

Coursework 9 (Scala)

This coursework is worth 10%. It is about a small programming language called brainf^{***}. The first part is due on 13 December at 11pm; the second, more advanced part, is due on 20 December at 11pm.

Important:

- Make sure the files you submit can be processed by just calling `scala <<filename.scala>>` on the commandline.¹ Use the template files provided and do not make any changes to arguments of functions or to any types. You are free to implement any auxiliary function you might need.
- **Do not leave any test cases running in your code because this might slow down your program!** Comment test cases out before submission, otherwise you might hit a time-out.
- Do not use any mutable data structures in your submissions! They are not needed. This means you cannot create new `Arrays` or `ListBuffers`, for example.
- Do not use `return` in your code! It has a different meaning in Scala than in Java.
- Do not use `var`! This declares a mutable variable. Only use `val`!
- Do not use any parallel collections! No `.par` therefore! Our testing and marking infrastructure is not set up for it.

Also note that the running time of each part will be restricted to a maximum of 30 seconds on my laptop.

Disclaimer

It should be understood that the work you submit represents your **own** effort! You have not copied from anyone else. An exception is the Scala code I showed during the lectures or uploaded to KEATS, which you can freely use.

Reference Implementation

As usual, this Scala assignment comes with a reference implementation in form of two `jar`-files. You can download them from KEATS. This allows you to run any test cases on your own computer. For example you can call Scala on the

¹All major OSes, including Windows, have a commandline. So there is no good reason to not download Scala, install it and run it on your own computer. Just do it!

Write another function `write`, which takes a memory, a memory pointer and an integer value as argument and updates the Map with the value at the given memory location. As usual the Map is not updated 'in-place' but a new map is created with the same data, except the value is stored at the given memory pointer. [1 Mark]

- (3) Write two functions, `jumpRight` and `jumpLeft` that are needed to implement the loop constructs of `brainf***`. They take a program (a `String`) and a program counter (an `Int`) as argument and move right (respectively left) in the string in order to find the **matching** opening/closing bracket. For example, given the following program with the program counter indicated by an arrow:

```
--[. .+>--],>,++
      ↑
```

then the matching closing bracket is in 9th position (counting from 0) and `jumpRight` is supposed to return the position just after this

```
--[. .+>--],>,++
      ↑
```

meaning it jumps to after the loop. Similarly, if you are in 8th position then `jumpLeft` is supposed to jump to just after the opening bracket (that is jumping to the beginning of the loop):

```
--[. .+>--],>,++      jumpLeft →      --[. .+>--],>,++
      ↑                                ↑
```

Unfortunately we have to take into account that there might be other opening and closing brackets on the 'way' to find the matching bracket. For example in the `brainf***` program

```
--[. .[+]--],>,++
      ↑
```

we do not want to return the index for the '-' in the 9th position, but the program counter for ',' in 12th position. The easiest to find out whether a bracket is matched is by using levels (which are the third argument in `jumpLeft` and `jumpLeft`). In case of `jumpRight` you increase the level by one whenever you find an opening bracket and decrease by one for a closing bracket. Then in `jumpRight` you are looking for the closing bracket on level 0. For `jumpLeft` you do the opposite. In this way you can find **matching** brackets in strings such as

```
--[. . [[-]+>[.]]--],>,++
      ↑
```

for which `jumpRight` should produce the position:

```
--[. . [[-]+>[.]]--],>,++
                        ↑
```

It is also possible that the position returned by `jumpRight` or `jumpLeft` is outside the string in cases where there are no matching brackets. For example

```
--[. . [[-]+>[.]]--],>,++       $\xrightarrow{\text{jumpRight}}$       --[. . [[-]+>[.]]-->,++
      ↑                                                    ↑
```

[2 Marks]

- (4) Write a recursive function `compute` that runs a `brainf***` program. It takes a program, a program counter, a memory pointer and a memory as arguments. If the program counter is outside the program string, the execution stops and `compute` returns the memory. If the program counter is inside the string, it reads the corresponding character and updates the program counter `pc`, memory pointer `mp` and memory `mem` according to the rules shown in Figure 1. It then calls recursively `compute` with the updated data. The most convenient way to implement the `brainf**k` rules in Scala is to use pattern-matching and to calculate a triple consisting of the updated `pc`, `mp` and `mem`.

Write another function `run` that calls `compute` with a given `brainfu**` program and memory, and the program counter and memory pointer set to 0. Like `compute`, it returns the memory after the execution of the program finishes. You can test your `brainf**k` interpreter with the Sierpinski triangle or the Hello world programs (they seem to be particularly useful for debugging purposes), or have a look at

<https://esolangs.org/wiki/Brainfuck>

[2 Marks]

'>'	<ul style="list-style-type: none"> • pc + 1 • mp + 1 • mem unchanged
'<'	<ul style="list-style-type: none"> • pc + 1 • mp - 1 • mem unchanged
'+'	<ul style="list-style-type: none"> • pc + 1 • mp unchanged • mem updated with mp -> mem(mp) + 1
'-'	<ul style="list-style-type: none"> • pc + 1 • mp unchanged • mem updated with mp -> mem(mp) - 1
'.'	<ul style="list-style-type: none"> • pc + 1 • mp and mem unchanged • print out mem(mp) as a character
','	<ul style="list-style-type: none"> • pc + 1 • mp unchanged • mem updated with mp -> input <p>the input is given by <code>Console.in.read().toByte</code></p>
'['	<p>if mem(mp) == 0 then</p> <ul style="list-style-type: none"> • pc = jumpRight(prog, pc + 1, 0) • mp and mem unchanged <p>otherwise if mem(mp) != 0 then</p> <ul style="list-style-type: none"> • pc + 1 • mp and mem unchanged
']'	<p>if mem(mp) != 0 then</p> <ul style="list-style-type: none"> • pc = jumpLeft(prog, pc - 1, 0) • mp and mem unchanged <p>otherwise if mem(mp) == 0 then</p> <ul style="list-style-type: none"> • pc + 1 • mp and mem unchanged
any other char	<ul style="list-style-type: none"> • pc + 1 • mp and mem unchanged

Figure 1: The rules for how commands in the brainf*** language update the program counter pc, the memory pointer mp and the memory mem.

Part 2 (4 Marks)

I am sure you agree while it is fun to look at bf-programs, like the Sierpinski triangle or the Mandelbrot program, being interpreted, it is much more fun to write a compiler for the bf-language.

Tasks (file `bfc.scala`)

- (5) Compilers in general attempt to make programs run faster by precomputing as much information as possible before running the program. In our case we can precompute the addresses where we need to jump at the beginning and end of loops.

For this write a function `jtable` that precomputes the “jump table” for a bf-program. This function takes a bf-program as an argument and returns a `Map[Int, Int]`. The purpose of this Map is to record the information that given on the pc-position `pc` is a '[' or a ']', then to which pc-position do we need to jump next?

For example for the program

```
+++++[->+++++++<]>--<+++[->+++++++<<]>>+<<-----[+>.>.<+<]
```

we obtain the Map (note the precise numbers might differ depending on white spaces etc. in the bf-program):

```
Map(69 -> 61, 5 -> 20, 60 -> 70, 27 -> 44, 43 -> 28, 19 -> 6)
```

This Map states that for the '[' on position 5, we need to jump to position 20, which is just after the corresponding ']'. Similarly, for the ']' on position 19, we need to jump to position 6, which is just after the '[' on position 5, and so on. The idea is to not calculate this information each time we hit a bracket, but just look up this information in the `jtable`.

Then adapt the `compute` and `run` functions from Part 1 in order to take advantage of the information stored in the `jtable`. This means whenever `jumpLeft` and `jumpRight` was called previously, you should look up the jump address in the `jtable`. Feel free to reuse the function `jumpLeft` and `jumpRight` for calculating the `jtable`. [1 Mark]

- (6) Compilers try to eliminate any “dead” code that could slow down programs and also perform what is often called *peephole optimisations*.³ While it is difficult for compilers to comprehend what is intended with whole programs, they are very good at finding out what small snippets of code do, and then try to generate faster code for such snippets.

³https://en.wikipedia.org/wiki/Peephole_optimization

In our case, dead code is everything that is not a bf-command. Therefore write a function `optimise` which deletes such dead code. Moreover this function should replace every substring of the form `[-]` by a new command `0`. The idea is that the loop `[-]` just resets the memory at the current location to 0. It is more efficient to do this in a single step, rather than incrementally as in the original bf-programs.

In the extended `compute3` and `run3` functions you should implement this command by writing 0 to `mem(mp)`, that is use `write(mem, mp, 0)` as the rule for the command `0`. The easiest way to modify a string in this way is to use the regular expression `""["^<>+-. ,\[\]]""`, which recognises everything that is not a bf-command. Similarly, the regular expression `""\[- \]""` finds all occurrences of `[-]` and by using the Scala method `.replaceAll` you can replace it with the string `"0"` standing for the new bf-command. [1 Mark]

- (7) Finally, real compilers try to take advantage of CPUs which often provide complex operations in hardware that can combine many smaller instructions into a single faster instruction.

In our case we can optimise the several single increments of a memory cell, for example `++++`, by a single “increment by 4”. For this optimisation we just have to make sure these single increments are all next to each other. Similarly optimisations should apply for the bf-commands `-`, `<` and `>`, which can all be replaced by extended versions that take the amount of the increment (decrement) into account. We will do this by introducing two-character bf-commands. For example

original bf-cmds	replacement
<code>+</code>	<code>+A</code>
<code>++</code>	<code>+B</code>
<code>+++</code>	<code>+C</code>
<code>...</code>	<code>...</code>
<code>+++... ++</code>	<code>+Z</code>
(these are 26 +’s)	

Similarly for `-`, `<` and `>`. If there are more than 26 +’s in a row, then more than on “two-character” bf-commands need to be generated (the idea is that more than 26 copies of a single bf-command in a row is a rare occurrence in actual bf-programs). All other bf-commands should be unaffected by this change.

For this write a function `combine` which replaces sequences of repeated increment and decrement commands by appropriate two-character commands. In the functions `compute4` and `run4`, the “combine” and the optimisation from (6) should be performed. Make sure that when a two-character bf-command is encountered you need to increase the `pc`-counter by two in order to process the next command. For example


```
combine(optimise(load_bff("benchmark.bf")))
```

generates the improved program

```
>A+B[<A+M>A-A]<A[[ ...
```

for the original benchmark program

```
>++[<+++++++>-]<[[ ...
```

As you can see, the compiler bets on saving so much time on the +B and +M steps so that the optimisations is worthwhile overall (of course for the >A's and so on, the compiler incurs a penalty). Luckily, after you have performed all optimisations in (5) - (7), you can expect that the `benchmark.bf` program runs four to five times faster. You can also test whether your compiler produces the correct result by for example testing

```
run(load_bff("sierpinski.bf")) == run4(load_bff("sierpinski.bf"))
```

which should return true for all the different compiler stages.

[2 Marks]