

# BU CAS CS 320, Fall 2009: Discussion Lab Notes

Andrei Lapets

December 6, 2009

# 1 Motivation

The major goal of this course is not just to become familiar with the diversity of modern programming practices, paradigms, and language features, but to learn ways to decide which of these are most appropriate for a given situation.

## 1.1 Overview of Different Programming Paradigms

Table 1 lists many of the features found in modern programming languages.

Table 1: Features of programming languages.

	C	C++	Java	C#	PHP	Python	Scheme	OCaml	ML	Haskell	SQL	Prolog
imperative												
object oriented												
garbage collection												
functional												
type system												
no side effects												
declarative												

Table 2 lists arguably suitable areas of application of the major language categories. It is important to note that this is *not* reflected in practice because the selection of a language for a given problem is often influenced by a variety of other factors. Thus, the table should be treated more as an informed editorial on the question.

Table 2: Well-suited application areas of programming language families.

	C/C++	Java/C#	Python/Scheme/Haskell/ML	SQL	Prolog
device drivers					
I/O and memory manipulation					
simulations					
user interfaces					
parallel programming					
parsing/data manipulation					
data description/representation					
purely mathematical applications					
fast prototyping					
ontologies					

It is particularly necessary at the current time in history to be well aware of the many characteristics of modern programming languages. While many of these characteristics were developed as early as the 1960s, it is only now that many of them have become feasible to implement and use in practice (due to improvements in hardware, such as for garbage collection) as well as vitally necessary (due to the advent of multi-user and multi-process operating systems, and the recent increases in scale of both raw data and distributed processing capability that have occurred with the advent of the internet and cloud computing).

## 2 Introduction to Haskell Syntax and Semantics

Recall that interfaces of modern programming languages (such as Java, Python, or Haskell [Je99]) typically consist of several components from the following (incomplete) list:

- Values and constants (no side effects, no evaluation)
- Expressions (no side effects)
- Statements (side effects, nodes on control flow graph)
- Definitions
- Modular collections of code (classes, modules)

Often, some of the above items are subcategories of other items (e.g. in Python both definitions and expressions are statements, and in C/C++ individual variable definitions are both statements and expressions).

### 2.1 Values and Expressions

Fundamentally, the Haskell language interface is organized into *values* and *expressions*, and code is organized at a high level into *modules* of *definitions* (which we will often call *equations*).

Table 3: Haskell values and expressions.

Values (constants, constructors)	Expressions	Definitions ( <i>name pattern = expression</i> )
5	1 + 4 * 2	x = 3
[1,2,3]	[x,y,z] ++ [a,b,c]	f x y = x*x + y*y
[]	f (1,2)	g (x,(y,z)) (w:ws) = x:y:z:w:ws
1:2:[]	if y == 2 then [] else [1]	
False	x	

Values are pieces of data in their irreducible form: once a value is obtained, no further work needs to be done. For example, constants, such as 2, as well as data constructors like `False`, `(:)` and `[]`, are values. Expressions, on the other hand, or some sort of computation process that may not yet be complete, so they may contain values, variables, and operators. Notice that the definition of an expression includes values: *all values are expressions*.

We now go over more detailed and exact definitions of the various concepts. Note that our definitions will ignore type restrictions for the time being. These are technically correct definitions of the concepts, we merely omit for the moment the requirement that they be *well-typed*.

#### 2.1.1 Values

Before values are defined, it is necessary to define *constructors*. Constructors can be viewed as building blocks. They are not computations, but descriptions of the shape of a piece of data. The following constructors for tuples and lists are built-in:

`(:)`, `[]`, `( , )`, `( , , )`, `( , , , )`, etc.

Notice that there is essentially an infinite number of tuple constructors, one for each possible number of components. It is important to note that the explicit list notation `[ , ... , ]` is essentially just a macro for the `(:)` constructor. The user can define her own constructors, and we cover this later. Table 4 defines Haskell values. Notice that the definition is recursive: there are several base cases, and one recursive case.

Table 4: Definition of values.

value	examples
numeric literals	1, 4.813
character literals	'#', 'a'
string literals (lists of characters)	"abc"
constructors that take no arguments	True, False, []
constructors that take arguments applied to values	(1, (2,3)), (:) 1 [], Left 3, Just "str"

### 2.1.2 Patterns

Modules consist of a list of equations of the form

$$\textit{name pattern} = \textit{expression}$$

The pattern is optional, but when it is present, it must satisfy the definition in Table 5.

Table 5: Definition of patterns.

pattern	examples
values	<i>see Table 4</i>
variables	x, test
constructors that take arguments applied to patterns	(:) x ((:) 2 []), (x, True, 2)

### 2.1.3 Expressions

Table 6: Definition of expressions.

expression	examples
values	<i>see Table 4</i>
variables	x, test, (+), (  ), not
constructors	True, False, Just, (:), []
parenthesized expressions	(e)
application of one expression to another	f 2, ((+) 2) 3
tuples of expressions	(e1, e2, e3)
conditional	if e1 then e2 else e3
local variable binding ("let")	let x = e1 in e2

Expressions consist of variables, values, constructors, and application of expressions to other expressions. Built-in operators are essentially just pre-defined variables. There are also three additional built-in expression forms (not essential for Turing completeness): tuples, conditional expressions, and expressions that allow binding of local variables.

Notice that application only allows for one function, and one variable. Thus, functions need to be applied to one variable at a time. Suppose that `f` takes three variables `x`, `y`, and `z`. Then when we write `f x y z` this is actually:

$$((f\ x)\ y)\ z$$

Any syntactic sequence of expressions separated by spaces is assumed to be left-associated application. Thus, if we want to write `f (g x)` we cannot write `f g x`, as that will be interpreted as `(f g) x`.

## 2.2 Evaluation

The process of turning expressions into values (i.e., computations into results) is called *evaluation*, and it's the job of the Haskell interpreter to take a single expressions as input and to return a corresponding value. Evaluation works in the usual way for expressions that contain numbers, arithmetic operators (`+`, `-`, etc.), boolean constants (`True`, `False`), and boolean operators (`==`, `<`, and so on).

## 2.3 Collections of Equations

A file of Haskell code (called a *module*) can be viewed as a set of equations relating variables to expressions. This makes Haskell a *declarative* language, since a module is simply a list of declarations about how variables correspond to expressions. Thus, an example of a module might look as follows:

```
module Simple
  x = 2
```

The Haskell interpreter uses this collection of equations to evaluate expressions with variables. If the interpreter runs into a variable while evaluating an expression, it finds the first equation in the module (starting from the top) whose *left*-hand side matches the variable, and replaces it with whatever expression is on the right-hand side of the equation.<sup>1</sup> This process is sometimes called *pattern matching*.

For example, if a user loads the `Simple` module and types `x` into the interactive prompt, the interpreter evaluates this to `2`. If the result of this sort of lookup is still an expression, the interpreter continues the evaluation process (evaluating operators, performing lookups). Of course, this process is not guaranteed to terminate because any expressions on the right-hand side can also contain variables.

**Exercise 2.1.** Change the module definition of `Simple` so that the user can provide the interpreter with an input that will cause it to run forever.

Of course, the left-hand sides of the equations in a module can be more sophisticated:

```
module Simple2
  f 1 = 0
  f x = x + x
```

If the left-hand side contains a variable followed by something else, the first variable is essentially treated as a function and whatever comes after it can either be a value or a variable. If it's a variable, as in `f x = ...` above, the left-hand side of the equation matches *anything* of the form `f ?`, such as `f 2` (which equation will `f 1` match?). In such a situation, when the part of the expression of this form is replaced with the right-hand

---

<sup>1</sup>This is an oversimplification, but suffices for the purposes of understanding at this point.

side of the equation, the variable is substituted with the expression being matched (e.g. if the expression being evaluated is `f 2`, then `x` is replaced with `2` in `x + x`).

Note that patterns are recursive: any variable is a pattern, any constant is a pattern, and any constructor applied to patterns is a pattern (e.g. `(p, p')` is a pattern if `p` and `p'` are patterns). Patterns cannot contain anything else, such as operators or functions.

## 2.4 Recursive Functions

It is possible to write down equations that, when combined with the evaluation process of the interpreter, will implement recursive functions. For example, the following two equations define a function that takes an input `n` (assume the input is greater than or equal to zero) and adds up the first `n` integers starting from `1`:

```
module Simple2

  f 0 = 0
  f n = n + f (n - 1)
```

## 2.5 Recursive Functions Involving Lists

**Exercise 2.2.** Define explicitly the function `length`, computing the length of a list.

**Exercise 2.3.** Define a function that takes one integer argument `x` and produces an infinite ascending list of integers starting at `x`.

**Exercise 2.4.** Trace the execution of `prefix 5` on an infinite list. Why does this terminate?

## 2.6 Converting Iteration to Recursion

Suppose we have iterative code of the following form, which computes the sum of the numbers from `1` to `100`:

```
function f (int max) {

  int sum = 0;
  int x = 0;
  while (x <= max) {
    sum = sum + x;
    x = x + 1;
  }

  return sum;

}
```

In order to convert a function that uses an iterative loop into a recursive one, one must first identify all the variables that occur within the scope of the body of the `while` loop. The recursive version of the function will need to have an argument for every one of the variables (in addition to the existing arguments from the iterative definition):

```
f max sum x = ...
```

Next, we identify what transformations are made within the body of the loop and reproduce them by making a recursive call. Notice that the statement `x = x + 1` is simulated by letting an instance of `x` on a new stack frame hold the value of `x + 1` from a previous stack frame:

```
f max sum x = ...
              f max (x + sum) (x + 1)
```

Finally, the terminating condition of the loop becomes the check for the base case. If the base case is reached, the appropriate result is returned.

```
f max sum x = if x <= max then
              f max (x + sum) (x + 1)
            else
              sum
```

The initial definitions of the variables can then be supplied as arguments to the function.

```
f_rec max = f max 0 0
```

## 3 Types

### 3.1 Types and Values

Within Haskell, one intuitive way to view types is as a set of values, where values are anything defined in Table 4. We specify the types for some values in Table 7.

Table 7: Types and values.

value	example	example type
numeric integer literal	2	Int or Integer
numeric non-integer literal	2.3	Double
character literal	'a'	Char
string literal	"abc"	[Char]
tuple of values	(True, 1)	(Bool, Int)
list of values	1:2:[]	[Int]
function	(+)	Int -> Int -> Int

### 3.2 Types Context and Value/Expression Context

In Haskell, type names and expression names (names of variables, functions, and constructors) are always kept in separate contexts. This also means that types never appear in expressions, and expressions never appear in types. Thus, it is possible to have in the same name space (such as a module) an expression variable and type (or type variable) that are identical.

Type names always begin with a capital letter (e.g. `Bool`). Constructor names also always begin with a capital letter. However, the two are in separate contexts.

### 3.3 Type Synonyms

In Haskell, it is possible to define *type synonyms*. These are alternative names for existing types. For example, if a module has the definition

```
type Label = Int
```

any value of type `Int` will also be a value of type `Label`, and vice versa.

### 3.4 Types and Expressions

We have seen how types are defined for values. How are types defined for expressions? Recall the following.

**Fact 3.1.** *For any Haskell expression  $e$ , when  $e$  is evaluated the evaluation process will either terminate (or converge) and produce a value, or it will never terminate (it will diverge).*

When we choose a way to assign types to expressions, it is useful to obey the following proposition.

**Proposition 3.2.** *For any Haskell expression  $e$ , if that expression is of type  $t$ , then either the evaluation of  $e$  diverges, or it converges to a value of type  $t$ .*

This is a kind of type safety condition. It ensures that determining the type of an expression is useful: if an expression has a certain type, it will always produce a value of that type.

## 4 Algebraic Data Types

Recall that types can be viewed as sets of values. The data type statement in Haskell allows users to define their own sets of values. For example, we can define a new set `Color` that contains the values `Red`, `Green`, and `Blue`:

```
data Color = Red | Green | Blue
```

Recall that with the `type` statement, it is possible to take an existing set of values, and give it a new name. The new type introduced by this statement is just another name for the existing set of values, and the two can be used interchangeably:

```
type Point = (Int, Int)
```

What if we want to build a new set from existing sets of values, but we want this set to be distinct? We can “label” or “wrap” values of an existing type with a data type constructor.

```
data Point = Point (Int, Int)
```

If we want to write a function that accept arguments which are values of this data type, we would need to use pattern matching.

```
movePointUp (Point (x,y)) = Point (x,y+1)
```

The general syntax for data type definitions is as follows:

```
data <data type name> <zero or more type arguments> =
  <constructor name> <type of 1st argument> ... <type of nth argument>
| <constructor name> <type of 1st argument> ... <type of nth argument>
| ...
| <constructor name> <type of 1st argument> ... <type of nth argument>
```



## 5 Folding over Algebraic Data Types

Many operations on algebraic data types can be expressed in terms of a *fold* operation over these data types. It is beneficial to formulate operations as folds because this makes them amenable to optimization, parallelization, and parametrization.

Basically, a fold over a value of a certain data type is a replacement of each constructor within that data type with some other expression (which could be another data constructor, a value, or a function). We consider two examples of fold functions for *non-recursive* data types.

```
data Maybe a = Just a | Nothing

maybe :: b -> (a -> b) -> Maybe a -> b
maybe n j (Nothing) = n
maybe n j (Just x)  = j x

data Either a b = Left a | Right b

either :: (a -> c) -> (b -> c) -> Either a b -> c
either l r (Left x)  = l x
either l r (Right x) = r x
```

### 5.1 Defining a Fold Function for a Data Type

We can describe a deterministic process for constructing a fold function for a given data type. Consider the following example type:

```
data G a =
  Line a (G a)
  | Fork a (G a) (G a)
  | Sink
```

To define a fold function, we first determine its type using the following process.

- (1) List the types of the constructors:

```
Line :: a -> G a -> G a
Fork :: a -> G a -> G a -> G a
Sink :: G a
```

- (2) Convert all instances of the data type in question to a fresh type variable. In this case, *b* is a fresh type variable.

```
Line :: a -> b -> b
Fork :: a -> b -> b -> b
Sink :: b
```

- (3) If there are *n* constructors, the `fold` function will take *n*+1 arguments, one for each of the constructors. The types of the arguments will be from the list of types generated in step (2) above. The last argument's type will be the actual data type over which the fold is performed. The output type will be the fresh type variable you introduced in step (2) above.

```
fold :: (a -> b -> b) -> (a -> b -> b -> b) -> b -> G a -> b
```

The actual `fold` definition is fairly straightforward given the type. The number of cases in the fold definition should correspond to the number of constructors in the data type:

```
fold l f s (Line x g)   = ...
fold l f s (Fork x g g') = ...
fold l f s (Sink)      = ...
```

Notice that the fold function has in its initial list of  $n$  arguments a replacement for each constructor. It should apply this replacement to the data inside each constructor. The intuition should be as follows (note that this code is not correct, it is only here to develop the ideas leading to the final version):

```
-- this code does not type check!
fold l f s (Line x g)   = l x g
fold l f s (Fork x g g') = f x g g'
fold l f s (Sink)      = s
```

If our data type is not recursive, our `fold` definition is complete. However, if it is recursive, the above code will not type check. This is because the replacement functions are expecting something of type `b` and not of type `G a`. Fortunately, we can recursively call the `fold` function to turn values of type `G a` into values of type `b`. Notice that we supply all the arguments to the recursive calls of the `fold` function.

```
fold l f s (Line x g)   = l x (fold l f s g)
fold l f s (Fork x g g') = f x (fold l f s g) (fold l f s g')
fold l f s (Sink)      = s
```

Our definition of a fold function is now complete.

## 6 Reasoning Formally about Haskell Programs

Haskell is a pure language in which evaluation is accomplished through algebraic manipulation of expressions according to a collection of equations. Consequently, it is possible to reason about Haskell programs in a formal, mathematical manner. In particular, it is possible to assemble algebraic proofs of logical statements about the behavior of Haskell programs. We demonstrate this by reproducing a collection of examples from a functional programming textbook [Tho99, Ch. 8]. Each example is a proof by induction of some logical statement about common Haskell functions.

We ensure that the proofs in each example are correct by using a variant of an experimental lightweight proof verification system [Lap09] that can verify symbolic manipulations involving conjunction, universal quantification, implication, and equality. To this end, proofs are written in a syntax that corresponds to a subset of Haskell syntax combined with a few new syntactic constructs: `Assume`, `Assert`, `\forall`, `/\`, `=>`. Each of these has the usual logical meaning (as discussed in the textbook [Tho99, Ch. 8]).

### 6.1 Example #1

In this example, we introduce two functions: `sum` and `doubleAll`.

```
Introduce sum, doubleAll.

Assume          sum    [] = 0
```

```

Assume \forall x,xs. sum (x:xs) = x + sum xs

Assume          doubleAll [] = []
Assume \forall z,zs. doubleAll (z:zs) = 2 * z : doubleAll zs

```

We also introduce the distributive property for addition and multiplication, as we will find it useful.

```

Assume \forall x,y,z. (z*x) + (z*y) = z*(x+y)

```

Suppose we want to prove that

$$\forall xs. \text{sum} (\text{doubleAll } xs) = 2 * \text{sum } xs$$

This means we need to prove the above statement for the  $xs = []$  case, and to show that if it holds for some  $xs$ , it also holds for  $x:xs$ .

We start the argument for the base case with an assumption, and continue from there with some basic substitutions based on the base definitions of the two functions.

```

Assert
      0 = 2 * 0
/\ sum [] = 2 * sum []
/\ sum (doubleAll []) = 2 * sum []

```

The recursive case is done in three steps. We take the inductive hypothesis and add an element to both sides. Then, we convert each side of the equation to produce the same side of our goal equation. Finally,

```

Assert \forall x,xs.
      sum (doubleAll xs) = 2 * sum xs -- This is the inductive hypothesis.
=>

-- We can add an element to both sides of the hypothesis equation.
(2*x) + sum (doubleAll xs) = (2*x) + 2 * sum xs

-- We first rewrite the left side.
/\ (2*x) + 2 * sum xs = 2 * (x + sum xs)
/\ sum (x:xs) = x + sum xs
/\ (2*x) + 2 * sum xs = 2 * (sum (x:xs))

-- Next, we rewrite the right side.
/\ sum ((2*x):doubleAll xs) = (2*x) + sum (doubleAll xs)
/\ doubleAll (x:xs) = (2*x):doubleAll xs
/\ sum (doubleAll (x:xs)) = (2*x) + sum (doubleAll xs)

-- Finally, we put the two sides back together.
/\ sum (doubleAll (x:xs)) = 2 * (sum (x:xs))

```

## 6.2 Example #2

For this example, we introduce two functions: `++` and `length`.

```

Assume \forall ys. [] ++ ys = ys

```

```
Assume \forall x,xs,ys. (x:xs) ++ ys = x:(xs ++ ys)
```

```
Assume length [] = 0
Assume \forall x,xs. length (x:xs) = 1 + length xs
```

We also introduce the left identity and associativity for addition, as we will find it useful.

```
Assume \forall x. 0 + x = x
Assume \forall x,y,z. x+(y+z) = (x+y)+z
```

Suppose we want to prove that

$$\forall xs,ys. \text{length } (xs ++ ys) = \text{length } xs + \text{length } ys$$

This means we need to prove the above statement for the  $xs = []$  case, and to show that if it holds for some  $xs$ , it also holds for  $x:xs$ .

The base case is a bit complicated; we even need to use our assumption about the left identity for addition.

```
Assert \forall ys.
  [] ++ ys = ys
  /\ length ([] ++ ys) = length ys
  /\ 0 + length ys = length ys
  /\ length ([] ++ ys) = 0 + length ys
  /\ length ([] ++ ys) = length [] + length ys
```

The recursive case is done in three steps, as in the last example.

```
Assert \forall x,xs,ys.
  length (xs++ys) = length xs + length ys
=>
  1 + length (xs++ys) = 1 + (length xs + length ys)

-- We first rewrite the left side.
/\ length (x:(xs++ys)) = 1 + length (xs++ys)
/\ (x:xs) ++ ys = x:(xs ++ ys)
/\ length ((x:xs)++ys) = 1 + length (xs++ys)

-- Then we rewrite the right side.
/\ 1 + (length xs + length ys) = (1 + length xs) + length ys
/\ length (x:xs) = 1 + length xs
/\ 1 + (length xs + length ys) = length (x:xs) + length ys

-- Finally, we put the two sides back together
/\ length ((x:xs)++ys) = length (x:xs) + length ys
```

### 6.3 General Strategies

When writing a sequence of equations as part of a formal proof, we have used one of several basic strategies. We enumerate and briefly illustrate them.

Introduce  $u,v,x,y,a,b,c,d,P$ .

- Use reflexivity, transitivity, or symmetry of equality.

```
Assert x + y = x + y
Assert ( x = y /\ y = z ) => x = z
Assert x = y => y = x
```

- Repeat an existing assumption.

```
Assume x = y
Assert x = y
```

- Instantiate (i.e. replace quantified variables within) an existing universal assumption.

```
Assume \forall t. P(t)
Assert \forall s. P(s) --instantiate with t = s
```

- Substitute a portion of an equation with an equivalent equation.

```
Assume x + y = a + b
Assume x = u
Assume a = v
Assert u + y = v + b --replace x with u and a with v
```

- Perform the same operation to both sides of an equation.

```
Assume x + y = a + b
Assert z + (x + y) = z + (a + b)
```

## 7 Type Classes

### 7.1 Motivation: Ad Hoc Polymorphism

In Haskell, type classes are a way to implement ad hoc polymorphism. This is distinct from the polymorphism we have seen before. For example, consider the following polymorphic function:

```
map :: (a -> b) -> [a] -> [b]
map f (x:xs) = f x : map f xs
map f []     = []
```

Notice that no matter how we instantiate the type variable `a` and `b`, the function will be the same: two equations, one for the recursive case, and one for the base case. However, what if we want to define a polymorphic function which has *different* definitions for each type to which it can be instantiated? For example, here are two data types and two different equality functions:

```
data Temp = C Double | F Double

(==) :: Temp -> Temp -> Bool
(==) (C d) (C d') = d == d'
(==) (F f) (F f') = f == f'
(==) _      _      = False

data Color = Red | Green | Blue

(==) :: Color -> Color -> Bool
```

```
(==) Red Red      = True
(==) Green Green  = True
(==) Blue Blue    = True
(==) _ _          = False
```

The above is not possible to do directly in Haskell. Those who are familiar with languages such as C++ and Java will recognize that what we are attempting to do is similar to *overloading*. In fact, operator overloading in C++ and Java is also an example of ad hoc polymorphism.

## 7.2 User-defined Type Classes

The examples in this section often supply definitions for functions and classes already found in the standard prelude for Haskell. This can cause conflicts, so it is best to import only those definitions which are of interest. The following will prevent any functions from being loaded from the prelude:

```
import Prelude()
```

If, for example, only the `Eq` type class is desired, the following line would work:

```
import Prelude(Eq, Bool)
```

You will need to explicitly list every function you use from the prelude if you choose to use this method. Most of the non-trivial examples in this section should work under the following form of the import statement:

```
import Prelude(Eq, (==), Bool, not, (||), (&&), Num, Int, (+), head, tail, foldr)
```

### 7.2.1 Basics

Haskell supports ad hoc polymorphism by allowing users to group *types* into collections called *type classes*. For example, there is a class `Eq` that contains all the types for which an equality function (`==`) is defined. Thus, conceptually, we can imagine the following to be true:

$$\text{Eq} = \{\text{Int}, \text{Double}, \text{Bool}, \text{Char}, \dots\}.$$

One way to put a user-defined data type into a class like `Eq` is to use the `deriving` keyword:

```
data Color = Red | Green | Blue
  deriving Eq
```

However, it is also possible to define one's own type class. For example, if we did not have `Eq` in the Haskell library, we would define it in the following way:

```
class Eq a where
  (==) :: a -> a -> Bool
```

This statement can be placed within any module. It says that in order for any type `a` to be in the `Eq` type class, the user *must* supply her own definition of a (`==`) function for that type. To put a type into the type class, a user can use the `instance` statement:

```
instance Eq Color where
  (==) Red Red      = True
  (==) Green Green  = True
  (==) Blue Blue    = True
  (==) _ _          = False
```

Haskell is able to determine *which* definition of (==) should be used using the types of the arguments supplied to the function (this is similar to how Java and C++ can determine which definition of an overloaded function should be used using the *signature*). For example, consider the following declaration:

```
data Tree = Leaf | Node Int Tree Tree

instance Eq Tree where
  (==) (Node i t1 t2) (Node i' t1' t2') = i == i' && t1 == t1' && t2 == t2'
  (==) Leaf           Leaf              = True
  (==) _              _                  = False
```

Notice that in `i == i'`, the instance of (==) for `Int` is called, while in `t1 == t1' && t2 == t2'`, the instance of (==) for `Tree` is called recursively.

## 7.2.2 Using Functions Defined in a Type Class

There are two situations in which we can use functions that are part of a type class: within new functions that are defined *within* that type class, and within normal functions found in our modules.

Suppose we have the following type class declaration (this is our own simplified version of the `Ord` class found in the Haskell libraries):

```
class HasOrder a where
  (<) :: a -> a -> Bool
```

Suppose we want any type in the `HasOrder` type class to also have the (`>=`) operator defined. We could require that both be defined, but we already know that (`>=`) can be defined in terms of (`<`). We can add to a type class function *definitions* that use the functions which are not yet defined (they will be defined in a class instance). So, we can extend the above type class declaration in the following way:

```
class HasOrder a where
  (<) :: a -> a -> Bool

  (>=) :: a -> a -> Bool
  (>=) x y = not (x < y)
```

Now, in order to put some type `a` into the `HasOrder` type class, all that needs to be defined is the (`<`) function, while the (`>=`) will always be defined in terms of (`<`).

```
instance HasOrder Color where
  (<) Red   Green = True
  (<) Green Blue = True
  (<) _     _     = False

  -- no need to define (>=)
```

Intuitively, (`>=`) is more like a normal polymorphic function: it is always the same, regardless of type. It just happens to use an ad hoc polymorphic function within its body. In fact, it is because of this that Haskell allows us to put the definition for (`>=`) *outside* the class declaration in the following manner:

```
(>=) :: (HasOrder a) => a -> a -> Bool
(>=) x y = not (x < y)
```

This corresponds to the second scenario: using ad hoc polymorphic functions defined within a type class declaration inside the bodies of normal functions within the module. In this case, we may need to add to the type of our new function a *class context* that specifies explicitly the type class membership requirements for all the type variables in our function's type. In the above example, `(HasOrder a) =>` is the class context. The class context always appears after the `::` symbol and before the actual type of the function. As another example, suppose we have the following function:

```
-- This is incorrect, and will not type check.
compare :: a -> Int
compare x y = if x < y then 1 else -1
```

The above explicit type annotation would not be correct, because the type `a` cannot be any type, it must be a type within our `HasOrder` type class. This is specified by providing a list of type class constraints before the actual type, separated by a `=>` symbol:

```
compare :: (HasOrder a) => a -> a -> Int
compare x y = if x < y then 1 else -1
```

If multiple type class constraints are necessary, they should be separated by commas.

```
f :: (Eq a, Num a) => a -> a -> a
f x y = if x == y then x + 1 else y + 1
```

### 7.2.3 Relationships between Type Classes

We can also add type class constraints to the class declaration itself.

```
class (Eq a) => HasOrder a where
  (<) :: a -> a -> Bool

  (<=) :: a -> a -> Bool
  (<=) x y = x == y || x < y

  (>=) :: a -> a -> Bool
  (>=) x y = not (x < y)

  (>) :: a -> a -> Bool
  (>) x y = not (x == y) && x >= y

  max :: a -> a -> a
  max x y = if x >= y then x else y
```

The above declaration can be interpreted conceptually as the requirement that `HasOrder`  $\subset$  `Eq` (since any type `a` in the `HasOrder` class must already be an instance of the `Eq` class). Consider another example: suppose we also have classes for finite and bounded data types. If a type is both finite and has an order over it, then it is also bounded. This is because we can simply take the maximum of all the values in the finite type.

```
class Finite a where
  values :: [a] --returns all values of type "a"
```



```
class (HasOrder a, Finite a) => Bounded a where
  upperbound :: a --returns largest value of type "a"
  upperbound = foldr max (head values) (tail values)
```

Finally, instance declarations can be polymorphic. As long as the type class on both sides of the => symbol are the same, the type expression on the right-hand side can be an instantiation of a polymorphic type that uses the variable constrained on the left-hand side of the => symbol. For example, we can declare that any instantiation of [a] is also in the Eq type class in the following manner:

```
instance (Eq a) => Eq [a] where
  (==) (x:xs) (y:ys) = x == y && xs == ys
  (==) []      []      = True
  (==) _      _       = False
```

**Exercise 7.1.** Complete the Eq instance declaration for the following data type.

```
data Tree a = Leaf | Node a (Tree a) (Tree a)

instance (Eq a) => Eq (Tree a) where
  (==) ???
```

## 8 Unification

Unification is a generalization of the idea of solving an equation. To unify two expressions is to find a way to substitute variables in those expressions so that they are equal. For example, to solve the following mathematical equation:

$$x + 2 = 5 + 2$$

is to *unify*  $x + 2$  and  $5 + 2$ . The result of unification in this case is the solution,  $x = 5$ . We call such an assignment of variables to values a *substitution*.

In this course, we will only consider unification over sets that have an *initial* structure (that is, data types): there are no operators like addition or multiplication as in the above example, only constructors. Unification (and, thus, the solution of equations) is much simpler and more tractable for sets with an initial structure.

### 8.1 Substitutions

A *substitution* is any mapping (in a mathematical sense) of variables to values. Substitutions can be represented in various ways. We will represent a substitution by a list of pairs. In each pair in the list, the first component specifies a variable name as a string and the second component specifies the value to which that variable is mapped. For example, the substitution

```
[("x", 1), ("y", 2)]
```

indicates that all instances of "x" map to 1 and all instances of "y" map to 2. To *apply* a substitution is to replace all variables in a data type instance (i.e. a value) with their corresponding value in the substitution. Thus, we represent substitutions as follows:

```
type Subst a = [(String, a)]
```

Unification is defined in terms of substitutions. We say a substitution  $s$  *unifies* two values  $v1$  and  $v2$  if

```
subst s v1 == subst s v2,
```

where ( $\equiv$ ) is derived equality on the type of those values and `subst s v` is the application of a substitution `s` on a value `v`.

## 8.2 Example

As an example, suppose we have already defined the functions `emp`, `sub`, and `get` for manipulating substitutions. Consider the following data type:

```
data Number = D Double | I Int | Var String
```

The substitution function for this data type is defined as follows:

```
subst s (D d)  = D d
subst s (I i)  = I i
subst s (Var x) = case get x s of Nothing -> Var x
                                     Just val -> val
```

The function `vars` is defined as follows:

```
vars (D d)  = []
vars (I i)  = []
vars (Var x) = [x]
```

Finally, the function `unify` for this data type is defined as follows:

```
unify (Var x)      n          = Just (sub x n)
unify n            (Var x)    = Just (sub x n)
unify (D d)        (D d')    = if d == d' then Just emp else Nothing
unify (I i)        (I i')    = if i == i' then Just emp else Nothing
unify n            n          = if t1 == t2 then Just emp else Nothing
```

The above example is redundant in order to act as a hint for other data types. Obviously, only the last of the three equations is necessary. Can you explain why? Notice that when two values are equal and contain no variables, the equation is trivially solved, and so the result of a unification is an empty substitution. In your homework, you will need to define unification for recursive data types, so you will need to make recursive calls to `unify` within the definitions of `unify` for those data types.

## 8.3 Homework Hints

The `vars` function should collect all the variables in the value supplied to it and return a list of those variables. For example:

```
vars (Node (Var "ab") (Var "cd")) = ["ab", "cd"]
vars (Leaf) = []
vars (Succ (Succ (Succ (Var "x")))) = ["x"]
```

The `solved` function must check that there are no variables *within* the values in the second component of each pair. For example:

```
solved [("y", Node (Var "z") Leaf), ("x", Node Leaf Leaf)] = False
solved [("x", Node Leaf Leaf)] = True
solved [] = True
```

The first result is `False` because there is a `Var "z"` within a value in the substitution, which means the substitution is not a solution (since we do not know what the value for `Var "z"` should be). The second result is `True` because it is a solution: it assigns a constant value (without any variables) to the variable `"x"`.

For an example of how `reduce` should work, suppose that your substitution looks as follows:

```
s = [
  ("x", Leaf),
  ("y", Leaf),
  ("z", Node (Var "x") (Var "y"))
]
```

Then `reduce` should compute the following:

```
reduce s = [
  ("x", subst s' Leaf),
  ("y", subst s' Leaf),
  ("z", subst s' (Node (Var "x") (Var "y")))
]
```

where `s'` consists only of the “solved” portion of the substitution, that is in this case:

```
s' = [
  ("x", Leaf),
  ("y", Leaf)
]
```

Notice that this means that `reduce s` in the above example produces:

```
reduce s = [
  ("x", Leaf),
  ("y", Leaf),
  ("z", Node Leaf Leaf)
]
```

Note that `reduce` won't always resolve the whole substitution in one step, e.g. if we have `("x", Var "y")`, `("y", Var "z")`, and `("z", Leaf)` in a substitution then it would take at least two applications of `reduce` to completely resolve the substitution. The variable instance `Var "z"` would first need to be substituted by `Leaf`, and only then could `Var "y"` be substituted by `Leaf` within `("x", Var "y")`.

The function `unify` should return a list with more than one pair if the two sides of the equation being solved have multiple occurrences of variables, e.g.

```
unify (Node (Var "x") (Var "y")) (Node Leaf Leaf) = Just [("x", Leaf), ("y", Leaf)]
```

The “emp” substitution is useful for representing solutions to equations that have no variables in them. Here are some more example inputs and outputs:

```
unify (Node Leaf Leaf) (Node Leaf Leaf) = Just emp --trivially solvable
unify (Node Leaf Leaf) (Leaf) = Nothing --never solvable
```

## 9 Representing and Interpreting a Programming Language

As an example, consider the following data type representing the syntax of a simple imperative programming language.

```
type Loc = Int

data Stmt =
  Assign Loc Int
  | Add Loc Loc Loc
  | If Loc Stmt Stmt
  | While Loc Stmt
  | Block [Stmt]
  | Return Loc
```

The type `Loc` represents a location in memory that is large enough to hold an integer. There are statements in the language for assigning an integer value to a location in memory, for the addition of the values in two memory locations (the result is inserted into the third location), for a conditional branch (based on whether the given location is 0), for a conditional loop, for a way to delineate blocks of code, and for a way to return the location of a result for the entire program.

In order to define an evaluation algorithm for this programming language, it would first be necessary to define a data structure for representing memory. In our case, something with the following interface is sufficient. We can assume that all memory locations are initially set to 0, and that `get` implements this behavior.

```
type Mem = [(Loc, Int)]
emp :: Mem
get :: Mem -> Loc -> Int
set :: Mem -> Loc -> Int -> Mem
```

We can now define an evaluation function for this language. For each statement, the evaluation function takes the memory and performs the actions specified by a single statement. It either updates the memory according to the statement, evaluates another statement based on the memory supplied, evaluates a block by calling a helper function, or returns the result from memory. Notice that in all cases except `Return`, the evaluation function returns a new version of the memory (possibly with changes). Also notice that while the result is represented as a location *within a program*, the *evaluated* result must be an actual integer.

```
eval :: Mem -> Stmt -> Either Mem Int
eval mem (Assign l i) = Left (set mem l i)
eval mem (Add l1 l2 l3) = Left (set mem l (get mem l2 + get mem l3))
eval mem (If l s1 s2) = eval mem (if (get mem l) /= 0 then s1 else s2)
eval mem (While l s) =
  if (get mem l) /= 0 then
    case (eval mem s) of
      Right i -> Right i
      Left mem' -> eval mem' (While l s)
  else
    Left mem --no code to evaluate, so no change to memory

eval mem (Block ss) = evals mem ss
eval mem (Return l) = Right (get mem l)
```

```

evals :: Mem -> [Stmt] -> Either Mem (Maybe Loc)
evals mem []      = Left mem
evals mem (s:ss) =
  case eval mem s of
    Left mem' -> evals mem' ss
    Right i   -> Right i

```

The helper function is fairly simple. If there are no statements in the list, the same memory is returned. Otherwise, the first statement is evaluated. If it returns a new copy of memory, the rest of the list is evaluated under the new memory. Otherwise, the result is returned as the result for the entire list of statements.

## 10 Understanding Inference Rules and Derivations

We will by convention use  $\Gamma$  to represent a particular instance of an *environment*, which is a map from variables to types. The easiest way to visualize a particular  $\Gamma$  is as an association list:

$$\Gamma = \{(x, \text{Int}), (\text{not}, \text{Bool} \rightarrow \text{Bool})\}.$$

The above environment  $\Gamma$  says that the *expression* variable  $x$  (note that this is just a variable in a program, such as one introduced by a `let` statement) has the type `Int`.

The symbol  $\vdash$  is called the *turnstile operator*, and relates particular environments  $\Gamma$  with particular claims about the type of an expression. For example, since  $(x, \text{Int})$  is in the above example instance  $\Gamma$ , we might say that under  $\Gamma$ ,  $x$  has type `Int`. In the form of a formula involving  $\vdash$ , this statement looks as follows:

$$\Gamma \vdash x :: \text{Int}.$$

Formulas of the above form are called *judgments* or, more specifically (when they involve types), *typing judgments*. In general, they take the following form:

$$\Gamma \vdash e :: \tau.$$

Logical relationships can exist between typing judgments. For example, we might have a type system in which the relationship described by the following proposition holds.

**Proposition 10.1.** *If  $\Gamma \vdash e_1 :: \tau_2 \rightarrow \tau_1$  and  $\Gamma \vdash e_2 :: \tau_2$ , then it is implied that  $\Gamma \vdash (e_1 e_2) :: \tau_1$ .*

Notice that the above logical proposition has both an instance of logical conjunction (“... and ...”) and an instance of logical implication (“if ... then it is implied that ...”). One conventional way to write such statements is to construct an *inference rule*. The above proposition would be represented as the inference rule below.

$$[\text{APP}] \frac{\Gamma \vdash e_1 :: \tau_2 \rightarrow \tau_1 \quad \Gamma \vdash e_2 :: \tau_2}{\Gamma \vdash (e_1 e_2) :: \tau_1}$$

The long bar represents implication, and if multiple items appear above the bar, these items are assumed to be combined into a conjunction. Each formula above the bar is called a *premise*, while the formula below the bar is called the *conclusion*.

$$\text{inference rule name} \rightarrow [\text{APP}] \frac{\overbrace{\Gamma \vdash e_1 :: \tau_2 \rightarrow \tau_1}^{\text{premise}} \quad \overbrace{\Gamma \vdash e_2 :: \tau_2}^{\text{premise}}}{\underbrace{\Delta \vdash (e_1 e_2) :: \tau_1}_{\text{conclusion}}} \leftarrow \text{this bar represents implication}$$

Converting an inference rule into a Haskell *type checking* function for Mini-Haskell is straightforward. The conclusion becomes the pattern on the left-hand side of an equation, the bar becomes the = notation, and the premises become the body of the function. Multiple premises are combined using the (&&) operator. For example, consider the following inference rule.

$$[\text{IF}] \frac{\Gamma \vdash e_1 :: \text{Bool} \quad \Gamma \vdash e_2 :: \tau \quad \Gamma \vdash e_3 :: \tau}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 :: \tau}$$

In Haskell, this rule would correspond to the following portion of a definition for a type checking function.

```
typeCheck :: Env Ty -> Exp -> Ty -> Bool
typeCheck gamma (If e1 e2 e3) t =
    typeCheck gamma e1 TyBool && typeCheck gamma e2 t && typeCheck gamma e3 t

-- other cases
typeCheck ...
```

The `typeCheck` function should return `True` if the user supplies to it an expression and the correct type for that expression, and `False` otherwise.

It is also not very difficult to convert type inference rules into a Haskell *type inference* function for Mini-Haskell. Essentially, we want to reproduce the `typeCheck` function above, but instead of asking the user to supply the type as an argument, we return it as a result. Since we are only interested in returning the correct type, it is no longer necessary to return a boolean value. However, some expressions cannot be typed (such as `1 && "a"`), so we must make the return type `Maybe Ty` instead of simply `Ty`. This also requires us to match the results of the function's recursive calls to patterns.

```
typeInfer :: Env Ty -> Exp -> Maybe Ty
typeInfer gamma (If e1 e2 e3) =
    case typeInfer gamma e1 of
        Nothing -> Nothing
        Just TyBool ->
            case (typeInfer gamma e2, typeInfer gamma e3) of
                (Just t, Just t') -> if t==t' then Just t else Nothing
                _ -> Nothing

typeInfer gamma (App e1 e2) =
    let t2 = typeInfer e2
    in case typeInfer e1 of
        (TyArrow t2' t1) -> if t2==t2' then Just t1 else Nothing
        _ -> Nothing
```

In the `If` case above, note how we first try to infer the type of the condition. Since this might fail, we must handle both possible outcomes. Notice also that while a type inference rule might have multiple occurrences of a type  $\tau$  within the premises (such as the type of both branches of the `if` statement in the `[IF]` inference rule), the function must obtain *distinct* results from each recursive call (such as `t` and `t'` in the first case) and then compare them using `(==)`.

# 11 Monads

## 11.1 Pure Functions

Suppose we are working a setting involving only pure functions that always return a result and have no effects (no errors, failures, options, effects on I/O devices, and so on). In this case, all our functions can have types of a straightforward form,

$$a \rightarrow b$$

for some types  $a$  and  $b$ . We can view all our programs as being constructed out of many functions that are combined with two operators, `apply` and `compose`:

```
apply :: a -> (a -> b) -> b
apply x g = g x

compose :: (a -> b) -> (b -> c) -> (a -> c)
compose f g = \x -> g (f x)
```

We can even represent all values of any type  $a$  as functions of the type  $() \rightarrow a$ , in which case we do not need a separate `apply` function. We can define `apply'` on such values in terms of `compose`:

```
apply' :: (() -> a) -> (a -> b) -> (() -> b)
apply' x f = compose x f
```

Thus, we can view programming as working with a space of functions with types of the form  $a \rightarrow b$  (note that this includes functions with different instantiations of  $a$  and  $b$ ) and a single operator on functions, `compose`.

## 11.2 Impure Functions

But what if we are not in a pure setting and our functions' types have other forms. Some functions might return errors, multiple results, or might update and return some sort of state value.

```
a -> Maybe b
a -> Error b
a -> [b]
a -> IO b
a -> State b
```

One convention way to deal with these situations is to introduce a *monad*. Suppose that instead of pure functions with types of the form  $a \rightarrow b$ , we must work with functions with types of the form  $a \rightarrow M b$  where  $M$  is some higher-order type (such as `Maybe` or `Error`). In order to program with such functions, we would need a `compose` operator for this space of functions that has the type

$$\text{compose} :: (a \rightarrow M b) \rightarrow (b \rightarrow M c) \rightarrow (a \rightarrow M c)$$

How do we define such a function? An intuitive approach is to define for each monad  $M$  (where  $M$  is something like `Maybe`) two function:

```
lift :: (a -> b) -> (M a -> M b)
join :: M (M a) -> M a
```

These two functions are straightforward to define for most monads  $M$ . As an example, we present the two functions for `Maybe`.

```
lift :: (a -> b) -> (Maybe a -> Maybe b)
lift f = \x -> case x of
  Nothing -> Nothing
  Just x -> Just (f x)

join :: Maybe (Maybe a) -> Maybe a
join (Just (Just x)) = Just x
join _                = Nothing
```

For another example, suppose  $M\ a = [a]$ . Then `lift = map` and `join = concat`.

Once these two functions are defined, it is possible to define `compose`.

```
compose :: (a -> M b) -> (b -> M c) -> (a -> M c)
compose f g = \x -> join ((lift g) (f x))
```

It is also possible to define a corresponding `apply` function, but in the case of monads, this function is called `bind` or `(>>=)`.

```
(>>=) :: M a -> (a -> M b) -> M b
(>>=) x f = join ((lift f) x)
```

However, notice that `(>>=)` alone cannot replace `compose`, because there is no way to build a value of the form  $M\ a$  in the first place. An additional function is needed, usually called `return`, of type  $a \rightarrow M\ a$ . These two functions can be used to define `compose`. In Haskell, the convention is to define these two functions for a monad instead of `compose`. This convention is embodied in the `Monad` class declaration, found in the Haskell libraries:

```
class Monad m where
  (>>=) :: m a -> (a -> m b) -> m b
  return :: a -> m a
```

We consider the `Maybe` monad again, this time defining `(>>=)` and `return` as part of an instance declaration for `Monad`.

```
instance Monad Maybe where
  (>>=) x f = case x of
    Just x -> f x
    Nothing -> Nothing

  return = Just
```

### 11.3 Representing Sequential Computations

Suppose we are again within a setting that involves only pure functions. One way to represent a sequenced computation in a functional language that is evaluated using a call-by-value evaluation scheme is as a collection of nested `let` statements:



```
let x1 = e1 in
let x2 = e2 in
let x3 = e3 in
f x1 x2 x3
```

We can rewrite this with parentheses as

```
let x1 = e1 in (let x2 = e2 in (let x3 = e3 in f x1 x2 x3)),
```

and using lambda abstractions as

```
(\x1 -> (\x2 -> (\x3 -> f x1 x2 x3) e3) e2) e1.
```

Notice that we apply a lambda to each of the expressions `e1`, `e2`, and `e3`. Suppose we now rewrite this again by using our explicit `apply` operator.

```
e1 'apply' (\x1 ->
e2 'apply' (\x2 ->
e3 'apply' (\x3 ->
  f x1 x2 x3
)))
```

What if we could create our own syntax? Then we could rearrange the above as:

```
do {
  x1 <- e1;
  x2 <- e2;
  x3 <- e3;
  f x1 x2 x3
}
```

This is precisely how Haskell `do` notation works, except that it works with `(>>=)` instead of `apply`. This is an advantage because the `(>>=)` operator, unlike `apply`, takes care of any side effects and failure outputs, allowing us to use functions that take only pure inputs even if the return types of our functions might have side effects or failures. However, this also means that the final result must be wrapped in a `return` statement so that it is of a monadic type. The result of the entire computation must be of a monadic type because any of the instances of `(>>=)` could result in a failure or side effect (the result may not even be returned). All of these possible outcomes, including the successful result, must have the same type.

## References

- [Je99] Simon Peyton Jones and John Hughes (editors). Haskell 98: A non-strict, purely functional language. Technical report, February 1999.
- [Lap09] Andrei Lapets. Improving the accessibility of lightweight formal verification systems. Technical Report BUCS-TR-2009-015, Computer Science Department, Boston University, April 30 2009.
- [Mit08] Neil Mitchell. Hoogle overview. *The Monad.Reader*, (12):27–35, November 2008.
- [Tho99] Simon Thompson. *The Haskell: The Craft of Functional Programming*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.