

The Haskell School of Music

by

Paul Hudak

**Yale University
Department of Computer Science**

Copyright © Paul Hudak

September 2008

All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the author.

Contents

Preface	iv
1 Computation by Calculation	1
1.1 Computation by Calculation in Haskell	2
1.2 Expressions, Values, and Types	6
1.3 Function Types and Type Signatures	8
1.4 Abstraction, Abstraction, Abstraction	9
1.4.1 Naming	10
1.4.2 Functional Abstraction	12
1.4.3 Data Abstraction	13
1.5 Code Reuse and Modularity	16
1.6 Beware of Programming with Numbers	17
2 Simple Music	20
2.1 Preliminaries	20
2.2 Notes and Music	22
2.3 Convenient Auxiliary Functions	24
2.4 Absolute Pitches	29
3 Polymorphic and Higher-Order Functions	32
3.1 Polymorphic Types	32
3.2 Abstraction Over Recursive Definitions	34
3.2.1 Map is Polymorphic	35
3.2.2 Using map	36
3.3 Append	37
3.4 Fold	39
3.4.1 Haskell's Folds	41
3.4.2 Why Two Folds?	41
3.4.3 Fold for Non-empty Lists	43

3.5	A Final Example: Reverse	43
3.6	Errors	45
4	More About Higher-Order Functions	48
4.1	Currying	48
4.2	Sections	51
4.3	Anonymous Functions	52
4.4	Function Composition	54
5	More Music	57
5.1	Delay and Repeat	57
5.2	Inversion and Retrograde	58
5.3	Polyrhythms	58
5.4	Symbolic Meter Changes	60
5.5	Computing Duration	60
5.6	Super-retrograde	60
5.7	Truncating Parallel Composition	61
5.8	Trills	62
5.9	Percussion	63
5.10	A Map for Music	65
5.11	A Fold for Music	66
5.12	Crazy Recursion	67
6	Interpretation and Performance	69
6.1	Abstract Performance	69
6.2	Players	73
6.2.1	Examples of Player Construction	76
7	Self-Similar Music	81
7.1	Self-Similar Melody	81
7.2	Self-Similar Harmony	84
7.3	Other Self-Similar Structures	85
8	Proof by Induction	87
8.1	Induction and Recursion	87
8.2	Examples of List Induction	88
8.3	Proving Function Equivalences	90
8.4	Useful Properties on Lists	93
8.4.1	Function Strictness	93
8.5	Induction on Other Data Types	98

8.5.1	A More Efficient Exponentiation Function	100
9	An Algebra of Music	104
9.1	Musical Equivalence	104
9.2	Some Simple Axioms	106
9.3	The Axiom Set	109
9.4	Soundness and Completeness	110
A	A Tour of the PreludeList Module	111
A.1	The PreludeList Module	112
A.2	Simple List Selector Functions	112
A.3	Index-Based Selector Functions	113
A.4	Predicate-Based Selector Functions	114
A.5	Fold-like Functions	115
A.6	List Generators	117
A.7	String-Based Functions	117
A.8	Boolean List Functions	118
A.9	List Membership Functions	119
A.10	Arithmetic on Lists	119
A.11	List Combining Functions	120
B	Built-in Types Are Not Special	122
C	Pattern-Matching Details	125

Preface

In 2000 I wrote a book called *The Haskell School of Expression – Learning Functional Programming through Multimedia*. In that book I used graphics, animation, music, and robotics as a way to motivate learning how to program, and specifically how to learn *functional programming* using Haskell, a purely functional programming language. Haskell is quite a bit different from conventional imperative or objected-oriented languages such as C, C++, Java, C#, and so on. It takes a different mind-set to program in such a language, and appeals to the mathematically inclined and to those who seek purity and elegance in their programs. Although Haskell was designed almost twenty years ago, it has only recently begun to catch on, not just because of its purity and elegance, but because with it you can solve real-world problems quickly and efficiently, and with great economy of code.

I have also had a long, informal, yet passionate interest in music, being an amateur jazz pianist and having played in several bands over the years. About ten years ago, in an effort to combine work with play, I wrote a Haskell library called *Haskore* for expressing high-level computer music concepts in a purely functional way. Indeed, three of the chapters in *The Haskell School of Expression* summarize the basic ideas of this work. Thus, when I recently became responsible for the Music Track in the new *Computing and the Arts* major at Yale, and became responsible for teaching not one, but two computer music courses in the new curriculum, it was natural to base the course material on Haskell. This current book is essentially a rewrite of *The Haskell School of Expression* with a focus entirely on music, based on and improving upon the ideas in *Haskore*.

Haskell was named after the logician Haskell B. Curry who, along with Alonzo Church, established the theoretical foundations of functional programming in the 1940's, when digital computers were mostly just a gleam in researchers' eyes. A curious historical fact is that Haskell Curry's father, Samuel Silas Curry, helped found and direct a school in Boston called the *School of Expression*. (This school eventually evolved into what is now

Curry College.) Since pure functional programming is centered around the notion of an *expression*, I thought that *The Haskell School of Expression* would be a good title for my first book. And it was thus quite natural to choose *The Haskell School of Music* for my second!

How To Read This Book

As mentioned earlier, there is a certain mind-set, a certain viewpoint of the world, and a certain approach to problem solving that collectively work best when programming in Haskell (this is true for any new programming paradigm). If you teach only Haskell language details to a C programmer, she is likely to write ugly, incomprehensible functional programs. But if you teach her how to think differently, how to see problems in a different light, functional solutions will come easily, and elegant Haskell programs will result. As Samuel Silas Curry once said:

All expression comes *from within outward*, from the center to the surface, from a hidden source to outward manifestation. The study of expression as a natural process brings you into contact with cause and makes you feel the source of reality.

What is especially beautiful about this quote is that music is a kind of expression, although Curry was more likely talking about speech. In addition, as has been noted by many, music has many ties to mathematics. So for me, combining the elegant mathematical nature of Haskell with that of music is as natural as singing a nursery tune.

Using a high-level language to express musical ideas is, of course, not new. But Haskell is unique in its insistence on purity (no side effects), and this alone makes it particularly suitable for expressing musical ideas. By focusing on *what* a musical entity is rather than on *how* we should create it, we allow musical ideas to take their natural form as Haskell expressions. Haskell's many abstraction mechanisms allow us to write musical programs that are elegant, concise, yet powerful. We will consistently attempt to let the music express itself as naturally as possible, without encoding it in terms of irrelevant language details.

Of course, my ultimate goal is to teach computer music concepts. But along the way you will also learn Haskell. There is no limit to what one might wish to do with computer music, and therefore the better you are at programming, the more success you will have. This is why I think that many languages designed specifically for computer music—although fun to

work with, easy to use, and cute in concept—will ultimately be too limited in expressiveness.

My general approach to introducing computer music concepts is to first provide an intuitive explanation, then a mathematically rigorous definition, and finally fully executable Haskell code. In the process I will introduce Haskell features as they are needed, rather than all at once. I believe that this interleaving of concepts and applications makes the material easier to digest.

Another characteristic of my approach is that I won't hide any details—I want Haskore to be as transparent as possible! There are no magical built-in operations, no special computer music commands or values. This works out well for several reasons. First, there is in fact nothing ugly or difficult to hide—so why hide anything at all? Second, by reading the code, you will better and more quickly understand Haskell. Finally, by stepping through the design process with me, you may decide that you prefer a different approach—there is, after all, no One True Way to express computer music ideas. Indeed, I expect that this process will position you well to write rich, creative musical applications on your own.

I encourage the seasoned programmer having experience only with conventional imperative and/or object-oriented languages to read this text with an open mind. Many things will be different, and will likely feel awkward. There will be a tendency to rely on old habits when writing new programs, and to ignore suggestions about how to approach things differently. If you can manage to resist those tendencies I am confident that you will have an enjoyable learning experience. Many of those who succeed in this process find that many ideas about functional programming can be applied to imperative and object-oriented languages as well, and that their imperative coding style changes for the better.

I also ask the experienced programmer to be patient while in the earlier chapters I explain things like “syntax,” “operator precedence,” etc., since it is my goal that this text should be readable by someone having only modest prior programming experience. With patience the more advanced ideas will appear soon enough.

If you are a novice programmer, I suggest taking your time with the book; work through the exercises, and don't rush things. If, however, you don't fully grasp an idea, feel free to move on, but try to re-read difficult material at a later time when you have seen more examples of the concepts in action. For the most part this is a “show by example” textbook, and you should try to execute as many of the programs in this text as you can, as well as every program that you write. Learn-by-doing is the corollary to

show-by-example.

Haskell Implementations

There are several good implementations of Haskell, all available free on the Internet through the Haskell Home Page at <http://haskell.org>. One that I especially recommend is *GHC*, an easy-to-use and easy-to-install Haskell compiler and interpreter (see <http://haskell.org/ghc>). GHC runs on a variety of platforms, including PC's (Windows XP and Vista), various flavors of Unix (Linux, FreeBSD, etc.), and Mac OS X. Any text editor can be used to create the source files, but I prefer to use emacs (see <http://www.gnu.org/software/emacs>), along with its Haskell mode (see <http://www.haskell.org/haskell-mode>). All of the source code from this textbook can be found at <http://plucky.cs.yale.edu/cs431>. Feel free to email me at paul.hudak@yale.edu with any comments, suggestions, or questions.

Happy Haskell Music Hacking!

Paul Hudak
New Haven
September 2008

Chapter 1

Computation by Calculation

Programming, in its broadest sense, is *problem solving*. It begins when we look out into the world and see problems that we want to solve, problems that we think can and should be solved using a digital computer. Understanding the problem well is the first—and probably the most important—step in programming, since without that understanding we may find ourselves wandering aimlessly down a dead-end alley, or worse, down a fruitless alley with no end. “Solving the wrong problem” is a phrase often heard in many contexts, and we certainly don’t want to be victims of that crime. So the first step in programming is answering the question, “What problem am I trying to solve?”

Once you understand the problem, then you must find a solution. This may not be easy, of course, and in fact you may discover several solutions, so we also need a way to measure success. There are various dimensions in which to do this, including correctness (“Will I get the right answer?”) and efficiency (“Will I have enough resources?”). But the distinction of which solution is better is not always clear, since the number of dimensions can be large, and programs will often excel in one dimension and do poorly in others. For example, there may be one solution that is fastest, one that uses the least amount of memory, and one that is easiest to understand. Deciding which to choose can be difficult, and is one of the more interesting challenges that you will face in programming.

The last measure of success mentioned above—clarity of a program—is somewhat elusive, most difficult to measure, and, quite frankly, sometimes difficult to rationalize. But in large software systems clarity is an especially important goal, since the most important maxim about such systems is that they are never really finished! The process of continuing work on a software

system after it is delivered to users is what software engineers call *software maintenance*, and is the most expensive phase of the so-called “software life-cycle.” Software maintenance includes fixing bugs in programs, as well as changing certain functionality and enhancing the system with new features in response to users’ experience.

Therefore taking the time to write programs that are highly legible—easy to understand and reason about—will facilitate the software maintenance process. To complete the emphasis on this issue, it is important to realize that the person performing software maintenance is usually not the person who wrote the original program. So when you write your programs, write them as if you are writing them for someone else to see, understand, and ultimately pass judgement on!

As we work through the many musical examples in this book, we will sometimes express them in several different ways (some of which are dead-ends!), taking the time to contrast them in style, efficiency, clarity, and functionality.¹ We do this not just for pedagogical purposes. *Such reworking of programs is the norm*, and you are encouraged to get into the habit of doing so. Don’t always be satisfied with your first solution to a problem, and always be prepared to go back and change—or even throw away—those parts of your program that you later discover do not fully satisfy your actual needs.

1.1 Computation by Calculation in Haskell

In this text we use the programming language *Haskell* to address many of the issues discussed in the last section. We will avoid the approach of explaining Haskell first and giving examples second. Rather, we will walk, step by step, along the path of understanding a problem, understanding the solution space, and understanding how to express a particular solution in Haskell. It is important to learn how to problem solve!

Along this path we will use whatever tools are appropriate for analyzing a particular problem, very often mathematical tools that should be familiar to the average college student, indeed most to the average high-school student. As we do this we will evolve our problems toward a particular view of computation that is especially useful: that of *computation by calculation*. You will find that such a viewpoint is not only powerful—we won’t shy away from difficult problems—it is also *simple*. Haskell supports well the idea of computation by calculation. Programs in Haskell can be viewed as *functions*

¹At times we will also explore different methods for proving properties of programs.

whose input is that of the problem being solved, and whose output is our desired result; and the behavior of functions can be understood easily as computation by calculation.

An example might help to demonstrate these ideas. Suppose we want to perform an arithmetic calculation such as $3 \times (9 + 5)$. In Haskell we would write this as $3 * (9 + 5)$, since most standard computer keyboards and text editors do not recognize the special symbol \times . To calculate the result, we proceed as follows:

$$\begin{aligned} &3 * (9 + 5) \\ \Rightarrow &3 * 14 \\ \Rightarrow &42 \end{aligned}$$

It turns out that this is not the only way to compute the result, as evidenced by this alternative calculation:²

$$\begin{aligned} &3 * (9 + 5) \\ \Rightarrow &3 * 9 + 3 * 5 \\ \Rightarrow &27 + 3 * 5 \\ \Rightarrow &27 + 15 \\ \Rightarrow &42 \end{aligned}$$

Even though this calculation takes two extra steps, it at least gives the correct answer. Indeed, an important property of each and every program in this textbook—in fact every program that can be written in the functional language Haskell—is that it will always yield the same answer when given the same inputs, regardless of the order we choose to perform the calculations.³ This is precisely the mathematical definition of a function: for the same inputs, it always yields the same output.

On the other hand, the first calculation above took less steps than the second, and so we say that it is more *efficient*. Efficiency in both space (amount of memory used) and time (number of steps executed) is important when searching for solutions to problems, but of course if we get the wrong answer, efficiency is a moot point. In general we will search first for any solution to a problem, and later refine it for better performance.

The above calculations are fairly trivial, of course. But we will be doing much more sophisticated operations soon enough. For starters—and to

²This assumes that multiplication distributes over addition in the number system being used, a point that we will return to later.

³As long as we don't choose a non-terminating sequence of calculations, another issue that we will return to later.

introduce the idea of a function—we could *generalize* the arithmetic operations performed in the previous example by defining a function to perform them for any numbers x , y , and z :

$$\text{simple } x \ y \ z = x * (y + z)$$

This equation defines *simple* as a function of three *arguments*, x , y , and z . In mathematical notation, we might see the above written slightly differently, namely:

$$\text{simple}(x, y, z) = x \times (y + z)$$

In any case, it should be clear that “*simple* 3 9 5” is the same as “ $3 * (9 + 5)$.” In fact the proper way to calculate the result is:

$$\begin{aligned} &\text{simple } 3 \ 9 \ 5 \\ &\Rightarrow 3 * (9 + 5) \\ &\Rightarrow 3 * 14 \\ &\Rightarrow 42 \end{aligned}$$

The first step in this calculation is an example of *unfolding* a function definition: 3 is substituted for x , 9 for y , and 5 for z on the right-hand side of the definition of *simple*. This is an entirely mechanical process, not unlike what the computer actually does to execute the program.

When we wish to say that an expression e evaluates (via zero, one, or possibly many more steps) to the value v , we will write $e \Longrightarrow v$ (this arrow is longer than that used earlier). So we can say directly, for example, that *simple* 3 9 5 \Longrightarrow 42, which should be read “*simple* 3 9 5 evaluates to 42.”

With *simple* now suitably defined, we can repeat the sequence of arithmetic calculations as often as we like, using different values for the arguments to *simple*. For example, *simple* 4 3 2 \Longrightarrow 20.

We can also use calculation to *prove properties* about programs. For example, it should be clear that for any a , b , and c , *simple* $a \ b \ c$ should yield the same result as *simple* $a \ c \ b$. For a proof of this, we calculate *symbolically*; that is, using the symbols a , b , and c rather than concrete numbers such as 3, 5, and 9:

$$\begin{aligned} &\text{simple } a \ b \ c \\ &\Rightarrow a * (b + c) \\ &\Rightarrow a * (c + b) \\ &\Rightarrow \text{simple } a \ c \ b \end{aligned}$$

We will use the same notation for these symbolic steps as for concrete ones. In particular, the arrow in the notation reflects the direction of our reasoning, and nothing more. In general, if $e1 \Rightarrow e2$, then it's also true that $e2 \Rightarrow e1$.

We will also refer to these symbolic steps as “calculations,” even though the computer will not typically perform them when executing a program (although it might perform them *before* a program is run if it thinks that it might make the program run faster). The second step in the calculation above relies on the commutativity of addition (namely that, for any numbers x and y , $x + y = y + x$). The third step is the reverse of an unfold step, and is appropriately called a *fold* calculation. It would be particularly strange if a computer performed this step while executing a program, since it does not seem to be headed toward a final answer. But for proving properties about programs, such “backward reasoning” is quite important.

When we wish to make the justification for each step clearer, whether symbolic or concrete, we will present a calculation with more detail, as in:

$$\begin{array}{l} \text{simple } a \ b \ c \\ \Rightarrow \{ \text{unfold} \} \\ a * (b + c) \\ \Rightarrow \{ \text{commutativity} \} \\ a * (c + b) \\ \Rightarrow \{ \text{fold} \} \\ \text{simple } a \ c \ b \end{array}$$

In most cases, however, this will not be necessary.

Proving properties of programs is another theme that will be repeated often in this text. As the world begins to rely more and more on computers to accomplish not just ordinary tasks such as writing term papers and sending email, but also life-critical tasks such as controlling medical procedures and guiding spacecraft, then the correctness of the programs that we write gains in importance. Proving complex properties of large, complex programs is not easy—and rarely if ever done in practice—but that should not deter us from proving simpler properties of the whole system, or complex properties of parts of the system, since such proofs may uncover errors, and if not, at least help us to gain confidence in our effort.

If you are someone who is already an experienced programmer, the idea of computing *everything* by calculation may seem odd at best, and naive at worst. How does one write to a file, draw a picture, play a sound, or respond to mouse-clicks? If you are wondering about these things, it is hoped that you have patience reading the early chapters, and that you find delight in reading the later chapters where the full power of this approach

begins to shine. We will avoid, however, most comparisons between Haskell and conventional programming languages such as C, C++, Java, or even Scheme or ML (two “almost functional” languages).

In many ways this first chapter is the most difficult chapter in the entire text, since it contains the highest density of new concepts. If you have trouble with some of the ideas here, keep in mind that we will return to almost every idea at later points in the text. And don’t hesitate to return to this chapter later to re-read difficult sections; they will likely be much easier to grasp at that time.

Exercise 1.1 Write out all of the steps in the calculation of the value of

simple (*simple* 2 3 4) 5 6

Exercise 1.2 Prove by calculation that $\text{simple } (a - b) \ a \ b \implies a^2 - b^2$.

Details: In the remainder of the text the need will often arise to explain some aspect of Haskell in more detail, without distracting too much from the primary line of discourse. In those circumstance we will off-set the comments and proceed them with the word “Details,” such as is done with this paragraph.

1.2 Expressions, Values, and Types

In Haskell, the entities that we perform calculations on are called *expressions*, and the entities that result from a calculation—i.e. “the answers”—are called *values*. It is helpful to think of a value just as an expression on which no more calculation can be carried out.

Examples of expressions include *atomic* (meaning, indivisible) values such as the integer 42 and the character ‘a’, which are examples of two *primitive* atomic values. In the next Chapter we will also see examples of *user-defined* atomic values, such as the pitch classes *C*, *Cs*, *Df*, etc. (denoting the musical notes C, C♯, D♭, etc.).

In addition, there are *structured* (meaning, made from smaller pieces) expressions such as the list [*C*, *Cs*, *Df*] and the pair (‘b’, 4) (lists and pairs are different in a subtle way, to be described later). Each of these structured expressions is also a value, since by themselves there is no calculation that can be carried out. As another example, $1 + 2$ is an expression, and one step of calculation yields the expression 3, which is a value, since no more calculations can be performed.

Sometimes, however, an expression has only a never-ending sequence of calculations. For example, if x is defined as:

$$x = x + 1$$

then here's what happens when we try to calculate the value of x :

$$\begin{aligned} x & \\ \Rightarrow x + 1 & \\ \Rightarrow (x + 1) + 1 & \\ \Rightarrow ((x + 1) + 1) + 1 & \\ \Rightarrow (((x + 1) + 1) + 1) + 1 & \\ \dots & \end{aligned}$$

This is clearly a never-ending sequence of steps, in which case we say that the expression does not terminate, or is *non-terminating*. In such cases the symbol \perp , pronounced “bottom,” is used to denote the value of the expression.

Every expression (and therefore every value) also has an associated *type*. You can think of types as sets of expressions (or values), in which members of the same set have much in common. Examples include the atomic types *Integer* (the set of all fixed-precision integers) and *Char* (the set of all characters), as well as the structured types $[Integer]$ and $[PitchClass]$ (the set of all lists of integers and pitch classes, respectively) and $(Char, Integer)$ (the set of all character/integer pairs). The association of an expression or value with its type is very important, and there is a special way of expressing it in Haskell. Using the examples of values and types above, we write:

$$\begin{aligned} 42 &:: Integer \\ 'a' &:: Char \\ [C, Cs, Df] &:: [PitchClass] \\ ('b', 4) &:: (Char, Integer) \end{aligned}$$

Details: Literal characters are written enclosed in single forward quotes, as in `'a'`, `'A'`, `'b'`, `' '`, `'!'`, `' '` (a space), etc. (There are some exceptions, however; see the Haskell Report for details.)

The `::` should be read “has type,” as in “42 has type *Integer*.”

Details: Note that the names of specific types are capitalized, such as *Integer* and *Char*, but the names of values are not, such as *simple* and *x*. This is not just a convention: it is required when programming in Haskell.

In addition, the case of the other characters matters, too. For example, *test*, *teSt*, and *tEST* are all distinct names for values, as are *Test*, *TeST*, and *TEST* for types.

Haskell’s *type system* ensures that Haskell programs are *well-typed*; that is, that the programmer has not mismatched types in some way. For example, it does not make much sense to add together two characters, so the expression `'a' + 'b'` is *ill-typed*. The best news is that Haskell’s type system will tell you if your program is well-typed *before you run it*. This is a big advantage, since most programming errors are manifested as typing errors.

1.3 Function Types and Type Signatures

What should the type of a function be? It seems that it should at least convey the fact that a function takes values of one type—*T1*, say—as input, and returns values of (possibly) some other type—*T2*, say—as output. In Haskell this is written $T1 \rightarrow T2$, and we say that such a function “maps values of type *T1* to values of type *T2*.” If there is more than one argument, the notation is extended with more arrows. For example, if our intent is that the function *simple* defined in the previous section has type $Integer \rightarrow Integer \rightarrow Integer \rightarrow Integer$, we can declare this fact by including a *type signature* with the definition of *simple*:

```
simple :: Integer → Integer → Integer → Integer
simple x y z = x * (y + z)
```

Details: When you write Haskell programs using a typical text editor, you will not see nice fonts and arrows as in $Integer \rightarrow Integer$. Rather, you will have to type `Integer -> Integer`.

Haskell’s type system also ensures that user-supplied type signatures such as this one are correct. Actually, Haskell’s type system is powerful enough to allow us to avoid writing any type signatures at all, in which case we say that the type system *infers* the correct types for us.⁴ Nevertheless, judicious placement of type signatures, as we did for *simple*, is a good habit, since type signatures are an effective form of documentation and help bring programming errors to light. Also, in almost every example in this text, we

⁴There are a few exceptions this rule, and in the case of *simple* the inferred type is actually a bit more general than that written above. Both of these points will be returned to later.

will make a habit of first talking about the types of expressions and functions as a way to better understand the problem at hand, organize our thoughts, and lay down the first ideas of a solution.

The normal use of a function is referred to as *function application*. For example, *simple* 3 9 5 is the application of the function *simple* to the arguments 3, 9, and 5.

Details: Some functions, such as $(+)$, are applied using what is known as *infix syntax*; that is, the function is written between the two arguments rather than in front of them (compare $x + y$ to $f\ x\ y$). Infix functions are often called *operators*, and are distinguished by the fact that they do not contain any numbers or letters of the alphabet. Thus $^!$ and $*\#$: are infix operators, whereas *thisIsAFunction* and *f9g* are not (but are still valid names for functions or other values). The only exception to this is that the symbol $'$ is considered to be alphanumeric; thus f' and *one's* are valid names, but not operators.

In Haskell, when referring to an operator as a value, it is enclosed in parentheses, such as when declaring its type, as in:

$$(+) :: Integer \rightarrow Integer \rightarrow Integer$$

Also, when trying to understand an expression such as $f\ x + g\ y$, there is a simple rule to remember: function application always has “higher precedence” than operator application, so that $f\ x + g\ y$ is the same as $(f\ x) + (g\ y)$.

Despite all of these syntactic differences, however, operators are still just functions.

Exercise 1.3 Identify the well-typed expressions in the following and, for each, give its proper type:

$[(2, 3), (4, 5)]$
 $[Cs, 42]$
 $(Df, -42)$
 $simple\ 'a'\ 'b'\ 'c'$
 $(simple\ 1\ 2\ 3, simple)$

1.4 Abstraction, Abstraction, Abstraction

The title of this section is the answer to the question: “What are the three most important ideas in programming?” Well, perhaps this is an overstatement, but the hope is that it has gotten your attention, at least. Webster defines the verb “abstract” as follows:

abstract, *vt* (1) remove, separate (2) to consider apart from application to a particular instance.

In programming we do this when we see a repeating pattern of some sort, and wish to “separate” that pattern from the “particular instances” in which it appears. Let’s refer to this process as the *abstraction principle*, and see how it might manifest itself in problem solving.

1.4.1 Naming

One of the most basic ideas in programming—for that matter, in every day life—is to *name* things. For example, we may wish to give a name to the value of π , since it is inconvenient to retype (or remember) the value of π beyond a small number of digits. In mathematics the greek letter π in fact *is* the name for this value, but unfortunately we don’t have the luxury of using greek letters on standard computer keyboards and text editors. So in Haskell we write:

```
pi :: Float
pi = 3.1415927
```

to associate the name *pi* with the number 3.1415927. The type signature in the first line declares *pi* to be a *floating-point number*, which mathematically—and in Haskell—is distinct from an integer.⁵ Now we can use the name *pi* in expressions whenever we want; it is an abstract representation, if you will, of the number 3.1415927. Furthermore, if we ever have a need to change a named value (which hopefully won’t ever happen for *pi*, but could certainly happen for other values), we would only have to change it in one place, instead of in the possibly large number of places where it is used.

Suppose now that we are working on a problem whose solution requires writing some expression more than once. For example, we might find ourselves computing something such as:

```
x :: Float
x = f (a - b + 2) + g y (a - b + 2)
```

The first line declares *x* to be a floating-point number, while the second is an equation that defines the value of *x*. Note on the right-hand side of this equation that the expression $a - b + 2$ is repeated—it has two instances—and thus, applying the abstraction principle, we wish to separate it from

⁵We will have more to say about floating-point numbers later in this chapter.

these instances. We already know how to do this—it’s called *naming*—so we might choose to rewrite the single equation above as two:

$$\begin{aligned}c &= a - b + 2 \\x &= f\ c + g\ y\ c\end{aligned}$$

If, however, the definition of c is not intended for use elsewhere in the program, then it is advantageous to “hide” the definition of c within the definition of x . This will avoid cluttering up the namespace, and prevents c from clashing with some other value named c . To achieve this, we simply use a **let** expression:

$$\begin{aligned}x &= \mathbf{let}\ c = a - b + 2 \\&\quad \mathbf{in}\ f\ c + g\ y\ c\end{aligned}$$

A **let** expression restricts the *visibility* of the names that it creates to the internal workings of the **let** expression itself. For example, if we write:

$$\begin{aligned}c &= 42 \\x &= \mathbf{let}\ c = a - b + 2 \\&\quad \mathbf{in}\ f\ c + g\ y\ c\end{aligned}$$

then there is no conflict of names—the “outer” c is completely different from the “inner” one enclosed in the **let** expression. Think of the inner c as analogous to the first name of someone in your household. If your brother’s name is “John” he will not be confused with John Thompson who lives down the street when you say, “John spilled the milk.”

Details: An equation such as $c = 42$ is called a *binding*. A simple rule to remember when programming in Haskell is never to give more than one binding for the same name in a context where the names can be confused, whether at the top level of your program or nestled within a **let** expression. For example, this is not allowed:

$$\begin{aligned}a &= 42 \\a &= 43\end{aligned}$$

nor is this:

$$\begin{aligned}a &= 42 \\b &= 43 \\a &= 44\end{aligned}$$

So you can see that naming—using either top-level equations or equations within a **let** expression—is an example of the abstraction principle in action.

1.4.2 Functional Abstraction

[I should replace the following with something musical, but for now it will have to do as is.]

Let's now consider a more complex example. Suppose we are computing the sum of the areas of three circles with radii $r1$, $r2$, and $r3$, as expressed by:

```
totalArea :: Float
totalArea = pi * r1^2 + pi * r2^2 + pi * r3^2
```

Details: $(^)$ is Haskell's integer exponentiation operator. In mathematics we would write $\pi \times r^2$ or just πr^2 instead of $pi * r^2$.

Although there isn't an obvious repeating expression here as there was in the last example, there is a repeating *pattern of operations*. Namely, the operations that square some given quantity—in this case the radius—and then multiply the result by π . To abstract a sequence of operations such as this, we use a *function*—which we will give the name *circleArea*—that takes the “given quantity”—the radius—as an argument. There are three instances of the pattern, each of which we can expect to replace with a call to *circleArea*. This leads to:

```
circleArea :: Float → Float
circleArea r = pi * r^2

totalArea = circleArea r1 + circleArea r2 + circleArea r3
```

Using the idea of unfolding described earlier, it is easy to verify that this definition is equivalent to the previous one.

This application of the abstraction principle is sometimes called *functional abstraction*, since the sequence of operations is abstracted as a function, in this case *circleArea*. Actually, it can be seen as a generalization of the previous kind of abstraction: *naming*. That is, *circleArea r1* is just a name for $pi * r1^2$, *circleArea r2* for $pi * r2^2$, and *circleArea r3* for $pi * r3^2$. Or in other words, a named quantity such as c or pi defined previously can be thought of as a function with no arguments.

Note that *circleArea* takes a radius (a floating-point number) as an argument and returns the area (also a floating-point number) as a result, as reflected in its type signature.

The definition of *circleArea* could also be hidden within *totalArea* using a **let** expression as we did in the previous example:

```
totalArea = let circleArea r = pi * r^2
            in circleArea r1 + circleArea r2 + circleArea r3
```

On the other hand, it is more likely that computing the area of a circle will be useful elsewhere in the program, so leaving the definition at the top level is probably preferable in this case.

1.4.3 Data Abstraction

The value of *totalArea* is the sum of the areas of three circles. But what if in another situation we must add the areas of five circles, or in other situations even more? In situations where the number of things is not certain, it is useful to represent them in a *list* whose length is arbitrary. So imagine that we are given an entire list of circle areas, whose length isn't known at the time we write the program. What now?

We will define a function *listSum* to add the elements of a list. Before doing so, however, there is a bit more to say about lists.

Lists are an example of a *data structure*, and when their use is motivated by the abstraction principle, we will say that we are applying *data abstraction*. Earlier we saw the example `[1,2,3]` as a list of integers, whose type is thus `[Integer]`. A list with *no* elements is—not surprisingly—written `[]`, and pronounced “nil.” To add a single element *x* to the front of a list *xs*, we write `x : xs`. (Note the naming convention used here; *xs* is the plural of *x*, and should be read that way.) In fact, the list `[1,2,3]` is equivalent to `1 : (2 : (3 : []))`, which can also be written `1 : 2 : 3 : []` since the infix operator `(:)` is “right associative.”

Details: In mathematics we rarely worry about whether the notation $a + b + c$ stands for $(a + b) + c$ (in which case $+$ would be “left associative”) or $a + (b + c)$ (in which case $+$ would “right associative”). This is because in situations where the parentheses are left out it's usually the case that the operator is *mathematically* associative, meaning that it doesn't matter which interpretation we choose. If the interpretation *does* matter, mathematicians will include parentheses to make it clear. Furthermore, in mathematics there is an implicit assumption that some operators have higher *precedence* than others; for example, $2 \times a + b$ is interpreted as $(2 \times a) + b$, not $2 \times (a + b)$.

In most programming languages, including Haskell, each operator is defined as having some precedence level and to be either left or right associative. For arithmetic operators, mathematical convention is usually followed; for example, $2 * a + b$ is interpreted as $(2 * a) + b$ in Haskell. The predefined list-forming operator `(:)` is defined to be right associative. Just

as in mathematics, this associativity can be over-ridden by using parentheses: thus $(a : b) : c$ is a valid Haskell expression (assuming that it is well-typed), and is very different from $a : b : c$. We will see later how to specify the associativity and precedence of new operators that we define.

Examples of pre-defined functions defined on lists in Haskell include *head* and *tail*, which return the “head” and “tail” of a list, respectively. That is, $\text{head } (x : xs) \Rightarrow x$ and $\text{tail } (x : xs) \Rightarrow xs$ (we will define these two functions formally in Section 3.1). Another example is the function $(++)$ which *concatenates*, or *appends*, together its two list arguments. For example, $[1, 2, 3] ++ [4, 5, 6] \Rightarrow [1, 2, 3, 4, 5, 6]$ ($(++)$ will be defined in Section 3.3).

Returning to the problem of defining a function to add the elements of a list, let’s first express what its type should be:

$$\text{listSum} :: [\text{Float}] \rightarrow \text{Float}$$

Now we must define its behavior appropriately. Often in solving problems such as this it is helpful to consider, one by one, all possible cases that could arise. To compute the sum of the elements of a list, what might the list look like? The list could be empty, in which case the sum is surely 0. So we write:

$$\text{listSum } [] = 0$$

The other possibility is that the list *isn’t* empty—i.e. it contains at least one element—in which case the sum is the first number plus the sum of the remainder of the list. So we write:

$$\text{listSum } (x : xs) = x + \text{listSum } xs$$

Combining these two equations with the type signature brings us to the complete definition of the function *listSum*:

$$\begin{aligned} \text{listSum} &:: [\text{Float}] \rightarrow \text{Float} \\ \text{listSum } [] &= 0 \\ \text{listSum } (x : xs) &= x + \text{listSum } xs \end{aligned}$$

Details: Although intuitive, this example highlights an important aspect of Haskell: *pattern matching*. The left-hand sides of the equations contain *patterns* such as $[]$ and $x : xs$. When a function is applied, these patterns are *matched* against the argument values in a fairly intuitive way ($[]$ only matches the empty list, and $x : xs$ will successfully match any list with at

least one element, while naming the first element x and the rest of the list xs). If the match succeeds, the right-hand side is evaluated and returned as the result of the application. If it fails, the next equation is tried, and if all equations fail, an error results. All of the equations that define a particular function must appear together, one after the other.

Defining functions by pattern matching is quite common in Haskell, and you should eventually become familiar with the various kinds of patterns that are allowed; see Appendix C for a concise summary.

This is called a *recursive* function definition since *listSum* “refers to itself” on the right-hand side of the second equation. Recursion is a very powerful technique that you will see used many times in this text. It is also an example of a general problem-solving technique where a large problem is broken down into many simpler but similar problems; solving these simpler problems one-by-one leads to a solution to the larger problem.

Here is an example of *listSum* in action:

```
listSum [1,2,3]
⇒ listSum (1 : (2 : (3 : [])))
⇒ 1 + listSum (2 : (3 : []))
⇒ 1 + (2 + listSum (3 : []))
⇒ 1 + (2 + (3 + listSum []))
⇒ 1 + (2 + (3 + 0))
⇒ 1 + (2 + 3)
⇒ 1 + 5
⇒ 6
```

The first step above is not really a calculation, but rather a rewriting of the list syntax. The remaining calculations consist of four unfold steps followed by three integer additions.

Given this definition of *listSum* we can rewrite the definition of *totalArea* as:

```
totalArea = listSum [circleArea r1, circleArea r2, circleArea r3]
```

This may not seem like much of an improvement, but if we were adding many such circle areas in some other context, it would be. Indeed, lists are arguably the most commonly used structured data type in Haskell. In the next chapter we will see a more convincing example of the use of lists; namely, to represent the vertices that make up a polygon. Since a polygon can have an arbitrary number of vertices, using a data structure such as a list seems like just the right approach.

In any case, how do we know that this version of *totalArea* behaves the same as the original one? By calculation, of course:

$$\begin{aligned} & \text{listSum } [\text{circleArea } r1, \text{circleArea } r2, \text{circleArea } r3] \\ \implies & \{ \text{unfold listSum (four successive times)} \} \\ & \text{circleArea } r1 + \text{circleArea } r2 + \text{circleArea } r3 + 0 \\ \implies & \{ \text{unfold circleArea (three places)} \} \\ & \pi * r1^2 + \pi * r2^2 + \pi * r3^2 + 0 \\ \Rightarrow & \{ \text{simple arithmetic} \} \\ & \pi * r1^2 + \pi * r2^2 + \pi * r3^2 \end{aligned}$$

1.5 Code Reuse and Modularity

There doesn't seem to be much repetition in our last definition for *totalArea*, so perhaps we're done. In fact, let's pause for a moment and consider how much progress we've made. We started with the definition:

$$\text{totalArea} = \pi * r1^2 + \pi * r2^2 + \pi * r3^2$$

and ended with:

$$\text{totalArea} = \text{listSum } [\text{circleArea } r1, \text{circleArea } r2, \text{circleArea } r3]$$

But additionally, we have introduced definitions for the auxiliary functions *circleArea* and *listSum*. In terms of size, our final program is actually larger than what we began with! So have we actually improved things?

From the standpoint of “removing repeating patterns,” we certainly have, and we could argue that the resulting program is easier to understand as a result. But there is more. Now that we have defined auxiliary functions such as *circleArea* and *listSum*, we can *reuse* them in other contexts. Being able to reuse code is also called *modularity*, since the reused components are like little modules, or bricks, that can form the foundation of many applications.⁶ We've already talked about reusing *circleArea*; and *listSum* is surely reusable: imagine a list of grocery item prices, or class sizes, or city populations, for each of which we must compute the total. In later chapters you will learn other concepts—most notably higher-order functions and polymorphism—that will substantially increase your ability to reuse code.

⁶“Code reuse” and “modularity” are important software engineering principles.

1.6 Beware of Programming with Numbers

In mathematics there are many different kinds of number systems. For example, there are integers, natural numbers (i.e. non-negative integers), real numbers, rational numbers, and complex numbers. These number systems possess many useful properties, such as the fact that multiplication and addition are commutative, and that multiplication distributes over addition. You have undoubtedly learned many of these properties in your studies, and have used them often in algebra, geometry, trigonometry, physics, etc.

Unfortunately, each of these number systems places great demands on computer systems. In particular, a number can in general require an *arbitrary amount of memory* to represent it—even an infinite amount! Clearly, for example, we cannot represent an irrational number such as π exactly; the best we can do is approximate it, or possibly write a program that computes it to whatever (finite) precision that we need in a given application. But even integers (and therefore rational numbers) present problems, since any given integer can be arbitrarily large.

Most programming languages do not deal with these problems very well. In fact, most programming languages do not have exact forms of any of these number systems. Haskell does slightly better than most, in that it has exact forms of integers (the type *Integer*) as well as rational numbers (the type *Rational*, defined in the Ratio Library). But in Haskell and most other languages there is no exact form of real numbers, for example, which are instead approximated by *floating-point numbers* with either single-word precision (*Float* in Haskell) or double-word precision (*Double*). What's worse, the behavior of arithmetic operations on floating-point numbers can vary somewhat depending on what CPU is being used, although hardware standardization in recent years has reduced the degree of this problem.

The bottom line is that, as simple as they may seem, great care must be taken when programming with numbers. Many computer errors, some quite serious and renowned, were rooted in numerical incongruities. The field of mathematics known as *numerical analysis* is concerned precisely with these problems, and programming with floating-point numbers in sophisticated applications often requires a good understanding of numerical analysis to devise proper algorithms and write correct programs.

As a simple example of this problem, consider the distributive law, expressed here as a calculation in Haskell and used earlier in this chapter in calculations involving the function *simple*:

$$a * (b + c) \Rightarrow a * b + a * c$$

For most floating-point numbers, this law is perfectly valid. For example, in the GHC implementation of Haskell, the expressions $pi * (3 + 4) :: Float$ and $pi * 3 + pi * 4 :: Float$ both yield the same result: 21.99115. But funny things can happen when the magnitude of $b + c$ differs significantly from the magnitude of either b or c . For example, the following two calculations are from GHC:

```
5 * (-0.123456 + 0.123457) :: Float => 4.991889e - 6
5 * (-0.123456) + 5 * (0.123457) :: Float => 5.00679e - 6
```

Although the error here is small, its very existence is worrisome, and in certain situations it could be disastrous. We will not discuss the nature of floating-point numbers much further in this text, but just remember that they are *approximations* to the real numbers. If real-number accuracy is important to your application, further study of the nature of floating-point numbers is probably warranted.

On the other hand, the distributive law (and many others) is valid in Haskell for the exact data types *Integer* and *Ratio Integer* (i.e. rationals). However, another problem arises: although the representation of an *Integer* in Haskell is not normally something that we are concerned about, it should be clear that the representation must be allowed to grow to an arbitrary size. For example, Haskell has no problem with the following number:

```
veryBigNumber :: Integer
veryBigNumber = 43208345720348593219876512372134059
```

and such numbers can be added, multiplied, etc. without any loss of accuracy. However, such numbers cannot fit into a single word of computer memory, most of which are limited to 32 bits. Worse, since the computer system does not know ahead of time exactly how many words will be required, it must devise a dynamic scheme to allow just the right number of words to be used in each case. The overhead of implementing this idea unfortunately causes programs to run slower.

For this reason, Haskell provides another integer data type called *Int* which has maximum and minimum values that depend on the word-size of the CPU being used. In other words, every value of type *Int* fits into one word of memory, and the primitive machine instructions for integers can be used to manipulate them very efficiently.⁷ Unfortunately, this means

⁷The Haskell Report requires that every implementation support *Ints* in the range -2^{29} to $2^{29} - 1$, inclusive. The GHC implementation running on a Pentium processor, for example, supports the range -2^{31} to $2^{31} - 1$.

that *overflow* or *underflow* errors could occur when an *Int* value exceeds either the maximum or minimum values. However, most implementations of Haskell (as well as most other languages) do not even tell you when this happens. For example, in GHC, the following *Int* value:

```
i :: Int
i = 1234567890
```

works just fine, but if you multiply it by two, GHC returns the value -1825831516 ! This is because twice *i* exceeds the maximum allowed value, so the resulting bits become nonsensical,⁸ and are interpreted in this case as a negative number of the given magnitude.

This is alarming! Indeed, why should anyone ever use *Int* when *Integer* is available? The answer, as mentioned earlier, is efficiency, but clearly care should be taken when making this choice. If you are indexing into a list, for example, and you are confident that you are not performing index calculations that might result in the above kind of error, then *Int* should work just fine, since a list longer than 2^{31} will not fit into memory anyway! But if you are calculating the number of microseconds in some large time interval, or counting the number of people living on earth, then *Integer* would most likely be a better choice. Choose your number data types wisely!

In this text we will use the data types *Integer*, *Int*, *Float*, *Double* and *Rational* for a variety of different applications; for a discussion of the other number types, consult the Haskell Report. As we use these data types, we will do so without much discussion—this is not, after all, a book on numerical analysis—but we will issue a warning whenever reasoning about floating-point numbers, for example, in a way that might not be technically sound.

⁸Actually, they are perfectly sensible in the following way: the 32-bit binary representation of *i* is 01001001100101100000001011010010, and twice that is 10010011001011000000010110100100. But the latter number is seen as negative because the 32nd bit (the highest-order bit on the CPU on which this was run) is a one, which means it is a negative number in “twos-complement” representation. The twos-complement of this number is in turn 0110110011010011111101001011100, whose decimal representation is 1825831516.

Chapter 2

Simple Music

```
module Haskore.Music where  
infixr 5:+:, :=:
```

In the previous chapter we introduced some of the fundamental ideas of functional programming in Haskell. In this chapter we begin to develop some *musical* ideas as well. As we do so, more Haskell features will be introduced.

2.1 Preliminaries

Sometimes it is useful to use a built-in Haskell data type to directly represent some concept of interest. For example, we may wish to use *Int* to represent octaves, where by convention octave 4 corresponds to the octave containing middle C on the piano. We can express this in Haskell using a *type synonym*:

```
type Octave = Int
```

A type synonym does not create a new data type—it just gives a new name to an existing type. Type synonyms can be defined not just for atomic types such as *Int*, but also for structured types such as pairs. For example, in music theory a pitch is normally defined as a pair, a pitch class and an octave. Assuming the existence of a data type called *PitchClass*, we can write the following type synonym:

```
type Pitch = (PitchClass, Octave)
```

For example, “concert A,” i.e. A above middle C (sometimes written A4) corresponds to the pitch (*A*, 4). For convenience we could define a Haskell variable with that value as follows:

```

a4 :: Pitch
a4 = (A, 4)    -- concert A

```

Details: This example also demonstrates the use of program *comments*. Any text to the right of “ -- ” till the end of the line is considered to be a comment, and is effectively ignored. Haskell also permits *nested* comments that have the form { - this is a comment - } and can appear anywhere in a program.

Another useful musical concept is *duration*. Rather than use either integers or floating-point numbers, we will use *rational* numbers to denote duration:

```

type Dur = Rational

```

Rational is the data type of rational numbers expressed as ratios of *Integers* in Haskell.

Rational numbers in Haskell are printed by GHCi in the form $n \% d$, where n is the numerator, and d is the denominator. Even a whole number, say the number 42, will print as $42 \% 1$ if it is a *Rational* number. To create a *Rational* number in our program, however, all we have to do is use the normal division operator, as in the following definition of a quarter note:

```

qn :: Dur
qn = 1 / 4    -- quarter note

```

So far so good. But what about *PitchClass*? We might try to use integers to represent pitch classes as well, but this is not very elegant—ideally we would like to write something that looks more like the conventional pitch class names C, C \sharp , Db, D, etc. The solution is to use an *algebraic data type* in Haskell:

```

data PitchClass = Cff | Cf | C | Dff | Cs | Df | Css | D | Eff | Ds
                | Ef | Fff | Dss | E | Es | Ff | F | Gff | Ess | Fs
                | Gf | Fss | G | Aff | Gs | Af | Gss | A | Bff | As
                | Bf | Ass | B | Bs | Bss
deriving (Eq, Ord, Show, Read, Enum)

```

Ignoring the line beginning with “**deriving**” for the moment, this data type declaration simply enumerates the 21 pitch class names (three for each of the note names A through G). Note that enharmonics (such as G \sharp and Ab) are listed separately, which may be important in certain applications.

Details: All constructors in a `data` declaration must be capitalized. In this way they are syntactically distinguished from ordinary values. This distinction is useful since only constructors can be used in the pattern matching that is part of a function definition, as will be described shortly.

Keep in mind that *PitchClass* is a completely new, user-defined data type that is not equal to any other.

2.2 Notes and Music

We can of course define other data types for other purposes. For example, we will want to define the notion of a *note* (the pairing of a pitch with a duration), and a *rest*. Both of these can be thought of as *primitive* musical values, and thus we write:

```
data Prim = Note Dur Pitch
          | Rest Dur
deriving (Show, Eq, Ord)
```

For example, *Note qn a4* is concert A played as a quarter note, and *Rest 1* is a whole-note rest.

This definition is not completely satisfactory, however, because we may wish to attach other information to a note, such as its loudness, or some other annotation or articulation. Furthermore, the pitch itself may actually be a percussive sound, having no true pitch at all. To fix this we will introduce an important concept in Haskell, namely *polymorphism*—the ability to parameterize over types. Instead of fixing the type of the pitch of a note, we will leave it unspecified through the use of a *type variable*, as follows:

```
data Primitive a = Note Dur a
                  | Rest Dur
deriving (Show, Eq, Ord)
```

Note the type variable *a*, which is used as an argument to *Primitive*, and then used in the body of the declaration—just like a variable in a function. *Primitive Pitch* is now the same as (or, technically, is now *isomorphic to*) the type *Prim*. Indeed, instead of defining *Prim* as above, we could now use a type synonym instead:

```
type Prim = Primitive Pitch
```

But *Primitive* is more flexible than *Prim*, since, for example, we could add loudness by pairing loudness with pitch, as in *Primitive (Pitch, Loudness)*. We will see more concrete instances of this idea later.

So far we only have a way to express primitive notes and rests—how do we combine many notes and rests into a larger composition? To achieve this we will define another polymorphic data type, perhaps the most important data type used in this book, which defines the fundamental structure of a musical entity:

```
data Music a = Primitive (Primitive a)    -- primitive value
              | Music a :+: Music a      -- sequential composition
              | Music a :=: Music a      -- parallel composition
              | Modify Control (Music a)  -- modifier
deriving (Show, Eq, Ord)
```

Details: The first line here looks odd: the name *Primitive* appears twice. The first occurrence, however, is the name of a new *constructor* in the *Music* data type, whereas the second is the name of the existing *data type* defined above. Haskell allows using the same name to define a constructor and a data type, since they can never be confused: the context in which they are used will always be sufficient to distinguish them.

Also note the use of *infix constructors* (*:+:*) and (*:=:*). Infix constructors are just like infix operators in Haskell, but they must begin with a colon. This distinction exists to make it easier to pattern match, and is analogous to the distinction between ordinary names (which must begin with a lower-case character) and constructor names (which must begin with an upper-case character).

It is convenient to represent these musical ideas as a recursive datatype because we wish to not only construct musical values, but also take them apart, analyze their structure, print them in a structure-preserving way, interpret them for performance purposes, etc. We will see many examples of these kinds of processes shortly.

This data type declaration essentially says that a value of type *Music a* has one of four possible forms:

- *Primitive p*, where *p* is a primitive value of type *Primitive a*, for some type *a*. For example:

```
ma4 :: Music Pitch
ma4 = Primitive (Note qn a4)
```

is the musical value corresponding to a quarter-note rendition of concert A.

- $m1 :+: m2$ is the *sequential* composition of $m1$ and $m2$; i.e. $m1$ and $m2$ are played in sequence.
- $m1 :=: m2$ is the *parallel* composition of $m1$ and $m2$; i.e. $m1$ and $m2$ are played simultaneously.
- *Modify* $cntrl\ m$ is an “annotated” version of m in which the control parameter $cntrl$ specifies some way in which m is to be modified.

Details: Note that *Music a* is defined in terms of *Music a*, and thus we say that is a *recursive* data type. It is also often called an *inductive* data type, since it is, in essence, an inductive definition of an infinite number of values, each of which can be arbitrarily complex.

The *Control* data type is defined as follows:

```
data Control =
    Tempo Rational      -- scale the tempo
  | Transpose AbsPitch  -- transposition
  | Instrument InstrumentName -- instrument label
  | Phrase [PhraseAttribute] -- phrase attributes
  | Player PlayerName   -- player label
deriving (Show, Eq, Ord)
```

```
type PlayerName = String
```

It allows one to annotate a *Music* value with a tempo change, a transposition, a phrase attribute, a player name, or an instrument. Instrument names are borrowed from the General MIDI standard, and are captured as an algebraic data type in Figure 2.1. Phrase attributes and the concept of a “player” are closely related, but a full explanation is deferred until Chapter 6.

2.3 Convenient Auxiliary Functions

For convenient we define a number of functions to make it easier to write certain kinds of musical values. For starters, we define:

```
note d p = Primitive (Note d p)
rest d = Primitive (Rest d)
```



```

data InstrumentName
= AcousticGrandPiano | BrightAcousticPiano | ElectricGrandPiano
| HonkyTonkPiano | RhodesPiano | ChorusedPiano
| Harpsichord | Clavinet | Celesta | Glockenspiel | MusicBox
| Vibraphone | Marimba | Xylophone | TubularBells
| Dulcimer | HammondOrgan | PercussiveOrgan
| RockOrgan | ChurchOrgan | ReedOrgan
| Accordion | Harmonica | TangoAccordion
| AcousticGuitarNylon | AcousticGuitarSteel | ElectricGuitarJazz
| ElectricGuitarClean | ElectricGuitarMuted | OverdrivenGuitar
| DistortionGuitar | GuitarHarmonics | AcousticBass
| ElectricBassFingered | ElectricBassPicked | FretlessBass
| SlapBass1 | SlapBass2 | SynthBass1 | SynthBass2
| Violin | Viola | Cello | Contrabass | TremoloStrings
| PizzicatoStrings | OrchestralHarp | Timpani
| StringEnsemble1 | StringEnsemble2 | SynthStrings1
| SynthStrings2 | ChoirAahs | VoiceOohs | SynthVoice
| OrchestraHit | Trumpet | Trombone | Tuba
| MutedTrumpet | FrenchHorn | BrassSection | SynthBrass1
| SynthBrass2 | SopranoSax | AltoSax | TenorSax
| BaritoneSax | Oboe | Bassoon | EnglishHorn | Clarinet
| Piccolo | Flute | Recorder | PanFlute | BlownBottle
| Shakuhachi | Whistle | Ocarina | Lead1Square
| Lead2Sawtooth | Lead3Calliope | Lead4Chiff
| Lead5Charang | Lead6Voice | Lead7Fifths
| Lead8BassLead | Pad1NewAge | Pad2Warm
| Pad3Polysynth | Pad4Choir | Pad5Bowed
| Pad6Metallic | Pad7Halo | Pad8Sweep
| FX1Train | FX2Soundtrack | FX3Crystal
| FX4Atmosphere | FX5Brightness | FX6Goblins
| FX7Echoes | FX8SciFi | Sitar | Banjo | Shamisen
| Koto | Kalimba | Bagpipe | Fiddle | Shanai
| TinkleBell | Agogo | SteelDrums | Woodblock | TaikoDrum
| MelodicDrum | SynthDrum | ReverseCymbal
| GuitarFretNoise | BreathNoise | Seashore
| BirdTweet | TelephoneRing | Helicopter
| Applause | Gunshot | Percussion
| Custom String
deriving (Show, Eq, Ord)

```

Figure 2.1: General MIDI Instrument Names

```

tempo r m = Modify (Tempo r) m
transpose i m = Modify (Transpose i) m
instrument i m = Modify (Instrument i) m
phrase pa m = Modify (Phrase pa) m
player pn m = Modify (Player pn) m

```

We can also create simple names for familiar notes, durations, and rests, as shown in Figures 2.2 and 2.3. Despite the large number of them, these names are sufficiently “unusual” that name clashes are unlikely.

As a simple example, here is a ii-V-I chord progression in C major:

```

t251 :: Music Pitch
t251 = let dMinor = d 3 wn :=: f 3 1 :=: a 3 wn
        gMajor = g 3 wn :=: b 3 1 :=: d 4 wn
        cMajor = c 3 bn :=: e 3 2 :=: g 3 bn
      in dMinor :+: gMajor :+: cMajor

```

Details: Note that more than one equation is allowed in a **let** expression. The first characters of each equation, however, must line up vertically, and if an equation takes more than one line then the subsequent lines must be to the right of the first characters. For example, this is legal:

```

let a = aLongName
      + anEvenLongerName
    b = 56
in...

```

but neither of these are:

```

let a = aLongName
    + anEvenLongerName
    b = 56
in...

```

```

let a = aLongName
      + anEvenLongerName
    b = 56
in...

```

(The second line of the first example is too far to the left, as is the third line in the second example.)

Although this rule, called the *layout rule*, may seem a bit *ad hoc*, it avoids having to use special syntax to denote the end of one equation and the

cf, c, cs, df, d, ds, ef, e, es, ff, f, fs, gf, g, gs, af, a, as, bf, b, bs::
Octave → Dur → Music Pitch

cff o d = note d (Cff, o)
cf o d = note d (Cf, o)
c o d = note d (C, o)
cs o d = note d (Cs, o)
css o d = note d (Css, o)
dff o d = note d (Dff, o)
df o d = note d (Df, o)
d o d = note d (D, o)
ds o d = note d (Ds, o)
dss o d = note d (Dss, o)
eff o d = note d (Eff, o)
ef o d = note d (Ef, o)
e o d = note d (E, o)
es o d = note d (Es, o)
ess o d = note d (Ess, o)
fff o d = note d (Fff, o)
ff o d = note d (Ff, o)
f o d = note d (F, o)
fs o d = note d (Fs, o)
fss o d = note d (Fss, o)
gff o d = note d (Gff, o)
gf o d = note d (Gf, o)
g o d = note d (G, o)
gs o d = note d (Gs, o)
gss o d = note d (Gss, o)
aff o d = note d (Aff, o)
af o d = note d (Af, o)
a o d = note d (A, o)
as o d = note d (As, o)
ass o d = note d (Ass, o)
bff o d = note d (Bff, o)
bf o d = note d (Bf, o)
b o d = note d (B, o)
bs o d = note d (Bs, o)
bss o d = note d (Bss, o)

Figure 2.2: Convenient note names.

```

bn, wn, hn, qn, en, sn, tn, sfn, dwn, dhn,
    dqn, den, dsn, dtn, ddhn, ddqn, dden :: Dur

bnr, wnr, hnr, qnr, enr, snr, tnr, dwnr, dhnr,
    dqnr, denr, dsnr, dtnr, ddhnr, ddqnr, ddenr :: Music Pitch

bn = 2; bnr = rest bn      -- brevis rest
wn = 1; wnr = rest wn      -- whole note rest
hn = 1 / 2; hnr = rest hn  -- half note rest
qn = 1 / 4; qnr = rest qn  -- quarter note rest
en = 1 / 8; enr = rest en  -- eighth note rest
sn = 1 / 16; snr = rest sn -- sixteenth note rest
tn = 1 / 32; tnr = rest tn -- thirty-second note rest
sf = 1 / 64; sfr = rest sf -- sixty-fourth note rest

dwn = 3 / 2; dwnr = rest dwn -- dotted whole note rest
dh = 3 / 4; dhr = rest dh   -- dotted half note rest
dq = 3 / 8; dqr = rest dq   -- dotted quarter note rest
de = 3 / 16; der = rest de  -- dotted eighth note rest
ds = 3 / 32; dsr = rest ds  -- dotted sixteenth note rest
dt = 3 / 64; dtr = rest dt  -- dotted thirty-second note rest

ddhn = 7 / 8; ddhnr = rest ddhn -- double-dotted half note rest
ddqn = 7 / 16; ddqnr = rest ddqn -- double-dotted quarter note rest
dden = 7 / 32; ddenr = rest dden -- double-dotted eighth note rest

```

Figure 2.3: Convenient rest names.

beginning of the next (such as a semicolon), thus enhancing readability. In practice, use of layout is rather intuitive. Just remember two things:

First, the first character following either **where** or **let** (and a few other keywords that we will see later) is what determines the starting column for the set of equations being written. Thus we can begin the equations on the same line as the keyword, the next line, or whatever.

Second, just be sure that the starting column is further to the right than the starting column associated with any immediately surrounding clause (otherwise it would be ambiguous). The “termination” of an equation happens when something appears at or to the left of the starting column associated with that equation.

In order to play this simple example, we can import the *play* function from Hasore’s MIDI library, and simply type:

```
play t251
```

at the GHC command line. Default instruments and tempos are used to then play the resulting composition.

2.4 Absolute Pitches

Treating pitches simply as integers is useful in many settings, so let’s use a type synonym to introduce a concept of “absolute pitch:”

```
type AbsPitch = Int
```

The absolute pitch of a (relative) pitch can be defined mathematically as 12 times the octave, plus the index of the pitch class. We can express this in Haskell as follows:

```
absPitch :: Pitch → AbsPitch
absPitch (pc, oct) = 12 * oct + pcToInt pc
```

Details: Note the use of *pattern-matching* to match the argument of *absPitch* to a pair.

pcToInt is simply a function that converts a particular pitch class to an index, easily expressed as:

$pcToInt :: PitchClass \rightarrow Int$

$pcToInt\ Cff = -2$

$pcToInt\ Cf = -1$

$pcToInt\ C = 0$

$pcToInt\ Cs = 1$

$pcToInt\ Css = 2$

$pcToInt\ Dff = 0$

$pcToInt\ Df = 1$

$pcToInt\ D = 2$

$pcToInt\ Ds = 3$

$pcToInt\ Dss = 4$

$pcToInt\ Eff = 2$

$pcToInt\ Ef = 3$

$pcToInt\ E = 4$

$pcToInt\ Es = 5$

$pcToInt\ Ess = 6$

$pcToInt\ Fff = 3$

$pcToInt\ Ff = 4$

$pcToInt\ F = 5$

$pcToInt\ Fs = 6$

$pcToInt\ Fss = 7$

$pcToInt\ Gff = 5$

$pcToInt\ Gf = 6$

$pcToInt\ G = 7$

$pcToInt\ Gs = 8$

$pcToInt\ Gss = 9$

$pcToInt\ Aff = 7$

$pcToInt\ Af = 8$

$pcToInt\ A = 9$

$pcToInt\ As = 10$

$pcToInt\ Ass = 11$

$pcToInt\ Bff = 9$

$pcToInt\ Bf = 10$

$pcToInt\ B = 11$

```

pcToInt Bs = 12
pcToInt Bss = 13

```

Converting an absolute pitch to a pitch is a bit more tricky, because of enharmonic equivalences. For example, the absolute pitch 15 might correspond to either $(Ds, 1)$ or $(Ef, 1)$. We take the approach of always returning a sharp in such ambiguous cases:

```

pitch :: AbsPitch → Pitch
pitch ap =
  let (oct, n) = divMod ap 12
  in ([ C, Cs, D, Ds, E, F, Fs, G, Gs, A, As, B ] !! n, oct)

```

Details: $(!!)$ is Haskell's zero-based list-indexing function; $list !! n$ returns the $(n + 1)$ th element in $list$. $divMod\ x\ n$ returns a pair (q, r) , where q is the integer quotient of x divided by n , and r is the value of x modulo n .

We can also define a function *trans*, which transposes pitches:

```

trans :: Int → Pitch → Pitch
trans i p = pitch (absPitch p + i)

```

Exercise 2.1 Show that $absPitch\ (pitch\ ap) = ap$, and, up to enharmonic equivalences, $pitch\ (absPitch\ p) = p$.

Exercise 2.2 Show that $trans\ i\ (trans\ j\ p) = trans\ (i + j)\ p$.

Chapter 3

Polymorphic and Higher-Order Functions

In the last chapter we learned a little about polymorphic data types. In this chapter we will also learn about *polymorphic functions*, which are essentially functions defined over polymorphic data types. The already familiar *list* is the most common example of a polymorphic data type, and we will study it in depth in this chapter. Although lists have no direct musical connection, they are perhaps the most commonly used data type in Haskell, and have many applications in computer music programming.

We will also learn about *higher-order functions*, which are functions that take one or more functions as arguments or return a function as a result (functions can also be placed in data structures, making the data constructors higher-order too). Together, polymorphic and higher-order functions substantially increase our expressive power and our ability to reuse code. We will see that both of these new ideas naturally follow the foundations that we have already built.

(A more detailed discussion of pre-defined polymorphic functions that operate on lists can be found in Chapter A.)

3.1 Polymorphic Types

In previous chapters we saw examples of lists containing several different kinds of elements—integers, characters, pitch classes, and so on—and you can well imagine situations requiring lists of other element types as well. Sometimes, however, we don't wish to be so particular about the precise type of the elements. For example, suppose we want to define a function

length that determines the number of elements in a list. We don't really care whether the list contains integers, pitch classes, or even other lists—we imagine computing the length in exactly the same way in each case. The obvious definition is:

$$\begin{aligned} \text{length } [] &= 0 \\ \text{length } (x : xs) &= 1 + \text{length } xs \end{aligned}$$

This recursive definition is self-explanatory. We can read the equations as saying: “The length of the empty list is 0, and the length of a list whose first element is x and remainder is xs is 1 plus the length of xs .”

But what should the type of *length* be? Intuitively, what we'd like to say is that, for *any* type a , the type of *length* is $[a] \rightarrow \text{Integer}$. In Haskell we write this simply as:

$$\text{length} :: [a] \rightarrow \text{Integer}$$

Details: Generic names for types, such as a above, are called *type variables*, and are uncapitalized to distinguish them from specific types such as *Integer*.

So *length* can be applied to a list containing elements of *any* type. For example:

$$\begin{aligned} \text{length } [1, 2, 3] &\Longrightarrow 3 \\ \text{length } [C, Cs, Df] &\Longrightarrow 3 \\ \text{length } [[1], [], [2, 3, 4]] &\Longrightarrow 3 \end{aligned}$$

Note that the type of the argument to *length* in the last example is $[[\text{Integer}]]$; that is, a list of lists of integers.

Here are two other examples of polymorphic list functions, which happen to be pre-defined in Haskell:

$$\begin{aligned} \text{head} &:: [a] \rightarrow a \\ \text{head } (x : _) &= x \end{aligned}$$

$$\begin{aligned} \text{tail} &:: [a] \rightarrow [a] \\ \text{tail } (_ : xs) &= xs \end{aligned}$$

Details: The $_$ on the left-hand side of these equations is called a *wild-card* pattern. It matches any value, and binds no variables. It is useful as a way of documenting the fact that we do not care about the value in that part of the pattern.

These two functions take the “head” and “tail,” respectively, of any non-empty list:

$$\begin{aligned} \text{head } [1, 2, 3] &\Rightarrow 1 \\ \text{head } ['a', 'b', 'c'] &\Rightarrow 'a' \\ \text{tail } [1, 2, 3] &\Rightarrow [2, 3] \\ \text{tail } ['a', 'b', 'c'] &\Rightarrow ['b', 'c'] \end{aligned}$$

Functions such as *length*, *head*, and *tail* are said to be *polymorphic* (*poly* means *many* and *morphic* refers to the structure, or *form*, of objects). Polymorphic functions arise naturally when defining functions on lists and other polymorphic data types, including the *Music* data type defined in the last chapter.

3.2 Abstraction Over Recursive Definitions

Suppose we have a list of pitches, and we wish to convert each of them to an absolute pitch. We might write a function:

$$\begin{aligned} \text{toAbsPitches} &:: [\text{Pitch}] \rightarrow [\text{AbsPitch}] \\ \text{toAbsPitches } [] &= [] \\ \text{toAbsPitches } (p : ps) &= \text{absPitch } p : \text{toAbsPitches } ps \end{aligned}$$

We might also want to convert a list of absolute pitches to a list of pitches:

$$\begin{aligned} \text{toPitches} &:: [\text{AbsPitch}] \rightarrow [\text{Pitch}] \\ \text{toPitches } [] &= [] \\ \text{toPitches } (a : as) &= \text{pitch } a : \text{toPitches } as \end{aligned}$$

These two functions are different, but share something in common: there is a repeating pattern of operations. But the pattern is not quite like any of the examples that we studied earlier, and therefore it is unclear how to apply the abstraction principle. What distinguishes this situation is that there is a repeating pattern of *recursion*.

In discerning the nature of a repeating pattern it’s sometimes helpful to identify those things that *aren’t* repeating—i.e. those things that are *changing*—since these will be the sources of *parameterization*: those values that must be passed as arguments to the abstracted function. In the case above, these changing values are the functions *absPitch* and *pitch*; let’s consider them instances of a new name, *f*. If we then simply rewrite either of the above functions as a new function—let’s call it *map*—that takes an extra argument *f*, we arrive at:

$$\begin{aligned} \text{map } f [] &= [] \\ \text{map } f (x : xs) &= f x : \text{map } f xs \end{aligned}$$

With this definition of *map*, we can now redefine *toAbsPitches* and *toPitches* as:

$$\begin{aligned} \text{toAbsPitches} &:: [Pitch] \rightarrow [AbsPitch] \\ \text{toAbsPitches } ps &= \text{map } \text{absPitch } ps \end{aligned}$$

$$\begin{aligned} \text{toPitches} &:: [AbsPitch] \rightarrow [Pitch] \\ \text{toPitches } as &= \text{map } \text{pitch } as \end{aligned}$$

Note that these definitions are non-recursive; the common pattern of recursion has been abstracted away and isolated in the definition of *map*. They are also very succinct; so much so, that it seems unnecessary to create new names for these functions at all! One of the powers of higher-order functions is that they permit concise yet easy-to-understand definitions such as this, and you will see many similar examples throughout the remainder of the text.

A proof that the new versions of these two functions are equivalent to the old ones can be done via calculation, but requires a proof technique called *induction*, because of the recursive nature of the original function definitions. We will discuss inductive proofs in detail, including these two examples, in Chapter 8.

3.2.1 Map is Polymorphic

What should the type of *map* be? Let's look first at its use in *toAbsPitches*: it takes the function *absPitch* :: *Pitch* → *AbsPitch* as its first argument, a list of *Pitches* as its second argument, and it returns a list of *AbsPitches* as its result. So its type must be:

$$\text{map} :: (Pitch \rightarrow AbsPitch) \rightarrow [Pitch] \rightarrow [AbsPitch]$$

Yet a similar analysis of its use in *toPitches* reveals that *map*'s type should be:

$$\text{map} :: (AbsPitch \rightarrow Pitch) \rightarrow [AbsPitch] \rightarrow [Pitch]$$

This apparent anomaly can be resolved by noting that *map*, like *length*, *head* and *tail*, does not really care what its list element types are, *as long as its functional argument can be applied to them*. Indeed, *map* is *polymorphic*, and its most general type is:

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

This can be read: “*map* is a function that takes a function from any type *a* to any type *b*, and a list of *a*’s, and returns a list of *b*’s.” The correspondence between the two *a*’s and between the two *b*’s is important: a function that converts *Int*’s to *Char*’s, for example, cannot be mapped over a list of *Char*’s. It is easy to see that in the case of *toAbsPitches*, *a* is instantiated as *Pitch* and *b* as *AbsPitch*, whereas in *toPitches*, *a* and *b* are instantiated as *AbsPitch* and *Pitch*, respectively.

Details: In Chapter 1 we mentioned that every expression in Haskell has an associated type. But with polymorphism, you might wonder if there is just one type for every expression. For example, *map* could have any of these types:

$$\begin{aligned} (a \rightarrow b) &\rightarrow [a] \rightarrow [b] \\ (\text{Integer} \rightarrow b) &\rightarrow [\text{Integer}] \rightarrow [b] \\ (a \rightarrow \text{Float}) &\rightarrow [a] \rightarrow [\text{Float}] \\ (\text{Char} \rightarrow \text{Char}) &\rightarrow [\text{Char}] \rightarrow [\text{Char}] \end{aligned}$$

and so on, depending on how it will be used. However, notice that the first of these types is in some fundamental sense more general than the other three. In fact, every expression in Haskell has a unique type known as its *principal type*: the least general type that captures all valid uses of the expression. The first type above is the principal type of *map*, since it captures all valid uses of *map*, yet is less general than, for example, the type $a \rightarrow b \rightarrow c$. As another example, the principal type of *head* is $[a] \rightarrow a$; the types $[b] \rightarrow a$, $b \rightarrow a$, or even *a* are too general, whereas something like $[\text{Integer}] \rightarrow \text{Integer}$ is too specific.¹

3.2.2 Using map

Now that we can picture *map* as a polymorphic function, it is useful to look back on some of the examples we have worked through to see if there are any situations where *map* might have been useful. For example, recall from Section 1.4.3 the definition of *totalArea*:

$$\text{totalArea} = \text{listSum} [\text{circleArea } r1, \text{circleArea } r2, \text{circleArea } r3]$$

It should be clear that this can be rewritten as:

¹The existence of unique principal types is the hallmark feature of the *Hindley-Milner type system* [Hin69, Mil78] that forms the basis of the type systems of Haskell, ML [MTH90] and many other functional languages [Hud89].

```
totalArea = listSum (map circleArea [r1, r2, r3])
```

A simple calculation is all that is needed to show that these are the same:

```
map circleArea [r1, r2, r3]
⇒ circleArea r1 : map circleArea [r2, r3]
⇒ circleArea r1 : circleArea r2 : map circleArea [r3]
⇒ circleArea r1 : circleArea r2 : circleArea r3 : map circleArea []
⇒ circleArea r1 : circleArea r2 : circleArea r3 : []
⇒ [circleArea r1, circleArea r2, circleArea r3]
```

For an interesting musical example, let's generate a whole-tone scale starting at a given pitch:

```
wts :: Pitch → [Music Pitch]
wts p = let ap = absPitch p
        f ap = note qn (pitch ap)
        in map f [mc, mc + 2 .. mc + 12]
```

Details: A list $[a, b \dots c]$ is called an *arithmetic sequence*, and is special syntax for the list $[a, a + d, a + 2 * d, \dots, c]$ where $d = b - a$.

3.3 Append

Let's now consider the problem of *concatenating* or *appending* two lists together; that is, creating a third list that consists of all of the elements from the first list followed by all of the elements of the second. Once again the type of list elements does not matter, so we will define this as a polymorphic infix operator $(++)$:

```
(++) :: [a] → [a] → [a]
```

For example, here are two uses of $(++)$ on different types:

```
[1, 2, 3] ++ [4, 5, 6] ⇒ [1, 2, 3, 4, 5, 6]
[C, E, G] ++ [D, F, A] ⇒ [C, E, G, D, F, A]
```

As usual, we can approach this problem by considering the various possibilities that could arise as input. But in the case of $(++)$ we are given *two* inputs—so which do we consider first? In general this is not an easy question, but in the case of $(++)$ we can get a hint about what to do by noting that the result contains firstly all of the elements from the first list. So let's consider the first list first: it could be empty, or non-empty. If it is empty the answer is easy:

$$[] \mathrel{++} ys = ys$$

and if it is not empty the answer is also straightforward:

$$(x : xs) \mathrel{++} ys = x : (xs \mathrel{++} ys)$$

Note the recursive use of $(++)$. Our full definition is thus:

$$\begin{aligned} (++) &:: [a] \rightarrow [a] \rightarrow [a] \\ [] \mathrel{++} ys &= ys \\ (x : xs) \mathrel{++} ys &= x : (xs \mathrel{++} ys) \end{aligned}$$

The Efficiency and Fixity of Append In Chapter 8 we will prove the following simple property about $(++)$:

$$(xs \mathrel{++} ys) \mathrel{++} zs = xs \mathrel{++} (ys \mathrel{++} zs)$$

That is, $(++)$ is *associative*.

But what about the efficiency of the left-hand and right-hand sides of this equation? It is easy to see via calculation that appending two lists together takes a number of steps proportional to the length of the first list (indeed the second list is not evaluated at all). For example:

$$\begin{aligned} &[1, 2, 3] \mathrel{++} xs \\ \Rightarrow &1 : ([2, 3] \mathrel{++} xs) \\ \Rightarrow &1 : 2 : ([3] \mathrel{++} xs) \\ \Rightarrow &1 : 2 : 3 : ([] \mathrel{++} xs) \\ \Rightarrow &1 : 2 : 3 : xs \end{aligned}$$

Therefore the evaluation of $xs \mathrel{++} (ys \mathrel{++} zs)$ takes a number of steps proportional to the length of xs plus the length of ys . But what about $(xs \mathrel{++} ys) \mathrel{++} zs$? The leftmost append will take a number of steps proportional to the length of xs , but then the rightmost append will require a number of steps proportional to the length of xs plus the length of ys , for a total cost of:

$$2 * \text{length } xs + \text{length } ys$$

Thus $xs \mathrel{++} (ys \mathrel{++} zs)$ is more efficient than $(xs \mathrel{++} ys) \mathrel{++} zs$. This is why the Standard Prelude defines the fixity of $(++)$ as:

infixr 5 ++

In other words, if you just write $xs \mathrel{++} ys \mathrel{++} zs$, you will get the most efficient association, namely the right association $xs \mathrel{++} (ys \mathrel{++} zs)$. In the next section we will see a more dramatic example of this property.

3.4 Fold

Suppose we wish to take a list of notes (each of type *Music*) and convert them into a *line*, or *melody*. We can define a recursive function to do this:

$$\begin{aligned} \text{line} &:: [\text{Music } a] \rightarrow \text{Music } a \\ \text{line } [] &= \text{rest } 0 \\ \text{line } (m : ms) &= m :+: \text{line } ms \end{aligned}$$

We might also wish to have a function *chord* that operates in an analogous way, but using $(:=)$ instead of $(:+)$:

$$\begin{aligned} \text{chord} &:: [\text{Music } a] \rightarrow \text{Music } a \\ \text{chord } [] &= \text{rest } 0 \\ \text{chord } (m : ms) &= m :=: \text{chord } ms \end{aligned}$$

In a completely different context we might wish to compute the highest pitch in a list of pitches:

$$\begin{aligned} \text{maxPitch} &:: [\text{Pitch}] \rightarrow \text{Pitch} \\ \text{maxPitch } [] &= \text{pitch } 0 \\ \text{maxPitch } (p : ps) &= p !!! \text{maxPitch } ps \end{aligned}$$

where $!!!$ is defined as:

$$p1 !!! p2 = \text{if } \text{absPitch } p1 > \text{absPitch } p2 \text{ then } p1 \text{ else } p2$$

Details: An expression **if** *pred* **then** *cons* **else** *alt* is called a *conditional expression*. If *pred* (called the *predicate*) is true, then *cons* (called the *consequence*) is the result; if *pred* is false, then *alt* (called the *alternative*) is the result.

Once again we have a situation where several definitions share something in common—a repeating recursive pattern. Using the process that we used to discover *map*, let’s first identify those things that are changing. There are two pairs: the *rest* 0 and *pitch* 0 values (for which we’ll use the generic name *init*, for “initial value”), and the $(:+)$ and $(!!!)$ operators (for which we’ll use the generic name *op*, for “operator”). If we now rewrite either of the above functions as a new function—lets call it *fold*—that takes extra arguments *op* and *init*, we arrive at:²

²The use of the name “*fold*” for this function is historical, and has little to do with the use of “fold” and “unfold” to describe steps in a calculation.

```
fold op init [] = init
fold op init (x : xs) = x 'op' fold op init xs
```

Details: Any normal binary function name can be used as an infix operator by enclosing it in backquotes; $x \text{ 'f' } y$ is equivalent to $f \ x \ y$. Using infix application here for *op* better reflects the structure of the repeating pattern that we are abstracting.

With this definition of *fold* we can now rewrite the definitions of *line*, *chord*, and *maxPitch* as:

```
line, chord :: [Music] → Music
line ms = fold (+:) (rest 0) ms
chord ms = fold (:=) (rest 0) ms

maxPitch :: [Pitch] → Pitch
maxPitch ps = fold (!!!) 0 ps
```

Details: Just as we can turn a function into an operator by enclosing it in backquotes, we can turn an operator into a function by enclosing it in parentheses. This is required in order to pass an operator as a value to another function, as in the examples above. (If we wrote *fold !!! 0 ps* instead of *fold (!!!) 0 ps* it would look like we were trying to compare *fold* to *0 ps*, which is nonsensical and ill-typed.)

In Chapter 8 we will use induction to prove that these new definitions are equivalent to the old ones.

As another example, recall the definition of *listSum* from Section 1.4.3:

```
listSum :: [Float] → Float
listSum [] = 0
listSum (x : xs) = x + listSum xs
```

We can now rewrite this more succinctly using *fold*:

```
listSum :: [Float] → Float
listSum xs = fold (+) 0 xs
```

fold, like *map*, is a highly useful—reusable—function, as we will see through several other examples later in the text. Indeed, it too is polymorphic, for note that it does not depend on the type of the list elements. Its most general type—somewhat trickier than that for *map*—is:

```
fold :: (a → b → b) → b → [a] → b
```

This allows us to use *fold* whenever we need to “collapse” a list of elements using a binary (i.e. two-argument) operator.

3.4.1 Haskell's Folds

Haskell actually defines two versions of *fold* in the Standard Prelude. The first is called *foldr* (“fold-from-the-right”) which is defined the same as our *fold*:

$$\begin{aligned} \text{foldr} &:: (a \rightarrow b \rightarrow b) \rightarrow b \rightarrow [a] \rightarrow b \\ \text{foldr } op \text{ init } [] &= \text{init} \\ \text{foldr } op \text{ init } (x : xs) &= x \text{ 'op' foldr } op \text{ init } xs \end{aligned}$$

A good way to think about *foldr* is that it replaces all occurrences of the list operator $(:)$ with its first argument (a function), and replaces $[]$ with its second argument. In other words:

$$\begin{aligned} \text{foldr } op \text{ init } (x1 : x2 : \dots : xn : []) \\ \implies x1 \text{ 'op' } (x2 \text{ 'op' } (\dots (xn \text{ 'op' } \text{init}) \dots)) \end{aligned}$$

This might help you to understand the type of *foldr* better, and also explains its name: the list is “folded from the right.” Stated another way, for any list *xs*, the following always holds:³

$$\text{foldr } (:) [] xs \implies xs$$

Haskell's second version of *fold* is called *foldl*:

$$\begin{aligned} \text{foldl} &:: (b \rightarrow a \rightarrow b) \rightarrow b \rightarrow [a] \rightarrow b \\ \text{foldl } op \text{ init } [] &= \text{init} \\ \text{foldl } op \text{ init } (x : xs) &= \text{foldl } op \text{ (init 'op' } x) xs \end{aligned}$$

A good way to think about *foldl* is to imagine “folding the list from the left:”

$$\begin{aligned} \text{foldl } op \text{ init } (x1 : x2 : \dots : xn : []) \\ \implies (\dots ((\text{init 'op' } x1) \text{ 'op' } x2) \dots) \text{ 'op' } xn \end{aligned}$$

3.4.2 Why Two Folds?

Note that if we had used *foldl* instead of *foldr* in the definitions given earlier then not much would change; *foldr* and *foldl* would give the same result. Indeed, judging from their types, it looks like the only difference between *foldr* and *foldl* is that the operator takes its arguments in a different order.

So why does Haskell define two versions of *fold*? It turns out that there are situations where using one is more efficient, and possibly “more defined,”

³We will formally prove this in Chapter 8.

than the other. (By more defined, we mean that the function terminates on more values of its input domain.)

Probably the simplest example of this is a generalization of the associativity of $(++)$ discussed in the last section. Suppose that we wish to collapse a list of lists into one list. The Standard Prelude defines the polymorphic function *concat* for this purpose:

$$\begin{aligned} \text{concat} &:: [[a]] \rightarrow [a] \\ \text{concat } xss &= \text{foldr } (++) [] xss \end{aligned}$$

For example:

$$\begin{aligned} \text{concat } [[1], [3, 4], [], [5, 6]] \\ \Rightarrow [1, 2, 3, 4, 5, 6] \end{aligned}$$

More importantly, from the earlier discussion it should be clear that this property holds:

$$\begin{aligned} \text{concat } [xs1, xs2, \dots, xsn] \\ \Rightarrow \text{foldr } (++) [] [xs1, xs2, \dots, xsn] \\ \Rightarrow xs1 ++ (xs2 ++ (\dots(xn ++ []))\dots) \end{aligned}$$

The total cost of this computation is proportional to the sum of the lengths of all of the lists. If each list has the same length *len*, then this cost is $n * len$.

On the other hand, if we had defined *concat* this way:

$$\text{slowConcat } xss = \text{foldl } (++) [] xss$$

then we have:

$$\begin{aligned} \text{slowConcat } [xs1, xs2, \dots, xsn] \\ \Rightarrow \text{foldl } (++) [] [xs1, xs2, \dots, xsn] \\ \Rightarrow (\dots(([] ++ x1) ++ x2)\dots) ++ xn \end{aligned}$$

If each list has the same length *len*, then the cost of this computation will be:

$$\begin{aligned} len + (len + len) + (len + len + len) + \dots + (n - 1) * len \\ \Rightarrow n * (n - 1) * len \end{aligned}$$

which is considerably worse than $n * len$. Thus the choice of *foldr* in the definition of *concat* is quite important.

Similar examples can be given to demonstrate that *foldl* is sometimes more efficient than *foldr*. On the other hand, in many cases the choice does not matter at all (consider, for example, $(+)$). The moral of all this is that care must be taken in the choice between *foldr* and *foldl* if efficiency is a concern.

3.4.3 Fold for Non-empty Lists

One might argue that both *line* and *maxPitch* should not be well defined on an empty list, and for that purpose the Standard Prelude provides functions *foldr1* and *foldl1*, which return an error if applied to an empty list. In certain contexts this may in fact be the preferred behavior for *line* and *maxPitch*, as well as a function *chord* that is similar to *line* except that it does parallel composition. So we could define:

```
line1, chord1 :: [Music] → Music
line1 ms = foldr1 (:+:) ms
chord1 ms = foldr1 (:=:) ms
```

```
maxPitch1 :: [Pitch] → Pitch
maxPitch1 ps = foldr1 (!!!) ps
```

3.5 A Final Example: Reverse

As a final example of a useful list function, consider the problem of *reversing* a list, which we will capture in a function called *reverse*. For example, *reverse* [1, 2, 3] is [3, 2, 1]. Thus *reverse* takes a single list argument, whose possibilities are the normal ones for a list: it is either empty, or it is not. And so we write:

```
reverse :: [a] → [a]
reverse [] = []
reverse (x : xs) = reverse xs ++ [x]
```

This, in fact, is a perfectly good definition for *reverse*—it is certainly clear—except for one small problem: it is terribly inefficient! To see why, first note that the number of steps needed to compute *xs* ++ *ys* is proportional to the length of *xs*. Now suppose that the list argument to *reverse* has length *n*. The recursive call to *reverse* will return a list of length *n* − 1, which is the first argument to (++). Thus the cost to reverse a list of length *n* will be proportional to *n* − 1 plus the cost to reverse a list of length *n* − 1. So the total cost is proportional to (n − 1) + (n − 2) + ⋯ + 1 = *n*(*n* − 1)/2, which in turn is proportional to the square of *n*.

Can we do better than this? Yes we can.

There is another algorithm for reversing a list, which goes something like this: take the first element, and put it at the front of an empty auxiliary list; then take the next element and add it to the front of the auxiliary list

(thus the auxiliary list now consists of the first two elements in the original list, but in reverse order); then do this again and again until you reach the end of the original list. At that point the auxiliary list will be the reverse of the original one.

This algorithm can be expressed recursively, but the auxiliary list implies that we need a function that takes *two* arguments—the original list and the auxiliary one—yet *reverse* only takes one. So we create an auxiliary function *rev*:

$$\begin{aligned} \text{reverse } xs &= \text{rev } [] \ xs \\ \textbf{where } \text{rev } acc \ [] &= acc \\ \text{rev } acc \ (x : xs) &= \text{rev } (x : acc) \ xs \end{aligned}$$

The auxiliary list is the first argument to *rev*, and is called *acc* since it behaves as an “accumulator” of the intermediate results. Note how it is returned as the final result once the end of the original list is reached.

A little thought should convince the reader that this function does not have the quadratic (n^2) behavior of the first algorithm, and indeed can be shown to execute a number of steps that is directly proportional to the length of the list, which we can hardly expect to improve upon.

But now, compare the definition of *rev* with the definition of *foldl*:

$$\begin{aligned} \text{foldl } op \ \text{init} \ [] &= \text{init} \\ \text{foldl } op \ \text{init} \ (x : xs) &= \text{foldl } op \ (\text{init} \ 'op' \ x) \ xs \end{aligned}$$

They are somewhat similar. In fact, suppose we were to slightly rewrite *rev*, yielding:

$$\begin{aligned} \text{rev } op \ acc \ [] &= acc \\ \text{rev } op \ acc \ (x : xs) &= \text{rev } op \ (acc \ 'op' \ x) \ xs \end{aligned}$$

Now *rev* looks exactly like *foldl*, and the question becomes whether or not there is a function that can be substituted for *op* that would make the latter definition of *rev* equivalent to the former one. Indeed there is:

$$\text{revOp } a \ b = b : a$$

For note that:

$$acc \ 'revOp' \ x \Rightarrow \text{revOp } acc \ x \Rightarrow x : acc$$

So *reverse* can be rewritten as:

```
reverse xs = rev revOp [] xs
  where rev op acc [] = acc
        rev op acc (x : xs) = rev op (acc 'op' x) xs
```

which is the same as:

```
reverse xs = foldl revOp [] xs
```

If all of this seems like magic, well, you are starting to see the beauty of functional programming!

3.6 Errors

In the last section we talked about the idea of “returning an error” when the argument to *foldr1* is the empty list. As you might imagine, there are other situations where an error result is also warranted.

There are many ways to deal with such situations, depending on the application, but sometimes we wish to literally stop the program, signalling to the user that some kind of an *error* has occurred. In Haskell this is done with the Standard Prelude function *error* :: *String* → *a*. Note that *error* is polymorphic, meaning that it can be used with any data type. The value of the expression *error s* is ⊥, the completely undefined, or “bottom” value. As an example of its use, here is the definition of *foldr1* from the Standard Prelude:

```
foldr1 :: (a → a → a) → [a] → a
foldr1 f [x] = x
foldr1 f (x : xs) = f x (foldr1 f xs)
foldr1 f [] = error "Prelude.foldr1: empty list"
```

Thus if the anomalous situation arises, the program will terminate immediately, and the string "Prelude.foldr1: empty list" will be printed.

Exercise 3.1 What is the principal type of each of the following expressions:

```
map map
map foldl
```

Exercise 3.2 Rewrite the definition of *length* non-recursively.

Exercise 3.3 Define a function that behaves as each of the following:

1. Doubles each number in a list. For example:

$$\text{doubleEach } [1, 2, 3] \Longrightarrow [2, 4, 6]$$

2. Pairs each element in a list with that number and one plus that number. For example:

$$\text{pairAndOne } [1, 2, 3] \Longrightarrow [(1, 2), (2, 3), (3, 4)]$$

3. Adds together each pair of numbers in a list. For example:

$$\text{addEachPair } [(1, 2), (3, 4), (5, 6)] \Longrightarrow [3, 7, 11]$$

In this exercise and the two that follow, give both recursive and (if possible) non-recursive definitions, and be sure to include type signatures.

Exercise 3.4 Define a function *maxList* that computes the maximum element of a list. Define *minList* analogously.

Exercise 3.5 Define a function that adds “pointwise” the elements of a list of pairs. For example:

$$\text{addPairsPointwise } [(1, 2), (3, 4), (5, 6)] \Longrightarrow (9, 12)$$

Exercise 3.6 Freddie the Frog wants to communicate privately with his girlfriend Francine by *encrypting* messages sent to her. Frog brains are not that large, so they agree on this simple strategy: each character in the text shall be converted to the character “one greater” than it, based on the representation described below (with wrap-around from 255 to 0). Define functions *encrypt* and *decrypt* that will allow Freddie and Francine to communicate using this strategy.

Hint: Characters are often represented inside a computer as some kind of an integer; in the case of Haskell, a 16-bit unicode representation is used. For this exercise, you will want to use two Haskell functions, *toEnum* and *fromEnum*. The first will convert an integer into a character, the second will convert a character into an integer.

Exercise 3.7 Suppose you are given a non-negative integer *amt* representing a sum of money, and a list of coin denominations $[v1, v2, \dots, vn]$, each being a positive integer. Your job is to make change for *amt* using the coins in the coin supply. Define a function *makeChange* to solve this problem. For example, your function may behave like this:

makeChange 99 [5, 1] \Rightarrow [19, 4]

where 99 is the amount and [5, 1] represents the types of coins (say, nickels and pennies in US currency) that we have. The answer [19, 4] means that we can make the exact change with 19 5-unit coins and 4 single-unit coins; this is the best (in terms of the total number of coins) possible solution.

To make things slightly easier, you may assume that the list representing the coin denominations is given in descending order, and that the single-unit coin is always one of the coin types.

[Need to add some musical exercises.]

Chapter 4

More About Higher-Order Functions

You have now seen several examples where functions are passed as arguments to other functions, such as with *fold* and *map*. In this chapter we will see several examples where functions are also returned as values. This will lead to several techniques for improving definitions that we have already written, techniques that we will use often in the remainder of the text.

4.1 Currying

The first improvement relates to the notation we have used to write function applications, such as *simple x y z*. Although we have seen the similarity of this to the mathematical notation *simple(x, y, z)*, in fact there is an important difference, namely that *simple x y z* is actually equivalent to $((\text{simple } x) \ y) \ z$. In other words, function application is *left associative*, taking one argument at a time.

Let's look at the expression $((\text{simple } x) \ y) \ z$ a bit closer: there is an application of *simple* to *x*, the result of which is applied to *y*; so $(\text{simple } x)$ must be a function! The result of this application, $((\text{simple } x) \ y)$, is then applied to *z*, so $((\text{simple } x) \ y)$ must also be a function!

Since each of these intermediate applications yields a function, it seems perfectly reasonable to define a function such as:

$$\text{multSumByFive} = \text{simple } 5$$

What is *simple 5*? From the above argument we know that it must be a function. And from the definition of *simple* in Section 1.1 we might guess that

this function takes two arguments, and returns 5 times their sum. Indeed, we can *calculate* this result as follows:

$$\begin{aligned} & \text{multSumByFive } a \ b \\ \Rightarrow & (\text{simple } 5) \ a \ b \\ \Rightarrow & \text{simple } 5 \ a \ b \\ \Rightarrow & 5 * (a + b) \end{aligned}$$

The intermediate step with parentheses is included just for clarity. This method of applying functions to one argument at a time, yielding intermediate functions along the way, is called *currying*, after the logician Haskell B. Curry who popularized the idea.¹ It is helpful to look at the types of the intermediate functions as arguments are applied:

$$\begin{aligned} \text{simple} & :: \text{Float} \rightarrow \text{Float} \rightarrow \text{Float} \rightarrow \text{Float} \\ \text{simple } 5 & :: \text{Float} \rightarrow \text{Float} \rightarrow \text{Float} \\ \text{simple } 5 \ a & :: \text{Float} \rightarrow \text{Float} \\ \text{simple } 5 \ a \ b & :: \text{Float} \end{aligned}$$

We use currying to improve some of our previous examples as follows. Suppose that I tell you that the expressions $f \ x$ and $g \ x$ are the same, for all values of x . Then it seems clear that the functions f and g are equivalent. So, if we want to define f in terms of g , instead of writing:

$$f \ x = g \ x$$

we can instead simply write:

$$f = g$$

Let's apply this reasoning to the definition of *line* from Section 3.4:

$$\text{line } ms = \text{fold } (:+:) \ (\text{Primitive } (\text{Rest } 0)) \ ms$$

Since function application is left associative, we can rewrite this as:

$$\text{line } ms = (\text{fold } (:+:) \ (\text{Primitive } (\text{Rest } 0)))) \ ms$$

But now applying the same reasoning here as we did for f and g means that we can write this simply as:

$$\text{line} = \text{fold } (:+:) \ (\text{Primitive } (\text{Rest } 0))$$

¹It was actually Schönfinkel who first called attention to this idea [Sch24], but the word “schönfinkelling” is rather a mouthful!

Similarly, the definitions of *maxPitch* and *listSum*

$$\begin{aligned} \text{maxPitch } ps &= \text{fold } (!!) \ 0 \ ps \\ \text{listSum } xs &= \text{foldl } (+) \ 0 \ xs \end{aligned}$$

can be rewritten as:

$$\begin{aligned} \text{maxPitch} &= \text{fold } (!!) \ 0 \\ \text{listSum} &= \text{foldl } (+) \ 0 \end{aligned}$$

We will refer to this kind of simplification as “currying simplification” or just “currying,” even though it actually has a more technical name, “eta contraction.”

Details: Some care should be taken when using this simplification idea.

In particular, note that an equation such as $f \ x = g \ x \ y \ x$ cannot be simplified to $f = g \ x \ y$, since then the x would become undefined!

Here is a more interesting example, in which currying simplification is used three times. Recall from Section 3.5 the definition of *reverse* using *foldl*:

$$\begin{aligned} \text{reverse } xs &= \text{foldl } \text{revOp } [] \ xs \\ \textbf{where } \text{revOp } acc \ x &= x : acc \end{aligned}$$

Using the polymorphic function *flip* which is defined in the Standard Prelude as:

$$\begin{aligned} \text{flip} &:: (a \rightarrow b \rightarrow c) \rightarrow (b \rightarrow a \rightarrow c) \\ \text{flip } f \ x \ y &= f \ y \ x \end{aligned}$$

it should be clear that *revOp* can be rewritten as:

$$\text{revOp } acc \ x = \text{flip } (:) \ acc \ x$$

But now currying simplification can be used twice to reveal that:

$$\text{revOp} = \text{flip } (:) \$$

This, along with a third use of currying, allows us to rewrite the definition of *reverse* simply as:

$$\text{reverse} = \text{foldl } (\text{flip } (:)) \ []$$

This is in fact the way *reverse* is defined in the Standard Prelude.

Exercise 4.1 Show that *flip* (*flip* *f*) is the same as *f*.

Exercise 4.2 What is the type of *ys* in:

```
xs = [1, 2, 3] :: [Float]
ys = map (+) xs
```

Exercise 4.3 Define a function *applyEach* that, given a list of functions, applies each to some given value. For example:

```
applyEach [simple 2 2, (+3)] 5 ==> [14, 8]
```

where *simple* is as defined in Section 1.1.

Exercise 4.4 Define a function *applyAll* that, given a list of functions [*f1*, *f2*, ..., *fn*] and a value *v*, returns the result *f1* (*f2* (...(*fn* *v*)...)). For example:

```
applyAll [simple 2 2, (+3)] 5 ==> 20
```

Exercise 4.5 Recall the discussion about the efficiency of (++) and *concat* in Chapter 3. Which of the following functions is more efficient, and why?

```
appendr, appendl :: [[a]] -> [a]
appendr = foldr (flip (++)) []
appendl = foldl (flip (++)) []
```

4.2 Sections

With a bit more syntax, we can also curry applications of infix operators such as (+). This syntax is called a *section*, and the idea is that, in an expression such as (*x* + *y*), you can omit either the *x* or the *y*, and the result (with the parentheses still intact) is a function of that missing argument. If *both* variables are omitted, it is a function of *two* arguments. In other words, the expressions (*x* +), (+ *y*) and (+) are equivalent, respectively, to the functions:

```
f1 y = x + y
f2 x = x + y
f3 x y = x + y
```

For example, suppose that we need to determine whether each number in a list is positive. Instead of writing:

```
posInts :: [Integer] → [Bool]
posInts xs = map test xs
           where test x = x > 0
```

we can simply write:

```
posInts :: [Integer] → [Bool]
posInts xs = map (>0) xs
```

which can be further simplified using currying:

```
posInts :: [Integer] → [Bool]
posInts = map (>0)
```

This is an extremely concise definition.

As you gain experience with higher-order functions you will not only be able to start writing definitions such as this directly, but you will also start *thinking* in “higher-order” terms. We will see many examples of this kind of reasoning throughout the text.

Exercise 4.6 Define a function *twice* that, given a function *f*, returns a function that applies *f* twice to its argument. For example:

```
(twice (+1)) 2 ⇒ 4
```

What is the principal type of *twice*? Describe what *twice twice* does, and give an example of its use. How about *twice twice twice* and *twice (twice twice)*?

Exercise 4.7 Generalize *twice* defined in the previous exercise by defining a function *power* that takes a function *f* and an integer *n*, and returns a function that applies the function *f* to its argument *n* times. For example:

```
power (+2) 5 1 ⇒ 11
```

Use *power* to define something (anything!) useful.

4.3 Anonymous Functions

The final way to define a function in Haskell is in some sense the most fundamental: it is called an *anonymous function*, or *lambda expressions* (since the concept is drawn directly from Church’s lambda calculus [Chu41]). The idea is that functions are values, just like numbers and characters and strings, and therefore there should be a way to create them without having to give

them a name. As a simple example, an anonymous function that increments its numeric argument by one can be written $\lambda x \rightarrow x + 1$. Anonymous functions are most useful in situations where you don't wish to name them, which is why they are called "anonymous."

Details: The typesetting that we use in this book prints an actual Greek lambda character, but in writing $\lambda x \rightarrow x + 1$ in your programs you will have to write `"\x -> x+1"` instead.

As another example, to add one and then divide by two every element of a list, we could write:

$$\text{map } (\lambda x \rightarrow (x + 1) / 2) \text{ } xs$$

An even better example is an anonymous function that pattern-matches its argument, as in:

$$\text{map } (\lambda(a, b) \rightarrow a + b) \text{ } xs$$

Details: Anonymous functions can only perform one match against an argument. That is, you cannot stack together several anonymous functions to define one function, as you can with equations.

Anonymous functions are considered most fundamental because definitions such as that for *simple* given in Chapter 1:

$$\text{simple } x \ y \ z = x * (y + z)$$

can be written instead as:

$$\text{simple} = \lambda x \ y \ z \rightarrow x * (y + z)$$

Details: $\lambda x \ y \ z \rightarrow exp$ is shorthand for $\lambda x \rightarrow \lambda y \rightarrow \lambda z \rightarrow exp$.

We can also use anonymous functions to explain precisely the behavior of sections. In particular, note that:

$$\begin{aligned} (x+) &\Rightarrow \lambda y \rightarrow x + y \\ (+y) &\Rightarrow \lambda x \rightarrow x + y \\ (+) &\Rightarrow \lambda x \ y \rightarrow x + y \end{aligned}$$

Exercise 4.8 Suppose we define a function *fix* as:

$$\text{fix } f = f (\text{fix } f)$$




Figure 4.1: Gluing Two Functions Together

What is the principal type of *fix*? (This is tricky!) Suppose further that we have a recursive function:

```
remainder :: Integer → Integer → Integer
remainder a b = if a < b then a
                else remainder (a - b) b
```

Rewrite this function using *fix* so that it is not recursive. (Also tricky!) Do you think that this process can be applied to *any* recursive function?

4.4 Function Composition

An example of polymorphism that has nothing to do with data structures arises from the desire to take two functions *f* and *g* and “glue them together,” yielding another function *h* that first applies *g* to its argument, and then applies *f* to that result. This is called function *composition*, and Haskell pre-defines a simple infix operator (*o*) to achieve it, as follows:

```
(o) :: (b → c) → (a → b) → a → c
(f o g) x = f (g x)
```

Details: The symbol for function composition is typeset in this book as *o*, which is the proper mathematical convention. When writing your programs, however, you will have to use a “period” , as in “*f . g*”.

Note the type of the operator (*o*); it is completely polymorphic. Note also that the result of the first function to be applied—some type *b*—must be the same as the type of the argument to the second function to be applied. Pictorially, if you think of a function as a black box that takes input at one end and returns some output at the other, function composition is like connecting two boxes together, end to end, as shown in Figure 4.1.

The ability to compose functions using (*o*) is extremely useful. For example, consider this function to compute the sum of the areas of circles with various radii:

```
totalCircleArea :: [Float] → Float
totalCircleArea radii = listSum (map circleArea radii)
```

We can first add parentheses to emphasize the application of interest:

```
totalCircleArea :: [Float] → Float
totalCircleArea radii = listSum ((map circleArea) radii)
```

then rewrite as a function composition:

```
totalCircleArea :: [Float] → Float
totalCircleArea radii = (listSum ∘ (map circleArea)) radii
```

and finally use currying to simplify:

```
totalCircleArea :: [Float] → Float
totalCircleArea = listSum ∘ map circleArea
```

Similarly, this definition:

```
totalSquareArea :: [Float] → Float
totalSquareArea sides = listSum (map squareArea sides)
```

can be rewritten as:

```
totalSquareArea :: [Float] → Float
totalSquareArea = listSum ∘ map squareArea
```

But let's also create additional compositions. A function that determines whether all of the elements in a list are greater than zero, and one that determines if at least one is greater than zero, can be written:

```
allOverZero, oneOverZero :: [Integer] → Bool
allOverZero = and ∘ posInts
oneOverZero = or ∘ posInts
```

Note that the auxiliary function *posInts* is simple enough that we could incorporate its definition directly, as in:

```
allOverZero, oneOverZero :: [Integer] → Bool
allOverZero = and ∘ map (>0)
oneOverZero = or ∘ map (>0)
```

Details: *and* :: [Bool] → Bool and *or* :: [Bool] → Bool are predefined functions that “and” and “or” together all of the elements in a list, returning a single Boolean result. The *Bool* type is predefined in Haskell simply as:

data *Bool* = *False* | *True*

In the remainder of this text we will not refrain from writing definitions such as this directly, using a small set of rich polymorphic functions such as *fold* and *map*, plus a few others drawn from the Prelude and Standard Libraries.

Exercise 4.9 Rewrite this example:

$\text{map } (\lambda x \rightarrow (x + 1) / 2) \text{ } xs$

using a composition of sections.

Exercise 4.10 Consider the expression:

$\text{map } f \text{ } (\text{map } g \text{ } xs)$

Rewrite this using function composition and a single call to *map*. Then rewrite the earlier example:

$\text{map } (\lambda x \rightarrow (x + 1) / 2) \text{ } xs$

as a “map of a map.”

Exercise 4.11 Go back to any exercises prior to this chapter, and simplify your solutions using ideas learned here.

Exercise 4.12 Using higher-order functions that we have now defined, fill in the two missing functions, *f1* and *f2*, in the evaluation below so that it is valid:

$f1 \text{ } (f2 \text{ } (*) \text{ } [1, 2, 3, 4]) \text{ } 5 \Rightarrow [5, 10, 15, 20]$

Chapter 5

More Music

```
module Haskore.MoreMusic where  
import Haskore.Music
```

In this chapter we will explore a number of simple musical ideas, and contribute to a growing collection of Haskell functions for expressing those ideas.

5.1 Delay and Repeat

Suppose that we wish to describe a melody m accompanied by an identical voice a perfect 5th higher. In Haskore we can simply write $m ::= \text{transpose } 7 \ m$. Similarly, a canon-like structure involving m can be expressed as $m ::= \text{delay } d \ m$, where:

$$\begin{aligned} \text{delay} &:: \text{Dur} \rightarrow \text{Music } a \rightarrow \text{Music } a \\ \text{delay } d \ m &= \text{rest } d \text{ :+} m \end{aligned}$$

More interestingly, Haskell's non-strict semantics also allows us to define *infinite* musical values. For example, a musical value may be repeated *ad nauseum* using this simple function:

$$\begin{aligned} \text{repeatM} &:: \text{Music } a \rightarrow \text{Music } a \\ \text{repeatM } m &= m \text{ :+} \text{repeatM } m \end{aligned}$$

Thus, for example, an infinite ostinato can be expressed in this way, and then used in different contexts that automatically extract only the portion that is actually needed. We will see more examples of this shortly.

Exercise 5.1 Define a function $\text{repM} :: \text{Int} \rightarrow \text{Music } a \rightarrow \text{Music } a$ such that $\text{repM } n \ m$ repeats m n times.

5.2 Inversion and Retrograde

The notions of inversion, retrograde, retrograde inversion, etc. as used in twelve-tone theory are also easily captured in Haskore. These terms are usually applied only to “lines” of notes, i.e. a melody (in twelve-tone theory it is called a “row”). The *retrograde* of a line is simply its reverse—i.e. the notes played in the reverse order. The *inversion* of a line is with respect to a given pitch (by convention usually the first pitch), where the intervals between successive pitches are inverted, i.e. negated. If the absolute pitch of the first note is ap , then each pitch p is converted into an absolute pitch $(ap - \text{absPitch } p) + ap$, in other words $2 * ap - \text{absPitch } p$.

To do all this in Haskell, let’s first define a transformation from a line created by *line* to a list:

```
lineToList :: Music a → [Music a]
lineToList n@(Primitive (Rest 0)) = []
lineToList (n :+: ns) = n : lineToList ns
lineToList _ = error "lineToList: argument not created by line"
```


Using this function it is then straightforward to define *invert*, from which the other functions are easily defined via composition:

```
retro, invert, retroInvert, invertRetro :: Music Pitch → Music Pitch
invert m = line (map inv l)
  where l@(Primitive (Note _ r) : _) = lineToList m
        inv (Primitive (Note d p)) =
            note d (pitch (2 * absPitch r - absPitch p))
        inv (Primitive (Rest d)) = rest d
retro = line ∘ reverse ∘ lineToList
retroInvert = retro ∘ invert
invertRetro = invert ∘ retro
```

Exercise 5.2 Show that $\text{retro} \circ \text{retro}$, $\text{invert} \circ \text{invert}$, and $\text{retroInvert} \circ \text{invertRetro}$ are the identity on values created by *line*.

5.3 Polyrhythms

For some rhythmical ideas, first note that if m is a line of three eighth notes, then *tempo* $(3/2)$ m is a *triplet* of eighth notes. In fact *tempo* can be used to create quite complex rhythmical patterns. For example, consider the

Figure 5.1: Nested Polyrhythms (top: *pr1*; bottom: *pr2*)

“nested polyrhythms” shown in Figure 5.1. They can be expressed naturally in Haskore as follows (note the use of the *where* clause in *pr2* to capture recurring phrases):

```

pr1, pr2 :: Pitch → Music Pitch
pr1 p = tempo (5 / 6)
      (tempo (4 / 3) (mkLn 1 p qn):+
        tempo (3 / 2) (mkLn 3 p en):+
          mkLn 2 p sn):+
          mkLn 1 p qn):+
        tempo (3 / 2) (mkLn 6 p en))

pr2 p = tempo (7 / 6) (m1:+
  tempo (5 / 4) (mkLn 5 p en):+
  m1:+
  tempo (3 / 2) m2)
  where m1 = tempo (5 / 4) (tempo (3 / 2) m2 :+ m2)
        m2 = mkLn 3 p en

mkLn n p d = line (take n (repeat (note d p)))

```

Details: *take n lst* is the first *n* elements of the list *lst*. For example, *take 3 [C, Cs, Df, D, Ds] ⇒ [C, Cs, Df]*. *repeat x* is the infinite list of the same value *x*. For example, *take 3 (repeat 42) ⇒ [42, 42, 42]*.

To play polyrhythms *pr1* and *pr2* in parallel using middle C and middle G, respectively, we do the following:

```
pr12 :: Music Pitch
pr12 = pr1 (C, 4) :=: pr2 (G, 4)
```

5.4 Symbolic Meter Changes

We can implement the notion of “symbolic meter changes” of the form “old-note = newnote” (quarter note = dotted eighth, for example) by defining an infix function:

```
(=:=) :: Dur → Dur → Music a → Music a
old :=: new = tempo (new / old)
```

Of course, using the new function is not much longer than using *Tempo* directly, but it may have nemonic value.

5.5 Computing Duration

It is often desirable to compute the *duration*, in whole notes, of a musical value; we can do so as follows:

```
dur :: Music a → Dur
dur (Primitive (Note d _)) = d
dur (Primitive (Rest d)) = d
dur (m1 :+: m2) = dur m1 + dur m2
dur (m1 :=: m2) = dur m1 'max' dur m2
dur (Modify (Tempo r) m) = dur m / r
dur (Modify _ m) = dur m
```

The duration of a primitive value is obvious. The duration of *m1 :+: m2* is the sum of the two, and the duration of *m1 :=: m2* is the maximum of the two. The only tricky part is the duration of a music value that is modified by the *Tempo* attribute—in this case the duration must be scaled appropriately.

5.6 Super-retrograde

Using *dur* we can define a function *revM* that reverses any *Music* value (and is thus considerably more useful than *retro* defined earlier). Note the tricky treatment of (*:=:*).

```

revM :: Music a → Music a
revM n@(Primitive _) = n
revM (Modify c m) = Modify c (revM m)
revM (m1 :+: m2) = revM m2 :+: revM m1
revM (m1 :=: m2) =
  let d1 = dur m1
      d2 = dur m2
  in if d1 > d2 then revM m1 :=: (rest (d1 - d2) :+: revM m2)
     else (rest (d2 - d1) :+: revM m1) :=: revM m2

```

5.7 Truncating Parallel Composition

Note that the duration of $m1 :=: m2$ is the maximum of the durations of $m1$ and $m2$ (and thus if one is infinite, so is the result). Sometimes we would rather have the result be of duration equal to the shorter of the two. This is not as easy as it sounds, since it may require interrupting the longer one in the middle of a note (or notes).

We will define a “truncating parallel composition” operator ($/=$), but first we will define an auxiliary function *cut* such that *cut* d m is the musical value m “cut short” to have at most duration d :

```

cut :: Dur → Music a → Music a
cut newDur m | newDur ≤ 0 = rest 0
cut newDur (Primitive (Note oldDur x)) = note (min oldDur newDur) x
cut newDur (Primitive (Rest oldDur)) = rest (min oldDur newDur)
cut newDur (m1 :=: m2) = cut newDur m1 :=: cut newDur m2
cut newDur (m1 :+: m2) = let m1' = cut newDur m1
                        m2' = cut (newDur - dur m1') m2
                        in m1' :+: m2'
cut newDur (Modify (Tempo r) m) = tempo r (cut (newDur * r) m)
cut newDur (Modify c m) = Modify c (cut newDur m)

```

Note that *cut* is equipped to handle a *Music* value of infinite length.

With *cut*, the definition of ($/=$) is now straightforward:

```

(/=) :: Music a → Music a → Music a
m1 /=: m2 = cut (min (dur m1) (dur m2)) (m1 :=: m2)

```

Unfortunately, whereas *cut* can handle infinite-duration music values, ($/=$) cannot.

Exercise 5.3 Define a version of ($/=:$) that shortens correctly when either or both of its arguments are infinite in duration.

5.8 Trills

A *trill* is an ornament that alternates rapidly between two (usually adjacent) pitches. We will define two versions of a trill function, both of which take the starting note and an interval for the trill note as arguments (the interval is usually one or two, but can actually be anything). One version will additionally have an argument that specifies how long each trill note should be, whereas the other will have an argument that specifies how many trills should occur. In both cases the total duration will be the same as the duration of the original note.

Here is the first trill function:

```
trill :: Int → Dur → Music Pitch → Music Pitch
trill i sDur (Primitive (Note tDur p)) =
  if sDur ≥ tDur then note tDur p
  else note sDur p
    :+: trill (negate i) sDur (note (tDur - sDur) (trans i p))
trill i d (Modify (Tempo r) m) = tempo r (trill i (d * r) m)
trill i d (Modify c m) = Modify c (trill i d m)
trill _ _ _ = error "trill: input must be a single note."
```

It is simple to define a version of this function that starts on the trill note rather than the start note:

```
trill' :: Int → Dur → Music Pitch → Music Pitch
trill' i sDur m = trill (negate i) sDur (transpose i m)
```

The second way to define a trill is in terms of the number of subdivided notes to be included in the trill. We can use the first trill function to define this new one:

```
trilln :: Int → Int → Music Pitch → Music Pitch
trilln i nTimes m = trill i (dur m / fromIntegral nTimes) m
```

This, too, can be made to start on the other note.

```
trilln' :: Int → Int → Music Pitch → Music Pitch
trilln' i nTimes m = trilln (negate i) nTimes (transpose i m)
```

Finally, a *roll* can be implemented as a trill whose interval is zero. This feature is particularly useful for percussion.

```
roll :: Dur → Music Pitch → Music Pitch
rolln :: Int → Music Pitch → Music Pitch
```

```
roll dur m = trill 0 dur m
rolln nTimes m = trilln 0 nTimes m
```

5.9 Percussion

Percussion is a difficult notion to represent in the abstract. On one hand, a percussion instrument is just another instrument, so why should it be treated differently? On the other hand, even common practice notation treats it specially, even though it has much in common with non-percussive notation. The MIDI standard is equally ambiguous about the treatment of percussion: on one hand, percussion sounds are chosen by specifying an octave and pitch, just like any other instrument; on the other hand these pitches have no tonal meaning whatsoever: they are just a convenient way to select from a large number of percussion sounds. Indeed, part of the General MIDI Standard is a set of names for commonly used percussion sounds.

Since MIDI is such a popular platform, we can at least define some handy functions for using the General MIDI Standard. We start by defining the data type shown in Figure 5.2, which borrows its constructor names from the General MIDI standard. The comments reflecting the “MIDI Key” numbers will be explained later, but basically a MIDI Key is the equivalent of an absolute pitch in Haskore terminology. So all we need is a way to convert these percussion sound names into a *Music* value; i.e. a *Note*:

```
perc :: PercussionSound → Dur → Music Pitch
perc ps dur = note dur (pitch (fromEnum ps + 35))
```

Details: *fromEnum* is a method in the *Enum* class, which is all about enumerations. A data type that is a member of this class can be *enumerated*—i.e. the elements of the data type can be listed in order. *fromEnum* maps each value to its index in this enumeration. Thus *fromEnum AcousticBassDrum* is 0, *fromEnum BassDrum1* is 1, and so on.

For example, here are eight bars of a simple rock or “funk groove” that uses *perc* and *roll*:

```

data PercussionSound =
  | AcousticBassDrum    -- MIDI Key 35
  | BassDrum1          -- MIDI Key 36
  | SideStick           -- ...
  | AcousticSnare | HandClap | ElectricSnare | LowFloorTom
  | ClosedHiHat | HighFloorTom | PedalHiHat | LowTom
  | OpenHiHat | LowMidTom | HiMidTom | CrashCymbal1
  | HighTom | RideCymbal1 | ChineseCymbal | RideBell
  | Tambourine | SplashCymbal | Cowbell | CrashCymbal2
  | Vibraslap | RideCymbal2 | HiBongo | LowBongo
  | MuteHiConga | OpenHiConga | LowConga | HighTimbale
  | LowTimbale | HighAgogo | LowAgogo | Cabasa
  | Maracas | ShortWhistle | LongWhistle | ShortGuiro
  | LongGuiro | Claves | HiWoodBlock | LowWoodBlock
  | MuteCuica | OpenCuica | MuteTriangle
  | OpenTriangle       -- MIDI Key 82
deriving (Show, Eq, Ord, Enum)

```

Figure 5.2: General MIDI Percussion Names

```

funkGroove
= let p1 = perc LowTom qn
    p2 = perc AcousticSnare en
in tempo 3 (instrument Percussion (cut 8 (repeatM
  ((p1 :+: qnr :+: p2 :+: qnr :+: p2 :+:
    p1 :+: p1 :+: qnr :+: p2 :+: enr)
    :=: roll en (perc ClosedHiHat 2))
  )))

```

Exercise 5.4 Define a function $\text{chrom} :: \text{Pitch} \rightarrow \text{Pitch} \rightarrow \text{Music Pitch}$ such that $\text{chrom } p1 \ p2$ is a chromatic scale of quarter-notes whose first pitch is $p1$ and last pitch is $p2$. If $p1 > p2$, the scale should be descending, otherwise it should be ascending. If $p1 \equiv p2$, then the scale should contain just one note. (A chromatic scale is one whose successive pitches are separated by one absolute pitch, i.e. one semitone).

Exercise 5.5 Abstractly, a scale can be described by the intervals between successive notes. For example, the 8-note major scale can be defined as the sequence of 7 intervals $[2, 2, 1, 2, 2, 2, 1]$, and the 12-note chromatic scale by the 11 intervals $[1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1]$. Define a function $\text{mkScale} :: \text{Pitch} \rightarrow [\text{Int}] \rightarrow \text{Music Pitch}$ such that $\text{mkScale } p \ \text{ints}$ is the scale beginning at pitch p and having the intervallic structure ints .

Exercise 5.6 Write the melody of "Frere Jacques" (or, "Are You Sleeping") in Haskore. Try to make it as succinct as possible. Then, using functions already defined, generate a four-part round, i.e. four identical voices, each delayed successively by two measures. Use a different instrument to realize each voice.

5.10 A Map for Music

Recall from Chapter 3 the definition of *map*:

$$\begin{aligned} \text{map} &:: (a \rightarrow b) \rightarrow [a] \rightarrow [b] \\ \text{map } f \ [] &= [] \\ \text{map } f \ (x : xs) &= f \ x : \text{map } f \ xs \end{aligned}$$

This function is defined on the list data type. Is there something analogous for *Music*? I.e. a function:

$$mMap :: (a \rightarrow b) \rightarrow \text{Music } a \rightarrow \text{Music } b$$

Such a function is indeed straightforward to define, but it helps to first define a map-like function for the *Primitive* type:

$$\begin{aligned} pMap &:: (a \rightarrow b) \rightarrow \text{Primitive } a \rightarrow \text{Primitive } b \\ pMap \ f \ (\text{Note } d \ x) &= \text{Note } d \ (f \ x) \\ pMap \ f \ (\text{Rest } d) &= \text{Rest } d \end{aligned}$$

With *pMap* in hand we can now define *mMap*:

$$\begin{aligned} mMap &:: (a \rightarrow b) \rightarrow \text{Music } a \rightarrow \text{Music } b \\ mMap \ f \ (\text{Primitive } x) &= \text{Primitive } (pMap \ f \ x) \\ mMap \ f \ (x \text{ :+ } y) &= mMap \ f \ x \text{ :+ } mMap \ f \ y \\ mMap \ f \ (x \text{ := } y) &= mMap \ f \ x \text{ := } mMap \ f \ y \\ mMap \ f \ (\text{Modify } c \ x) &= \text{Modify } c \ (mMap \ f \ x) \end{aligned}$$

Just as *map f xs* for lists replaces each polymorphic element *x* in *xs* with *f x*, *mMap f m* for *Music* replaces each polymorphic element *p* in *m* with *f p*.

As an example of how *mMap* can be used, suppose that we introduce a *Volume* type for a note simply as:

```
type Volume = Integer
```

and then wish to convert a value of type *Music Pitch* to a value of type *Music (Pitch, Volume)* – that is, we wish to pair each pitch with a volume attribute. We can define a function to do so as follows:

$$\begin{aligned} \text{addVolume} &:: \text{Volume} \rightarrow \text{Music Pitch} \rightarrow \text{Music (Pitch, Volume)} \\ \text{addVolume } v &= \text{mMap } (\lambda p \rightarrow (p, v)) \end{aligned}$$

Exercise 5.7 Using *mMap*, define a function:

$$\text{scaleVolume} :: \text{Rational} \rightarrow \text{Music (Pitch, Volume)} \rightarrow \text{Music (Pitch, Volume)}$$

such that *scaleVolume s m* scales the volume of each note in *m* by a factor of *s*.

5.11 A Fold for Music

We can also define a fold-like operator for *Music*. But whereas the list data type has only two constructors (the nullary constructor `[]` and the binary constructor `(:)`), *Music* has *four* constructors, and thus we define:

$$\begin{aligned} \text{mFold} &:: (b \rightarrow b \rightarrow b) \rightarrow (b \rightarrow b \rightarrow b) \rightarrow (\text{Primitive } a \rightarrow b) \\ &\rightarrow (\text{Control} \rightarrow b \rightarrow b) \rightarrow \text{Music } a \rightarrow b \\ \text{mFold } (+ :) (=:) f g m &= \\ \text{let } \text{rec} &= \text{mFold } (+ :) (=:) f g \\ \text{in case } m &\text{ of} \\ \text{Primitive } p &\rightarrow f p \\ m1 :+ m2 &\rightarrow \text{rec } m1 + : \text{rec } m2 \\ m1 := m2 &\rightarrow \text{rec } m1 := \text{rec } m2 \\ \text{Modify } c m &\rightarrow g c (\text{rec } m) \end{aligned}$$

This somewhat unwieldy function basically takes apart a *Music* value and puts it back together with different constructors. Indeed, note that:

$$\text{mFold } (:+ :) (=:) \text{Primitive Modify} \equiv \text{id}$$

Exercise 5.8 Prove the above property.

To see how *mFold* might be used, note first of all that it is more general than *mMap*—indeed, *mMap* can be defined in terms of *mFold* like this:

$$\begin{aligned} \text{mMap}' &:: (a \rightarrow b) \rightarrow \text{Music } a \rightarrow \text{Music } b \\ \text{mMap}' f &= \text{mFold } (:+ :) (=:) g \text{Modify where} \\ g (\text{Note } d x) &= \text{Primitive } (\text{Note } d (f x)) \\ g (\text{Rest } d) &= \text{Primitive } (\text{Rest } d) \end{aligned}$$

More interestingly, we can use *mFold* to redefine things like the *dur* function from Section 5.5:

```

dur' :: Music a → Dur
dur' = mFold (+) max getDur modDur where
  getDur (Note d _) = d
  getDur (Rest d) = d
  modDur (Tempo r) d = d / r
  modDur _ d = d

```

Exercise 5.9 Redefine *revM* from Section 5.6 using *mFold*.

Exercise 5.10 Define a function *insideOut* that inverts the role of serial and parallel composition in a *Music* value. Using *insideOut*, see if you can,
 (a) find a non-trivial value $m :: \text{Music Pitch}$ such that $m \equiv \text{insideOut } m$,
 (b) find a value $m :: \text{Music Pitch}$ such that:

$$m :+: \text{insideOut } m :+: m$$

sounds interesting. (You are free to define what “sounds interesting” means.)

Exercise 5.11 Find a simple piece of music written by your favorite composer, and transcribe it into Haskore. In doing so, look for repeating patterns, transposed phrases, etc. and reflect this in your code, thus revealing deeper structural aspects of the music than that found in common practice notation.

5.12 Crazy Recursion

With all the functions and data types that we have defined, and the power of recursion and higher-order functions at our finger tips, we can start to do some wild and crazy things. Here is just one such idea.

Let’s define a function to recursively apply transformations *f* (to elements in a sequence) and *g* (to accumulated phrases) some specified number of times:

```

rep :: (Music a → Music a) → (Music a → Music a) → Int
    → Music a → Music a
rep f g 0 m = rest 0
rep f g n m = m :=: g (rep f g (n - 1) (f m))

```

With this simple function we can create some interesting phrases of music with very little code. For example, let's use *rep* three times, nested together, to create a “cascade” of sounds.

```
run = rep (transpose 5) (delay tn) 8 (c 4 tn)
cascade = rep (transpose 4) (delay en) 8 run
cascades = rep id (delay sn) 2 cascade
```

and then make the cascade run up, and then down:

```
final = cascades :+: revM cascades
```

What happens if we reverse the *f* and *g* arguments?

```
run' = rep (delay tn) (transpose 5) 8 (c 4 tn)
cascade' = rep (delay en) (transpose 4) 8 run'
cascades' = rep (delay sn) id 2 cascade'
final' = cascades' :+: revM cascades'
```

Exercise 5.12 Do something wild and crazy with Haskore.

Chapter 6

Interpretation and Performance

```
module Haskore.Performance
  where

import Haskore.Music
import Haskore.MoreMusic

instance Show (a → b) where
  showsPrec p f = showString "<<function>>"
```

6.1 Abstract Performance

So far, our presentation of musical values in Haskell has been entirely structural, i.e. *syntactic*. But what do these musical values actually *mean*, i.e. what is their *semantics*, or *interpretation*? The formal process of giving a semantic interpretation to syntactic constructs is very common in computer science, especially in programming language theory. But it is obviously also common in music: the interpretation of music is the very essence of musical performance. However, in conventional music this process is usually informal, appealing to aesthetic judgments and values. What we would like to do is make the process formal in Haskore—but still flexible, so that more than one interpretation is possible, just as in music.

To begin, we need to say exactly what an abstract *performance* is. Our approach is to consider a performance to be a time-ordered sequence of musical *events*, where each event captures the playing of one individual

note. In Haskellese:

```
type Performance = [Event]

data Event = Event { eTime :: Time, eInst :: InstrumentName, ePitch :: AbsPitch,
                    eDur :: DurT, eVol :: Volume, pFields :: [Float] }
deriving (Eq, Ord, Show)

type Time = Rational
type DurT = Rational
type Volume = Integer
```

An event $Event\{eTime = s, eInst = i, ePitch = p, eDur = d, eVol = v\}$ captures the fact that at start time s , instrument i sounds pitch p with volume v for a duration d (where now duration is measured in seconds, rather than beats). (The $pField$ of an event is for special instruments that require extra parameters, and will not be discussed much further in this chapter.)

An abstract performance is the lowest of our music representations not yet committed to MIDI, csound, or some other low-level computer music representation. In Chapter ?? we will discuss how to map a performance into MIDI.

Details: The data declaration for *Event* uses Haskell's *field label* syntax, also called *record* syntax, and is equivalent to:

```
data Event = Event Time InstrumentName AbsPitch DurT Volume [Float]
deriving (Eq, Ord, Show)
```

except that the former also defines “field labels” $eTime$, $eInst$, $ePitch$, $eDur$, $eVol$, and $pFields$, which can be used both to create and select from *Event* values. For example, this equation:

$$e = Event\ 0\ Cello\ 27\ (1 / 4)\ 50\ []$$

is equivalent to:

$$e = Event\{eTime = 0, ePitch = 27, eDur = 1 / 4, \\ eInst = Cello, eVol = 50, pFields = []\}$$

The latter is more descriptive, however, and the order of the fields does not matter (indeed the order here is not the same as above).

Field labels can be used to select fields from an *Event* value; for example, $eInst\ e \Rightarrow Cello$, $eDur\ e \Rightarrow 1 / 4$, and so on. They can also be used to selectively *update* fields of an existing *Event* value. For example:

$$e\{eInst = Flute\} \Rightarrow Event\ 0\ Flute\ 27\ (1 / 4)\ 50\ []$$

Finally, they can be used selectively in pattern matching:

$$f\ (Event\{eDur = d, ePitch = p\}) = \dots d \dots p \dots$$

Field labels do not change the basic nature of a data type; they are simply a convenient syntax for referring to the components of a data type by name rather than by position.

To generate a complete performance of, i.e. give an interpretation to, a musical value, we must know the time to begin the performance, and the proper instrument, volume, key and tempo. In addition, to give flexibility to our interpretations, we must also know what *player* to use; that is, we need a mapping from the *PlayerNames* in a *Music* value to the actual players to be used.¹ We capture these ideas in Haskell as a “context” and “player map,” respectively:

```
data Context a = Context { cTime :: Time, cPlayer :: Player a,
                           cInst :: InstrumentName, cDur :: DurT,
                           cKey :: Key, cVol :: Volume }

deriving Show
type PMap a = PlayerName → Player a
type Key = AbsPitch
```

Finally, we are ready to give an interpretation to a piece of music, which we do by defining a function *perform*, which is conceptually perhaps the most important function defined in this book, and is shown in Figure 6.1.

Some things to note about *perform*:

1. The *Context* is the running “state” of the performance, and gets updated in several different ways. For example, the interpretation of the *Tempo* constructor involves scaling *dt* appropriately and updating the *DurT* field of the context.
2. The interpretation of notes and phrases is player dependent. Ultimately a single note is played by the *playNote* function, which takes the player as an argument. Similarly, phrase interpretation is also player dependent, reflected in the use of *interpPhrase*. Precisely how these two functions work is described in Section 6.2.

¹We don’t need a mapping from *InstrumentNames* to instruments, since this is handled in the translation from a performance into MIDI, which is discussed in Chapter ??.

```

perform :: PMap a → Context a → Music a → Performance
perform pmap
  c@Context{ cTime = t, cPlayer = pl, cDur = dt, cKey = k } m =
  case m of
    Primitive (Note d p) → playNote pl c d p
    Primitive (Rest d) → []
    m1 :+: m2 → perform pmap c m1 ++
      perform pmap (c{ cTime = t + dur m1 * dt }) m2
    m1 :=: m2 → merge (perform pmap c m1) (perform pmap c m2)
    Modify (Tempo r) m → perform pmap (c{ cDur = dt / r }) m
    Modify (Transpose p) m → perform pmap (c{ cKey = k + p }) m
    Modify (Instrument i) m → perform pmap (c{ cInst = i }) m
    Modify (Player pn) m → perform pmap (c{ cPlayer = pmap pn }) m
    Modify (Phrase pa) m → interpPhrase pl pmap c pa m

```

Figure 6.1: An abstract *perform* function

3. The *DurT* component of the context is the duration, in seconds, of one whole note. To make it easier to compute, we can define a “metronome” function that, given a standard metronome marking (in beats per minute) and the note type associated with one beat (quarter note, eighth note, etc.) generates the duration of one whole note:

```

metro :: Int → Dur → DurT
metro setting dur = 60 / (fromIntegral setting * dur)

```

Thus, for example, *metro* 96 *qn* creates a tempo of 96 quarter notes per minute.

4. In the treatment of *(:+:)*, note that the sub-sequences are appended together, with the start time of the second argument delayed by the duration of the first. The function *dur* (defined in Section 5.5) is used to compute this duration. However, this results in a quadratic time complexity for *perform*. A more efficient solution is to have *perform* compute the duration directly, returning it as part of its result. This version of *perform* is shown in Figure ??.
5. The sub-sequences derived from the arguments to *(:=:)* are merged into a time-ordered stream. The definition of *merge* is given below.

merge :: *Performance* → *Performance* → *Performance*

```

merge a@(e1 : es1) b@(e2 : es2) =
    if e1 < e2 then e1 : merge es1 b
    else e2 : merge a es2
merge [] es2 = es2
merge es1 [] = es1

```

Note that *merge* compares entire events rather than just start times. This is to ensure that it is commutative, a desirable condition for some of the proofs used later in the text. Here is a more efficient version that will work just as well in practice:

```

merge a@(e1 : es1) b@(e2 : es2) =
    if eTime e1 < eTime e2 then e1 : merge es1 b
    else e2 : merge a es2
merge [] es2 = es2
merge es1 [] = es1

```

6.2 Players

Recall from Section 2.2 the definition of the *Control* data type:

```

data Control =
    Tempo Rational      -- scale the tempo
  | Transpose AbsPitch  -- transposition
  | Instrument InstrumentName -- instrument label
  | Phrase [PhraseAttribute] -- phrase attributes
  | Player PlayerName   -- player label
deriving (Show, Eq, Ord)

```

```

type PlayerName = String

```

We mentioned, but did not define, the *PhraseAttribute* data type, shown now fully in Figure 6.3. These attributes give us great flexibility in the interpretation process, because they can be interpreted by different players in different ways. For example, how should “legato” be interpreted in a performance? Or “diminuendo?” Different players interpret things in different ways, of course, but even more fundamental is the fact that a pianist, for example, realizes legato in a way fundamentally different from the way a

```

perform :: PMap a → Context a → Music a → Performance
perform pmap c m = fst (perf pmap c m)

perf :: PMap a → Context a → Music a → (Performance, DurT)
perf pmap
  c@Context{ cTime = t, cPlayer = pl, cDur = dt, cKey = k } m =
  case m of
    Primitive (Note d p) → (playNote pl c d p, d * dt)
    Primitive (Rest d) → ([], d * dt)
    m1 :+: m2 → let (pf1, d1) = perf pmap c m1
                  (pf2, d2) = perf pmap (c{ cTime = t + d1 }) m2
                  in (pf1 ++ pf2, d1 + d2)
    m1 :=: m2 → let (pf1, d1) = perf pmap c m1
                  (pf2, d2) = perf pmap c m2
                  in (merge pf1 pf2, max d1 d2)
    Modify (Tempo r) m → perf pmap (c{ cDur = dt / r }) m
    Modify (Transpose p) m → perf pmap (c{ cKey = k + p }) m
    Modify (Instrument i) m → perf pmap (c{ cInst = i }) m
    Modify (Player pn) m → perf pmap (c{ cPlayer = pmap pn }) m
    Modify (Phrase pas) m → interpPhrase pl pmap c pas m

```

Figure 6.2: The “real” *perform* function

```

data PhraseAttribute = Dyn Dynamic
                        | Tmp Tempo
                        | Art Articulation
                        | Orn Ornament
deriving (Eq, Ord, Show)

data Dynamic = Accent Rational | Crescendo Rational | Diminuendo Rational
              | StdLoudness StdLoudness | Loudness Rational
deriving (Eq, Ord, Show)

data StdLoudness = PPP | PP | P | MP | SF | MF | NF | FF | FFF
deriving (Eq, Ord, Show, Enum)

data Tempo = Ritardando Rational | Accelerando Rational
deriving (Eq, Ord, Show)

data Articulation = Staccato Rational | Legato Rational | Slurred Rational
                   | Tenuto | Marcato | Pedal | Fermata | FermataDown | Breath
                   | DownBow | UpBow | Harmonic | Pizzicato | LeftPizz
                   | BartokPizz | Swell | Wedge | Thumb | Stopped
deriving (Eq, Ord, Show)

data Ornament = Trill | Mordent | InvMordent | DoubleMordent
               | Turn | TrilledTurn | ShortTrill
               | Arpeggio | ArpeggioUp | ArpeggioDown
               | Instruction String | Head NoteHead
deriving (Eq, Ord, Show)

data NoteHead = DiamondHead | SquareHead | XHead | TriangleHead
                | TremoloHead | SlashHead | ArtHarmonic | NoHead
deriving (Eq, Ord, Show)

```

Figure 6.3: Phrase Attributes

violinist does, because of differences in their instruments. Similarly, *diminuendo* on a piano and *diminuendo* on a harpsichord are different concepts.

With a slight stretch of the imagination, we can even consider a “notator” of a score as a kind of player: exactly how the music is rendered on the written page may be a personal, stylized process. For example, how many, and which staves should be used to notate a particular instrument?

In any case, to handle these issues, Haskore has a notion of a *player* that “knows” about differences with respect to performance and notation. A Haskore player is a 4-tuple consisting of a name and three functions: one for interpreting notes, one for phrases, and one for producing a properly notated score.

```
data Player a = MkPlayer{ pName :: PlayerName,
                          playNote :: NoteFun a,
                          interpPhrase :: PhraseFun a,
                          notatePlayer :: NotateFun a }

deriving Show

type NoteFun a = Context a → Dur → a → Performance
type PhraseFun a = PMap a → Context a → [PhraseAttribute]
                               → Music a → (Performance, DurT)
type NotateFun a = ()
```

The last line above is because notation is currently not implemented.

6.2.1 Examples of Player Construction

In order to provide the most flexibility, we define attributes for individual notes:

```
data NoteAttribute = Volume Integer    -- by convention: 0=min, 100=max
                  | Fingering Integer
                  | Dynamics String
                  | PFields [Float]

deriving (Eq, Show)
```

Our goal then is to define a player for music values of type:

```
type Note1 = (Pitch, [NoteAttribute])
type Music1 = Music Note1
```

A “default player” called *defPlayer* (not to be confused with a “deaf player”!) is defined for use when none other is specified in the score; it also functions

```

defPlayer :: Player (Pitch, [NoteAttribute])
defPlayer = MkPlayer{ pName = "Default",
                      playNote = defPlayNote defNasHandler,
                      interpPhrase = defInterpPhrase defPasHandler,
                      notatePlayer = defNotatePlayer () }

defPlayNote :: (Context (Pitch, [a]) → a → Event → Event)
              → NoteFun (Pitch, [a])
defPlayNote nasHandler
  c@(Context cTime cPlayer cInst cDur cKey cVol) d (p, nas) =
    [foldr (nasHandler c)
      (Event{ eTime = cTime, eInst = cInst,
              ePitch = absPitch p + cKey,
              eDur = d * cDur, eVol = cVol,
              pFields = [] })
     nas]

defNasHandler :: Context a → NoteAttribute → Event → Event
defNasHandler c (Volume v) ev = ev{ eVol = v }
defNasHandler c (PFields pfs) ev = ev{ pFields = pfs }
defNasHandler _ _ ev = ev

defInterpPhrase :: (PhraseAttribute → Performance → Performance)
                 → PhraseFun a
defInterpPhrase pasHandler pmap context pas m =
  let (pf, dur) = perf pmap context m
  in (foldr pasHandler pf pas, dur)

defPasHandler :: PhraseAttribute → Performance → Performance
defPasHandler (Dyn (Accent x)) =
  map (λe → e{ eVol = round (x * fromIntegral (eVol e)) })
defPasHandler (Art (Staccato x)) = map (λe → e{ eDur = x * eDur e })
defPasHandler (Art (Legato x)) = map (λe → e{ eDur = x * eDur e })
defPasHandler _ = id

defNotatePlayer :: a → ()
defNotatePlayer _ = ()

```

Figure 6.4: Definition of default Player *defPlayer*.

as a base from which other players can be derived. *defPlayer* responds only to the *Volume* note attribute and to the *Accent*, *Staccato*, and *Legato* phrase attributes. It is defined in Figure 6.4. Before reading this code, recall how players are invoked by the *perform* function defined in the last section; in particular, note the calls to *playNote* and *interpPhrase*. Then note:

1. *defPlayNote* is the only function (even in the definition of *perform*) that actually generates an event. It also modifies that event based on an interpretation of each note attribute by the function *defHasHandler*.
2. *defNasHandler* only recognizes the *Volume* attribute, which it uses to set the event volume accordingly.
3. *defInterpPhrase* calls (mutually recursively) *perform* to interpret a phrase, and then modifies the result based on an interpretation of each phrase attribute by the function *defPasHandler*.
4. *defPasHandler* only recognizes the *Accent*, *Staccato*, and *Legato* phrase attributes. For each of these it uses the numeric argument as a “scaling” factor of the volume (for *Accent*) and duration (for *Staccato* and *Lagato*). Thus *Modify (Phrase [Legato (5 / 4)]) m* effectively increases the duration of each note in *m* by 25% (without changing the tempo).

It should be clear that much of the code in Figure 6.4 can be re-used in defining a new player. For example, to define a player *weird* that interprets note attributes just like *defPlayer* but behaves differently with respect to phrase attributes, we could write:

```
weird :: Player (Pitch, [NoteAttribute])
weird = MkPlayer { pName = "newPlayer",
                  playNote = defPlayNote defNasHandler,
                  interpPhrase = defInterpPhrase myPasHandler,
                  notatePlayer = defNotatePlayer () }
```

and then supply a suitable definition of *myPasHandler*. That definition could also re-use code, in the following sense: suppose we wish to add an interpretation for *Crescendo*, but otherwise have *myPasHandler* behave just like *defPasHandler*.

```
myPasHandler :: PhraseAttribute → Performance → Performance
myPasHandler (Dyn (Crescendo x)) pf = ...
myPasHandler pa pf = defPasHandler pa pf
```

Exercise 6.1 Fill in the ... in the definition of *myPasHandler* according to the following strategy: Gradually scale the volume of each event in the performance by a factor of 1 through $1 + x$, using linear interpolation.

Exercise 6.2 Choose some of the other phrase attributes and provide interpretations of them, such as *Diminuendo*, *Slurred*, *Trill*, etc. (The *trill* functions from section 5.8 may be useful here.)

Figure 6.5 defines a relatively sophisticated player called *fancyPlayer* that knows all that *defPlayer* knows, and much more. Note that *Slurred* is different from *Legato* in that it doesn't extend the duration of the *last* note(s). The behavior of *Ritardando* x can be explained as follows. We'd like to "stretch" the time of each event by a factor from 0 to x , linearly interpolated based on how far along the musical phrase the event occurs. I.e., given a start time t_0 for the first event in the phrase, total phrase duration D , and event time t , the new event time t' is given by:

$$t' = (1 + \frac{t - t_0}{D}x)(t - t_0) + t_0$$

Further, if d is the duration of the event, then the end of the event $t + d$ gets stretched to a new time t'_d given by:

$$t'_d = (1 + \frac{t + d - t_0}{D}x)(t + d - t_0) + t_0$$

The difference $t'_d - t'$ gives us the new, stretched duration d' , which after simplification is:

$$d' = (1 + \frac{2(t - t_0) + d}{D}x)d$$

Accelerando behaves in exactly the same way, except that it shortens event times rather than lengthening them. And, a similar but simpler strategy explains the behaviors of *Crescendo* and *Diminuendo*.

```

fancyPlayer :: Player (Pitch, [NoteAttribute])
fancyPlayer = MkPlayer{ pName = "Fancy",
                        playNote = defPlayNote defNasHandler,
                        interpPhrase = fancyInterpPhrase,
                        notatePlayer = defNotatePlayer ()}

fancyInterpPhrase :: PhraseFun a
fancyInterpPhrase pmap c [] m = perf pmap c m
fancyInterpPhrase pmap c@Context{ cTime = t, cPlayer = pl, cInst = i,
                                cDur = dt, cKey = k, cVol = v }
    (pa : pas) m =
let pfd@(pf, dur) = fancyInterpPhrase pmap c pas m
    loud x = fancyInterpPhrase pmap c (Dyn (Loudness x) : pas) m
    stretch x = let t0 = eTime (head pf); r = x / dur
                 upd (e@Event{ eTime = t, eDur = d }) =
                     let dt = t - t0
                         t' = (1 + dt * r) * dt + t0
                         d' = (1 + (2 * dt + d) * r) * d
                     in e{ eTime = t', eDur = d' }
                 in (map upd pf, (1 + x) * dur)
    inflate x = let t0 = eTime (head pf); r = x / dur
                upd (e@Event{ eTime = t, eVol = v }) =
                    e{ eVol = (round o fromRational) (1 + (t - t0) * r) * v }
                in (map upd pf, dur)
in case pa of
    Dyn (Accent x) → (map (λe → e{ eVol = round (x * fromIntegral (eVol e)) }) pf, dur)
    Dyn (StdLoudness l) →
        case l of
            PPP → loud 40; PP → loud 50; P → loud 60
            MP → loud 70; SF → loud 80; MF → loud 90
            NF → loud 100; FF → loud 110; FFF → loud 120
    Dyn (Loudness x) → fancyInterpPhrase pmap c{ cVol = (round o fromRational) x } pas m
    Dyn (Crescendo x) → inflate x; Dyn (Diminuendo x) → inflate (-x)
    Tmp (Ritardando x) → stretch x; Tmp (Accelerando x) → stretch (-x)
    Art (Staccato x) → (map (λe → e{ eDur = x * eDur e }) pf, dur)
    Art (Legato x) → (map (λe → e{ eDur = x * eDur e }) pf, dur)
    Art (Slurred x) →
        let lastStartTime = foldr (λe t → max (eTime e) t) 0 pf
            setDur e = if eTime e < lastStartTime
                        then e{ eDur = x * eDur e }
                        else e
        in (map setDur pf, dur)
    Art _ → pfd
    Orn _ → pfd
-- Design Bug: To do these right we need to keep the KEY SIGNATURE
-- around so that we can determine, for example, what the trill note is.
-- Alternatively, provide an argument to Trill to carry this info.

```

Figure 6.5: Definition of Player *fancyPlayer*.

Chapter 7

Self-Similar Music

```
module Haskore.SelfSimilar where  
import Haskore
```

In this chapter we will explore the notion of *self-similar* music—i.e. musical structures that have patterns that repeat themselves recursively in interesting ways. There are many approaches to generating self-similar structures, the most general being *fractals*, which have been used to generate not just music, but also graphical images. We will delay a general treatment of fractals, however, and will instead focus on more specialized notions of self-similarity, notions that we conceive of musically, and then manifest as Haskell programs.

7.1 Self-Similar Melody

Here is the first notion of self-similar music that we will consider: Begin with a very simple melody of n notes. Now duplicate this melody n times, playing each in succession, but first perform the following transformations: transpose the i th melody by an amount proportional to the pitch of the i th note in the original melody, and scale its tempo by a factor proportional to the duration of the i th note. For example, Figure 7.1 shows the result of applying this process once to a four-note melody. Now imagine that this process is repeated infinitely often. For a melody whose notes are all shorter than a whole note, it yields an infinitely dense melody of infinitesimally shorter notes. To make the result playable, however, we will stop the process at some pre-determined level.

How can this be represented in Haskell? A *tree* seems like it would be a logical choice; let's call it a *Cluster*:



Figure 7.1: An Example of Self-Similar Music

```
data Cluster = Cluster SNote [Cluster]
type SNote = (Dur, AbsPitch)
```

This particular kind of tree happens to be called a *rose tree* []. An *SNote* is just a “simple note,” a duration paired with an absolute pitch. We prefer to stick with absolute pitches in creating the self-similar structure, and will convert the result into a normal *Music* value only after we are done.

The sequence of *SNotes* at each level of the cluster is the melodic fragment for that level. The very top cluster will contain a “dummy” note, whereas the next level will contain the original melody, the next level will contain one iteration of the process described above (e.g. the melody in Figure 7.1), and so forth.

To achieve this we will define a function *selfSim* that takes the initial melody as argument and generates an infinitely deep cluster:

```
selfSim :: [SNote] → Cluster
selfSim pat = Cluster (0,0) (map mkCluster pat)
  where mkCluster note
        = Cluster note (map (mkCluster ∘ addmult note) pat)

addmult (d0, p0) (d1, p1) = (d0 * d1, p0 + p1)
```

Note that *selfSim* itself is not recursive, but *mkCluster* is.

Next, we define a function to skim off the notes at the *n*th level, or *n*th “fringe,” of a cluster:

```
fringe :: Int → Cluster → [SNote]
```

```

fringe 0 (Cluster note cls) = [note]
fringe n (Cluster note cls) = concatMap (fringe (n - 1)) cls

```

Details: *concatMap* is defined in the Standard Prelude as:

```

concatMap :: (a -> [b]) -> [a] -> [b]
concatMap f = concat o map f

```

Also recall that *concat* appends together a list of lists, and is defined in the Prelude as:

```

concat :: [[a]] -> [a]
concat = foldr (++) []

```

All that is left to do is convert this into a *Music* value that we can play:

```

simToMusic :: [SNote] -> Music Pitch
simToMusic ss = let mkNote (d, ap) = note d (pitch ap)
                 in line (map mkNote ss)

```

We can define this with a bit more elegance as follows:

```

simToMusic :: [SNote] -> Music Pitch
simToMusic = line o map mkNote

mkNote :: (Dur, AbsPitch) -> Music Pitch
mkNote (d, ap) = note d (pitch ap)

```

The increased modularity will allow us to reuse *mkNote* later in the chapter.

Putting it all together, we can define a function that takes an initial pattern, a level, a number of pitches to transpose the result, and a tempo scaling factor, to yield a final result:

```

ss pat n tr te = transpose tr (tempo te (simToMusic (fringe n (selfSim pat))))

```

Here are some example compositions:

```

p1 :: [SNote]
p1 = [(hn, 3), (qn, 4), (qn, 0), (wn, 6)]

ss1 = ss p1 4 50 (3 / 2)

ss1a = let l1 = instrument Flute ss1
        l2 = instrument AcousticBass (transpose (-12) (revM ss1))

```

in $l1 :=: l2$

-- Note that the flute and bass lines are the reverse of one another.

$p2 = [(dqn, 0), (qn, 4)]$
 $p3 = [(6 / 10, 2), (13 / 10, 5), (wn, 0), (9 / 10, 7)]$
 $p4 = [(hn, 3), (hn, 8), (hn, 22), (qn, 4), (qn, 7), (qn, 21),$
 $(qn, 0), (qn, 5), (qn, 15), (wn, 6), (wn, 9), (wn, 19)]$

$ss2 = ss\ p2\ 6\ 50\ (1 / 30)$
 $ss3 = ss\ p3\ 4\ 50\ 20$
 $ss4 = ss\ p4\ 3\ 50\ 8$

Exercise 7.1 Experiment with this idea further, using other melodic seeds, exploring different depths of the clusters, and so on.

Exercise 7.2 Note that *concat* is defined as *foldr* $(++)$ $[]$, which means that it takes a number of steps proportional to the sum of the lengths of the lists being concatenated; we cannot do any better than this. (If *foldl* were used instead, the number of steps would be proportional to the number of lists times their average length.)

However, *fringe* is not very efficient, for the following reason: *concat* is being used over and over again, like this:

$concat\ [concat\ [...],\ concat\ [...],\ concat\ [...]]$

This causes a number of steps proportional to the depth of the tree times the length of the sub-lists; clearly not optimal.

Define a version of *fringe* that is linear in the total length of the final list.

7.2 Self-Similar Harmony

In the last section we used a melody as a seed, and created longer melodies from it. Another idea is to stack the melodies vertically. Specifically, suppose we redefine *fringe* in such a way that it does not concatenate the sub-clusters together:

$fringe' :: Int \rightarrow Cluster \rightarrow [[SNote]]$
 $fringe'\ 0\ (Cluster\ note\ cls) = [[note]]$
 $fringe'\ n\ (Cluster\ note\ cls) = map\ (fringe'\ (n - 1))\ cls$

Note that this strategy is only applied to the top level—below that we use fringe. Thus the type of the result is $[[S\text{Note}]]$, i.e. a list of lists of notes.

We can convert the individual lists into melodies, and play the melodies all together, like this:

$$\begin{aligned} \text{simToMusic}' &:: [[S\text{Note}]] \rightarrow \text{Music Pitch} \\ \text{simToMusic}' &= \text{chord} \circ \text{map} (\text{line} \circ \text{map} \text{mkNote}) \end{aligned}$$

Finally, we can define a function akin to *ss* defined earlier:

$$\text{ss}' \text{ pat } n \text{ tr } te = \text{transpose } tr (\text{tempo } te (\text{simToMusic}' (\text{fringe}' n (\text{selfSim } \text{pat}))))$$

Using the same patterns as used earlier, here are some sample compositions:

$$\begin{aligned} \text{ss}1' &= \text{ss}' \text{ p1 } 4 \text{ 50 } (3 / 2) \\ \text{ss}2' &= \text{ss}' \text{ p2 } 4 \text{ 50 } 4 \\ \text{ss}3' &= \text{ss}' \text{ p3 } 4 \text{ 50 } 20 \\ \text{ss}4' &= \text{ss}' \text{ p4 } 3 \text{ 50 } 8 \end{aligned}$$

And a new one, based on a major triad:

$$\begin{aligned} \text{ss}5 &= \text{ss } \text{p5 } 4 \text{ 45 } (1 / 500) \\ \text{ss}5' &= \text{ss}' \text{ p5 } 4 \text{ 45 } (1 / 500) \\ \text{p5} &= [(en, 4), (sn, 7), (en, 0)] \end{aligned}$$

Note the need to scale the tempo back drastically, due to the short durations of the starting notes.

7.3 Other Self-Similar Structures

The reader will observe that our notion of “self-similar harmony” did *not* involve changing the structure of the *Cluster* data type, nor the algorithm for computing the sub-structures (as captured in *selfSim*). All that we did was interpret the result differently. This is a common characteristic of algorithmic music composition—the same mathematical or computational structure is interpreted in different ways to yield musically different results.

For example, instead of the above strategy for playing melodies in parallel, we could play entire levels of the *Cluster* in parallel, where the number of levels that we choose is given as a parameter. We leave this idea and others as exercises for the reader.

Exercise 7.3 Devise a version of *simToMusic* that constructs a *Music* value as outlined above. Specifically, given a parameter *n*, *simToMusic n pat* plays the first *n* levels of the cluster generated by *pat* in parallel.

Exercise 7.4 Devise some other variant of self-similar music, and encode it in Haskell. In particular, consider structures that are different from those generated by the *selfSim* function.

Chapter 8

Proof by Induction

In this chapter we will study a powerful proof technique based on *mathematical induction*. With it we will be able to prove complex and important properties of programs that cannot be accomplished with proof-by-calculation alone. The inductive proof method is one of the most powerful and common methods for proving program properties.

8.1 Induction and Recursion

Induction is very closely related to *recursion*. In fact, in certain contexts the terms are used interchangeably; in others, one is preferred over the other primarily for historical reasons. Think of them as being duals of one another: induction is used to describe the process of starting with something small and simple, and building up from there, whereas recursion describes the process of starting with something large and complex, and working backward to the simplest case.

For example, although we have previously used the phrase *recursive data type*, in fact data types are often described *inductively*, such as a list:

A *list* is either empty, or it is a pair consisting of a value and another list.

On the other hand, we usually describe functions that manipulate lists, such as *map* and *foldr*, as being recursive. This is because when you apply a function such as *map*, you apply it initially to the whole list, and work backwards toward [].

But these differences between induction and recursion run no deeper: they are really just two sides of the same coin.

This chapter is about *inductive properties* of programs (but based on the above argument could just as rightly be called *recursive properties*) that are not usually proven via calculation alone. Proving inductive properties usually involves the inductive nature of data types and the recursive nature of functions defined on the data types.

As an example, suppose that p is an inductive property of a list. In other words, $p(l)$ for some list l is either true or false (no middle ground!). To prove this property inductively, we do so based on the length of the list: starting with length 0, we first prove $p([])$ (using our standard method of proof-by-calculation).

Now for the key step: assume for the moment that $p(xs)$ is true for any list xs whose length is less than or equal to n . Then if we can prove (via calculation) that $p(x : xs)$ is true for any x —i.e. that p is true for lists of length $n + 1$ —then the claim is that p is true for lists of *any* (finite) length.

Why is this so? Well, from the first step above we know that p is true for length 0, so the second step tells us that it's also true for length 1. But if it's true for length 1 then it must also be true for length 2; similarly for lengths 3, 4, etc. So p is true for lists of any length!

(It is important to realize, however, that a property being true for every finite list does not necessarily imply that it is true for every infinite list. The property “the list is finite” is a perfect example of this! We will see how to prove properties of infinite lists in Chapter ??.)

To summarize, to prove a property p by induction on the length of a list, we proceed in two steps:

1. Prove $p([])$ (this is called the *base step*).
2. Assume that $p(xs)$ is true (this is called the *induction hypothesis*, and prove that $p(x : xs)$ is true (this is called the *induction step*).

8.2 Examples of List Induction

Ok, enough talk, let's see this idea in action. Recall in Section 3.1 the following property about *foldr*:

$$(\forall xs) \text{ foldr } (:) [] xs \Longrightarrow xs$$

We will prove this by induction on the length of xs . Following the ideas above, we begin with the base step by proving the property for length 0; i.e. for $xs = []$:

$$\text{foldr } (:) [] [] \Rightarrow []$$

This step is immediate from the definition of *foldr*. Now for the induction step: we first *assume* that the property is true for all lists *xs* of length *n*, and then prove the property for list *x : xs*. Again proceeding by calculation:

$$\begin{aligned} & \text{foldr } (:) [] (x : xs) \\ & \Rightarrow x : \text{foldr } (:) [] xs \\ & \Rightarrow x : xs \end{aligned}$$

And we are done; the induction hypothesis is what justifies the second step.

Now let's do something a bit harder. Suppose we are interested in proving the following property:

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

Our first problem is to decide which list to perform the induction over. A little thought (in particular, a look at how the definitions of *length* and *(++)* are structured) should convince you that *xs* is the right choice. (If you do not see this, you are encouraged to try the proof by induction over the length of *ys*!) Again following the ideas above, we begin with the base step by proving the property for length 0; i.e. for *xs* = []:

$$\begin{aligned} & \text{length } ([] ++ ys) \\ & \Rightarrow \text{length } ys \\ & \Rightarrow 0 + \text{length } ys \\ & \Rightarrow \text{length } [] + \text{length } ys \end{aligned}$$

For the induction step, we first assume that the property is true for all lists *xs* of length *n*, and then prove the property for list *x : xs*. Again proceeding by calculation:

$$\begin{aligned} & \text{length } ((x : xs) ++ ys) \\ & \Rightarrow \text{length } (x : (xs ++ ys)) \\ & \Rightarrow 1 + \text{length } (xs ++ ys) \\ & \Rightarrow 1 + (\text{length } xs + \text{length } ys) \\ & \Rightarrow (1 + \text{length } xs) + \text{length } ys \\ & \Rightarrow \text{length } (x : xs) + \text{length } ys \end{aligned}$$

And we are done. The transition from the 3rd line to the 4th is where we used the induction hypothesis.

8.3 Proving Function Equivalences

At this point it is a simple matter to return to Chapter 3 and supply the proofs that functions defined using *map* and *fold* are equivalent to the recursive versions. In particular, let's prove first that:

$$\text{toAbsPitches } ps = \text{map } \text{absPitch } ps$$

for any finite list ps , where:

$$\begin{aligned} \text{toAbsPitches } [] &= [] \\ \text{toAbsPitches } (p : ps) &= \text{absPitch } p : \text{toAbsPitches } ps \end{aligned}$$

We proceed by induction, starting with the base case $ps = []$:

$$\begin{aligned} \text{toAbsPitches } [] \\ \Rightarrow [] \\ \Rightarrow \text{map } \text{absPitch } [] \end{aligned}$$

Next we assume that $\text{toAbsPitches } ps = \text{map } \text{absPitch } ps$ holds, and try to prove that $\text{toAbsPitches } (p : ps) = \text{map } \text{absPitch } (p : ps)$ (note the use of the induction hypothesis in the second step):

$$\begin{aligned} \text{toAbsPitches } (p : ps) \\ \Rightarrow \text{absPitch } p : \text{toAbsPitches } ps \\ \Rightarrow \text{absPitch } p : \text{map } \text{absPitch } ps \\ \Rightarrow \text{map } \text{absPitch } (p : ps) \end{aligned}$$

The proof that $\text{toPitches } aps = \text{map } \text{pitch } aps$ is very similar, and is left as an exercise.

For a proof involving *foldr*, recall from Section 3.4 this recursive definition of *line*:

$$\begin{aligned} \text{line } [] &= \text{rest } 0 \\ \text{line } (m : ms) &= m :+ \text{line } ms \end{aligned}$$

and this non-recursive version:

$$\text{line} = \text{foldr } (:+) (\text{rest } 0)$$

We can prove that these definitions are equivalent by induction. First the base case:

$$\begin{aligned} \text{line } [] \\ \Rightarrow \text{rest } 0 \\ \Rightarrow \text{foldr } (:+) (\text{rest } 0) [] \end{aligned}$$

Then the induction step:

$$\begin{aligned}
 & \text{line } (m : ms) \\
 \Rightarrow & m :+ \text{line } ms \\
 \Rightarrow & m :+ \text{foldr } (:+) \text{ (rest 0) } ms \\
 \Rightarrow & \text{foldr } (:+) \text{ (rest 0) } (m : ms)
 \end{aligned}$$

The proofs of equivalence of the definitions of *chord*, *maxPitch*, and *listSum* from Section 3.4 are similar, and left as exercises.

These proofs were in fact quite easy. For something more challenging, consider the definition of *reverse* given in Section 3.5:

$$\begin{aligned}
 \text{reverse1 } [] &= [] \\
 \text{reverse1 } (x : xs) &= \text{reverse1 } xs ++ [x]
 \end{aligned}$$

and the version given in Section 4.1:

$$\text{reverse2 } xs = \text{foldl } (\text{flip } (:)) [] xs$$

We would like to show that these are the same; i.e. that $\text{reverse1 } xs = \text{reverse2 } xs$ for any finite list xs . In carrying out this proof two new ideas will be demonstrated, the first being that induction can be used to prove the equivalence of two programs. The second is the need for an *auxiliary property* which is proved independently of the main result.

The base case is easy, as it often is:

$$\begin{aligned}
 & \text{reverse1 } [] \\
 \Rightarrow & [] \\
 \Rightarrow & \text{foldl } (\text{flip } (:)) [] [] \\
 \Rightarrow & \text{reverse2 } []
 \end{aligned}$$

Assume now that $\text{reverse1 } xs = \text{reverse2 } xs$. The induction step proceeds as follows:

$$\begin{aligned}
 & \text{reverse1 } (x : xs) \\
 \Rightarrow & \text{reverse1 } xs ++ [x] \\
 \Rightarrow & \text{reverse2 } xs ++ [x] \\
 \Rightarrow & \text{foldl } (\text{flip } (:)) [] xs ++ [x] \\
 \Rightarrow & ???
 \end{aligned}$$

But now what do we do? Intuitively, it seems that the following property, which we will call property (1), should hold:

$$\begin{aligned}
 & \text{foldl } (\text{flip } (:)) [] xs ++ [x] \\
 \Rightarrow & \text{foldl } (\text{flip } (:)) [] (x : xs)
 \end{aligned}$$

in which case we could complete the proof as follows:

$$\begin{aligned}
 & \dots \\
 & \Rightarrow \text{foldl } (\text{flip } (:)) [] \text{ } xs ++ [x] \\
 & \Rightarrow \text{foldl } (\text{flip } (:)) [] (x : xs) \\
 & \Rightarrow \text{reverse2 } (x : xs)
 \end{aligned}$$

The ability to see that if we could just prove one thing, then perhaps we could prove another, is a useful skill in conducting proofs. In this case we have reduced the overall problem to one of proving property (1), which simplifies the structure of the proof, although not necessarily the difficulty. These auxiliary properties are often called *lemmas* in mathematics, and in many cases their proofs become the most important contributions, since they are often at the heart of a problem.

In fact if you try to prove property (1) directly, you will run into a problem, namely that it is not *general* enough. So first let's generalize property (1) (while renaming x to y), as follows:

$$\begin{aligned}
 & \text{foldl } (\text{flip } (:)) \text{ } ys \text{ } xs ++ [y] \\
 & \Rightarrow \text{foldl } (\text{flip } (:)) (ys ++ [y]) \text{ } xs
 \end{aligned}$$

Let's call this property (2). If (2) is true for any finite xs and ys , then property (1) is also true, because:

$$\begin{aligned}
 & \text{foldl } (\text{flip } (:)) [] \text{ } xs ++ [x] \\
 & \Rightarrow \{ \text{property (2)} \} \\
 & \text{foldl } (\text{flip } (:)) ([] ++ [x]) \text{ } xs \\
 & \Rightarrow \{ \text{unfold } (++) \} \\
 & \text{foldl } (\text{flip } (:)) [x] \text{ } xs \\
 & \Rightarrow \{ \text{fold } (\text{flip } (:)) \} \\
 & \text{foldl } (\text{flip } (:)) (\text{flip } (:)) [] \text{ } x \text{ } xs \\
 & \Rightarrow \{ \text{fold foldl} \} \\
 & \text{foldl } (\text{flip } (:)) [] (x : xs)
 \end{aligned}$$

You are encouraged to try proving property (1) directly, in which case you will likely come to the same conclusion, namely that the property needs to be generalized. This is not always easy to see, but is sometimes an important step is constructing a proof, because, despite being somewhat counterintuitive, it is often the case that making a property more general (and therefore more powerful) makes it easier to prove.

In any case, how do we prove property (2)? Using induction, of course! Setting xs to $[],$ the base case is easy:

$$\begin{aligned}
& \text{foldl } (\text{flip } (:)) \text{ } ys [] \text{ } ++ [y] \\
& \Rightarrow \{ \text{unfold foldl} \} \\
& ys ++ [y] \\
& \Rightarrow \{ \text{fold foldl} \} \\
& \text{foldl } (\text{flip } (:)) (ys ++ [y]) []
\end{aligned}$$

and the induction step proceeds as follows:

$$\begin{aligned}
& \text{foldl } (\text{flip } (:)) \text{ } ys (x : xs) \text{ } ++ [y] \\
& \Rightarrow \{ \text{unfold foldl} \} \\
& \text{foldl } (\text{flip } (:)) (\text{flip } (:)) \text{ } ys x \text{ } xs \text{ } ++ [y] \\
& \Rightarrow \{ \text{unfold flip} \} \\
& \text{foldl } (\text{flip } (:)) (x : ys) \text{ } xs \text{ } ++ [y] \\
& \Rightarrow \{ \text{induction hypothesis} \} \\
& \text{foldl } (\text{flip } (:)) ((x : ys) \text{ } ++ [y]) \text{ } xs \\
& \Rightarrow \{ \text{unfold } (++) \} \\
& \text{foldl } (\text{flip } (:)) (x : (ys \text{ } ++ [y])) \text{ } xs \\
& \Rightarrow \{ \text{fold foldl} \} \\
& \text{foldl } (\text{flip } (:)) (ys \text{ } ++ [y]) (x : xs)
\end{aligned}$$

8.4 Useful Properties on Lists

There are many useful properties of functions on lists that require inductive proofs. Tables 8.1 and 8.2 list a number of them involving functions used in this text, but their proofs are left as exercises (except for one; see below). You may assume that these properties are true, and use them freely in proving other properties of your programs. In fact, some of these properties can be used to simplify the proof that *reverse1* and *reverse2* are the same; see if you can find them!¹

(Note, by the way, that in the first rule for *map* in Figure 8.1, the type of $\lambda x \rightarrow x$ on the left-hand side is $a \rightarrow b$, whereas on the right-hand side it is $[a] \rightarrow [b]$; i.e. these are really two different functions.)

8.4.1 Function Strictness

Note that the last rule for *map* in Figure 8.1 is only valid for *strict* functions. A function f is said to be strict if $f \perp = \perp$. Recall from Section 1.2 that \perp is the value associated with a non-terminating computation. So another

¹More thorough discussions of these properties and their proofs may be found in [BW88, Bir98].

Properties of *map*:

$$\begin{aligned}
\text{map } (\lambda x \rightarrow x) &= \lambda x \rightarrow x \\
\text{map } (f \circ g) &= \text{map } f \circ \text{map } g \\
\text{map } f \circ \text{tail} &= \text{tail} \circ \text{map } f \\
\text{map } f \circ \text{reverse} &= \text{reverse} \circ \text{map } f \\
\text{map } f \circ \text{concat} &= \text{concat} \circ \text{map } (\text{map } f) \\
\text{map } f (xs \mathrel{++} ys) &= \text{map } f xs \mathrel{++} \text{map } f ys
\end{aligned}$$

For all strict f :

$$f \circ \text{head} = \text{head} \circ \text{map } f$$

Properties of the *fold* functions:

1. If op is associative, and $e \mathrel{‘op’} x = x$ and $x \mathrel{‘op’} e = x$ for all x , then for all finite xs :

$$\text{foldr } op \ e \ xs = \text{foldl } op \ e \ xs$$

2. If the following are true:

$$\begin{aligned}
x \mathrel{‘op1’} (y \mathrel{‘op2’} z) &= (x \mathrel{‘op1’} y) \mathrel{‘op2’} z \\
x \mathrel{‘op1’} e &= e \mathrel{‘op2’} x
\end{aligned}$$

then for all finite xs :

$$\text{foldr } op1 \ e \ xs = \text{foldl } op2 \ e \ xs$$

3. For all finite xs :

$$\text{foldr } op \ e \ xs = \text{foldl } (\text{flip } op) \ e \ (\text{reverse } xs)$$

Table 8.1: Some Useful Properties of *map* and *fold*.

Properties of $(++)$:

For all xs , ys , and zs :

$$\begin{aligned}(xs ++ ys) ++ zs &= xs ++ (ys ++ zs) \\ xs ++ [] &= [] ++ xs = xs\end{aligned}$$

Properties of $take$ and $drop$:

For all finite non-negative m and n , and finite xs :

$$\begin{aligned}take\ n\ xs ++ drop\ n\ xs &= xs \\ take\ m \circ take\ n &= take\ (\min\ m\ n) \\ drop\ m \circ drop\ n &= drop\ (m + n) \\ take\ m \circ drop\ n &= drop\ n \circ take\ (m + n)\end{aligned}$$

For all finite non-negative m and n such that $n \geq m$:

$$drop\ m \circ take\ n = take\ (n - m) \circ drop\ m$$

Properties of $reverse$:

For all finite xs :

$$\begin{aligned}reverse\ (reverse\ xs) &= xs \\ head\ (reverse\ xs) &= last\ xs \\ last\ (reverse\ xs) &= head\ xs\end{aligned}$$

Table 8.2: Useful Properties of Other Functions Over Lists

way to think about a strict function is that it is one that, when applied to a non-terminating computation, results in a non-terminating computation. For example, the successor function $(+1)$ is strict, because $(+1) \perp = \perp + 1 = \perp$. In other words, if you apply $(+1)$ to a non-terminating computation, you end up with a non-terminating computation.

Not all functions in Haskell are strict, and we have to be careful to say on which argument a function is strict. For example, $(+)$ is strict on both of its arguments, which is why the section $(+1)$ is also strict. On the other hand, the constant function:

$$\text{const } x \ y = x$$

is strict on its first argument (why?), but not its second, because $\text{const } x \ \perp = x$, for any x .

Details: Understanding strictness requires a careful understanding of Haskell's pattern-matching rules. For example, consider the definition of (\wedge) from the Standard Prelude:

$$\begin{aligned} (\wedge) &:: \text{Bool} \rightarrow \text{Bool} \rightarrow \text{Bool} \\ \text{True} \wedge x &= x \\ \text{False} \wedge _ &= \text{False} \end{aligned}$$

When choosing a pattern to match, Haskell starts with the top, left-most pattern, and works to the right and downward. So in the above, (\wedge) first evaluates its left argument. If that value is *True*, then the first equation succeeds, and the second argument gets evaluated because that is the value that is returned. But if the first argument is *False*, the second equation succeeds. In particular, *it does not bother to evaluate the second argument at all*, and simply returns *False* as the answer. This means that (\wedge) is strict in its first argument, but not its second.

A more detailed discussion of pattern matching is found in Appendix C.

Let's now look more closely at the last law for *map*, which says that for all strict f :

$$f \circ \text{head} = \text{head} \circ \text{map } f$$

Let's try to prove this property, starting with the base case, but ignoring for now the strictness constraint on f :

$$\begin{aligned} &f (\text{head } []) \\ &\Rightarrow f \ \perp \end{aligned}$$

$head []$ is an error, which you will recall has value \perp . So you can see immediately that the issue of strictness might play a role in the proof, because without knowing anything about f , there is no further calculation to be done here. Similarly, if we start with the right-hand side:

$$\begin{aligned} & head (map f []) \\ \Rightarrow & head [] \\ \Rightarrow & \perp \end{aligned}$$

It should be clear that for the base case to be true, it must be that $f \perp = \perp$; i.e., f must be strict. Thus we have essentially “discovered” the constraint on the theorem through the process of trying to prove it! (This is not an uncommon phenomenon.)

The induction step is less problematic:

$$\begin{aligned} & f (head (x : xs)) \\ \Rightarrow & f x \\ \Rightarrow & head (f x : map f xs) \\ \Rightarrow & head (map f (x : xs)) \end{aligned}$$

and we are done.

Exercise 8.1 From Chapter 3, prove that:

- $toPitches = map pitch$
- $chord = fold (:=) (rest 0)$
- $maxPitch = fold (!!) 0$
- $listSum xs = fold (+) 0$

Exercise 8.2 Prove as many of the properties in Tables 8.1 and 8.2 as you can.

Exercise 8.3 Which of the following functions are strict (if the function takes more than one argument, specify on which arguments it is strict): *reverse*, *simple*, *map*, *tail*, *area*, (\wedge) , $(True \wedge)$, $(False \wedge)$, and the following function:

$$\begin{aligned} & ifFun :: Bool \rightarrow a \rightarrow a \rightarrow a \\ & ifFun pred cons alt = \mathbf{if} \, pred \, \mathbf{then} \, cons \, \mathbf{else} \, alt \end{aligned}$$

[Replace the following section with properties about musical functions.
In particular:

$$\text{revM } (\text{revM } m) = m$$

...

Also prove or leave as exercise the fact that the two version of perform are the same.]

8.5 Induction on Other Data Types

Proof by induction is not limited to lists. For example, we can use it to reason about natural numbers.² Suppose we define an exponentiation function as follows:

$$\begin{aligned} (\wedge) &:: \text{Integer} \rightarrow \text{Integer} \rightarrow \text{Integer} \\ x \wedge 0 &= 1 \\ x \wedge n &= x * x \wedge (n - 1) \end{aligned}$$

Details: $(*)$ is defined in the Standard Prelude to have precedence level 7, and recall that if no **infix** declaration is given for an operator it defaults to precedence level 9, which means that (\wedge) has precedence level 9, which is higher than that for $(*)$. Therefore no parentheses are needed to disambiguate the last line in the definition above, which corresponds nicely to mathematical convention.

Now suppose that we want to prove that:

$$(\forall x, n \geq 0, m \geq 0) \quad x \wedge (n + m) = x \wedge n * x \wedge m$$

We proceed by induction on n , beginning with $n = 0$:

$$\begin{aligned} &x \wedge (0 + m) \\ &\Rightarrow x \wedge m \\ &\Rightarrow 1 * (x \wedge m) \\ &\Rightarrow x \wedge 0 * x \wedge m \end{aligned}$$

Next we assume that the property is true for numbers less than or equal to n , and prove it for $n + 1$:

²Indeed, one could argue that a proof by induction over finite lists is really an induction over natural numbers, since it is an induction over the *length* of the list, which is a natural number.

$$\begin{aligned}
& x^{\wedge}((n+1) + m) \\
& \Rightarrow x * x^{\wedge}(n + m) \\
& \Rightarrow x * (x^{\wedge}n * x^{\wedge}m) \\
& \Rightarrow (x * x^{\wedge}n) * x^{\wedge}m \\
& \Rightarrow x^{\wedge}(n + 1) * x^{\wedge}m
\end{aligned}$$

and we are done.

Or are we? What if, in the definition of $(^{\wedge})$, x or n is *negative*? Since a negative integer is not a natural number, we could dispense with the problem by saying that these situations fall beyond the bounds of the property we are trying to prove. But let's look a little closer. If x is negative, the property we are trying to prove still holds (why?). But if n is negative, $x^{\wedge}n$ will not terminate (why?). As diligent programmers we may wish to defend against the latter situation by writing:

$$\begin{aligned}
& (^{\wedge}) :: Integer \rightarrow Integer \rightarrow Integer \\
& x^{\wedge}0 = 1 \\
& x^{\wedge}n \mid n < 0 = \text{error "negative exponent"} \\
& \quad \mid \text{otherwise} = x * x^{\wedge}(n - 1)
\end{aligned}$$

If we consider non-terminating computations and ones that produce an error to both have the same value, namely *botom*, then these two versions of $(^{\wedge})$ are equivalent. Pragmatically, however, the latter is clearly superior.

Note that the above definition will test for $n < 0$ on every recursive call, when actually the only call in which it could happen is the first. Therefore a slightly more efficient version of this program would be:

$$\begin{aligned}
& (^{\wedge}) :: Integer \rightarrow Integer \rightarrow Integer \\
& x^{\wedge}n \mid n < 0 = \text{error "negative exponent"} \\
& \quad \mid \text{otherwise} = f \ x \ n \\
& \quad \textbf{where} \ f \ x \ 0 = 1 \\
& \quad \quad f \ x \ n = x * f \ x \ (n - 1)
\end{aligned}$$

Proving the property stated earlier for this version of the program is straightforward, with one minor distinction: what we really need to prove is that the property is true for f ; that is:

$$(\forall x, n \geq 0, m \geq 0) \ f \ x \ (n + m) = f \ x \ n * f \ x \ m$$

from which the proof for the whole function follows trivially.

8.5.1 A More Efficient Exponentiation Function

But in fact there is a more serious inefficiency in our exponentiation function: we are not taking advantage of the fact that, for any even number n , $x^n = (x * x)^{n/2}$. Using this fact, here is a more clever way to accomplish the exponentiation task, using the names $(^!)$ and ff for our functions to distinguish them from the previous versions:

```
(^!) :: Integer -> Integer -> Integer
x ^! n | n < 0 = error "negative exponent"
      | otherwise = ff x n
      where ff x n | n == 0 = 1
                  | even n = ff (x * x) (n `quot` 2)
                  | otherwise = x * ff x (n - 1)
```

Details: *quot* is Haskell's *quotient* operator, which returns the integer quotient of the first argument divided by the second, rounded toward zero.

You should convince yourself that, intuitively at least, this version of exponentiation is not only correct, but also more efficient. More precisely, $(^)$ executes a number of steps proportional to n , whereas $(^!)$ executes a number of steps proportional to the \log_2 of n . The Standard Prelude defines $(^)$ similarly to the way in which $(^!)$ is defined here.

Since intuition is not always reliable, let's *prove* that this version is equivalent to the old. That is, we wish to prove that $x^n = x ^! n$ for all x and n .

A quick look at the two definitions reveals that what we really need to prove is that $f x n = ff x n$, from which it follows immediately that $x^n = x ^! n$. We do this by induction on n , beginning with the base case $n = 0$:

$$f x 0 \Rightarrow 1 \Rightarrow ff x 0$$

so the base step holds trivially. The induction step, however, is considerably more complicated. We must consider two cases: $n + 1$ is either even, or it is odd. If it is odd, we can show that:

$$\begin{aligned} & f x (n + 1) \\ & \Rightarrow x * f x n \\ & \Rightarrow x * ff x n \\ & \Rightarrow ff x (n + 1) \end{aligned}$$

and we are done (note the use of the induction hypothesis in the second step).

If $n + 1$ is even, we might try proceeding in a similar way:

$$\begin{aligned} & f \ x \ (n + 1) \\ \Rightarrow & x * f \ x \ n \\ \Rightarrow & x * ff \ x \ n \end{aligned}$$

But now what shall we do? Since n is odd, we might try unfolding the call to ff :

$$\begin{aligned} & x * ff \ x \ n \\ \Rightarrow & x * (x * ff \ x \ (n - 1)) \end{aligned}$$

but this doesn't seem to be getting us anywhere. Furthermore, *folding* the call to ff (as we did in the odd case) would involve *doubling* n and taking the square root of x , neither of which seems like a good idea!

We could also try going in the other direction:

$$\begin{aligned} & ff \ x \ (n + 1) \\ \Rightarrow & ff \ (x * x) \ ((n + 1) \text{ 'quot' } 2) \\ \Rightarrow & f \ (x * x) \ ((n + 1) \text{ 'quot' } 2) \end{aligned}$$

The use of the induction hypothesis in the second step needs to be justified, because the first argument to f has changed from x to $x * x$. But recall that the induction hypothesis states that for *all* values x , and all natural numbers up to n , $f \ x \ n$ is the same as $ff \ x \ n$. So this is OK.

But even allowing this, we seem to be stuck again!


Instead of pushing this line of reasoning further, let's pursue a different tact based on the (valid) assumption that if m is even, then:

$$m = m \text{ 'quot' } 2 + m \text{ 'quot' } 2$$

Let's use this fact together with the property that we proved in the last section:

$$\begin{aligned} & f \ x \ (n + 1) \\ \Rightarrow & f \ x \ ((n + 1) \text{ 'quot' } 2 + (n + 1) \text{ 'quot' } 2) \\ \Rightarrow & f \ x \ ((n + 1) \text{ 'quot' } 2) * f \ x \ ((n + 1) \text{ 'quot' } 2) \end{aligned}$$

Next, as with the proof in the last section involving *reverse*, let's make an assumption about a property that will help us along. Specifically, what if we could prove that $f \ x \ n * f \ x \ n$ is equal to $f \ (x * x) \ n$? If so, we could proceed as follows:

Figure 8.1: Proof that $f\ x\ n * f\ x\ n = f\ (x * x)\ n$.

$$\begin{aligned}
 & f\ x\ ((n + 1)\ 'quot'\ 2) * f\ x\ ((n + 1)\ 'quot'\ 2) \\
 & \Rightarrow f\ (x * x)\ ((n + 1)\ 'quot'\ 2) \\
 & \Rightarrow ff\ (x * x)\ ((n + 1)\ 'quot'\ 2) \\
 & \Rightarrow ff\ x\ (n + 1)
 \end{aligned}$$

and we are finally done. Note the use of the induction hypothesis in the second step, as justified earlier. The proof of the auxiliary property is not difficult, but also requires induction; it is shown in Figure 8.1.

Aside from improving efficiency, one of the pleasant outcomes of proving that (\wedge) and $(\wedge!)$ are equivalent is that *anything that we prove about one function will be true for the other*. For example, the validity of the property that we proved earlier:

$$x^{\wedge}(n + m) = x^{\wedge}n * x^{\wedge}m$$

immediately implies the validity of:

$$x^{\wedge!}(n + m) = x^{\wedge!}n * x^{\wedge!}m$$

Although $(\wedge!)$ is more efficient than (\wedge) , it is also more complicated, so it makes sense to try proving new properties for (\wedge) , since the proofs will likely be easier.

The moral of this story is that you shouldn't throw away old code that is simpler but less efficient than a newer version. That old code can serve

at least two good purposes: First, if it is simpler, it is likely to be easier to understand, and thus serves a useful role in documenting your effort. Second, as we have just discussed, if it is provably equivalent to the new code, then it can be used to simplify the task of proving properties about the new code.

Exercise 8.4 The function $(^!)$ can be made more efficient by noting that in the last line of the definition of ff , n is odd, and therefore $n - 1$ must be even, so the test for n being even on the next recursive call could be avoided. Redefine $(^!)$ so that it avoids this (minor) inefficiency.

Exercise 8.5 Consider this definition of the *factorial* function:³

```

fac1 :: Integer → Integer
fac1 0 = 1
fac1 n = n * fac1 (n - 1)

```

and this alternative definition:

```

fac2 :: Integer → Integer
fac2 n = fac' n 1
where fac' 0 x = x
        fac' n x = fac' (n - 1) (n * x)

```

Prove that $fac1\ n = fac2\ n$ for all non-negative integers n .

³The factorial function is defined mathematically as:

$$factorial(n) = \begin{cases} 1 & \text{if } n = 0 \\ n * factorial(n - 1) & \text{otherwise} \end{cases}$$

Chapter 9

An Algebra of Music

In this chapter we will explore a number of properties of the *Music* data type and functions defined on it, properties that collectively form an *algebra of music*. With this algebra we can reason about, transform, and optimize computer music programs in a meaning preserving way.

9.1 Musical Equivalence

Suppose we have two values $m1 :: \text{Music Pitch}$ and $m2 :: \text{Music Pitch}$, and we want to know if they are equal. If we treat them simply as Haskell values, we could easily write a function that compares their structures recursively to see if they are the same at every level, all the way down to the *Primitive* rests and notes. This is in fact what the Haskell function (\equiv) does. For example, if:

```
m = c 4 en :+: c 5 en
m1 = m :+: m
m2 = revM (revM m :+: revM m)
```

Then $m1 \equiv m2$ is *True*.

Unfortunately, this is not always good enough from a musical point of view. For example, we would expect the following two musical values to *sound* the same, regardless of the actual values of $m1$, $m2$, and $m3$:

```
(m1 :+: m2) :+: m3
m1 :+: (m2 :+: m3)
```

In other words, we expect the operator ($:+:$) to be *associative*.

The problem is that, as data structures, these two values are *not* equal in general, in fact there are no finite values that can be assigned to $m1$, $m2$, and $m3$ to make them equal.¹

The obvious way out of this dilemma is to define a new notion of equality that captures the fact that the *performances* are the same—i.e. if two things sound the same, they must be musically equivalent. And thus we define a formal notion of musical equivalence:

Definition: Two musical values $m1$ and $m2$ are *equivalent*, written $m1 \equiv m2$, if and only if:

$$(\forall pmap, c) \text{ perf } pmap \ c \ m1 = \text{perf } pmap \ c \ m2$$

We will study a number of properties in this chapter that capture musical equivalences, similar in spirit to the associativity of $(+:)$ above. Each of them can be thought of as an *axiom*, and the set of valid axioms collectively forms an *algebra of music*. By proving the validity of each axiom we not only confirm our intuitions about how music is interpreted, but also gain confidence that our *perform* function actually does the right thing. Furthermore, with these axioms in hand, we can *transform* musical values in meaning-preserving ways.

Speaking of the *perform* function, recall from Chapter 6 that we defined *two* versions of *perform*, and the definition above uses the function *perf*, which includes the duration of a musical value in its result. The following Lemma captures the connection between these functions:

Lemma 9.1.1 For all $pmap$, c , and $m2$:

$$\text{perf } pmap \ c \ m2 = (\text{perform } pmap \ c \ m2, \text{dur } m2)$$

where *perform* is the function defined in Figure 6.1.

To see the importance of including duration in the definition of equivalence, we first note that if two musical values are equivalent, we should be able to substitute one for the other in any valid musical context. But if duration is not taken into account, then all rests are equivalent (because their performances are just the empty list). This means that, for example, $m1 :+: \text{rest } 1 :+: m2$ is equivalent to $m1 :+: \text{rest } 2 :+: m2$, which is surely not what we want.

¹If $m1 = m1 :+: m2$ and $m3 = m2 :+: m3$ then the two expressions are equal, but these are infinite values that cannot even be performed.

Note that we could have defined *perf* as above, i.e. in terms of *perform* and *dur*, but as mentioned in Section 6.1 it would have been computationally inefficient to do so. On the other hand, if the Lemma above is true, then our proofs might be simpler if we first proved the property using *perform*, and then using *dur*. That is, to prove $m1 \equiv m2$ we need to prove:

$$\text{perf } \text{pmap } c \ m1 = \text{perf } \text{pmap } c \ m2$$

Instead of doing this directly using the definition of *perf*, we could instead prove both of the following:

$$\begin{aligned} \text{perform } \text{pmap } c \ m1 &= \text{perform } \text{pmap } c \ m2 \\ \text{dur } m1 &= \text{dur } m2 \end{aligned}$$

9.2 Some Simple Axioms

Let's look at a few simpler axioms, and see how we can prove each of them using the proof techniques that we have developed so far.

Axiom 9.2.1 For any $r1$, $r2$, and m :

$$\text{Modify } (\text{Tempo } r1) (\text{Modify } (\text{Tempo } r2) \ m) \equiv \text{Modify } (\text{Tempo } (r1 * r2)) \ m$$

In other words, *tempo scaling is multiplicative*.

We can prove this by calculation, starting with the definition of musical equivalence. For clarity we will first prove the property for *perform*, and then for *dur*, as suggested in the last section:

let $dt = cDur \ c$

$$\begin{aligned} &\text{perform } \text{pmap } c (\text{Modify } (\text{Tempo } r1) (\text{Modify } (\text{Tempo } r2) \ m)) \\ &\Rightarrow \{ \text{unfold } \text{perform} \} \\ &\text{perform } \text{pmap } (c\{cDur = dt / r1\}) (\text{Modify } (\text{Tempo } r2) \ m) \\ &\Rightarrow \{ \text{unfold } \text{perform} \} \\ &\text{perform } \text{pmap } (c\{cDur = (dt / r1) / r2\}) \ m \\ &\Rightarrow \{ \text{arithmetic} \} \\ &\text{perform } \text{pmap } (c\{cDur = dt / (r1 * r2)\}) \ m \\ &\Rightarrow \{ \text{fold } \text{perform} \} \\ &\text{perform } \text{pmap } c (\text{Modify } (\text{Tempo } (r1 * r2)) \ m) \end{aligned}$$

$$\begin{aligned} &\text{dur } (\text{Modify } (\text{Tempo } r1) (\text{Modify } (\text{Tempo } r2) \ m)) \\ &\Rightarrow \{ \text{unfold } \text{dur} \} \end{aligned}$$

$$\begin{aligned}
& \text{dur } (\text{Modify } (\text{Tempo } r2) \ m) / r1 \\
& \Rightarrow \{ \text{unfold dur} \} \\
& (\text{dur } m / r2) / r1 \\
& \Rightarrow \{ \text{arithmetic} \} \\
& \text{dur } m / (r1 * r2) \\
& \Rightarrow \{ \text{fold dur} \} \\
& \text{dur } (\text{Modify } (\text{Tempo } (r1 * r2)) \ m)
\end{aligned}$$

Here is another useful axiom and its proof:

Axiom 9.2.2 For any r , $m1$, and $m2$:

$$\text{Modify } (\text{Tempo } r) \ (m1 :+ m2) \equiv \text{Modify } (\text{Tempo } r) \ m1 :+ \text{Modify } (\text{Tempo } r) \ m2$$

In other words, *tempo scaling distributes over sequential composition*.

Proof:

$$\begin{aligned}
& \text{let } t = cTime \ c; dt = cDur \ c \\
& \quad t1 = t + \text{dur } m1 * (dt / r) \\
& \quad t2 = t + (\text{dur } m1 / r) * dt \\
& \quad t3 = t + \text{dur } (\text{Modify } (\text{Tempo } r) \ m1) * dt \\
& \text{perform pmap } c \ (\text{Modify } (\text{Tempo } r) \ (m1 :+ m2)) \\
& \Rightarrow \{ \text{unfold perform} \} \\
& \text{perform pmap } (c \{ cDur = dt / r \}) \ (m1 :+ m2) \\
& \Rightarrow \{ \text{unfold perform} \} \\
& \text{perform pmap } (c \{ cDur = dt / r \}) \ m1 \\
& \quad \text{++ perform pmap } (c \{ cTime = t1, cDur = dt / r \}) \ m2 \\
& \Rightarrow \{ \text{fold perform} \} \\
& \text{perform pmap } c \ (\text{Modify } (\text{Tempo } r) \ m1) \\
& \quad \text{++ perform pmap } (c \{ cTime = t1 \}) \ (\text{Modify } (\text{Tempo } r) \ m2) \\
& \Rightarrow \{ \text{arithmetic} \} \\
& \text{perform pmap } c \ (\text{Modify } (\text{Tempo } r) \ m1) \\
& \quad \text{++ perform pmap } (c \{ cTime = t2 \}) \ (\text{Modify } (\text{Tempo } r) \ m2) \\
& \Rightarrow \{ \text{fold dur} \} \\
& \text{perform pmap } c \ (\text{Modify } (\text{Tempo } r) \ m1) \\
& \quad \text{++ perform pmap } (c \{ cTime = t3 \}) \ (\text{Modify } (\text{Tempo } r) \ m2) \\
& \Rightarrow \{ \text{fold perform} \} \\
& \text{perform pmap } c \ (\text{Modify } (\text{Tempo } r) \ m1 :+ \text{Modify } (\text{Tempo } r) \ m2) \\
& \text{dur } (\text{Modify } (\text{Tempo } r) \ (m1 :+ m2))
\end{aligned}$$

$$\begin{aligned}
&\Rightarrow \text{dur } (m1 :+: m2) / r \\
&\Rightarrow (\text{dur } m1 + \text{dur } m2) / r \\
&\Rightarrow \text{dur } m1 / r + \text{dur } m2 / r \\
&\Rightarrow \text{dur } (\text{Modify } (\text{Tempo } r) m1) + \text{dur } (\text{Modify } (\text{Tempo } r) m2) \\
&\Rightarrow \text{dur } (\text{Modify } (\text{Tempo } r) m1 :+: \text{Modify } (\text{Tempo } r) m2)
\end{aligned}$$

An even simpler axiom is given by:

Axiom 9.2.3 For any m , $\text{Modify } (\text{Tempo } 1) m \equiv m$.

In other words, *unit tempo scaling is the identity function for type Music*.

Proof:

$$\begin{aligned}
&\text{let } dt = cDur \ c \\
&\text{perform } pmap \ c \ (\text{Modify } (\text{Tempo } 1) m) \\
&\Rightarrow \{ \text{unfold } perform \} \\
&\text{perform } pmap \ (c \{ cDur = dt / 1 \}) \ m \\
&\Rightarrow \{ \text{arithmetic} \} \\
&\text{perform } pmap \ c \ m \\
&\text{dur } (\text{Modify } (\text{Tempo } 1) m) \\
&\Rightarrow \text{dur } m / 1 \\
&\Rightarrow \text{dur } m
\end{aligned}$$

Note that the above three proofs, being used to establish axioms, all involve the definition of *perform*. In contrast, we can also establish *theorems* whose proofs involve only the axioms. For example, Axioms 1, 2, and 3 are all needed to prove the following:

Theorem 9.2.1 For any r , $m1$, and $m2$:

$$\text{Modify } (\text{Tempo } r) m1 :+: m2 \equiv \text{Modify } (\text{Tempo } r) (m1 :+: \text{Modify } (\text{Tempo } (1 / r)) m2)$$

Proof:

$$\begin{aligned}
&\text{Modify } (\text{Tempo } r) m1 :+: m2 \\
&\Rightarrow \{ \text{Axiom } 3 \} \\
&\text{Modify } (\text{Tempo } r) m1 :+: \text{Modify } (\text{Tempo } 1) m2 \\
&\Rightarrow \{ \text{arithmetic} \} \\
&\text{Modify } (\text{Tempo } r) m1 :+: \text{Modify } (\text{Tempo } (r * (1 / r))) m2
\end{aligned}$$

$$\begin{aligned}
&\Rightarrow \{ \textit{Axiom 1} \} \\
&\textit{Modify} (\textit{Tempo } r) \textit{ m1} \textit{ :+} \textit{ Modify} (\textit{Tempo } r) (\textit{Modify} (\textit{Tempo } (1 / r)) \textit{ m2}) \\
&\Rightarrow \{ \textit{Axiom 2} \} \\
&\textit{Modify} (\textit{Tempo } r) (\textit{m1} \textit{ :+} \textit{ Modify} (\textit{Tempo } (1 / r)) \textit{ m2})
\end{aligned}$$

9.3 The Axiom Set

There are many other useful axioms, but we do not have room to include all of their proofs here. They are listed below, which include the axioms from the previous section as special cases, and the proofs are left as exercises.

Axiom 9.3.1 *Tempo* is *multiplicative* and *Transpose* is *additive*. That is, for any $r1$, $r2$, $p1$, $p2$, and m :

$$\begin{aligned}
&\textit{Modify} (\textit{Tempo } r1) (\textit{Modify} (\textit{Tempo } r2) \textit{ m}) \equiv \textit{Modify} (\textit{Tempo } (r1 * r2)) \textit{ m} \\
&\textit{Modify} (\textit{Trans } p1) (\textit{Modify} (\textit{Trans } p2) \textit{ m}) \equiv \textit{Modify} (\textit{Trans } (p1 + p2)) \textit{ m}
\end{aligned}$$

Axiom 9.3.2 Function composition is *commutative* with respect to both tempo scaling and transposition. That is, for any $r1$, $r2$, $p1$ and $p2$:

$$\begin{aligned}
&\textit{Modify} (\textit{Tempo } r1) \circ \textit{Modify} (\textit{Tempo } r2) \equiv \textit{Modify} (\textit{Tempo } r2) \circ \textit{Modify} (\textit{Tempo } r1) \\
&\textit{Modify} (\textit{Trans } p1) \circ \textit{Modify} (\textit{Trans } p2) \equiv \textit{Modify} (\textit{Trans } p2) \circ \textit{Modify} (\textit{Trans } p1) \\
&\textit{Modify} (\textit{Tempo } r1) \circ \textit{Modify} (\textit{Trans } p1) \equiv \textit{Modify} (\textit{Trans } p1) \circ \textit{Modify} (\textit{Tempo } r1)
\end{aligned}$$

Axiom 9.3.3 Tempo scaling and transposition are *distributive* over both sequential and parallel composition. That is, for any r , p , $m1$, and $m2$:

$$\begin{aligned}
&\textit{Modify} (\textit{Tempo } r) (\textit{m1} \textit{ :+} \textit{ m2}) \equiv \textit{Modify} (\textit{Tempo } r) \textit{ m1} \textit{ :+} \textit{ Modify} (\textit{Tempo } r) \textit{ m2} \\
&\textit{Modify} (\textit{Tempo } r) (\textit{m1} \textit{ :=} \textit{ m2}) \equiv \textit{Modify} (\textit{Tempo } r) \textit{ m1} \textit{ :=} \textit{ Modify} (\textit{Tempo } r) \textit{ m2} \\
&\textit{Modify} (\textit{Trans } p) (\textit{m1} \textit{ :+} \textit{ m2}) \equiv \textit{Modify} (\textit{Trans } p) \textit{ m1} \textit{ :+} \textit{ Modify} (\textit{Trans } p) \textit{ m2} \\
&\textit{Modify} (\textit{Trans } p) (\textit{m1} \textit{ :=} \textit{ m2}) \equiv \textit{Modify} (\textit{Trans } p) \textit{ m1} \textit{ :=} \textit{ Modify} (\textit{Trans } p) \textit{ m2}
\end{aligned}$$

Axiom 9.3.4 Sequential and parallel composition are *associative*. That is, for any $m0$, $m1$, and $m2$:

$$\begin{aligned}
&\textit{m0} \textit{ :+} (\textit{m1} \textit{ :+} \textit{ m2}) \equiv (\textit{m0} \textit{ :+} \textit{ m1}) \textit{ :+} \textit{ m2} \\
&\textit{m0} \textit{ :=} (\textit{m1} \textit{ :=} \textit{ m2}) \equiv (\textit{m0} \textit{ :=} \textit{ m1}) \textit{ :=} \textit{ m2}
\end{aligned}$$

Axiom 9.3.5 Parallel composition is *commutative*. That is, for any $m0$ and $m1$:

$$\textit{m0} \textit{ :=} \textit{ m1} \equiv \textit{m1} \textit{ :=} \textit{ m0}$$

Axiom 9.3.6 *Rest 0* is a *unit* for *Tempo* and *Trans*, and a *zero* for sequential and parallel composition. That is, for any r , p , and m :

$$\begin{aligned} \text{Tempo } r \text{ (Rest 0)} &\equiv \text{Rest 0} \\ \text{Trans } p \text{ (Rest 0)} &\equiv \text{Rest 0} \\ m \text{ :+ : Rest 0} &\equiv m \equiv \text{Rest 0 :+ : } m \\ m \text{ := : Rest 0} &\equiv m \equiv \text{Rest 0 := : } m \end{aligned}$$

Axiom 9.3.7 There is a duality between (:+ :) and (:+ :) , namely that, for any $m0$, $m1$, $m2$, and $m3$ such that $\text{dur } m0 = \text{dur } m2$:

$$(m0 \text{ :+ : } m1) \text{ := : } (m2 \text{ :+ : } m3) \equiv (m0 \text{ := : } m2) \text{ :+ : } (m1 \text{ := : } m3)$$

Exercise 9.1 Establish the validity of each of the above axioms.

Exercise 9.2 Recall the function *revM* defined in Chapter 2, and note that, in general, $\text{revM } (\text{revM } m)$ is not equal to m . However, the following is true:

$$\text{revM } (\text{revM } m) \equiv m$$

Prove this fact by calculation.

9.4 Soundness and Completeness

TBD

Appendix A

A Tour of the PreludeList Module

The use of lists is particularly common when programming in Haskell, and thus, not surprisingly, there are many pre-defined polymorphic functions for lists. The list data type itself, plus some of the most useful functions on it, are contained in the Standard Prelude’s *PreludeList* module, which we will look at in detail in this chapter. There is also a Standard Library module called *List* that has additional useful functions. It is a good idea to become familiar with both modules.

Although this chapter may feel like a long list of “Haskell features,” the functions described here capture many common patterns of list usage that have been discovered by functional programmers over many years of trials and tribulations. In many ways higher-order declarative programming with lists takes the place of lower-level imperative control structures in more conventional languages. By becoming familiar with these list functions you will be able to more quickly and confidently develop your own applications using lists. Furthermore, if all of us do this, we will have a common vocabulary with which to understand each others’ programs. Finally, by reading through the code in this module you will develop a good feel for how to write proper function definitions in Haskell.

It is not necessary for you to understand the details of every function, but you should try to get a sense for what is available so that you can return later when your programming needs demand it. In the long run you are well-advised to read the rest of the Standard Prelude as well as the various Standard Libraries, to discover a host of other functions and data types that you might someday find useful in your own work.

A.1 The PreludeList Module

To get a feel for the *PreludeList* module, let's first look at its module declaration:

```
module PreludeList (  
    map, (++), filter, concat,  
    head, last, tail, init, null, length, (!!),  
    foldl, foldl1, scanl, scanl1, foldr, foldr1, scanr, scanr1,  
    iterate, repeat, replicate, cycle,  
    take, drop, splitAt, takeWhile, dropWhile, span, break,  
    lines, words, unlines, unwords, reverse, and, or,  
    any, all, elem, notElem, lookup,  
    sum, product, maximum, minimum, concatMap,  
    zip, zip3, zipWith, zipWith3, unzip, unzip3)  
where  
  
import qualified Char (isSpace)  
  
infixl 9!  
infixr 5++  
infix 4 ∈, ∉
```

We will not discuss all of the functions listed above, but will cover most of them (and some were discussed in previous chapters).

A.2 Simple List Selector Functions

head and *tail* extract the first element and remaining elements, respectively, from a list, which must be non-empty. *last* and *init* are the dual functions that work from the end of a list, rather than from the beginning.

```
head :: [a] → a  
head (x : _) = x  
head [] = error "PreludeList.head: empty list"  
  
last :: [a] → a  
last [x] = x  
last (_ : xs) = last xs  
last [] = error "PreludeList.last: empty list"
```



```

tail :: [a] → [a]
tail (_: xs) = xs
tail [] = error "PreludeList.tail: empty list"

init :: [a] → [a]
init [x] = []
init (x: xs) = x: init xs
init [] = error "PreludeList.init: empty list"

```

Although *head* and *tail* were previously discussed in Section 3.1, the definitions here include an equation describing their behaviors under erroneous situations—such as selecting the head of an empty list—in which case the *error* function is called. It is a good idea to include such an equation for any definition in which you have not covered every possible case in pattern-matching; i.e. if it is possible that the pattern-matching could “run off the end” of the set of equations. The string argument that you supply to the *error* function should be detailed enough that you can easily track down the precise location of the error in your program.

Details: If such an error equation is omitted, and then during pattern-matching all equations fail, most Haskell systems will invoke the *error* function anyway, but most likely with a string that will be less informative than one you can supply on your own.

The *null* function tests to see if a list is empty.

```

null :: [a] → Bool
null [] = True
null (_: _) = False

```

A.3 Index-Based Selector Functions

To select the *n*th element from a list, with the first element being the 0th element, we can use the indexing function (!!):

```

(!!) :: [a] → Int → a
(x: _) !! 0 = x
(_: xs) !! n | n > 0 = xs !! (n - 1)
(_: _) !! _ = error "PreludeList.!!: negative index"
[] !! _ = error "PreludeList.!!: index too large"

```

Details: Note the definition of two error conditions; be sure that you understand under what conditions these two equations would succeed. In particular, recall that equations are matched in top-down order: the first to match is the one that is chosen.

take *n xs* returns the prefix of *xs* of length *n*, or *xs* itself if *n* > *length xs*. Similarly, *drop* *n xs* returns the suffix of *xs* after the first *n* elements, or [] if *n* > *length xs*. Finally, *splitAt* *n xs* is equivalent to (*take* *n xs*, *drop* *n xs*).

```

take :: Int → [a] → [a]
take 0 _ = []
take _ [] = []
take n (x : xs) | n > 0 = x : take (n - 1) xs
take _ _ = error "PreludeList.take: negative argument"

drop :: Int → [a] → [a]
drop 0 xs = xs
drop _ [] = []
drop n (_ : xs) | n > 0 = drop (n - 1) xs
drop _ _ = error "PreludeList.drop: negative argument"

splitAt :: Int → [a] → ([a], [a])
splitAt 0 xs = ([], xs)
splitAt _ [] = ([], [])
splitAt n (x : xs) | n > 0 = (x : xs', xs'')
                                where (xs', xs'') = splitAt (n - 1) xs
splitAt _ _ = error "PreludeList.splitAt: negative argument"

length :: [a] → Int
length [] = 0
length (_ : l) = 1 + length l

```

For example:

```

take 3 [0, 1 .. 5] ⇒ [0, 1, 2]
drop 3 [0, 1 .. 5] ⇒ [3, 4, 5]
splitAt 3 [0, 1 .. 5] ⇒ ([0, 1, 2], [3, 4, 5])

```

A.4 Predicate-Based Selector Functions

takeWhile *p xs* returns the longest (possibly empty) prefix of *xs*, all of whose elements satisfy the predicate *p*. *dropWhile* *p xs* returns the remaining

suffix. Finally, $\text{span } p \text{ } xs$ is equivalent to $(\text{takeWhile } p \text{ } xs, \text{dropWhile } p \text{ } xs)$, while $\text{break } p$ uses the negation of p .

$$\begin{aligned} \text{takeWhile} &:: (a \rightarrow \text{Bool}) \rightarrow [a] \rightarrow [a] \\ \text{takeWhile } p \text{ } [] &= [] \\ \text{takeWhile } p \text{ } (x : xs) & \\ &\quad | \text{ } p \text{ } x = x : \text{takeWhile } p \text{ } xs \\ &\quad | \text{ otherwise } = [] \end{aligned}$$

$$\begin{aligned} \text{dropWhile} &:: (a \rightarrow \text{Bool}) \rightarrow [a] \rightarrow [a] \\ \text{dropWhile } p \text{ } [] &= [] \\ \text{dropWhile } p \text{ } xs@(x : xs') & \\ &\quad | \text{ } p \text{ } x = \text{dropWhile } p \text{ } xs' \\ &\quad | \text{ otherwise } = xs \end{aligned}$$

$$\begin{aligned} \text{span}, \text{break} &:: (a \rightarrow \text{Bool}) \rightarrow [a] \rightarrow ([a], [a]) \\ \text{span } p \text{ } [] &= ([], []) \\ \text{span } p \text{ } xs@(x : xs') & \\ &\quad | \text{ } p \text{ } x = (x : xs', xs'') \text{ where } (xs', xs'') = \text{span } p \text{ } xs \\ &\quad | \text{ otherwise } = (xs, []) \end{aligned}$$

$$\text{break } p = \text{span } (\neg \circ p)$$

filter removes all elements not satisfying a predicate:

$$\begin{aligned} \text{filter} &:: (a \rightarrow \text{Bool}) \rightarrow [a] \rightarrow [a] \\ \text{filter } p \text{ } [] &= [] \\ \text{filter } p \text{ } (x : xs) & \quad | \text{ } p \text{ } x = x : \text{filter } p \text{ } xs \\ & \quad | \text{ otherwise } = \text{filter } p \text{ } xs \end{aligned}$$

A.5 Fold-like Functions

foldl1 and foldr1 are variants of foldl and foldr that have no starting value argument, and thus must be applied to non-empty lists.

$$\begin{aligned} \text{foldl} &:: (a \rightarrow b \rightarrow a) \rightarrow a \rightarrow [b] \rightarrow a \\ \text{foldl } f \text{ } z \text{ } [] &= z \\ \text{foldl } f \text{ } z \text{ } (x : xs) &= \text{foldl } f \text{ } (f \text{ } z \text{ } x) \text{ } xs \end{aligned}$$

$$\begin{aligned} \text{foldl1} &:: (a \rightarrow a \rightarrow a) \rightarrow [a] \rightarrow a \\ \text{foldl1 } f \text{ } (x : xs) &= \text{foldl } f \text{ } x \text{ } xs \end{aligned}$$

```
foldr1 :: (a -> a -> a) -> [a] -> a
foldr1 f [x] = x
foldr1 f (x : xs) = f x (foldr1 f xs)
foldr1 _ [] = error "PreludeList.foldr1: empty list"
```

foldl1 and *foldr1* are best used in cases where an empty list makes no sense for the application. For example, computing the maximum or minimum element of a list does not make sense if the list is empty. Thus *foldl1 max* is a proper function to compute the maximum element of a list.

scanl is similar to *foldl*, but returns a list of successive reduced values from the left:

$$\text{scanl } f \ z \ [x1, x2, \dots] \equiv [z, z \text{ 'f' } x1, (z \text{ 'f' } x1) \text{ 'f' } x2, \dots]$$

For example:

$$scanl\ (+)\ 0\ [1, 2, 3] \Rightarrow [0, 1, 3, 6]$$

Note that $last (scanl\ f\ z\ xs) = foldl\ f\ z\ xs$. $scanl1$ is similar, but without the starting element:

$$scanl1\ f\ [x1, x2, \dots] \equiv [x1, x1\ 'f'\ x2, \dots]$$

Here are the full definitions:

```

scanl :: (a → b → a) → a → [b] → [a]
scanl f q xs = q : (case xs of
                        [] → []
                        x : xs → scanl f (f q x) xs)

scanl1 :: (a → a → a) → [a] → [a]
scanl1 f (x : xs) = scanl f x xs
scanl1 _ [] = error "PreludeList.scanl1: empty list"

```

$$\begin{aligned} scanr &:: (a \rightarrow b \rightarrow b) \rightarrow b \rightarrow [a] \rightarrow [b] \\ scanr\ f\ q0\ [] &= [q0] \\ scanr\ f\ q0\ (x : xs) &= f\ x\ q : qs \\ &\textbf{where } qs @ (q : _) = scanr\ f\ q0\ xs \end{aligned}$$

```

scanr1 :: (a → a → a) → [a] → [a]
scanr1 f [x] = [x]
scanr1 f (x : xs) = f x q : qs
      where qs@(q : _) = scanr1 f xs
scanr1 _ [] = error "PreludeList.scanr1: empty list"

```

A.6 List Generators

There are some functions which are very useful for generating lists from scratch in interesting ways. To start, *iterate f x* returns an *infinite list* of repeated applications of *f* to *x*. That is:

```
iterate f x ⇒ [x, f x, f (f x), ...]
```

The “infinite” nature of this list may at first seem alarming, but in fact is one of the more powerful and useful features of Haskell.

[say more]

```

iterate :: (a → a) → a → [a]
iterate f x = x : iterate f (f x)

```

repeat x is an infinite list, with *x* the value of every element. *replicate n x* is a list of length *n* with *x* the value of every element. And *cycle* ties a finite list into a circular one, or equivalently, the infinite repetition of the original list.

```

repeat :: a → [a]
repeat x = xs where xs = x : xs

replicate :: Int → a → [a]
replicate n x = take n (repeat x)

cycle :: [a] → [a]
cycle [] = error "Prelude.cycle: empty list"
cycle xs = xs' where xs' = xs ++ xs'

```

A.7 String-Based Functions

Recall that strings in Haskell are just lists of characters. Manipulating strings (i.e. text) is a very common practice, so it makes sense that Haskell would have a few pre-defined functions to make this easier for you.

lines breaks a string at every newline character (written as `'\n'` in Haskell), thus yielding a *list* of strings, each of which contains no newline characters. Similarly, *words* breaks a string up into a list of words, which were delimited by white space. Finally, *unlines* and *unwords* are the inverse operations: *unlines* joins lines with terminating newline characters, and *unwords* joins words with separating spaces. (Because of the potential presence of multiple spaces and newline characters, however, these pairs of functions are not true inverses of each other.)

```

lines :: String → [String]
lines "" = []
lines s = let (l, s') = break (≡ '\n') s
           in l : case s' of
                    [] → []
                    (_ : s'') → lines s''

words :: String → [String]
words s = case dropWhile Char.isSpace s of
            "" → []
            s' → w : words s''
            where (w, s'') = break Char.isSpace s'

unlines :: [String] → String
unlines = concatMap (++ "\n")

unwords :: [String] → String
unwords [] = ""
unwords ws = foldr1 (λw s → w ++ ' ' : s) ws

```

reverse reverses the elements in a finite list.

```

reverse :: [a] → [a]
reverse = foldl (flip (:)) []

```

A.8 Boolean List Functions

and and *or* compute the logical “and” and “or,” respectively, of all of the elements in a list of Boolean values.

```

and, or :: [Bool] → Bool
and = foldr (∧) True
or = foldr (∨) False

```

Applied to a predicate and a list, *any* determines if any element of the list satisfies the predicate. An analogous behavior holds for *all*.

$$\begin{aligned} any, all &:: (a \rightarrow Bool) \rightarrow [a] \rightarrow Bool \\ any\ p &= or \circ map\ p \\ all\ p &= and \circ map\ p \end{aligned}$$

A.9 List Membership Functions

elem is the list membership predicate, usually written in infix form, e.g., $x \in xs$ (which is why it was given a fixity declaration at the beginning of the module). *notElem* is the negation of this function.

$$\begin{aligned} elem, notElem &:: (Eq\ a) \Rightarrow a \rightarrow [a] \rightarrow Bool \\ elem\ x &= any\ (\equiv x) \\ notElem\ x &= all\ (\not\equiv x) \end{aligned}$$

It is common to store “key/value” pairs in a list, and to access the list by finding the value associated with a given key (for this reason the list is often called an *association list*). The function *lookup* looks up a key in an association list, returning *Nothing* if it is not found, or *Just y* if *y* is the value associated with the key.

$$\begin{aligned} lookup &:: (Eq\ a) \Rightarrow a \rightarrow [(a, b)] \rightarrow Maybe\ b \\ lookup\ key\ [] &= Nothing \\ lookup\ key\ ((x, y) : xys) & \\ &\quad | key \equiv x = Just\ y \\ &\quad | otherwise = lookup\ key\ xys \end{aligned}$$

A.10 Arithmetic on Lists

sum and *product* compute the sum and product, respectively, of a finite list of numbers.

$$\begin{aligned} sum, product &:: (Num\ a) \Rightarrow [a] \rightarrow a \\ sum &= foldl\ (+)\ 0 \\ product &= foldl\ (*)\ 1 \end{aligned}$$

maximum and *minimum* return the maximum and minimum value, respectively from a non-empty, finite list whose element type is ordered.

```

maximum, minimum :: (Ord a) => [a] -> a
maximum [] = error "Prelude.maximum: empty list"
maximum xs = foldl1 max xs

minimum [] = error "Prelude.minimum: empty list"
minimum xs = foldl1 min xs

```

Note that even though *foldl1* is used in the definition, a test is made for the empty list to give an error message that more accurately reflects the source of the problem.

A.11 List Combining Functions

map and $(++)$ were defined in previous chapters, but are repeated here for completeness:

```

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

(++) :: [a] -> [a] -> [a]
[] ++ ys = ys
(x : xs) ++ ys = x : (xs ++ ys)

```

concat appends together a list of lists:

```

concat :: [[a]] -> [a]
concat xss = foldr (++) [] xss

```

concatMap does what it says: it concatenates the result of mapping a function down a list.

```

concatMap :: (a -> [b]) -> [a] -> [b]
concatMap f = concat o map f

```

zip takes two lists and returns a list of corresponding pairs. If one input list is short, excess elements of the longer list are discarded. *zip3* takes three lists and returns a list of triples. (“Zips” for larger tuples are contained in the List Library.)

```

zip :: [a] -> [b] -> [(a, b)]
zip = zipWith (,)

```


$zip3 :: [a] \rightarrow [b] \rightarrow [c] \rightarrow [(a, b, c)]$
 $zip3 = zipWith3 \, (,,)$

Details: The functions $(,)$ and $(,,)$ are the pairing and tripling functions, respectively:

$(,) \Rightarrow \lambda x \, y \rightarrow (x, y)$
 $(,,) \Rightarrow \lambda x \, y \, z \rightarrow (x, y, z)$

The *zipWith* family generalises the *zip* and *map* families (or, in a sense, combines them) by applying a function (given as the first argument) to each pair (or triple, etc.) of values. For example, *zipWith* $(+)$ is applied to two lists to produce the list of corresponding sums.

$zipWith :: (a \rightarrow b \rightarrow c) \rightarrow [a] \rightarrow [b] \rightarrow [c]$
 $zipWith \, z \, (a : as) \, (b : bs)$
 $\quad = z \, a \, b : zipWith \, z \, as \, bs$
 $zipWith \, _ \, _ \, _ = []$

$zipWith3 :: (a \rightarrow b \rightarrow c \rightarrow d) \rightarrow [a] \rightarrow [b] \rightarrow [c] \rightarrow [d]$
 $zipWith3 \, z \, (a : as) \, (b : bs) \, (c : cs)$
 $\quad = z \, a \, b \, c : zipWith3 \, z \, as \, bs \, cs$
 $zipWith3 \, _ \, _ \, _ \, _ = []$

The following two functions perform the inverse operations of *zip* and *zip3*, respectively.

$unzip :: [(a, b)] \rightarrow ([a], [b])$
 $unzip = foldr \, (\lambda(a, b) \sim (as, bs) \rightarrow (a : as, b : bs)) \, ([], [])$

$unzip3 :: [(a, b, c)] \rightarrow ([a], [b], [c])$
 $unzip3 = foldr \, (\lambda(a, b, c) \sim (as, bs, cs) \rightarrow (a : as, b : bs, c : cs)) \, ([], [], [])$

Appendix B

Built-in Types Are Not Special

Throughout this text we have introduced many “built-in” types such as lists, tuples, integers, and characters. We have also shown how new user-defined types can be defined. Aside from special syntax, you might be wondering if the built-in types are in any way more special than the user-defined ones. The answer is *no*. The special syntax is for convenience and for consistency with historical convention, but has no semantic consequence.

We can emphasize this point by considering what the type declarations would look like for these built-in types if in fact we were allowed to use the special syntax in defining them. For example, the *Char* type might be written as:

```
data Char = 'a' | 'b' | 'c' | ...    -- This is not valid
          | 'A' | 'B' | 'C' | ...    -- Haskell code!
          | '1' | '2' | '3' | ...
```

These constructor names are not syntactically valid; to fix them we would have to write something like:

```
data Char = Ca | Cb | Cc | ...
          | CA | CB | CC | ...
          | C1 | C2 | C3 | ...
```

Even though these constructors are actually more concise, they are quite unconventional for representing characters, and thus the special syntax is used instead.

In any case, writing “pseudo-Haskell” code in this way helps us to see through the special syntax. We see now that *Char* is just a data type consisting of a large number of nullary (meaning they take no arguments) constructors. Thinking of *Char* in this way makes it clear why, for example, we can pattern-match against characters; i.e., we would expect to be able to do so for any of a data type’s constructors.

Similarly, using pseudo-Haskell, we could define *Int* and *Integer* by:

```
-- more pseudo-code:
data Int = (-2^29) | ... | -1 | 0 | 1 | ... | (2^29 - 1)
data Integer = ... - 2 | -1 | 0 | 1 | 2...
```

(Recall that -2^{29} to $2^{29}-1$ is the minimum range for the *Int* data type.) *Int* is clearly a much larger enumeration than *Char*, but it’s still finite! In contrast, the pseudo-code for *Integer* (the type of arbitrary precision integers) is intended to convey an *infinite* enumeration (and in that sense only, the *Integer* data type *is* somewhat special).

Haskell has a data type called *unit* which has exactly one value: `()`. The name of this data type is also written `()`. This is trivially expressed in Haskell pseudo-code:

```
data () = ()    -- more pseudo-code
```

Tuples are also easy to define playing this game:

```
data (a, b) = (a, b)    -- more pseudo-code
data (a, b, c) = (a, b, c)
data (a, b, c, d) = (a, b, c, d)
```

and so on. Each declaration above defines a tuple type of a particular length, with parentheses playing a role in both the expression syntax (as data constructor) and type-expression syntax (as type constructor). By “and so on” we mean that there are an infinite number of such declarations, reflecting the fact that tuples of all finite lengths are allowed in Haskell.

The list data type is also easily handled in pseudo-Haskell, and more interestingly, it is recursive:

```
data [a] = [] | a : [a]    -- more pseudo-code
infixr 5:
```

We can now see clearly what we described about lists earlier: `[]` is the empty list, and `(:)` is the infix list constructor; thus `[1, 2, 3]` must be equivalent to the list `1 : 2 : 3 : []`. (Note that `(:)` is right associative.) The type of `[]` is `[a]`, and the type of `(:)` is `a → [a] → [a]`.

Details: The way $(:)$ is defined here is actually legal syntax—infix constructors are permitted in **data** declarations, and are distinguished from infix operators (for pattern-matching purposes) by the fact that they must begin with a colon (a property trivially satisfied by “:”).

At this point the reader should note carefully the differences between tuples and lists, which the above definitions make abundantly clear. In particular, note the recursive nature of the list type whose elements are homogeneous and of arbitrary length, and the non-recursive nature of a (particular) tuple type whose elements are heterogeneous and of fixed length. The typing rules for tuples and lists should now also be clear:

For $(e1, e2, \dots, en)$, $n \geq 2$, if Ti is the type of ei , then the type of the tuple is $(T1, T2, \dots, Tn)$.

For $[e1, e2, \dots, en]$, $n \geq 0$, each ei must have the same type T , and the type of the list is $[T]$.

Appendix C

Pattern-Matching Details

In this chapter we will look at Haskell’s pattern-matching process in greater detail.

Haskell defines a fixed set of patterns for use in case expressions and function definitions. Pattern matching is permitted using the constructors of any type, whether user-defined or pre-defined in Haskell. This includes tuples, strings, numbers, characters, etc. For example, here’s a contrived function that matches against a tuple of “constants:”

```
contrived :: ([a], Char, (Int, Float), String, Bool) → Bool
contrived ([], 'b', (1, 2.0), "hi", True) = False
```

This example also demonstrates that *nesting* of patterns is permitted (to arbitrary depth).

Technically speaking, *formal parameters* to functions are also patterns—it’s just that they *never fail to match a value*. As a “side effect” of a successful match, the formal parameter is bound to the value it is being matched against. For this reason patterns in any one equation are not allowed to have more than one occurrence of the same formal parameter.

A pattern that may fail to match is said to be *refutable*; for example, the empty list `[]` is refutable. Patterns such as formal parameters that never fail to match are said to be *irrefutable*. There are three other kinds of irrefutable patterns, which are summarized below.

As-Patterns Sometimes it is convenient to name a pattern for use on the right-hand side of an equation. For example, a function that duplicates the first element in a list might be written as:

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

Note that $x : xs$ appears both as a pattern on the left-hand side, and as an expression on the right-hand side. To improve readability, we might prefer to write $x : xs$ just once, which we can achieve using an *as-pattern* as follows:¹

$$f\ s@ (x : xs) = x : s$$

Technically speaking, as-patterns always result in a successful match, although the sub-pattern (in this case $x : xs$) could, of course, fail.

Wildcards Another common situation is matching against a value we really care nothing about. For example, the functions *head* and *tail* can be written as:

$$\begin{aligned} \text{head}\ (x : _) &= x \\ \text{tail}\ (_ : xs) &= xs \end{aligned}$$

in which we have “advertised” the fact that we don’t care what a certain part of the input is. Each wildcard will independently match anything, but in contrast to a formal parameter, each will bind nothing; for this reason more than one are allowed in an equation.

Lazy Patterns There is one other kind of pattern allowed in Haskell. It is called a *lazy pattern*, and has the form $\sim pat$. Lazy patterns are *irrefutable*: matching a value v against $\sim pat$ always succeeds, regardless of pat . Operationally speaking, if an identifier in pat is later “used” on the right-hand-side, it will be bound to that portion of the value that would result if v were to successfully match pat , and \perp otherwise.

Lazy patterns are useful in contexts where infinite data structures are being defined recursively. For example, infinite lists are an excellent vehicle for writing *simulation* programs, and in this context the infinite lists are often called *streams*.

Pattern-Matching Semantics

So far we have discussed how individual patterns are matched, how some are refutable, some are irrefutable, etc. But what drives the overall process? In what order are the matches attempted? What if none succeed? This section addresses these questions.

¹Another advantage to doing this is that a naive implementation might otherwise completely reconstruct $x : xs$ rather than re-use the value being matched against.

Pattern matching can either *fail*, *succeed* or *diverge*. A successful match binds the formal parameters in the pattern. Divergence occurs when a value needed by the pattern diverges (i.e. is non-terminating) or results in an error (\perp). The matching process itself occurs “top-down, left-to-right.” Failure of a pattern anywhere in one equation results in failure of the whole equation, and the next equation is then tried. If all equations fail, the value of the function application is \perp , and results in a run-time error.

For example, if *bot* is a divergent or erroneous computation, and if $[1, 2]$ is matched against $[0, \textit{bot}]$, then 1 fails to match 0, so the result is a failed match. But if $[1, 2]$ is matched against $[\textit{bot}, 0]$, then matching 1 against *bot* causes divergence (i.e. \perp).

The only other twist to this set of rules is that top-level patterns may also have a boolean *guard*, as in this definition of a function that forms an abstract version of a number’s sign:

$$\begin{aligned} \textit{sign } x \mid x > 0 &= 1 \\ &\mid x \equiv 0 = 0 \\ &\mid x < 0 = -1 \end{aligned}$$

Note here that a sequence of guards is given for a single pattern; as with patterns, these guards are evaluated top-down, and the first that evaluates to *True* results in a successful match.

An Example The pattern-matching rules can have subtle effects on the meaning of functions. For example, consider this definition of *take*:

$$\begin{aligned} \textit{take } 0 _ &= [] \\ \textit{take } _ [] &= [] \\ \textit{take } n (x : xs) &= x : \textit{take } (n - 1) xs \end{aligned}$$

and this slightly different version (the first 2 equations have been reversed):

$$\begin{aligned} \textit{take1 } _ [] &= [] \\ \textit{take1 } 0 _ &= [] \\ \textit{take1 } n (x : xs) &= x : \textit{take1 } (n - 1) xs \end{aligned}$$

Now note the following:

$$\begin{aligned} \textit{take } 0 \textit{ bot} &\implies [] \\ \textit{take1 } 0 \textit{ bot} &\implies \perp \\ \textit{take } \textit{bot} [] &\implies \perp \\ \textit{take1 } \textit{bot} [] &\implies [] \end{aligned}$$

We see that *take* is “more defined” with respect to its second argument, whereas *take1* is more defined with respect to its first. It is difficult to say in this case which definition is better. Just remember that in certain applications, it may make a difference. (The Standard Prelude includes a definition corresponding to *take*.)

Case Expressions

Pattern matching provides a way to “dispatch control” based on structural properties of a value. However, in many circumstances we don’t wish to define a *function* every time we need to do this. Haskell’s *case expression* provides a way to solve this problem. Indeed, the meaning of pattern matching in function definitions is specified in the Haskell Report in terms of case expressions, which are considered more primitive. In particular, a function definition of the form:

$$\begin{aligned} f p_{11} \dots p_{1k} &= e_1 \\ &\dots \\ f p_{n1} \dots p_{nk} &= e_n \end{aligned}$$

where each p_{ij} is a pattern, is semantically equivalent to:

$$\begin{aligned} f \ x1 \ x2 \ \dots \ xk &= \mathbf{case} \ (x1, \dots, xk) \ \mathbf{of} \ (p_{11}, \dots, p_{1k}) \rightarrow e_1 \\ &\dots \\ &\quad (p_{n1}, \dots, p_{nk}) \rightarrow e_n \end{aligned}$$

where the x_i are new identifiers. For example, the definition of *take* given earlier is equivalent to:

$$\begin{aligned} take \ m \ ys &= \mathbf{case} \ (m, ys) \ \mathbf{of} \\ &\quad (0, _) \rightarrow [] \\ &\quad (_, []) \rightarrow [] \\ &\quad (n, x : xs) \rightarrow x : take \ (n - 1) \ xs \end{aligned}$$

For type correctness, the types of the right-hand sides of a case expression or set of equations comprising a function definition must all be the same; more precisely, they must all share a common principal type.

The pattern-matching rules for case expressions are the same as we have given for function definitions.

Bibliography

- [Bir98] R. Bird. *Introduction to Functional Programming using Haskell (second edition)*. Prentice Hall, London, 1998.
- [BW88] R. Bird and P. Wadler. *Introduction to Functional Programming*. Prentice Hall, New York, 1988.
- [Chu41] A. Church. *The Calculi of Lambda Conversion*. Princeton University Press, Princeton, NJ, 1941.
- [Hin69] R. Hindley. The principal type scheme of an object in combinatory logic. *Transactions of the American Mathematical Society*, 146:29–60, December 1969.
- [Hud89] P. Hudak. Conception, evolution, and application of functional programming languages. *ACM Computing Surveys*, 21(3):359–411, 1989.
- [Mil78] R.A. Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348–375, December 1978.
- [MTH90] R. Milner, M. Tofte, and R. Harper. *The Definition of Standard ML*. The MIT Press, Cambridge, MA, 1990.
- [Sch24] M. Schönfinkel. Über die bausteine der mathematischen logik. *Mathematische Annalen*, 92:305, 1924.