like this:

```
% ruby xfield.rb 1 3 < xf.1
one      1.0
two      2.0
three    3.0
four     4.0
twenty   20.0
%
```

Note that field numbering is 1-based—the first field is 1, not 0.

`xfield` can be used to reorder fields:

```
% ruby xfield.rb 3 2 1 < xf.1
1.0      1      one
2.0      2      two
3.0      3      three
4.0      4      four
20.0     20     twenty
%
```

`xfield` supports negative indexing, just like Ruby arrays:

```
% ruby xfield.rb -1 1 < xf.1
1.0      one
2.0      two
3.0      three
4.0      four
20.0     twenty
% ruby xfield.rb -1 1 2 -2 < xf.1
1.0      one      1      1
2.0      two      2      2
3.0      three    3      3
4.0      four     4      4
20.0     twenty   20     20
%
```

If a field reference is out of bounds, the string "<NONE>" is used:

```
% ruby xfield.rb 1 10 2 < xf.1
one      <NONE>  1
two      <NONE>  2
three    <NONE>  3
four     <NONE>  4
twenty   <NONE> 20
%
```

By default, output fields are separated by a single tab. (Use "\\t".) This default separator can be overridden with the `-s` flag. The separator must be at least one character in length.

```
% ruby xfield.rb -s... 1 3 1 <xf.1
one...1.0...one
two...2.0...two
three...3.0...three
four...4.0...four
twenty...20.0...twenty
%
```

An delimiter other than space can be specified with the `-d` option. To extract login ids and real names from `/etc/passwd`, one might use this:

```
% ruby xfield.rb -d: 1 5 < /etc/passwd | head
root      root
rootcsh   Csh root on lectura
root      Root on Lectura
...lots more..
```

Note that the `-s` and `-d` options are single arguments—there's no space between `-s/-d` and the following string. The behavior of something like `xfield -s ... -d = 1` is undefined.

Non-numeric arguments other than the `-s` and `-d` flags are considered to be text to be included in each output line. If a textual argument (not a number) falls between two field specifications (two numbers), that text is used instead of the separator:

```
% ruby xfield.rb int= 2 ", real=" 3 ", english=" 1 < xf.1
int=1, real=1.0, english=one
int=2, real=2.0, english=two
int=3, real=3.0, english=three
int=4, real=4.0, english=four
int=20, real=20.0, english=twenty
%
```

Note the use of quotation marks to form an argument that contains blanks. The shell strips off the quotation marks so that the resulting arguments passed to the program do not have quotes. See the *Implementation notes for xfield* section below for more on this.

Here's that rule again:

If a textual argument (not a number) falls between two field specifications (two numbers), that text is used instead of the separator:

Below are some more examples showing the rule in action. A blank line has been inserted in this write-up

after the output of each command.

```
% cat xf.2
one two three four

% ruby xfield.rb -s. 1 2 3 < xf.2
one.two.three

% ruby xfield.rb -s. A 1 2 3 B C < xf.2
Aone.two.threeBC

% ruby xfield.rb -s. A 1 B C 2 3 D < xf.2
AoneBCtwo.threeD

% ruby xfield.rb -s. A 1 B C 2 D 3 E < xf.2
AoneBCtwoDthreeE
```

Below are some cases that bring all the elements into play. Trailing blank lines have been added for readability.

```
% cat xf.3
xxxxxxxxAxxxxxxxxxxxxBxC
DxExF
xG1xG2
xxxxHIxxxJKxxxLMNxxxOPQRSx

% ruby xfield.rb -dx -s- 1 2 3 < xf.3
A-B-C
D-E-F
G1-G2-<NONE>
HI-JK-LMN

% ruby xfield.rb -dx -s- -1 ... -2 -3 @ < xf.3
C...B-A@
F...E-D@
G2...G1-<NONE>@
OPQRS...LMN-JK@

% ruby xfield.rb -s/ -de 1 2 < xf.1
on/      1      1.0
two      2      2.0/<NONE>
thr/     3      3.0
four     4      4.0/<NONE>
tw/nty   20     20.0
```

If there are no input lines, `xfield` produces no output:

```
% ruby xfield.rb -s/ -d: 1 x 2 3 < /dev/null
%
```

If no fields are specified, an error is printed and execution terminates

```
% ruby xfield.rb
xfield: no fields specified
% ruby xfield.rb -d.
xfield: no fields specified
%
```

To terminate execution, use code like this:

```
if ...no fields found... then
  puts "xfield: no fields specified"
  exit 1
end
```

The only error handling required for `xfield` is the "no fields" situation just mentioned. In all other odd situations, "the behavior is undefined", which means any behavior is ok—run-time errors, curious results, etc. are no problem if the user misuses `xfield`.

For example, `xfield -s -d:, 1`, is invalid for two reasons: (1) `-s` should be immediately followed by a separator string, like `-s,` or `-s///`. (2) `-d` should be followed by a single character, not two.

Implementation notes for `xfield`

gets vs. STDIN.gets

The `gets` method does a little more than simply reading lines from standard input. *If command line arguments are specified gets will consider those arguments to be file names. It will then try to open those files and produce the lines from each in turn.* That's really handy in some cases but it gets in the way for `xfield`. To avoid this behavior, **don't use just `line = gets` to read lines**. Instead, do this:

```
while line = STDIN.gets do
```

This limits `gets` to the contents of standard input.

Delimiter-specific behavior in `String#split`

It astounds me but the fact is that `split` behaves differently when the delimiter is a space:

```
>> " a b c ".split(" ")
=> ["a", "b", "c"]

>> ".a..b..c.".split(".")
=> [ "", "a", "", "b", "", "c" ]
```

Command line argument handling

The command line arguments specified when a Ruby program is run are made available as strings in `ARGV`, an array. Here is `echo.rb`, a Ruby program that prints the command line arguments:

```
printf("%d arguments:\n", ARGV.length)
for i in 0...ARGV.length do      # three dots goes to length -1
  printf("argument %d is '%s'\n", i, ARGV[i])
end
```

Execution:

```
% ruby echo.rb -s -s2 -abc x y
5 arguments:
argument 0 is '-s'
argument 1 is '-s2'
```



```
argument 2 is '-abc'
argument 3 is 'x'
argument 4 is 'y'
```

Quotes and backslashes specified on the command line are processed by the shell. In the usual case, the program doesn't "see" them. Example:

```
% ruby echo.rb int= 2 ", real=" 3 ", english="
5 arguments:
argument 0 is 'int='
argument 1 is '2'
argument 2 is ', real='
argument 3 is '3'
argument 4 is ', english='
% ruby echo.rb " " ' x ' \ y\ ""
4 arguments:
argument 0 is ' '
argument 1 is ' x '
argument 2 is ' y '
argument 3 is ''
%
%
```

The shell does provide some mechanisms to allow quotes and backslashes to be transmitted in arguments:

```
% ruby echo.rb "" \\x\\
2 arguments:
argument 0 is ''
argument 1 is '\\x\\'
%
%
```

Additionally, the shell intercepts the < and > redirection operators—the program never sees them:

```
% ruby echo.rb 1 2 3 < lg.1
3 arguments:
argument 0 is '1'
argument 1 is '2'
argument 2 is '3'
% ruby echo.rb 1 2 3 < lg.1 >out
% cat out
3 arguments:
argument 0 is '1'
argument 1 is '2'
argument 2 is '3'
%
%
```

The above examples were produced with a UNIX shell but you'll see the same behavior with Windows Command Prompt.

BOTTOM LINE: Don't add code to your solution that attempts to process those shell metacharacters—that's the job of the shell, not your program!

An admonishment/HINT about argument handling

I've seen many students turn command line argument handling into an incredibly complicated mess. Don't do that! Here's an easy way to process arguments in `xfield`: Iterate over the elements in `ARGV`, the variable that references an array of strings that holds the command line arguments. If the argument starts

with "-s" or "-d" then save the rest of the string for later use. If `argument.to_i` produces something other than zero, then add the value (as an integer) to an array that specifies what's to be printed for each line. If `argument.to_i` produces zero, add `argument` to that same array. That's about 15 lines of simple code.

If the user specifies multiple `-s` and/or `-d` arguments, just let the last one "win". For example,

```
ruby xfield.rb -s1 -d: -d. -s2 -d, -s. 1
```

is equivalent to

```
ruby xfield.rb -d, -s. 1
```

A HINT on handling the textual argument/seperator rule

One way to handle the textual argument/seperator rule is to simply make a pass over the argument array and if two consecutive numbers are encountered, put the separator between them, as if the user had done that in the first place. For example, if the separator is ".", the specification

```
1 3 x 4 x -2 1
```

would be transformed into

```
1 . 3 x 4 x -2 . 1
```

A hint in a hint: Think about representing the above specification with this Ruby list:

```
[0, ".", 2, "x", 3, "x", -2, ".", 0]
```

Note that there is a combination of integers and strings.

Problem 5. (2 points) `answers.txt`

Create a text file named `answers.txt` with answers to the following questions. **DO NOT submit a Word document, PDF, rich text file, etc.**—I want plain text.

- (1 point) Ruby has two sets of logical operators: `&&`, `||`, `!` and lower-precedence equivalents `and`, `or`, `not`. Present a brief argument in favor of one set, like most languages, or Ruby's choice to provide two sets.
- (1 point) What's a good reason that Java provides `s.charAt(n)` instead of simply using `s[n]`?
- (1 point extra credit) Estimate how long it took you to complete this assignment. Other comments about the assignment are welcome, too—I appreciate all feedback, good or bad.

Deliverables

The deliverables for this assignment are these files: `longest.rb`, `seqwords.rb`, `minmax.rb`, `xfield.rb` and `answers.txt`. That list can be found on `lectura` in `/home/cs372/fall06/a3/delivs`.

Use the `turnin` tag `372_3` to submit your solutions.

Corrections and FAQs, late submissions, turnin, retests, etc.

Refer to the write-up for the first assignment for details on these topics and similar ones.

Miscellaneous

The output of your solutions should exactly match the output shown in this write-up. The data files used in the examples can be found in `/home/cs372/fall106/a3` on `lectura`.

Keep in mind the point value of each problem; don't invest an inordinate amount of time in a problem before you ask for help with it. Remember that the purpose of the assignments is to build understanding of the course material by applying it to solve problems—assignments are not intended to kill you. Try to think in terms of spending no more than seven hours on this assignment. Seek the help of Poorna and me as needed to meet your time budget.

As stated in the syllabus, a ten point homework problem corresponds to one point on your final average.

Restrictions notwithstanding you can use any elements of Ruby that you desire but the assignment is written with the intention that it can be completed easily using only the material presented on slides 1-76.

If you wish, you may incorporate code from lectures in your solutions.

Feel free to use comments to document your code as you wish but note that no comments, not even your name, are required.

Solutions will be scored only on correctness and compliance with the cited restrictions. If you're worried about whether a solution complies with the restrictions, mail it to `cs372-staff` and we'll look it over, assuming time permits. (Avoid the last-minute rush!) **BE VERY CAREFUL that you don't get the addresses mixed up and mail your solution to the whole class.**

I hate to have to mention it but keep in mind that I don't give cheaters a second chance to waste everybody's time. If you give your code to somebody else and they turn it in, you'll both likely fail the class, and more. (See the syllabus for the details.)