

Yet Another Haskell Tutorial

Hal Daume III

<http://www.isi.edu/~hdaume/htut/>

Copyright (c) Hal Daume III, 2002–2004.

The preprint version of this tutorial is intended to be available to the entire Haskell community, so we grant permission to copy and distribute it for any purpose, provided that it is reproduced in its entirety, including this notice. Modified versions may not be distributed without prior consent of the author, and must still maintain a copy of this notice. The author retains the right to change or modify this copyright at any time, as well as to make the book no longer free of charge.

About This Report

The goal of the *Yet Another Haskell Tutorial* is to provide a complete introduction to the Haskell programming language. It assumes no knowledge of the Haskell language or familiarity with functional programming in general. However, general familiarity with programming concepts (such as algorithms) will be helpful. This is not intended to be an introduction to programming in general; rather, to programming in Haskell. Sufficient familiarity with your operating system and a text editor is also necessary (this report only discusses installation on configuration on Windows and *Nix system; other operating systems may be supported – consult the documentation of your chosen compiler for more information on installing on other platforms).

What is Haskell?

Haskell is called a lazy, pure functional programming language. It is called *lazy* because expressions which are not needed to determine the answer to a problem are not evaluated. The opposite of *lazy* is *strict*, which is the evaluation strategy of most common programming languages (C, C++, Java, even ML). A strict language is one in which every expression is evaluated, whether the result of its computation is important or not. (This is probably not entirely true as optimizing compilers for strict languages often do what's called “dead code elimination” – this removes unused expressions from the program.) It is called *pure* because it does not allow side effects (A side effect is something that affects the “state” of the world. For instance, a function that prints something to the screen is said to be side-effecting, as is a function which affects the value of a global variable.) – of course, a programming language without side effects would be horribly useless; Haskell uses a system of *monads* to isolate all impure computations from the rest of the program and perform them in the safe way (see Chapter 9 for a discussion of monads proper or Chapter 5 for how to do input/output in a pure language).

Haskell is called a *functional* language because the evaluation of a program is equivalent to evaluating a function in the pure mathematical sense. This also differs from standard languages (like C and Java) which evaluate a sequence of statements, one after the other (this is termed an *imperative* language).

The History of Haskell

The history of Haskell is best described using the words of the authors. The following text is quoted from the published version of the Haskell 98 Report:

In September of 1987 a meeting was held at the conference on Functional Programming Languages and Computer Architecture (FPCA '87) in Portland, Oregon, to discuss an unfortunate situation in the functional programming community: there had come into being more than a dozen non-strict, purely functional programming languages, all similar in expressive power and semantic underpinnings. There was a strong consensus at this meeting that more widespread use of this class of functional languages was being hampered by the lack of a common language. It was decided that a committee should be formed to design such a language, providing faster communication of new ideas, a stable foundation for real applications development, and a vehicle through which others would be encouraged to use functional languages. This document describes the result of that committee's efforts: a purely functional programming language called Haskell, named after the logician Haskell B. Curry whose work provides the logical basis for much of ours.

The committee's primary goal was to design a language that satisfied these constraints:

1. It should be suitable for teaching, research, and applications, including building large systems.
2. It should be completely described via the publication of a formal syntax and semantics.
3. It should be freely available. Anyone should be permitted to implement the language and distribute it to whomever they please.
4. It should be based on ideas that enjoy a wide consensus.
5. It should reduce unnecessary diversity in functional programming languages.

The committee intended that Haskell would serve as a basis for future research in language design, and hoped that extensions or variants of the language would appear, incorporating experimental features.

Haskell has indeed evolved continuously since its original publication. By the middle of 1997, there had been four iterations of the language design (the latest at that point being Haskell 1.4). At the 1997 Haskell Workshop in Amsterdam, it was decided that a stable variant of Haskell was needed; this stable language is the subject of this Report, and is called "Haskell 98".

Haskell 98 was conceived as a relatively minor tidy-up of Haskell 1.4, making some simplifications, and removing some pitfalls for the unwary.

It is intended to be a “stable” language in sense the *implementors are committed to supporting Haskell 98 exactly as specified, for the foreseeable future.*

The original Haskell Report covered only the language, together with a standard library called the Prelude. By the time Haskell 98 was stabilised, it had become clear that many programs need access to a larger set of library functions (notably concerning input/output and simple interaction with the operating system). If these program were to be portable, a set of libraries would have to be standardised too. A separate effort was therefore begun by a distinct (but overlapping) committee to fix the Haskell 98 Libraries.

Why Use Haskell?

Clearly you’re interested in Haskell since you’re reading this tutorial. There are many motivations for using Haskell. My personal reason for using Haskell is that I have found that I write more bug-free code in less time using Haskell than any other language. I also find it very readable and extensible.

Perhaps most importantly, however, I have consistently found the Haskell community to be incredibly helpful. The language is constantly evolving (that’s not to say it’s instable; rather that there are numerous extensions that have been added to some compilers which I find very useful) and user suggestions are often heeded when new extensions are to be implemented.

Why Not Use Haskell?

My two biggest complaints, and the complaints of most Haskellers I know, are: (1) the generated code tends to be slower than equivalent programs written in a language like C; and (2) it tends to be difficult to debug.

The second problem tends not be to a very big issue: most of the code I’ve written is not buggy, as most of the common sources of bugs in other languages simply don’t exist in Haskell. The first issue certainly has come up a few times in my experience; however, CPU time is almost always cheaper than programmer time and if I have to wait a little longer for my results after having saved a few days programming and debugging.

Of course, this isn’t the case of all applications. Some people may find that the speed hit taken for using Haskell is unbearable. However, Haskell has a standardized *foreign-function interface* which allow you to link in code written in other languages, for when you need to get the most speed out of your code. If you don’t find this sufficient, I would suggest taking a look at the language O’Caml, which often *outperforms* even C++, yet also has many of the benefits of Haskell.

Target Audience

There have been many books and tutorials written about Haskell; for a (nearly) complete list, visit the <http://haskell.org/bookshelf> (Haskell Bookshelf) at the Haskell homepage. A brief survey of the tutorials available yields:

- *A Gentle Introduction to Haskell* is an introduction to Haskell, given that the reader is familiar with functional programming en large.
- *Haskell Companion* is a short reference of common concepts and definitions.
- *Online Haskell Course* is a short course (in German) for beginning with Haskell.
- *Two Dozen Short Lessons in Haskell* is the draft of an excellent textbook that emphasizes user involvement.
- *Haskell Tutorial* is based on a course given at the 3rd International Summer School on Advanced Functional Programming.
- *Haskell for Miranda Programmers* assumes knowledge of the language Miranda.
- *PLEAC-Haskell* is a tutorial in the style of the Perl Cookbook.

Though all of these tutorials is excellent, they are on their own incomplete: The “Gentle Introduction” is far too advanced for beginning Haskellers and the others tend to end too early, or not cover everything. Haskell is full of pitfalls for new programmers and experienced non-functional programmers alike, as can be witnessed by reading through the archives of the Haskell mailing list.

It became clear that there is a strong need for a tutorial which is introductory in the sense that it does not assume knowledge of functional programming, but which is advanced in the sense that it *does* assume some background in programming. Moreover, none of the known tutorials introduce input/output and interactivity soon enough (not even until the 248th page, as in the case of the Hudak book). This tutorial is not for beginning programmers; some experience and knowledge of programming and computers is assumed (though the appendix does contain some background information).

The Haskell language underwent a standardization process and the result is called Haskell 98. The majority of this book will cover the Haskell 98 standard. Any deviations from the standard will be noted (for instance, many compilers offer certain extensions to the standard which are useful; some of these may be discussed).

The goals of this tutorial are:

- to be *practical* above all else
- to provide a comprehensive, free introduction to the Haskell language
- to point out common pitfalls and their solutions
- to provide a good sense of how Haskell can be used in the real world

Acknowledgements

It would be inappropriate not to give credit also to the original designers of Haskell. Those are: Arvind, Lennart Augustsson, Dave Barton, Brian Boutel, Warren Burton, Jon Fairbairn, Joseph Fasel, Andy Gordon, Maria Guzman, Kevin Hammond, Ralf Hinze, Paul Hudak, John Hughes, Thomas Johnsson, Mark Jones, Dick Kieburtz, John Launchbury, Erik Meijer, Rishiyur Nikhil, John Peterson, Simon Peyton Jones, Mike Reeve, Alastair Reid, Colin Runciman, Philip Wadler, David Wise, Jonathan Young.

Finally, I would like to specifically thank Simon Peyton Jones, Simon Marlow, John Hughes, Alastair Reid, Koen Classen, Manuel Chakravarty, Sigbjorn Finne and Sven Panne, all of whom have made my life learning Haskell all the more enjoyable by always being supportive. There were doubtless others who helped and are not listed, but these are those who come to mind.

- Hal Daumé III

Contents

1	Introduction	3
2	Getting Started	5
2.1	Hugs	5
2.1.1	Where to get it	6
2.1.2	Installation procedures	6
2.1.3	How to run it	6
2.1.4	Program options	7
2.1.5	How to get help	7
2.2	Glasgow Haskell Compiler	7
2.2.1	Where to get it	7
2.2.2	Installation procedures	8
2.2.3	How to run the compiler	8
2.2.4	How to run the interpreter	8
2.2.5	Program options	9
2.2.6	How to get help	9
2.3	NHC	9
2.3.1	Where to get it	9
2.3.2	Installation procedures	9
2.3.3	How to run it	9
2.3.4	Program options	9
2.3.5	How to get help	9
2.4	Editors	9
3	Language Basics	11
3.1	Arithmetic	13
3.2	Pairs, Triples and More	14
3.3	Lists	15
3.3.1	Strings	17
3.3.2	Simple List Functions	17
3.4	Source Code Files	20
3.5	Functions	22
3.5.1	Let Bindings	26
3.5.2	Infix	27

3.6	Comments	27
3.7	Recursion	28
3.8	Interactivity	31
4	Type Basics	37
4.1	Simple Types	37
4.2	Polymorphic Types	39
4.3	Type Classes	40
4.3.1	Motivation	40
4.3.2	Equality Testing	41
4.3.3	The Num Class	41
4.3.4	The Show Class	41
4.4	Function Types	42
4.4.1	Lambda Calculus	42
4.4.2	Higher-Order Types	42
4.4.3	That Pesky IO Type	44
4.4.4	Explicit Type Declarations	45
4.4.5	Functional Arguments	46
4.5	Data Types	47
4.5.1	Pairs	47
4.5.2	Multiple Constructors	49
4.5.3	Recursive Datatypes	51
4.5.4	Binary Trees	51
4.5.5	Enumerated Sets	52
4.5.6	The Unit type	53
4.6	Continuation Passing Style	53
5	Basic Input/Output	57
5.1	The RealWorld Solution	57
5.2	The CPS Solution	58
5.3	Actions	59
5.4	The IO Library	63
5.5	A File Reading Program	65
6	Modules	69
6.1	Exports	69
6.2	Imports	71
6.3	Hierarchical Imports	72
6.4	Literate Versus Non-Literate	73
6.4.1	Bird-scripts	73
6.4.2	LaTeX-scripts	74

7	Advanced Features	75
7.1	Sections and Infix Operators	75
7.2	Local Declarations	76
7.3	Partial Application	78
7.4	Pattern Matching	82
7.5	Guards	85
7.6	Instance Declarations	87
7.6.1	The <code>Eq</code> Class	87
7.6.2	The <code>Show</code> Class	88
7.6.3	Other Important Classes	89
7.6.4	Class Contexts	91
7.6.5	Deriving Classes	91
7.7	Datatypes Revisited	92
7.7.1	Named Fields	92
7.8	More Lists	94
7.8.1	Standard List Functions	94
7.8.2	List Comprehensions	96
7.9	Arrays	98
7.10	Finite Maps	99
7.11	Layout	101
7.12	The Final Word on Lists	101
8	Advanced Types	105
8.1	Type Synonyms	105
8.2	Newtypes	106
8.3	Datatypes	107
8.3.1	Strict Fields	107
8.4	Classes	110
8.4.1	Pong	110
8.4.2	Computations	111
8.5	Instances	115
8.6	Kinds	117
8.7	Class Hierarchies	119
8.8	Default	119
9	Monads	121
9.1	Do Notation	122
9.2	Definition	124
9.3	A Simple State Monad	126
9.4	Common Monads	132
9.5	Monadic Combinators	136
9.6	MonadPlus	139
9.7	Monad Transformers	141
9.8	Parsing Monads	145
9.8.1	A Simple Parsing Monad	146
9.8.2	Parsec	152

10 Advanced Techniques	159
10.1 Exceptions	159
10.2 Mutable Arrays	159
10.3 Mutable References	159
10.4 The ST Monad	159
10.5 Concurrency	159
10.6 Regular Expressions	159
10.7 Dynamic Types	159
11 Theory	161
11.1 Bottom	161
12 For Further Information	163
A Brief Complexity Theory	165
B Recursion and Induction	167
C “Real World” Examples	169
D Solutions To Exercises	171

Chapter 1

Introduction

This tutorial contains a whole host of example code, all of which should have been included in its distribution. If not, please refer to the links off of the Haskell web site (haskell.org) to get it. This book is formatted to make example code stand out from the rest of the text.

```
Code will look like this.
```

Occasionally, we will refer to interaction between you and the operating system and/or the interactive shell (more on this in Section 2).

```
Interaction will look like this.
```

Strewn throughout the tutorial, we will often make additional notes to something written. These are often for making comparisons to other programming languages or adding helpful information.

```
■ NOTE ■ Notes will appear like this.
```

If we're covering a difficult or confusing topic and there is something you should watch out for, we will place a warning.

```
■ WARNING ■ Warnings will appear like this.
```

Finally, we will sometimes make reference to built-in functions (so-called Prelude-functions). This will look something like this:

$$\text{map} :: (a \rightarrow b) \rightarrow [a] \rightarrow [b]$$

Within the body text, Haskell keywords will appear like this: **where**, identifiers as `map`, types as `String` and classes as `Eq`.

Chapter 2

Getting Started

There are three well known Haskell system: Hugs, GHC and NHC. Hugs is exclusively an interpreter, meaning that you cannot compile stand-alone programs with it, but can test and debug programs in an interactive environment. GHC is both an interpreter (like Hugs) and a compiler which will produce stand-alone programs. NHC is exclusively a compiler. Which you use is entirely up to you. I've tried to make a list of some of the differences in the following list but of course this is far from exhaustive:

Hugs - very fast; implements almost all of Haskell 98 (the standard) and most extensions; built-in support for module browsing; cannot create stand-alones; written in C; works on almost every platform; build in graphics library.

GHC - interactive environment is slower than Hugs, but allows function definitions in the environment (in Hugs you have to put them in a file); implements all of Haskell 98 and extensions; good support for interfacing with other languages; in a sense the “de facto” standard.

NHC - less used and no interactive environment, but produces smaller and often faster executables than does GHC; supports Haskell 98 and some extensions.

I, personally, have all of them installed and use them for different purposes. I tend to use GHC to compile (primarily because I'm most familiar with it) and the Hugs interactive environment, since it is much faster. As such, this is what I would suggest. However, that is a fair amount to download and install, so if you had to go with just one, I'd get GHC, since it contains both a compiler and interactive environment.

Following is a description of how to download and install each of this as of the time this tutorial was written. It may have changed – see <http://haskell.org> (the Haskell website) for up-to-date information.

2.1 Hugs

Hugs supports almost all of the Haskell 98 standard (it lacks some of the libraries), as well as a number of advanced/experimental extensions, including: multi-parameter

type classes, extensible records, rank-2 polymorphism, existentials, scoped type variables, and restricted type synonyms.

2.1.1 Where to get it

The official Hugs web page is at:

`http://haskell.org/hugs` (`http://haskell.org/hugs`)

If you go there, there is a link titled “downloading” which will send you to the download page. From that page, you can download the appropriate version of Hugs for your computer.

2.1.2 Installation procedures

Once you’ve downloaded Hugs, installation differs depending on your platform, however, installation for Hugs is more or less identical to installation for any program on your platform.

For Windows when you click on the “msi” file to download, simply choose “Run This Program” and the installation will begin automatically. From there, just follow the on-screen instructions.

For RPMs use whatever RPM installation program you know best.

For source first gunzip the file, then untar it. Presumably if you’re using a system which isn’t otherwise supported, you know enough about your system to be able to run configure scripts and make things by hand.

2.1.3 How to run it

On Unix machines, the Hugs interpreter is usually started with a command line of the form: `hugs [option — file] ...`

On Windows, Hugs may be started by selecting it from the start menu or by double clicking on a file with the `.hs` or `.lhs` extension. (This manual assumes that Hugs has already been successfully installed on your system.)

Hugs uses options to set system parameters. These options are distinguished by a leading `+` or `-` and are used to customize the behaviour of the interpreter. When Hugs starts, the interpreter performs the following tasks:

- Options in the environment are processed. The variable `HUGSFLAGS` holds these options. On Windows 95/NT, the registry is also queried for Hugs option settings.
- Command line options are processed.
- Internal data structures are initialized. In particular, the heap is initialized, and its size is fixed at this point; if you want to run the interpreter with a heap size other than the default, then this must be specified using options on the command line, in the environment or in the registry.

- The prelude file is loaded. The interpreter will look for the prelude file on the path specified by the `-P` option. If the prelude, located in the file `Prelude.hs`, cannot be found in one of the path directories or in the current directory, then Hugs will terminate; Hugs will not run without the prelude file.
- Program files specified on the command line are loaded. The effect of a command `hugs f1 ... fn` is the same as starting up Hugs with the `hugs` command and then typing `:load f1 ... fn`. In particular, the interpreter will not terminate if a problem occurs while it is trying to load one of the specified files, but it will abort the attempted load command.

The environment variables and command line options used by Hugs are described in the following sections.

2.1.4 Program options

To list all of the options would take too much space. The most important option at this point is `+98` or `-98`. When you start hugs with `+98` it is in Haskell 98 mode, which turns off all extensions. When you start in `-98`, you are in Hugs mode and all extensions are turned on. If you've downloaded someone else's code and you're having trouble loading it, first make sure you have the `98` flag set properly.

Further information on the Hugs options is in the manual:

<http://cvs.haskell.org/Hugs/pages/hugsman/started.html>).

2.1.5 How to get help

To get Hugs specific help, go to the Hugs web page. To get general Haskell help, go to the Haskell web page.

2.2 Glasgow Haskell Compiler

The Glasgow Haskell Compiler (GHC) is a robust, fully-featured, optimising compiler and interactive environment for Haskell 98; GHC compiles Haskell to either native code or C. It implements numerous experimental language extensions to Haskell 98; for example: concurrency, a foreign language interface, multi-parameter type classes, scoped type variables, existential and universal quantification, unboxed types, exceptions, weak pointers, and so on. GHC comes with a generational garbage collector, and a space and time profiler.

2.2.1 Where to get it

Go to the official GHC web page <http://haskell.org/ghc> (GHC) to download the latest release. The current version as of the writing of this tutorial is 5.04.2 and can be downloaded off of the GHC download page (follow the "Download" link). From that page, you can download the appropriate version of GHC for your computer.

2.2.2 Installation procedures

Once you've downloaded GHC, installation differs depending on your platform; however, installation for GHC is more or less identical to installation for any program on your platform.

For Windows when you click on the “msi” file to download, simply choose “Run This Program” and the installation will begin automatically. From there, just follow the on-screen instructions.

For RPMs use whatever RPM installation program you know best.

For source first gunzip the file, then untar it. Presumably if you're using a system which isn't otherwise supported, you know enough about your system to be able to run configure scripts and make things by hand.

For a more detailed description of the installation procedure, look at the GHC users manual under “Installing GHC”.

2.2.3 How to run the compiler

Running the compiler is fairly easy. Assuming that you have a program written with a `main` function in a file called `Main.hs`, you can compile it simply by writing:

```
% ghc --make Main.hs -o main
```

The “`--make`” option tells GHC that this is a program and not just a library and you want to build it and all modules it depends on. “`Main.hs`” stipulates the name of the file to compile; and the “`-o main`” means that you want to put the output in a file called “`main`”.

■ NOTE ■ In Windows, you should say “`-o main.exe`” to tell Windows that this is an executable file.

You can then run the program by simply typing “`main`” at the prompt.

2.2.4 How to run the interpreter

GHCi is invoked with the command “`ghci`” or “`ghc -interactive`”. One or more modules or filenames can also be specified on the command line; this instructs GHCi to load the specified modules or filenames (and all the modules they depend on), just as if you had said `:load modules` at the GHCi prompt.

2.2.5 Program options

To list all of the options would take too much space. The most important option at this point is “-fglasgow-exts”. When you start GHCi without “-fglasgow-exts” it is in Haskell 98 mode, which turns off all extensions. When you start with “-fglasgow-exts”, all extensions are turned on. If you’ve downloaded someone else’s code and you’re having trouble loading it, first make sure you have this flag set properly.

Further information on the GHC and GHCi options are in the manual off of the GHC web page.

2.2.6 How to get help

To get GHC(i) specific help, go to the GHC web page. To get general Haskell help, go to the Haskell web page.

2.3 NHC

About NHC . . .

2.3.1 Where to get it

2.3.2 Installation procedures

2.3.3 How to run it

2.3.4 Program options

2.3.5 How to get help

2.4 Editors

With good text editor, programming is fun. Of course, you can get along with simplistic editor capable of just cut-n-paste, but good editor is capable of doing most of the chores for you, letting you concentrate on what you are writing. With respect to programming in Haskell, good text editor should have as much as possible of the following features:

- Syntax highlighting for source files
- Indentation of source files
- Interaction with Haskell interpreter (be it Hugs or GHCi)
- Computer-aided code navigation
- Code completion

At the time of writing, several options were available: Emacs/XEmacs support Haskell via *haskell-mode* and accompanying Elist code (available from <http://www.haskell.org/haskell-mode/>), and

What's else available? . . .

(X)Emacs seem to do the best job, having all the features listed above. Indentation is aware about Haskell's 2-dimensional layout rules (see Section 7.11, very smart and have to be seen in action to be believed. You can quickly jump to the definition of chosen function with the help of "Definitions" menu, and name of the currently edited function is always displayed in the modeline.

Chapter 3

Language Basics

In this chapter we present the basic concepts of Haskell. In addition to familiarizing you with the interactive environments and showing you how to compile a basic program, we introduce the basic syntax of Haskell, which will probably be quite alien if you are used to languages like C and Java.

However, before we talk about specifics of the language, we need to establish some general properties of Haskell. Most importantly, Haskell is a *lazy* language, which means that no computation takes place until the result of that computation is used.

This means, for instance, that you can define infinitely large data structures, provided that you never use the entire structure. For instance, using imperative-esque pseudo-code, we could create an infinite list containing the number 1 in each position by doing something like:

```
List makeList()  
{  
  List current = new List();  
  current.value = 1;  
  current.next = makeList();  
  return current;  
}
```

By looking at this code, we can see what it's trying to do: it creates a new list, sets its value to 1 and then recursively calls itself to make the rest of the list. Of course, if you actually wrote this code and called it, the program would never terminate, because `makeList` would keep calling itself ad infinitum.

This is because we assume this imperative-esque language is *strict*, the opposite of *lazy*. Strict languages are often referred to as “call by value” while lazy languages are referred to as “call by name.” In the above pseudo-code, when we “run” `makeList` on the fifth line, we attempt to get a *value* out of it. This leads to an infinite loop.

The equivalent code in Haskell is:

```
makeList = 1 : makeList
```

This program reads: we’re defining something called `makeList` (this is what goes on the left-hand side of the equals sign). On the right-hand side, we give the definition of `makeList`. In Haskell (we’ll talk more about this soon), the colon operator is used to create lists. This right-hand side says that the value of `makeList` is the element 1 stuck on to the beginning of the value of `makeList`.

However, since Haskell is lazy (or “call by name”), we do not actually attempt to evaluate what `makeList` is at this point: we simply remember that if ever in the future we need the second element of `makeList`, we need to just look at `makeList`.

Now, if you attempt to write `makeList` to a file, or print it to the screen, or calculate the sum of its elements, the operation won’t terminate because it would have to evaluate an infinitely long list. However, if you simply use a finite portion of the list, say the first 10 elements, the fact that the list is infinitely long doesn’t matter. If you only use the first 10 elements, only the first 10 elements are ever calculated. This is what it means that Haskell is lazy.

case sensitive values types Second, Haskell is case sensitive. Many languages are, but Haskell actually uses case to give meaning. Haskell distinguishes between *values* (things like numbers (1, 2, 3, . . .), strings (“abc”, “hello”, . . .), characters (‘a’, ‘b’, ‘ ’, . . .), even functions (for instance, the function which squares a value, or the square-root function)) and *types* (the categories to which values belong).

By itself, this is not unusual. Most languages have some system of types. What is unusual is that Haskell *requires* that the names given to functions and values begin with a lower-case letter and that the names given to types begin with an upper-case letter. The moral is: if your otherwise correct program won’t compile, be sure you haven’t named your function `FOO`, or something else beginning with a capital letter.

side effects Being a functional language, Haskell eschews side effects. A side effect is essentially something which happens in the course of executing a function which is not related to the output produced by that function.

For instance, in a language like C or Java, you are able to modify “global” variables from within a function. This is a side effect because the modification of this global variable is not related to the output produced by the function. Furthermore, modifying the state of the real world is considered a side effect: printing something to the screen, reading a file, etc., are all side effecting operations.

pure Functions which do not have side effects are called *pure*. An easy test for whether a function is pure or not is to ask yourself a simple question: given the same arguments, will this function always produce the same result.

All of this means that if you’re used to writing code in an imperative language (like C or Java), you’re going to have to start thinking differently. Most importantly, if you have a value `x`, you must *not* think about `x` as a register or a memory location or anything like that. `x` is simply a name, just as “Hal” is a name. You cannot arbitrarily decide to store a different person in my name any more than you can arbitrarily decide to store a different value in `x`. This means that code which might look like the following C code is invalid in Haskell (not only is it invalid, but it has no counterpart):

```
int x = 5;
x = x + 1;
```

A call like $x = x + 1$ is called *destructive update* because we are destroying whatever was in x before and replacing it with a new value. Destructive update does not exist in Haskell. destructive update

By not allowing destructive updates (or any other such side effecting operations), Haskell code is very easy to reason about. That is, when we define a function f and call that function with a particular argument a in the beginning of a program and then, at the end of the program, again call f with the same argument a , we know we will get the same result out. This is because we know that a cannot have changed and because we know that f only depends on a (for instance, it didn't increment a global counter). This property is called *referential transparency* and basically states that if two functions f and g produce the same values for the same arguments, then we may replace f with g (and vice-versa). referential transparency

■ NOTE ■ The exact definition of referential transparency is not generally agreed upon. The one given above is the one I like best. They all carry the same interpretation; the differences lie in how they are formalized.

3.1 Arithmetic

Let's begin our foray into Haskell with simple arithmetic. Start up your favorite interactive shell (Hugs or GHCi; see Chapter 2 for installation instructions). The shell will output to the screen a few lines talking about itself and what it's doing and then should finish with the cursor on a line reading:

```
Prelude>
```

From here, you can begin to evaluate *expressions*. An expression is basically something which has a value. For instance, the number 5 is an expression (its value is 5). Values can be built up from other values; for instance, $5 + 6$ is an expression (its value is 11). In fact, most simple arithmetic operations are supported by Haskell, including plus (+), minus (-), times (*), divided-by (/), exponentiation (^) and square-root (**sqrt**). You can experiment with these by asking the interactive shell to evaluate expressions and give you their value. In this way, a Haskell shell can be used as a powerful calculator. Try some of the following: expressions

```
Prelude> 5*4+3
23
Prelude> 5^5-2
3123
Prelude> sqrt 2
1.4142135623730951
Prelude> 5*(4+3)
35
```


We can see that in addition to the standard arithmetic operations, Haskell also allows grouping by parentheses, hence the difference between the values of $5*4+3$ and $5*(4+3)$. The reason for this is that the “understood” grouping of the first expression is operator precedence $(5*4)+3$, due to *operator precedence*.

Note also that parentheses aren’t required around function arguments. For instance, we simply wrote `sqrt 2`, not `sqrt (2)` as would be required in most other languages. You could write it with the parentheses, but in Haskell since function application is so common, parentheses aren’t required.

■ **WARNING** ■ Even though parentheses are not always needed, sometimes it is better to leave them in anyway; afterall, other people will probably have to read your code and if extra parentheses make the intent of the code clearer, use them.

Now try entering `2^5000`. Does it work?

■ **NOTE** ■ If you’re familiar with programming in other languages, you may find it odd that `sqrt 2` comes back with a decimal point (i.e., is a floating point number) even though the argument to the function seems to be an integer. This interchangeability of numeric types is due to Haskell’s system of *type classes* and will be discussed in grueling detail in Section 4.3).

Exercises

Exercise 1 *We’ve seen that multiplication binds more tightly than division. Can you think of a way to determine whether function application binds more or less tightly than multiplication?*

3.2 Pairs, Triples and More

In addition to single values, we might want to talk about multiple values at the same time. For instance, we may want to refer to a position by its x/y coordinate, which would be a pair of integers. To make a pair of integers is simple: you enclose the pair in parenthesis and separate them with a comma. Try the following:

```
Prelude> (5,3)
(5,3)
```

heterogeneous Here, we have a pair of integers, 5 and 3. Pairs are allowed to be heterogeneous, meaning that you can store, for instance, an integer in the first position and a string in the second.

There are two predefined functions which allow you to extract the first and second elements of a pair. They are, respectively, `fst` and `snd`. You can see how they work as follows:

```
Prelude> fst (5, "hello")
5
Prelude> snd (5, "hello")
"hello"
```

In addition to pairs, you can define triples, quadruples and so on. To define a triple and a quadruple respectively, we write:

```
Prelude> (1,2,3)
(1,2,3)
Prelude> (1,2,3,4)
(1,2,3,4)
```

And so on. In general, pairs, triples, and so on are called *tuples* and can store fixed amounts of heterogeneous data.

■ NOTE ■ The functions `fst` and `snd` won't work on anything longer than a pair; if you try to use them on a larger tuple, you will get a message complaining that there was a type error. What this means will be explained in Chapter 4.

Exercises

Exercise 2 Use a combination of `fst` and `snd` to extract the character out of the tuple `((1, 'a'), "foo")`.

3.3 Lists

The primary limitation of tuples is that they hold only a fixed number of elements: pairs hold two, triples hold three, and so on. A data structure which can hold an arbitrary number of elements is a **list**. Lists are assembled in a very similar fashion to tuples, except they use square brackets instead of parentheses. We can define a list like:

```
Prelude> [1,2]
[1,2]
Prelude> [1,2,3]
[1,2,3]
```

Lists don't need to have any elements. The empty list is simply `[]`.

Unlike tuples, we can very easily add an element on to the beginning of the list using the colon operator. The colon is called the “cons” operator; the process of adding an element is called “consing.” The etymology of this is that we're saying that we are *constructing* a new list from an element and an old list. We can see the cons operator in action in the following examples:

```
Prelude> 0:[1,2]
[0,1,2]
Prelude> 5:[1,2,3,4]
[5,1,2,3,4]
```

We can actually build any list simply by using the cons operator (the colon) and the empty list:

```
Prelude> 5:1:2:3:4:[ ]
[5,1,2,3,4]
```

In fact, the `[5,1,2,3,4]` syntax is simply syntactic sugar for the expression using the explicit cons operators and empty list. If we write something using the `[5,1,2,3,4]` notation, the compiler simply translates it to the other.

One further difference between lists and tuples is that while tuples are heterogeneous, lists must be homogenous. That means that you cannot have a list that holds both integers and strings. If you try to, you will get a type error.

Of course, lists don't have to just contain integers or strings; they can also contain tuples or even other lists. Tuples, similarly, can contain lists and other tuples. Try some of the following:

```
Prelude> [(1,1),(2,4),(3,9),(4,16)]
[(1,1),(2,4),(3,9),(4,16)]
Prelude> ([1,2,3,4],[5,6,7])
([1,2,3,4],[5,6,7])
```

There are two basic list functions: `head` and `tail`. The `head` function returns the first element off of a (non-empty) list and the `tail` function returns all but the first element off of a (non-empty) list, respectively.

To get the length of a list, you use the `length` function:

```
Prelude> length [1,2,3,4,10]
5
Prelude> head [1,2,3,4,10]
1
Prelude> length (tail [1,2,3,4,10])
4
```

3.3.1 Strings

In Haskell, a `String` is simply a list of `Chars`. So we can create the string “Hello” as:

```
Prelude> 'H':'e':'l':'l':'o':[]  
"Hello"
```

Lists (and, of course, strings) can be concatenated using the `++` operator:

```
Prelude> "Hello " ++ "World"  
"Hello World"
```

Additionally, non-string values can be converted to strings using the `show` function and strings can be converted to non-string values using the `read` function. Of course, if you try to read a value that’s malformed, you’ll get an error:

```
Prelude> "Five squared is " ++ show (5*5)  
"Five squared is 25"  
Prelude> read "5" + 3  
8  
Prelude> read "Hello" + 3  
Program error: Prelude.read: no parse
```

The exact error message is implementation dependent. However, the interpreter has inferred that you’re trying to add three to something. This means that when we execute `read "Hello"` we should get a number out. However, “Hello” cannot be parsed as a number, so an error is reported.

3.3.2 Simple List Functions

Much of the computation in Haskell programs is done by processing lists. There are three primary list-processing functions: `map`, `filter` and `foldr` (also `foldl`).

The `map` function takes as arguments a list of values and a function which should be applied to each of the values. For instance, there is a built-in function `toUpper` which takes as input a `Char` and produces a `Char` that is the upper-case version of the original argument. So, to convert an entire string (which is simply a list of characters) to upper case, we can map the `toUpper` function across the entire list:

```
Prelude> map toUpper "Hello World"  
"HELLO WORLD"
```

When you map across a list, the length of the list never changes – only the individual values in the list change.

To remove elements from the list, you can use the `filter` function. This function allows you to remove certain elements from a list depending on their value, but not their context. For instance, the function `isLower` tells you whether a given character is lower case. We can filter out all non-lowercase characters using this:

```
Prelude> filter isLower "Hello World"
"elloorld"
```

The function `foldr` takes a little more getting used to. `foldr` takes three arguments: a function, an initial value, and a list. The best way to think about `foldr` is that it replaces occurrences of the list cons operator `:` with the function parameter and replaces the empty list constructor `[]` with the initial value. Thus, if we have a list:

```
3 : 8 : 12 : 5 : []
```

And we apply `foldr (+) 0` to it, we get:

```
3 + 8 + 12 + 5 + 0
```

Which sums the list. We can test this:

```
Prelude> foldr (+) 0 [3,8,12,5]
28
```

We can do the same sort of thing to calculate the product of all the elements on a list:

```
Prelude> foldr (*) 1 [4,8,5]
160
```

We said earlier that folding is like replacing `:` with a particular function and `[]` with an initial element. This raises a question as to what happens when the function isn't associative (a function `.` is associative if $a \cdot (b \cdot c) = (a \cdot b) \cdot c$). When we write $4 \cdot 8 \cdot 5 \cdot 1$, we need to specify where to put the parentheses. Namely, do we mean $((4 \cdot 8) \cdot 5) \cdot 1$ or $4 \cdot (8 \cdot ((5 \cdot 1)))$? `foldr` assumes the function is right-associative (i.e., the correct bracketing is the latter). Thus, when we use it on a non-associative function (like minus), we can see the effect:

```
Prelude> foldr (-) 1 [4,8,5]
0
```

The exact derivation of this looks something like:

```
foldr (-) 1 [4,8,5]
==> 4 - (foldr (-) 1 [8,5])
==> 4 - (8 - foldr (-) 1 [5])
==> 4 - (8 - (5 - foldr (-) 1 []))
==> 4 - (8 - (5 - 1))
==> 4 - (8 - 4)
==> 4 - 4
==> 0
```

The `foldl` function goes the other way and effectively produces the opposite bracketing. `foldl` looks the same when applied, so we could have done summing just as well with `foldl`:

```
Prelude> foldl (+) 0 [3,8,12,5]
28
```

However, we get different results when using the non-associative function minus:

```
Prelude> foldl (-) 1 [4,8,5]
-16
```

This is because `foldl` uses the opposite bracketing. The way it accomplishes this is essentially by going all the way down the list, taking the last element and combining it with the initial value via the provided function. It then takes the second-to-last element in the list and combines it to this new value. It does so until there is no more list left.

The derivation here proceeds in the opposite fashion:

```
foldl (-) 1 [4,8,5]
==> foldl (-) (1 - 4) [8,5]
==> foldl (-) ((1 - 4) - 8) [5]
==> foldl (-) (((1 - 4) - 8) - 5) []
==> ((1 - 4) - 8) - 5
==> ((-3) - 8) - 5
==> (-11) - 5
==> -16
```

Note that once the `foldl` goes away, the parenthesization is exactly the opposite of the `foldr`.

■ NOTE ■ `foldl` is often more efficient than `foldr` for reasons we will discuss in Section 7.8. However, `foldr` can work on infinite lists while `foldl` cannot. This is because before `foldl` does anything, it has to go to the end of the list. On the other hand, `foldr` starts producing output immediately. For instance, `foldr (:) [] [1,2,3,4,5]` simply returns the same list. Even if the list were infinite, it would produce output. A similar function using `foldl` would fail to produce any output.

If this discussion of the folding functions doesn't entirely make sense, that's okay. We'll talk about them further in Section 7.8.

Exercises

Exercise 3 Use `map` to convert a string into a list of booleans, each element in the new list representing whether the original element was a lower-case character or not. That is, it should take the string “aBCde” and return `[True,False,False,True,True]`.

Exercise 4 Use the functions mentioned in this section (you will need two of them) to compute the number of lower-case letters in a string. For instance, on “aBCde” it should return 3.

Exercise 5 We’ve seen how to calculate sums and products using folding functions. Given that the function `max` returns the maximum of two numbers, write a function using a fold which will return the maximum value in a list (and zero if the list is empty). So, when applied to `[5,10,2,8,1]` it will return 10. Assume that the values in the list are always ≥ 0 .

Explain to yourself why it works.

Exercise 6 Write a function which takes a list of pairs of length at least 2 and returns the first component of the second element in the list. So, when provided with `[(5,'b'),(1,'c'),(6,'a')]`, it will return 1.

3.4 Source Code Files

As programmers, we don’t want to simply evaluate small little expressions like these – we want to sit down, write code in our editor of choice, save it, and then use it.

We already saw in Sections 2.2 and 2.3 how to write a Hello World program and compile it. Here, we show how to do use functions defined in a source-code file in the interactive environment. To do this, create a file called `Test.hs` and enter the following code:

```
module Test
  where

x = 5
y = (6, "Hello")
z = x * fst y
```

This is a very simple “program” written in Haskell. It defines a module called “Test” (in general module names should match file names; see Section 6 for more on this). In this module, there are three definitions: `x`, `y` and `z`. Once you’ve written and saved this file, in the directory in which you saved it, load this up in your favorite interpreter, by executing either of the following:

```
% hugs Test.hs
% ghci Test.hs
```

This will start Hugs or GHCi, respectively, and load the file. Alternatively, if you already have one of them loaded, you can use the “:load” command (or just “:l”) to load a module, as:

```
Prelude> :l Test.hs
...
Test>
```

Between the first and last line, various data will be printed telling you what the interpreter is doing. If any errors crop up, you probably mistyped something in the file; double check and then try again.

You’ll notice that where it used to say “Prelude” it now says “Test”. That means that Test is the current module. Which means, you’re probably thinking, “Prelude” must also be a module. Exactly correct. The Prelude module (usually simply referred to as “the Prelude”) is always loaded and contains the standard definitions (for instance,

the Prelude

the `:` operator for lists, or `+` or `*`, `fst`, `snd` and so on).

Now that we’ve loaded Test, we can use things which were defined in it. For example:

```
Test> x
5
Test> y
(6, "Hello")
Test> z
30
```

Perfect, just as we expected!

One final issue regards how to compile programs to stand-alone executables. In order for a program to be an executable, it *must* have the module name “Main” and must contain a function called “main”. So, if you go in to Test.hs and rename it to “Main” (change the line that reads `module Test` to `module Main`), we simply need to add a main function. Try this:

```
main = putStrLn "Hello World"
```

Now, save the file and we can compile it (refer back to Section 2 for information on how to do this for your compiler. For example, in GHC, you would say:

```
% ghc --make Test.hs -o test
```

■ NOTE ■ For Windows, it would be “-o test.exe”

This will create a file called “test” (or on Windows, “test.exe”) which you can then run.


```
% ./test
Hello World
```

■ NOTE ■ Or, on Windows:

```
C:\> test.exe
Hello World
```

3.5 Functions

Now that we've seen how to write code in a file, we can start writing functions. As you might have expected, functions are central to Haskell, as it is a *functional* language. This means that the evaluation of a program is simply the evaluation of a function.

We can write a simple function to square a number and enter it into our `Test.hs` file. We might define this as follows:

```
square x = x * x
```

In this function definition, we say that we're defining a function `square` who takes one argument (aka parameter) which we call `x`. We then say that the *value* of `square x` is equal to `x * x`.

Haskell also supports standard conditional expressions. For instance, we could define a function that returns `-1` if its argument is less than 0, `0` if its argument is 0 and `1` if its argument is greater than 0 (this is called the *signum* function):

```
signum x =
  if x < 0
  then -1
  else if x > 0
  then 1
  else 0
```

if/then/else The `if/then/else` construct in Haskell is very similar to that of most other programming languages, however you must have both a **then** and an **else** clause. It evaluates the condition (in this case `x < 0` and, if this evaluates to `True`, it evaluates the **then** condition; if the condition evaluated to `False`, it evaluates the **else** condition.

You can test this program by editing the file and loading it back into your interpreter. If `Test` is already the current module, instead of typing `:l Test.hs` again, you can simply type `:reload` or just `:r` to reload the current file. This is usually much faster.

case/of Haskell, like many other languages, also supports **case** constructions. These are

used when there are multiple values you want to check against (case expressions are actually quite a bit more powerful than this – see Section 7.4 for all the gory details).

Suppose we wanted to define a function that had a value of 1 if its argument were 0, a value of 5 if its argument were 1, a value of 2 if its argument were 2 and a value of -1 in all other instances. Writing this function using **if** statements would be long and very unreadable; so we write it using a **case** statement as follows (we call this function `f`):

```
f x =
  case x of
    0 -> 1
    1 -> 5
    2 -> 2
    _ -> -1
```

In this program, we’re defining `f` to take an argument `x` and then inspect the value of `x`. If it matches 0, the value of `f` is 1. If it matches 1, the value of `f` is 5. If it matches 2, then the value of `f` is 2 and if it hasn’t matched anything (the underscore can be thought of as a “wildcard” – it will match anything) by that point, the value of `f` is -1 .

The indentation here is important. Haskell uses a system called “layout” to structure its code (the programming language Python uses a similar system). The layout system allows you to write code without the explicit semicolons and braces that other languages like C and Java require.

■ **WARNING** ■ Because whitespace matters in Haskell, you need to be careful about whether you are using tabs or spaces. If you can configure your editor to never use tabs, that’s probably better. If not, make sure your tabs are always 8 spaces long, or you’re likely to run in to problems.

The general rule for layout is that an open-brace is inserted after the keywords **where**, **let**, **do** and **of**, and the column position at which the next command appears is remembered. From then on, every new line which is indented the same amount automatically gets a semicolon. If a following line is indented less, a close-brace is inserted. This may sound complicated, but if you follow the general rule of indenting after each of those keywords, you’ll never have to remember it (see Section 7.11 for a more complete discussion of layout).

Some people prefer not to use layout and write the braces and semicolons explicitly. This is perfectly acceptable. In this style, the above function might look like:

```
f x = case x of
      { 0 -> 1 ; 1 -> 5 ; 2 -> 2 ; _ -> 1 }
```

Of course, if you write the braces and semicolons explicitly, you’re free to structure the code as you wish. The following is also equally valid:

```
f x =  
  case x of { 0 -> 1 ;  
             1 -> 5 ; 2 -> 2  
             ; _ -> 1 }
```

Though structuring your code like this only serves to make it unreadable (in this case).

Functions can also be defined piece-wise, meaning that you can write one version of your function for certain parameters and then another version for other parameters. For instance, the above function `f` could also be written as:

```
f 0 = 1  
f 1 = 5  
f 2 = 2  
f _ = -1
```

Here, the order is important. If we had put the last line first, it would have matched every argument and `f` would return `-1` regardless of its argument (most compilers will warn you about this, though, saying something about overlapping patterns). If we had not included this last line, `f` would produce an error if anything other than `0`, `1` or `2` were applied to it (most compilers will warn you about this, too, saying something about incomplete patterns). This style of piece-wise definition is very popular and will be used quite frequently throughout this tutorial. These two definitions of `f` are actually equivalent – this piece-wise version is translated into the case expression.

More complicated functions can be built up from simpler functions using *function composition*. Function composition is simply taking the result of the application of one function and using that as an argument to another. We've already seen this way back when we were doing arithmetic (Section 3.1), when we wrote `5*4+3`. In this, we were evaluating `5*4` and then applying `+3` to the result. We can do the same thing with our `square` and `f` functions:

```
Test> square (f 1)  
25  
Test> square (f 2)  
4  
Test> f (square 1)  
5  
Test> f (square 2)  
-1
```

The result of each of these function applications is fairly straightforward. The parentheses around the inner function are necessary; otherwise in the first line, the interpreter would think you're trying to get the value of "square f" which has no meaning. Function application like this is fairly standard in most programming languages. There

is another, more mathematically oriented, way to express function composition, using the `.` (just a single period) function. This `.` function is supposed to look like the \circ operator in mathematics.

■ NOTE ■ In mathematics we write $f \circ g$ to mean “f following g”, in Haskell we write `f . g` also to mean “f following g.” The meaning of $f \circ g$ is simply that $(f \circ g)(x) = f(g(x))$. That is, applying the value x to the function $f \circ g$ is the same as applying it to g , taking the result, and then applying that to f .

The `.` function (called the function composition function), takes two functions and makes them in to one. For instance, if we write `(square . f)`, this means to create a new function which takes an argument, applies `f` to that argument and then applies `square` to the result. Conversely, `(f . square)` means to create a new function which takes an argument, applies `square` to that argument and then applies `f` to the result. We can see this by testing it as before:

```
Test> (square . f) 1
25
Test> (square . f) 2
4
Test> (f . square) 1
5
Test> (f . square) 2
-1
```

Here, we must enclose the function composition in parentheses; otherwise, it will think we’re trying in the first line to compose `square` with the value `f 1`, which makes no sense since `f 1` isn’t even a function.

It would probably be wise to take a little time-out to look at some of the functions which are defined in the Prelude. Undoubtedly at some point you will accidentally rewrite some already-existing function (I’ve done it more times than I can count), but if we can keep this to a minimum, that would save a lot of time. Here are some simple functions, some of which we’ve already seen:

- `sqrt` the square root function
- `id` the identity function: `id x = x`
- `fst` extracts the first element from a pair
- `snd` extracts the second element from a pair
- `null` tells you whether or not a list is empty
- `head` returns the first element on a non-empty list
- `tail` returns everything but the first element of a non-empty list
- `++` concatenates two lists
- `==` checks to see if two things are equal
- `/=` checks to see if two things are unequal

Here, we show example usages of each of these functions:

```

Prelude> sqrt 2
1.41421
Prelude> id "hello"
"hello"
Prelude> id 5
5
Prelude> fst (5,2)
5
Prelude> snd (5,2)
2
Prelude> null []
True
Prelude> null [1,2,3,4]
False
Prelude> head [1,2,3,4]
1
Prelude> tail [1,2,3,4]
[2,3,4]
Prelude> [1,2,3] ++ [4,5,6]
[1,2,3,4,5,6]
Prelude> [1,2,3] == [1,2,3]
True
Prelude> 'a' /= 'b'
True
Prelude> head []
Program error: {head []}

```

We can see that applying `head` to an empty list gives an error (the exact error message depends on whether you're using GHCi or Hugs – the shown error message is from Hugs).

3.5.1 Let Bindings

Often we wish to provide local declarations for use in our functions. For instance, if you remember back to your grade school mathematics courses, there is the following equation used to find the roots (zeros) of a polynomial of the form $ax^2 + bx + c = 0$: $x = (-b \pm \sqrt{b^2 - 4ac})/2a$. We could write the following function to compute the two values of x :

```

roots a b c =
    ((-b + sqrt(b*b - 4*a*c)) / (2*a)),
    ((-b - sqrt(b*b - 4*a*c)) / (2*a))

```

To remedy this problem, Haskell allows for local bindings. That is, we can create values inside of a function which only that function can see. For instance, we could create a local binding for `sqrt (b*b-4*a*c)` and call it, say, `det` and then use that in both places where `sqrt (b*b - 4*a*c)` occurred. We can do this using a **let/in** declaration:

```
roots a b c =
  let det = sqrt (b*b - 4*a*c)
  in  ((-b + det) / (2*a),
      (-b - det) / (2*a))
```

In fact, you can provide multiple declarations inside a `let`. Just make sure they're indented the same amount, or you will have layout problems:

```
roots a b c =
  let det = sqrt (b*b - 4*a*c)
      twice_a = 2*a
  in  ((-b + det) / twice_a,
      (-b - det) / twice_a)
```

3.5.2 Infix

Infix functions are ones which are composed of symbols, rather than letters. For instance, `(+)`, `(*)`, `(++)` are all infix functions. You can use them in non-infix mode by enclosing them in parentheses. Hence the two following expressions are the same:

```
Prelude> 5 + 10
15
Prelude> (+) 5 10
15
```

Similarly, non-infix functions (like `map`) can be made infix by enclosing them in backquotes (the ticks on the tilde key on American keyboards):

```
Prelude> map toUpper "Hello World"
"HELLO WORLD"
Prelude> toUpper `map` "Hello World"
"HELLO WORLD"
```

3.6 Comments

There are two types of comments in Haskell: line comments and block comments. Line comments begin with the token `--` and extend until the end of the line. Block comments begin with `{-` and extend to a corresponding `-}`. Block comments can be nested.

■ NOTE ■ The `--` in Haskell corresponds to `//` in C++ or Java, and `{-` and `-}` correspond to `/*` and `*/`.

Comments are used to explain in English what is going on in your program and are completely ignored by compilers and interpreters. For example:

```
module Test2
  where

main =
  putStrLn "Hello World"  -- write a string
                          -- to the screen

{- f is a function which takes an integer and
   produces integer.  {- this is an embedded
   comment -} the original comment extends to the
   matching end-comment token: -}
f x =
  case x of
    0 -> 1    -- 0 maps to 1
    1 -> 5    -- 1 maps to 5
    2 -> 2    -- 2 maps to 2
    _ -> -1   -- everything else maps to -1
```

This example programs shows the use of both line comments and (embedded) block comments.

3.7 Recursion

In imperative languages like C and Java, the most basic control structure is a loop (like a for loop). However, for loops don't make much sense in Haskell because they require destructive update (the index variable is constantly being updated). Instead, Haskell uses recursion.

A function is *recursive* if it calls itself (see Appendix B for more). Recursive functions exist also in C and Java, but are used less than they are in functional languages. The prototypical recursive function is the factorial function. In an imperative language, you might write this as something like:

```
int factorial(int n) {
  int fact = 1;
  for (int i=2; i <= n; i++)
    fact = fact * i;
  return fact;
}
```

While this code fragment will successfully compute factorials for positive integers, it somehow ignores the basic definition of factorial, usually given as:

$$n! = \begin{cases} 1 & n = 1 \\ n * (n - 1)! & \text{otherwise} \end{cases}$$

This definition itself is exactly a recursive definition: namely the value of $n!$ depends on the value of $(n - 1)!$. If you think of $!$ as a function, then it is calling itself. We can translate this definition almost verbatim into Haskell code:

```
factorial 1 = 1
factorial n = n * factorial (n-1)
```

This is likely the simplest recursive function you'll ever see, but it is correct.

■ NOTE ■ Of course, an imperative recursive version could be written:

```
int factorial(int n) {
    if (n == 1)
        return 1;
    else
        return n * factorial(n-1);
}
```

but this is likely to be much slower than the loop version in C.

Recursion can be a difficult thing to master. It is completely analogous to the concept of induction in mathematics (see Chapter B for a more formal treatment of this). However, usually a problem can be thought of having one or more base-cases and one or more recursive-cases. In the case of `factorial`, there is one base case (when $n = 1$) and one recursive case (when $n > 1$). For designing your own recursive algorithms it is often useful to try to differentiate these two cases.

Turning now to the task of exponentiation, suppose we have two positive integers a and b and we want to calculate a^b . This problem has a single base-case: namely when b is 1. The recursive case is when $b > 1$. We can write a general form as:

$$a^b = \begin{cases} a & b = 1 \\ a * a^{b-1} & \text{otherwise} \end{cases}$$

Again, this translates directly into Haskell code:

```
exponent a 1 = a
exponent a b = a * exponent a (b-1)
```

Just as we can define recursive functions on integers, so can we define recursive functions on lists. In this case, usually the base case is the empty list `[]` and the recursive case is a cons list (i.e., a value consed on to another list).

Consider the task of calculating the length of a list. We can again break this down into two cases: either we have an empty list or we have a non-empty list. Clearly the length of an empty list is zero. Furthermore, if we have a cons list, then the length of this list is just the length of its tail plus one. Thus, we can define a length function as (whenever we provide alternative definitions for standard Haskell functions, we prefix them with `my_` so the compiler doesn't get confused):

```
my_length [] = 0
my_length (x:xs) = 1 + my_length xs
```

Similarly, we can consider the `filter` function. Again, the base case is the empty list and the recursive case is a cons list. However, this time, we're choosing whether to keep an element depending on whether or not a particular predicate holds. We can define the filter function as:

```
my_filter p [] = []
my_filter p (x:xs) =
  if p x
  then x : my_filter p xs
  else my_filter p xs
```

In this code, when presented with an empty list, we simply return an empty list. This is because `filter` cannot add elements: it can only remove them.

When presented with a list of the form `(x:xs)`, we need to decide whether to keep the value `x` or not. To do this, we use an `if` statement and the predicate `p`. If `p x` is true, then we return a list which begins with `x` followed by the result of filtering the tail of the list. If `p x` is false, then we exclude `x` and return the result of filtering the tail of the list.

We can also define `map` and both fold functions using explicit recursion. See the exercises for the definition of `map` and Chapter 7 for the folds.

Exercises

Exercise 7 *The fibonacci sequence is defined by:*

$$F_n = \begin{cases} 1 & n = 1 \text{ or } n = 2 \\ F_{n-2} + F_{n-1} & \text{otherwise} \end{cases}$$

Write a recursive function `fib` which takes a positive integer `n` as a parameter and calculates F_n .

Exercise 8 *Define a recursive function `mult` which takes two positive integers `a` and `b` and returns `a*b`, but only uses addition (i.e., no fair just using multiplication). Begin by making a mathematical definition in the style of the previous exercise and the rest of this section.*

Exercise 9 *Define a recursive function `my_map` which behaves identically to the standard function `map`.*

3.8 Interactivity

If you are familiar with books on other (imperative) languages, you might be wondering by now why you haven't seen many of the standard programs written in tutorials of other languages, like one which asks a user for his name and then says "Hi" to the user by name. The reason for this is simple: Being a pure functional language, it is not entirely clear how one should handle things like user input.

After all, suppose you have a function which reads a string from the keyboard. If you call this function twice and the user types something the first time and something else the second time, then you no longer have a function, since it would return two different values. The solution to this was found in the depths of category theory, a branch of formal mathematics: monads. We're not yet ready to talk about monads formally, but for now, think of them simply as a convenient way to express things like input/output. We'll discuss them in this context much more in Chapter 5 and then about monads for monads' sake in Chapter 9. For now, let's just see how you can write interactive programs. For this section, completely forget that you ever heard the word "monad."

monads

Suppose we want to write a function that's interactive. The way to do this is to use the **do** keyword. This allows us to specify the order of operations (remember than normally in Haskell, since it's a lazy language, the order in which operations are evaluated is unspecified). So, to write a simple program which asks a user for his name and then address him directly, enter the following code into "Name.hs":

do notation

```
module Main
  where

import IO

main = do
  hSetBuffering stdin LineBuffering
  putStrLn "Please enter your name: "
  name <- getLine
  putStrLn ("Hello, " ++ name ++ ", how are you?")
```

You can then either load this code in your interpreter and execute "main" by simply typing "main" or you can compile it and run it from the command line. I'll show the results of the interactive approach:

```
Main> main
Please enter your name:
Hal
Hello, Hal, how are you?
Main>
```

And there's interactivity. Let's go back and look at the code a little, though. We name the module "Main" so we can compile it. We name the primary function "main"

so that the compile knows that this is the function to run when the program is run. On the fourth line, we import the `IO` library so that we get access to `IO` functions. On the seventh line, we start with `do`, telling Haskell that we're executing a sequence of commands.

The first command is `hSetBuffering stdin LineBuffering` which you should probably ignore for now (incidentally, this is only required by `GHC` – in `Hugs` you can get by without it). The necessity for this is because when `GHC` reads input, it expects to read it in rather large blocks. A typical person's name is nowhere near large enough to fill this block. Thus, when we try to read from `stdin`, it waits until it's gotten a whole block. We want to get rid of this, so we tell it to use `LineBuffering` instead of block buffering.

The next command is `putStrLn`, which prints a string to the screen. On the ninth line, we say `name <- getLine`. This would normally be written `name = getLine` but using the arrow instead of the equality sign shows that `getLine` isn't a real function and can return different values. This command means "run the action `getLine` and store the results in `name`."

The last line constructs a string using what we read in on the previous line and then prints it to the screen.

Another example of a function which isn't really a function would be one that returns a random value. In this context, a function that does this is called `randomRIO`. Using this, we can write a "guess the number" program. Enter the following code into "Guess.hs":

```
module Main
  where

import IO
import Random

main = do
  hSetBuffering stdin LineBuffering
  num <- randomRIO (1::Int, 100)
  putStrLn "I'm thinking of a number between 1 and 100"
  doGuessing num

doGuessing num = do
  putStrLn "Enter your guess:"
  guess <- getLine
  if (read guess) < num
    then do putStrLn "Too low!"
            doGuessing num
    else if read guess > num
          then do putStrLn "Too high!"
                  doGuessing num
    else do putStrLn "You Win!"
```

Let's examine this code. On the fifth line we write "import Random" to tell the compiler that we're going to be using some random functions (these aren't built into the Prelude). In the first line of `main`, we get a random number in the range (1, 100). We need to write `:: Int` to tell the compiler that we're using integers here, not floating point numbers or other numbers. We'll talk more about this in Section 4. On the next line, we tell the user what's going on and then on the last line of `main` we tell the compiler to execute the command `doGuessing`.

The `doGuessing` function takes as an argument the number the user is trying to guess. First, it asks the user to guess and then accepts their guess (which is a `String`) from the keyboard. The `if` statement checks first to see if their guess is too low. However, since `guess` is a string and `num` is an integer, we first need to convert `guess` to an integer by reading it. If they guessed too low, we inform them and then start `doGuessing` over again. If they didn't guess too low, we check to see if they guessed too high. If they did, we tell them and start `doGuessing` again. Otherwise, they didn't guess too low and they didn't guess too high, so they must have gotten it correct, we tell them that they won and exit. The fact that we exit is implicit in the fact that there are no commands following this. We don't need an explicit `return ()` statement.

You can either compile this code or load it into your interpreter and will get something like:

```
Main> main
I'm thinking of a number between 1 and 100
Enter your guess:
50
Too low!
Enter your guess:
75
Too low!
Enter your guess:
85
Too high!
Enter your guess:
80
Too high!
Enter your guess:
78
Too low!
Enter your guess:
79
You Win!
```

The recursive action we just saw doesn't actually return a value that we look at. In the case when it does, the "obvious" way to write the command is actually incorrect. Here, we will give the incorrect version, explain why it is wrong, then give the correct version.

Let's say we're writing a simple program which repeatedly asks the user to type in a few words. If at any point they enter the empty word (i.e., they just hit enter without typing anything), the program prints out everything they've typed up until that point and then exits. The primary function (actually, an action) in this program is one which asks them for a word, checks to see if it's empty, and then either continues or ends. The *incorrect* formulation of this might look something like:

```
askForWords = do
  putStrLn "Please enter a word:"
  word <- getLine
  if word == ""
    then return []
    else return (word : askForWords)
```

Before reading ahead, see if you can figure out what is wrong with the above code.

The error is on the last line, specifically with the term `word : askForWords`. Remember that when using `:`, we are making a list out of an element (in this case `word`) and another list (in this case, `askForWords`). However, `askForWords` is *not* a list. It is an action, which, when run, will produce a list. That means that before we can attach anything to the front, we need to run the action and get the result. In this case, we want to do something like:

```
askForWords = do
  putStrLn "Please enter a word:"
  word <- getLine
  if word == ""
    then return []
    else do
      rest <- askForWords
      return (word : rest)
```

Here, we first run `askForWords` and get the result and store it in the variable `rest`. Then, we return the list created from `word` and `rest`.

By now you should have a good understanding of how to write simple functions, compile them, test things in the interactive environment, and manipulate lists.

Exercises

Exercise 10 Write a program which will repeatedly ask the user for numbers until they type in zero, at which point it will tell them the sum of all the numbers, the product of all the numbers, and, for each number, its factorial. For instance, a session might look like:

```
Give me a number (or 0 to stop):
5
Give me a number (or 0 to stop):
```

```
8
Give me a number (or 0 to stop):
2
Give me a number (or 0 to stop):
0
The sum is 15
The product is 80
5 factorial is 120
8 factorial is 40320
2 factorial is 2
```

Hint: write an IO action which reads a number and, if it's zero, returns the empty list. If it's not zero, it recurses itself and then makes a list out of the number it just read and the result of the recursive call.

Chapter 4

Type Basics

Haskell uses a system of *static type checking*. This means that every expression in Haskell is assigned a *type*. For instance 'a' would have type Char, for “character.” Then, if you have a function which expects an argument of a certain type and you give it the wrong type, a compile-time error will be generated (that is, you will not be able to compile the program). This vastly reduces the number of bugs that can creep into your program.

Furthermore, Haskell uses a system of *type inference*. This means that you don't even need to specify the type of expressions. For comparison, in C, when you define a variable, you need to specify its type (for instance, int, char, etc.). In Haskell, you needn't do this – the type will be inferred from context.

■ NOTE ■ If you want, you certainly are allowed to explicitly specify the type of an expression; this often helps debugging. In fact, it is sometimes considered good style to explicitly specify the types of outermost functions.

Both Hugs and GHCi allow you to apply type inference to an expression to find its type. This is done by using the `:t` command. For instance, start up your favorite shell and try the following:

```
Prelude> :t 'c'  
'c' :: Char
```

This tells us that the expression 'c' has type Char (the double colon `::` is used throughout Haskell to specify types).

4.1 Simple Types

There are a slew of built-in types, including `Int` (for integers, both positive and negative), `Double` (for floating point numbers), `Char` (for single characters), `String` (for

strings), and others. We have already seen an expression of type Char; let's examine one of type String:

```
Prelude> :t "Hello"
"Hello" :: String
```

You can also enter more complicated expressions, for instance, a test of equality:

```
Prelude> :t 'a' == 'b'
'a' == 'b' :: Bool
```

You should note that even though this expression is false, it still has a type, namely the type Bool.

■NOTE■ Bool is short for Boolean (usually pronounced “boo-lee-uhn”, though I’ve heard “boo-leen” once or twice) and has two possible values: True and False.

You can observe the process of type checking and type inference by trying to get the shell to give you the type of an ill-typed expression. For instance, the equality operator requires that the type of both of its arguments are of the same type. We can see that Char and String are of different types by trying to compare a character to a string:

```
Prelude> :t 'a' == "a"
ERROR - Type error in application
*** Expression      : 'a' == "a"
*** Term           : 'a'
*** Type           : Char
*** Does not match : [Char]
```

The first line of the error (the line containing “Expression”) tells us the expression in which the type error occurred. The second line tells us which part of this expression is ill-typed. The third line tells us the inferred type of this term and the fourth line tells us what it needs to have matched. In this case, it says that type Char doesn't match the type [Char] (a list a characters – a string in Haskell is represented as a list of characters).

As mentioned before, you can explicitly specify the type of an expression using the :: operator. For instance, instead of "a" in the previous example, we could have written ("a"::String). In this case, this has no effect since there's only one possible interpretation of "a". However, consider the case of numbers. You can try:

```
Prelude> :t 5 :: Int
5 :: Int
Prelude> :t 5 :: Double
5 :: Double
```

Here, we can see that the number 5 can be instantiated as either an `Int` or a `Double`. What if we don't specify the type?

```
Prelude> :t 5
5 :: Num a => a
```

Not quite what you expected? What this means, briefly, is that if some type `a` is an *instance* of the `Num` class, then the type of the expression `5` can be of type `a`. If that made no sense, that's okay for now. In Section 4.3 we talk extensively about type classes (which is what this is). The way to read this, though, is to say "a being an instance of `Num` implies `a`."

Exercises

Exercise 11 *Figure out for yourself, and then verify the types of the following expressions, if they have a type. Also note if the expression is a type error:*

1. `'h':'e':'l':'l':'o':[]`
2. `[5,'a']`
3. `(5,'a')`
4. `(5::Int) + 10`
5. `(5::Int) + (10::Double)`

4.2 Polymorphic Types

Haskell employs a polymorphic type system. This essentially means that you can have *type variables*, which we have alluded to before. For instance, note that a function like `tail` doesn't care what the elements in the list are:

```
Prelude> tail [5,6,7,8,9]
[6,7,8,9]
Prelude> tail "hello"
"ello"
Prelude> tail ["the","man","is","happy"]
["man","is","happy"]
```

This is possible because `tail` has a polymorphic type: $[\alpha] \rightarrow [\alpha]$. That means it can take as an argument *any* list and return a value which is a list of the same type.

The same analysis can explain the type of `fst`:

```
Prelude> :t fst
forall a b . (a,b) -> a
```

Here, GHCi has made explicit the *universal quantification* of the type values. That is, it is saying that for all types a and b , `fst` is a function from (a, b) to a .

Exercises

Exercise 12 *Figure out for yourself, and then verify the types of the following expressions, if they have a type. Also note if the expression is a type error:*

1. `snd`
2. `head`
3. `null`
4. `head . tail`
5. `head . head`

4.3 Type Classes

We saw last section some strange typing having to do with the number five. Before we delve too deeply into the subject of type classes, let's take a step back and see some of the motivation.

4.3.1 Motivation

In many languages (C++, Java, etc.), there exists a system of overloading. That is, a function can be written that takes parameters of differing types. For instance, the canonical example is the equality function. If we want to compare two integers, we should use an integer comparison; if we want to compare two floating point numbers, we should use a floating point comparison; if we want to compare two characters, we should use a character comparison. In general, if we want to compare two things which have type α , we want to use an α -compare. We call α a *type variable* since it is a variable whose value is a type.

■ NOTE ■ In general, type variables will be written using the first part of the Greek alphabet: $\alpha, \beta, \gamma, \delta, \dots$

Unfortunately, this presents some problems for static type checking, since the type checker doesn't know which types a certain operation (for instance, equality testing) will be defined for. There are as many solutions to this problem as there are statically typed languages (perhaps a slight exaggeration, but not so much so). The one chosen in Haskell is the system of type classes. Whether this is the "correct" solution or the "best" solution of course depends on your application domain. It is, however, the one we have, so you should learn to love it.

4.3.2 Equality Testing

Returning to the issue of equality testing, what we want to be able to do is define a function `==` (the equality operator) which takes two parameters, each of the same type (call it α), and returns a boolean. But this function may not be defined for *every* type; just for some. Thus, we associate this function `==` with a type class, which we call **Eq**. If a specific type α belongs to a certain type class (that is, all functions associated with that class are implemented for α), we say that α is an *instance* of that class. For instance, **Int** is an instance of **Eq** since equality is defined over integers.

4.3.3 The Num Class

In addition to overloading operators like `==`, Haskell has overloaded numeric constants (i.e., 1, 2, 3, etc.). This was done so that when you type in a number like 5, the compiler is free to say 5 is an integer or floating point number as it sees fit. It defines the **Num** class to contain all of these numbers and certain minimal operations over them (addition, for instance). The basic numeric types (`Int`, `Double`) are defined to be instances of **Num**.

We have only skimmed the surface of the power (and complexity) of type classes here. There will be much more discussion of them in Section 8.4, but we need some more background before we can get there. Before we do that, we need to talk a little more about functions.

4.3.4 The Show Class

Another of the standard classes in Haskell is the **Show** class. Types which are members of the **Show** class have functions which convert values of that type to a string. This function is called `show`. For instance `show` applied to the integer 5 is the string `"5"`; `show` applied to the character `'a'` is the three-character string `"'a'"` (the first and last characters are apostrophes). `show` applied to a string simply puts quotes around it. You can test this in the interpreter:

```
Prelude> show 5
"5"
Prelude> show 'a'
"'a'"
Prelude> show "Hello World"
 "\"Hello World\" "
```

■ NOTE ■ The reason the backslashes appear in the last line is because the interior quotes are “escaped”, meaning that they are part of the string, not part of the interpreter printing the value. The actual string doesn’t contain the backslashes.

Some types are not instances of **Show**; functions for example. If you try to show a function (like `sqrt`), the compiler or interpreter will give you some cryptic error mes-

sage, complaining about a missing instance declaration or an illegal class constraint.

4.4 Function Types

In Haskell, functions are *first class values*, meaning that just as `1` or `'c'` are values which have a type, so are functions like `square` or `++`. Before we talk too much about functions, we need to make a short diversion into very theoretical computer science (don't worry, it won't be too painful) and talk about the lambda calculus.

4.4.1 Lambda Calculus

The name “Lambda Calculus”, while perhaps daunting, describes a fairly simple system for representing functions. The way we would write a squaring function in lambda calculus is: $\lambda x.x * x$, which means that we take a value, which we will call x (that's what “ $\lambda x.$ ” means) and then multiply it by itself. The λ is called “lambda abstraction.” In general, lambdas can only have one parameter. If we want to write a function that takes two numbers, doubles the first and adds it to the second, we would write: $\lambda x \lambda y.2 * x + y$. When we apply a value to a lambda expression, we remove the outermost λ and replace every occurrence of the lambda variable with the value. For instance, if we evaluate $(\lambda x.x * x)5$, we remove the lambda and replace every occurrence of x with `5`, yielding $(5 * 5)$ which is `25`.

In fact, Haskell is largely based on an extension of the lambda calculus, and these two expressions can be written directly in Haskell (we simply replace the λ with a backslash and the $.$ with an arrow; also we don't need to repeat the lambdas; and, of course, in Haskell we have to give them names if we're defining functions):

```
square = \x -> x*x
f = \x y -> 2*x + y
```

You can also evaluate lambda expressions in your interactive shell:

```
Prelude> (\x -> x*x) 5
25
Prelude> (\x y -> 2*x + y) 5 4
14
```

We can see in the second example that we need to give the lambda abstraction two arguments, one corresponding to `x` and the other corresponding to `y`.

4.4.2 Higher-Order Types

“Higher-Order Types” is the name given to functions. The type given to functions mimics the lambda calculus representation of the functions. For instance, the definition of `square` gives $\lambda x.x * x$. To get the type of this, we first ask ourselves what the type of `x`

is. Say we decide x is an `Int`. Then, we notice that the function `square` takes an `Int` and produces a value $x * x$. We know that when we multiply two `Int`s together, we get another `Int`, so the type of the results of `square` is also an `Int`. Thus, we say the type of `square` is `Int → Int`.

We can apply a similar analysis to the function `f` above. The value of this function (remember, functions *are* values) is something which takes a value x and given that value, produces a *new* value, which takes a value y and produces $2 * x + y$. For instance, if we take `f` and apply only one number to it, we get $(\lambda x \lambda y. 2x + y) 5$ which becomes our new value $\lambda y. 2(5) + y$, where all occurrences of x have been replaced with the applied value, 5.

So we know that `f` takes an `Int` and produces a value of some type, of which we're not sure. But we know the type of this value is the type of $\lambda y. 2(5) + y$. We apply the above analysis and find out that this expression has type `Int → Int`. Thus, `f` takes an `Int` and produces something which has type `Int → Int`. So the type of `f` is `Int → (Int → Int)`.

■ NOTE ■ The parentheses are not necessary; in function types, if you have $\alpha \rightarrow \beta \rightarrow \gamma$ it is assumed that $\beta \rightarrow \gamma$ is grouped. If you want the other way, with $\alpha \rightarrow \beta$ grouped, you need to put parentheses around them.

This isn't entirely accurate. As we saw before, numbers like 5 aren't really of type `Int`, they are of type `Num` $a \Rightarrow a$.

We can easily find the type of Prelude functions using “:t” as before:

```
Prelude> :t head
head :: [a] -> a
Prelude> :t tail
tail :: [a] -> [a]
Prelude> :t null
null :: [a] -> Bool
Prelude> :t fst
fst :: (a,b) -> a
Prelude> :t snd
snd :: (a,b) -> b
```

We read this as: “head” is a function that takes a list containing values of type “a” and gives back a value of type “a”; “tail” takes a list of “a”s and gives back another list of “a”s; “null” takes a list of “a”s and gives back a boolean; “fst” takes a pair of type “(a,b)” and gives back something of type “a”, and so on.

■ NOTE ■ Saying that the type of `fst` is $(a, b) \rightarrow a$ does not necessarily mean that it simply gives back the first element; it only means that it gives back something with the *same type* as the first element.

We can also get the type of operators like `+` and `*` and `++` and `:`; however, in order to do this we need to put them in parentheses. In general, any function which is used *infix* (meaning in the middle of two arguments rather than before them) must be put in parentheses when getting its type.

```
Prelude> :t (+)
(+) :: Num a => a -> a -> a
Prelude> :t (*)
(*) :: Num a => a -> a -> a
Prelude> :t (++)
(++) :: [a] -> [a] -> [a]
Prelude> :t (:)
(:) :: a -> [a] -> [a]
```

The types of `+` and `*` are the same, and mean that `+` is a function which, for some type `a` which is an instance of `Num`, takes a value of type `a` and produces another function which takes a value of type `a` and produces a value of type `a`. In short hand, we might say that `+` takes two values of type `a` and produces a value of type `a`, but this is less precise.

The type of `++` means, in shorthand, that, for a given type `a`, `++` takes two lists of `a` and produces a new list of `a`s. Similarly, `:` takes a value of type `a` and another value of type `[a]` (list of `a`s) and produces another value of type `[a]`.

4.4.3 That Pesky IO Type

You might be tempted to try getting the type of a function like `putStrLn`:

```
Prelude> :t putStrLn
putStrLn :: String -> IO ()
Prelude> :t readFile
readFile :: FilePath -> IO String
```

What in the world is that `IO` thing? It's basically Haskell's way of representing that these functions aren't really functions. They're called "IO Actions" (hence the `IO`). The immediate question which arises is: okay, so how do I get rid of the `IO`. In brief, you can't directly remove it. That is, you cannot write a function with type `IO String -> String`. The only way to use things with an `IO` type is to combine them with other functions using (for example), the **do** notation.

For example, if you're reading a file using `readFile`, presumably you *want* to do something with the string it returns (otherwise, why would you read the file in the first place). Suppose you have a function `f` which takes a `String` and produces an `Int`. You can't directly apply `f` to the result of `readFile` since the input to `f` is `String` and the output of `readFile` is `IOString` and these don't match. However, you can combine these as:

```
main = do
  s <- readFile "somefile"
  let i = f s
  putStrLn (show i)
```

Here, we use the arrow convention to “get the string out of the IO action” and then apply `f` to the string (called `s`). We then, for example, print `i` to the screen. Note that the **let** here doesn’t have a corresponding **in**. This is because we are in a **do** block. Also note that we don’t write `i <- f s` because `f` is just a normal function, not an IO action.

4.4.4 Explicit Type Declarations

It is sometimes desirable to explicitly specify the types of some elements or functions, for one (or more) of the following reasons:

- Clarity
- Speed
- Debugging

Some people consider it good software engineering to specify the types of all top-level functions. If nothing else, if you’re trying to compile a program and you get type errors that you cannot understand, if you declare the types of some of your functions explicitly, it may be easier to figure out where the error is.

Type declarations are written separately from the function definition. For instance, we could explicitly type the function `square` as in the following code (an explicitly declared type is called a *type signature*):

```
square :: Num a => a -> a
square x = x*x
```

These two lines do not even have to be next to each other. However, the type that you specify must match the inferred type of the function definition (or be more specific). In this definition, you could apply `square` to anything which is an instance of `Num`: `Int`, `Double`, etc. However, if you knew apriori that `square` were only going to be applied to value of type `Int`, you could *refine* its type as:

```
square :: Int -> Int
square x = x*x
```

Now, you could only apply `square` to values of type `Int`. Moreover, with this definition, the compiler doesn’t have to generate the general code specified in the original

function definition since it knows you will only apply `square` to `Int`s, so it may be able to generate faster code.

If you have extensions turned on (“-98” in Hugs or “-fglasgow-exts” in GHC(i)), you can also add a type signature to expressions and not just functions. For instance, you could write:

```
square (x :: Int) = x*x
```

which tells the compiler that `x` is an `Int`; however, it leaves the compiler alone to infer the type of the rest of the expression. What is the type of `square` in this example? Make your guess then you can check it either by entering this code into a file and loading it into your interpreter or by asking for the type of the expression:

```
Prelude> :t (\(x :: Int) -> x*x)
```

since this lambda abstraction is equivalent to the above function declaration.

4.4.5 Functional Arguments

In Section 3.3 we saw examples of functions taking other functions as arguments. For instance, `map` took a function to apply to each element in a list, `filter` took a function that told it which elements of a list to keep, and `foldl` took a function which told it how to combine list elements together. As with every other function in Haskell, these are well-typed.

Let’s first think about the `map` function. It’s job is to take a list of elements and produce another list of elements. These two lists don’t necessarily have to have the same types of elements. So `map` will take a value of type `[a]` and produce a value of type `[b]`. How does it do this? It uses the user-supplied function to convert. In order to convert an `a` to a `b`, this function must have type `a → b`. Thus, the type of `map` is $(a \rightarrow b) \rightarrow [a] \rightarrow [b]$, which you can verify in your interpreter with “:t”.

We can apply the same sort of analysis to `filter` and discern that it has type $(a \rightarrow \text{Bool}) \rightarrow [a] \rightarrow [a]$. As we presented the `foldl` function, you might be tempted to give it type $(a \rightarrow a \rightarrow a) \rightarrow a \rightarrow [a] \rightarrow a$, meaning that you take a function which combines two `a`s into another one, an initial value of type `a`, a list of `a`s to produce a final value of type `a`. In fact, `foldl` has a more general type: $(a \rightarrow b \rightarrow a) \rightarrow a \rightarrow [b] \rightarrow a$. So it takes a function which turn an `a` and a `b` into an `a`, an initial value of type `a` and a list of `b`s. It produces an `a`.

To see this, we can write a function `count` which counts how many members of a list satisfy a given constraint. You can of course you `filter` and `length` to do this, but we will also do it using `foldr`:

```
module Count
  where
import Char
```

```
count1 p l = length (filter p l)
count2 p l = foldr (\x c -> if p x then c+1 else c) 0 l
```

The functioning of `count1` is simple. It filters the list `l` according to the predicate `p`, then takes the length of the resulting list. On the other hand, `count2` uses the initial value (which is an integer) to hold the current count. For each element in the list `l`, it applies the lambda expression shown. This takes two arguments, `c` which holds the current count and `x` which is the current element in the list that we're looking at. It checks to see if `p` holds about `x`. If it does, it returns the new value `c+1`, increasing the count of elements for which the predicate holds. If it doesn't, it just returns `c`, the old count.

Exercises

Exercise 13 *Figure out for yourself, and then verify the types of the following expressions, if they have a type. Also note if the expression is a type error:*

1. `\x -> [x]`
2. `\x y z -> (x,y:z:[])`
3. `\x -> x + 5`
4. `\x -> "hello, world"`
5. `\x -> x 'a'`
6. `\x -> x x`
7. `\x -> x + x`

4.5 Data Types

Tuples and lists are nice, common ways to define structured values. However, it is often desirable to be able to define our own data structures and functions over them. So-called “datatypes” are defined using the **data** keyword.

4.5.1 Pairs

For instance, a definition of a pair of elements (much like the standard, build-in pair type) could be:

```
data Pair a b = Pair a b
```

Let's walk through this code one word at a time. First we say "data" meaning that we're defining a datatype. We then give the name of the datatype, in this case, "Pair." The "a" and "b" that follow "Pair" are type parameters, just like the "a" is the type of the function map. So up until this point, we've said that we're going to define a data structure called "Pair" which is parameterized over two types, a and b.

After the equals sign, we specify the *constructors* of this data type. In this case, there is a single constructor, "Pair" (this doesn't necessarily have to have the same name as the type, but in this case it seems to make more sense). After this pair, we again write "a b", which means that in order to construct a Pair we need two values, one of type a and one of type b.

This definition introduces a function, `Pair :: a -> b -> Pair a b` that you can use to construct Pairs. If you enter this code into a file and load it, you can see how these are constructed:

```
Datatypes> :t Pair
Pair :: a -> b -> Pair a b
Datatypes> :t Pair 'a'
Pair 'a' :: a -> Pair Char a
Datatypes> :t Pair 'a' "Hello"
:t Pair 'a' "Hello"
Pair 'a' "Hello" :: Pair Char [Char]
```

So, by giving Pair two values, we have completely constructed a value of type Pair. We can write functions involving pairs as:

```
pairFst (Pair x y) = x
pairSnd (Pair x y) = y
```

In this, we've used the *pattern matching* capabilities of Haskell to look at a pair and extract values from it. In the definition of `pairFst` we take an entire Pair and extract the first element; similarly for `pairSnd`. We'll discuss pattern matching in much more detail in Section 7.4.

Exercises

Exercise 14 Write a data type declaration for `Triple`, a type which contains three elements, all of different types. Write functions `tripleFst`, `tripleSnd` and `tripleThr` to extract respectively the first, second and third elements.

Exercise 15 Write a datatype `Quadruple` which holds four elements. However, the first two elements must be the same type and the last two elements must be the same type. Write a function `firstTwo` which returns a list containing the first two elements and a function `lastTwo` which returns a list containing the last two elements. Write type signatures for these functions

4.5.2 Multiple Constructors

We have seen an example of the data type with one constructor: `Pair`. It is also possible (and extremely useful) to have multiple constructors.

Let us consider a simple function which searches through a list for an element satisfying a given predicate and then returns the first element satisfying that predicate. What should we do if none of the elements in the list satisfy the predicate? A few options are listed below:

- Raise an error
- Loop indefinitely
- Write a check function
- Return the first element
- ...

Raising an error is certainly an option (see Section 10.1 to see how to do this). The problem is that it is difficult/impossible to recover from such errors. Looping indefinitely is possible, but not terribly useful. We could write a sister function which checks to see if the list contains an element satisfying a predicate and leave it up to the user to always use this function first. We could return the first element, but this is very ad-hoc and difficult to remember.

The fact that there is no basic option to solve this problem simply means we have to think about it a little more. What are we trying to do? We're trying to write a function which might succeed and might not. Furthermore, if it does succeed, it returns some sort of value. Let's write a datatype:

```
data Maybe a = Nothing
             | Just a
```

This is one of the most common datatypes in Haskell and is defined in the Prelude.

Here, we're saying that there are two possible ways to create something of type `Maybe a`. The first is to use the nullary constructor `Nothing`, which takes no arguments (this is what "nullary" means). The second is to use the constructor `Just`, together with a value of type `a`.

The `Maybe` type is useful in all sorts of circumstances. For instance, suppose we want to write a function (like `head`) which returns the first element of a given list. However, we don't want the program to die if the given list is empty. We can accomplish this with a function like:

```
firstElement :: [a] -> Maybe a
firstElement []      = Nothing
firstElement (x:xs) = Just x
```

The type signature here says that `firstElement` takes a list of `a`s and produces something with type `Maybe a`. In the first line of code, we match against the empty list `[]`. If this match succeeds (i.e., the list is, in fact, empty), we return `Nothing`. If the first match fails, then we try to match against `x:xs` which must succeed. In this case, we return `Just x`.

For our `findElement` function, we represent failure by the value `Nothing` and success with value `a` by `Just a`. Our function might look something like this:

```
findElement :: (a -> Bool) -> [a] -> Maybe a
findElement p [] = Nothing
findElement p (x:xs) =
    if p x then Just x
    else findElement p xs
```

The first line here gives the type of the function. In this case, our first argument is the predicate (and takes an element of type `a` and returns `True` if and only if the element satisfies the predicate); the second argument is a list of `a`s. Our return value is *maybe* an `a`. That is, if the function succeeds, we will return `Just a` and if not, `Nothing`.

Another useful datatype is the `Either` type, defined as:

```
data Either a b = Left a
                | Right b
```

This is a way of expressing alternation. That is, something of type `Either a b` is *either* a value of type `a` (using the `Left` constructor) or a value of type `b` (using the `Right` constructor).

Exercises

Exercise 16 Write a datatype `Tuple` which can hold one, two, three or four elements, depending on the constructor (that is, there should be four constructors, one for each number of arguments). Also provide functions `tuple1` through `tuple4` which take a tuple and return `Just` the value in that position, or `Nothing` if the number is invalid (i.e., you ask for the `tuple4` on a tuple holding only two elements).

Exercise 17 Based on our definition of `Tuple` from the previous exercise, write a function which takes a `Tuple` and returns either the value (if it's a one-tuple), a Haskell-pair (i.e., `('a', 5)`) if it's a two-tuple, a Haskell-triple if it's a three-tuple or a Haskell-quadruple if it's a four-tuple. You will need to use the `Either` type to represent this.

4.5.3 Recursive Datatypes

We can also define *recursive datatypes*. These are datatypes whose definitions are based on themselves. For instance, we could define a list datatype as:

```
data List a = Nil
           | Cons a (List a)
```

In this definition, we have defined what it means to be of type `List a`. We say that a list is either empty (`Nil`) or it's the `Cons` of a value of type `a` and another value of type `List a`. This is almost identical to the actual definition of the list datatype in Haskell, except that uses special syntax where `[]` corresponds to `Nil` and `:` corresponds to `Cons`. We can write our own `length` function for our lists as:

```
listLength Nil = 0
listLength (Cons x xs) = 1 + listLength xs
```

This function is slightly more complicated and uses *recursion* to calculate the length of a `List`. The first line says that the length of an empty list (a `Nil`) is 0. This much is obvious. The second line tells us how to calculate the length of a non-empty list. A non-empty list must be of the form `Cons x xs` for some values of `x` and `xs`. We know that `xs` is another list and we know that whatever the length of the current list is, it's the length of its tail (the value of `xs`) plus one (to account for `x`). Thus, we apply the `listLength` function to `xs` and add one to the result. This gives us the length of the entire list.

Exercises

Exercise 18 Write functions `listHead`, `listTail`, `listFoldl` and `listFoldr` which are equivalent to their Prelude twins, but function on our `List` datatype. Don't worry about exceptional conditions on the first two.

4.5.4 Binary Trees

We can define datatypes that are more complicated than lists. Suppose we want to define a structure that looks like a binary tree. A binary tree is a structure that has a single root node; each node in the tree is either a "leaf" or a "branch." If it's a leaf, it holds a value; if it's a branch, it holds a value and a left child and a right child. Each of these children is another node. We can define such a data type as:

```
data BinaryTree a
  = Leaf a
  | Branch (BinaryTree a) a (BinaryTree a)
```

In this datatype declaration we say that a `BinaryTree` of `as` is either a `Leaf` which holds an `a`, or it's a branch with a left child (which is a `BinaryTree` of `as`), a

node value (which is an `a`), and a right child (which is also a `BinaryTree` of `as`). It is simple to modify the `listLength` function so that instead of calculating the length of lists, it calculates the number of nodes in a `BinaryTree`. Can you figure out how? We can call this function `treeSize`. The solution is given below:

```
treeSize (Leaf x) = 1
treeSize (Branch left x right) =
  1 + treeSize left + treeSize right
```

Here, we say that the size of a leaf is 1 and the size of a branch is the size of its left child, plus the size of its right child, plus one.

Exercises

Exercise 19 Write a function `elements` which returns the elements in a `BinaryTree` in a bottom-up, left-to-right manner (i.e., the first element returned in the left-most leaf, followed by its parent's value, followed by the other child's value, and so on). The result type should be a normal Haskell list.

Exercise 20 Write a fold function for `BinaryTrees` and rewrite `elements` in terms of it (call the new one `elements2`).

4.5.5 Enumerated Sets

You can also use datatypes to define things like enumerated sets, for instance, a type which can only have a constrained number of values. We could define a color type:

```
data Color
  = Red
  | Orange
  | Yellow
  | Green
  | Blue
  | Purple
  | White
  | Black
```

This would be sufficient to deal with simple colors. Suppose we were using this to write a drawing program, we could then write a function to convert between a `Color` and a RGB triple. We can write a `colorToRGB` function, as:

```
colorToRGB Red    = (255,0,0)
colorToRGB Orange = (255,128,0)
colorToRGB Yellow = (255,255,0)
colorToRGB Green  = (0,255,0)
colorToRGB Blue   = (0,0,255)
```

```
colorToRGB Purple = (255,0,255)
colorToRGB White = (255,255,255)
colorToRGB Black = (0,0,0)
```

If we wanted also to allow the user to define his own custom colors, we could change the `Color` datatype to something like:

```
data Color
  = Red
  | Orange
  | Yellow
  | Green
  | Blue
  | Purple
  | White
  | Black
  | Custom Int Int Int -- R G B components
```

And add a final definition for `colorToRGB`:

```
colorToRGB (Custom r g b) = (r,g,b)
```

4.5.6 The Unit type

A final useful datatype defined in Haskell (from the Prelude) is the unit type. Its definition is:

```
data () = ()
```

The only true value of this type is `()`. This is essentially the same as a *void* type in a language like C or Java and will be useful when we talk about IO in Chapter 5.

We'll dwell much more on data types in Sections 7.4 and 8.3.

4.6 Continuation Passing Style

There is a style of functional programming called “Continuation Passing Style” (also simply “CPS”). The idea behind CPS is to pass around as a function argument what to do next. I will handwave through an example which is too complex to write out at this point and then give a real example, though one with less motivation.

Consider the problem of parsing. The idea here is that we have a sequence of tokens (words, letters, whatever) and we want to ascribe structure to them. The task of converting a string of Java tokens to a Java abstract syntax tree is an example of a

parsing problem. So is the task of taking an English sentence and creating a parse tree (though the latter is quite a bit harder).

Suppose we're parsing something like C or Java where functions take arguments in parentheses. But for simplicity, assume they are not separated by commas. That is, a function call looks like `myFunction(x y z)`. We want to convert this into something like a pair containing first the string "myFunction" and then a list with three string elements: "x", "y" and "z".

The general approach to solving this would be to write a function which parses function calls like this one. First it would look for an identifier ("myFunction"), then for an open parenthesis, then for zero or more identifiers, then for a close parenthesis.

One way to do this would be to have two functions:

```
parseFunction  ::
  [Token] -> Maybe ((String, [String]), [Token])

parseIdentifier ::
  [Token] -> Maybe (String, [Token])
```

The idea would be that if we call `parseFunction`, if it doesn't return `Nothing`, then it returns the pair described earlier, together with whatever is left after parsing the function. Similarly, `parseIdentifier` will parse one of the arguments. If it returns `Nothing`, then it's not an argument; if it returns `Just` something, then that something is the argument paired with the rest of the tokens.

What the `parseFunction` function would do is to parse an identifier. If this fails, it fails itself. Otherwise, it continues and tries to parse a open parenthesis. If that succeeds, it repeatedly calls `parseIdentifier` until that fails. It then tries to parse a close parenthesis. If that succeeds, then it's done. Otherwise, it fails.

There is, however, another way to think about this problem. The advantage to this solution is that functions no longer need to return the remaining tokens (which tends to get ugly). Instead of the above, we write functions:

```
parseFunction  ::
  [Token] -> ((String, [String]) -> [Token] -> a) ->
  ([Token] -> a) -> a

parseIdentifier ::
  [Token] -> (String -> [Token] -> a) ->
  ([Token] -> a) -> a
```

Let's consider `parseIdentifier`. This takes three arguments: a list of tokens and two *continuations*. The first continuation is what to do when you succeed. The second continuation is what to do if you fail. What `parseIdentifier` does, then, is try to read an identifier. If this succeeds, it calls the first continuation with that identifier and the remaining tokens as arguments. If reading the identifier fails, it calls the second continuation with all the tokens.

Now consider `parseFunction`. Recall that it wants to read an identifier, an open parenthesis, zero or more identifiers and a close parenthesis. Thus, the first thing it does is call `parseIdentifier`. The first argument it gives is the list of tokens. The first continuation (which is what `parseIdentifier` should do if it succeeds) is in turn a function which will look for an open parenthesis, zero or more arguments and a close parenthesis. The second argument (the failure argument) is just going to be the failure function given to `parseFunction`.

Now, we simply need to define this function which looks for an open parenthesis, zero or more arguments and a close parenthesis. This is easy. We write a function which looks for the open parenthesis and then calls `parseIdentifier` with a success continuation that looks for more identifiers, and a “failure” continuation which looks for the close parenthesis (note that this failure doesn’t really mean failure – it just means there are no more arguments left).

I realize this discussion has been quite abstract. I would willingly give code for all this parsing, but it is perhaps too complex at the moment. Instead, consider the problem of folding across a list. If we’re willing to assume that the list is nonempty (actually, this code can be retrofitted to also work for empty lists), we can write a CPS-style fold as:

```
fold f z l = fold' f (\y -> f z y) l
fold' f g [x] = g x
fold' f g (x:xs) = g (fold' f (\y -> f x y) xs)
```

In this code, `fold` basically turns the accumulator argument `z` into a continuation. This continuation `g` says “give me a value and then I’ll combine it with what I had before and give you the result). The function `fold'` first matches against a list with exactly one element: `x`. If this match succeeds, it says to `g`: “Here is a value, `x`. Give me the result.” Of course, `g` obliges.

In the recursive case, `fold'` builds a new continuation that says: “give me a value and I’ll combine it with what I see right now.” It then recurses on itself and gives the result of this recursion to the original function `g` as a result.

You can test for yourself that this fold immitates one of `foldl/foldr`: try to figure out which one without using an interpreter. Also, try writing `map` or `filter` using CPS.

Chapter 5

Basic Input/Output

As we mentioned earlier, it is difficult to think of a good, clean way to integrate things like input/output into a pure functional language. Before we give the solution, let's take a step back and think about why such a thing is difficult.

Any IO library should provide a host of functions, containing (at a minimum) things like:

- print a string to the screen
- read a string from a keyboard
- write data to a file
- read data from a file

There are two issues here. Let us consider the first two examples first and think about what their types should be. Certainly the first item (I hesitate to call it a “function”) should take a `String` argument and produce something, but what should it produce? It could produce a unit `()`, since there is essentially no return value from printing a string. The second item similarly should return a string, but how this happens is unclear.

We want both of these items to be functions. But they are by definition not functions. The item which reads a string from the keyboard cannot be a function, as it will not return the same thing every time. And if the first function simply returns `()` every time, there should be no problem with replacing it with a function `f _ = ()` because of referential transparency. But clearly this does not have the desired effect.

5.1 The RealWorld Solution

In a sense, the reason that these items are not functions is that they interact with the “real world.” Their values depend directly on the real world. Supposing we had a type `RealWorld`, we might write these functions as having type:

```
printAString :: RealWorld -> String -> RealWorld
readAString  :: RealWorld -> (RealWorld, String)
```

That is, `printAString` takes a current state of the world and a string to print; it then modifies the state of the world in such a way that that string is now printed and returns this new value. Similarly, `readAString` takes a current state of the world and returns a *new* state of the world paired with the `String` that was typed.

This would be a possible way to do IO, though it is more than somewhat unweildy. In this style, our “Name.hs” program from Section 3.8 would look something like (assuming an initial `RealWorld` state were an argument to `main`):

```
main rW =
  let rW' = printAString rW "Please enter your name: "
      (rW'',name) = readAString rW'
  in  printAString rW''
      ("Hello, " ++ name ++ ", how are you?")
```

This is not only hard to read, but prone to error if you accidentally use the wrong version of the real world. It also doesn’t model the fact that the program below makes no sense:

```
main rW =
  let rW' = printAString rW "Please enter your name: "
      (rW'',name) = readAString rW'
  in  printAString rW'
      ("Hello, " ++ name ++ ", how are you?")
```

In this program the reference to `rW''` on the last line has been changed to a reference to `rW'`. It is completely unclear what this program should do. Clearly it must read a string in order to get a value for `name` to be printed. But that means that the real world has been updated. However, we then try to ignore this update by using an “old version” of the real world. There is clearly something very bad going on here.

5.2 The CPS Solution

A solution to this “old real world reference” problem is to use a continuation passing style. That is, the `printAString` function takes a `String` as well as something to do after this string has been printed. Thus, we never have access to the “old world.” This results in types like:

```
printAString2 ::
  String -> (RealWorld -> a) -> RealWorld -> a
```

```
readAString2 ::
  (String -> RealWorld -> a) -> RealWorld -> a
```

These functions could then be used to build our class name program as:

```
main rw =
  printAString2 "Please enter your name: "
  (readAString2 (\name ->
    printAString2
      ("Hello, " ++ name ++ ", how are you?")))) rw
```

Here, first we print the query string then pass a function which reads a string and passes the read string to a function which prints the greeting. The advantage to this style is that we now have no way of referencing the actual real world and so we don't have to worry about people accidentally (or maliciously) using old states of the world.

We can even define these functions in terms of the old ones (presumably a compiler would do this internally and simply not allow us access to the “basic” functions):

```
printAString2 s f rw = f (printAString rw s)

readAString2 f rw =
  let (rw', s) = readAString rw
  in f s rw'
```

The disadvantage to this style is that it tends to be difficult to read. Even this simple program is difficult to understand at first glance (or perhaps even second glance).

Suffice it to say that doing IO operations in a pure lazy functional language is not trivial.

The breakthrough came when Phil Wadler decided that monads would be a good way to think about IO computations. In fact, monads are able to express much more than just the simple operations described above. We can use them to express things like concurrence, exceptions, IO, non-determinism and much more. Moreover, there is nothing special about them: they can be defined *within* Haskell with no special handling from the compiler (though compilers often choose to optimize monadic operations).

5.3 Actions

As pointed out before, we cannot think of things like “print a string to the screen” or “read data from a file” as functions, since they are not (in the pure mathematical sense). Therefore, we give them another name: *actions*. Not only do we give them a special name, we give them a special type. One particularly useful action is `putStrLn` which prints a string to the screen. This action has type:

```
putStrLn :: String -> IO ()
```

As expected, `putStrLn` takes a string argument. What it returns is of type `IO ()`. This means that this function is actually an action (that is what the `IO` means). Furthermore, when this action is *evaluated* (or “run”), the result will have type `()`.

■ NOTE ■ Actually, this type means that `putStrLn` is an action *within the IO monad*, but we will gloss over this for now.

You can probably already guess the type of `getLine`:

```
getLine :: IO String
```

This means that `getLine` is an `IO` action which, when run, will have type `String`.

The question immediately arises: how do you “run” an action. This is something which is up to the compiler. You cannot actually run an action yourself; instead, a program is itself a single action which is run when the compiled program is executed. Thus, the compiler requires that the `main` function have type `IO ()`, which means that it is an `IO` action which returns nothing. The compiled code then executes this action.

However, while you are not allowed to run actions yourself, you *are* allowed to combine actions. In fact, we have already seen one way to do this using the **do** notation (how to “really” do this will be unveiled in Chapter 9). Let’s consider the original name program, repeated here for convenience:

```
main = do
  hSetBuffering stdin LineBuffering
  putStrLn "Please enter your name: "
  name <- getLine
  putStrLn ("Hello, " ++ name ++ ", how are you?")
```

We can consider the `do` notation as a way to combine a sequence of actions. Moreover, the `<-` notation is a way to get the value out of an action. So, in this program, we’re sequencing four actions: setting buffering, a `putStrLn`, a `getLine` and another `putStrLn`. The `putStrLn` action has type `String → IO ()` so we provide it a `String` so the fully applied action has type `IO ()`. This is something which we are allowed to execute.

The `getLine` action has type `IO String`, so it is okay to execute it directly. However, in order to get the value out of the action, we write `name <- getLine` which basically means “run `getLine` and put the results in the variable called `name`.”

Normal Haskell constructions like **if/then/else** and **case/of** can be used within the **do** notation, but you do need to be somewhat careful. For instance, in our “guess the number” program, we have:

```
do ...
  if (read guess) < num
    then do putStrLn "Too low!"
          doGuessing num
```

```

else if read guess > num
    then do putStrLn "Too high!"
           doGuessing num
    else do putStrLn "You Win!"

```

If we think about how the **if/then/else** construction works, it essentially takes three arguments: the condition, the “then” branch, and the “else” branch. The condition needs to have type `Bool` and the two branches can have any type, provided they have the *same* type. The type of the entire **if/then/else** construction is then the type of the two branches.

In the outermost comparison, we have `(read guess) < num` as the condition. This clearly has the correct type. Let us inspect the “then” branch. The code here is:

```

do putStrLn "Too low!"
  doGuessing num

```

Here, we are sequencing two actions: `putStrLn` and `doGuessing`. The first has type `IO ()` which is fine. The second also has type `IO ()` which is fine. The type result of the entire computation is precisely the type of the final computation. Thus, the type of the “then” branch is also `IO ()`. A similar argument shows that the type of the “else” branch is also `IO ()`. This means the type of the entire **if/then/else** construction is `IO ()` which is just what we want.

■ NOTE ■ In this code, the last line is “else do putStrLn “You Win!””. This is actually somewhat overly verbose. In fact, `else putStrLn “You Win!”` would have been sufficient, since `do` is only necessary to sequence actions. Since we have only one action here, it is superfluous.

It is *incorrect* to think to yourself “Well, I already started a **do** block; I don’t need another one” and hence write something like:

```

do if (read guess) < num
    then putStrLn "Too low!"
       doGuessing num
    else ...

```

Here, since we didn’t repeat the **do**, the compiler doesn’t know that the `putStrLn` and `doGuessing` calls are supposed to be sequenced and will think you’re trying to call `putStrLn` with three arguments; namely the string, the function `doGuessing` and the integer `num`. It will certainly complain (though the error may be somewhat difficult to comprehend at this point).

We can write the same `doGuessing` function using a **case** statement. To do this, we first introduce the Prelude function `compare` which takes two values of the same type (in the `Ord` class) and returns one of `GT`, `LT`, `EQ` depending on whether the first is greater than, less than, or equal to the second.


```
doGuessing num = do
  putStrLn "Enter your guess:"
  guess <- getLine
  case compare (read guess) num of
    LT -> do putStrLn "Too low!"
          doGuessing num
    GT -> do putStrLn "Too high!"
          doGuessing num
    EQ -> putStrLn "You Win!"
```

Here, again, the **dos** after the **->s** are necessary on the first two options because we are sequencing actions.

If you're used to programming in an imperative language like C or Java, you might think that **return** will exit you out of the current function. This is not so in Haskell. In Haskell, **return** simply takes a normal value (for instance, of type `IO Int`) and makes it into an action which returns that value (for instance, of type `Int`). In particular, in an imperative language you might write this function as:

```
void doGuessing(int num) {
  print "Enter your guess:";
  int guess = atoi(readLine());
  if (guess == num) {
    print "You win!";
    return ();
  }

  // we won't get here if guess == num
  if (guess < num) {
    print "Too low!";
    doGuessing(num);
  } else {
    print "Too high!";
    doGuessing(num);
  }
}
```

Here, because we have the `return ()` in the first `if` match, we expect the code to exit there (and in mode imperative languages, it does). However, the equivalent code in Haskell, which might look something like:

```
doGuessing num = do
  putStrLn "Enter your guess:"
  guess <- getLine
  case compare (read guess) num of
```

```

EQ -> do putStrLn "You win!"
      return ()

-- we don't expect to get here unless guess == num
if (read guess < num)
  then do print "Too low!";
          doGuessing
  else do print "Too high!";
          doGuessing

```

This will not behave as you expect. First of all, if you do guess correctly, it will first print “You win!”, but it won’t exit and it will check whether `guess` is less than `num`. Of course it is not, so the `else` branch is taken and it will print “Too high!” and then ask you to guess again.

On the other hand, if you guess incorrectly, it will try to evaluate the case statement and get either `LT` or `GT` as the result of the compare. In either case, it won’t have a pattern which matches and the program will fail immediately.

Exercises

Exercise 21 Write a program that asks the user for his or her name. If their name is one of *Simon*, *John* or *Phil*, tell them you think Haskell is a great programming language. If their name is *Koen*, tell them you think debugging Haskell is fun (*Koen Classen* is one of the people who works on Haskell debugging); otherwise, tell them you don’t know who they are.

Write two different versions of this program, one using *if* statements, one using a *case* statement.

5.4 The IO Library

The IO Library (available by **importing** the IO module) contains many definitions, the most common of which are listed below:

```

data IOMode = ReadMode   | WriteMode
             | AppendMode | ReadWriteMode

openFile    :: FilePath -> IOMode -> IO Handle
hClose      :: Handle   -> IO ()

hIsEOF      :: Handle   -> IO Bool

hGetChar    :: Handle   -> IO Char
hGetLine    :: Handle   -> IO String
hGetContents :: Handle  -> IO String

```

```

getChar      :: IO Char
getLine     :: IO String
getContents  :: IO String

hPutChar    :: Handle -> Char -> IO ()
hPutStr     :: Handle -> String -> IO ()
hPutStrLn   :: Handle -> String -> IO ()

putChar     :: Char -> IO ()
putStr      :: String -> IO ()
putStrLn    :: String -> IO ()

readFile    :: FilePath -> IO String
writeFile   :: FilePath -> String -> IO ()

bracket     ::
  IO a -> (a -> IO b) -> (a -> IO c) -> IO c

```

■ NOTE ■ The type `FilePath` is a *type synonym* for `String`. That is, there is no difference between `FilePath` and `String`. So, for instance, the `readFile` function takes a `String` (the file to read) and returns an action which, when run, produces the contents of that file. See Section 8.1 for more about type synonyms.

Most of these functions are self-explanatory. The `openFile` and `hClose` functions open and close a file, respectively, using the `IOMode` argument as the mode for opening the file. `hIsEOF` tests for end-of file. `hGetChar` and `hGetLine` read a character or line from a file, respectively. `hGetContents` reads in the entire file. The `getChar`, `getLine` and `getContents` variants read from standard input. `hPutChar` prints a character to a file; `hPutStr` prints a string and `hPutStrLn` prints a string with a newline character at the end. The variants without the `h` prefix work on standard output. The `readFile` and `writeFile` functions read an entire file without having to open it first.

The `bracket` function is used to perform actions safely. Consider a function which opens a file, writes a character to it, and then closes the file. When writing such a function, one needs to be careful to ensure that if there were an error at some point, the file is still successfully closed. The `bracket` function makes this easy. It takes three arguments. The first is the action to perform at the beginning. The second is the action to perform at the end, regardless of whether there's an error or not. The third is the action to perform in the middle which might result in an error. For instance, our character-writing function might look like:

```

writeChar :: FilePath -> Char -> IO ()
writeChar fp c =
  bracket

```

```
(openFile fp ReadMode)
hClose
(\h -> hPutChar h c)
```

This will open the file, write the character and then close the file. However, if writing the character fails, `hClose` will still be executed and the exception will be reraised afterwards. That way you don't need to worry too much about catching the exceptions and closing all your handles.

5.5 A File Reading Program

We can write a simple program that allows a user to read and write files. The interface is admittedly poor and it does not catch all errors (try reading a non-existent file). Nevertheless it should give a fairly complete example of how to use IO. Enter the following code into "FileRead.hs" and compile/run.

```
module Main
  where

import IO

main = do
  hSetBuffering stdin LineBuffering
  doLoop

doLoop = do
  putStrLn "Enter a command rFN wFN or q to quit:"
  command <- getLine
  case command of
    'q':_ -> return ()
    'r':filename -> do putStrLn ("Reading " ++ filename)
                      doRead filename
                      doLoop
    'w':filename -> do putStrLn ("Writing " ++ filename)
                      doWrite filename
                      doLoop
    _ -> doLoop

doRead filename =
  bracket (openFile filename ReadMode) hClose
    (\h -> do contents <- hGetContents h
              putStrLn "The first 100 chars:"
              putStrLn (take 100 contents))

doWrite filename = do
```

```

putStrLn "Enter text to go into the file:"
contents <- getLine
bracket (openFile filename WriteMode) hClose
      (\h -> hPutStrLn h contents)

```

What does this program do? First, it issues a short string of instructions and reads a command. It then performs a **case** switch on the command and checks first to see if the first character is a 'q'. If it is, it returns a value of unit type. The `return` function is a function which takes a value of type `a` and returns an action of type `IO a`. Thus, the type of `return ()` is `IO ()`.

If the first character of the command wasn't a 'q', it checks to see if it was an 'r' followed by some string which is bound to the variable `filename`. It then tells you that it's reading the file, does the read, and runs `doLoop` again. The check for 'w' is nearly identical. Otherwise, it matches `_`, the wildcard character, and loops to `doLoop`.

The `doRead` function uses the `bracket` function to make sure there are no problems reading the file. It opens a file in `ReadMode`, gets its contents and prints the first 100 characters (the `take` function takes an integer `n` and a list and returns the first `n` elements of the list).

The `doWrite` function asks for some text, reads it from the keyboard and then writes it to the file specified.

■ **NOTE** ■ Both `doRead` and `doWrite` could have been made simpler by using `readFile` and `writeFile`, but they were written in the extended fashion to show how the more complex functions are used.

The only major problem with this program is that it will die if you try to read a file that already exists, or if you specify some bad filename like `*#. @`. You may think that the calls to `bracket` in `doRead` and `doWrite` should take care of this, but they don't. They only catch exceptions within the main body, not within the startup or shutdown functions (`openFile` and `hClose` in these cases). We would need to catch exceptions raised by `openFile` in order to make this complete. We will do this when we talk about exceptions in more detail in Section 10.1.

Exercises

Exercise 22 Write a program which first asks whether the user wants to read from a file, write to a file or quit. If they respond quit, the program should exit. If they respond read, the program should ask them for a file name and print that file out (if it doesn't exist, it should not crash). If they respond write, it should ask them for a file name and then ask them for text to write to the file, with "." signaling done. All but the "." should be written to the file.

For example, running this program might produce:

```

Do you want to [read] a file, [write] a file or [quit]?
read
Enter a file name to read:

```

```
foo
...contents of foo...
Do you want to [read] a file, [write] a file or [quit]?
write
Enter a file name to write:
foo
Enter text (dot on a line by itself to end):
this is some
text for
foo
.
Do you want to [read] a file, [write] a file or [quit]?
read
Enter a file name to read:
foo
this is some
text for
foo
Do you want to [read] a file, [write] a file or [quit]?
read
Enter a file name to read:
foof
Sorry, that file does not exist.
Do you want to [read] a file, [write] a file or [quit]?
blech
I don't understand the command blech.
Do you want to [read] a file, [write] a file or [quit]?
quit
Goodbye!
```


Chapter 6

Modules

In Haskell, program subcomponents are divided into modules. Each module sits in its own file and the name of the module should match the name of the file (without the “.hs” extension, of course), if you wish to ever use that module in a larger program.

For instance, suppose I am writing a game of poker. I may wish to have a separate module called “Cards” to handle the generation of cards, the shuffling and the dealing functions, and then use this “Cards” module in my “Poker” modules. That way, if I ever go back and want to write a blackjack program, I don’t have to rewrite all the code for the cards; I can simply import the old “Cards” module.

6.1 Exports

Suppose as suggested we are writing a cards module. I have left out the implementation details, but suppose the skeleton of our module looks something like this:

```
module Cards
  where

data Card = ...
data Deck = ...

newDeck :: ... -> Deck
newDeck = ...

shuffle :: ... -> Deck -> Deck
shuffle = ...

-- 'deal deck n' deals 'n' cards from 'deck'
deal :: Deck -> Int -> [Card]
deal deck n = dealHelper deck n []
```



```
dealHelper = ...
```

In this code, the function `deal` calls a helper function `dealHelper`. The implementation of this helper function is very dependent on the exact data structures you used for `Card` and `Deck` so we don't want other people to be able to call this function. In order to do this, we create an *export list*, which we insert just after the module name declaration:

```
module Cards ( Card(),
              Deck(),
              newDeck,
              shuffle,
              deal
            )
  where
  ...
```

Here, we have specified exactly what functions the module exports, so people who use this module won't be able to access our `dealHelper` function. The `()` after `Card` and `Deck` specify that we are exporting the *type* but none of the constructors. For instance if our definition of `Card` were:

```
data Card = Card Suit Face
data Suit = Hearts
          | Spades
          | Diamonds
          | Clubs
data Face = Jack
          | Queen
          | King
          | Ace
          | Number Int
```

Then users of our module would be able to *use* things of type `Card`, but wouldn't be able to construct their own `Cards` and wouldn't be able to extract any of the suit/face information stored in them.

If we wanted users of our module to be able to access all of this information, we would have to specify it in the export list:

```
module Cards ( Card(Card),
              Suit(Hearts, Spades, Diamonds, Clubs),
              Face(Jack, Queen, King, Ace, Number),
              ...
            )
```

```

    where
...

```

This can get frustrating if you're exporting datatypes with many constructors, so if you want to export them all, you can simply write `(..)`, as in:

```

module Cards ( Card(..),
               Suit(..),
               Face(..),
               ...
             )
  where
...

```

And this will automatically export all the constructors.

6.2 Imports

There are a few idiosyncracies in the module import system, but as long as you stay away from the corner cases, you should be fine. Suppose, as before, you wrote a module called “Cards” which you saved in the file “Cards.hs”. You are now writing your poker module and you want to *import* all the definitions from the “Cards” module. To do this, all you need to do is write:

```

module Poker
  where
import Cards

```

This will enable to you use any of the functions, types and constructors exported by the module “Cards”. You may refer to them simply by their name in the “Cards” module (as, for instance, `newDeck`), or you may refer to them explicitly as imported from “Cards” (as, for instance, `Cards.newDeck`). It may be the case that two module export functions or types of the same name. In these cases, you can import one of the modules **qualified** which means that you would no longer be able to simply use the `newDeck` format but must use the longer `Cards.newDeck` format, to remove ambiguity. If you wanted to import “Cards” in this qualified form, you would write:

```

import qualified Cards

```

Another way to avoid problems with overlapping function definitions is to import only certain functions from modules. Suppose we knew the only function from “Cards” that we wanted was `newDeck`, we could import only this function by writing:

```
import Cards (newDeck)
```

On the other hand, suppose we knew that that the `deal` function overlapped with another module, but that we didn't need the "Cards" version of that function. We could hide the definition of `deal` and import everything else by writing:

```
import Cards hiding (deal)
```

Finally, suppose we want to import "Cards" as a qualified module, but don't want to have to type `Cards.` out all the time and would rather just type, for instance, `C.` – we could do this using the `as` keyword:

```
import qualified Cards as C
```

These options can be mixed and matched – you can give explicit import lists on qualified/as imports, for instance.

6.3 Hierarchical Imports

Though technically not part of the Haskell 98 standard, most Haskell compilers support hierarchical imports. This was designed to get rid of clutter in the directories in which modules are stored. Hierarchical imports allow you to specify (to a certain degree) where in the directory structure a module exists. For instance, if you have a "haskell" directory on your computer and this directory is in your compiler's path (see your compiler notes for how to set this; in GHC it's "-i", in Hugs it's "-P"), then you can specify module locations in subdirectories to that directory.

Suppose instead of saving the "Cards" module in your general haskell directory, you created a directory specifically for it called "Cards". The full path of the `Cards.hs` file is then `haskell/Cards/Cards.hs` (or, for Windows `haskell\Cards\Cards.hs`). If you then change the name of the `Cards` module to "`Cards.Cards`", as in:

```
module Cards.Cards(...)  
  where  
  
  ...
```

You could then import it in any module, regardless of this module's directory, as:

```
import Cards.Cards
```

If you start importing these module qualified, I highly recommend using the `as` keyword to shorten the names, so you can write:

```
import qualified Cards.Cards as Cards
... Cards.newDeck ...
```

instead of:

```
import qualified Cards.Cards
... Cards.Cards.newDeck ...
```

which tends to get ugly.

6.4 Literate Versus Non-Literate

The idea of literate programming is a relatively simple one, but took quite a while to become popularized. When we think about programming, we think about the *code* being the default mode of entry and *comments* being secondary. That is, we write code without any special annotation, but comments are annotated with either `--` or `{- ... -}`. Literate programming swaps these preconceptions.

There are two types of literate programs in Haskell; the first uses so-called Bird-scripts and the second uses L^AT_EX-style markup. Each will be discussed individually. No matter which you use, literate scripts must have the extension `lhs` instead of `hs` to tell the compiler that the program is written in a literate style.

6.4.1 Bird-scripts

In a Bird-style literate program, comments are default and code is introduced with a leading greater-than sign (`>`). Everything else remains the same. For example, our Hello World program would be written in Bird-style as:

This is a simple (literate!) Hello World program.

```
> module Main
>     where
```

All our main function does is print a string:

```
> main = putStrLn "Hello World"
```

Note that the spaces between the lines of code and the “comments” are necessary (your compiler will probably complain if you are missing them). When compiled or loaded in an interpreter, this program will have the exact same properties as the non-literate version from Section 3.4.

6.4.2 LaTeX-scripts

LaTeX is a text-markup language very popular in the academic community for publishing. If you are unfamiliar with LaTeX, you may not find this section terribly useful.

Again, a literate Hello World program written in LaTeX-style would look like:

This is another simple (literate!) Hello World program.

```
\begin{code}
module Main
  where
\end{code}
```

All our main function does is print a string:

```
\begin{code}
main = putStrLn "Hello World"
\end{code}
```

In LaTeX-style scripts, the blank lines are *not* necessary.

Chapter 7

Advanced Features

Discussion

7.1 Sections and Infix Operators

We've already seen how to double the values of elements in a list using `map`:

```
Prelude> map (\x -> x*2) [1,2,3,4]
[2,4,6,8]
```

However, there is a more concise way to write this:

```
Prelude> map (*2) [1,2,3,4]
[2,4,6,8]
```

This type of thing can be done for any infix function:

```
Prelude> map (+5) [1,2,3,4]
[6,7,8,9]
Prelude> map (/2) [1,2,3,4]
[0.5,1.0,1.5,2.0]
Prelude> map (2/) [1,2,3,4]
[2.0,1.0,0.666667,0.5]
```

You might be tempted to try to subtract values from elements in a list by mapping `-2` across a list. This won't work, though, because while the `+` in `+2` is parsed as the standard plus operator (as there is no ambiguity), the `-` in `-2` is interpreted as the unary minus, not the binary minus. Thus `-2` here is the *number* `-2`, not the function $\lambda x.x - 2$.

In general, these are called sections. For binary infix operators (like `+`), we can cause the function to become prefix by enclosing it in parentheses. For example:

```
Prelude> (+) 5 3
8
Prelude> (-) 5 3
2
```

Additionally, we can provide either of its argument to make a section. For example:

```
Prelude> (+5) 3
8
Prelude> (/3) 6
2.0
Prelude> (3/) 6
0.5
```

Non-infix functions can be made infix by enclosing them in backquotes (“```”). For example:

```
Prelude> (+2) `map` [1..10]
[3,4,5,6,7,8,9,10,11,12]
```

7.2 Local Declarations

Recall back from Section 3.5, there are many computations which require using the result of the same computation in multiple places in a function. There, we considered the function for computing the roots of a quadratic polynomial:

```
roots a b c =
  ((-b + sqrt(b*b - 4*a*c)) / (2*a),
   (-b - sqrt(b*b - 4*a*c)) / (2*a))
```

In addition to the **let** bindings introduced there, we can do this using a **where** clause. **where** clauses come immediately after function definitions and introduce a new level of layout (see Section 7.11). We write this as:

```
roots a b c =
  ((-b + det) / (2*a), (-b - det) / (2*a))
  where det = sqrt(b*b-4*a*c)
```

Any values defined in a **where** clause *shadow* any other values with the same name. For instance, if we had the following code block:

```

det = "Hello World"

roots a b c =
  ((-b + det) / (2*a), (-b - det) / (2*a))
  where det = sqrt(b*b-4*a*c)

f _ = det

```

The value of `roots` doesn't notice the top-level declaration of `det`, since it is shadowed by the local definition (the fact that the types don't match doesn't matter either). Furthermore, since `f` cannot "see inside" of `roots`, the only thing it knows about `det` is what is available at the top level, which is the string "Hello World." Thus, `f` is a function which takes any argument to that string.

Where clauses can contain any number of subexpressions, but they must be aligned for layout. For instance, we could also pull out the `2*a` computation and get the following code:

```

roots a b c =
  ((-b + det) / (a2), (-b - det) / (a2))
  where det = sqrt(b*b-4*a*c)
        a2 = 2*a

```

Sub-expressions in **where** clauses must come after function definitions. Sometimes it is more convenient to put the local definitions before the actual expression of the function. This can be done by using **let/in** clauses. We have already seen **let** clauses; **where** clauses are virtually identical to their **let** clause cousins except for their placement. The same `roots` function can be written using **let** as:

```

roots a b c =
  let det = sqrt (b*b - 4*a*c)
      a2 = 2*a
  in ((-b + det) / a2, (-b - det) / a2)

```

Using a **where** clause, it looks like:

```

roots a b c = ((-b + det) / a2, (-b - det) / a2)
  where
    det = sqrt (b*b - 4*a*c)
    a2 = 2*a

```

These two types of clauses can be mixed (i.e., you can write a function which has both a **let** clause and a **where** clause). This is strongly advised *against*, as it tends to make code difficult to read. However, if you choose to do it, values in the **let** clause shadow those in the **where** clause. So if you define the function:


```
f x =
  let y = x+1
  in y
  where y = x+2
```

The value of `f 5` is 6, not 7. Of course, I plead with you to never ever write code that looks like this. No one should have to remember this rule and by shadowing **where**-defined values in a **let** clause only makes your code difficult to understand.

In general, whether you should use **let** clauses or **where** clauses is largely a matter of personal preference. Usually, the names you give to the subexpressions should be sufficiently expressive that without reading their definitions any reader of your code should be able to figure out what they do. In this case, **where** clauses are probably more desirable because they allow the reader to see immediately what a function does. However, in real life, values are often given cryptic names. In which case **let** clauses may be better. Either is probably okay, though I think **where** clauses are more common.

7.3 Partial Application

Partial application is when you take a function which takes n arguments and you supply it with $< n$ of them. When discussing sections in Section 7.1, we saw a form of “partial application” in which functions like `+` were partially applied. For instance, in the expression `map (+1) [1, 2, 3]`, the section `(+1)` is a partial application of `+`. This is because `+` really takes two arguments, but we’ve only given it one.

Partial application is very common in function definitions and sometimes goes by the name “eta reduction”. For instance, suppose we are writing a function `toLowerCaseString` which converts a whole string into lower case. We could write this as:

```
toLowerCaseString s = map toLower s
```

Here, there is no partial application (though you could argue that applying no arguments to `toLowerCase` could be considered partial application). However, we notice that the application of `s` occurs at the end of both `toLowerCaseString` and of `map toLower`. In fact, we can remove it by performing eta reduction, to get:

```
toLowerCaseString = map toLower
```

Now, we have a partial application of `map`: it expects a function and a list, but we’ve only given it the function.

This all is related to type type of `map`, which is $(a \rightarrow b) \rightarrow ([a] \rightarrow [b])$, when parentheses are all included. In our case, `toLowerCase` is of type `Char → Char`. Thus, if we supply this function to `map`, we get a function of type `[Char] → [Char]`, as desired.

Now, consider the task of converting a string to lowercase and remove all non letter characters. We might write this as:

```
lcaseLetters s = map toLower (filter isAlpha s)
```

But note that we can actually write this in terms of function composition:

```
lcaseLetters s = (map toLower . filter isAlpha) s
```

And again, we're left with an eta reducible function:

```
lcaseLetters = map toLower . filter isAlpha
```

Writing functions in this style is very common among advanced Haskell users. In fact it has a name: point-free programming (not to be confused with *pointless* programming). It is called point free because in the original definition of `lcaseLetters`, we can think of the value `s` as a point on which the function is operating. By removing the point from the function definition, we have a point-free function.

point-free programming

A function similar to `(.)` is `($)`. Whereas `(.)` is function composition, `($)` is function application. The definition of `($)` from the Prelude is very simple:

`$`
function application

```
f $ x = f x
```

However, this function is given very low fixity, which means that it can be used to replace parentheses. For instance, we might write a function:

```
foo x y = bar y (baz (fluff (ork x)))
```

However, using the function application function, we can rewrite this as:

```
foo x y = bar y $ baz $ fluff $ ork x
```

This moderately resembles the function composition syntax. The `($)` function is also useful when combined with other infix functions. For instance, we cannot write:

```
Prelude> putStrLn "5+3=" ++ show (5+3)
```

because this is interpreted as `(putStrLn "5+3=") ++ (show (5+3))`, which makes no sense. However, we can fix this by writing instead:

```
Prelude> putStrLn $ "5+3=" ++ show (5+3)
```

Which works fine.

Consider now the task of extracting from a list of tuples all the ones whose first component is greater than zero. One way to write this would be:

```
fstGt0 l = filter (\ (a,b) -> a>0) l
```

We can first apply eta reduction to the whole function, yielding:

```
fstGt0 = filter (\ (a,b) -> a>0)
```

Now, we can rewrite the lambda function to use the `fst` function instead of the pattern matching:

```
fstGt0 = filter (\x -> fst x > 0)
```

Now, we can use function composition between `fst` and `>` to get:

```
fstGt0 = filter (\x -> ((>0) . fst) x)
```

And finally we can eta reduce:

```
fstGt0 = filter ((>0).fst)
```

This definition is simultaneously shorter and easier to understand than the original. We can clearly see exactly what it is doing: we're filtering a list by checking whether something is greater than zero. What are we checking? The `fst` element.

While converting to point free style often results in clearer code, this is of course not always the case. For instance, converting the following map to point free style yields something nearly uninterpretable:

```
foo = map (\x -> sqrt (3+4*(x^2)))
foo = map (sqrt . (3+) . (4*) . (^2))
```

There are a handful of combinators defined in the Prelude which are useful for point free programming:

- `uncurry` takes a function of type $a \rightarrow b \rightarrow c$ and converts it into a function of type $(a, b) \rightarrow c$. This is useful, for example, when mapping across a list of pairs:

```
Prelude> map (uncurry (*)) [(1,2),(3,4),(5,6)]
[2,12,30]
```

- `curry` is the opposite of `uncurry` and takes a function of type $(a, b) \rightarrow c$ and produces a function of type $a \rightarrow b \rightarrow c$.
- `flip` reverse the order of arguments to a function. That is, it takes a function of type $a \rightarrow b \rightarrow c$ and produces a function of type $b \rightarrow a \rightarrow c$. For instance, we can sort a list in reverse order by using `flip compare`:

```
Prelude> List.sortBy compare [5,1,8,3]
[1,3,5,8]
Prelude> List.sortBy (flip compare) [5,1,8,3]
[8,5,3,1]
```

This is the same as saying:

```
Prelude> List.sortBy (\a b -> compare b a) [5,1,8,3]
[8,5,3,1]
```

only shorter.

Of course, not all functions can be written in point free style. For instance:

```
square x = x*x
```

Cannot be written in point free style, without some other combinators. For instance, if we can define other functions, we can write:

```
pair x = (x,x)
square = uncurry (*) . pair
```

But in this case, this is not terribly useful.

Exercises

Exercise 23 Convert the following functions into point-free style, if possible.

```
func1 x l = map (\y -> y*x) l
func2 f g l = filter f (map g l)
func3 f l = l ++ map f l
func4 l = map (\y -> y+2)
           (filter (\z -> z `elem` [1..10])
            (5:1))
func5 f l = foldr (\x y -> f (y,x)) 0 l
```

You might have been tempted to try to write `func2` as `filter f . map g`, trying to eta-reduce off the `g`. In this case, this isn't possible. This is because the function composition operator `(.)` has type $(b \rightarrow c) \rightarrow (a \rightarrow b) \rightarrow (a \rightarrow c)$. In this case, we're trying to use `map` as the second argument. But `map` takes two arguments, while `(.)` expects a function which takes only one.

7.4 Pattern Matching

Pattern matching is one of the most powerful features of Haskell (and most functional programming languages). It is most commonly used in conjunction with **case** expressions, which we have already seen in Section 3.5. Let's return to our `Color` example from Section 4.5. I'll repeat the definition we already had for the datatype:

```
data Color
  = Red
  | Orange
  | Yellow
  | Green
  | Blue
  | Purple
  | White
  | Black
  | Custom Int Int Int -- R G B components
  deriving (Show,Eq)
```

We then want to write a function that will convert between something of type `Color` and a triple of `Int`s, which correspond to the RGB values, respectively. Specifically, if we see a `Color` which is `Red`, we want to return `(255,0,0)`, since this is the RGB value for red. So we write that (remember that piecewise function definitions are just **case** statements):

```
colorToRGB Red = (255,0,0)
```

If we see a `Color` which is `Orange`, we want to return `(255,128,0)`; and if we see `Yellow`, we want to return `(255,255,0)`, and so on. Finally, if we see a custom color, which is comprised of three components, we want to make a triple out of these, so we write:

```
colorToRGB Orange = (255,128,0)
colorToRGB Yellow = (255,255,0)
colorToRGB Green  = (0,255,0)
colorToRGB Blue   = (0,0,255)
colorToRGB Purple = (255,0,255)
colorToRGB White  = (255,255,255)
colorToRGB Black  = (0,0,0)
colorToRGB (Custom r g b) = (r,g,b)
```

Then, in our interpreter, if we type:

```
Color> colorToRGB Yellow
(255,255,0)
```

What is happening is this: we create a value, call it x , which has value `Red`. We then apply this to `colorToRGB`. We check to see if we can “match” x against `Red`. This match fails because according to the definition of `Eq Color`, `Red` is not equal to `Yellow`. We continue down the definitions of `colorToRGB` and try to match `Yellow` against `Orange`. This fails, too. We then try to match `Yellow` against `Yellow`, which succeeds, so we use this function definition, which simply returns the value `(255, 255, 0)`, as expected.

Suppose instead, we used a custom color:

```
Color> colorToRGB (Custom 50 200 100)
(50,200,100)
```

We apply the same matching process, failing on all values from `Red` to `Black`. We then get to try to match `Custom 50 200 100` against `Custom r g b`. We can see that the `Custom` part matches, so then we go see if the subelements match. In the matching, the variables `r`, `g` and `b` are essentially wild cards, so there is no trouble matching `r` with 50, `g` with 200 and `b` with 100. As a “side-effect” of this matching, `r` gets the value 50, `g` gets the value 200 and `b` gets the value 100. So the entire match succeeded and we look at the definition of this part of the function and bundle up the triple using the matched values of `r`, `g` and `b`.

We can also write a function to check to see if a `Color` is a custom color or not:

```
isCustomColor (Custom _ _ _) = True
isCustomColor _ = False
```

When we apply a value to `isCustomColor` it tries to match that value against `Custom _ _ _`. This match will succeed if the value is `Custom x y z` for any `x`, `y` and `z`. The `_` (underscore) character is a “wildcard” and will match anything, but will not do the binding that would happen if you put a variable name there. If this match succeeds, the function returns `True`; however, if this match fails, it goes on to the next line, which will match anything and then return `False`.

For some reason we might want to define a function which tells us whether a given color is “bright” or not, where my definition of “bright” is that one of its RGB components is equal to 255 (admittedly and arbitrary definition, but it’s simply an example). We could define this function as:

```
isBright = isBright' . colorToRGB
  where isBright' (255,_,_) = True
        isBright' (_,255,_) = True
        isBright' (_,_,255) = True
        isBright' _         = False
```

Let’s dwell on this definition for a second. The `isBright` function is the composition of our previously defined function `colorToRGB` and a helper function `isBright'`, which tells us if a given RGB value is bright or not. We could replace the

first line here with `isBright c = isBright' (colorToRGB c)` but there is no need to explicitly write the parameter here, so we don't. Again, this function composition style of programming takes some getting used to, so I will try to use it frequently in this tutorial.

The `isBright'` helper function takes the RGB triple produced by `colorToRGB`. It first tries to match it against `(255,_,_)` which succeeds if the value has 255 in its first position. If this match succeeds, `isBright'` returns `True` and so does `isBright`. The second and third line of definition check for 255 in the second and third position in the triple, respectively. The fourth line, the *fallthrough*, matches everything else and reports it as not bright.

We might want to also write a function to convert between RGB triples and `Colors`. We could simply stick everything in a `Custom` constructor, but this would defeat the purpose; we want to use the `Custom` slot only for values which don't match the predefined colors. However, we don't want to allow the user to construct custom colors like `(600,-40,99)` since these are invalid RGB values. We could throw an error if such a value is given, but this can be difficult to deal with. Instead, we use the `Maybe` datatype. This is defined (in the Prelude) as:

```
data Maybe a = Nothing
             | Just a
```

The way we use this is as follows: our `rgbToColor` function returns a value of type `Maybe Color`. If the RGB value passed to our function is *invalid*, we return `Nothing`, which corresponds to a failure. If, on the other hand, the RGB value is valid, we create the appropriate `Color` value and return `Just` that. The code to do this is:

```
rgbToColor 255  0  0 = Just Red
rgbToColor 255 128 0 = Just Orange
rgbToColor 255 255 0 = Just Yellow
rgbToColor  0 255  0 = Just Green
rgbToColor  0  0 255 = Just Blue
rgbToColor 255  0 255 = Just Purple
rgbToColor 255 255 255 = Just White
rgbToColor  0  0  0 = Just Black
rgbToColor  r  g  b =
  if 0 <= r && r <= 255 &&
    0 <= g && g <= 255 &&
    0 <= b && b <= 255
  then Just (Custom r g b)
  else Nothing  -- invalid RGB value
```

The first eight lines match the RGB arguments against the predefined values and, if they match, `rgbToColor` returns `Just` the appropriate color. If none of these matches, the last definition of `rgbToColor` matches the first argument against `r`, the

second against `g` and the third against `b` (which causes the side-effect of binding these values). It then checks to see if these values are valid (each is greater than or equal to zero and less than or equal to 255). If so, it returns `Just (Custom r g b)`; if not, it returns `Nothing` corresponding to an invalid color.

Using this, we can write a function that checks to see if a right RGB value is valid:

```
rgbIsValid = rgbIsValid' . rgbToColor
  where rgbIsValid' (Just _) = True
        rgbIsValid' _      = False
```

Here, we compose the helper function `rgbIsValid'` with our function `rgbToColor`. The helper function checks to see if the value returned by `rgbToColor` is `Just` anything (the wildcard). If so, it returns `True`. If not, it matches anything and returns `False`.

Pattern matching isn't magic, though. You can only match against datatypes; you cannot match against functions. For instance, the following is invalid:

```
f x = x + 1
g (f x) = x
```

Even though the intended meaning of `g` is clear (i.e., $g\ x = x - 1$), the compiler doesn't know in general that `f` has an inverse function, so it can't perform matches like this.

7.5 Guards

Guards can be thought of as an extension to the pattern matching facility. They enable you to allow piecewise function definitions to be taken according to arbitrary boolean expressions. Guards appear after all arguments to a function but before the equals sign, and are begun with a vertical bar `|`. We could use guards to write a simple function which returns a string telling you the result of comparing two elements:

```
comparison x y | x < y = "The first is less"
                | x > y = "The second is less"
                | otherwise = "They are equal"
```

You can read the vertical bar as “such that.” So we say that the value of `comparison x y` “such that” `x` is less than `y` is “The first is less.” The value such that `x` is greater than `y` is “The second is less” and the value **otherwise** is “They are equal”. The keyword `otherwise` is simply defined to be equal to `True` and thus matches anything that falls through that far. So, we can see that this works:


```

Guards> comparison 5 10
"The first is less"
Guards> comparison 10 5
"The second is less"
Guards> comparison 7 7
"They are equal"

```

One thing to note about guards is that they are tested *after* pattern matching, not in conjunction with pattern matching. This means that once a pattern matches, if none of the guards succeed, further pattern matches will not be attempted. So, if we had instead defined:

```

comparison2 x y | x < y = "The first is less"
                | x > y = "The second is less"
comparison2 _ _ = "They are equal"

```

The intention would be that if both of the guards failed, it would “fall through” to the final match and say that they were equal. This is not what happens, though.

```

Guards> comparison2 7 7
*** Exception: Guards.hs:8: Non-exhaustive patterns
    in function comparison2

```

If we think about what is happening in the compiler this makes sense. When we apply two sevens to `comparison2`, they are matched against `x` and `y`, which succeeds and the values are bound. Pattern matching then stops completely, since it has succeeded. The guards are then activated and `x` and `y` are compared. Neither of the guards succeeds, so an error is raised.

One nicety about guards is that **where** clauses are common to all guards. So another possible definition for our `isBright` function from the previous section would be:

```

isBright2 c | r == 255 = True
            | g == 255 = True
            | b == 255 = True
            | otherwise = False
  where (r,g,b) = colorToRGB c

```

The function is equivalent to the previous version, but performs its calculation slightly differently. It takes a color, `c`, and applies `colorToRGB` to it, yielding an RGB triple which is matched (using pattern matching!) against `(r,g,b)`. This match succeeds and the values `r`, `g` and `b` are bound to their respective values. The first guard checks to see if `r` is 255 and, if so, returns true. The second and third guard check `g` and `b` against 255, respectively and return true if they match. The last guard fires as a last resort and returns `False`.

7.6 Instance Declarations

In order to declare a type to be an instance of a class, you need to provide an instance declaration for it. Most classes provide what's called a "minimal complete definition." This means the functions which must be implemented for this class in order for its definition to be satisfied. Once you've written these functions for your type, you can declare it an instance of the class.

7.6.1 The `Eq` Class

The `Eq` class has two members (i.e., two functions):

```
(==) :: Eq a => a -> a -> Bool
(/=) :: Eq a => a -> a -> Bool
```

The first of these type signatures reads that the function `==` is a function which takes two `a`s which are members of `Eq` and produces a `Bool`. The type signature of `/=` (not equal) is identical. A minimal complete definition for the `Eq` class requires that either one of these functions be defined (if you define `==`, then `/=` is defined automatically by negating the result of `==`, and vice versa). These declarations must be provided inside the instance declaration.

This is best demonstrated by example. Suppose we have our color example, re-peted here for convenience:

```
data Color
  = Red
  | Orange
  | Yellow
  | Green
  | Blue
  | Purple
  | White
  | Black
  | Custom Int Int Int -- R G B components
```

We can define `Color` to be an instance of `Eq` by the following declaration:

```
instance Eq Color where
  Red == Red = True
  Orange == Orange = True
  Yellow == Yellow = True
  Green == Green = True
  Blue == Blue = True
  Purple == Purple = True
  White == White = True
```

```

Black == Black = True
(Custom r g b) == (Custom r' g' b') =
    r == r' && g == g' && b == b'
_ == _ = False

```

The first line here begins with the keyword **instance** telling the compiler that we're making an instance declaration. It then specifies the class, **Eq**, and the type, **Color** which is going to be an instance of this class. Following that, there's the **where** keyword. Finally there's the method declaration.

The first eight lines of the method declaration are basically identical. The first one, for instance, says that the value of the expression `Red == Red` is equal to `True`. Lines two through eight are identical. The declaration for custom colors is a bit different. We pattern match `Custom` on both sides of `==`. On the left hand side, we bind `r`, `g` and `b` to the components, respectively. On the right hand side, we bind `r'`, `g'` and `b'` to the components. We then say that these two custom colors are equal precisely when `r == r'`, `g == g'` and `b == b'` are all equal. The fallthrough says that any pair we haven't previously declared as equal are unequal.

7.6.2 The Show Class

The **Show** class is used to display arbitrary values as strings. This class has three methods:

```

show :: Show a => a -> String
showsPrec :: Show a => Int -> a -> String -> String
showList :: Show a => [a] -> String -> String

```

A minimal complete definition is either `show` or `showsPrec` (we will talk about `showsPrec` later – it's in there for efficiency reasons). We can define our **Color** datatype to be an instance of **Show** with the following instance declaration:

```

instance Show Color where
    show Red = "Red"
    show Orange = "Orange"
    show Yellow = "Yellow"
    show Green = "Green"
    show Blue = "Blue"
    show Purple = "Purple"
    show White = "White"
    show Black = "Black"
    show (Custom r g b) =
        "Custom " ++ show r ++ " " ++
        show g ++ " " ++ show b

```

This declaration specifies exactly how to convert values of type **Color** to **Strings**. Again, the first eight lines are identical and simply take a **Color** and produce a string.

The last line for handling custom colors matches out the RGB components and creates a string by concatenating the result of showing the components individually (with spaces in between and “Custom” at the beginning).

7.6.3 Other Important Classes

There are a few other important classes which I will mention briefly because either they are commonly used or because we will be using them shortly. I won’t provide example instance declarations; how you can do this should be clear by now.

The Ord Class

The ordering class, the functions are:

```
compare :: Ord a => a -> a -> Ordering
(<=) :: Ord a => a -> a -> Bool
(>) :: Ord a => a -> a -> Bool
(>=) :: Ord a => a -> a -> Bool
(<) :: Ord a => a -> a -> Bool
min :: Ord a => a -> a -> a
max :: Ord a => a -> a -> a
```

The almost any of the functions alone is a minimal complete definition; it is recommended that you implement `compare` if you implement only one, though. This function returns a value of type `Ordering` which is defined as:

```
data Ordering = LT | EQ | GT
```

So, for instance, we get:

```
Prelude> compare 5 7
LT
Prelude> compare 6 6
EQ
Prelude> compare 7 5
GT
```

In order to declare a type to be an instance of `Ord` you must already have declared it an instance of `Eq` (in other words, `Ord` is a *subclass* of `Eq` – more about this in Section 8.4).

The Enum Class

The `Enum` class is for enumerated types; that is, for types where each element has a successor and a predecessor. It’s methods are:

```

pred :: Enum a => a -> a
succ :: Enum a => a -> a
toEnum :: Enum a => Int -> a
fromEnum :: Enum a => a -> Int
enumFrom :: Enum a => a -> [a]
enumFromThen :: Enum a => a -> a -> [a]
enumFromTo :: Enum a => a -> a -> [a]
enumFromThenTo :: Enum a => a -> a -> a -> [a]

```

The minimal complete definition contains both `toEnum` and `fromEnum`, which converts from and to ints. The `pred` and `succ` functions give the predecessor and successor, respectively. The `enum` functions enumerate lists of elements. For instance, `enumFrom x` lists all elements after `x`; `enumFromThen x step` lists all elements starting at `x` in steps of size `step`. The `To` functions end the enumeration at the given element.

The Num Class

The **Num** class provides the standard arithmetic operations:

```

(-) :: Num a => a -> a -> a
(*) :: Num a => a -> a -> a
(+) :: Num a => a -> a -> a
negate :: Num a => a -> a
signum :: Num a => a -> a
abs :: Num a => a -> a
fromInteger :: Num a => Integer -> a

```

All of these are obvious except for perhaps `negate` which is the unary minus. That is, `negate x` means $-x$.

The Read Class

The **Read** class is the opposite of the **Show** class. It is a way to take a string and read in from it a value of arbitrary type. The methods for **Read** are:

```

readsPrec :: Read a => Int -> String -> [(a, String)]
readList :: String -> [(a, String)]

```

The minimal complete definition is `readsPrec`. The most important function related to this is `read`, which uses `readsPrec` as:

```

read s = fst (head (readsPrec 0 s))

```

This will fail if parsing the string fails. You could define a `maybeRead` function as:

```
maybeRead s =
  case readsPrec 0 s of
    [(a, _)] -> Just a
    _ -> Nothing
```

How to write and use `readsPrec` directly will be discussed further in the examples.

7.6.4 Class Contexts

Suppose we are defining the `Maybe` datatype from scratch. The definition would be something like:

```
data Maybe a = Nothing
             | Just a
```

Now, when we go to write the instance declarations, for, say, `Eq`, we need to know that `a` is an instance of `Eq` otherwise we can't write a declaration. We express this as:

```
instance Eq a => Eq (Maybe a) where
  Nothing == Nothing = True
  (Just x) == (Just x') == x == x'
```

This first line can be read “That `a` is an instance of `Eq` *implies* (`=>`) that `Maybe a` is an instance of `Eq`.”

7.6.5 Deriving Classes

Writing obvious `Eq`, `Ord`, `Read` and `Show` classes like these is tedious and should be automated. Luckily for us, it is. If you write a datatype that's “simple enough” (almost any datatype you'll write unless you start writing fixed point types), the compiler can automatically *derive* some of the most basic classes. To do this, you simply add a **deriving** clause to after the datatype declaration, as in:

```
data Color
  = Red
  | ...
  | Custom Int Int Int -- R G B components
  deriving (Eq, Ord, Show, Read)
```

This will automatically create instances of the `Color` datatype of the named classes. Similarly, the declaration:

```
data Maybe a = Nothing
             | Just a
             deriving (Eq, Ord, Show, Read)
```

derives these classes just when `a` is appropriate.

All in all, you are allowed to derive instances of **Eq**, **Ord**, **Enum**, **Bounded**, **Show** and **Read**. There is considerable work in the area of “polytypic programming” or “generic programming” which, among other things, would allow for instance declarations for *any* class to be derived. This is much beyond the scope of this tutorial; instead, I refer you to the literature.

7.7 Datatypes Revisited

I know by this point you’re probably terribly tired of hearing about datatypes. They are, however, incredibly important, otherwise I wouldn’t devote so much time to them. Datatypes offer a sort of notational convenience if you have, for instance, a datatype that holds many many values. These are called named fields.

7.7.1 Named Fields

Consider a datatype whose purpose is to hold configuration settings. Usually when you extract members from this type, you really only care about one or possibly two of the many settings. Moreover, if many of the settings have the same type, you might often find yourself wondering “wait, was this the fourth or *fifth* element?” One thing you could do would be to write accessor functions. Consider the following made-up configuration type for a terminal program:

```
data Configuration =
    Configuration String      -- user name
                  String      -- local host
                  String      -- remote host
                  Bool         -- is guest?
                  Bool         -- is super user?
                  String       -- current directory
                  String       -- home directory
                  Integer      -- time connected
    deriving (Eq, Show)
```

You could then write accessor functions, like (I’ve only listed a few):

```
getUserName (Configuration un _ _ _ _ _ _) = un
getLocalHost (Configuration _ lh _ _ _ _ _) = lh
getRemoteHost (Configuration _ _ rh _ _ _ _) = rh
```

```
getIsGuest (Configuration _ _ _ ig _ _ _ _) = ig
...
```

You could also write update functions to update a single element. Of course, now if you add an element to the configuration, or remove one, all of these functions now have to take a different number of arguments. This is highly annoying and is an easy place for bugs to slip in. However, there's a solution. We simply give names to the fields in the datatype declaration, as follows:

```
data Configuration =
  Configuration { username      :: String,
                 localhost     :: String,
                 remotehost    :: String,
                 isguest       :: Bool,
                 issuperuser   :: Bool,
                 currentdir    :: String,
                 homedir       :: String,
                 timeconnected :: Integer
                 }
```

This will automatically generate the following accessor functions for us:

```
username :: Configuration -> String
localhost :: Configuration -> String
...
```

Moreover, it gives us very convenient update methods. Here is a short example for a “post working directory” and “change directory” like functions that work on Configurations:

```
changeDir :: Configuration -> String -> Configuration
changeDir cfg newDir =
  -- make sure the directory exists
  if directoryExists newDir
  then -- change our current directory
       cfg{currentdir = newDir}
  else error "directory does not exist"

postWorkingDir :: Configuration -> String
-- retrieve our current directory
postWorkingDir cfg = currentdir cfg
```

So, in general, to update the field x in a datatype y to z , you write $y\{x=z\}$. You can change more than one; each should be separated by commas, for instance, $y\{x=z, a=b, c=d\}$.

You can of course continue to pattern match against `Configuration`s as you did before. The named fields are simply syntactic sugar; you can still write something like:

```
getUsername (Configuration un _ _ _ _ _ _) = un
```

But there is little reason to. Finally, you can pattern match against named fields as in:

```
getHostData (Configuration {localhost=lh,remotehost=rh})
  = (lh,rh)
```

This matches the variable `lh` against the `localhost` field on the `Configuration` and the variable `rh` against the `remotehost` field on the `Configuration`. These matches of course succeed. You could also constring the matches by putting values instead of variable names in these positions, as you would for standard datatypes.

You can create values of `Configuration` in the old way as shown in the first definition below, or in the named-field style, as shown in the second definition below:

```
initCFG =
  Configuration "nobody" "nowhere" "nowhere"
               False False "/" "/" 0
initCFG' =
  Configuration
    { username="nobody",
      localhost="nowhere",
      remotehost="nowhere",
      isguest=False,
      issuperuser=False,
      currentdir="/",
      homedir="/",
      timeconnected=0 }
```

Though the second is probably much more understandable unless you litter your code with comments.

7.8 More Lists

todo: put something here

7.8.1 Standard List Functions

Recall that the definition of the built-in Haskell list datatype is equivalent to:

```
data List a = Nil
            | Cons a (List a)
```

With the exception that `Nil` is called `[]` and `Cons x xs` is called `x:xs`. This is simply to make pattern matching easier and code smaller. Let's investigate how some of the standard list functions may be written. Consider `map`. A definition is given below:

```
map _ [] = []
map f (x:xs) = f x : map f xs
```

Here, the first line says that when you `map` across an empty list, no matter what the function is, you get an empty list back. The second line says that when you `map` across a list with `x` as the head and `xs` as the tail, the result is `f` applied to `x` consed onto the result of mapping `f` on `xs`.

The `filter` can be defined similarly:

```
filter _ [] = []
filter p (x:xs) | p x = x : filter p xs
                | otherwise = filter p xs
```

How this works should be clear. For an empty list, we return an empty list. For a non empty list, we return the filter of the tail, perhaps with the head on the front, depending on whether it satisfies the predicate `p` or not.

We can define `foldr` as:

```
foldr _ z [] = z
foldr f z (x:xs) = f x (foldr f z xs)
```

Here, the best interpretation is that we are replacing the empty list (`[]`) with a particular value and the list constructor (`:`) with some function. On the first line, we can see the replacement of `[]` for `z`. Using backquotes to make `f` infix, we can write the second line as:

```
foldr f z (x:xs) = x `f` (foldr f z xs)
```

From this, we can directly see how `:` is being replaced by `f`.

Finally, `foldl`:

```
foldl _ z [] = z
foldl f z (x:xs) = foldl f (f z x) xs
```

This is slightly more complicated. Remember, `z` can be thought of as the current state. So if we're folding across a list which is empty, we simply return the current state. On the other hand, if the list is not empty, it's of the form `x:xs`. In this case, we get a new state by applying `f` to the current state `z` and the current list element `x` and then recursively call `foldl` on `xs` with this new state.

There is another class of functions: the `zip` and `unzip` functions, which respectively take multiple lists and make one or take one lists and split them apart. For instance, `zip` does the following:

```
Prelude> zip "hello" [1,2,3,4,5]
[( 'h',1),('e',2),('l',3),('l',4),('o',5)]
```

Basically, it pairs the first elements of both lists and makes that the first element of the new list. It then pairs the second elements of both lists and makes that the second element, etc. What if the lists have unequal length? It simply stops when the shorter one stops. A reasonable definition for `zip` is:

```
zip [] _ = []
zip _ [] = []
zip (x:xs) (y:ys) = (x,y) : zip xs ys
```

The `unzip` function does the opposite. It takes a zipped list and returns the two “original” lists:

```
Prelude> unzip [( 'f',1),('o',2),('o',3)]
("foo",[1,2,3])
```

There are a whole slew of `zip` and `unzip` functions, named `zip3`, `unzip3`, `zip4`, `unzip4` and so on; the `...3` functions use triples instead of pairs; the `...4` functions use 4-tuples, etc.

Finally, the function `take` takes an integer n and a list and returns the first n elements off the list. Correspondingly, `drop` takes an integer n and a list and returns the result of throwing away the first n elements off the list. Neither of these functions produces an error; if n is too large, they both will just return shorter lists.

7.8.2 List Comprehensions

There is some syntactic sugar for dealing with lists whose elements are members of the `Enum` class (see Section 7.6), such as `Int` or `Char`. If we want to create a list of all the elements from 1 to 10, we can simply write:

```
Prelude> [1..10]
[1,2,3,4,5,6,7,8,9,10]
```

We can also introduce an amount to step by:

```
Prelude> [1,3..10]
[1,3,5,7,9]
Prelude> [1,4..10]
[1,4,7,10]
```

These expressions are short hand for `enumFromTo` and `enumFromThenTo`, respectively. Of course, you don’t need to specify an upper bound. Try the following (but be ready to hit `Control+C` to stop the computation!):

```
Prelude> [1..]
[1,2,3,4,5,6,7,8,9,10,11,12{Interrupted!}]
```

Probably yours printed a few thousand more elements than this. As we said before, Haskell is lazy. That means that a list of all numbers from 1 on is perfectly well formed and that's exactly what this list is. Of course, if you attempt to print the list (which we're implicitly doing by typing it in the interpreter), it won't halt. But if we only evaluate an initial segment of this list, we're fine:

```
Prelude> take 3 [1..]
[1,2,3]
Prelude> take 3 (drop 5 [1..])
[6,7,8]
```

This comes in useful if, say, we want to assign an ID to each element in a list. Without laziness we'd have to write something like this:

```
assignID :: [a] -> [(a,Int)]
assignID l = zip l [1..length l]
```

Which means that the list will be traversed twice. However, because of laziness, we can simply write:

```
assignID l = zip l [1..]
```

And we'll get exactly what we want. We can see that this works:

```
Prelude> assignID "hello"
[( 'h' ,1), ( 'e' ,2), ( 'l' ,3), ( 'l' ,4), ( 'o' ,5)]
```

Finally, there is some useful syntactic sugar for `map` and `filter`, based on standard set-notation in mathematics. In math, we would write something like $\{f(x) \mid x \in s \wedge p(x)\}$ to mean the set of all values of f when applied to elements of s which satisfy p . This is equivalent to the Haskell statement `map f (filter p s)`. However, we can also use more math-like notation and write `[f x | x <- s, p x]`. While in math the ordering of the statements on the side after the pipe is free, it is not so in Haskell. We could not have put `p x` before `x <- s` otherwise the compiler wouldn't know yet what x was. We can use this to do simple string processing. Suppose we want to take a string, remove all the lower-case letters and convert the rest of the letters to upper case. We could do this in either of the following two equivalent ways:

```
Prelude> map toLower (filter isUpper "Hello World")
"hw"
Prelude> [toLower x | x <- "Hello World", isUpper x]
"hw"
```

These two are equivalent, and, depending on the exact functions you're using, one might be more readable than the other. There's more you can do here, though. Suppose you want to create a list of pairs, one for each point between (0,0) and (5,7) below the diagonal. Doing this manually with lists and maps would be cumbersome and possibly difficult to read. It couldn't be easier with list comprehensions:

```
Prelude> [(x,y) | x <- [1..5], y <- [x..7]]
[(1,1), (1,2), (1,3), (1,4), (1,5), (1,6), (1,7), (2,2), (2,3),
 (2,4), (2,5), (2,6), (2,7), (3,3), (3,4), (3,5), (3,6), (3,7),
 (4,4), (4,5), (4,6), (4,7), (5,5), (5,6), (5,7)]
```

If you reverse the order of the `x <-` and `y <-` clauses, the order in which the space is traversed will be reversed (of course, in that case, `y` could no longer depend on `x` and you would need to make `x` depend on `y` but this is trivial).

7.9 Arrays

Lists are nice for many things. It is easy to add elements to the beginning of them and to manipulate them in various ways that change the length of the list. However, they are bad for random access, having average complexity $\mathcal{O}(n)$ to access an arbitrary element (if you don't know what $\mathcal{O}(\dots)$ means, you can either ignore it or take a quick detour and read Appendix A, a two-page introduction to complexity theory). So, if you're willing to give up fast insertion and deletion because you need random access, you should use arrays instead of lists.

In order to use arrays you must import the `Array` module. There are a few methods for creating arrays, the `array` function, the `listArray` function, and the `accumArray` function. The `array` function takes a pair which is the bounds of the array, and an association list which specifies the initial values of the array. The `listArray` function takes bounds and then simply a list of values. Finally, the `accumArray` function takes an accumulation function, an initial value and an association list and accumulates pairs from the list into the array. Here are some examples of arrays being created:

```
Arrays> array (1,5) [(i,2*i) | i <- [1..5]]
array (1,5) [(1,2), (2,4), (3,6), (4,8), (5,10)]
Arrays> listArray (1,5) [3,7,5,1,10]
array (1,5) [(1,3), (2,7), (3,5), (4,1), (5,10)]
Arrays> accumArray (+) 2 (1,5) [(i,i) | i <- [1..5]]
array (1,5) [(1,3), (2,4), (3,5), (4,6), (5,7)]
```

When arrays are printed out (via the `show` function), they are printed with an association list. For instance, in the first example, the association list says that the value of the array at 1 is 2, the value of the array at 2 is 4, and so on.

You can extract an element of an array using the `!` function, which takes an array and an index, as in:

```
Arrays> (listArray (1,5) [3,7,5,1,10]) ! 3
5
```

Moreover, you can update elements in the array using the `//` function. This takes an array and an association list and updates the positions specified in the list:

```
Arrays> (listArray (1,5) [3,7,5,1,10]) //
      [(2,99),(3,-99)]
array (1,5) [(1,3),(2,99),(3,-99),(4,1),(5,10)]
```

There are a few other functions which are of interest:

- `bounds` returns the bounds of an array
- `indices` returns a list of all indices of the array
- `elems` returns a list of all the values in the array in order
- `assocs` returns an association list for the array

If we define `arr` to be `listArray (1,5) [3,7,5,1,10]`, the result of these functions applied to `arr` are:

```
Arrays> bounds arr
(1,5)
Arrays> indices arr
[1,2,3,4,5]
Arrays> elems arr
[3,7,5,1,10]
Arrays> assocs arr
[(1,3),(2,7),(3,5),(4,1),(5,10)]
```

Note that while arrays are $\mathcal{O}(1)$ access, they are not $\mathcal{O}(1)$ update. They are in fact $\mathcal{O}(n)$ update, since in order to maintain purity, the array must be *copied* in order to make an update. Thus, functional arrays are pretty much only useful when you're filling them up once and then only reading. If you need fast access and update, you should probably use `FiniteMaps`, which are discussed in Section 7.10 and have $\mathcal{O}(\log n)$ access and update.

7.10 Finite Maps

The `FiniteMap` datatype (which is available in the `FiniteMap` module, or `Data.FiniteMap` module in the hierarchical libraries) is a purely functional implementation of balanced trees. Finite maps can be compared to lists or arrays in terms of the time it takes to perform various operations on those datatypes of a fixed size, n . A brief comparison is:

	List	Array	FiniteMap
insert	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(\log n)$
update	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(\log n)$
delete	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(\log n)$
find	$\mathcal{O}(n)$	$\mathcal{O}(1)$	$\mathcal{O}(\log n)$
map	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n \log n)$

As we can see, lists provide fast insertion (but slow everything else), arrays provide fast lookup (but slow everything else) and finite maps provide moderately fast everything (except mapping, which is a bit slower than lists or arrays).

The type of a finite map is for the form `FiniteMapkeyelt` where `key` is the type of the keys and `elt` is the type of the elements. That is, finite maps are lookup tables from type `key` to type `elt`.

The basic finite map functions are:

```
emptyFM    :: FiniteMap key elt
addToFM    :: FiniteMap key elt -> key -> elt ->
             FiniteMap key elt
delFromFM  :: FiniteMap key elt -> key ->
             FiniteMap key elt
elemFM     :: key -> FiniteMap key elt -> Bool
lookupFM   :: FiniteMap key elt -> key -> Maybe elt
```

In all these cases, the type `key` must be an instance of `Ord` (and hence also an instance of `Eq`).

There are also functions `listToFM` and `fmToList` to convert lists to and from finite maps. Try the following:

```
Prelude> :m FiniteMap
FiniteMap> let fm = listToFM
             [( 'a',5),('b',10),('c',1),('d',2) ]
FiniteMap> let myFM = addToFM fm 'e' 6
FiniteMap> fmToList fm
[( 'a',5),('b',10),('c',1),('d',2) ]
FiniteMap> fmToList myFM
[( 'a',5),('b',10),('c',1),('d',2),('e',6) ]
FiniteMap> lookupFM myFM 'e'
Just 6
FiniteMap> lookupFM fm 'e'
Nothing
```

You can also experiment with the other commands. Note that you cannot show a finite map, as they are not instances of `Show`:

```
FiniteMap> show myFM
```

```
<interactive>:1:
  No instance for (Show (FiniteMap Char Integer))
  arising from use of `show' at <interactive>:1
  In the definition of `it': show myFM
```

In order to inspect the elements, you first need to use `fmToList`.

7.11 Layout

7.12 The Final Word on Lists

You are likely tired of hearing about lists at this point, but they are so fundamental to Haskell (and really all of functional programming) that it would be terrible not to talk about them some more.

It turns out that `foldr` is actually quite a powerful function: it can compute an *primitive recursive* function. A primitive recursive function is essentially one which can be calculated using only “for” loops, but not “while” loops. primitive recursive

In fact, we can fairly easily define `map` in terms of `foldr`:

```
map2 f = foldr (\a b -> f a : b) []
```

Here, `b` is the accumulator (i.e., the result list) and `a` is the element being currently considered. In fact, we can simplify this definition through a sequence of steps:

```
foldr (\a b -> f a : b) []
==> foldr (\a b -> (:) (f a) b) []
==> foldr (\a -> (:) (f a)) []
==> foldr (\a -> ((:) . f) a) []
==> foldr ((:) . f) []
```

This is directly related to the fact that `foldr (:) []` is the identity function on lists. This is because, as mentioned before, `foldr f z` can be thought of as replacing the `[]` in lists by `z` and the `:` by `f`. In this case, we’re keeping both the same, so it is the identity function.

In fact, you can convert any function of the following style into a `foldr`:

```
myfunc [] = z
myfunc (x:xs) = f x (myfunc xs)
```

By writing the last line with `f` in infix form, this should be obvious:

```
myfunc [] = z
myfunc (x:xs) = x `f` (myfunc xs)
```


Clearly, we are just replacing `[]` with `z` and `:` with `f`. Consider the `filter` function:

```
filter p [] = []
filter p (x:xs) =
  if p x
  then x : filter p xs
  else filter p xs
```

This function also follows the form above. Based on the first line, we can figure out that `z` is supposed to be `[]`, just like in the `map` case. Now, suppose that we call the result of calling `filter p xs` simply `b`, then we can rewrite this as:

```
filter p [] = []
filter p (x:xs) =
  if p x then x : b else b
```

Given this, we can transform `filter` into a fold:

```
filter p = foldr (\a b -> if p a then a:b else b) []
```

Let's consider a slightly more complicated function: `++`. The definition for `++` is:

```
(++) [] ys = ys
(++) (x:xs) ys = x : (xs ++ ys)
```

Now, the question is whether we can write this in fold notation. First, we can apply eta reduction to the first line to give:

```
(++) [] = id
```

Through a sequence of steps, we can also eta-reduce the second line:

```
(++) (x:xs) ys = x : ((++) xs ys)
==> (++) (x:xs) ys = (x:) ((++) xs ys)
==> (++) (x:xs) ys = ((x:) . (++) xs) ys
==> (++) (x:xs) = (x:) . (++) xs
```

Thus, we get that an eta-reduced definition of `++` is:

```
(++) [] = id
(++) (x:xs) = (x:) . (++) xs
```

Now, we can try to put this into fold notation. First, we notice that the base case converts `[]` into `id`. Now, if we assume `(++) xs` is called `b` and `x` is called `a`, we can get the following definition in terms of `foldr`:

```
(++) = foldr (\a b -> (a:) . b) id
```

This actually makes sense intuitively. If we only think about applying `++` to one argument, we can think of it as a function which takes a list and creates a function which, when applied, will prepend this list to another list. In the lambda function, we assume we have a function `b` which will do this for the rest of the list and we need to create a function which will do this for `b` as well as the single element `a`. In order to do this, we first apply `b` and then further add `a` to the front.

We can further reduce this expression to a point-free style through the following sequence:

```
==> (++) = foldr (\a b -> (a:) . b) id
==> (++) = foldr (\a b -> (.) (a:) b) id
==> (++) = foldr (\a -> (.) (a:)) id
==> (++) = foldr (\a -> (.) ((:) a)) id
==> (++) = foldr (\a -> ((.) . (:)) a) id
==> (++) = foldr ((.) . (:)) id
```

This final version is point free, though not necessarily understandable. Presumably the original version is clearer.

As a final example, consider `concat`. We can write this as:

```
concat [] = []
concat (x:xs) = x ++ concat xs
```

It should be immediately clear that the `z` element for the fold is `[]` and that the recursive function is `++`, yielding:

```
concat = foldr (++) []
```

Exercises

Exercise 24 The function `and` takes a list of booleans and returns `True` if and only if all of them are `True`. It also returns `True` on the empty list. Write this function in terms of `foldr`.

Exercise 25 The function `concatMap` behaves such that `concatMap f` is the same as `concat . map f`. Write this function in terms of `foldr`.

Chapter 8

Advanced Types

As you've probably ascertained by this point, the type system is integral to Haskell. While this chapter is called "Advanced Types", you will probably find it to be more general than that and it must not be skipped simply because you're not interested in the type system.

8.1 Type Synonyms

Type synonyms exist in Haskell simply for convenience: their removal would not make Haskell any less powerful.

Consider the case when you are constantly dealing with lists of three-dimensional points. For instance, you might have a function with type `[(Double, Double, Double)] → Double → [(Double, Double, Double)]`. Since you are a good software engineer, you want to place type signatures on all your top-level functions. However, typing `[(Double, Double, Double)]` all the time gets very tedious. To get around this, you can define a type synonym:

```
type List3D = [(Double, Double, Double)]
```

Now, the type signature for your functions may be written `List3D → Double → List3D`.

We should note that type synonyms cannot be self-referential. That is, you cannot have:

```
type BadType = Int -> BadType
```

This is because this is an "infinite type." Since Haskell removes type synonyms very early on, any instance of `BadType` will be replaced by `Int → BadType`, which will result in an infinite loop.

Type synonyms can also be parameterized. For instance, you might want to be able to change the types of the points in the list of 3D points. For this, you could define:

```
type List3D a = [(a, a, a)]
```

Then your references to `[(Double, Double, Double)]` would become `List3D Double`.

8.2 Newtypes

Consider the problem in which you need to have a type which is very much like `Int`, but its ordering is defined differently. Perhaps you wish to order `Ints` first by even numbers then by odd numbers (that is, all odd numbers are greater than any even number and within the odd/even subsets, ordering is standard).

Unfortunately, you cannot define a new instance of `Ord` for `Int` because then Haskell won't know which one to use. What you want is to define a type which is *isomorphic* to `Int`.

■ NOTE ■ “Isomorphic” is a common term in mathematics which basically means “structurally identical.” For instance, in graph theory, if you have two graphs which are identical except they have different labels on the nodes, they are isomorphic. In our context, two types are isomorphic if they have the same underlying structure.

One way to do this would be to define a new datatype:

```
data MyInt = MyInt Int
```

We could then write appropriate code for this datatype. The problem (and this is very subtle) is that this type is not truly isomorphic to `Int`: it has one more value. When we think of the type `Int`, we usually think that it takes all values of integers, but it really has one more value: `⊥` (pronounced “bottom”), which is used to represent erroneous or undefined computations. Thus, `MyInt` has not only values `MyInt 0`, `MyInt 1` and so on, but also `MyInt ⊥`. However, since datatypes can themselves be undefined, it has an additional value: `⊥` which differs from `MyInt ⊥` and this makes the types non-isomorphic. (See Section 11.1 for more information on bottom.)

Disregarding that subtlety, there may be efficiency issues with this representation: now, instead of simply storing an integer, we have to store a pointer to an integer and have to follow that pointer whenever we need the value of a `MyInt`.

To get around these problems, Haskell has a **newtype** construction. A **newtype** is a cross between a datatype and a type synonym: it has a constructor like a datatype, but it can have only one constructor and this constructor can have only one argument. For instance, we can define:

```
newtype MyInt = MyInt Int
```

But we cannot define any of:

```
newtype Bad1 = Bad1a Int | Bad1b Double
newtype Bad2 = Bad2 Int Double
```

Of course, the fact that we cannot define `Bad2` as above is not a big issue: we can simply define the following by pairing the types:

```
newtype Good2 = Good2 (Int, Double)
```

Now, suppose we've defined `MyInt` as a **newtype**. This enables use to write our desired instance of `Ord` as:

```
instance Ord MyInt where
  MyInt i < MyInt j
    | odd i && odd j = i < j
    | even i && even j = i < j
    | even i          = True
    | otherwise       = False
  where odd x = (x `mod` 2) == 0
        even = not . odd
```

Like datatype, we can still derive classes like `Show` and `Eq` over newtypes (in fact, I'm implicitly assuming we have derived `Eq` over `MyInt` – where is my assumption in the above code?).

Moreover, in recent versions of GHC (see Section 2.2), on newtypes, you are allowed to derive *any* class of which the base type (in this case, `Int`) is an instance. For example, we could derive `Num` on `MyInt` to provide arithmetic functions over it.

Pattern matching over newtypes is exactly as in datatypes. We can write constructor and destructor functions for `MyInt` as follows:

```
mkMyInt i = MyInt i
unMyInt (MyInt i) = i
```

8.3 Datatypes

We've already seen datatypes used in a variety of contexts. This section concludes some of the discussion and introduces some of the common datatypes in Haskell. It also provides a more theoretical underpinning to what datatypes actually are.

8.3.1 Strict Fields

One of the great things about Haskell is that computation is performed lazily. However, sometimes this leads to inefficiencies. One way around this problem is to use datatypes with strict fields. Before we talk about the solution, let's spend some time to get a bit more comfortable with how bottom works in to the picture (for more theory, see Section 11.1).

Suppose we've defined the unit datatype (this one of the simplest datatypes you can define):

```
data Unit = Unit
```

This datatype has exactly one constructor, `Unit`, which takes no arguments. In a strict language like ML, there would be exactly one value of type `Unit`: namely, `Unit`. This is not quite so in Haskell. In fact, there are *two* values of type `Unit`. One of them is `Unit`. The other is bottom (written `_|_`).

You can think of bottom as representing a computation which won't halt. For instance, suppose we define the value:

```
foo = foo
```

This is perfectly valid Haskell code and simply says that when you want to evaluate `foo`, all you need to do is evaluate `foo`. Clearly this is an “infinite loop.”

What is the type of `foo`? Simply `a`. We cannot say anything more about it than that. The fact that `foo` has type `a` in fact tells us that it must be an infinite loop (or some other such strange value). However, since `foo` has type `a` and thus can have any type, it can also have type `Unit`. We could write, for instance:

```
foo :: Unit
foo = foo
```

Thus, we have found a second value with type `Unit`. In fact, we have found all values of type `Unit`. Any other non-terminating function or error-producing function will have exactly the same effect as `foo` (though Haskell provides some more utility with the function `error`).

This means, for instance, that there are actually *four* values with type `Maybe Unit`. They are: `_|_`, `Nothing`, `Just _|_` and `Just Unit`. However, it could be the fact that you, as a programmer, know that you will never come across the third of these. Namely, you want the argument to `Just` to be *strict*. This means that if the argument to `Just` is bottom, then the entire structure becomes bottom. You use an exclamation point to specify a constructor as strict. We can define a strict version of `Maybe` as:

```
data SMaybe a = SNothing | SJust !a
```

There are now only three values of `SMaybe`. We can see the difference by writing the following program:

```
module Main where

import System

data SMaybe a = SNothing | SJust !a deriving Show
```

```

main = do
  [cmd] <- getArgs
  case cmd of
    "a" -> printJust  undefined
    "b" -> printJust  Nothing
    "c" -> printJust  (Just undefined)
    "d" -> printJust  (Just ())

    "e" -> printSJust undefined
    "f" -> printSJust SNothing
    "g" -> printSJust (SJust undefined)
    "h" -> printSJust (SJust ())

printJust :: Maybe () -> IO ()
printJust Nothing = putStrLn "Nothing"
printJust (Just x) = do putStr "Just "; print x

printJust :: SMaybe () -> IO ()
printSJust SNothing = putStrLn "Nothing"
printSJust (SJust x) = do putStr "Just "; print x

```

Here, depending on what command line argument is passed, we will do something different. The outputs for the various options are:

```

\% ./strict a
Fail: Prelude.undefined

\% ./strict b
Nothing

\% ./strict c
Just
Fail: Prelude.undefined

\% ./strict d
Just ()

\% ./strict e
Fail: Prelude.undefined

\% ./strict f
Nothing

\% ./strict g
Fail: Prelude.undefined

```



```
\% ./strict h
Just ()
```

The thing worth noting here is the difference between cases “c” and “g”. In the “c” case, the `Just` is printed, because this is printed *before* the undefined value is evaluated. However, in the “g” case, since the constructor is strict, as soon as you match the `SJust`, you also match the value. In this case, the value is undefined, so the whole thing fails before it gets a chance to do *anything*.

8.4 Classes

We have already encountered type classes a few times, but only in the context of previously existing type classes. This section is about how to define your own. We will begin the discussion by talking about Pong and then move on to a useful generalization of computations.

8.4.1 Pong

The discussion here will be motivated by the construction of the game Pong (see Appendix ?? for the full code). In Pong, there are three things drawn on the screen: the two paddles and the ball. While the paddles and the ball are different in a few respects, they share many commonalities, such as position, velocity, acceleration, color, shape, and so on. We can express these commonalities by defining a class for Pong entities, which we call **Entity**. We make such a definition as follows:

```
class Entity a where
  getPosition :: a -> (Int,Int)
  getVelocity :: a -> (Int,Int)
  getAcceleration :: a -> (Int,Int)
  getColor :: a -> Color
  getShape :: a -> Shape
```

This code defines a typeclass **Entity**. This class has five methods: `getPosition`, `getVelocity`, `getAcceleration`, `getColor` and `getShape` with the corresponding types.

The first line here uses the keyword **class** to introduce a new typeclass. We can read this typeclass definition as “There is a typeclass ‘Entity’; a type ‘a’ is an instance of Entity if it provides the following five functions: . . .”. To see how we can write an instance of this class, let us define a player (paddle) datatype:

```
data Paddle =
  Paddle { paddlePosX, paddlePosY,
          paddleVelX, paddleVelY,
          paddleAccX, paddleAccY :: Int,
          paddleColor :: Color,
```

```
paddleHeight :: Int,
playerNumber :: Int }
```

Given this data declaration, we can define `Paddle` to be an instance of `Entity`:

```
instance Entity Paddle where
  getPosition p = (paddlePosX p, paddlePosY p)
  getVelocity p = (paddleVelX p, paddleVelY p)
  getAcceleration p = (paddleAccX p, paddleAccY p)
  getColor = paddleColor
  getShape = Rectangle 5 . paddleHeight
```

The actual Haskell types of the class functions all have included the context `Entity a =>`. For example, `getPosition` has type `Entity a => a -> (Int, Int)`. However, it will turn out that many of our routines will need entities to also be instances of `Eq`. We can therefore choose to make `Entity` a subclass of `Eq`: namely, you can only be an instance of `Entity` if you are already an instance of `Eq`. To do this, we change the first line of the class declaration to:

```
class Eq a => Entity a where
```

Now, in order to define `Paddles` to be instances of `Entity` we will first need them to be instances of `Eq` – we can do this by deriving the class.

8.4.2 Computations

Let's think back to our original motivation for defining the `Maybe` datatype from Section ???. We wanted to be able to express that functions (i.e., computations) can fail.

Let us consider the case of performing search on a graph. Allow us to take a small aside to set up a small graph library:

```
data Graph v e = Graph [(Int,v)] [(Int,Int,e)]
```

The `Graph` datatype takes two type arguments which correspond to vertex and edge labels. The first argument to the `Graph` constructor is a list (set) of vertices; the second is the list (set) of edges. We will assume these lists are always sorted and that each vertex has a unique id and that there is at most one edge between any two vertices.

Suppose we want to search for a path between two vertices. Perhaps there is no path between those vertices. To represent this, we will use the `Maybe` datatype. If it succeeds, it will return the list of vertices traversed. Our search function could be written (naively) as follows:

```
search :: Graph v e -> Int -> Int -> Maybe [Int]
search g@(Graph vl el) src dst
  | src == dst = Just [src]
```

```

| otherwise = search' el
where search' [] = Nothing
      search' ((u,v,_) : es)
          | src == u =
            case search g v dst of
              Just p  -> Just (u:p)
              Nothing -> search' es
          | otherwise = search' es

```

This algorithm works as follows (try to read along): to search in a graph g from src to dst , first we check to see if these are equal. If they are, we have found our way and just return the trivial solution. Otherwise, we want to traverse the edge-list. If we're traversing the edge-list and it is empty, we've failed, so we return `Nothing`. Otherwise, we're looking at an edge from u to v . If u is our source, then we consider this step and recursively search the graph from v to dst . If this fails, we try the rest of the edges; if this succeeds, we put our current position before the path found and return. If u is not our source, this edge is useless and we continue traversing the edge-list.

This algorithm is terrible: namely, if the graph contains cycles, it can loop indefinitely. Nevertheless, it is sufficient for now. Be sure you understand it well: things only get more complicated.

Now, there are cases where the `Maybe` datatype is not sufficient: perhaps we wish to include an error message together with the failure. We could define a datatype to express this as:

```
data Failable a = Success a | Fail String
```

Now, failures come with a failure string to express what went wrong. We can rewrite our search function to use this datatype:

```

search2 :: Graph v e -> Int -> Int -> Failable [Int]
search2 g@(Graph vl el) src dst
  | src == dst = Success [src]
  | otherwise = search' el
where search' [] = Fail "No path"
      search' ((u,v,_) : es)
          | src == u =
            case search2 g v dst of
              Success p -> Success (u:p)
              _         -> search' es
          | otherwise = search' es

```

This code is a straightforward translation of the above.

There is another option for this computation: perhaps we want not just one path, but all possible paths. We can express this as a function which returns a list of lists of vertices. The basic idea is the same:

```

search3 :: Graph v e -> Int -> Int -> [[Int]]
search3 g@(Graph vl el) src dst
  | src == dst = [[src]]
  | otherwise = search' el
  where search' [] = []
        search' ((u,v,_) : es)
          | src == u =
              map (u:) (search3 g v dst) ++
              search' es
          | otherwise = search' es

```

The code here has gotten a little shorter, thanks to the standard prelude `map` function, though it is essentially the same.

We may ask ourselves what all of these have in common and try to gobble up those commonalities in a class. In essence, we need some way of representing success and some way of representing failure. Furthermore, we need a way to combine two successes (in the first two cases, the first success is chosen; in the third, they are strung together). Finally, we need to be able to augment a previous success (if there was one) with some new value. We can fit this all into a class as follows:

```

class Computation c where
  success :: a -> c a
  failure :: String -> c a
  augment :: c a -> (a -> c b) -> c b
  combine :: c a -> c a -> c a

```

In this class declaration, we're saying that `c` is an instance of the class **Computation** if it provides four functions: `success`, `failure`, `augment` and `combine`. The `success` function takes a value of type `a` and returns it wrapped up in `c`, representing a successful computation. The `failure` function takes a `String` and returns a computation representing a failure. The `combine` function takes two previous computation and produces a new one which is the combination of both. The `augment` function is a bit more complex.

The `augment` function takes some previously given computation (namely, `c a`) and a function which takes the value of that computation (the `a`) and returns a `b` and produces a `b` inside of that computation. Note that in our current situation, giving `augment` the type `c a -> (a -> a) -> c a` would have been sufficient, since `a` is always `[Int]`, but we make it this more general time just for generality.

How `augment` works is probably best shown by example. We can define `Maybe`, `Failable` and `[]` to be instances of **Computation** as:

```

instance Computation Maybe where
  success = Just
  failure = const Nothing

```

```
augment (Just x) f = f x
augment Nothing _ = Nothing
combine Nothing y = y
combine x _ = x
```

Here, success is represented with `Just` and failure ignores its argument and returns `Nothing`. The `combine` function takes the first success we found and ignores the rest. The function `augment` checks to see if we succeeded before (and thus had a `Just` something) and, if we did, applies `f` to it. If we failed before (and thus had a `Nothing`), we ignore the function and return `Nothing`.

```
instance Computation Failable where
  success = Success
  failure = Fail
  augment (Success x) f = f x
  augment (Fail s) _ = Fail s
  combine (Fail _) y = y
  combine x _ = x
```

These definitions are obvious. Finally:

```
instance Computation [] where
  success a = [a]
  failure = const []
  augment l f = concat (map f l)
  combine = (++)
```

Here, the value of a successful computation is a singleton list containing that value. Failure is represented with the empty list and to combine previous successes we simply concatenate them. Finally, augmenting a computation amounts to mapping the function across the list of previous computations and concatenate them. we apply the function to each element in the list and then concatenate the results.

Using these computations, we can express all of the above versions of search as:

```
searchAll g@(Graph vl el) src dst
  | src == dst = success [src]
  | otherwise = search' el
  where search' [] = failure "no path"
        search' ((u,v,_) : es)
          | src == u = (searchAll g v dst `augment`
                       (success . (u:)))
                    `combine` search' es
          | otherwise = search' es
```

In this, we see the uses of all the functions from the class `Computation`.

If you've understood this discussion of computations, you are in a very good position as you have understood the concept of *monads*, probably the most difficult concept in Haskell. In fact, the `Computation` class is almost exactly the `Monad` class, except that `success` is called `return`, `failure` is called `fail` and `augment` is called `>>=` (read "bind"). The `combine` function isn't actually required by monads, but is found in the `MonadPlus` class for reasons which will become obvious later.

If you didn't understand everything here, read through it again and then wait for the proper discussion of monads in Chapter 9.

8.5 Instances

We have already seen how to declare instances of some simple classes; allow us to consider some more advanced classes here. There is a `Functor` class defined in the `Functor` module.

■ NOTE ■ The name "functor", like "monad" comes from category theory. There, a functor is like a function, but instead of mapping elements to elements, it maps structures to structures.

The definition of the functor class is:

```
class Functor f where
  fmap :: (a -> b) -> f a -> f b
```

The type definition for `fmap` (not to mention its name) is very similar to the function `map` over lists. In fact, `fmap` is essentially a generalization of `map` to arbitrary structures (and, of course, lists are already instances of `Functor`). However, we can also define other structures to be instances of functors. Consider the following datatype for binary trees:

```
data BinTree a = Leaf a
               | Branch (BinTree a) (BinTree a)
```

We can immediately identify that the `BinTree` type essentially "raises" a type `a` into trees of that type. There is a naturally associated functor which goes along with this raising. We can write the instance:

```
instance Functor BinTree where
  fmap f (Leaf a) = Leaf (f a)
  fmap f (Branch left right) =
    Branch (fmap f left) (fmap f right)
```

Now, we've seen how to make something like `BinTree` an instance of `Eq` by using the **deriving** keyword, but here we will do it by hand. We want to make `BinTree` as instances of `Eq` but obviously we cannot do this unless `a` is itself an instance of `Eq`. We can specify this dependence in the instance declaration:

```
instance Eq a => Eq (BinTree a) where
  Leaf a == Leaf b = a == b
  Branch l r == Branch l' r' = l == l' && r == r'
  _ == _ = False
```

The first line of this can be read “if `a` is an instance of `Eq`, then `BinTree a` is also an instance of `Eq`”. We then provide the definitions. If we did not include the “`Eq a =>`” part, the compiler would complain because we're trying to use the `==` function on `a` as in the second line.

The “`Eq a =>`” part of the definition is called the “context.” We should note that there are some restrictions on what can appear in the context and what can appear in the declaration. For instance, we're not allowed to have instance declarations that don't contain type constructors on the right hand side. To see why, consider the following declarations:

```
class MyEq a where
  myeq :: a -> a -> Bool

instance Eq a => MyEq a where
  myeq = (==)
```

As it stands, there doesn't seem to be anything wrong with this definition. However, if elsewhere in a program we had the definition:

```
instance MyEq a => Eq a where
  (==) = myeq
```

In this case, if we're trying to establish if some type is an instance of `Eq`, we could reduce it to trying to find out if that type is an instance of `MyEq`, which we could in turn reduce to trying to find out if that type is an instance of `Eq`, and so on. The compiler protects itself against this by refusing the first instance declaration.

This is commonly known as the *closed-world assumption*. That is, we're assuming, when we write a definition like the first one, that there won't be any declarations like the second. However, this assumption is invalid because there's nothing to prevent the second declaration (or some equally evil declaration). The closed world assumption can also bite you in cases like:

```
class OnlyInts a where
  foo :: a -> a -> Bool
```

```
instance OnlyInts Int where
    foo == (==)

bar :: OnlyInts a => a -> Bool
bar = foo 5
```

We've again made the closed-world assumption: we've assumed that the only instance of `OnlyInts` is `Int`, but there's no reason another instance couldn't be defined elsewhere, ruining our definition of `bar`.

8.6 Kinds

Let us take a moment and think about what types are available in Haskell. We have simple types, like `Int`, `Char`, `Double` and so on. We then have type constructors like `Maybe` which take a type (like `Char`) and produce a new type, `Maybe Char`. Similarly, the type constructor `[]` (lists) takes a type (like `Int`) and produces `[Int]`. We have more complex things like `→` (function arrow) which takes *two* types (say `Int` and `Bool`) and produces a new type `Int → Bool`.

In a sense, these types themselves have type. Types like `Int` have some sort of basic type. Types like `Maybe` have a type which takes something of basic type and returns something of basic type. And so forth.

Talking about the types of types becomes unwieldy and highly ambiguous, so we call the types of types “kinds.” What we have been calling “basic types” have kind “*”. Something of kind * is something which can have an actual value. There is also a single kind constructor, `→` with which we can build more complex kinds.

Consider `Maybe`. This takes something of kind * and produces something of kind *. Thus, the kind of `Maybe` is `* -> *`. Recall the definition of `Pair` from Section 4.5.1:

```
data Pair a b = Pair a b
```

Here, `Pair` is a type constructor which takes two arguments, each of kind * and produces a type of kind *. Thus, the kind of `Pair` is `* -> (* -> *)`. However, we again assume associativity so we just write `* -> * -> *`.

Let us make a slightly strange datatype definition:

```
data Strange c a b =
    MkStrange (c a) (c b)
```

Before we analyze the kind of `Strange`, let's think about what it does. It is essentially a pairing constructor, though it doesn't pair actual elements, but elements within another constructor. For instance, think of `c` as `Maybe`. Then `MkStrange` pairs `Maybes` of the two types `a` and `b`. However, `c` need not be `Maybe` but could instead be `[]`, or many other things.

What do we know about c , though? We know that it must have kind $* \rightarrow *$. This is because we have $c \ a$ on the right hand side. The type variables a and b each have kind $*$ as before. Thus, the kind of `Strange` is $(* \rightarrow *) \rightarrow * \rightarrow * \rightarrow *$. That is, it takes a constructor (c) of kind $* \rightarrow *$ together with two types of kind $*$ and produces something of kind $*$.

A question may arise regarding how we know a has kind $*$ and not some other kind k . In fact, the inferred kind for `Strange` is $(k \rightarrow *) \rightarrow k \rightarrow k \rightarrow *$. However, this requires polymorphism on the kind level, which is too complex, so we make a default assumption that $k = *$.

■ NOTE ■ There are extensions to GHC which allow you to specify the kind of constructors directly. For instance, if you wanted a different kind, you could write this explicitly:

```
data Strange (c :: (* -> *) -> *) a b = MkStrange (c a) (c b)
```

to give a different kind to `Strange`.

The notation of kinds suggests that we can perform partial application, as we can for functions. And, in fact, we can. For instance, we could have:

```
type MaybePair = Strange Maybe
```

The kind of `MaybePair` is, not surprisingly, $* \rightarrow * \rightarrow *$.

We should note here that all of the following definitions are acceptable:

```
type MaybePair1      = Strange Maybe
type MaybePair2 a    = Strange Maybe a
type MaybePair3 a b  = Strange Maybe a b
```

These all appear to be the same, but they are in fact not identical as far as Haskell's type system is concerned. The following are all valid type definitions using the above:

```
type MaybePair1a = MaybePair1
type MaybePair1b = MaybePair1 Int
type MaybePair1c = MaybePair1 Int Double

type MaybePair2b = MaybePair2 Int
type MaybePair2c = MaybePair2 Int Double

type MaybePair3c = MaybePair3 Int Double
```

But the following are *not* valid:

```
type MaybePair2a = MaybePair2
type MaybePair3a = MaybePair3
type MaybePair3b = MaybePair3 Int
```

This is because while it is possible to partially apply type constructors on datatypes, it is not possible on type synonyms. For instance, the reason `MaybePair2a` is invalid is because `MaybePair2` is defined as a type synonym with one argument and we have given it none. The same applies for the invalid `MaybePair3` definitions.

8.7 Class Hierarchies

8.8 Default

what is it?

Chapter 9

Monads

Learning about monads is the hardest thing you will have to do when learning Haskell. I will distinguish two different subcomponents of learning about monads: (1) learning how to use already-existing monads and (2) learning how to write your own. If you want to use Haskell, you must be able to use already-existing monads. Only if you want to be a super Haskell guru must you learn how to write your own. However, if you do get the hang of writing your own monads, your life will become a lot less painful.

So far we've seen two uses of monads. The first was IO actions. We've seen that by using monads, we can abstract away from the problems plaguing both the RealWorld and CPS solutions presented in Chapter 5. The second was using monads to represent different types of computations in Section 8.4.2. In both cases, we needed a way to sequence operations and saw that a sufficient definition (at least for computations) was:

```
class Computation c where
  success :: a -> c a
  failure :: String -> c a
  augment :: c a -> (a -> c b) -> c b
  combine :: c a -> c a -> c a
```

Let us see if this will enable us to also do IO. Essentially, we need a way to represent taking a value out of an action and performing some new operation on it (as in the example from Section 4.4.3, rephrased slightly):

```
main = do
  s <- readFile "somefile"
  putStrLn (show (f s))
```

But this is exactly what `augment` does. In fact, using `augment`, we can write the above code as:

```
main = -- note the lack of a "do"
      readFile "somefile" `augment` \s ->
      putStrLn (show (f s))
```

This certainly seems to be sufficient. And, in fact, it turns out to be more than sufficient.

The definition of a monad is a slightly trimmed down version of our `Computation` class. It contains four methods, but one has a default definition in term of another:

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

In this definition, `return` is equivalent to our `success`; `fail` is equivalent to our `failure`; and `>>=` (read: “bind”) is equivalent to our `augment`. The `>>` (read: “then”) method is simply a version of `>>=` which ignores the `a`. This will turn out to be useful, although it can be defined in terms of `>>=`: we define `a >> x ≡ a >>= _ -> x`.

9.1 Do Notation

We have hinted that there is a connection between monads and the **do** notation. Here, we make that relationship concrete. As it turns out, there is nothing special about the syntactic sugar **do** notation. It is simply syntactic sugar for monadic operations.

As we mentioned earlier, using our `Computation` class, we could define our above program as:

```
main =
  readFile "somefile" `augment` \s ->
  putStrLn (show (f s))
```

But we now know that `augment` is called `>>=` in the monadic world. Thus, this program really reads (and this is completely valid Haskell at this point: if you defined a function `f :: Show a => String -> a`, you could compile and run this program):

```
main =
  readFile "somefile" >>= \s ->
  putStrLn (show (f s))
```

This suggests that we can translate:

```
do x <- f
  g x
```

into `f >>= \x -> g x`. This is exactly what the compiler does. Talking about **do** becomes easier if we abstract away from layout (see Section ?? for how to do this). There are four translation rules:

1. `do {e} → e`
2. `do {e; es} → e >> do {es}`
3. `do {let decls; es} → let decls in do {es}`
4. `do {p <- e; es} → let ok p = do {es} ; ok _ = fail "..."`
`in e >>= ok`

Again, we will take these one at a time.

Translation Rule 1

The first translation rule, `do {e} → e`, states (as we have stated before) that when performing a single action, having a **do** or not is irrelevant. This is essentially the base case for an inductive definition of **do**. The base case has one action (namely `e` here); the other three translation rules handle the cases where there is more than one action.

Translation Rule 2

This states that `do {e; es} → e >> do {es}`. This tells us what to do if we have an action (`e`) followed by a list of actions (`es`). Here, we make use of the `>>` function defined in monads. All this rule says is that to **do** `{e; es}`, first we perform the action `e`, throw away the result and then **do** `es`.

For instance, if `e` is `putStrLn s` for some string `s`, then the translation of `do {e; es}` is to perform `e` (i.e., print the string) and then **do** `es`. This is clearly what we want.

Translation Rule 3

This states that `do {let decls; es} → let decls in do {es}`. This basically tells us how to deal with **lets** inside of a **do** statement. All we do is lift the `let` declarations within the **let** out and **do** whatever comes after the declarations.

Translation Rule 4

This states that `do {p <- e; es} → let ok p = do {es} ; ok _ = fail "..."`
`in e >>= ok`. Again, it is not exactly obvious what is going on here. However, an alternate formulation of this rule that is roughly equivalent is: `do {p <- e;`

$es\} \rightarrow e \gg= \backslash p \rightarrow es$. Here, it is clear what is happening. We run the action e , and then send the results into es , but first give the result the name p .

The reason for the complex definition is that p doesn't need to simply be a variable; it could be some complex pattern. For instance, the following is valid code:

```
foo = do ('a':'b':'c':x:xs) <- getLine
        putStrLn (x:xs)
```

In this, we're assuming that the results of the action `getLine` will begin with the string "abc" and will have at least one more character. The question becomes what should happen if this pattern match fails. The compiler could simply throw an error like usual for failed pattern matches, but since we're within a monad, we have access to a special `fail` function, and we'd prefer to fail using that function, rather than the "catch all" `error` function. Thus, the translation as defined allows the compiler to fill in the . . . with an appropriate error message about the pattern matching having failed. Other than that, though, the two definitions are equivalent.

9.2 Definition

There are three rules which all monads must obey (and it is up to *you* to ensure your monad laws monads obey these rules) called the "Monad Laws":

1. $\text{return } a \gg= f \equiv f \ a$
2. $f \gg= \text{return} \equiv f$
3. $f \gg= (\backslash x \rightarrow g \ x \gg= h) \equiv (f \gg= g) \gg= h$

Let's look at each of these individually.

Law 1

This states that $\text{return } a \gg= f \equiv f \ a$. Suppose we think about monads as computations. This means that if we create a trivial computation which simply returns the value a irregardless of anything else going on (this is the `return a` part) and then bind it together with some other computation f , then this is equivalent to simply performing the computation f on a directly.

For example, suppose f is the function `putStrLn` and a is the string "Hello World". This states binding a computation whose result is "Hello World" to `putStrLn` is the same as simply printing it to the screen. This seems to make sense.

In `do` notation, this law states that the following two programs are equivalent:

```
law1a = do
  x <- return a
  f x
```

```
law1b = do
  f a
```

Law 2

This states that $f \gg= \text{return} \equiv f$. That is, f is some computation and the law states that if we perform the computation f and then pass the result on to the trivial `return` function, then this is the same as simply performing the computation.

This should be obvious. Think of f as `getLine` (reads a string from the keyboard). This states that reading a string and then returning the value read is exactly the same as just reading the string.

In **do** notation, this law states that the following two programs are equivalent:

```
law2a = do
  x <- f
  return x

law2b = do
  f
```

Law 3

This states that $f \gg= (\lambda x \rightarrow g\ x \gg= h) \equiv (f \gg= g) \gg= h$. This is not as easy to grasp at first glance as the other two. It is essentially an associativity law for monads. associative

■ NOTE ■ Outside the world of monads, a function \cdot is associative if $(f \cdot g) \cdot h = f \cdot (g \cdot h)$. For instance, $+$ and $*$ are associative, since bracketing on these functions doesn't make a difference. On the other hand, $-$ and $/$ are not associative since, for example, $5 - (3 - 1) \neq (5 - 3) - 1$.

If we throw away the messiness with the lambdas, we get that this law states: $f \gg= (g \gg= h) \equiv (f \gg= g) \gg= h$. The intuition behind this law is that when we're string together actions, it doesn't matter how we group them.

For a concrete example, take f to be `getLine`. Take g to be an action which takes a value as input, prints it to the screen and reads another string via `getLine` and then returns that string. Take h to be `putStrLn`.

Let's consider what $(\lambda x \rightarrow g\ x \gg= h)$ does. It takes a value called x , and runs g on it, feeding the results into h . In this instance, this means that it's going to take a value, print it, read another value and then print that. Thus, the entire left hand side of the law reads a string and then does what we've just described.

On the other hand, consider $(f \gg= g)$. This action reads a string from the keyboard, prints it, and then reads another string, returning that newly read string as a

result. When we bind this with `h` as on the right hand side of the law, we get an action which does this, and then prints the results.

Clearly, these two actions are the same.

While this explanation is quite complicated, and the text of the law is also quite complicated, the actual meaning is simple: if we have three actions and we compose them in the same order, it doesn't matter where we put the parentheses. The rest is just notation.

In `do` notation, this requires equivalence of the following two programs:

```
law3a = do
  x <- f
  do g <- x
    h y

law3b = do
  y <- do x <- f
        g x
  h y
```

9.3 A Simple State Monad

One of the simplest monads we can craft is a state-passing monad. As we know, usually in Haskell all state information must be passed to functions explicitly as arguments. Using monads, we can effectively hide some state information.

Suppose we have a function `f` of type `a → b` and we need to add state to this function. In general, if state is of type `state`, we can encode this by changing the type of `f` to `a → state → (state, b)`. That is, the new version of `f` takes the original parameter of type `a` as well as a new state parameter. And in addition to returning the value of type `b`, it also returns an updated state, encoded in a tuple.

For instance, suppose we have a binary tree defined as:

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

Now, we can write a simple `map` function to apply some function to each value in the leaves:

```
mapTree :: (a -> b) -> Tree a -> Tree b
mapTree f (Leaf a) = Leaf (f a)
mapTree f (Branch lhs rhs) =
  Branch (mapTree f lhs) (mapTree f rhs)
```

This is all well and good until we need to write a function which numbers the leaves left to right. In a sense, we need to add state to the `mapTree` function, which keeps track of how many leaves we've numbered so far. We can augment the function to something like:

```
mapTreeState :: (a -> state -> (state, b)) ->
              Tree a -> state -> (state, Tree b)
mapTreeState f (Leaf a) state =
  let (state', b) = f a state
  in (state', Leaf b)
mapTreeState f (Branch lhs rhs) state =
  let (state' , lhs') = mapTreeState f lhs state
      (state'' , rhs') = mapTreeState f rhs state'
  in (state'', Branch lhs' rhs')
```

This is beginning to get a bit unwieldy, and the type signature is getting harder and harder to understand. What we want to do is abstract away the state passing part. That is, the differences between `mapTree` and `mapTreeState` are: (1) the augmented `f` type, (2) we replaced the type `-> Tree b` with `-> state -> (state, Tree b)`. Notice that both types changed in exactly the same way. We can abstract this away with a type synonym declaration:

```
type State st a = st -> (st, a)
```

To go along with this type, we write two functions:

```
returnState :: a -> State st a
returnState a = \st -> (st, a)

bindState :: State st a -> (a -> State st b) -> State st b
bindState m k = \st ->
  let (st', a) = m st
      m'      = k a
  in m' st'
```

Let's examine each of these in turn. The first function, `returnState`, takes a value of type `a` and creates something of type `State st a`. If we think of the `st` as the state, and the value of type `a` as the value, then this is a function which doesn't change the state and returns the value `a`.

The `bindState` function looks distinctly like the interior `let` declarations in `mapTreeState`. It takes two arguments. The first argument is an action which returns something of type `a` with state `st`. The second is a function which takes this `a` and produces something of type `b` also with the same state. The result of `bindState` is essentially the result of transforming the `a` into a `b`.

The definition of `bindState` takes an initial state, `st`. It first applies this to the `State st a` argument called `m`. This gives back a new state `st'` and a value `a`. It then lets the function `k` act on `a`, producing something of type `State st b`, called `m'`. We finally run `m'` with the new state `st'`.

We write a new function, `mapTreeStateM` and give it the type:

```
mapTreeStateM :: (a -> State st b) -> Tree a -> State st (Tree b)
```

Using these “plumbing” functions, `returnState` and `bindState`, we can write this function without ever having to explicitly talk about the state:

```
mapTreeStateM f (Leaf a) =
  f a `bindState` \b ->
  returnState (Leaf b)
mapTreeStateM f (Branch lhs rhs) =
  mapTreeStateM f lhs `bindState` \lhs' ->
  mapTreeStateM f rhs `bindState` \rhs' ->
  returnState (Branch lhs' rhs')
```

In the `Leaf` case, we apply `f` to `a` and then *bind* the result to a function which takes the result and returns a `Leaf` with the new value.

In the `Branch` case, we recurse on the left-hand-side, binding the result to a function which recurses on the right-hand-side, binding that to a simple function which returns the newly created `Branch`.

As you have probably guessed by this point, `State st a` is a monad and `returnState` is analogous to the overloaded `return` method and `bindState` is analogous to the overloaded `>>=` method. In fact, we can verify that `State st a` obeys the monad laws:

Law 1 states: `return a >>= f ≡ f a`. Let’s calculate on the left hand side, substituting our names:

```
returnState a `bindState` f
==>
  \st -> let (st', a) = (returnState a) st
          m'      = f a
          in m' st'
==>
  \st -> let (st', a) = (\st -> (st, a)) st
          in (f a) st'
==>
  \st -> let (st', a) = (st, a)
          in (f a) st'
==>
  \st -> (f a) st
==>
  f a
```

In the first step, we simply substitute the definition of `bindState`. In the second step, we simplify the last two lines and substitute the definition of `returnState`. In the third step, we apply `st` to the lambda function. In the fourth step, we rename `st'` to `st` and remove the **let**. In the last step, we eta reduce.

Moving on to *Law 2*, we need to show that $f \gg= \text{return} \equiv f$. This is shown as follows:

```
f `bindState` returnState
==>
\st -> let (st', a) = f st
        in (returnState a) st'
==>
\st -> let (st', a) = f st
        in (\st -> (st, a)) st'
==>
\st -> let (st', a) = f st
        in (st', a)
==>
\st -> f st
==>
f
```

Finally, we need to show that `State` obeys the third law: $f \gg= (\lambda x \rightarrow g x \gg= h) \equiv (f \gg= g) \gg= h$. This is much more involved to show, so we will only sketch the proof here. Notice that we can write the left-hand-side as:

```
\st -> let (st', a) = f st
        in (\x -> g x `bindState` h) a st'
==>
\st -> let (st', a) = f st
        in (g a `bindState` h) st'
==>
\st -> let (st', a) = f st
        in (\st' -> let (st'', b) = g a
                    in h b st'') st'
==>
\st -> let (st' , a) = f st
        (st'', b) = g a st'
        (st''', c) = h b st''
        in (st''', c)
```

The interesting thing to note here is that we have both action applications on the same **let** level. Since **let** is associative, this means that we can convert this to either bracketing and it won't matter. Of course, this is an informal, "hand waving" argument and it would take a few more derivations to actually prove it, but this gives the general idea.

Now that we know that `State st a` is actually a monad, we'd like to make it an instance of the `Monad` class. Unfortunately, the straightforward way of doing this doesn't work. We cannot write:

```
instance Monad (State st) where { ... }
```

This is because you cannot make instances out of non-fully-applied type synonyms. What we need to do instead is to convert the type synonym into a **newtype**, as:

```
newtype State st a = State (st -> (st, a))
```

Unfortunately, this means that we need to do some packing and unpacking of the `State` constructor in the `Monad` instance declaration, but it's not too bad:

```
instance Monad (State state) where
  return a = State (\state -> (state, a))
  State run >>= action = State run'
  where run' st =
        let (st', a) = run st
            State run'' = action a
        in run'' st'
```

`mapTreeM` Now, we can write our `mapTreeM` function as:

```
mapTreeM :: (a -> State state b) -> Tree a ->
           State state (Tree b)
mapTreeM f (Leaf a) = do
  b <- f a
  return (Leaf b)
mapTreeM f (Branch lhs rhs) = do
  lhs' <- mapTreeM f lhs
  rhs' <- mapTreeM f rhs
  return (Branch lhs' rhs')
```

which is significantly cleaner than before. In fact, if we remove the type signature, we get the more general type:

```
mapTreeM :: Monad m => (a -> m b) -> Tree a ->
           m (Tree b)
```

That is, `mapTreeM` can be run in *any* monad, not just our `State` monad.

Now, the nice thing about encapsulating stuff like this is that we can provide functions to get and change the current state. These look like:

```
getState
putState
```

```

getState :: State state state
getState = State (\state -> (state, state))

putState :: state -> State state ()
putState new = State (\_ -> (new, ()))

```

Here, `getState` is a monadic operation which takes the current state, passes it through unchanged, and then also returns it as the value. The `putState` function takes a new state and produces an action which ignores the current state and inserts the new one.

Now, we can write our `numberTree` function as:

```

numberTree :: Tree a -> State Int (Tree (a, Int))
numberTree tree = mapTreeM number tree
  where number v = do
        cur <- getState
        putState (cur+1)
        return (v,cur)

```

Finally, we need to be able to run the action by providing an initial state:

```

runStateM :: State state a -> state -> a
runStateM (State f) st = snd (f st)

```

Now, we can provide an example Tree:

```

testTree =
  Branch
    (Branch
      (Leaf 'a')
      (Branch
        (Leaf 'b')
        (Leaf 'c')))
    (Branch
      (Leaf 'd')
      (Leaf 'e'))

```

And number it:

```

State> runStateM (numberTree testTree) 1
Branch (Branch (Leaf ('a',1)) (Branch (Leaf ('b',2))
  (Leaf ('c',3)))) (Branch (Leaf ('d',4))
  (Leaf ('e',5)))

```

This may seem like an awful lot of work to do something simple. However, note the new power of `mapTreeM`. We can also print out the leaves of the tree in a left-to-right fashion as:

```
State> mapTreeM print testTree
'a'
'b'
'c'
'd'
'e'
```

This crucially relies on the fact that `mapTreeM` has the more general type involving arbitrary monads, not just the state monad. Furthermore, we can write an action which will make each leaf value equal to its old value as well as all the values preceding:

```
fluffLeaves tree = mapTreeM fluff tree
  where fluff v = do
          cur <- getState
          putState (v:cur)
          return (v:cur)
```

and can see it in action:

```
State> runStateM (fluffLeaves testTree) []
Branch (Branch (Leaf "a") (Branch (Leaf "ba")
  (Leaf "cba"))) (Branch (Leaf "dcba")
  (Leaf "edcba"))
```

In fact, you don't even need to write your own monad instance and datatype. All this is built in in the `Control.Monad.State` module. There, our `runStateM` is called `evalState`; our `getState` is called `get` and our `putState` is called `put`.

This module also contains a *state transformer monad*, which we will discuss in Section 9.7.

9.4 Common Monads

It turns out that many of our favorite datatypes are actually themselves monads. Consider, for instance, lists. They have a monad definition which looks something like:

```
instance Monad [] where
  return x = [x]
  l >>= f = concatMap f l
  fail _ = []
```

This enables us to use lists in `do` notation. For instance, given the definition:

```
cross l1 l2 = do
  x <- l1
  y <- l2
  return (x,y)
```

we get a cross-product function:

```
Monads> cross "ab" "cde"
[( 'a', 'd' ), ( 'a', 'e' ), ( 'a', 'f' ), ( 'b', 'd' ), ( 'b', 'e' ), ( 'b', 'f' )]
```

It is not a coincidence that this looks very much like the list comprehension form: list comprehensions

```
Prelude> [(x,y) | x <- "ab", y <- "def"]
[( 'a', 'd' ), ( 'a', 'e' ), ( 'a', 'f' ), ( 'b', 'd' ), ( 'b', 'e' ), ( 'b', 'f' )]
```

List comprehension form is simply an abbreviated form of a monadic statement using lists. In fact, in older versions of Haskell, the list comprehension form could be used for *any* monad, not just lists. However, in Haskell these days, this is no longer allowed.

The Maybe type is also a monad, with failure being represented as `Nothing` and `Maybe` with success as `Just`. We get the following instance declaration:

```
instance Monad Maybe where
  return a      = Just a
  Nothing >>= f = Nothing
  Just x  >>= f = f x
  fail _      = Nothing
```

We can use the *same* cross product function that we did for lists on `Maybe`s. This is because the **do** notation works for any monad and there's nothing specific to lists about the cross function.

```
Monads> cross (Just 'a') (Just 'b')
Just ('a','b')
Monads> cross (Nothing :: Maybe Char) (Just 'b')
Nothing
Monads> cross (Just 'a') (Nothing :: Maybe Char)
Nothing
Monads> cross (Nothing :: Maybe Char) (Nothing :: Maybe Char)
Nothing
```

What this means is that if we write a function (like `searchAll` from Section 8.4) `searchAll` only in terms of monadic operators, we can use it with any monad, depending on what we mean. Using real monadic functions (not **do** notation), the `searchAll` function looks something like:


```
searchAll g@(Graph vl el) src dst
  | src == dst = return [src]
  | otherwise  = search' el
  where search' [] = fail "no path"
        search' ((u,v,_):es)
          | src == u  = searchAll g v dst >>= \path ->
                    return (u:path)
          | otherwise = search' es
```

The type of this function is `Monad m => Graph v e -> Int -> Int -> m [Int]`. This means that no matter what monad we're using at the moment, this will be able to perform the calculation. Suppose we have the following graph:

```
gr = Graph [(0, 'a'), (1, 'b'), (2, 'c'), (3, 'd')]
       [(0,1,'l'), (0,2,'m'), (1,3,'n'), (2,3,'m')]
```

This represents a graph with four nodes, labelled **a**, **b**, **c** and **d**. There is an edge from **a** to both **b** and **c**. There is also an edge from both **b** and **c** to **d**. Using the `Maybe` monad, we can compute the path from **a** to **d**:

```
Monads> searchAll gr 0 3 :: Maybe [Int]
Just [0,1,3]
```

we provide the type signature so that the interpreter knows what monad we're using. If we try to search in the opposite direction, there is no path which is represented as `Nothing` in the `Maybe` monad:

```
Monads> searchAll gr 3 0 :: Maybe [Int]
Nothing
```

Note that the string “no path” has disappeared since there's no way for the `Maybe` monad to record this.

If we perform the same impossible search in the list monad, we get the empty list, indicating no path:

```
Monads> searchAll gr 3 0 :: [[Int]]
[]
```

If we perform the possible search, we get back a list containing the first path:

```
Monads> searchAll gr 0 3 :: [[Int]]
[[0,1,3]]
```

You may have expected it to return *all* paths, but as coded it does not. See Section 9.6 for more about using lists to represent nondeterminism.

If we use the IO monad, we can actually get at the error message, since IO knows how to keep track of that stuff:

```
Monads> searchAll gr 0 3 :: IO [Int]
Monads> it
[0,1,3]
Monads> searchAll gr 3 0 :: IO [Int]
*** Exception: user error
Reason: no path
```

In the first case, we needed to type `it` to get GHCi to actually evaluate the search.

There is one problem with this implementation of `searchAll`: if it finds an edge which does not lead to a solution, it won't be able to backtrack. This has to do with the recursive call to `searchAll` inside of `search'`. Consider, for instance, what happens if `searchAll g v dst` doesn't find a path. There's no way for this implementation to recover. For instance, if we remove the edge from node **b** to node **d**, we should still be able to find a path from **a** to **d**, but this algorithm can't find it. We define:

```
gr2 = Graph [(0, 'a'), (1, 'b'), (2, 'c'), (3, 'd')]
        [(0,1,'l'), (0,2,'m'), (2,3,'m')]
```

And then try to search:

```
Monads> searchAll gr2 0 3
*** Exception: user error
Reason: no path
```

To fix this, we need a function like `combine` from our `Computation` class. We will see how to do this in Section 9.6.

Exercises

Exercise 26 *Verify that Maybe obeys the three monad laws.*

Exercise 27 *The type Either String is a monad which can keep track of errors. Write an instance for it and then try doing this search in this monad.*

Hint: Your instance declaration should begin: instance Monad (Either String) where.

9.5 Monadic Combinators

The `Monad/Control.Monad` library contains a few very useful monadic combinators which haven't yet been discussed thoroughly. The ones we will discuss in this section, together with their types, are (`m` is always an instance of `Monad`):

- `(=<<) :: (a -> m b) -> m a -> m b`
- `mapM :: (a -> m b) -> [a] -> m [b]`
- `mapM_ :: (a -> m b) -> [a] -> m ()`
- `filterM :: (a -> m Bool) -> [a] -> m [a]`
- `foldM :: (a -> b -> m a) -> a -> [b] -> m a`
- `sequence :: [m a] -> m [a]`
- `sequence_ :: [m a] -> m ()`
- `liftM :: (a -> b) -> m a -> m b`
- `when :: Bool -> m () -> m ()`
- `join :: m (m a) -> m a`

In general, functions with an underscore at the end are equivalent to the ones without, except they do not return any value.

The `=<<` function is exactly the same as `>>=`, except it takes its arguments in the opposite order. For instance, in the IO monad, we can write equivalently either of the following:

```
Monads> writeFile "foo" "hello world!" >>
          (readFile "foo" >>= putStrLn)
hello world!
Monads> writeFile "foo" "hello world!" >>
          (putStrLn =<< readFile "foo")
hello world!
```

`mapM` The `mapM`, `filterM` and `foldM` are our old friends `map`, `filter` and `foldr`
`filterM` wrapped up inside monads. These functions are incredibly useful (particularly `foldM`)
`foldM` when working with monads. We can use `mapM_`, for instance, to print a list of things
to the screen:

```
Monads> mapM_ print [1,2,3,4,5]
1
2
3
4
5
```

We can use `foldM` to sum a list and print the intermediate sum at each step:

```
Monads> foldM (\a b -> putStrLn (show a ++ "+" ++
                                show b ++ "=" ++ show (a+b)) >>
                                return (a+b)) 0 [1..5]
0+1=1
1+2=3
3+3=6
6+4=10
10+5=15
Monads> it
15
```

The `sequence` and `sequence_` functions simply “execute” a list of actions. For instance:

```
Monads> sequence [print 1, print 2, print 'a']
1
2
'a'
*Monads> it
[(),(),()]
*Monads> sequence_ [print 1, print 2, print 'a']
1
2
'a'
*Monads> it
()
```

We can see that the underscored version doesn’t return each value, while the non-underscored version returns the list of the return values.

The `liftM` function “lifts” a non-monadic function to a monadic function. (Do not confuse this with the `lift` function used for monad transformers in Section 9.7.) This is useful for shortening code (among other things). For instance, we might want to write a function which prepends each line in a file with its line number. We can do this with:

```
numberFile :: FilePath -> IO ()
numberFile fp = do
  text <- readFile fp
  let l = lines text
      n = zipWith (\n t -> show n ++ ' ' : t) [1..] l
  mapM_ putStrLn n
```

However, we can shorten this using `liftM`:

```
numberFile :: FilePath -> IO ()
numberFile fp = do
  l <- lines `liftM` readFile fp
  let n = zipWith (\n t -> show n ++ ' ' : t) [1..] l
      mapM_ putStrLn n
```

In fact, you can apply any sort of (pure) processing to a file using `liftM`. For instance, perhaps we also want to split lines into words; we can do this with:

```
...
w <- (map words . lines) `liftM` readFile fp
...
```

Note that the parentheses are required, since the `(.)` function has the same fixity as `liftM`.

Lifting pure functions into monads is also useful in other monads. For instance `liftM` can be used to apply function inside of `Just`. For instance:

```
Monads> liftM (+1) (Just 5)
Just 6
*Monads> liftM (+1) Nothing
Nothing
```

when The `when` function executes a monadic action only if a condition is met. For instance, we might only want to print non-empty lines:

```
Monads> mapM_ (\l -> when (not $ null l) (putStrLn l))
              [ "", "abc", "def", "", "", "ghi" ]
abc
def
ghi
```

Of course, the same could be accomplished with `filter`, but sometimes `when` is more convenient.

join Finally, the `join` function is the monadic equivalent of `concat` on lists. In fact, when `m` is the list monad, `join` is exactly `concat`. In other monads, it accomplishes a similar task:

```
Monads> join (Just (Just 'a'))
Just 'a'
Monads> join (Just (Nothing :: Maybe Char))
Nothing
Monads> join (Nothing :: Maybe (Maybe Char))
Nothing
```

```
Monads> join (return (putStrLn "hello"))
hello
Monads> return (putStrLn "hello")
Monads> join [[1,2,3],[4,5]]
[1,2,3,4,5]
```

These functions will turn out to be even more useful as we move on to advanced topics in Chapter 10.

9.6 MonadPlus

Given only the `>>=` and `return` functions, it turns out to be impossible to write a function like `combine` with type `c a → c a → c a`. However, such a function is so generally useful that it exists in another class called **MonadPlus**. In addition to having a `combine` function, instance of **MonadPlus** also have a “zero” element which is the identity under the “plus” (i.e., `combine`) action. The definition is:

`combine`
`MonadPlus`

```
class Monad m => MonadPlus m where
  mzero :: m a
  mplus :: m a -> m a -> m a
```

In order to gain access to **MonadPlus**, you need to import the `Monad` module (or `Control.Monad` in the hierarchical libraries).

In Section 9.4, we say that `Maybe` and `list` were both monads. In fact, they are also both instances of **MonadPlus**. In the case of `Maybe`, the zero element is `Nothing`; in the case of `lists`, it is the empty list. The `mplus` operation on `Maybe` is `Nothing` if both elements are `Nothing`; otherwise it is the first `Just` value. For `lists`, `mplus` is the same as `++`.

`Maybe`
`lists`

That is, the instance declarations look like:

```
instance MonadPlus Maybe where
  mzero = Nothing
  mplus Nothing y = y
  mplus x _ = x

instance MonadPlus [] where
  mzero = []
  mplus x y = x ++ y
```

We can use this class to reimplement the search function we’ve been exploring such that it will explore all possible paths. The new function looks like:

```
searchAll2 g@(Graph v1 e1) src dst
  | src == dst = return [src]
```

```

| otherwise = search' el
where search' [] = fail "no path"
      search' ((u,v,_):es)
          | src == u =
            (searchAll2 g v dst >>= \path ->
             return (u:path)) `mplus`
            search' es
          | otherwise = search' es

```

Now, when we're going through the edge list in `search'` and we come across a matching edge, not only do we explore this path but we also continue to explore the out-edges of the current node in the recursive call to `search'`.

The IO monad is not an instance of `MonadPlus`, we're not able to execute the search with this monad. We can see that when using lists as the monad, we (a) get all possible paths in `gr` and (b) get a path in `gr2`.

```

MPlus> searchAll2 gr 0 3 :: [[Int]]
[[0,1,3],[0,2,3]]
MPlus> searchAll2 gr2 0 3 :: [[Int]]
[[0,2,3]]

```

You might be tempted to implement this as:

```

searchAll2 g@(Graph v1 el) src dst
| src == dst = return [src]
| otherwise = search' el
where search' [] = fail "no path"
      search' ((u,v,_):es)
          | src == u = do
            path <- searchAll2 g v dst
            rest <- search' es
            return ((u:path) `mplus` rest)
          | otherwise = search' es

```

But note that this doesn't do what we want. Here, if the recursive call to `searchAll2` fails, we don't try to continue down and execute `search' es`. The call to `mplus` must be at the top level in order for it to work.

Exercises

Exercise 28 Suppose that we changed the order of arguments to `mplus`. I.e., the matching case of `search'` looked like:

```

search' es `mplus`
(searchAll2 g v dst >>= \path ->
 return (u:path))

```

How would you expect this to change the results when using the list monad on `gr`? Why?

9.7 Monad Transformers

Often we want to “piggyback” monads on top of each other. For instance, there might be a case where you need access both to IO functions through the IO monad and state functions through some state monad. In order to accomplish this, we introduce a `MonadTrans` class which essentially lifts one monad on top of another. This class has a simple method: `lift`. The class declaration for `MonadTrans` is:

`MonadTrans`
`lift`

```
class MonadTrans t where
  lift :: Monad m => m a -> t m a
```

The idea here is that `t` is the outer monad and that `m` lives inside of it. In order to execute a command of type `Monad m => m a`, we first `lift` it into the transformer.

The simplest example of a transformer (and arguably the most useful) is the state transformer monad, which is a state monad wrapped around an arbitrary monad. Before, we defined a state monad as:

state monad

```
newtype State state a = State (state -> (state, a))
```

Now, instead of using a function of type `state -> (state, a)` as the monad, we assume there’s some other monad `m` and make the internal action into something of type `state -> m (state, a)`. This gives rise to the following definition for a state transformer:

state transformer

```
newtype StateT state m a =
  StateT (state -> m (state, a))
```

For instance, we can think of `m` as IO. In this case, our state transformer monad is able to execute actions in the IO monad. First, we make this an instance of `MonadTrans`:

```
instance MonadTrans (StateT state) where
  lift m = StateT (\s -> do a <- m
                    return (s,a))
```

Here, lifting a function from the realm of `m` to the realm of `StateT state` simply involves keeping the state (the `s` value) constant and executing the action.

Of course, we also need to make `StateT` a monad itself. This is relatively straightforward, provided that `m` is already a monad:


```
instance Monad m => Monad (StateT state m) where
  return a = StateT (\s -> return (s,a))
  StateT m >>= k = StateT (\s -> do
    (s', a) <- m s
    let StateT m' = k a
        m' s')
  fail s = StateT (\_ -> fail s)
```

The idea behind the definition of `return` is that we keep the state constant and simply return the state/a pair in the enclosed monad. Note that the use of `return` in the definition of `return` refers to the enclosed monad, not the state transformer.

In the definition of `bind`, we create a new `StateT` which takes a state `s` as an argument. First, it applies this state to the first action (`StateT m`) and gets the new state and answer out. It then runs the `k` action on this new state and gets a new transformer out. It finally applies the new state to this transformer. This definition is quite parallel to the definition of `bind` for the standard `State` monad described earlier.

The `fail` function simply calls `fail` on the enclosed monad, since state transformers don't natively know how to deal with failure.

Of course, in order to actually use this monad, we need to provide function `getT`, `putT` and `evalStateT`. These are analogous to `getState`, `putState` and `runStateM` from Section 9.3:

`evalStateT`

```
getT :: Monad m => StateT s m s
getT = StateT (\s -> return (s, s))

putT :: Monad m => s -> StateT s m ()
putT s = StateT (\_ -> return (s, ()))

evalStateT :: Monad m => StateT s m a -> s -> m a
evalStateT (StateT m) state = do
  (s', a) <- m state
  return a
```

These functions should be straightforward. Note, however, that the result of `evalStateT` is actually a monadic action in the enclosed monad. This is typical of monad transformers: they don't know how to actually run things in their enclosed monad (they only know how to `lift` actions). Thus, what you get out is a monadic action in the inside monad (in our case, `IO`), which you then need to run yourself.

We can use state transformers to reimplement a version of our `mapTreeM` function from Section 9.3. The only change here is that when we get to a leaf, we print out the value of the leaf; when we get to a branch we just print out "Branch".

```
mapTreeM action (Leaf a) = do
  lift (putStrLn ("Leaf " ++ show a))
```

```

b <- action a
return (Leaf b)
mapTreeM action (Branch lhs rhs) = do
  lift (putStrLn "Branch")
  lhs' <- mapTreeM action lhs
  rhs' <- mapTreeM action rhs
  return (Branch lhs' rhs')

```

The only difference between this function and the one from Section 9.3 is the calls to `lift (putStrLn ...)` as the first line. The `lift` tells us that we're going to be executing a command in an enclosed monad. In this case, the enclosed monad is `IO`, since the command lifted is `putStrLn`.

The type of this function is actually relatively complex:

```

mapTreeM :: (MonadTrans t, Monad (t IO), Show a) =>
  (a -> t IO a1) -> Tree a -> t IO (Tree a1)

```

Ignoring for a section the class constraints, this says that `mapTreeM` takes an action and a tree and returns a tree. This is pretty much the same as before. In this, we require that `t` is a monad transformer (since we apply `lift` in it); we require that `t IO` is a monad, since we use `putStrLn` we know that the enclosed monad is `IO`; finally, we require that `a` is an instance of `show` – this is simply because we use `show` to show the value of leaves.

Now, we simply change `numberTree` to use this version of `mapTreeM` and the new versions of `get` and `put` and we end up with:

```

numberTree tree = mapTreeM number tree
  where number v = do
    cur <- getT
    putT (cur+1)
    return (v,cur)

```

Using this, we can run our monad:

```

MTrans> evalStateT (numberTree testTree) 0
Branch
Branch
Leaf 'a'
Branch
Leaf 'b'
Leaf 'c'
Branch
Leaf 'd'
Leaf 'e'
*MTrans> it

```

```
Branch (Branch (Leaf ('a',0))
              (Branch (Leaf ('b',1)) (Leaf ('c',2))))
      (Branch (Leaf ('d',3)) (Leaf ('e',4)))
```

One problem not specified in our discussion of `MonadPlus` is that our search cycles algorithm will fail to terminate on graphs with cycles. Consider:

```
gr3 = Graph [(0, 'a'), (1, 'b'), (2, 'c'), (3, 'd')]
          [(0,1,'l'), (1,0,'m'), (0,2,'n'),
           (1,3,'o'), (2,3,'p')]
```

In this graph, there is a back edge from node **b** back to **a**. If we attempt to run `searchAll2`, no matter what monad we use, we will fail to terminate. Moreover, if we move this erroneous edge to the end of the list (and call this `gr4`), the result of `searchAll2 gr4 0 3` will contain an infinite number of paths: presumably we only want ones which don't contain cycles.

In order to get around this problem, we need to introduce state. Namely, we need to keep track of which nodes we have visited so that we don't visit them any more.

We can do this as follows:

```
searchAll5 g@(Graph vl el) src dst
  | src == dst = do
    visited <- getT
    putT (src:visited)
    return [src]
  | otherwise = do
    visited <- getT
    putT (src:visited)
    if src `elem` visited
      then mzero
      else search' el
  where
    search' [] = mzero
    search' ((u,v,_):es)
      | src == u =
        (do path <- searchAll5 g v dst
         return (u:path)) `mplus`
      search' es
      | otherwise = search' es
```

Here, we implicitly use a state transformer (see the calls to `getT` and `putT`) to keep track of visited states. We only continue to recurse when we encounter a state we haven't yet visited. Furthermore, when we recurse, we add the current state to our set of visited states.

Now, we can run the state transformer and get out only the correct paths, even on the cyclic graphs:

```
MTrans> evalStateT (searchAll5 gr3 0 3) [] :: [[Int]]
[[0,1,3],[0,2,3]]
MTrans> evalStateT (searchAll5 gr4 0 3) [] :: [[Int]]
[[0,1,3],[0,2,3]]
```

Here, the empty list provided as an argument to `evalStateT` is the initial state: i.e., the initial visited list. In our case, it is empty.

We can also provide an `execStateT` method which, instead of giving us back a result, gives us back the final state. This function looks like:

```
execStateT :: Monad m => StateT s m a -> s -> m s
execStateT (StateT m) state = do
  (s', a) <- m state
  return s'
```

This is not so useful in our case, as it will give back to us exactly the reverse of `evalStateT` (try it and find out!), but can be useful in general (if, for instance, we need to know how many numbers are used in `numberTree`).

Exercises

Exercise 29 Write a function `searchAll6` based on the code for `searchAll2` which, at every entry to the main function (not the recursion over the edge list), prints the search being conducted. For instance, the output generated for `searchAll6 gr 0 3` should look like:

```
Exploring 0 -> 3
Exploring 1 -> 3
Exploring 3 -> 3
Exploring 2 -> 3
Exploring 3 -> 3
MTrans> it
[[0,1,3],[0,2,3]]
```

In order to do this, you will have to define your own list monad transformer and make appropriate instances of it.

Exercise 30 Combine the `searchAll5` function (from this section) with the `searchAll6` function (from the previous exercise) into a single function called `searchAll7`. This function should perform IO as in `searchAll6` but also keep track of state using a state transformer.

9.8 Parsing Monads

todo: Introduction...

9.8.1 A Simple Parsing Monad

Consider the task of parsing. A simple parsing monad is much like a state monad where the state is the unparsed string. We can represent this exactly as:

```
newtype Parser a = Parser
  { runParser :: String -> Either String (String, a) }
```

We again use `Left err` to be an error condition. This yields standard instance of **Monad** and **MonadPlus**:

```
instance Monad Parser where
  return a = Parser (\xl -> Right (xl,a))
  fail s = Parser (\xl -> Left s)
  Parser m >>= k = Parser $ \xl ->
    case m xl of
      Left s -> Left s
      Right (xl', a) ->
        let Parser n = k a
        in n xl'

instance MonadPlus Parser where
  mzero = Parser (\xl -> Left "mzero")
  Parser p `mplus` Parser q = Parser $ \xl ->
    case p xl of
      Right a -> Right a
      Left err -> case q xl of
                    Right a -> Right a
                    Left _ -> Left err
```

primitives Now, we want to build up a library of parsing “primitives.” The most basic primitive is a parser that will read a specific character. This function looks like:

```
char :: Char -> Parser Char
char c = Parser char'
  where char' [] = Left ("expecting " ++ show c ++
                        " got EOF")
        char' (x:xs)
          | x == c = Right (xs, c)
          | otherwise = Left ("expecting " ++
                              show c ++ " got " ++
                              show x)
```

Here, the parser succeeds only if the first character of the input is the expected character.

We can use this parser to build up a parser for the string “Hello”:

```
helloParser :: Parser String
helloParser = do
  char 'H'
  char 'e'
  char 'l'
  char 'l'
  char 'o'
  return "Hello"
```

This shows how easy it is to combine these parsers. We don't need to worry about the underlying string: the monad takes care of that for us. All we need to do is combine these things. We can test this parser by using `runParser` and supplying input:

`runParser`

```
Parsing> runParser helloParser "Hello"
Right ("","Hello")
Parsing> runParser helloParser "Hello World!"
Right (" World!","Hello")
Parsing> runParser helloParser "hello World!"
Left "expecting 'H' got 'h'"
```

We can have a slightly more general function which will match any character fitting a description:

```
matchChar :: (Char -> Bool) -> Parser Char
matchChar c = Parser matchChar'
  where matchChar' [] =
          Left ("expecting char, got EOF")
        matchChar' (x:xs)
          | c x      = Right (xs, x)
          | otherwise =
            Left ("expecting char, got " ++
                  show x)
```

Using this, we can write a case-insensitive “Hello” parser:

```
ciHelloParser = do
  c1 <- matchChar (\elem "Hh")
  c2 <- matchChar (\elem "Ee")
  c3 <- matchChar (\elem "Ll")
  c4 <- matchChar (\elem "Ll")
  c5 <- matchChar (\elem "Oo")
  return [c1,c2,c3,c4,c5]
```

Of course, we could have used something like `matchChar ((=='h') . toLower)`, but this works just as well. We can test this:

```
Parsing> runParser ciHelloParser "hELlO world!"
Right (" world!","hELlO")
```

Finally, we can have a function which will match any character at all:

```
anyChar :: Parser Char
anyChar = Parser anyChar'
  where anyChar' [] =
          Left ("expecting character, got EOF")
        anyChar' (x:xs) = Right (xs, x)
```

many On top of these primitives we usually build some combinators. The `many` combinator, for instance, will take a parser which parses things of type `a` and will make it into a parser which parses things of type `[a]` (this is a Kleene-star operator):

```
many :: Parser a -> Parser [a]
many (Parser p) = Parser many'
  where many' xl =
          case p xl of
            Left err -> Right (xl, [])
            Right (xl',a) ->
              let Right (xl'', rest) = many' xl'
              in Right (xl'', a:rest)
```

The idea here is that first we try to apply the given parser, `p`. If this fails, we *succeed*, but return the empty list. If `p` succeeds, we recurse and keep trying to apply `p` until it fails. We then return the list we've built up.

In general, there would be many more functions of this sort, and they would be hidden away in a library so that users could actually look inside the `Parser` type. However, using them, you could build up, for instance, a parser which parses (non-negative) integers:

```
int :: Parser Int
int = do
  t1 <- matchChar isDigit
  tr <- many (matchChar isDigit)
  return (read (t1:tr))
```

The idea here is that first we match a digit (the `isDigit` function comes from the module `Char/Data.Char`) and then match as many digits more as we can. We then read the result and return it. We can test this parser as before:

```
Parsing> runParser int "54"
Right ("",54)
```

```
*Parsing> runParser int "54abc"
Right ("abc",54)
*Parsing> runParser int "a54abc"
Left "expecting char, got 'a'"
```

Now, suppose we want to parse a Haskell-style list of ints. This becomes somewhat difficult because at some point we're either going to parse a comma or a close brace, but we don't know when this will happen. This is where the fact that `Parser` is an instance of `MonadPlus` comes in handy: first we try one, then we try the other.

Consider the following code:

```
intList :: Parser [Int]
intList = do
  char '['
  intList' `mplus` (char ']') >> return []
  where intList' = do
        i <- int
        r <- (char ',' >> intList') `mplus`
              (char ']') >> return []
        return (i:r)
```

The first thing this code does is parse and open brace. Then, using `mplus`, it tries one of two things: parsing using `intList'`, or parsing a close brace and returning an empty list.

The `intList'` function assumes we're not yet at the end of the list and so it first parses and int. It then parses the rest of the list. However, it doesn't know whether we're at the end yet or not, so it again using `mplus`. On the one hand, it tries to parse a comma and then recurse; on the other, it parses a close brace and returns the empty list. Either way, it simply prepends the int it parsed itself to the beginning.

One thing that you should be careful of is the order in which you supply arguments to `mplus`. Consider the following parser:

```
tricky =
  mplus (string "Hal") (string "Hall")
```

You might expect this parser to parse both the words "Hal" and "Hall"; however, it only parses the former. You can see this with:

```
Parsing> runParser tricky "Hal"
Right ("","Hal")
Parsing> runParser tricky "Hall"
Right ("l","Hal")
```

This is because it tries to parse "Hal", which succeeds and then it doesn't bother trying to parse "Hall".

You can attempt to fix this by providing a parser primitive which detects end-of-file (really, end-of-string) as:

```
eof :: Parser ()
eof = Parser eof'
  where eof' [] = Right ([], ())
        eof' xl = Left ("Expecting EOF, got " ++
                        show (take 10 xl))
```

You might then rewrite `tricky` using this as:

```
tricky2 = do
  s <- mplus (string "Hal") (string "Hall")
  eof
  return s
```

But this also doesn't work, as we can easily see:

```
Parsing> runParser tricky2 "Hal"
Right ("",())
Parsing> runParser tricky2 "Hall"
Left "Expecting EOF, got \"l\""
```

This is because, again, the `mplus` doesn't know that it needs to parse the whole input. So when you provide it with "Hall", it parses just "Hal" and leaves the last "l" around to be parsed later. This causes `eof` to produce an error message.

The correct way to implement this is:

```
tricky3 =
  mplus (do s <- string "Hal"
          eof
          return s)
        (do s <- string "Hall"
          eof
          return s)
```

We can see that this works:

```
Parsing> runParser tricky3 "Hal"
Right ("","Hal")
Parsing> runParser tricky3 "Hall"
Right ("","Hall")
```

This works precisely because each side of the `mplus` knows that it must read the end.

In this case, fixing the parser to accept both “Hal” and “Hall” was fairly simple, due to the fact that we assumed we would be reading an end-of-file immediately after. Things can get a bit more complicated if we cannot disambiguate immediately.

Exercises

Exercise 31 Write a parser `intListSpace` which will parse int lists, but will allow arbitrary white space (spaces, tabs or newlines) between the commas and brackets.

Given this monadic parser, it is fairly easy to add information regarding source position. For instance, if we’re parsing a large file, it might be helpful to report the line number on which an error occurred. We could do this simply by extending the `Parser` type, modifying the instances and the primitives: line numbers

```

newtype Parser a = Parser
  { runParser :: Int -> String ->
    Either String (Int, String, a) }

instance Monad Parser where
  return a = Parser (\n xl -> Right (n,xl,a))
  fail s = Parser (\n xl -> Left (show n ++
    ": " ++ s))

  Parser m >>= k = Parser $ \n xl ->
    case m n xl of
      Left s -> Left s
      Right (n', xl', a) ->
        let Parser m2 = k a
        in m2 n' xl'

instance MonadPlus Parser where
  mzero = Parser (\n xl -> Left "mzero")
  Parser p `mplus` Parser q = Parser $ \n xl ->
    case p n xl of
      Right a -> Right a
      Left err -> case q n xl of
        Right a -> Right a
        Left _ -> Left err

matchChar :: (Char -> Bool) -> Parser Char
matchChar c = Parser matchChar'
  where matchChar' n [] =
    Left ("expecting char, got EOF")
        matchChar' n (x:xs)
          | c x =
            Right (n+if x=='\n' then 1 else 0

```

```

        , xs, x)
    | otherwise =
      Left ("expecting char, got " ++
           show x)

```

The definitions for `char` and `anyChar` are ellided since they can be written in terms of `matchChar`. The `many` function needs to be modified only to include the new state.

Now, when we run a parser, it will tell us what line number the error is on (if there is an error):

```

Parsing2> runParser helloParser 1 "Hello"
Right (1,"","Hello")
Parsing2> runParser int 1 "a54"
Left "1: expecting char, got 'a'"
Parsing2> runParser intList 1 "[1,2,3,a]"
Left "1: expecting ']' got '1'"

```

We can use the `intListSpace` parser from the prior exercise to see that this does in fact work:

```

Parsing2> runParser intListSpace 1
          "[1 ,2 , 4 \n\n ,a\n]"
Left "3: expecting char, got 'a'"
Parsing2> runParser intListSpace 1
          "[1 ,2 , 4 \n\n\n ,a\n]"
Left "4: expecting char, got 'a'"
Parsing2> runParser intListSpace 1
          "[1 ,\n2 , 4 \n\n\n ,a\n]"
Left "5: expecting char, got 'a'"

```

We can see that the line number on which the error occurs increases as we add additional newlines before the erroneous “a”.

9.8.2 Parsec

As you continue developing your parser, you might want to add more and more features. Luckily, someone (Graham Hutton and Daan Leijen) has already done this for us in the Parsec library. This section is intended to be an introduction to the Parsec library; it by no means covers the whole library, but it should be enough to get you going.

Like our library, Parsec provides a few basic functions to build parsers from characters. These are: `char`, which is the same as our `char`; `anyChar`, which is the same as our `anyChar`; `satisfy`, which is the same as our `matchChar`; `oneOf`, which takes a list of `Chars` and matches any of them; and `noneOf` which is the opposite of `oneOf`.

The primary function Parsec uses to run a parser is `parse`. However, in addition to a parser, this function takes a string representing the same of the file you're parsing. This is so that it can give better error message. We can try parsing with the above functions:

```
ParsecI> parse (char 'a') "stdin" "a"
Right 'a'
ParsecI> parse (char 'a') "stdin" "ab"
Right 'a'
ParsecI> parse (char 'a') "stdin" "b"
Left "stdin" (line 1, column 1):
unexpected "b"
expecting "a"
ParsecI> parse (char 'H' >> char 'a' >> char 'l')
           "stdin" "Hal"
Right 'l'
ParsecI> parse (char 'H' >> char 'a' >> char 'l')
           "stdin" "Hap"
Left "stdin" (line 1, column 3):
unexpected "p"
expecting "l"
```

Here, we can see a few differences between our parser and Parsec: first of all, we don't get back the rest of the string when we run `parse`. Secondly, the error messages produced are much better.

In addition to the basic character parsing functions, Parsec provides primitives for: `spaces`, which is the same as ours; `space` which parses a single space; `letter`, which parses a letter; `digit`, which parses a digit; `string`, which is the same as ours; and a few others.

We can write our `int` and `intList` functions in Parsec as:

```
int :: CharParser st Int
int = do
  il <- digit
  ir <- many digit
  return (read (il:ir))

intList :: CharParser st [Int]
intList = do
  char '['
  intList' `mplus` (char ']') >> return []
  where intList' = do
        i <- int
        r <- (char ',' >> intList') `mplus`
              (char ']') >> return []
        return (i:r)
```

spaces
space
letter
digit
string

First, note the type signatures. The `st` type variable is simply a state variable which we are not using. In the `int` function, we use the `many` function (built in to `Parsec`) together with the `digit` function (also built in to `Parsec`). The `intList` function is actually identical to before.

Note, however, that using `mplus` explicitly is not the preferred method of combining parsers: `Parsec` provides a `<|>` function which is a synonym of `mplus`, but looks nicer:

```
intList :: CharParser st [Int]
intList = do
  char '['
  intList' <|> (char ']' >> return [])
  where intList' = do
        i <- int
        r <- (char ',' >> intList') <|>
              (char ']' >> return [])
        return (i:r)
```

We can test this:

```
ParsecI> parse intList "stdin" "[3,5,2,10]"
Right [3,5,2,10]
ParsecI> parse intList "stdin" "[3,5,a,10]"
Left "stdin" (line 1, column 6):
unexpected "a"
expecting digit
```

In addition to these basic combinators, `Parsec` provides a few other useful ones:

- `choice` takes a list of parsers and performs an or (`<|>`) between all of them.
- `option` takes a default value of type `a` and a parser which returns something of type `a`. It then tries to parse with the parser, but using the default value as the return if the parsing fails.
- `optional` takes a parser which returns `()` and makes a parser which might parse or might not.
- `between` takes three parsers: an open parser, a close parser and a between parser. It runs them in order and returns the value of the between parser.
- `notFollowedBy` takes a parser and returns one which succeeds only if the given parser would have failed.

Suppose we want to parse a simple calculator language which includes only plus and times. Furthermore, for simplicity, assume each embedded expression must be enclosed in parentheses. We can give a datatype for this language as:

```
data Expr = Value Int
          | Expr :+: Expr
          | Expr **: Expr
          deriving (Eq, Ord, Show)
```

And then write a parser for this language as:

```
parseExpr :: Parser Expr
parseExpr = choice
  [ do i <- int; return (Value i)
  , between (char '(') (char ')') $ do
    e1 <- parseExpr
    op <- oneOf "+*"
    e2 <- parseExpr
    case op of
      '+' -> return (e1 :+: e2)
      '*' -> return (e1 **: e2)
  ]
```

Here, the parser alternates between two options (we could have used `<|>`, but wanted to show the `choice` combinator in action). The first simply parses an `int` and then wraps it up in the `Value` constructor. The second option using `between` to parse stuff between parentheses. What is parsed is first an expression, then one of plus or times, then another expression. Depending on what the operator is, it returns either `e1 :+: e2` or `e1 **: e2`.

We can modify this parser so that instead of computing an `Expr`, it simply computes the value:

```
parseValue :: Parser Int
parseValue = choice
  [int
  ,between (char '(') (char ')') $ do
    e1 <- parseValue
    op <- oneOf "+*"
    e2 <- parseValue
    case op of
      '+' -> return (e1 + e2)
      '*' -> return (e1 * e2)
  ]
```

We can use this as:

```
ParsecI> parse parseValue "stdin" "(3*(4+3))"
Right 21
```

Now, suppose we want to introduce bindings into our language. That is, we want to also be able to say “let $x = 5$ in” inside of our expressions and then use the variables we’ve defined. In order to do this, we need to use the `getState` and `setState` (or `updateState`) functions built in to Parsec.

```

parseValueLet :: CharParser (FiniteMap Char Int) Int
parseValueLet = choice
  [ int
  , do string "let "
    c <- letter
    char '='
    e <- parseValueLet
    string " in "
    updateState (\fm -> addToFM fm c e)
    parseValueLet
  , do c <- letter
    fm <- getState
    case lookupFM fm c of
      Nothing -> unexpected ("variable " ++ show c ++
                             " unbound")
      Just i -> return i
  , between (char '(') (char ')') $ do
    e1 <- parseValueLet
    op <- oneOf "+*"
    e2 <- parseValueLet
    case op of
      '+' -> return (e1 + e2)
      '*' -> return (e1 * e2)
  ]

```

The `int` and recursive cases remain the same. We add two more cases, one to deal with let-bindings, the other to deal with usages.

In the let-bindings case, we first parse a “let” string, followed by the character (the `letter` function is a Parsec primitive which parses alphabetic characters) we’re binding, followed by its value (a `parseValueLet`). Then we parse the “in” and update the state to include this binding. Finally, we continue and parse the rest.

In the usage case, we simply parse the character and then look it up in the state. However, if it doesn’t exist, we use the Parsec primitive `unexpected` to report an error.

We can see this parser in action using the `runParser` command, which enables us to provide an initial state:

```

ParsecI> runParser parseValueLet emptyFM "stdin"
           "let c=5 in ((5+4)*c)"
Right 45
*ParsecI> runParser parseValueLet emptyFM "stdin"

```

```
Right 63          "let c=5 in ((5+4)*let x=2 in (c+x))"
*ParsecI> runParser parseValueLet emptyFM "stdin"
                "((let x=2 in 3+4)*x)"
Right 14
```

Note that the bracketing does not affect the definitions of the variables. For instance, in the last example, the usage of “x” is in some sense outside the scope of the definition. However, our parser doesn’t notice this since it operates in a strictly left-to-right fashion. In order to fix this, bindings would have to be removed (see the exercises).

Exercises

Exercise 32 *Modify the `parseValueLet` parser so that it obeys bracketing. In order to do this, you will need to change the state to something like `FiniteMap Char [Int]`, where the `[Int]` is a stack of definitions.*

Chapter 10

Advanced Techniques

10.1 Exceptions

10.2 Mutable Arrays

10.3 Mutable References

10.4 The ST Monad

10.5 Concurrency

10.6 Regular Expressions

10.7 Dynamic Types

Chapter 11

Theory

11.1 Bottom

Chapter 12

For Further Information

Appendix A

Brief Complexity Theory

Complexity Theory is the study of how long a program will take to run, depending on the size of its input. There are many good introductory books to complexity theory and the basics are explained in any good algorithms book. I'll keep the discussion here to a minimum.

The idea is to say how well a program scales with more data. If you have a program that runs quickly on very small amounts of data but chokes on huge amounts of data, it's not very useful (unless you know you'll only be working with small amounts of data, of course). Consider the following Haskell function to return the sum of the elements in a list:

```
sum [] = 0
sum (x:xs) = x + sum xs
```

How long does it take this function to complete? That's a very difficult question; it would depend on all sorts of things: your processor speed, your amount of memory, the exact way in which the addition is carried out, the length of the list, how many other programs are running on your computer, and so on. This is far too much to deal with, so we need to invent a simpler model. The model we use is sort of an arbitrary "machine step." So the question is "how many machine steps will it take for this program to complete?" In this case, it only depends on the length of the input list.

If the input list is of length 0, the function will take either 0 or 1 or 2 or some very small number of machine steps, depending exactly on how you count them (perhaps 1 step to do the pattern matching and 1 more to return the value 0). What if the list is of length 1. Well, it would take however much time the list of length 0 would take, plus a few more steps for doing the first (and only element).

If the input list is of length n , it will take however many steps an empty list would take (call this value y) and then, for each element it would take a certain number of steps to do the addition and the recursive call (call this number x). Then, the total time this function will take is $nx + y$ since it needs to do those additions n many times. These x and y values are called *constant values*, since they are independent of n , and actually dependent only on exactly how we define a machine step, so we really don't

want to consider them all that important. Therefore, we say that the complexity of this sum function is $\mathcal{O}(n)$ (read “order n ”). Basically saying something is $\mathcal{O}(n)$ means that for some constant factors x and y , the function takes $nx + y$ machine steps to complete.

Consider the following sorting algorithm for lists (commonly called “insertion sort”):

```

sort [] = []
sort [x] = [x]
sort (x:xs) = insert (sort xs)
              where insert [] = [x]
                  insert (y:ys) | x <= y    = x : y : ys
                                | otherwise = y : insert ys

```

The way this algorithm works is as follow: if we want to sort an empty list or a list of just one element, we return them as they are, as they are already sorted. Otherwise, we have a list of the form $x:xs$. In this case, we sort xs and then want to insert x in the appropriate location. That’s what the `insert` function does. It traverses the now-sorted tail and inserts x wherever it naturally fits.

Let’s analyze how long this function takes to complete. Suppose it takes $f(n)$ steps to sort a list of length n . Then, in order to sort a list of n -many elements, we first have to sort the tail of the list first, which takes $f(n-1)$ time. Then, we have to insert x into this new list. If x has to go at the end, this will take $\mathcal{O}(n-1) = \mathcal{O}(n)$ steps. Putting all of this together, we see that we have to do $\mathcal{O}(n)$ amount of work $\mathcal{O}(n)$ many times, which means that the entire complexity of this sorting algorithm is $\mathcal{O}(n^2)$. Here, the squared is not a constant value, so we cannot throw it out.

What does this mean? Simply that for really long lists, the `sum` function won’t take very long, but that the `sort` function will take quite some time. Of course there are algorithms that run much more slowly than simply $\mathcal{O}(n^2)$ and there are ones that run more quickly than $\mathcal{O}(n)$.

Consider the random access functions for lists and arrays. In the worst case, accessing an arbitrary element in a list of length n will take $\mathcal{O}(n)$ time (think about accessing the last element). However with arrays, you can access any element immediately, which is said to be in *constant* time, or $\mathcal{O}(1)$, which is basically as fast as any algorithm can go.

There’s much more in complexity theory than this, but this should be enough to allow you to understand all the discussions in this tutorial. Just keep in mind that $\mathcal{O}(1)$ is faster than $\mathcal{O}(n)$ is faster than $\mathcal{O}(n^2)$, etc.

Appendix B

Recursion and Induction

Informally, a function is recursive if its definition depends on itself. The prototypical example is factorial, whose definition is:

$$fact(n) = \begin{cases} 1 & n = 0 \\ n * fact(n - 1) & n > 0 \end{cases}$$

Here, we can see that in order to calculate $fact(5)$, we need to calculate $fact(4)$, but in order to calculate $fact(4)$, we need to calculate $fact(3)$, and so on.

Recursive function definitions always contain a number of non-recursive base cases and a number of recursive cases. In the case of factorial, we have one of each. The base case is when $n = 0$ and the recursive case is when $n > 0$.

One can actually think of the natural numbers themselves as recursive (in fact, if you ask set theorists about this, they'll say this *is* how it is). That is, there is a zero element and then for every element, it has a successor. That is $1 = succ(0)$, $2 = succ(1)$, \dots , $573 = succ(573)$, \dots and so on forever. We can actually implement this system of natural numbers in Haskell:

```
data Nat = Zero | Succ Nat
```

This is a recursive type definition. Here, we represent one as `Succ Zero` and three as `Succ (Succ (Succ Zero))`. One thing we might want to do is be able to convert back and forth between `Nats` and `Ints`. Clearly, we can write a base case as:

```
natToInt Zero = 0
```

In order to write the recursive case, we realize that we're going to have something of the form `Succ n`. We can make the assumption that we'll be able to take `n` and produce an `Int`. Assuming we can do this, all we need to do is add one to this result. This gives rise to our recursive case:

```
natToInt (Succ n) = natToInt n + 1
```

There is a close connection between recursion and mathematical induction. Induction is a proof technique which typically breaks problems down into base cases and “inductive” cases, very analogous to our analysis of recursion.

Let’s say we want to prove the statement $n! \geq n$ for all $n \geq 0$. First we formulate a base case: namely, we wish to prove the statement when $n = 0$. When $n = 0$, $n! = 1$ by definition. Since $n! = 1 > 0 = n$, we get that $0! \geq 0$ as desired.

Now, suppose that $n > 0$. Then $n = k + 1$ for some value k . We now invoke the *inductive hypothesis* and claim that the statement holds for $n = k$. That is, we assume that $k! \geq k$. Now, we use k to formulate the statement for our value of n . That is, $n! \geq n$ if and only iff $(k + 1)! \geq (k + 1)$. We now apply the definition of factorial and get $(k + 1)! = (k + 1) * k!$. Now, we know $k! \geq k$, so $(k + 1) * k! \geq k + 1$ if and only if $k + 1 \geq 1$. But we know that $k \geq 0$, which means $k + 1 \geq 1$. Thus it is proven.

It may seem a bit counter-intuitive that we are assuming that the claim is true for k in our proof that it is true for n . You can think of it like this: we’ve proved the statement for the case when $n = 0$. Now, we know it’s true for $n = 0$ so using this we use our inductive argument to show that it’s true for $n = 1$. Now, we know that it is true for $n = 1$ so we reuse our inductive argument to show that it’s true for $n = 2$. We can continue this argument as long as we want and then see that it’s true for all n .

It’s much like pushing down dominoes. You know that when you push down the first domino, it’s going to knock over the second one. This, in turn will knock over the third, and so on. The base case is like pushing down the first domino, and the inductive case is like showing that pushing down domino k will cause the $k + 1$ st domino to fall.

In fact, we can use induction to prove that our `natToInt` function does the right thing. First we prove the base case: does `natToInt Zero` evaluate to 0? Yes, obviously it does. Now, we can assume that `natToInt n` evaluates to the correct value (this is the inductive hypothesis) and ask whether `natToInt (Succ n)` produces the correct value. Again, it is obvious that it does, by simply looking at the definition.

Let’s consider a more complex example: addition of Nats. We can write this concisely as:

```
addNat Zero m = m
addNat (Succ n) m = addNat n (Succ m)
```

Now, let’s prove that this does the correct thing. First, as the base case, suppose the first argument is `Zero`. We know that $0 + m = m$ regardless of what m is; thus in the base case the algorithm does the correct thing. Now, suppose that `addNat n m` does the correct thing for all m and we want to show that `addNat (Succ n) m` does the correct thing. We know that $(n + 1) + m = n + (m + 1)$ and thus since `addNat n (Succ m)` does the correct thing (by the inductive hypothesis), our program is correct.

Appendix C

“Real World” Examples

In this section we have attempted to include the code for some real world applications written in Haskell. We have attempted to comment the code sufficiently that it can be understood with only the knowledge of Haskell which comes from this tutorial. The code for all of these programs is available off the tutorial website <http://www.isi.edu/~hdaume/htut/> (.)

Appendix D

Solutions To Exercises

Solution 1

It binds more tightly; actually, function application binds more tightly than anything else. To see this, we can do something like:

```
Prelude> sqrt 3 * 3  
5.19615
```

If multiplication bound more tightly, the result would have been 3.

Solution 2

Solution: `snd (fst ((1, 'a'), "foo"))`. This is because first we want to take the first half the the tuple: `(1, 'a')` and then out of this we want to take the second half, yielding just `'a'`.

If you tried `fst (snd ((1, 'a'), "foo"))` you will have gotten a type error. This is because the application of `snd` will leave you with `fst "foo"`. However, the string "foo" isn't a tuple, so you cannot apply `fst` to it.

Solution 3

Solution: `map isLower "aBCde"`

Solution 4

Solution: `length (filter isLower "aBCde")`

Solution 5

`foldr max 0 [5,10,2,8,1]`. You could also use `foldl`. The `foldr` case is easier to explain: we replace each cons with an application of `max` and the empty list with 0. Thus, the inner-most application will take the maximum of 0 and the last element of the list (if it exists). Then, the next-most inner application will return the maximum of whatever was the maximum before and the second-to-last element. This will continue on, carrying to current maximum all the way back to the beginning of the list.

In the foldl case, we can think of this as looking at each element in the list in order. We start off our “state” with 0. We pull off the first element and check to see if it’s bigger than our current state. If it is, we replace our current state with that number and then continue. This happens for each element and thus eventually returns the maximal element.

Solution 6

```
fst (head (tail [(5,'b'),(1,'c'),(6,'a')]))
```

Solution 7

We can define a fibonacci function as:

```
fib 1 = 1
fib 2 = 1
fib n = fib (n-1) + fib (n-2)
```

We could also write it using explicit **if** statements, like:

```
fib n =
  if n == 1 || n == 2
  then 1
  else fib (n-1) + fib (n-2)
```

Either is acceptable, but the first is perhaps more natural in Haskell.

Solution 8

We can define:

$$a * b = \begin{cases} a & b = 1 \\ a + a * (b - 1) & \text{otherwise} \end{cases}$$

And then type out code:

```
mult a 1 = a
mult a b = a + mult a (b-1)
```

Note that it doesn’t matter which of a and b we do the recursion on. We could just as well have defined it as:

```
mult 1 b = b
mult a b = b + mult (a-1) b
```

Solution 9

We can define `my_map` as:

```
my_map f [] = []
my_map f (x:xs) = f x : my_map f xs
```

Recall that the `my_map` function is supposed to apply a function `f` to every element in the list. In the case that the list is empty, there are no elements to apply the function to, so we just return the empty list.

In the case that the list is non-empty, it is an element `x` followed by a list `xs`. Assuming we've already properly applied `my_map` to `xs`, then all we're left to do is apply `f` to `x` and then stick the results together. This is exactly what the second line does.

Solution 10

The code below appears in `Numbers.hs`. The only tricky parts are the recursive calls in `getNums` and `showFactorials`.

```
module Main
  where

import IO

main = do
  nums <- getNums
  putStrLn ("The sum is " ++ show (sum nums))
  putStrLn ("The product is " ++ show (product nums))
  showFactorials nums

getNums = do
  putStrLn "Give me a number (or 0 to stop):"
  num <- getLine
  if read num == 0
    then return []
    else do rest <- getNums
           return ((read num :: Int):rest)

showFactorials [] = return ()
showFactorials (x:xs) = do
  putStrLn (show x ++ " factorial is " ++
            show (factorial x))
  showFactorials xs

factorial 1 = 1
factorial n = n * factorial (n-1)
```

The idea for `getNums` is just as spelled out in the hint. For `showFactorials`, we consider first the recursive call. Suppose we have a list of numbers, the first of

which is `x`. First we print out the string showing the factorial. Then we print out the rest, hence the recursive call. But what should we do in the case of the empty list? Clearly we are done, so we don't need to do anything at all, so we simply return `()`.

Note that this must be `return ()` instead of just `()` because if we simply wrote `showFactorials [] = ()` then this wouldn't be an IO action, as it needs to be. For more clarification on this, you should probably just keep reading the tutorial.

Solution 11

String or `[Char]`

- 1.
2. type error: lists are homogenous
3. `Num a => (a, Char)`
4. `Int`
5. type error: cannot add values of different types

Solution 12

The types:

1. `(a, b) -> b`
2. `[a] -> a`
3. `[a] -> Bool`
4. `[a] -> a`
5. `[[a]] -> a`

Solution 13

The types:

1. `a -> [a]`. This function takes an element and returns the list containing only that element.
2. `a -> b -> b -> (a, [b])`. The second and third argument must be of the same type, since they go into the same list. The first element can be of any type.
3. `Num a => a -> a`. Since we apply `(+)` to `a`, it must be an instance of `Num`.
4. `a -> String`. This ignores the first argument, so it can be any type.
5. `(Char -> a) -> a`. In this expression, `x` must be a function which takes a `Char` as an argument. We don't know anything about what it produces, though, so we call it `a`.

6. Type error. Here, we assume x has type a . But x is applied to itself, so it must have type $b \rightarrow c$. But then it must have type $(b \rightarrow c) \rightarrow c$, but then it must have type $((b \rightarrow c) \rightarrow c) \rightarrow c$ and so on, leading to an infinite type.
7. **Num** $a \Rightarrow a \rightarrow a$. Again, since we apply $(+)$, this must be an instance of **Num**.

Solution 14

The definitions will be something like:

```
data Triple a b c = Triple a b c

tripleFst (Triple x y z) = x
tripleSnd (Triple x y z) = y
tripleThr (Triple x y z) = z
```

Solution 15

The code, with type signatures, is:

```
data Quadruple a b = Quadruple a a b b

firstTwo :: Quadruple a b -> [a]
firstTwo (Quadruple x y z t) = [x,y]

lastTwo :: Quadruple a b -> [b]
lastTwo (Quadruple x y z t) = [z,t]
```

We note here that there are only two type variables, a and b associated with `Quadruple`.

Solution 16

The code:

```
data Tuple a b c d e = One a
                    | Two a b
                    | Three a b c
                    | Four a b c d

tuple1 (One a      ) = Just a
tuple1 (Two a b    ) = Just a
tuple1 (Three a b c) = Just a
tuple1 (Four a b c d) = Just a

tuple2 (One a      ) = Nothing
tuple2 (Two a b    ) = Just b
tuple2 (Three a b c) = Just b
```

```

tuple2 (Four  a b c d) = Just b

tuple3 (One   a       ) = Nothing
tuple3 (Two  a b     ) = Nothing
tuple3 (Three a b c   ) = Just c
tuple3 (Four  a b c d) = Just c

tuple4 (One   a       ) = Nothing
tuple4 (Two  a b     ) = Nothing
tuple4 (Three a b c   ) = Nothing
tuple4 (Four  a b c d) = Just d

```

Solution 17

The code:

```

fromTuple (One   a       ) = Left  (Left  a       )
fromTuple (Two  a b     ) = Left  (Right (a,b)    )
fromTuple (Three a b c   ) = Right (Left  (a,b,c)  )
fromTuple (Four  a b c d) = Right (Right (a,b,c,d))

```

Here, we use embedded Eithers to represent the fact that there are four (instead of two) options.

Solution 18

The code:

```

listHead (Cons x xs) = x
listTail (Cons x xs) = xs

listFoldl f y Nil = y
listFoldl f y (Cons x xs) = listFoldl f (f y x) xs

listFoldr f y Nil = y
listFoldr f y (Cons x xs) = f x (foldr f z xs)

```

Solution 19

The code:

```

elements (Leaf x) = [x]
elements (Branch lhs x rhs) =
  elements lhs ++ [x] ++ elements rhs

```

Solution 20

The code:

```

foldTree :: (a -> b -> b) -> b -> BinaryTree a -> b
foldTree f z (Leaf x) = f x z
foldTree f z (Branch lhs x rhs) =
    foldTree f (f x (foldTree f z rhs)) lhs

elements2 = foldTree (:) []

```

or:

```

elements2 tree = foldTree (\a b -> a:b) [] tree

```

The first `elements2` is simply a more compact version of the second.

Solution 21

Using **if**, we get something like:

```

main = do
    putStrLn "Please enter your name:"
    name <- getLine
    if name == "Simon" || name == "John" || name == "Phil"
    then putStrLn "Haskell is great!"
    else if name == "Koen"
    then putStrLn "Debugging Haskell is fun!"
    else putStrLn "I don't know who you are."

```

Note that we don't need to repeat the **dos** inside the **ifs**, since these are only one action commands.

We could also be a bit smarter and use the `elem` command which is built in to the Prelude:

```

main = do
    putStrLn "Please enter your name:"
    name <- getLine
    if name `elem` ["Simon", "John", "Phil"]
    then putStrLn "Haskell is great!"
    else if name == "Koen"
    then putStrLn "Debugging Haskell is fun!"
    else putStrLn "I don't know who you are."

```

Of course, we needn't put all the `putStrLn`s inside the **if** statements. We could instead write:

```

main = do
    putStrLn "Please enter your name:"

```

```

name <- getLine
putStrLn
  (if name `elem` ["Simon", "John", "Phil"]
   then "Haskell is great!"
   else if name == "Koen"
        then "Debugging Haskell is fun!"
        else "I don't know who you are.")

```

Using `case`, we get something like:

```

main = do
  putStrLn "Please enter your name:"
  name <- getLine
  case name of
    "Simon" -> putStrLn "Haskell is great!"
    "John"  -> putStrLn "Haskell is great!"
    "Phil"  -> putStrLn "Haskell is great!"
    "Koen"  -> putStrLn "Debugging Haskell is fun!"
    _      -> putStrLn "I don't know who you are."

```

Which, in this case, is actually not much cleaner.

Solution 22

The code might look something like:

```

module DoFile where

import IO

main = do
  putStrLn "Do you want to [read] a file, ...?"
  cmd <- getLine
  case cmd of
    "quit" -> return ()
    "read" -> do doRead; main
    "write" -> do doWrite; main
    _      -> do putStrLn
                  ("I don't understand the command "
                   ++ cmd ++ ".")
              main

doRead = do
  putStrLn "Enter a file name to read:"
  fn <- getLine
  bracket (openFile fn ReadMode) hClose

```

```

        (\h -> do txt <- hGetContents h
                  putStrLn txt)

doWrite = do
  putStrLn "Enter a file name to write:"
  fn <- getLine
  bracket (openFile fn WriteMode) hClose
    (\h -> do putStrLn
              "Enter text (...):"
              writeLoop h)

writeLoop h = do
  l <- getLine
  if l == "."
  then return ()
  else do hPutStrLn h l
          writeLoop h

```

The only interesting things here are the calls to `bracket`, which ensure that the program lives on, regardless of whether there's a failure or not; and the `writeLoop` function. Note that we need to pass the handle returned by `openFile` (through `bracket` to this function, so it knows where to write the input to).

Solution 23

Function `func3` cannot be converted into point-free style. The others look something like:

```

func1 x = map (*x)

func2 f g = filter f . map g

func4 = map (+2) . filter ('elem' [1..10]) . (5:)

func5 = foldr (uncurry $ flip f) 0

```

\end{Soln}
\begin{Soln}{24}

We can start out with a recursive definition:

```

\begin{code}
and [] = True
and (x:xs) = x && and xs

```

From here, we can clearly rewrite this as:

```
and = foldr (&&) True
```

Solution 25

We can write this recursively as:

```
concatMap f [] = []
concatMap f (x:xs) = f x ++ concatMap f xs
```

This hints that we can write this as:

```
concatMap f = foldr (\a b -> f a ++ b) []
```

Now, we can do point elimination to get:

```
foldr (\a b -> f a ++ b) []
==> foldr (\a b -> (++) (f a) b) []
==> foldr (\a -> (++) (f a)) []
==> foldr (\a -> ((++) . f) a) []
==> foldr ((++) . f) []
```

Solution 26

The first law is: $\text{return } a \gg= f \equiv f \cdot a$. In the case of `Maybe`, we get:

```
return a >>= f
==> Just a >>= \x -> f x
==> (\x -> f x) a
==> f a
```

The second law is: $f \gg= \text{return} \equiv f$. Here, we get:

```
f >>= return
==> f >>= \x -> return x
==> f >>= \x -> Just x
```

At this point, there are two cases depending on whether `f` is `Nothing` or not. In the first case, we get:

```
==> Nothing >>= \x -> Just x
==> Nothing
==> f
```

In the second case, `f` is `Just a`. Then, we get:

```
==> Just a >>= \x -> Just x
==> (\x -> Just x) a
==> Just a
==> f
```

And the second law is shown. The third law states: $f \gg= (\backslash x \rightarrow g x \gg= h) \equiv (f \gg= g) \gg= h$.

If f is `Nothing`, then the left-hand-side clearly reduces to `Nothing`. The right-hand-side reduces to `Nothing >>= h` which in turn reduces to `Nothing`, so they are the same.

Suppose f is `Just a`. Then the LHS reduces to $g a \gg= h$ and the RHS reduces to $(Just a \gg= \backslash x \rightarrow g x) \gg= h$ which in turn reduces to $g a \gg= h$, so these two are the same.

Solution 27

The idea is that we wish to use the `Left` constructor to represent errors on the `Right` constructor to represent successes. This leads to an instance declaration like:

```
instance Monad (Either String) where
  return x      = Right x
  Left s >>= _  = Left s
  Right x >>= f = f x
  fail s       = Left s
```

If we try to use this monad to do search, we get:

```
Monads> searchAll gr 0 3 :: Either String [Int]
Right [0,1,3]
Monads> searchAll gr 3 0 :: Either String [Int]
Left "no path"
```

which is exactly what we want.

Solution 28

The order to `mplus` essentially determines the search order. When the recursive call to `searchAll2` comes first, we are doing depth-first search. When the recursive call to `search'` comes first, we are doing breadth-first search. Thus, using the list monad, we expect the solutions to come in the other order:

```
MPlus> searchAll3 gr 0 3 :: [[Int]]
[[0,2,3],[0,1,3]]
```

Just as we expected.

Solution 29

This is a very difficult problem; if you found that you were stuck immediately, please just read as much of this solution as you need to try it yourself.

First, we need to define a list transformer monad. This looks like:


```
newtype ListT m e = ListT { unListT :: m [e] }
```

The `ListT` constructor simply wraps a monadic action (in monad `m`) which returns a list.

We now need to make this a monad:

```
instance Monad m => Monad (ListT m) where
  return x = ListT (return [x])
  fail s = ListT (return [] )
  ListT m >>= k = ListT $ do
    l <- m
    l' <- mapM (unListT . k) l
    return (concat l')
```

Here, success is designated by a monadic action which returns a singleton list. Failure (like in the standard list monad) is represented by an empty list: of course, it's actually an empty list returned from the enclosed monad. Binding happens essentially by running the action which will result in a list `l`. This has type `[e]`. We now need to apply `k` to each of these elements (which will result in something of type `ListT m [e2]`). We need to get rid of the `ListTs` around this (by using `unListT`) and then concatenate them to make a single list.

Now, we need to make it an instance of **MonadPlus**

```
instance Monad m => MonadPlus (ListT m) where
  mzero = ListT (return [])
  ListT m1 `mplus` ListT m2 = ListT $ do
    l1 <- m1
    l2 <- m2
    return (l1 ++ l2)
```

Here, the zero element is a monadic action which returns an empty list. Addition is done by executing both actions and then concatenating the results.

Finally, we need to make it an instance of **MonadTrans**:

```
instance MonadTrans ListT where
  lift x = ListT (do a <- x; return [a])
```

Lifting an action into `ListT` simply involves running it and getting the value (in this case, `a`) out and then returning the singleton list.

Once we have all this together, writing `searchAll6` is fairly straightforward:

```
searchAll6 g@(Graph v1 e1) src dst
  | src == dst = do
    lift $ putStrLn $
```

```

    "Exploring " ++ show src ++ " -> " ++ show dst
  return [src]
| otherwise = do
  lift $ putStrLn $
    "Exploring " ++ show src ++ " -> " ++ show dst
  search' el
where
  search' [] = mzero
  search' ((u,v,_) : es)
    | src == u =
      (do path <- searchAll6 g v dst
         return (u:path)) `mplus`
      search' es
    | otherwise = search' es

```

The only change (besides changing the recursive call to call `searchAll6` instead of `searchAll2`) here is that we call `putStrLn` with appropriate arguments, lifted into the monad.

If we look at the type of `searchAll6`, we see that the result (i.e., after applying a graph and two ints) has type `MonadTrans t, MonadPlus (t IO) => t IO [Int]`. In theory, we could use this with any appropriate monad transformer; in our case, we want to use `ListT`. Thus, we can run this by:

```

MTrans> unListT (searchAll6 gr 0 3)
Exploring 0 -> 3
Exploring 1 -> 3
Exploring 3 -> 3
Exploring 2 -> 3
Exploring 3 -> 3
MTrans> it
[[0,1,3],[0,2,3]]

```

This is precisely what we were looking for.

Solution 30

This exercise is actually simpler than the previous one. All we need to do is incorporate the calls to `putT` and `getT` into `searchAll6` and add an extra `lift` to the IO calls. This extra `lift` is required because now we're stacking two transformers on top of IO instead of just one.

```

searchAll7 g@(Graph vl el) src dst
  | src == dst = do
    lift $ lift $ putStrLn $
      "Exploring " ++ show src ++ " -> " ++ show dst
    visited <- getT

```

```

putT (src:visited)
return [src]
| otherwise = do
  lift $ lift $ putStrLn $
    "Exploring " ++ show src ++ " -> " ++ show dst
  visited <- getT
  putT (src:visited)
  if src `elem` visited
  then mzero
  else search' el
where
  search' [] = mzero
  search' ((u,v,_):es)
    | src == u =
      (do path <- searchAll7 g v dst
         return (u:path)) `mplus`
      search' es
    | otherwise = search' es

```

The type of this has grown significantly. After applying the graph and two ints, this has type `Monad (t IO), MonadTrans t, MonadPlus (StateT [Int] (t IO)) => StateT [Int] (t IO) [Int]`.

Essentially this means that we've got something that's a state transformer wrapped on top of some other arbitrary transformer (`t`) which itself sits on top of `IO`. In our case, `t` is going to be `ListT`. Thus, we run this beast by saying:

```

MTrans> unListT (evalStateT (searchAll7 gr4 0 3) [])
Exploring 0 -> 3
Exploring 1 -> 3
Exploring 3 -> 3
Exploring 0 -> 3
Exploring 2 -> 3
Exploring 3 -> 3
MTrans> it
[[0,1,3],[0,2,3]]

```

And it works, even on `gr4`.

Solution 31

First we write a function `spaces` which will parse out whitespaces:

```

spaces :: Parser ()
spaces = many (matchChar isSpace) >> return ()

```

Now, using this, we simply sprinkle calls to `spaces` through `intList` to get `intListSpace`:

```

intListSpace :: Parser [Int]
intListSpace = do
  char '['
  spaces
  intList' `mplus` (char ']' >> return [])
  where intList' = do
        i <- int
        spaces
        r <- (char ',' >> spaces >> intList')
        `mplus`
        (char ']' >> return [])
        return (i:r)

```

We can test that this works:

```

Parsing> runParser intListSpace "[1 ,2 , 4 \n\n ,5\n]"
Right ("",[1,2,4,5])
Parsing> runParser intListSpace "[1 ,2 , 4 \n\n ,a\n]"
Left "expecting char, got 'a'"

```

Solution 32

We do this by replacing the state functions with push and pop functions as follows:

```

parseValueLet2 :: CharParser (FiniteMap Char [Int]) Int
parseValueLet2 = choice
  [ int
  , do string "let "
    c <- letter
    char '='
    e <- parseValueLet2
    string " in "
    pushBinding c e
    v <- parseValueLet2
    popBinding c
    return v
  , do c <- letter
    fm <- getState
    case lookupFM fm c of
      Nothing    -> unexpected ("variable " ++
                                show c ++
                                " unbound")
      Just (i:_) -> return i
  , between (char '(') (char ')') $ do
    e1 <- parseValueLet2

```

```

op <- oneOf "+*"
e2 <- parseValueLet2
case op of
  '+' -> return (e1 + e2)
  '*' -> return (e1 * e2)
]
where
  pushBinding c v = do
    fm <- getState
    case lookupFM fm c of
      Nothing -> setState (addToFM fm c [v])
      Just l -> setState (addToFM fm c (v:l))
  popBinding c = do
    fm <- getState
    case lookupFM fm c of
      Just [_] -> setState (delFromFM fm c)
      Just (_:l) -> setState (addToFM fm c l)

```

The primary difference here is that instead of calling `updateState`, we use two local functions, `pushBinding` and `popBinding`. The `pushBinding` function takes a variable name and a value and adds the value onto the head of the list pointed to in the state `FiniteMap`. The `popBinding` function looks at the value and if there is only one element on the stack, it completely removes the stack from the `FiniteMap`; otherwise it just removes the first element. This means that if something is in the `FiniteMap`, the stack is never empty.

This enables us to modify only slightly the usage case; this time, we simply take the top element off the stack when we need to inspect the value of a variable.

We can test that this works:

```

ParsecI> runParser parseValueLet2 emptyFM "stdin"
          "((let x=2 in 3+4)*x)"
Left "stdin" (line 1, column 20):
unexpected variable 'x' unbound

```

Index

- `()`, 87
- `*`, 32
- `+`, 32
- `++`, 38
- `-`, 32
- `--`, 52
- `..`, 50
- `...`, 140
- `/`, 32
- `//`, 144
- `;`, 36
- `::`, 65
- `==`, 65
- `[]`, 36
- `^`, 32
- `$\Gamma E30F$` , 72
- `λ` , 72
- `{--}`, 52
- `->`, 122
- as**, 111
- derive**, 134
- do**, 57
- hiding**, 111
- import**, 110
- qualified**, 110

- `accumArray`, 143
- actions, 95–100
- arithmetic, 32–34
- array, 143
- arrays, 143–145
- assocs, 144

- Bird scripts, 113
- boolean, 66
- bounds, 144
- bracket, 101

- brackets, 36
- buffering, 58

- comments, 52–53
- common sub-expression, 117
- Compilers, *see* GHC,NHC,Interpreters
- concatenate, 38
- cons, 36
- constructors, 81
- continuation passing style, 87–90
- CPS, *see* continuation passing style

- destructive update, 31
- do notation, 173–176
- do** notation, 57
- drop, 140

- Editors, 27–28
 - Emacs, 27
- `elems`, 144
- Enum, 131–132
- enumerated types, 84
- `enumFromThenTo`, 141
- `enumFromTo`, 141
- Eq, 127–129
- equality, 65
- equals, 70
- evaluation order, 32
- exceptions, 182
- exports, 107–110
- expressions, 32
- extensions
 - enabling
 - in GHC, 25
 - in Hugs, 23

- fallthrough, 123
- false, *see* boolean

- files, 42–44
- filter, 38, 138
- FiniteMap, 145–147
- foldl, 38, 139
- foldr, 38, 138
- fst, 35
- functional, 3
- functions, 44–52
 - anonymous, *see* lambda
 - as arguments, 78–80
 - associative, 40
 - composition, 50
 - type, 71–80
- getChar, 101
- getContents, 101
- getLine, 58, 101
- GHC, 19, 23–26
- guards, 125–127
- Haskell 98, 8
- Haskell Bookshelf, 7
- hClose, 101
- head, 37, 74
- hGetChar, 101
- hGetContents, 101
- hGetLin, 101
- hIsEOF, 101
- hPutChar, 101
- hPutStr, 101
- hPutStrLn, 101
- Hugs, 19–23
- ICFP, 9
- immutable, 31
- imports, 110–111
- indices, 144
- induction, 54
- infix, 75, 115–116
- input, *see* IO
- instance, 67
- interactive programs, 56–63
- Interpreters, *see* GHC,Hugs
- IO, 91–106
 - library, 100–103
- isLower, 38
- lambda, 72
- lambda calculus, 72
- LaTeX scripts, 114
- layout, 147
- lazy, 2, 29
- length, 37, 79
- let
 - and **do**, 175
- list
 - comprehensions, 143
- listArray, 143
- lists, 35–42, 137–143
 - comprehensions, 140
 - cons, 36
 - empty, 36
- literate style, 112–114
- local bindings, 117
- loop, 53
- map, 38, 138
- maps, 145–147
- Maybe, 85–87
- modules, 107–114
 - hierarchical, 111–112
- monads, 94
 - and **do**, 173–176
 - definition of, 170–173
 - laws, 171
 - plus, 181
 - state, 176–181
 - transformer, 181
- mutable, *see* immutable
- named fields, 135–137
- NHC, 19, 26–27
- null, 74
- Num, 132
- numeric types, 70
- openFile, 101
- operator precedence, 33
- Ord, 130–131
- output, *see* IO
- pairs, *see* tuples
- parentheses, 33

- pattern matching, 82, 119–125
- pure, 2
- putChar, 101
- putStr, 101
- putStrLn, 44, 101

- random numbers, 59
- randomRIO, 59
- Read, 132–133
- read, 38
- readFile, 101
- recursion, 53–56
- referential transparency, 31
- robot, 9

- sections, 115–116
- shadowing, 117
- Show, 129–130
- show, 38, 71
- snd, 35
- sqrt, 32
- standard, 8
- state, 2
- strict, 2, 152–155
- strings, 37–38
 - converting from/to, 38

- tail, 37, 74
- take, 140
- toUpper, 38
- true, *see* boolean
- tuples, 34–35
- type, 64, 168
 - checking, 64
 - classes, 69–71, 155–162
 - instances, 127–134, 162–165
 - datatypes, 80–87, 134–137, 152–155
 - strict, 152–155
 - default, 168
 - errors, 66
 - explicit declarations, 77–78
 - hierarchy, 168
 - higher-order, 73–75
 - inference, 64
 - IO, 75–77
 - kinds, 165–168
 - newtype, 149–152
 - polymorphic, 67–69
 - signatures, 77
 - synonyms, 101, 148–149

- Unit, 87
- unzip, 139
- user input, *see* interactive programs

- wildcard, 122
- writeFile, 101

- zip, 139