

Advanced `bash` scripting

(block course)

Solutions to the exercises



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Michael F. Herbst

michael.herbst@iwr.uni-heidelberg.de

<https://michael-herbst.com>

Interdisziplinäres Zentrum für wissenschaftliches Rechnen
Ruprecht-Karls-Universität Heidelberg

6th – 10th November 2017

Contents

Contents	i
Solutions to the exercises	1
Solution to 1.6	1
Solution to 1.7	1
Solution to 1.8	2
Solution to 2.1	3
Solution to 2.2	3
Solution to 2.3	4
Solution to 2.4	4
Solution to 2.5	4
Solution to 2.6	5
Solution to 2.7	5
Solution to 2.8	6
Solution to 2.9	6
Solution to 2.10	7
Solution to 3.2	7
Solution to 3.3	7
Solution to 3.4	8
Solution to 3.8	9
Solution to 3.9	9
Solution to 3.10	10
Solution to 3.11	10
Solution to 4.1	11
Solution to 4.2	11
Solution to 4.4	12
Solution to 4.5	13
Solution to 4.6	15
Solution to 4.7	16
Solution to 4.8	16
Solution to 4.13	17
Solution to 4.14	18
Solution to 4.15	18
Solution to 4.16	20
Solution to 4.17	21
Solution to 4.18	21
Solution to 4.19	21

Solution to 5.1	22
Solution to 5.2	23
Solution to 5.3	24
Solution to 5.4	25
Solution to 5.5	26
Solution to 5.6	27
Solution to 5.7	28
Solution to 6.2	29
Solution to 6.3	30
Solution to 6.4	31
Solution to 6.5	32
Solution to 6.6	33
Solution to 6.7	37
Solution to 7.2	38
Solution to 7.3	38
Solution to 7.4	38
Solution to 7.5	39
Solution to 7.6	39
Solution to 7.7	40
Solution to 7.8	40
Solution to 8.2	41
Solution to 8.3	42
Solution to 8.4	44
Solution to 8.5	44
Solution to 8.9	45
Solution to 8.10	45

Licensing and redistribution **47**

Solutions to the exercises

Solution to 1.6

A few things you should have observed:

- Using the `-L` flag we can change the language of the manual page shown. Most notably `man -LC command` will always show the manual page in English. Sometimes the content of the manpages is different depending on the language used. Most often the English manpage offers the best documentation.
- Different sections contain documentation about different topic, i.e.
 - section 1** Executable programs or shell commands
 - section 2** System calls (needed for Unix programming)
 - section 3** Library calls (needed for Unix programming)
 - section 4** Special files and device files
 - section 5** File formats and conventions of special files in the system
 - section 6** Games
 - section 7** Miscellaneous
 - section 8** System administration commands
 - section 9** Kernel routines (needed for Unix programming)
- For us the most important is the first section, i.e. the section documenting executables and shell commands
- By prepending the section number as first argument to `man`. E.g. try `man 2 mkdir` vs `man 1 mkdir`. Here `man 1 mkdir` gives the documentation for the `mkdir` command, the other for the system call.

Solution to 1.7

grep'ing in Project Gutenberg

- The two solutions are

```
1 < pg74.txt grep hunger | wc -l
2 < pg74.txt grep -c hunger
```

where the second one should be preferred, since it does the counting already in `grep`. This means that we need to call one program less \Rightarrow Usually better for performance.

- The options do the following:
 - `-A n` Add n lines of input after each matching line
 - `-B n` Add n lines of input before each matching line
 - `-n` Print line numbers next to each matching input line as well
 - `-H` Print file name next to each matching input line as well

-w A line only is displayed if exactly the keyword exists. Usually it is sufficient if the search string is *contained* in the line only.

- Run the command

```
1 < pg74.txt grep -nwA 1 hunger | grep -w soon
```

in order to find the numbers 8080 and 8081.

Solution to 1.8

Possible solutions are:

- Here we need to invert the file first in order for **head** to select the last 10 lines (which are now the first 10). Then another inversion using **tac** gives back the original order, i.e.

```
1 < resources/digitfile tac | head | tac
```

- The trick is to use **tail -n +2**, i.e.

```
1 tail +n2 resources/matrices/3.mtx
```

- We can use the **-v** flag of **grep** in order to invert the result, i.e. now all non-matching lines are printed:

```
1 < resources/matrices/3.mtx grep -v %
```

- Use **cut** to extract the third field and **sort -u** to get a sorted list of the values with all duplicates removed. Now piping this to **wc -l** gives the number of lines in the output of **sort -u**, i.e. the number of distinct values:

```
1 < resources/matrices/3.mtx grep -v % | cut -d "_" -f 3 | ↙  
↔sort -u | wc -l
```

- Now we need **sort** without the **-u**. We get the smallest as the first in the sorted output:

```
1 < resources/matrices/3.mtx grep -v % | cut -d "_" -f 3 | ↙  
↔sort | head -n1
```

- Running the equivalent command

```
1 < resources/matrices/bcsstm01.mtx grep -v % | cut -d "_" -f ↙  
↔3 | sort | head -n1
```

gives the result 0. Looking at the file we realise, that there is actually another, negative value, which should be displayed here. The problem is that **sort** does lexicographic ordering by default. To force it into numeric ordering, which furthermore includes the interpretation of special strings like **1E-09**, we need the flag **-g**. The correct result is displayed with

```
1 < resources/matrices/bcsstm01.mtx grep -v % | cut -d "_" -f ↙  
↔3 | sort -g | head -n1
```

- Running

```
1 < resources/matrices/lund_b.mtx grep -v % | cut -d " " -f 3 ✓
   ↪ | sort -g | head -n1
```

gives an empty output. This happens since the file contains lines like

```
1 9 8 5.5952377000000e+01
```

where there are two spaces used between 2nd and 3rd column. The problem is that `cut` splits data at *each* of the delimiter characters — `<space>` in this case. In other words it considers the third field to be empty and will take the data `5.5952377000000e+01` to be in field 4. For us this means that there are empty lines present in the output of `cut`, which `sort` first and are printed by `head`.

- Using `awk`, we would run

```
1 < resources/matrices/lund_b.mtx grep -v % | awk '{print $3}' ✓
   ↪ | sort -g | head -n1
```

which gives the correct result.

Solution to 2.1

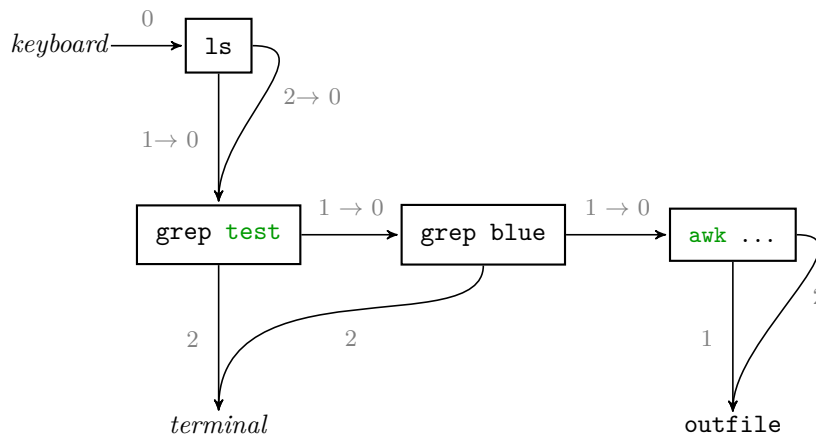
One way of doing this could be:

```
cd !$ Enter
Ctrl + R i End ← ← ← blue Enter
ls|!:3-4 Enter
mkdir !:2_red !:3_blue Enter
```

This makes $5 + 11 + 9 + 22 = 47$.

Solution to 2.2

Using the same kind of redirection diagrams as in the notes, we get



where `awk ...` denotes the `awk '{print $2}'` command.

Solution to 2.3

Exploring `tee` for logging:

- The commandline proposed does not work as intended. Error output of `some_program` will still be written to the terminal and error messages of `tee` and `grep` both reach the `log.summary` file.

This commandline, however, does work exactly as intended

```
1 some_program |& tee log.full | grep keyword > log.summary
```

here both `stdin` and `stderr` of `some_program` reach `tee` and get subsequently filtered.

- Each time the program executes, both `tee` as well as the normal output redirector `>` will cause the logfiles to be filled from scratch with the output from the current program run. In other words all logging from the previous executions is lost.

We can prevent this from happening using the `-a` (append) flag for `tee` and the redirector `>>`. Hence we should run

```
1 some_program |& tee -a log.full | grep keyword >> log.summary
```

Solution to 2.4

Some notes:

- Running `< in cat > out` is exactly like copying the file `in` to `out` as mentioned before.
- Running `< in cat > in` gives rise to the `in` file to be empty.
This is because the shell actually opens the file handles to read/write data before calling the program for which input or output redirection was requested. This means that in fact the file handle to write the output to `in` is already opened before `cat` is called and hence `in` is already at the time `cat` looks at it (because the non-appending output file handle deletes everything). Overall therefore no data can be read from `in` and thus the `in` file is empty after execution.
- `stdin` is connected to the keyboard, `stdout` and `stderr` are connected to the terminal. Therefore everything we type(`stdin` of `cat`) is copied verbatim to the terminal (`stdout` of `cat`). The shell just seems to “hang” because `cat` waits for our input via the keyboard and thus blocks the execution of further commands. `[Ctrl] + [D]` sends an “EOF” character and hence signals that there is no more input to come. This quits `cat` and returns to the command prompt.

Solution to 2.5

- For the first case, pressing `[Ctrl] + [D]` signals `cat` that end-of-file, i.e. EOF, has been reached. This causes `cat` to quit processing and since no error has occurred since startup 0 is returned. Since no data was read by `cat`, also no output will be produced at all.
- `[Ctrl] + [C]` aborts the program and hence returns a non-zero exit code. Again no data was read so no data is returned.

Solution to 2.6

- `true` is a program that — without producing any output — always terminates with return code 0.
- `false` is a program that produces no output and always terminates with return code 1.

For the first set of commands the exit code is

- 0, since `false` returns 1 and hence `true` is executed, which returns 0.
- 0, since `true` triggers the execution of `false`, which in turn triggers the execution of `true`
- 1, since `false` returns 1, so nothing else is executed and the return code is 1.
- 0, since `false` causes `true` to be executed, which returns 0. So the final `false` is not executed and the return code is 0.

Running the commandlines in the shell, we get

- 0
- 0
- 1
- 1
- 0

In a pipe sequence, the return code is solely determined by the last command executed. In other words the return code of all other commands in the pipe is lost¹

Solution to 2.7

This problem is meant to be a summary of the different types of syntax containing `&` and `|`.

- A usual pipe: The output of `echo test` on *stdout*, i.e. “test” gets piped into `grep test`, which filters for the string “test”. Since this string is contained in `echo`’s output, we see it on the terminal and the return code is 0
- Recall that `&` sends the command to its LHS into the background. So the `echo` happens, which we see on the screen. At the same time `grep test` is executed, which does not have its *stdin* connected to a file or another program’s *stdout*. In other words it has its *stdin* connected to the keyboard and it waits for the user to input data (The terminal “hangs”). Depending on what we type (or if we type anything at all) the return code of `grep` is different.
- A pipe where both the *stdout* as well as the *stderr* are piped to `grep`. The effect is the same as in the first example, since `echo` produces no output on *stderr*. I.e. we get “test” on the terminal and return code 0.
- We print once again “test” onto the terminal by executing `echo test`. Since this

¹Not entirely true ... there exists a special array variable `PIPESTATUS` which actually captures the return code of each command of the pipe. How to use it is, however, out of the scope of this course.

is successful (zero return code) `grep test` is also executed. Similar to the second case, `stdin` of `grep` is connected to the keyboard and waits for user input. The exit code of `grep` — and hence the whole commandline — depends on what is typed.

- The `echo test` prints “test” onto the terminal and since this is successful, nothing else happens.

Solution to 2.8

Possible solutions are

- The question asks explicitly to just search for the word “the”, so we need to use `grep -w`:

```
1 < pg1661.txt grep -w the && echo success || echo error
```

- We need to provide `grep` with the `-q` argument as well:

```
1 < pg1661.txt grep -wq the && echo success || echo error
```

This code executes a lot quicker, since `grep` can use a different algorithm for the search: Once it found a single match, it can quit the search and return 0.

- We need to use `grep` twice. Otherwise we get a “0” printed if there is no match:

```
1 <pg1661.txt grep -wq Heidelberg && <pg1661.txt grep -wc ✓  
↪Heidelberg || echo "no matches"
```

- The results are

word	output
Holmes	460
a	2287
Baker	42
it	1209
room	168

- We can use the command `wc -w pg1661.txt` or equivalently `< pg1661.txt wc -w` to achieve this task.

Solution to 2.9

- Since the return code of the commands to the left of `&&` or `||` determine if the command to the right is executed, we best look at the command list left-to-right as well:

- The directory `3/3` does not exist and hence the first `cd` gives return code 1.
- Therefore `cd 4/2` is executed (`||` executes following command if preceding command has non-zero return code)
- The command list to the left of the first `&&`, i.e. `cd 3/3 || cd 4/2` has return code 0 (the last command executed was `cd 4/2`, which succeeded)
- Hence `cd ../4` is executed which fails since the directory `4/4` does not exist below `resources/directories`

- In other words the list `cd 3/3 || cd 4/2 && cd ../4` has return code 1 and thus `cd ../3` gets executed
- This succeeds and thus the command to the right of `&&` is executed, i.e. we `cat` the file
- We need to suppress the error messages of the failing `cd` commands. These are `cd 3/3` and `cd ../4`. In other words the shortest commandline would be


```
1 cd 3/3 2>/dev/null || cd 4/2 && cd ../4 >/dev/null || cd ↵
      ↵ ../3 && cat file
```
- We now first make the directory `3/3`. So the execution changes slightly:
 - `cd 3/3` now succeeds and thus `cd 4/2` is not executed; we are in directory `resources/directories/3/3`.
 - The last command executed was the succeeding `cd 3/3`, such that the return code of the command list `cd 3/3 || cd 4/2` is zero.
 - We attempt to execute `cd ../4`, which fails as the dir `resources/directories/3/4` does not exist.
 - Hence we execute `cd ../3`, which is successful and “changes” the directory to `resources/directories/3/3`.
 - Finally the `pwd` command is also executed, since `cd ../3` was successful.

Solution to 2.10

We easily find out that the commands

```
1 kill time fg history pwd exit
```

have documentation which can be accessed using `help command`. This means that they are shell builtins.

Solution to 3.2

This would give a quine:

```
1 #!/bin/bash
2 cat $0
```

`3_simple_scripts/sol/quine.sh`

Solution to 3.3

The solution for the generalised version is:

```
1 #!/bin/bash
2
3 # first print everything non-comment
4 < "$1" grep -v "%" > "$2"
```

```

5
6 # now everything comment, note the append operator >>
7 < "$1" grep "%" >> "$2"

```

3_simple_scripts/sol/comment_move.sh

or alternatively

```

1 #!/bin/bash
2
3 # First copy everything *from* the second line
4 < resources/matrices/3.mtx tail -n +2 > output.mtx
5
6 # Then only the *first* line, again note the append.
7 < resources/matrices/3.mtx head -n 1 >> output.mtx

```

3_simple_scripts/sol/comment_move_alternative.sh

Solution to 3.4

Since `cat` takes data on *stdin* and copies it to *stdout* without modification, we can cache all data a script gets on *stdin* in a variable `CACHE` using the simple line

```

1 CACHE=$(cat)

```

Once this has been achieved we just use `echo` to access this data again and `grep` inside it twice:

```

1 #!/bin/bash
2
3 # Store the keyword we get as first arg:
4 KEYWORD=$1
5
6 # Read every data cat can get on stdin into the
7 # variable CACHE.
8 # Since cat's stdin gets fed from the script's stdin,
9 # this effectively reads the stdin of the script into
10 # the variable CACHE
11 CACHE=$(cat)
12
13 # Now echo the data again, i.e. transfer the data from the
14 # variable to stdin of grep.
15 # Grep for the keyword and store it in the cache again.
16 CACHE=$(echo "$CACHE" | grep "$KEYWORD")
17 #
18 # The above two commands can be done at once using:
19 # CACHE=$(grep "$KEYWORD")
20
21 # We need the so-called quoting here, i.e. the " character
22 # before and after the parameter expansion for reasons
23 # explained in the next section.
24
25 # Print the first line of the results using head:
26 echo "$CACHE" | head -n1

```

```

27
28 # Print the last line of the results using tail:
29 echo "$CACHE" | tail -n1
30
31 # Now print an empty line
32 echo
33
34 # and now all the matches:
35 echo "$CACHE"

```

3_simple_scripts/sol/grep_print.sh

Solution to 3.8

```

1 #!/bin/bash
2 # Script to extract matching lines from a few project
3 # gutenber books and show the results
4 # $1: Keyword to search for
5 #
6 cd resources
7 ILLIAD=$(<"Project_Gutenberg_selection/The_Iliad.txt" grep -i "$1")
8 YELLOW=$(<"Project\ Gutenberg\ selection/"The_Yellow_Wallpaper.txt" grep -i "$1")
9
10 cd "Project_Gutenberg_selection"
11 OTHERS=$(<"Dracula.txt" grep -H "$1"; <"The_Count_of_Monte_Cristo.txt" grep -H "$1")
12 COUNT=$(echo "$OTHERS" | grep -c ^)
13
14 echo Searching for the keyword "$1":
15 echo "    Iliad: $ILLIAD"
16 echo "    Yellow Wallpaper: $YELLOW"
17 echo We found $COUNT more findings in
18 echo "$OTHERS"

```

3_simple_scripts/sol/ex_quoting.sh

Solution to 3.9

When using `echo "$VAR" | wc -l` the results are

- 2 (correct)
- 1 (correct)
- 1 (wrong, since string is empty)

On the other hand `echo -n "$VAR" | grep -c ^` gives

- 2 (correct)
- 1 (correct)
- 0 (correct)

Therefore this method should be preferred.

Solution to 3.10

A short excerpt of an output of `ls --recursive`:

```
1 ./resources/directories/5:
2 1__2__3__4__6
3
4 ./resources/directories/5/1:
5 file
6
7 ./resources/directories/5/2:
8 file
9
10 ./resources/directories/5/3:
11 file
12
13 ./resources/directories/5/4:
14 file
15
16 ./resources/directories/5/6:
17 file
```

It shows the following features:

- Each subdirectory is denoted by the relative path to it
- For each subdirectory we get a list of files it contains.
- Most notably the path of the subdirectories always ends in a ":"

If we now assume that no file or directory contains a ":" in its name, we can `grep` for ":" in order to get a list of all subdirectories. Since by our assumption no file or dir contains a ":" we can use `cut -d: -f1` in order to get rid of the trailing ":" in the output. A final `grep` of exactly this output achieves the desired filtering function. Overall we get the script

```
1 #!/bin/bash
2
3 # filter in the output of recursive ls
4 # for a pattern
5 #
6 # overall prints those directory paths that match the pattern
7 ls --recursive | grep ":" | cut -d: -f1 | grep "$1"
3_simple_scripts/sol/recursive_ls.sh
```

Solution to 3.11

The solution makes use of the fact that `grep -n` separates the line number and the text of the matching line by a ":". So by sending the output of `grep` to a `cut -d: -f1` we can just extract the numbers of the matching lines.

```
1 #!/bin/bash
2 FILENAME="$1"
```

```

3 KEYWORD="$2"
4
5 # Grep in the file for each of the keywords
6 # use the -w flag to only match words
7 # and the -c flag to only count the matches
8 # Since grep will return with exit code 1 if no match
9 # was found, this aborts the script prematurely in this case.
10 COUNT_KEYWORD=$(grep -cw "$KEYWORD" "$FILENAME") || exit 1
11
12 # Grep in the file for each of the keywords again
13 # now use the -n flag to get the line number of the matches
14 # use the -w flag to only match words
15 #
16 # if one considers the output of grep -n, one notices, that
17 # the line numbers and the text of the line are
18 # separated by :
19 # so using cut we can only extract the line numbers:
20 LINES_KEYWORD=$(grep -wn "$KEYWORD" "$FILENAME" | cut -d: -f1)
21 # now each of the former variables contains a list of
22 # line numbers with matching text
23
24 # Now just print the data as requested
25 echo $COUNT_KEYWORD $LINES_KEYWORD
26
27 # The exit 0 is not needed, since a successfully executed script
28 # always will exit with return code 0

```

3_simple_scripts/sol/grepscript.sh

Solution to 4.1

Here we use the positional parameters \$1 to \$3 and the [command in order to achieve our goal:

```

1 #!/bin/bash
2
3 # print the first 3 arguments in reverse:
4 echo "$3_$2_$1"
5
6 # if there is help, we print some help statement:
7 [ "$1" == "-h" -o "$2" == "-h" -o "$3" == "-h" ] && echo "You ↙
   ↪asked for some help"
8
9 # alternative:
10 [ "$1" == "-h" ] || [ "$2" == "-h" ] || [ "$3" == "-h" ] && echo ↙
   ↪"-h passed"

```

4_control_io/sol/arg_reverse.sh

Solution to 4.2

We use `test` to determine the file type and take appropriate action:

```

1 #!/bin/bash
2
3 # $1 is a file and an executable:
4 [ -f "$1" -a -x "$1" ] && echo "File_$1_is_executable."
5
6 # $1 is just a plain file:
7 [ -f "$1" ] && echo "File_$1_has_size_$(<"$1" wc -c)_bytes."
8
9 # $1 is a directory:
10 [ -d "$1" ] && cd "$1" && ls

```

4_control_io/sol/test_file.sh

Solution to 4.4

The solution requires two nested `while` loops. We use `((ICNT++))` and `((JCNT++))` to increase the counter variables `ICNT` and `JCNT` by one in each iteration.

```

1 #!/bin/bash
2
3 # check if user wants help:
4 if [ "$1" == "-h" -o "$2" == "-h" -o "$3" == "-h" ]; then
5     echo "The_script_needs_two_integers_and_optionally_a_file_name"
6     exit 0
7 fi
8
9 # store parameters as I, J and File
10 I="$1"
11 J="$2"
12 FILE="$3"
13
14 # check if I and J are not negative:
15 if [ $I -lt 0 -o $J -lt 0 ]; then
16     echo "Both_$I_and_$J_need_to_be_non-negative_integers." >&2
17     exit 1
18 fi
19
20 ICNT=1 # counter for I
21
22 # loop over directories:
23 while [ $ICNT -le $I ];do
24     # create directory ICNT
25     mkdir $ICNT
26
27     JCNT=1 # counter for J
28     # loop over files:
29     while [ $JCNT -le $J ];do
30         # name of the file to generate
31         NAME="$ICNT/$JCNT"
32
33         # if the file exists, we throw an error and exit
34         if [ -f "$NAME" ]; then
35             echo "The_file_$NAME_already_exists."

```

```

36     exit 1
37 fi
38
39 # if user specified a file copy it:
40 if [ -f "$FILE" ]; then
41     cp "$FILE" "$NAME"
42 else
43     # else just create a new file
44     touch "$NAME"
45 fi
46
47 # increase JCNT by one.
48 ((JCNT++))
49 done
50
51 # increase ICNT by one.
52 ((ICNT++))
53 done

```

4_control_io/sol/while_files.sh

Solution to 4.5

A solution only implementing the 1-argument and 2-argument version of seq is:

```

1 #!/bin/bash
2
3 if [ "$1" == "-h" ]; then
4     echo "Some help"
5
6     # exit the shell with return code 0
7     exit 0
8 fi
9
10 # set the default for the 1-argument case:
11 FROMNUMBER=1 # the number at which the sequence starts
12 TONUMBER="$1" # the number until the sequence goes
13
14 if [ "$#" == 2 ]; then
15     # overwrite defaults for 2-argument case:
16     FROMNUMBER="$1"
17     TONUMBER="$2"
18 fi
19
20 # check the assumptions are not violated:
21 if ! [ $FROMNUMBER -le $TONUMBER ]; then
22     # assumption is violated => exit
23     echo "$FROMNUMBER is not less or equal to $TONUMBER" >&2
24     exit 1
25 fi
26
27 N=$FROMNUMBER
28 while [ $N -lt $TONUMBER ];do

```



```

29 echo $N
30 ((N++))
31 done
32 echo $N

```

4_control_io/sol/seq2.sh

If the 3-argument version should be supported as well we arrive at:

```

1 #!/bin/bash
2
3 if [ "$1" == "-h" ]; then
4     echo "Some_help"
5
6     # exit the shell with return code 0
7     exit 0
8 fi
9
10 #-----
11 # checking:
12
13 # set the default for the 1-argument case:
14 FROMNUMBER=1 # the number at which the sequence starts
15 TONUMBER="$1" # the number until the sequence goes
16 STEP=1
17
18 if [ "$#" -gt 3 ];then
19     # we can only consume up to 3 args:
20     echo "Extra_argument" >&2
21     exit 1
22 fi
23
24 if [ "$#" -eq 3 ]; then
25     # third arg not zero -> 3-arg case
26     # overwrite defaults accordingly
27     FROMNUMBER="$1"
28     STEP="$2"
29     TONUMBER="$3"
30 elif [ "$#" -eq 2 ];then
31     # third arg is zero
32     # (otherwise we did not get here)
33     # but second is set, hence:
34     # overwrite defaults for 2-argument case:
35     FROMNUMBER="$1"
36     TONUMBER="$2"
37     STEP=1
38 fi
39
40 # one arg case is default, but what if $1 is also empty
41 if [ -z "$1" ] ; then
42     echo "Need_at_least_one_arg" >&2
43     exit 1
44 fi
45
46 # check the assumptions are not violated:

```

```

47 if ! [ $FROMNUMBER -le $TONUMBER ]; then
48     # assumption is violated => exit
49     echo "$FROMNUMBER is not less or equal to $TONUMBER" >&2
50     exit 1
51 fi
52
53 #-----
54 # do the seq:
55
56 N=$FROMNUMBER
57 while [ $N -lt $TONUMBER ];do
58     echo $N
59
60     # do the increment STEP times:
61     # using a loop running STEP times:
62     I=0
63     while [ $I -lt $STEP ]; do
64         ((N++)) # increment our number
65         ((I++)) # also increment the I
66     done
67 done
68 echo $N

```

4_control_io/sol/seq3.sh

Solution to 4.6

A script which has the required functionality is:

```

1 #!/bin/bash
2
3 # The directory where the Project Gutenberg books are located:
4 DIR="resources/gutenberg"
5
6 if [ "$1" == "-h" ]; then
7     echo "Please supply an argument with the pattern to search for."
8     exit 0
9 fi
10
11 if [ -z "$1" ]; then
12     echo "Please provide a pattern as first arg"
13     exit 1
14 fi
15
16 #-----
17
18 # go through all gutenberg files
19 # they all end in txt
20 for book in $DIR/*.txt; do
21     # grep for the pattern
22     # we need the " because $book is a path
23     # and because $1 is user input
24     # and store the count in the variable

```

```

25 MATCHES=$(< "$book" grep -ic "$1")
26
27 # suppress books without matches:
28 [ $MATCHES -eq 0 ] && continue
29
30 # find out the number of lines:
31 LINES=$(< "$book" grep -c ^)
32
33 # print it using tabs as separators
34 echo -e "$book\t$MATCHES\t$LINES"
35 done

```

4_control_io/sol/book_parse.sh

Solution to 4.7

One way to achieve the required substitutions is to exploit word splitting. Recall that word splitting takes place at all `<tab>`, `<newline>` or `<space>` characters.

- For the first part we need a commandline that performs word splitting on the content of `VAR` and inserts a `<newline>` between each word. This can be done e.g. by executing

```

1 for i in $VAR;
2   echo $i
3 done

```

- For the second part we need to insert a `<space>` between all the words after word splitting. This is the effect of

```

1 echo $VAR

```

Note that we deliberately leave `$VAR` unquoted here.

Solution to 4.8

We need to use the `while-case-shift` paradigm in order to parse the commandline in the required way

```

1 #!/bin/bash
2
3 QUIET=0    # are we in quiet mode -> 1 for yes
4 FILE=     # variable to contain the file
5
6 while [ "$1" ]; do
7   case "$1" in
8     -h|--help)
9       echo "Help!!"
10      exit 0
11      ;;
12     -q|--quiet)
13       QUIET=1
14      ;;

```

```

15     -f)
16         shift
17         FILE="$1"
18         ;;
19     *)
20         echo "Unknown argument: $1" >&2
21         exit 1
22     esac
23     shift
24 done
25
26 # check whether the file is valid
27 if [ -f "$FILE" ]; then
28     echo "File: $FILE"
29 else
30     echo "Not a valid file: $FILE"
31     exit 1
32 fi
33
34 # if we are not quiet:
35 if [ "$QUIET" != "1" ]; then
36     echo "Welcome to this script"
37 fi
38
39 # exit the script
40 exit 0

```

4_control_io/sol/commandline_parsing.sh

Solution to 4.13

We use a few calls to `read` in order to read single lines from the script's *stdin*. Then we output the relevant lines to *stdout* or *stderr*.

```

1 #!/bin/bash
2 # read first line of stdin
3 read line
4
5 # read second line of stdin
6 read line
7
8 # read third line of stdin
9 # and print to stdout
10 read line
11 echo $line
12
13 # read fourth line on stdin
14 # and print to stderr
15 read line
16 echo $line >&2

```

4_control_io/sol/read_third.sh

Solution to 4.14

We use `read` together with the `-p` (prompt) flag in order to ask the user for two numbers and then use the code from the exercise on page 13 to do the counting and printing:

```
1 #!/bin/bash
2
3 # Ask for two numbers:
4 read -p "Enter a first number: " N
5 read -p "Enter a second number: " M
6
7 if [ $N -lt $M ]; then
8     START=$N
9     END=$M
10 else
11     START=$M
12     END=$N
13 fi
14
15 C=$START
16 while [ $C -lt $END ]; do
17     echo $C
18     ((C++))
19 done
20 echo $C
```

4_control_io/sol/read_seq.sh

Solution to 4.15

The final version of the script could look like this

```
1 #!/bin/bash
2
3 # the default values for from and to
4 FROM=1
5 TO="end"
6
7 # while-case-shift to parse arguments:
8 while [ "$1" ]; do
9     case "$1" in
10         --help)
11             echo "Script can take two arguments --from and --to."
12             echo "Each have a line number following"
13             echo "The script then moves the --from line to the --to line"
14             exit 0
15             ;;
16         --from)
17             shift
18             FROM=$1
19             ;;
20         --to)
21             shift
```

```

22     TO=$1
23     ;;
24     *)
25     echo "Unknown argument: $1" >&2
26     exit 1
27     esac
28     shift
29 done
30
31 if [ "$TO" != "end" ] && [ "$TO" -le "$FROM" ]; then
32     echo "The --to line(=$TO) needs to be larger than the --from
↪line(=$FROM)." >&2
33     exit 1
34 fi
35
36 # line count
37 LC=0
38
39 # line cache (for the line that should be moved)
40 CACHE=
41
42 # var to keep track if cache is filled or not
43 # just needed to spot errors more quickly
44 CACHEFILLED=n
45
46 # while read line to read stdin line-by-line
47 while read line; do
48     # increase line count
49     ((LC++))
50
51     # if the current line is the from line
52     # just store the line in a cache
53     if [ $LC -eq $FROM ]; then
54         # fill the cache:
55         CACHE=$line
56         CACHEFILLED=y
57
58         # no printing of this line
59         # just continue to next line
60         continue
61
62     # if TO is not "end"
63     # and it is equal to the current line number
64     elif [ "$TO" != "end" ] && [ $LC -eq $TO ]; then
65         # check first if we have something in the cache:
66
67         if [ "$CACHEFILLED" != "y" ];then
68             # this means some error
69             echo "Expected cache to be filled in line $LC" >&2
70             echo "This is not the case, however." >&2
71             exit 1
72         fi
73

```

```

74 # print the cached line
75 echo "$CACHE"
76
77 # reset state of the cache
78 # just done to spot errors more quickly
79 CACHE=""
80 CACHEFILLED=n
81 fi
82
83 # print current line:
84 echo "$line"
85 # note that quoting is needed such that
86 # characters like tab are kept and not
87 # removed by word splitting
88 done
89
90 # we still have something in the cache?
91 if [ "$CACHEFILLED" != "n" ]; then
92     if [ "$TO" == "end" ]; then
93         # just print it after everything:
94         echo "$CACHE"
95         exit 0
96     fi
97
98     # if we are getting here this means that
99     # the CACHE is still filled even though
100    # TO is a number and not "end"
101    # so TO is too large:
102    echo "The argument supplied to --to(=$TO) is not correct." >&2
103    echo "We got less number on stdin than the value given to --to"
104    ↪ " >&2
105    exit 1
106 fi
107 exit 0

```

4_control_io/sol/swap_lines_general.sh

Solution to 4.16

The solution just takes 5 lines of bash code:

```

1 #!/bin/bash
2 CACHE=
3 while read line; do
4     # insert line by line into the CACHE, but
5     # in reverse order.
6     # quoting is important here to not loose any
7     # newlines due to word splitting
8     CACHE=$(echo "$line"; echo "$CACHE")
9 done
10 # print the result: Again quoting is needed

```

```
11 echo "$CACHE"
```

4_control_io/sol/tac.sh

Solution to 4.17

We need to use a slightly modified version of `while read` line, where we pass multiple arguments to `read`:

```
1 #!/bin/bash
2 # read line by line and extract the first second
3 # and third column to BIN, BIN2 and COL.
4 # Extract all the other columns to TRASH
5 while read BIN BIN2 COL TRASH; do
6     # just print the third column
7     echo $COL
8 done
```

4_control_io/sol/mtx_third.sh

Now if we run our script, redirecting the file `resources/matrices/lund_b.mtx` to its `stdin`

```
1 < resources/matrices/lund_b.mtx 4_control_io/sol/mtx_third.sh
```

we realise that it can deal with the multiple spaces which are used in some lines to separate the columns. In other words, compared to `cut` it gives the correct result when the third column of the `mtx` files is to be extracted.

Solution to 4.18

We can achieve exactly what is asked for in a `bash` three-liner:

```
1 #!/bin/bash
2
3 # search for all files using find
4 # and process them line by line using
5 # while read line:
6 find . -type f | while read file; do
7     # now grep inside the files
8     # we use -n -H in order to keep an overview
9     # which file and which lines did match
10
11     grep -n -H "$1" "$file"
12 done
```

4_control_io/sol/grep_all.sh

Solution to 4.19

Since the directories are separated by a ":" in `PATH`, a good `IFS` to use is `..`.


```

1 #!/bin/bash
2
3 # we change the field separator to :
4 OIFS="$IFS"
5 IFS=":"
6
7 # if the user did not provide a command as first arg
8 # we complain:
9 if [ -z "$1" ]; then
10     echo "Please provide a command as first arg" >&2
11     exit 1
12 fi
13
14 # now make use of the new IFS and go through all
15 # directories in PATH
16 for dir in $PATH; do
17     # does an executable $dir/$1 exist?
18     if [ -x "$dir/$1" ]; then
19         # yes -> we are done
20         echo "$dir/$1"
21         exit 0
22     fi
23 done
24 IFS="$OIFS"
25
26 # there still has not been an executable found:
27 exit 1

```

4_control_io/sol/which.sh

Solution to 5.1

The return codes are

- 1 because the assignment $B=0$ inside the arithmetic evaluation returns zero, so running $((B=0))$ is equivalent to running $((0))$, which is C-false. Hence the return code is 1.
- 0 because we just do a simple `echo` of the last value of the arithmetic evaluation $((B=0))$, which is 0. So the command is equivalent to `echo 0`, i.e. it prints "0" onto the terminal and exits with return code 0 as well.
- 0: Here we take the output of `echo $((B=0))$` — which is "0" — and `grep` for "0" within it. This character is of course found and hence the return code is 0 again.
- 0, since -1 is nonzero, i.e. C-true.
- 1, since 0 is, well zero, which is interpreted as C-false.
- 1, since the last subexpression, i.e. 0 is, well zero.
- 0, since the last subexpression 3 is nonzero.
- 0: By just running

```

1 for((C=100,A=99 ; C%A-3 ; C++,A-- )); do echo "C:␣$C"; echo ✓

```

```
↔ "A:␣$A"; done
```

on a shell, we get the output

```
1 C:␣100
2 A:␣99
```

which means that the loop is only run once.

If we look at the 3 fields of the C-like for loop, we see that A is initialised to 99 and C to 100. After each iteration C gets increased by one and A gets decreased by one. The iteration stops if $C\%A-3$ is equal to 0 (C-false), i.e. if

$$C\%A = 3$$

This is the case *after* the first iteration, since this gives C equal to 101 and A equal to 99.

Now we know that the loop body ($(B=(B+1)\%2)$) is only executed once. Since B has not been set, it defaults to zero. Executing the statement under arithmetic evaluation hence assigns B with

$$(B + 1)\%2 = 1\%2 = 1$$

which is not C-false. Therefore the final $((B))$ returns 0, which is also the return code of the whole expression.

- 1: $((B=1001\%10))$ gives rise to no output, such that the first statement

```
1 ((B=1001%10)) | grep 4
```

fails. Note that B is assigned with $1001\%10 = 1$, however.

We continue with the second statement

```
1 ((C=$(echo "0" | grep 2)+4, 2%3 ))
```

the command substitution `echo "0" | grep 2` gives rise to no output, hence the resulting string is interpreted as zero. This means that C gets assigned to 4. The return code of the statement is determined by $2\%3$, which is 2, i.e. the return code is 0.

We proceed to execute the final statement

```
1 echo $((4-5 && C-3+B)) | grep 2
```

$4-5$ is -1 and hence C-true and $C-3+B$ gives $4 - 3 + 1 = 2$, hence also C-true. In other words $4-5 \ \&\& \ C-3+B$ is true and $\$((4-5 \ \&\& \ C-3+B))$ is the string "1". This means, however, that `grep` cannot find the character 2 in the output and overall the return code of this last expression is 1.

Solution to 5.2

If one runs the code provided here on the shell, one realises, that for proper integer numbers the result of `echo $\$(A+0)$` and the result of `echo $A` is identical. Exactly this behaviour was used in the following script:

```

1 #!/bin/bash
2
3 # store the first argument in A
4 A=$1
5
6 # check whether it is an integer by the trick
7 # we just learned about:
8 if [ "$((A))" == "$A" ]; then
9     # compute the cube and echo it
10    echo "$((A*A*A))"
11 else
12    echo "Argument $1 is not a valid integer." >&2
13    exit 1
14 fi

```

5_variables/sol/cube.sh

Solution to 5.3

One fairly naive solution is

```

1 #!/bin/bash
2 N=$1
3
4 # check if input is a positive number.
5 # note that this also checks whether the input is actually an ↙
6 ↘integer
7 # since strings that cannot be converted to an integer properly are
8 # interpreted as zero in the following arithmetic evaluation:
9 if (( N <= 0 )); then
10    echo Please provide a positive number as first argument
11    exit 1
12 fi
13 # have a loop over all integers C less than or equal to N
14 C=1
15 while (( ++C <= N )); do
16    S=1 # integer we use to test divisibility
17    isprime=1 # flag which is 1 if C is a prime, else
18              # it is 0
19    while (( ++S, S*S <= C )); do
20        # loop over all S from 1 to sqrt(C)
21        if (( C%S==0 )); then
22            # S divides C, hence C is not a prime
23            isprime=0
24
25            # break the inner loop: No need to
26            # keep looking for divisors of C
27            break
28        fi
29    done
30
31    # if C is a prime, print it

```

```

32 (( isprime==1 )) && echo $C
33 done

```

5_variables/sol/primes.sh

Solution to 5.4

The first version making use of a temporary file can be achieved like this

```

1 #!/bin/bash
2
3 # check that the argument provided is not zeros:
4 if [ -z "$1" ]; then
5     echo "Please provide a pattern as first arg" >&2
6     exit 1
7 fi
8
9 # delete the temporary file if it is still here:
10 rm -f tEMPorary_File
11
12 # create an empty temporary file
13 touch tEMPorary_File
14
15 # call book_parse.sh and analyse resulting table line-by-line
16 4_control_io/sol/book_parse.sh "$1" | while read FILE MATCH ↙
17     ↪NUMBER; do
18
19     # read already splits the table up into the 3 columns
20
21     # calculate the xi value:
22     XI=$(echo "$MATCH/$NUMBER" | bc -l)
23
24     # echo the xi value followed by a tab and the
25     # filename to the temporary file
26     echo -e "$XI\t$FILE" >> tEMPorary_File
27 done
28
29 # sort the temporary file:
30 # -n    numeric sort
31 # -r    reverse sort: largest values first
32 sort -nr tEMPorary_File | \
33 # print the three highest scoring books
34 head -n 3 tEMPorary_File
35
36 # remove temporary file again:
37 rm tEMPorary_File

```

5_variables/sol/book_analyse.sh

If we want to omit the reading and writing to/from disk, we have to do everything in one pipe. One solution for this could be

```

1 #!/bin/bash
2
3 # check that the argument provided is not zeros:

```

```

4 if [ -z "$1" ]; then
5     echo "Please provide a pattern as first arg" >&2
6     exit 1
7 fi
8
9 # call book_parse.sh and analyse resulting table line-by-line
10 4_control_io/sol/book_parse.sh "$1" | while read FILE MATCH ✓
    ↪NUMBER; do
11     # read already splits the table up into the 3 columns
12
13     # calculate the xi value:
14     XI=$(echo "$MATCH/$NUMBER" | bc -l)
15
16     # echo the xi value followed by a tab and the
17     # filename to stdout of the loop
18     echo -e "$XI\t$FILE"
19 done | \
20 # sort stdout of the loop
21 sort -nr | \
22 # filter the first three matches
23 head -n 3 | \
24 # format the output a little:
25 while read XI FILE; do
26     echo -e "$FILE\t(score:\t$XI)"
27 done

```

5_variables/sol/book_analyse_notemp.sh

Solution to 5.5

We first parse the arguments and check whether there is anything to do (if there are no numbers supplied, we are done). Then we build up the expression for `bc` in `BCLINE` and `echo` it to `bc` to get the result.

```

1 #!/bin/bash
2
3 MEAN=n # if y the mean should be calculated
4 # else the sum only
5
6 # first arg has to be -s or -m:
7 case "$1" in
8     -m) MEAN=y
9         ;;
10    -s) MEAN=n
11        ;;
12    *)
13        echo "Expected -s (for sum) or -m (for mean) as first arg" >&2
14        exit 1
15 esac
16 shift # remove first arg
17
18 if [ -z "$1" ]; then
19     # if new first arg, i.e. original second arg is empty

```

```

20 # we have no numbers on the commandline
21 # hence the result is 0 in both cases:
22 echo 0
23 exit 0
24 fi
25
26 # We build up the expression for bc in this variable:
27 # note that we know that $1 is nonzero and we can hence
28 # initialise BCLINE with it
29 BCLINE=$1
30
31 # count how many numbers we were given:
32 COUNT=1
33
34 # remove the arg we dealt with:
35 shift
36
37 # go over all other arguments
38 # one by one:
39 for num in $@; do
40 # build up BCLINE
41 BCLINE="$BCLINE+$num"
42 ((COUNT++))
43 done
44
45 # amend BCLINE if we are calculating the MEAN:
46 if [ "$MEAN" == "y" ]; then
47 BCLINE="($BCLINE)/$COUNT"
48 fi
49
50 # calculate it with bc
51 # and print result to stdout
52 echo "$BCLINE" | bc -l
53 exit $?

```

5_variables/sol/sum_mean.sh

Solution to 5.6

We have to take care to exclude both the first comment line as well as the first non-comment line from being manipulated at all. Apart from these lines all other, however, have to be touched. This script uses a so-called first-run flag and as well as the `while read` line paradigm to achieve this:

```

1 #!/bin/bash
2
3 NUM=$1
4
5 if [ -z "$NUM" ]; then
6 echo "Need a number as first arg." >&2
7 exit 1
8 fi

```

```

9
10 # read the comment line and copy to stdout
11 read line
12 echo "$line"
13
14 # initialise a firstrun flag (see below)
15 FIRSTLINE=1
16
17 # read all remaining data from stdin using grep
18 # ignore all other comment lines but parse the
19 # non-comment ones:
20 grep -v "%" | while read ROW COL VAL; do
21     # if this is the first non-comment line
22     # then it is special, we have to copy it as is
23     if (( FIRSTLINE )); then
24         FIRSTLINE=0
25         echo "$ROW_$COL_$VAL"
26         continue
27     fi
28
29     # for all other rows:
30     echo "$ROW_$COL_$(echo "$NUM*$VAL" | bc -l)"
31 done

```

5_variables/sol/mtx_multiplier.sh

Solution to 5.7

One solution is:

```

1 #!/bin/bash
2
3 # read stdin line by line:
4 while read line; do
5     # var containing the reversed line:
6     LINEREV=""
7
8     # do the reversal in a loop from
9     # I=0 to I= length of line -1
10    for ((I=0; I < ${#line}; ++I)); do
11        # the substring expansion
12        # ${line:I:1}
13        # extracts exactly the (I+1)th
14        # character from line
15        LINEREV="${line:I:1}$LINEREV"
16    done
17    echo "$LINEREV"
18 done

```

5_variables/sol/rev.sh

Another solution is:

```

1 #!/bin/bash

```

```

2
3 # read stdin line by line:
4 while read line; do
5     # do the reversal in a loop from
6     # I=0 to I= length of line -1
7     LENGTH=${#line}
8     for ((I=0; I < LENGTH; ++I)); do
9         # Use again a substring expansion
10        # but instead we use the index expression
11        #     LENGTH-I-1
12        # to access the characters from the RHS
13        # of the line to the LHS of the line
14        echo -n "${line:LENGTH-I-1:1}"
15    done
16    echo
17 done

```

5_variables/sol/rev_other.sh

Solution to 6.2

One solution is

```

1 #!/bin/bash
2
3 DIRECTORY=$1
4 echo "Directory: _____largest_file"
5 echo "-----"
6 for f in $DIRECTORY/*; do
7     # Only go through directories => skip the '$subdir'
8     # if it is not a directory
9     [ ! -d "$f" ] && continue
10
11    ( # Subshell for cd
12      cd "$f"
13
14      MAXSIZE=0
15      MAXFILE="<No_file_found>"
16      for file in *; do
17          # This time skip if '$file' is not
18          # a valid file.
19          [ ! -f "$file" ] && continue
20
21          # The filesize can be determined using
22          # wc -c == number of bytes
23          SIZE=$(wc -c "$file" | cut -f1 -d " ")
24          if [ "$SIZE" -gt "$MAXSIZE" ]; then
25              MAXSIZE="$SIZE"
26              MAXFILE="$file"
27          fi
28      done
29
30      # Print the findings

```



```

31     echo "$f:␣␣␣␣$MAXFILE"
32 )
33 done

```

6_functions_subshells/sol/largest_file.sh

Solution to 6.3

The script contains the following problems:

- Line 10: We alter the `ERROR` flag, which is checked later on to determine if the script execution is to be aborted. This change becomes lost because it happens in a subshell. We should use grouping `{ ... }` instead.
- Line 31: The accumulation of matching lines happens within the implicit subshell started by the pipe. So after the `done`, `MATCHING` is empty again. It is better to fill `MATCHING` directly by a command substitution.
- Line 39: Better use `echo -n "$MATCHING" | grep -c ^` instead of `wc -l` (See exercise).

A better version of the script would be

```

1 #!/bin/bash
2 # initial note:
3 #   this script is deliberately made cumbersome
4 #   this script is bad style. DO NOT COPY
5 KEYWORD=$1
6
7 ERROR=0      # Error flag
8 [ ! -f "bash_course.pdf" ] && {
9     echo "Please␣run␣at␣the␣top␣of␣the␣bash␣course␣repository" >&2
10    ERROR=1
11 }
12
13 # change to the resources directory
14 if ! cd resources/; then
15     echo "Could␣not␣change␣to␣resources␣directory" >&2
16     echo "Are␣we␣in␣the␣right␣directory?"
17     ERROR=1
18 fi
19
20 [ $ERROR -eq 1 ] && (
21     echo "A␣fatal␣error␣occurred"
22     exit 1
23 )
24
25 # List of all matching files
26 # VERSION1: making minimal changes:
27 MATCHING=$(ls matrices/*.mtx gutenber*/*.txt | while read line; do
28     if < "$line" grep -q "$KEYWORD"; then
29         echo "$line"
30     fi
31 done)

```

```

32
33 # VERSION2: Even more simple and more reliable
34 MATCHING=$(for line in matrices/*.mtx gutenber/*.*.txt; do
35     if < "$line" grep -q "$KEYWORD"; then
36         echo "$line"
37     fi
38 done)
39
40 # count the number of matches:
41 COUNT=$(echo -n "$MATCHING" | grep -c ^)
42
43 if [ $COUNT -gt 0 ]; then
44     echo "We found $COUNT matches!"
45     exit 0
46 else
47     echo "No match" >&2
48     exit 1
49 fi

```

6_functions_subshells/sol/subshell_exercise_corrected.sh

Solution to 6.4

We use the function `list_files` that deals with a directory and all subdirectories recursively. A little care has to be taken when printing the paths such that the “/” appears at the right places.

```

1 #!/bin/bash
2
3 list_files() {
4     # $1: prefix to append when listing the files
5     DIR="$1"
6
7     # deal with all files in current directory:
8     for file in *; do
9         # file is a regular file => list it
10        if [ -f "$file" ]; then
11            # print prepending prefix
12            echo "$DIR$file"
13        elif [ -d "$file" ]; then
14            # file is a directory:
15            # recursively call this fctn:
16            (
17                # go into subshell
18                # this keeps track of
19                # the working directory
20                cd "$file"
21                list_files "$DIR$file/"
22            )
23        fi
24        # do nothing for all other types of
25        # files
26    done

```

27 }

6_functions_subshells/sol/find_file.sh

Solution to 6.5

Instead of using one single line with all commands, we use functions to split the tasks up into logical parts and name these parts sensibly.

```
1 #!/bin/bash
2
3 # check that the argument provided is not zeros:
4 if [ -z "$1" ]; then
5     echo "Please provide a pattern as first arg" >&2
6     exit 1
7 fi
8
9 calculate_xi() {
10     # analyse the output from book_parse.sh
11     # calculate the xi values and print a table
12     # of xi values followed by a tab and the filename to stdout
13
14     while read FILE MATCH NUMBER; do
15         # read already splits the table up into the 3 columns
16
17         # calculate the xi value:
18         XI=$(echo "$MATCH/$NUMBER" | bc -l)
19
20         # echo the xi value followed by a tab and the
21         # filename to stdout of the loop
22         echo -e "$XI\t$FILE"
23     done
24 }
25
26 filter_3_largest() {
27     # filter the output of calculate_xi such that only the 3
28     # books with the largest xi values are passed from stdin
29     # to stdout
30
31     # sort stdin and filter for first 3 matches
32     sort -nr | head -n 3
33 }
34
35 print_results() {
36     # Take a table in the format produced by calculate_xi and
37     # print the rows in a formatted way
38
39     while read XI FILE; do
40         echo -e "$FILE\t(score:\t$XI)"
41     done
42 }
43
44 #-----
```

```

45
46 4_control_io/sol/book_parse.sh "$1" | \
47   calculate_xi | filter_3_largest | print_results
        6_functions_subshells/sol/book_analyse_fun.sh

```

Solution to 6.6

After the subtract operation has been implemented as well, we arrive at

```

1  #!/bin/bash
2
3  # global variable SEL to make selection between
4  # addition and multiplication
5  SEL=
6
7  #-----
8
9  add() {
10     # add two numbers
11     # $1: first number
12     # $2: second number
13     # echos result on stdout
14     echo $((($1+$2))
15 }
16
17 multiply() {
18     # multiply two numbers
19     # $1: first number
20     # $2: second number
21     # echos result on stdout
22     echo $((($1*$2))
23 }
24
25 subtract() {
26     # Subtract two numbers
27     # $1: first number
28     # $2: number subtracted from first number
29     # echos result on stdout
30     echo $((($1-$2))
31 }
32
33 operation() {
34     # selects for add or multiply depending on
35     # SEL
36     # $1: first operand for operator (add or multiply)
37     # $2: second operand for operator (add or multiply)
38     # echos the result on stdout
39
40     # this will call add if $SEL == "add"
41     # or it will call multiply if $SEL == "multiply"
42     # or subtract if $SEL == "subtract"
43     local FIRST=$1

```

```

44 local SECOND=$2
45 $SEL $FIRST $SECOND
46 }
47
48 calculate3() {
49 # it calls operation with 3 and $1
50 # such that we either add, subtract or multiply (depending on ↙
51 ↪SEL) 3 and $1
52 # echos the result on stdout
53
54 operation $1 3
55 }
56
57 map() {
58 # $1: a command
59
60 local COMMAND=$1
61 shift
62
63 # loop over all arguments left on the commandline
64 # and execute the command in COMMAND with this
65 # argument
66 for val in $@; do
67     $COMMAND $val
68 done
69 }
70
71 usage() {
72 echo "$0 [ -h | --help | --add3 | --multiply3 ] <arguments>"
73 echo "Script to do some operation to all arguments"
74 echo
75 echo "Options:"
76 echo "--add3 adds 3 to all arguments"
77 echo "--multiply3 multiplies 3 to all arguments"
78 echo "--subtract3 subtracts 3 from all arguments"
79 }
80 #-----
81
82 # $1 selects method
83
84 case "$1" in
85 --help|-h)
86     usage
87     exit 0
88 ;;
89 --add3)
90     SEL=add
91 ;;
92 --multiply3)
93     SEL=multiply
94 ;;
95 --subtract3)

```

```

96     SEL=subtract
97     ;;
98 *)
99     echo "Unknown argument: \"$1\"" >&2
100    echo "Usage:" >&2
101    usage >&2
102    exit 1
103 esac
104
105 # remove the first arg we dealt with
106 shift
107
108 # deliberately no quotes below to get rid of linebreak
109 # in the results:
110 echo $(map calculate3 $@)

```

6_functions_subshells/sol/functional.sh

It takes very little effort to add extra operators, since the script only needs to be changed at two places: We need to add the function and we need to add an extra case in order to get SEL set accordingly.

One could go even further: The functions `add`, `multiply` and `subtract` are very similar. So one could use the tool `eval` in order to write a generating function which automatically defines these aforementioned functions. Then we arrive at

```

1 #!/bin/bash
2
3 # global variable SEL to make selection between
4 # addition and multiplication
5 SEL=
6
7 #-----
8
9 generator() {
10 # function to generate a function that takes
11 # two numbers and echos the result of applying
12 # an operation to these numbers on stdout
13 #
14 # $1: name of the function to generate
15 # $2: operator to use in the operation
16 NAME=$1
17 OP=$2
18
19 eval "${NAME}() { ONE=$1; TWO=$2; echo $((ONE$OPTWO)); }"
20 }
21
22 generator "add" "+" # generate add function
23 generator "multiply" "*" # generate multiply
24 generator "subtract" "-" # generate subtract
25
26 operation() {
27 # selects for add or multiply depending on
28 # SEL
29 # $1: first operand for operator (add or multiply)

```

```

30 # $2: second operand for operator (add or multiply)
31 # echos the result on stdout
32
33 # this will call add if $SEL == "add"
34 # or it will call multiply if $SEL == "multiply"
35 # or subtract if $SEL == "subtract"
36 local FIRST=$1
37 local SECOND=$2
38 $SEL $FIRST $SECOND
39 }
40
41 calculate3() {
42 # it calls operation with 3 and $1
43 # such that we either add, subtract or multiply (depending on ↩ ✓
44 ↩SEL) 3 and $1
45 # echos the result on stdout
46
47 operation $1 3
48 }
49
50 map() {
51 # $1: a command
52
53 local COMMAND=$1
54 shift
55
56 # loop over all arguments left on the commandline
57 # and execute the command in COMMAND with this
58 # arguement
59 for val in $@; do
60 $COMMAND $val
61 done
62 }
63
64 usage() {
65 echo "$0 [ -h | --help | --add3 | --multiply3 ] <arguments>"
66 echo "Script to do some operation to all arguments"
67 echo "Options:"
68 echo "--add3 adds 3 to all arguments"
69 echo "--multiply3 multiplies 3 to all arguments"
70 echo "--subtract3 subtracts 3 from all arguments"
71 }
72
73 # -----
74
75 # $1 selects method
76
77 case "$1" in
78 --help|-h)
79 usage
80 exit 0
81 ;;

```

```

82  --add3)
83      SEL=add
84      ;;
85  --multiply3)
86      SEL=multiply
87      ;;
88  --subtract3)
89      SEL=subtract
90      ;;
91  *)
92      echo "Unknown argument: \ "$1\" " >&2
93      echo "Usage: \ " >&2
94      usage >&2
95      exit 1
96 esac
97
98 # remove the first arg we dealt with
99 shift
100
101 # deliberately no quotes below to get rid of linebreak
102 # in the results:
103 echo $(map calculate3 $@)

```

6_functions_subshells/sol/functional_generator.sh

Note, however, that `eval` is a dangerous command and should never be used on anything that contains data, which the user of your script can set. In other words: Only use it if you know what it does and how it works!

Solution to 6.7

In order to make the script from the other exercise sourceable, we just need to insert the code

```

1  return 0 &>/dev/null

```

before the `case` statement, e.g. in line 80 (of the version not using the generator). The script, which is sourceable, can be found in `6_functions_subshells/sol/functional_sourceable.sh`. Note that it still can be executed normally and runs as expected.

If we want to use `functional_sourceable.sh` in the script `6_functions_subshells/source_exercise.sh`, we need to change it slightly:

```

1  #!/bin/bash
2
3  # check first if sourced script exists:
4  if [ ! -f 6_functions_subshells/sol/functional_sourceable.sh ]; then
5      echo "This script only works if executed from the top ✓
6      ↪ directory" >&2
7      echo "of the tarball containing the solution scripts." >&2
8      echo "Please change the working directory accordingly and" >&2
9      echo "execute again." >&2
10     exit 1

```



```

10 fi
11
12 # source the other script
13 . 6_functions_subshells/sol/functional_sourcable.sh
14
15 # add 4 and 5 and print result to stdout:
16 add 4 5
17
18 # multiply 6 and 7 and print result to stdout:
19 multiply 6 7

```

6_functions_subshells/sol/source_exercise_amended.sh

Due to the relative path to the sourced script we used in this modified version of `6_functions_subshells/source_exercise.sh`, the script only works if executed from the top directory of the tarball, which contains the solution scripts.

Solution to 7.2

The matching part:

- `..` matches any string that contains any two character substring, i.e. any string with two or more letters. This is everything except `g` and the empty string.
- `^..$` matches a string with exactly two characters, i.e. `ab` and `67`.
- `[a-e]` matches any string that contains at least one of the characters `a` to `e`, i.e. `ab` and `7b7`.
- `^.7*$` matches any string which starts with an arbitrary character and then has zero or more `7`s following. This is `g`, `67`, `67777`, `7777` and `77777`.
- `^(.7)*$` matches any string which has zero or more consecutive substrings consisting of an arbitrary character and a `7`. This is `67`, `o7x7g7`, `7777` and the empty string. Note that e.g. `77777` does not match: If we “use” the pattern `.7` three times we get `^.7.7.7$` and `77777` has one character too little to be a match for this.

Solution to 7.3

The crossword:

	a?[3[:space:]]+b?	b[~eaf0-2]
[a-f][0-3]	a3	b3
[:xdigit:]+b+	3b	bb

Solution to 7.4

- ab*c or c\$ or c
- ab+c or bc\$ or bc
- ^a.*c or c\$

- d) `^_q` or `q.`
- e) `^a|w` or `....` or `r|w`

Solution to 7.5

- A single digit can be matched by the regex `[0-9]`.
- The list of digits we get by running

```
1 < resources/digitfile grep -o '[0-9]'
```

is just the list of all digits which are contained in the file `resources/digitfile` in exactly the order they occur.

- We can run

```
1 < resources/digitfile grep -o '[0-9]' | sort -n | uniq -c
```

to get the required table

```
1 00000010
2 00000071
3 00000022
4 00000033
5 00000034
6 00000075
7 00000017
8 00000039
```

In other words there is 1 zero, 7 ones, 2 twos, ..., 3 nines.

Solution to 7.6

- Regexes for the parts:
 - sign: `"[+-]"`
 - prefactor: `"[01]\.[0-9]*"`
 - exponent: `"[0-9]+"`
- So altogether the scientific numbers need to match:

```
1 ([+-]?)([01]\.[0-9]*)e([+-]?)([0-9]+)
```

where the parenthesis () are only provided to show the individual parts, i.e.

```
1 [+-]?[01]\.[0-9]*e[+-]?[0-9]+
```

would be valid as well. Executing this on the `digitfile` gives

```
$ < resources/digitfile grep -E ✓
  ↪ '[+-]?[01]\.[0-9]*e[+-]?[0-9]+'
```

```

1 1.759e+15
2 1.5e+5da_is_a_scientific_number
3 -1.34e+04

```

- Introducing the fault tolerance implies:
 - We replace the plain requirement for “e” by the bracket expansion “[eEdD]”.
 - Instead of “[01]\.[0-9]*”, we require a number with an optional decimal part, i.e. “[0-9]+(\.[0-9]*)?”

Hence overall

```
1 [+ -]?[0-9]+(\.[0-9]*)?[eEdD][+ -]?[0-9]+
```

```
$ < resources/digitfile grep -E ✓
↪ '[+ -]?[0-9]+(\.[0-9]*)?[eEdD][+ -]?[0-9]+'

```

```

1 1.759e+15
2 -9.3e-5
3 19e-5_is_not_properly_formatted_either.
4 1.5e+5da_is_a_scientific_number
5 -1.34e+04

```

Solution to 7.7

By running

```
1 < resources/matrices/bcsstm01.mtx grep -E ✓
↪ '[+ -]?[0-9]+(\.[0-9]*)?[eEdD][+ -]?[0-9]+'

```

we can easily verify that the proposed pattern gives indeed the values in the third column. As usually we get the largest of these values by piping the result to `sort -r -n | head -n1`:

```
1 < resources/matrices/bcsstm01.mtx grep -E ✓
↪ '[+ -]?[0-9]+(\.[0-9]*)?[eEdD][+ -]?[0-9]+' | sort -r | head ✓
↪ -n1

```

Solution to 7.8

The whole problem can be solved using the command lines

```

1 < resources/chem_output/qchem.out head -n48 > file
2 < file sed -r '/Q-Chem/d; s/[A-Z]\.-?//g; s/,/\n/g' | sed ✓
↪ 's/^[[:space:]]*//; /$d' | sort

```

The `sed` commands in more detail:

- `/Q-Chem/d`: Delete all lines containing Q-Chem

- `s/[A-Z]\.-?//g`: Replace all initials by nothing. Since `sed` tries to match as much as possible, the `-?` makes sure that first names with a “-” are removed completely as well. E.g.

```
1 T.-C.
```

gets replaced by the empty string by the means of two substitutions in this step.

- `s/,/\n/g`: All commas get replaced by a line break.
- `s/^[[:space:]]*//`: Replace leading whitespace by nothing, i.e. remove it.
- `/^$/d`: Remove empty lines.

Note that we need two `sed` invocations here because `sed` does not take proper note of the extra line break we introduce with the substitution `s/,/\n/g`. This can be explained as follows:

`sed` processes all rules for each line going from top to bottom, right to left. So even though we introduce new line breaks by the substitution, `sed` considers the resulting string still as a logical line and all regexes are applied to the logical line instead of the actual lines. Using such the procedure, which was suggested by the exercise, we cannot deal with this in any other way but piping it to another `sed`, which now honours the new line breaks.

Note that a careful inspection of the problem reveals that the one-liner

```
1 < file sed -r '/Q-Chem|^$/d; s/( *[A-Z]\.-? ?|, *$)//g; ↵
   ↵s/,/\n/g'
```

does the trick as well, just using a single `sed`.

Solution to 8.2

One possible solution is:

```
1 #!/bin/bash
2
3 if [ ! -r "$1" ]; then
4     echo "Cannot read file: $1" >&2
5     exit 1
6 fi
7
8 < "$1" awk '{ print $2 " " $3 " = " $2+$3 }'
           8_awk/sol/print_add.sh
```

If we execute this like

```
1 < resources/matrices/3.mtx 8_awk/sol/print_add.sh
```

or like

```
1 < resources/matrices/lund_b.mtx 8_awk/sol/print_add.sh
```

it prints the correct results, thus dealing well with the multiple separators in `resources/matrices/lund_b.mtx`.

Solution to 8.3

We use `echo` in order to transfer the numbers to `awk`, let `awk` do the computation and print the result on `stdout` (straight from `awk` itself):

```
1 #!/bin/bash
2
3 # global variable SEL to make selection between
4 # addition and multiplication
5 SEL=
6
7 #-----
8
9 add() {
10 # add two numbers
11 # $1: first number
12 # $2: second number
13 # echos result on stdout
14 echo "$1_ $2" | awk '{ print $1+$2 }'
15 }
16
17 multiply() {
18 # multiply two numbers
19 # $1: first number
20 # $2: second number
21 # echos result on stdout
22 echo "$1_ $2" | awk '{ print $1*$2 }'
23 }
24
25 subtract() {
26 # Subtract two numbers
27 # $1: first number
28 # $2: number subtracted from first number
29 # echos result on stdout
30 echo "$1_ $2" | awk '{ print $1-$2 }'
31 }
32
33 operation() {
34 # selects for add or multiply depending on
35 # SEL
36 # $1: first operand for operator (add or multiply)
37 # $2: second operand for operator (add or multiply)
38 # echos the result on stdout
39
40 # this will call add if $SEL == "add"
41 # or it will call multiply if $SEL == "multiply"
42 # or subtract if $SEL == "subtract"
43 local FIRST=$1
44 local SECOND=$2
45 $SEL $FIRST $SECOND
46 }
47
48 calculate3() {
49 # it calls operation with 3 and $1
```

```

50 # such that we either add, subtract or multiply (depending on ↙
    ↪SEL) 3 and $1
51 # echos the result on stdout
52
53 operation $1 3
54 }
55
56 map() {
57 # $1: a command
58
59 local COMMAND=$1
60 shift
61
62 # loop over all arguments left on the commandline
63 # and execute the command in COMMAND with this
64 # arguement
65 for val in $@; do
66     $COMMAND $val
67 done
68 }
69
70 usage() {
71 echo "$0 [ -h | --help | --add3 | --multiply3 ] <arguments> "
72 echo "Script to do some operation to all arguments "
73 echo
74 echo "Options:"
75 echo "--add3 adds 3 to all arguments"
76 echo "--multiply3 multiplies 3 to all arguments"
77 echo "--subtract3 subtracts 3 from all arguments"
78 }
79
80 #-----
81 # make script sourcable:
82 return 0 &>/dev/null
83 #-----
84
85 # $1 selects method
86
87 case "$1" in
88     --help|-h)
89         usage
90         exit 0
91         ;;
92     --add3)
93         SEL=add
94         ;;
95     --multiply3)
96         SEL=multiply
97         ;;
98     --subtract3)
99         SEL=subtract
100        ;;
101    *)

```

```

102     echo "Unknown argument: \" $1 \" " >&2
103     echo "Usage: \" >&2
104     usage >&2
105     exit 1
106 esac
107
108 # remove the first arg we dealt with
109 shift
110
111 # deliberately no quotes below to get rid of linebreak
112 # in the results:
113 echo $(map calculate3 $@)

```

8_awk/sol/functional_awk.sh

Solution to 8.4

The BEGIN rule initialises the variable c to zero, which is the default anyway. Therefore it can be omitted.

```

1 #!/bin/bash
2 awk '
3   # BEGIN { count=0 }
4   { count=count + 1 }
5   END { print count }
6 '

```

8_awk/sol/awk_wc.sh

Solution to 8.5

One possible solution is:

```

1 #!/bin/bash
2 awk '
3   # initialise inside_table
4   # the flag we use to keep track whether we are inside or outside
5   # of a Davidson table
6   BEGIN { inside_table=0 }
7
8   # whenever we encounter the " Starting Davidson", we
9   # change the flag to indicate that we are inside the table.
10  /^[[:space:]]*Starting Davidson/ { inside_table=1; itercount=0 }
11
12  # save current iteration number in itercount
13  # but only if the first field contains a digit.
14  # and we are not in the "Guess" step
15  $1 ~ /[0-9]/ && inside_table == 1 && $0 !~ /Guess/ { ↵
16      ↵itercount=$1 }
17
18  # here the last stored itercount is the actual number
19  # of iteration steps performed

```

```

19 # we print it and reset the flag to 0
20 /^[[[:space:]]*Davidson Summary:/ { inside_table=0; print ✓
    ↪itercount }
21 '

```

8_awk/sol/davidson_extract.sh

Solution to 8.9

The `uniq` command can be implemented like this:

```

1 #!/bin/bash
2 awk '
3     $0 != prev { print $0; }
4     { prev=$0 }
5 '

```

8_awk/sol/awk_uniq_2line.sh

or alternatively one line shorter:

```

1 #!/bin/bash
2 awk '
3     $0 != prev {print $0; prev=$0 }
4 '

```

8_awk/sol/awk_uniq.sh

A solution, which implements `uniq -c` is

```

1 #!/bin/bash
2 awk '
3     # initialise prev on first run:
4     prev == "" { prev=$0; c=1; next }
5
6     # current line is same as previous:
7     # increase counter by one:
8     $0 == prev { c++ }
9
10    # current line is not same as previous:
11    # print the statistics for previous line and
12    # reset prev and c
13    $0 != prev {print c "␣" prev; prev=$0; c=1 }
14
15    # print statistics for the last set of
16    # equal lines:
17    END { print c "␣" prev }
18 '

```

8_awk/sol/awk_uniqc.sh

Solution to 8.10

One possible solution is:


```

1 #/bin/bash
2 awk '
3 # initialise inside_block
4 # the flag we use to keep track whether we are inside or outside
5 # an excited states block
6 BEGIN { inside_block=0 }
7
8 # whenever we encounter the " Excited state ", we
9 # change the flag to indicate that we are inside the table.
10 # also we store the state number, which sits in the third field
11 /^ *Excited state [0-9]+/ { inside_block=1; state_number=$3 }
12
13 # if we find the "Term symbol" line inside the block, we store
14 # the term symbol which sits in $3 $4 and $5
15 inside_block==1 && /^ *Term symbol/ { term_symbol=$3 " " $4 " "
    ↪ "$5 }
16
17 # if we find the "Excitation energy" line, we store the ↪
    ↪ excitation energy
18 # and print the table, since we do not care about the rest of the
19 # block. Next we reset the inside_block flag for the next ↪
    ↪ block to come.
20 inside_block==1 && /^ *Excitation energy/ {
21     excitation_energy=$3
22
23     # print the data tab-separated (for analysis with e.g. cut)
24     print state_number "\t" term_symbol "\t" excitation_energy
25
26     inside_block=0
27 }
28 '

```

8_awk/sol/exstates_extract.sh

Licensing and redistribution

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/4.0/>.



An electronic version of this document is available from <https://michael-herbst.com/teaching/advanced-bash-scripting-2017/>. If you use any part of my work, please include a reference to this URL and as well cite

Michael F. Herbst. Advanced bash scripting 2017. November 2017. URL <http://doi.org/10.5281/zenodo.1045332>.