INFOMCPD Assignment 0: Getting started with Haskell

Tom Smeding, Gabriele Keller

December 18, 2024

1 Installing the Haskell compiler

To do the assignments it is most convenient to be able to develop on your own computer. The best way to install GHC, the primary Haskell compiler, is to use GHCup: go to https: //www.haskell.org/ghcup/ and follow the instructions for your platform. During the GHCup installation, be sure to install GHC and Cabal, but not Stack¹; if you want interactive compiler feedback in your editor, also install HLS (the Haskell Language Server). (For VSCode there is a Haskell plugin that works well with HLS.)

If the GHCup installation does not work, first ask us (we might know what is going wrong and be able to fix it). If we cannot fix it, an alternative is to install Stack using the instructions at https://docs.haskellstack.org/en/stable/; if you do this, you will need the stack instructions instead of the cabal instructions in the rest of this document.

On Windows, a common cause for GHCup installation failure is overly aggressive anti-virus software; this may be the cause if the initial part of the script seems to work but downloading the actual software later fails. If you run into this, a fix is probably to remove your current anti-virus and install Microsoft Defender instead.

Haskell tutorial: the rest of this assignment will assume that you have some basic knowledge of how Haskell works, or at least that you know how to look things up. A good tutorial is http://learn.hfm.io/; note that there is a small amount of overlap with the end of that tutorial and this assignment.

For debugging tips, see Section 7 at the end of this document.

2 Submission and grading

The submission system for this assignment is at: https://domjudge.uu.nl/cpd/domjudge. You should log in with your Solis account as usual.

- The sub-assignments (0.1, 0.2 and 0.3) have **different deadlines**, but these deadlines are **not** enforced by DomJudge. Please carefully check the deadline of each assignment, and make sure to submit before the deadline.
- Please test your program before submitting: the automatic grading system is not a substitute for making sure your code works correctly. Thinking about useful test cases helps understanding what is really going on in your code.
- Nevertheless, the number of attempts you make does not influence your grade; all that matters is that you eventually (but before the deadline of the corresponding assignment)

¹While installing Stack is not harmful at all, Cabal and Stack have roughly the same purpose and having both can be confusing at times, especially with language-server-based editor support.

get a CORRECT submission.

This assignment has three parts: 0.1, 0.2 and 0.3. All three parts are mandatory:

- Assignment 0.1 (Section 4): 4.5% of your final grade (due 2024-11-29 23:59)
- Assignment 0.2 (Section 5): 4.5% of your final grade (due 2024-12-06 23:59)
- Assignment 0.3 (Section 6): 6% of your final grade (due 2024-12-13 23:59)

All together, assignment 0 (this document) is 15% of your final grade, giving you a 1.5 grade if you would do nothing else in the course. Equivalently, assignment 0 is worth 30% of your practical grade, which is in turn 50% of your final grade.

You get the points for a part if you submit a working program that conforms to the specification. Programming style is not graded, though you can of course ask for feedback on style and approach during the lab sessions.

Why? This assignment is intended to let you get somewhat familiar with Haskell in a simple setting, so that the larger assignment later in the course will allow you to focus more on the content and less on learning the language. We made it mandatory on purpose: if you do not know Haskell well yet, it is a good exercise or a good refresher; if you are already adept at Haskell, it should take you very little time. Hence, in neither case do you have an excuse. :)

If you use generative AI tools to help you in this assignment, or any assignment in this course, the responsibility is on *you* to ensure that you actually learn all you can from the assignment. For skills (and programming is a skill), we humans learn much better by *doing* than by reading passively. If you know how to do something yourself, you are much better equipped to judge attempts by others.

3 Reverse Polish Notation (RPN)

Since the first assignment (Section 4) is about arithmetic expressions in postfix notation, also called Reverse Polish Notation, here is a brief introduction about how those work. Normal arithmetic expressions have operators like + and * between their arguments; this is called *infix* notation. In infix, the sum of 1 and 2 is written 1 + 2, and the product of that and 3 is written (1 + 2) * 3. Note that we needed to use parentheses to disambiguate the order of operations; without parentheses and without precedence rules, it would be ambiguous whether 1 + 2 * 3 meant (1 + 2) * 3 or 1 + (2 * 3).

In postfix notation, the operators are written *after* their arguments. For example, the sum of 1 and 2 is written $1\ 2\ +$, and the product of that and 3 is written $1\ 2\ +\ 3\ *$. Note that we do not need to use parentheses here; indeed, looking from the back, the multiplication * needs two arguments. The second argument is clearly the 3, and the first argument must then be the next previous thing — the addition of 1 and 2.

An alternative interpretation of postfix notation is by evaluation using a stack. For example, evaluation of the expression $2 \ 3 \ 7 + 4 \ * +$ would proceed as follows:

Initial stack:			
After 2:	2		
After 3:	2	3	
After 7:	2	3	7
After +:	2	10]
After 4:	2	10	4
After * :	2	40]
After +:	42		

You will not need to write a stack-based evaluator in this style, but perhaps this helps to clarify the meaning of a postfix expression.

4 Assignment 0.1: RPN

You are going to write an interpreter for postfix (RPN) expressions. You will only have to support non-negative integers, and only addition (+) and multiplication (*). In addition to the examples given in Section 3, here are some more RPN expressions that you will need to be able to handle:

Input expression	Evaluation	
12 3 * 6 +	$(12 \cdot 3) + 6 = 42$	-
1 2 3 4 5 + + + +	1 + (2 + (3 + (4 + 5))) = 15	-
123 0 + 2 2 * +	$(123+0) + (2 \cdot 2) = 127$	
123 0+2 2* +	$(123+0) + (2 \cdot 2) = 127$	- You need to handle this too

Your program will be split in three parts:

- 1. You will write one function to split the input into tokens: numbers, plus signs and multiplication signs.
- 2. Then you will write a function to parse the input expression (in *concrete syntax*) to an internal representation (the *abstract syntax*).
- 3. Finally, you will write a function that evaluates the expression in this internal, abstract representation to its integer result.

Due to the simplicity of this task, going through the intermediate representation might feel redundant; please follow this two-part structure anyway, because it allows you get practice with some relevant concepts in this course (and because you can re-use some parts in assignment 0.2).

4.1 Setting up your development environment

Create a file called, for example, RPN.hs. Open it in your editor, and to the side of your editor, open a terminal (or cmd on Windows) in the directory where your file is stored. Then type ghci RPN.hs in the terminal (if you installed via Stack instead of GHCup, type stack repl RPN.hs). Make sure that opens correctly and asks for input with a prompt that looks like ghci> or *Main>.

Next, put foo = "hello!" your file, and type :r (followed by Enter) in your terminal (i.e. in ghci). It should reload the module, and afterwards you should be able to type foo (and Enter), which will show "hello!". You can remove the definition of foo from the file now that you have verified that ghci works.

Recommendation: Put the following line all the way at the top of your program:

{-# OPTIONS -Wall #-}

This enables warnings in ghci, such as when your pattern matches are not exhaustive (i.e. when they do not cover all possibilities). This is very helpful while developing.

4.2 Lexing

Note: There is a template file for this assignment here: https://utrechtuniversity.github. io/infomcpd/assignments/asg-0-1-template.hs

Important: Even if you don't use the template, you must start your file with module RPN where. This module name is necessary for grading to proceed correctly.

Splitting up the input string into tokens is often called *tokenisation* or *lexing*. Define a data type called **Token** that has three constructors: one for numbers (with an **Int** field), one for plus signs and one for multiplication signs. The second and third constructor of your **Token** type do not need any fields; this is almost like an enumeration type.

Don't forget to make your Token type showable! You do this by appending deriving (Show) after your data type definition, like so:

```
data Token = ...
deriving (Show)
```

This ensures that you will have a *much* nicer debugging experience in ghci.

After you have defined Token, define a function called lexer with the type signature String -> [Token] that will take the input string and produce a list of tokens. If some invalid input is detected, the function is allowed to error out using error; you do not need to do proper error handling in this assignment.

Be sure to use pattern matching to walk over the input, and recursion to parse "the rest" of the input after you have parsed one token. Functions that may come in handy:

- isDigit :: Char -> Bool ² (you need to import Data.Char for this)
- read :: String -> Int ³; actually this function has a much more general type, but you can use it at this type as well. If you want to try read in ghci, write something like read "123" :: Int to constrain its return type to Int.
- span :: (a -> Bool) -> [a] -> ([a], [a]) ⁴: this function returns the prefix of the given list that satisfies the predicate, as well as the rest of the list.

You can test your function in ghci (remember to type :r first to reload the file): for example, try entering lexer "1 2 +" in ghci, or see what error is thrown on invalid input such as in lexer "1!" or lexer "abc".

ghci> lexer "1 2 +"
[TNum 1,TNum 2,TPlus]
ghci>

(You do not need to name the constructors of your Token data type like this — though these are common choices.)

4.3 Parsing

After having split the input into tokens, we will now parse them into an expression tree. Define a data type called Expr with three constructors: one for numbers again, and two for addition

²https://hackage.haskell.org/package/base-4.17.2.1/docs/Data-Char.html#v:isDigit

³https://hackage.haskell.org/package/base-4.17.2.1/docs/Prelude.html#v:read

⁴https://hackage.haskell.org/package/base-4.17.2.1/docs/Prelude.html#v:span

and multiplication, each with two fields of type Expr. This type is an abstract representation of an arithmetic expression: there is no concept of parentheses or any kind of syntax anymore. Furthermore, it is unambiguous: there is no ambiguity in the order of operations like we had with (1 + 2) * 3 versus 1 + (2 * 3).

The actual parsing will be done by a helper function⁵ parser' :: [Expr] -> [Token] -> Expr, where the [Expr] argument is a *stack* of expressions, and the [Token] argument is the list of tokens still to parse. The parsing model is similar to the evaluation example in Section 3, except here the stack contains expressions instead of integers. The top-level parsing function, parser :: [Token] -> Expr, will simply initialise the stack with the empty stack, [], and delegate the rest to parser'.

In **parser'**, recurse over the input token list, pattern matching on (the top of) the stack whenever you need some expressions from it (i.e. when you encounter the tokens for plus signs and multiplication signs). You can push something on the stack simply by passing the enlarged stack to the recursive call to **parser'**.

Again, test your function (or some half-done version of your function) in ghci by running it on some inputs (and see what error is thrown on invalid input such as parser (lexer "1 +") or parser (lexer "1 1")). You can either go through your lexer and run parser (lexer "1 2 +") or you can pass a list of tokens directly: the list notation that you get from running lexer "1 2 +" is valid Haskell syntax and can be typed in directly as an argument to parser.

4.4 Evaluation

After you have written the parser, all that is left is to evaluate the arithmetic expression in Expr. For this, write a recursive eval function of type Expr -> Int. This function will be quite short.⁶

A note: in a way, this eval function follows the *big-step semantics* of these simple arithmetic expressions. What exactly this means you will learn later on in the course.

After this, you have two data types (Token and Expr) and three functions (lexer, parser and eval). (Actually, you also have parser', but that is simply a helper function for parser.) With that, this task is done and you can submit your file in DomJudge.

4.5 Submission

Sign in to DomJudge (via the link in Section 2), and click 'Submit' at the top of the screen. Select your file (RPN.hs for assignment 0.1), and be sure to select the correct problem **and** language. DomJudge will automatically select "Compile for 0.1" for you, but it will do so regardless of which "problem" you select; remember to set both fields to the correct value.

If you submission is correct, you will see the result CORRECT. If you get COMPILER-ERROR, that means that either your code did not compile, or your submission was incorrect. (We use the "compiler output" in the system to give feedback when your submission gives the wrong answer.) Click the submission in the table to see the feedback, if any.

If you get WRONG-ANSWER, you (somehow) got through all the public test cases but still failed one of the additional, private test cases that we check your submission with. Test your code thoroughly, and if you can't figure out what you're doing wrong, ask in the lab sessions.

⁵We pronounce **parser**' as "parser-prime".

⁶For the mathematicians in the room: eval is a nice example of a catamorphism.

5 Assignment 0.2: Infix

You received complaints from your users that they want to write expressions in infix form, not in prefix form. So you get to work.

Note: There is a template file for this assignment here: https://utrechtuniversity.github. io/infomcpd/assignments/asg-0-2-template.hs

Important: Start your file with module Infix where — the module name 'Infix' is necessary for grading to proceed correctly.

You will update the lexer and the parser to now parse infix expressions with non-negative numbers, addition, multiplication and parentheses. For this, your Token type will need to be expanded as well. However, the expression returned from your new parser will be expressed in the same Expr data type as in Assignment 0.1; indeed, prefix versus infix is just a choice of *concrete syntax*. The *abstract syntax*, here the Expr data type, does not change, and neither does the *semantics*, your evaluator eval.

Name your updated token type **TokenInfix**, and name your updated lexer and parser functions as follows:

```
lexerInfix :: String -> [TokenInfix]
parserInfix :: [TokenInfix] -> Expr
```

You must handle operator precedence (order of operations): 1 * 2 + 3 should parse the same as (1 * 2) + 3, and 1 + 2 * 3 should parse the same as 1 + (2 * 3). However, the operators do not need to be left-associative: you can parse 1 + 2 + 3 + 4 either as ((1 + 2) + 3) + 4 or as 1 + (2 + (3 + 4)); both are allowed (because it does not matter for + and *).

Input expression	Result
3 + 12	15
3*12	36
3 * (12 + 2)	42
(3 * 12) + 2	38
3 * 12 + 2	38
12 + 3 * 6	30
1 + 2 + 3 + 4 + 5	15
1 + 2+ 3 *4+5	20
(1 + ((2) * 7)) + 3)*(4+5)	162

Some examples that you can test your code on:

Your previous lexer and parser functions are allowed to remain in your submission for assignment 0.2, but they don't have to. Submission works the same as for 0.1; remember to set both the problem *and* the "language".

6 Assignment 0.3: Lambda

Your users are again complaining; they wanted a programming language, not a calculator. So you get to work again. This time, you will be writing a very simple interpreter for the lambda calculus with numbers, addition and multiplication — but as before, you are going to give your users exactly what they asked for, and nothing more. Indeed, your users did not ask for an *easy to use* programming language.

This assignment is separate from the last two; it will be a single file again, but you should start from scratch.

Note: There is a template file for this assignment here: https://utrechtuniversity.github. io/infomcpd/assignments/asg-0-3-template.hs

Important: Start your file with module Lambda where — the module name 'Lambda' is necessary for grading to proceed correctly.

The syntax of your language is as follows:

 $\frac{v \operatorname{Var} e \operatorname{Expr}}{#ve \operatorname{Expr}} \operatorname{LAMBDA} \qquad \frac{e_1 \operatorname{Expr} e_2 \operatorname{Expr}}{`e_1 e_2 \operatorname{Expr}} \operatorname{APP} \qquad \frac{v \operatorname{Var}}{v \operatorname{Expr}} \operatorname{Var} \qquad \frac{n \operatorname{Num}}{n \operatorname{Expr}} \operatorname{Num}$

where the Var judgement matches on a sequence of 1 or more letters (a–z A–Z) and the Num judgement matches on a sequence of 1 or more digits (0-9).

For example, the following example programs have the listed meaning in normal lambda calculus notation:

Input expression	Meaning in normal notation
fx	f x
``f x y	f x y
`f`gx	f (g x)
# x ``add x 1	$\lambda x. \text{ add } x \ 1$
#x``add x 1	$\lambda x. \text{ add } x \ 1$
` #x ``add x 1 10	$(\lambda x. \text{ add } x \ 1) \ 10$
`f #x ``add x 1	$f (\lambda x. \text{ add } x 1)$
` #f `f 10 #x ``mul 2 x	$(\lambda f. f \ 10) \ (\lambda x. \ \mathrm{mul} \ 2 \ x)$
`#inc `inc 10 #x ``add 1 x	($\lambda inc. inc 10$) ($\lambda x. add 1 x$)

And some more examples, all of which evaluate to an integer:

Input expression	Result
`` add 2 3	5
`` mul 11 3	33
`` mul 6 `` add 3 4	42
` #inc `inc 10 #x ``add 1 x	11
``#r#t```r t`add 1 0#f```f f f f#f#x`f`f x	65536

Split your implementation into three parts: a lexer (that splits up the input String into a list of tokens, [Token]), a parser (that parses a list of tokens [Token] into an expression Expr), and an evaluator (that evaluates an Expr to a Value). Note that because not all expressions evaluate to integers any more in this language, the evaluator cannot simply return an Int.

Instead, define a data type Value with two constructors, one for integers and one for functions. This data type looks as follows:

data Value = VInt Int | VFun (Value -> Value)
 deriving (Show)

You must use this data type — though you are allowed to change the constructor names VInt and VFun if you want.

The idea of this data type is that an expression that results in an integer should evaluate to a VInt, and an expression that results in a function should evaluate to a VFun. A VFun contains a *function* value that takes a Value representing the argument to the function, and returns the Value that this function evaluates to when passed that argument. For example, when lexing, parsing and then evaluating the program $#x \ add x 1$, your evaluator should return a VFun value that contains a function f for which f (VInt 7) returns VInt 8.

You will notice that if you put deriving (Show) on the data type definition, the Haskell compiler will reject your code, saying that there is no instance of Show defined for Value -> Value. Indeed, by default, Haskell defines no method for printing functions, as that is kind of hard to do in general. However, you can cop-out by writing:

import Text.Show.Functions ()

which adds a dummy instance that shows any function value as <function>. With that, deriving (Show) on Value should be accepted.

6.1 The evaluator

Since this language has support for variables, the evaluator has a more interesting type than the one in assignments 0.1 and 0.2. Import the following modules:

import Data.Map.Strict (Map)
import qualified Data.Map.Strict as Map

This allows you to use the Map type in your code, with the Map operations written as for example Map.insert. The documentation for this type and the supported operations can be found at https://hackage.haskell.org/package/containers-0.6.8/docs/Data-Map-Strict.html. Particularly relevant operations are empty, insert and lookup⁷.

A value of type Map k v is a finite map from keys of type k to values of type v — also known as dictionaries, tree-maps, hashmaps (although in this case they are in fact tree-maps, not hash maps), associative arrays, etc. You will store the variables currently in scope, together with their values, in a Map in your evaluator.

Define the core of your evaluator with this type:

eval' :: Map String Value -> Expr -> Value

and then define a (very short) wrapper that initialises evaluation with an empty environment:

eval :: Expr -> Value

The top-level interpreter. Finally, define the following function:

interpret :: String -> Value

that composes your lexer, your parser and your evaluator. The **interpret** function is the function that will be evaluated when the grading your submission for Assignment 0.3. Again you may assume that all input is valid, i.e. you may throw errors using **error** if you do encounter invalid input.

6.2 Hints

If you have trouble writing the case for lambda expressions in eval: note that the result of a lambda expression is surely a function, not an integer. What function? Well, it should do whatever the body of the lambda says. Fortunately, it turns out you have a convenient function lying around that implements the "meaning" of an expression in a certain environment.

⁷Consume the result of Map.lookup using a case expression.

7 Debugging

The functional purity of Haskell is nice for application and algorithm development, but for debugging it can be a royal nuisance. Thus there exists the module Debug.Trace⁸. The basic function is trace :: String -> a -> a, where trace s x (if its result is demanded — this is relevant because of laziness) first evaluates x, then prints s to standard output as a side effect, and then simply returns x.

If you want to use trace to trace all calls to a certain function, a typical pattern is the following:

The call trace ("fibo " ++ show n) False will always evaluate to False, but not before the trace message is printed; so the undefined is never run, but the message is printed anyway.

⁸https://hackage.haskell.org/package/base/docs/Debug-Trace.html