Haskell 2010
Revised Language Report

Simon Marlow
(editor)

DRAFT

Contents

Preface

> *Some half dozen persons have written technically on combinatory logic, and most of these, including ourselves, have published something erroneous. Since some of our fellow sinners are among the most careful and competent logicians on the contemporary scene, we regard this as evidence that the subject is refractory. Thus fullness of exposition is necessary for accuracy; and excessive condensation would be false economy here, even more than it is ordinarily.*
>
> — Haskell B. Curry and Robert Feys
> in the Preface to *Combinatory Logic* [1], May 31, 1956

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

Goals

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

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

Haskell 2010: language and libraries

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

Haskell has indeed evolved continuously since its original publication. By the middle of 1997, there had been five versions of the language design (from Haskell 1.0 - 1.4). At the 1997 Haskell Workshop in Amsterdam, it was decided that a stable variant of Haskell was needed; this became "Haskell 98" and was published in February 1999. The fixing of minor bugs led to the *Revised* Haskell 98 Report in 2002.

At the 2005 Haskell Workshop, the consensus was that so many extensions to the official language were widely used (and supported by multiple implementations), that it was worthwhile to define another iteration of the language standard, essentially to codify (and legitimise) the status quo.

The Haskell Prime effort was thus conceived as a relatively conservative extension of Haskell 98, taking on board new features only where they were well understood and widely agreed upon. It too was

intended to be a "stable" language, yet reflecting the considerable progress in research on language design in recent years.

After several years exploring the design space, it was decided that a single monolithic revision of the language was too large a task, and the best way to make progress was to evolve the language in small incremental steps, each revision integrating only a small number of well-understood extensions and changes. Haskell 2010 is the first revision to be created in this way, and new revisions are expected once per year.

Extensions to Haskell 98

The most significant language changes in Haskell 2010 relative to Haskell 98 are listed here.

New language features:

- A Foreign Function Interface (FFI).
- Hierarchical module names, e.g. `Data.Bool`.
- Pattern guards.

Removed language features:

- The $(n + k)$ pattern syntax.

Haskell Resources

The Haskell web site http://haskell.org gives access to many useful resources, including:

- Online versions of the language and library definitions.
- Tutorial material on Haskell.
- Details of the Haskell mailing list.
- Implementations of Haskell.
- Contributed Haskell tools and libraries.
- Applications of Haskell.
- User-contributed wiki pages.
- News and events in the Haskell community.

You are welcome to comment on, suggest improvements to, and criticise the language or its presentation in the report, via the Haskell mailing list, or the Haskell wiki.

Building the language

Haskell was created, and continues to be sustained, by an active community of researchers and application programmers. Those who served on the Language and Library committees, in particular, devoted a huge amount of time and energy to the language. Here they are, with their affiliation(s) for the relevant period:

Arvind (MIT)
Lennart Augustsson (Chalmers University)
Dave Barton (Mitre Corp)
Brian Boutel (Victoria University of Wellington)
Warren Burton (Simon Fraser University)
Manuel M T Chakravarty (University of New South Wales)
Duncan Coutts (Well-Typed LLP)
Jon Fairbairn (University of Cambridge)
Joseph Fasel (Los Alamos National Laboratory)
John Goerzen
Andy Gordon (University of Cambridge)

Maria Guzman (Yale University)
Kevin Hammond [editor] (University of Glasgow)
Bastiaan Heeren (Utrecht University)
Ralf Hinze (University of Bonn)
Paul Hudak [editor] (Yale University)
John Hughes [editor] (University of Glasgow; Chalmers University)
Thomas Johnsson (Chalmers University)
Isaac Jones (Galois, inc.)
Mark Jones (Yale University, University of Nottingham, Oregon Graduate Institute)
Dick Kieburtz (Oregon Graduate Institute)
John Launchbury (University of Glasgow; Oregon Graduate Institute; Galois, inc.)
Andres Löh (Utrecht University)
Ian Lynagh (Well-Typed LLP)
Simon Marlow [editor] (Microsoft Research)
John Meacham
Erik Meijer (Utrecht University)
Ravi Nanavati (Bluespec, inc.)
Rishiyur Nikhil (MIT)
Henrik Nilsson (University of Nottingham)
Ross Paterson (City University, London)
John Peterson [editor] (Yale University)
Simon Peyton Jones [editor] (University of Glasgow; Microsoft Research Ltd)
Mike Reeve (Imperial College)
Alastair Reid (University of Glasgow)
Colin Runciman (University of York)
Don Stewart (Galois, inc.)
Martin Sulzmann (Informatik Consulting Systems AG)
Audrey Tang
Simon Thompson (University of Kent)
Philip Wadler [editor] (University of Glasgow)
Malcolm Wallace (University of York)
Stephanie Weirich (University of Pennsylvania)
David Wise (Indiana University)
Jonathan Young (Yale University)

Those marked [editor] served as the co-ordinating editor for one or more revisions of the language.

In addition, dozens of other people made helpful contributions, some small but many substantial. They are as follows: Hans Aberg, Kris Aerts, Sten Anderson, Richard Bird, Tom Blenko, Stephen Blott, Duke Briscoe, Paul Callaghan, Magnus Carlsson, Mark Carroll, Franklin Chen, Olaf Chitil, Chris Clack, Guy Cousineau, Tony Davie, Craig Dickson, Chris Dornan, Laura Dutton, Chris Fasel, Pat Fasel, Sigbjorn Finne, Michael Fryers, Peter Gammie, Andy Gill, Mike Gunter, Cordy Hall, Mark Hall, Thomas Hallgren, Matt Harden, Klemens Hemm, Fergus Henderson, Dean Herington, Bob Hiromoto, Nic Holt, Ian Holyer, Randy Hudson, Alexander Jacobson, Patrik Jansson, Robert Jeschofnik, Orjan Johansen, Simon B. Jones, Stef Joosten, Mike Joy, Wolfram Kahl, Stefan Kahrs, Antti-Juhani Kaijanaho, Jerzy Karczmarczuk, Kent Karlsson, Martin D. Kealey, Richard Kelsey, Siau-Cheng Khoo, Amir Kishon, Feliks Kluzniak, Jan Kort, Marcin Kowalczyk, Jose Labra, Jeff Lewis, Mark Lillibridge, Björn Lisper, Sandra Loosemore, Pablo Lopez, Olaf Lubeck, Christian Maeder, Ketil Malde, Michael Marte, Jim Mattson, John Meacham, Sergey Mechveliani, Gary Memovich, Randy Michelsen, Rick Mohr, Andy Moran, Graeme Moss, Arthur Norman, Nick North, Chris Okasaki, Bjarte M. Østvold, Paul Otto, Sven Panne, Dave

Parrott, Larne Pekowsky, Rinus Plasmeijer, Ian Poole, Stephen Price, John Robson, Andreas Rossberg, George Russell, Patrick Sansom, Michael Schneider, Felix Schroeter, Julian Seward, Nimish Shah, Christian Sievers, Libor Skarvada, Jan Skibinski, Lauren Smith, Raman Sundaresh, Josef Svenningsson, Ken Takusagawa, Wolfgang Thaller, Satish Thatte, Tom Thomson, Tommy Thorn, Dylan Thurston, Mike Thyer, Mark Tullsen, David Tweed, Pradeep Varma, Keith Wansbrough, Tony Warnock, Michael Webber, Carl Witty, Stuart Wray, and Bonnie Yantis.

Finally, aside from the important foundational work laid by Church, Rosser, Curry, and others on the lambda calculus, it is right to acknowledge the influence of many noteworthy programming languages developed over the years. Although it is difficult to pinpoint the origin of many ideas, the following languages were particularly influential: Lisp (and its modern-day incarnations Common Lisp and Scheme); Landin's ISWIM; APL; Backus's FP [2]; ML and Standard ML; Hope and Hope$^+$; Clean; Id; Gofer; Sisal; and Turner's series of languages culminating in Miranda[1]. Without these forerunners Haskell would not have been possible.

Simon Marlow
Cambridge, April 2010

[1]Miranda is a trademark of Research Software Ltd.

Preface to the Revised Report

Over 16 years have passed since the publication of the Haskell 2010 language report, and the Haskell that we write today has evolved in many ways, some of which are incompatible with the language specified in the Haskell 2010 language report. In order to account for this change we saw the need to compile a *revised version* of the Haskell 2010 language report which does not add any novel language features but brings the report in line with the version of Haskell 2010 implemented by compilers and maintainers of the standard libraries. This revised report is the result of that effort.

Organization of the Revised Report

In distinction to the previous version of the Haskell 2010 language report, this revision separates out documentation of the standard libraries. These standard libraries are now available in the form of a zero-dependency Haskell package which makes the API specification machine readable and allows to generate documentation using standard Haskell tools like Haddock.

Changes to the Report

Chapter 1

Introduction

Haskell is a general purpose, purely functional programming language incorporating many recent innovations in programming language design. Haskell provides higher-order functions, non-strict semantics, static polymorphic typing, user-defined algebraic datatypes, pattern-matching, list comprehensions, a module system, a monadic I/O system, and a rich set of primitive datatypes, including lists, arrays, arbitrary and fixed precision integers, and floating-point numbers. Haskell is both the culmination and solidification of many years of research on non-strict functional languages.

This report defines the syntax for Haskell programs and an informal abstract semantics for the meaning of such programs. We leave as implementation dependent the ways in which Haskell programs are to be manipulated, interpreted, compiled, etc. This includes such issues as the nature of programming environments and the error messages returned for undefined programs (i.e. programs that formally evaluate to $\bot$).

1.1 Program Structure

In this section, we describe the abstract syntactic and semantic structure of Haskell, as well as how it relates to the organization of the rest of the report.

1. At the topmost level a Haskell program is a set of *modules*, described in Chapter 5. Modules provide a way to control namespaces and to re-use software in large programs.
2. The top level of a module consists of a collection of *declarations*, of which there are several kinds, all described in Chapter 4. Declarations define things such as ordinary values, datatypes, type classes, and fixity information.
3. At the next lower level are *expressions*, described in Chapter 3. An expression denotes a *value* and has a *static type*; expressions are at the heart of Haskell programming "in the small."
4. At the bottom level is Haskell's *lexical structure*, defined in Chapter 2. The lexical structure captures the concrete representation of Haskell programs in text files.

This report proceeds bottom-up with respect to Haskell's syntactic structure.

The chapters not mentioned above are Chapter 6, which describes the standard built-in datatypes and classes in Haskell, and Chapter 7, which discusses the I/O facility in Haskell (i.e. how Haskell programs communicate with the outside world). Also, there are several chapters describing the Prelude, the concrete syntax, literate programming, the specification of derived instances, and pragmas supported by most Haskell compilers.

Examples of Haskell program fragments in running text are given in typewriter font:

```
let x = 1
    z = x+y
in  z+1
```

"Holes" in program fragments representing arbitrary pieces of Haskell code are written in italics, as in `if` e_1 `then` e_2 `else` e_3. Generally the italicized names are mnemonic, such as e for expressions, d for declarations, t for types, etc.

1.2 The Haskell Kernel

Haskell has adopted many of the convenient syntactic structures that have become popular in functional programming. In this Report, the meaning of such syntactic sugar is given by translation into simpler constructs. If these translations are applied exhaustively, the result is a program written in a small subset of Haskell that we call the Haskell *kernel*.

Although the kernel is not formally specified, it is essentially a slightly sugared variant of the lambda calculus with a straightforward denotational semantics. The translation of each syntactic structure into the kernel is given as the syntax is introduced. This modular design facilitates reasoning about Haskell programs and provides useful guidelines for implementors of the language.

1.3 Values and Types

An expression evaluates to a *value* and has a static *type*. Values and types are not mixed in Haskell. However, the type system allows user-defined datatypes of various sorts, and permits not only parametric polymorphism (using a traditional Hindley-Milner type structure) but also *ad hoc* polymorphism, or *overloading* (using *type classes*).

Errors in Haskell are semantically equivalent to $\perp$ ("bottom"). Technically, they are indistinguishable from nontermination, so the language includes no mechanism for detecting or acting upon errors. However, implementations will probably try to provide useful information about errors. See Section 3.1.

1.4 Namespaces

There are six kinds of names in Haskell: those for *variables* and *constructors* denote values; those for *type variables*, *type constructors*, and *type classes* refer to entities related to the type system; and *module names* refer to modules. There are two constraints on naming:

1. Names for variables and type variables are identifiers beginning with lowercase letters or underscore; the other four kinds of names are identifiers beginning with uppercase letters.
2. An identifier must not be used as the name of a type constructor and a class in the same scope.

These are the only constraints; for example, `Int` may simultaneously be the name of a module, class, and constructor within a single scope.

Chapter 2

Lexical Structure

In this chapter, we describe the low-level lexical structure of Haskell. Most of the details may be skipped in a first reading of the report.

2.1 Notational Conventions

These notational conventions are used for presenting syntax:

[*pattern*]	optional
{*pattern*}	zero or more repetitions
(*pattern*)	grouping
$pat_1 \mid pat_2$	choice
$pat_{\langle pat' \rangle}$	difference—elements generated by *pat* except those generated by *pat'*
`fibonacci`	terminal syntax in typewriter font

Because the syntax in this section describes *lexical* syntax, all whitespace is expressed explicitly; there is no implicit space between juxtaposed symbols. BNF-like syntax is used throughout, with productions having the form:

$$nonterm \quad \rightarrow \quad alt_1 \mid alt_2 \mid \ldots \mid alt_n$$

Care must be taken in distinguishing metalogical syntax such as | and [...] from concrete terminal syntax (given in typewriter font) such as `|` and `[...]`, although usually the context makes the distinction clear.

Haskell uses the Unicode [3] character set. However, source programs are currently biased toward the ASCII character set used in earlier versions of Haskell.

This syntax depends on properties of the Unicode characters as defined by the Unicode consortium. Haskell compilers are expected to make use of new versions of Unicode as they are made available.

2.2 Lexical Program Structure

$$\begin{array}{lcl} program & \rightarrow & \{\ lexeme \mid whitespace\ \} \\ lexeme & \rightarrow & qvarid \mid qconid \mid qvarsym \mid qconsym \end{array}$$

	|	*literal* | *special* | *reservedop* | *reservedid*
literal	→	*integer* | *float* | *char* | *string*
special	→	(|) | , | ; | [|] | ` | { | }
whitespace	→	*whitestuff* {*whitestuff*}
whitestuff	→	*whitechar* | *comment* | *ncomment*
whitechar	→	*newline* | *vertab* | *space* | *tab* | *uniWhite*
newline	→	*return linefeed* | *return* | *linefeed* | *formfeed*
return	→	a carriage return
linefeed	→	a line feed
vertab	→	a vertical tab
formfeed	→	a form feed
space	→	a space
tab	→	a horizontal tab
uniWhite	→	any Unicode character defined as whitespace
comment	→	*dashes* $\left[any_{\langle symbol \rangle} \{ any \} \right]$ *newline*
dashes	→	-- {-}
opencom	→	{-
closecom	→	-}
ncomment	→	*opencom ANYseq* {*ncomment ANYseq*} *closecom*
ANYseq	→	$\{ ANY \}_{\langle \{ ANY \} (opencom \mid closecom) \{ ANY \} \rangle}$
ANY	→	*graphic* | *whitechar*
any	→	*graphic* | *space* | *tab*
graphic	→	*small* | *large* | *symbol* | *digit* | *special* | " | '
small	→	*ascSmall* | *uniSmall* | _
ascSmall	→	a | b | ... | z
uniSmall	→	any Unicode lowercase letter
large	→	*ascLarge* | *uniLarge*
ascLarge	→	A | B | ... | Z
uniLarge	→	any uppercase or titlecase Unicode letter
symbol	→	*ascSymbol* | $uniSymbol_{\langle special \mid \_ \mid " \mid ' \rangle}$
ascSymbol	→	! | # | $ | % | & | * | + | . | / | < | = | > | ? | @
	|	\ | ^ | | | - | ~ | :
uniSymbol	→	any Unicode symbol or punctuation
digit	→	*ascDigit* | *uniDigit*
ascDigit	→	0 | 1 | ... | 9
uniDigit	→	any Unicode decimal digit
octit	→	0 | 1 | ... | 7
hexit	→	*digit* | A | ... | F | a | ... | f

Lexical analysis should use the maximal munch'' rule: at each point, the longest possible lexeme satisfying the *lexeme* production is read. So, although `case` is a reserved word, `cases` is not. Similarly, although = is reserved, == and ~= are not.

Any kind of *whitespace* is also a proper delimiter for lexemes.

Characters not in the category *ANY* are not valid in Haskell programs and should result in a lexing error.

2.3 Comments

Comments are valid whitespace.

An ordinary comment begins with a sequence of two or more consecutive dashes (e.g. `--`) and extends to the following newline. *The sequence of dashes must not form part of a legal lexeme.* For example, "`-->`" or "`|--`" do *not* begin a comment, because both of these are legal lexemes; however "`--foo`" does start a comment.

A nested comment begins with "`{-`" and ends with "`-}`". No legal lexeme starts with "`{-`"; hence, for example, "`{---`" starts a nested comment despite the trailing dashes.

The comment itself is not lexically analysed. Instead, the first unmatched occurrence of the string "`-}`" terminates the nested comment. Nested comments may be nested to any depth: any occurrence of the string "`{-`" within the nested comment starts a new nested comment, terminated by "`-}`". Within a nested comment, each "`{-`" is matched by a corresponding occurrence of "`-}`".

In an ordinary comment, the character sequences "`{-`" and "`-}`" have no special significance, and, in a nested comment, a sequence of dashes has no special significance.

Nested comments are also used for compiler pragmas, as explained in Chapter 12.

If some code is commented out using a nested comment, then any occurrence of `{-` or `-}` within a string or within an end-of-line comment in that code will interfere with the nested comments.

2.4 Identifiers and Operators

$$
\begin{array}{lcl}
\mathit{varid} & \rightarrow & (\mathit{small}\ \{\mathit{small} \mid \mathit{large} \mid \mathit{digit} \mid \texttt{'}\})_{\langle \mathit{reservedid} \rangle} \\
\mathit{conid} & \rightarrow & \mathit{large}\ \{\mathit{small} \mid \mathit{large} \mid \mathit{digit} \mid \texttt{'}\} \\
\mathit{reservedid} & \rightarrow & \texttt{case} \mid \texttt{class} \mid \texttt{data} \mid \texttt{default} \mid \texttt{deriving} \mid \texttt{do} \mid \texttt{else} \\
 & \mid & \texttt{foreign} \mid \texttt{if} \mid \texttt{import} \mid \texttt{in} \mid \texttt{infix} \mid \texttt{infixl} \\
 & \mid & \texttt{infixr} \mid \texttt{instance} \mid \texttt{let} \mid \texttt{module} \mid \texttt{newtype} \mid \texttt{of} \\
 & \mid & \texttt{then} \mid \texttt{type} \mid \texttt{where} \mid \texttt{\_}
\end{array}
$$

An identifier consists of a letter followed by zero or more letters, digits, underscores, and single quotes. Identifiers are lexically distinguished into two namespaces (Section 1.4): those that begin with a lowercase letter (variable identifiers) and those that begin with an upper-case letter (constructor

identifiers). Identifiers are case sensitive: `name`, `naMe`, and `Name` are three distinct identifiers (the first two are variable identifiers, the last is a constructor identifier).

Underscore, "_", is treated as a lowercase letter, and can occur wherever a lowercase letter can. However, "_" all by itself is a reserved identifier, used as wild card in patterns. Compilers that offer warnings for unused identifiers are encouraged to suppress such warnings for identifiers beginning with underscore. This allows programmers to use "`_foo`" for a parameter that they expect to be unused.

$$
\begin{array}{lcl}
varsym & \rightarrow & \left(symbol_{\langle : \rangle}\{symbol\}\right)_{\langle reservedop \mid dashes \rangle} \\
consym & \rightarrow & (: \{symbol\})_{\langle reservedop \rangle} \\
reservedop & \rightarrow & .. \mid : \mid :: \mid = \mid \backslash \mid \mid \mid <- \mid -> \mid @ \mid \sim \mid =>
\end{array}
$$

Operator symbols are formed from one or more symbol characters, as defined above, and are lexically distinguished into two namespaces (Section 1.4):

- An operator symbol starting with a colon is a constructor.
- An operator symbol starting with any other character is an ordinary identifier.

Notice that a colon by itself, "`:`", is reserved solely for use as the Haskell list constructor; this makes its treatment uniform with other parts of list syntax, such as "`[]`" and "`[a,b]`".

Other than the special syntax for prefix negation, all operators are infix, although each infix operator can be used in a *section* to yield partially applied operators (see Section 3.5). All of the standard infix operators are just predefined symbols and may be rebound.

In the remainder of the report six different kinds of names will be used:

$varid$			(variables)
$conid$			(constructors)
$tyvar$	$\rightarrow$	$varid$	(type variables)
$tycon$	$\rightarrow$	$conid$	(type constructors)
$tycls$	$\rightarrow$	$conid$	(type classes)
$modid$	$\rightarrow$	$\{conid\ .\}\ conid$	(modules)

Variables and type variables are represented by identifiers beginning with small letters, and the others by identifiers beginning with capitals; also, variables and constructors have infix forms, the other four do not. Module names are a dot-separated sequence of *conid*s. Namespaces are also discussed in Section 1.4.

A name may optionally be *qualified* in certain circumstances by prepending them with a module identifier. This applies to variable, constructor, type constructor and type class names, but not type variables or module names. Qualified names are discussed in detail in Chapter 5.

$$
\begin{array}{lcl}
qvarid & \rightarrow & [modid\ .]\ varid \\
qconid & \rightarrow & [modid\ .]\ conid \\
qtycon & \rightarrow & [modid\ .]\ tycon \\
qtycls & \rightarrow & [modid\ .]\ tycls \\
qvarsym & \rightarrow & [modid\ .]\ varsym \\
qconsym & \rightarrow & [modid\ .]\ consym
\end{array}
$$

Since a qualified name is a lexeme, no spaces are allowed between the qualifier and the name. Sample lexical analyses are shown below.

This	Lexes as this
`f.g`	`f . g` (three tokens)
`F.g`	`F.g` (qualified 'g')
`f..`	`f ..` (two tokens)
`F..`	`F..` (qualified '.')
`F.`	`F .` (two tokens)

The qualifier does not change the syntactic treatment of a name; for example, `Prelude.+` is an infix operator with the same fixity as the definition of + in the Prelude (Section 4.4.2).

2.5 Numeric Literals

$$
\begin{array}{lcl}
\textit{decimal} & \rightarrow & \textit{digit}\{\textit{digit}\} \\
\textit{octal} & \rightarrow & \textit{octit}\{\textit{octit}\} \\
\textit{hexadecimal} & \rightarrow & \textit{hexit}\{\textit{hexit}\} \\
\textit{integer} & \rightarrow & \textit{decimal} \\
 & | & \texttt{0o}\ \textit{octal} \mid \texttt{0O}\ \textit{octal} \\
 & | & \texttt{0x}\ \textit{hexadecimal} \mid \texttt{0X}\ \textit{hexadecimal} \\
\textit{float} & \rightarrow & \textit{decimal}\ .\ \textit{decimal}\ [\textit{exponent}] \\
 & | & \textit{decimal}\ \textit{exponent} \\
\textit{exponent} & \rightarrow & (\texttt{e} \mid \texttt{E})[+ \mid \texttt{-}]\ \textit{decimal}
\end{array}
$$

There are two distinct kinds of numeric literals: integer and floating. Integer literals may be given in decimal (the default), octal (prefixed by `0o` or `0O`) or hexadecimal notation (prefixed by `0x` or `0X`). Floating literals are always decimal. A floating literal must contain digits both before and after the decimal point; this ensures that a decimal point cannot be mistaken for another use of the dot character. Negative numeric literals are discussed in Section 3.4. The typing of numeric literals is discussed in Section 6.4.1.

2.6 Character and String Literals

$$
\begin{array}{lcl}
\textit{char} & \rightarrow & \texttt{'}\ \left(\textit{graphic}_{\langle \texttt{'} \mid \texttt{\textbackslash} \rangle} \mid \textit{space} \mid \textit{escape}_{\langle \texttt{\textbackslash\&} \rangle}\right)\ \texttt{'} \\
\textit{string} & \rightarrow & \texttt{"}\ \left\{\textit{graphic}_{\langle \texttt{"} \mid \texttt{\textbackslash} \rangle} \mid \textit{space} \mid \textit{escape} \mid \textit{gap}\right\}\ \texttt{"} \\
\textit{escape} & \rightarrow & \texttt{\textbackslash}\ (\textit{charesc} \mid \textit{ascii} \mid \textit{decimal} \mid \texttt{o}\ \textit{octal} \mid \texttt{x}\ \textit{hexadecimal}) \\
\textit{charesc} & \rightarrow & \texttt{a} \mid \texttt{b} \mid \texttt{f} \mid \texttt{n} \mid \texttt{r} \mid \texttt{t} \mid \texttt{v} \mid \texttt{\textbackslash} \mid \texttt{"} \mid \texttt{'} \mid \texttt{\&} \\
\textit{ascii} & \rightarrow & \texttt{\^{}}\ \textit{cntrl} \mid \texttt{NUL} \mid \texttt{SOH} \mid \texttt{STX} \mid \texttt{ETX} \mid \texttt{EOT} \mid \texttt{ENQ} \mid \texttt{ACK} \\
 & | & \texttt{BEL} \mid \texttt{BS} \mid \texttt{HT} \mid \texttt{LF} \mid \texttt{VT} \mid \texttt{FF} \mid \texttt{CR} \mid \texttt{SO} \mid \texttt{SI} \mid \texttt{DLE}
\end{array}
$$

| | DC1 | DC2 | DC3 | DC4 | NAK | SYN | ETB | CAN

| | EM | SUB | ESC | FS | GS | RS | US | SP | DEL

cntrl → *ascLarge* | @ | [| \ |] | ^ | _

gap → \ *whitechar* { *whitechar* } \

Character literals are written between single quotes, as in `'a'`, and strings between double quotes, as in `"Hello"`.

Escape codes may be used in characters and strings to represent special characters. Note that a single quote `'` may be used in a string, but must be escaped in a character; similarly, a double quote `"` may be used in a character, but must be escaped in a string. `\` must always be escaped. The category *charesc* also includes portable representations for the characters "alert" (`\a`), "backspace" (`\b`), "form feed" (`\f`), "new line" (`\n`), "carriage return" (`\r`), "horizontal tab" (`\t`), and "vertical tab" (`\v`).

Escape characters for the Unicode character set, including control characters such as `\^X`, are also provided. Numeric escapes such as `\137` are used to designate the character with decimal representation 137; octal (e.g. `\o137`) and hexadecimal (e.g. `\x37`) representations are also allowed.

Consistent with the "maximal munch" rule, numeric escape characters in strings consist of all consecutive digits and may be of arbitrary length. Similarly, the one ambiguous ASCII escape code, `"\SOH"`, is parsed as a string of length 1. The escape character `\&` is provided as a "null character" to allow strings such as `"\137\&\9"` and `"\SO\&\H"` to be constructed (both of length two). Thus `"\&"` is equivalent to `""` and the character `\&` is disallowed. Further equivalences of characters are defined in Section 6.1.2.

A string may include a "gap"—two backslants enclosing white characters—which is ignored. This allows one to write long strings on more than one line by writing a backslant at the end of one line and at the start of the next. For example,

```
"Here is a backslant \\ as well as \137, \
    \a numeric escape character, and \^X, a control character."
```

String literals are actually abbreviations for lists of characters (see Section 3.7).

2.7 Layout

Haskell permits the omission of the braces and semicolons used in several grammar productions, by using *layout* to convey the same information. This allows both layout-sensitive and layout-insensitive styles of coding, which can be freely mixed within one program. Because layout is not required, Haskell programs can be straightforwardly produced by other programs.

The effect of layout on the meaning of a Haskell program can be completely specified by adding braces and semicolons in places determined by the layout. The meaning of this augmented program is now layout insensitive.

Informally stated, the braces and semicolons are inserted as follows. The layout (or "off-side") rule takes effect whenever the open brace is omitted after the keyword `where`, `let`, `do`, or `of`. When this happens, the indentation of the next lexeme (whether or not on a new line) is remembered and the omitted open brace is inserted (the whitespace preceding the lexeme may include comments). For each subsequent line, if it contains only whitespace or is indented more, then the previous item is continued (nothing is inserted); if it is indented the same amount, then a new item begins (a semicolon is inserted); and if it

is indented less, then the layout list ends (a close brace is inserted). If the indentation of the non-brace lexeme immediately following a `where`, `let`, `do` or `of` is less than or equal to the current indentation level, then instead of starting a layout, an empty list "`{}`" is inserted, and layout processing occurs for the current level (i.e. insert a semicolon or close brace). A close brace is also inserted whenever the syntactic category containing the layout list ends; that is, if an illegal lexeme is encountered at a point where a close brace would be legal, a close brace is inserted. The layout rule matches only those open braces that it has inserted; an explicit open brace must be matched by an explicit close brace. Within these explicit open braces, *no* layout processing is performed for constructs outside the braces, even if a line is indented to the left of an earlier implicit open brace.

Section 10.3 gives a more precise definition of the layout rules.

Given these rules, a single newline may actually terminate several layout lists. Also, these rules permit:

```
f x = let a = 1; b = 2
          g y = exp2
      in exp1
```

making `a`, `b` and `g` all part of the same layout list.

As an example, Figure 1 shows a (somewhat contrived) module and Figure 2 shows the result of applying the layout rule to it. Note in particular: (a) the line beginning `}};pop`, where the termination of the previous line invokes three applications of the layout rule, corresponding to the depth (3) of the nested `where` clauses, (b) the close braces in the `where` clause nested within the tuple and `case` expression, inserted because the end of the tuple was detected, and (c) the close brace at the very end, inserted because of the column 0 indentation of the end-of-file token.

```
module AStack( Stack, push, pop, top, size ) where
data Stack a = Empty
             | MkStack a (Stack a)

push :: a -> Stack a -> Stack a
push x s = MkStack x s

size :: Stack a -> Int
size s = length (stkToLst s)  where
           stkToLst  Empty         = []
           stkToLst (MkStack x s)  = x:xs where xs = stkToLst s

pop :: Stack a -> (a, Stack a)
pop (MkStack x s)
  = (x, case s of r -> i r where i x = x) -- (pop Empty) is an error

top :: Stack a -> a
top (MkStack x s) = x                    -- (top Empty) is an error
```

Figure 1: A sample program

```
module AStack( Stack, push, pop, top, size ) where
{data Stack a = Empty
              | MkStack a (Stack a)

;push :: a -> Stack a -> Stack a
;push x s = MkStack x s

;size :: Stack a -> Int
;size s = length (stkToLst s)  where
           {stkToLst  Empty         = []
           ;stkToLst (MkStack x s)  = x:xs where {xs = stkToLst s

}};pop :: Stack a -> (a, Stack a)
;pop (MkStack x s)
  = (x, case s of {r -> i r where {i x = x}}) -- (pop Empty) is an error

;top :: Stack a -> a
;top (MkStack x s) = x                        -- (top Empty) is an error
}
```

Figure 2: Sample program with layout expanded

Chapter 3

Expressions

In this chapter, we describe the syntax and informal semantics of Haskell *expressions*, including their translations into the Haskell kernel, where appropriate. Except in the case of `let` expressions, these translations preserve both the static and dynamic semantics. Free variables and constructors used in these translations always refer to entities defined by the `Prelude`. For example, "`concatMap`" used in the translation of list comprehensions (Section 3.11) means the `concatMap` defined by the `Prelude`, regardless of whether or not the identifier "`concatMap`" is in scope where the list comprehension is used, and (if it is in scope) what it is bound to.

exp	→	*infixexp* :: [*context* =>] *type*	(expression type signature)
	|	*infixexp*	
infixexp	→	*lexp qop infixexp*	
	|	- *infixexp*	(prefix negation)
	|	*lexp*	
lexp	→	\ $apat_1 ... apat_n$ -> *exp*	(lambda abstraction, $n \geq 1$)
	|	`let` *decls* `in` *exp*	(let expression)
	|	`if` *exp* [;] `then` *exp* [;] `else` *exp*	(conditional)
	|	`case` *exp* `of` { *alts* }	(case expression)
	|	`do` { *stmts* }	(do expression)
	|	*fexp*	
fexp	→	[*fexp*] *aexp*	(function application)
aexp	→	*qvar*	(variable)
	|	*gcon*	(general constructor)
	|	*literal*	
	|	(*exp*)	(parenthesized expression)
	|	(exp_1 , ... , exp_k)	(tuple, $k \geq 2$)
	|	[exp_1 , ... , exp_k]	(list, $k \geq 1$)
	|	[exp_1[, exp_2] .. [exp_3]]	(arithmetic sequence)
	|	[*exp* | $qual_1$, ... , $qual_n$]	(list comprehension, $n \geq 1$)
	|	(*infixexp qop*)	(left section)
	|	($qop_{\langle - \rangle}$ *infixexp*)	(right section)
	|	*qcon* { $fbind_1$, ... , $fbind_n$ }	(labeled construction, $n \geq 0$)
	|	$aexp_{\langle qcon \rangle}$ { $fbind_1$, ... , $fbind_n$ }	(labeled update, $n \geq 1$)

Expressions involving infix operators are disambiguated by the operator's fixity (see Section 4.4.2). Consecutive unparenthesized operators with the same precedence must both be either left or right associative to avoid a syntax error. Given an unparenthesized expression "$x\ qop^{(a,i)}\ y\ qop^{(b,j)}\ z$" (where $qop^{(a,i)}$ means an operator with associativity a and precedence i), parentheses must be added around either $x\ qop^{(a,i)}\ y$ or $y\ qop^{(b,j)}\ z$ when $i = j$ unless $a = b = \mathrm{l}$ or $a = b = \mathrm{r}$.

An example algorithm for resolving expressions involving infix operators is given in Section 10.6.

Negation is the only prefix operator in Haskell; it has the same precedence as the infix - operator defined in the Prelude (see Section 4.4.2, Figure 1).

The grammar is ambiguous regarding the extent of lambda abstractions, let expressions, and conditionals. The ambiguity is resolved by the meta-rule that each of these constructs extends as far to the right as possible.

Sample parses are shown below.

This	Parses as
`f x + g y`	`(f x) + (g y)`
`- f x + y`	`(- (f x)) + y`
`let {...} in x + y`	`let {...} in (x + y)`
`z + let {...} in x + y`	`z + (let {...} in (x + y))`
`f x y :: Int`	`(f x y) :: Int`
`\ x -> a+b :: Int`	`\x -> ((a+b) :: Int)`

For the sake of clarity, the rest of this section will assume that expressions involving infix operators have been resolved according to the fixities of the operators.

3.1 Errors

Errors during expression evaluation, denoted by $\bot$ ("bottom"), are indistinguishable by a Haskell program from non-termination. Since Haskell is a non-strict language, all Haskell types include $\bot$. That is, a value of any type may be bound to a computation that, when demanded, results in an error. When evaluated, errors cause immediate program termination and cannot be caught by the user. The Prelude provides two functions to directly cause such errors:

```
error     :: String -> a
undefined :: a
```

A call to `error` terminates execution of the program and returns an appropriate error indication to the operating system. It should also display the string in some system-dependent manner. When `undefined` is used, the error message is created by the compiler.

Translations of Haskell expressions use `error` and `undefined` to explicitly indicate where execution time errors may occur. The actual program behavior when an error occurs is up to the implementation. The messages passed to the `error` function in these translations are only suggestions; implementations may choose to display more or less information when an error occurs.

3.2 Variables, Constructors, Operators, and Literals

aexp	→	*qvar*	(variable)
	|	*gcon*	(general constructor)
	|	*literal*	
gcon	→	()	
	|	[]	
	|	(, {,})	
	|	*qcon*	
var	→	*varid* | (*varsym*)	(variable)
qvar	→	*qvarid* | (*qvarsym*)	(qualified variable)
con	→	*conid* | (*consym*)	(constructor)
qcon	→	*qconid* | (*gconsym*)	(qualified constructor)
varop	→	*varsym* | ` *varid* `	(variable operator)
qvarop	→	*qvarsym* | ` *qvarid* `	(qualified variable operator)
conop	→	*consym* | ` *conid* `	(constructor operator)
qconop	→	*gconsym* | ` *qconid* `	(qualified constructor operator)
op	→	*varop* | *conop*	(operator)
qop	→	*qvarop* | *qconop*	(qualified operator)
gconsym	→	: | *qconsym*	

Haskell provides special syntax to support infix notation. An *operator* is a function that can be applied using infix syntax (Section 3.4), or partially applied using a *section* (Section 3.5).

An *operator* is either an *operator symbol*, such as `+` or `$$`, or is an ordinary identifier enclosed in grave accents (backquotes), such as `` `op` ``. For example, instead of writing the prefix application `op x y`, one can write the infix application `x `op` y`. If no fixity declaration is given for `op` then it defaults to highest precedence and left associativity (see Section 4.4.2).

Dually, an operator symbol can be converted to an ordinary identifier by enclosing it in parentheses. For example, `(+) x y` is equivalent to `x + y`, and `foldr (*) 1 xs` is equivalent to `foldr (\x y -> x*y) xs`.

Special syntax is used to name some constructors for some of the built-in types, as found in the production for *gcon* and *literal*. These are described in Section 6.1.

An integer literal represents the application of the function `fromInteger` to the appropriate value of type `Integer`. Similarly, a floating point literal stands for an application of `fromRational` to a value of type `Rational` (that is, `Ratio Integer`).

> **Translation:** The integer literal i is equivalent to `fromInteger` i, where `fromInteger` is a method in class `Num` (see Section 6.4.1).
>
> The floating point literal f is equivalent to `fromRational` $(n$ `Ratio.%` $d)$, where `fromRational` is a method in class `Fractional` and `Ratio.%` constructs a rational from two integers, as defined in the `Ratio` library. The integers n and d are chosen so that $n/d = f$.

3.3 Curried Applications and Lambda Abstractions

fexp	→	[*fexp*] *aexp*	(function application)
lexp	→	\ $apat_1$ ... $apat_n$ -> *exp*	(lambda abstraction $n \geq 1$)

Function application is written $e_1\ e_2$. Application associates to the left, so the parentheses may be omitted in `(f x) y`. Because e_1 could be a data constructor, partial applications of data constructors are allowed.

Lambda abstractions are written \ $p_1 ... p_n$ -> e, where the p_i are *patterns*. An expression such as `\x:xs->x` is syntactically incorrect; it may legally be written as `\(x:xs)->x`.

The set of patterns must be *linear*—no variable may appear more than once in the set.

Translation: The following identity holds:

$$\backslash\ p_1 \ldots p_n \texttt{ -> } e = \backslash\ x_1 \ldots x_n \texttt{ -> case } (x_1, \ldots, x_n) \texttt{ of } (p_1, \ldots, p_n) \texttt{ -> } e$$

where the x_i are new identifiers.

Given this translation combined with the semantics of case expressions and pattern matching described in Section 3.17.3, if the pattern fails to match, then the result is $\bot$.

3.4 Operator Applications

infixexp	→	*lexp qop infixexp*	
	|	- *infixexp*	(prefix negation)
	|	*lexp*	
qop	→	*qvarop* | *qconop*	(qualified operator)

The form e_1 *qop* e_2 is the infix application of binary operator *qop* to expressions e_1 and e_2.

The special form $-e$ denotes prefix negation, the only prefix operator in Haskell, and is syntax for `negate` (e). The binary `-` operator does not necessarily refer to the definition of `-` in the Prelude; it may be rebound by the module system. However, unary `-` will always refer to the `negate` function defined in the Prelude. There is no link between the local meaning of the `-` operator and unary negation.

Prefix negation has the same precedence as the infix operator `-` defined in the Prelude (see Table 1). Because `e1-e2` parses as an infix application of the binary operator `-`, one must write `e1(-e2)` for the alternative parsing. Similarly, `(-)` is syntax for `\x y -> x-y`, as with any infix operator, and does not denote `\x -> -x`—one must use `negate` for that.

Translation: The following identities hold:

$$e_1\ op\ e_2 = (op)\ e_1\ e_2$$
$$-e = \texttt{negate } (e)$$

3.5 Sections

$$
\begin{array}{lcll}
aexp & \to & \texttt{(}\ infixexp\ qop\ \texttt{)} & \text{(left section)} \\
 & | & \texttt{(}\ qop_{\langle - \rangle}\ infixexp\ \texttt{)} & \text{(right section)}
\end{array}
$$

Sections are written as $(op\ e)$ or $(e\ op)$, where *op* is a binary operator and e is an expression. Sections are a convenient syntax for partial application of binary operators.

Syntactic precedence rules apply to sections as follows. $(op\ e)$ is legal if and only if $(x\ op\ e)$ parses in the same way as $(x\ op\ (e))$; and similarly for $(e\ op$. For example, `(*a+b)` is syntactically invalid, but `(+a*b)` and `(*(a+b))` are valid. Because `(+)` is left associative, `(a+b+)` is syntactically correct, but `(+a+b)` is not; the latter may legally be written as `(+(a+b))`. As another example, the expression

```
(let n = 10 in n +)
```

is invalid because, by the let/lambda meta-rule (Section 3), the expression

```
(let n = 10 in n + x)
```

parses as

```
(let n = 10 in (n + x))
```

rather than

```
((let n = 10 in n) + x)
```

Because `-` is treated specially in the grammar, $(-\ exp)$ is not a section, but an application of prefix negation, as described in the preceding section. However, there is a `subtract` function defined in the Prelude such that (`subtract` exp) is equivalent to the disallowed section. The expression $(+(-\ exp))$ can serve the same purpose.

Translation: The following identities hold:

$$(op\ e) = \backslash\ x \texttt{ -> } x\ op\ e$$
$$(e\ op) = \backslash\ x \texttt{ -> } e\ op\ x$$

where *op* is a binary operator, e is an expression, and x is a variable that does not occur free in e.

3.6 Conditionals

$$lexp \to \texttt{if}\ exp\ [;]\ \texttt{then}\ exp\ [;]\ \texttt{else}\ exp$$

A *conditional expression* has the form `if` e_1 `then` e_2 `else` e_3 and returns the value of e_2 if the value of e_1 is `True`, e_3 if e_1 is `False`, and $\bot$ otherwise.

Translation: The following identity holds:

$$\texttt{if } e_1 \texttt{ then } e_2 \texttt{ else } e_3 = \texttt{case } e_1 \texttt{ of } \{\texttt{True -> } e_2 \texttt{ ; False -> } e_3\}$$

where `True` and `False` are the two nullary constructors from the type `Bool`, as defined in the Prelude. The type of e_1 must be `Bool`; e_2 and e_3 must have the same type, which is also the type of the entire conditional expression.

3.7 Lists

$$\begin{array}{lcl}
infixexp & \rightarrow & exp_1\ qop\ exp_2 \\
aexp & \rightarrow & [\ exp_1\ ,\ \ldots\ ,\ exp_k\ ] \quad (k \geq 1) \\
 & | & gcon \\
gcon & \rightarrow & [] \\
 & | & qcon \\
qcon & \rightarrow & (\ gconsym\ ) \\
qop & \rightarrow & qconop \\
qconop & \rightarrow & gconsym \\
gconsym & \rightarrow & :
\end{array}$$

Lists are written $[e_1, \ldots, e_k]$, where $k \geq 1$. The list constructor is `:`, and the empty list is denoted `[]`. Standard operations on lists are given in the Prelude (see Section 6.1.3, and Chapter 9 notably Section 9.1).

Translation: The following identity holds:

$$[e_1, \ldots, e_k] = e_1 : (e_2 : (\ldots (e_k : [\,])))$$

where `:` and `[]` are constructors for lists, as defined in the Prelude (see Section 6.1.3). The types of e_1 through e_k must all be the same (call it t), and the type of the overall expression is `[t]` (see Section 4.1.2).

The constructor "`:`" is reserved solely for list construction; like `[]`, it is considered part of the language syntax, and cannot be hidden or redefined. It is a right-associative operator, with precedence level 5 (Section 4.4.2).

3.8 Tuples

$$\begin{array}{lcl}
aexp & \rightarrow & (\ exp_1\ ,\ \ldots\ ,\ exp_k\ ) \quad (k \geq 2) \\
 & | & qcon \\
qcon & \rightarrow & (\ ,\{,\})
\end{array}$$

Tuples are written $(e_1, ..., e_k)$, and may be of arbitrary length $k \geq 2$. The constructor for an n-tuple is denoted by $(, ...,)$, where there are $n - 1$ commas. Thus `(a,b,c)` and `(,,) a b c` denote the same value. Standard operations on tuples are given in the Prelude (see Section 6.1.4 and Chapter 9).

> **Translation:** $(e_1, ..., e_k)$ for $k \geq 2$ is an instance of a k-tuple as defined in the Prelude, and requires no translation. If t_1 through t_k are the types of e_1 through e_k, respectively, then the type of the resulting tuple is $(t_1, ..., t_k)$ (see Section 4.1.2).

3.9 Unit Expressions and Parenthesized Expressions

$$
\begin{array}{lcl}
aexp & \rightarrow & gcon \\
 & | & (\ exp \ ) \\
gcon & \rightarrow & ()
\end{array}
$$

The form (e) is simply a *parenthesized expression*, and is equivalent to e. The *unit expression* `()` has type `()` (see Section 4.1.2). It is the only member of that type apart from $\bot$, and can be thought of as the "nullary tuple" (see Section 6.1.5).

> **Translation:** (e) is equivalent to e.

3.10 Arithmetic Sequences

$$aexp \rightarrow [\ exp_1 [, \ exp_2] \ .. \ [exp_3] \ ]$$

The *arithmetic sequence* $[e_1, e_2..e_3]$ denotes a list of values of type t, where each of the e_i has type t, and t is an instance of class `Enum`.

> **Translation:** Arithmetic sequences satisfy these identities:
>
> $$
> \begin{array}{lcl}
> \texttt{[} \ e_1 \ \texttt{..]} & = & \texttt{enumFrom} \ e_1 \\
> \texttt{[} \ e_1, e_2 \ \texttt{..]} & = & \texttt{enumFromThen} \ e_1 \ e_2 \\
> \texttt{[} \ e_1 \ \texttt{..} \ e_2 \ \texttt{]} & = & \texttt{enumFromTo} \ e_1 \ e_2 \\
> \texttt{[} \ e_1, e_2 \ \texttt{..} \ e_3 \ \texttt{]} & = & \texttt{enumFromThenTo} \ e_1 \ e_2 \ e_3
> \end{array}
> $$
>
> where `enumFrom`, `enumFromThen`, `enumFromTo`, and `enumFromThenTo` are class methods in the class `Enum` as defined in the Prelude (see Figure 3).

The semantics of arithmetic sequences therefore depends entirely on the instance declaration for the type `t`. See Section 6.3.4 for more details of which `Prelude` types are in `Enum` and their semantics.

3.11 List Comprehensions

aexp	→	[*exp* | *qual*$_1$, ... , *qual*$_n$]	(list comprehension, $n \geq 1$)
qual	→	*pat* <- *exp*	(generator)
	|	`let` *decls*	(local declaration)
	|	*exp*	(boolean guard)

A *list comprehension* has the form $[e \mid q_1, ..., q_n]$, $n \geq 1$, where the q_i qualifiers are either

- *generators* of the form p <- e, where p is a pattern (see Section 3.17) of type t and e is an expression of type $[t]$
- *local bindings* that provide new definitions for use in the generated expression e or subsequent boolean guards and generators
- *boolean guards*, which are arbitrary expressions of type `Bool`.

Such a list comprehension returns the list of elements produced by evaluating e in the successive environments created by the nested, depth-first evaluation of the generators in the qualifier list. Binding of variables occurs according to the normal pattern matching rules (see Section 3.17), and if a match fails then that element of the list is simply skipped over. Thus:

```
[ x |  xs   <- [ [(1,2),(3,4)], [(5,4),(3,2)] ],
      (3,x) <- xs ]
```

yields the list `[4,2]`. If a qualifier is a boolean guard, it must evaluate to `True` for the previous pattern match to succeed. As usual, bindings in list comprehensions can shadow those in outer scopes; for example:

$$[x \mid x \text{ <- } x, x \text{ <- } x] = [z \mid y \text{ <- } x, z \text{ <- } y]$$

Translation: List comprehensions satisfy these identities, which may be used as a translation into the kernel:

$$\begin{array}{lcl}
[e \mid \texttt{True}] & = & [e] \\
[e \mid q] & = & [e \mid q, \texttt{True}] \\
[e \mid b, Q] & = & \texttt{if } b \texttt{ then } [e, Q] \texttt{ else } [] \\
[e \mid p \text{ <- } l, Q] & = & \texttt{let ok} = [e \mid Q] \\
 & & \qquad \texttt{ok _} = [] \\
 & & \texttt{in concatMap ok } l \\
[e \mid \texttt{let } decls, Q] & = & \texttt{let } decls \texttt{ in } [e \mid Q]
\end{array}$$

where e ranges over expressions, p over patterns, l over list-valued expressions, b over boolean expressions, *decls* over declaration lists, q over qualifiers, and Q over sequences of qualifiers. `ok` is a fresh variable. The function `concatMap`, and boolean value `True`, are defined in the Prelude.

As indicated by the translation of list comprehensions, variables bound by `let` have fully polymorphic types while those defined by <- are lambda bound and are thus monomorphic (see Section 4.5.4).

3.12 Let Expressions

lexp → `let` *decls* `in` *exp*

Let expressions have the general form `let` $\{d_1; ...; d_n\}$ `in` e, and introduce a nested, lexically-scoped, mutually-recursive list of declarations (`let` is often called `letrec`). The scope of the declarations is the expression e and the right hand side of the declarations. Declarations are described in Chapter 4. Pattern bindings are matched lazily; an implicit ~ makes these patterns irrefutable. For example,

$$\texttt{let } (x, y) = \texttt{undefined in } e$$

does not cause an execution-time error until `x` or `y` is evaluated.

Translation: The dynamic semantics of the expression `let` $\{d_1; ...; d_n\}$ `in` e_0 are captured by this translation: After removing all type signatures, each declaration d_i is translated into an equation of the form $p_i = e_i$, where p_i and e_i are patterns and expressions respectively, using the translation in Section 4.4.3. Once done, these identities hold, which may be used as a translation into the kernel:

`let` $\{p_1 = e_1; ...; p_n = e_n\}$ `in` e_0	=	`let`$(\sim p_1, ..., \sim p_n) = (e_1, ..., e_n)$ `in` e_0
`let` $p = e_1$ `in` e_0	=	`case` e_1 `of` $\sim p$ -> e_0
		where no variable in p appears free in e_1
`let` $p = e_1$ `in` e_0	=	`let` $p =$ `fix` $(\backslash \sim p$ -> $e_1)$ `in` e_0

where `fix` is the least fixpoint operator. Note the use of the irrefutable patterns ~p. This translation does not preserve the static semantics because the use of `case` precludes a fully polymorphic typing of the bound variables. The static semantics of the bindings in a `let` expression are described in Section 4.4.3.

3.13 Case Expressions

lexp	→	`case` *exp* `of` { *alts* }	
alts	→	alt_1 ; ... ; alt_n	$(n \geq 1)$
alt	→	*pat* -> *exp* [`where` *decls*]	
	|	*pat gdpat* [`where` *decls*]	
	|		(empty alternative)
gdpat	→	*guards* -> *exp* [*gdpat*]	
guards	→	| $guard_1$, ... , $guard_n$	$(n \geq 1)$
guard	→	*pat* <- *infixexp*	(pattern guard)
	|	`let` *decls*	(local declaration)
	|	*infixexp*	(boolean guard)

A *case expression* has the general form

$$\texttt{case}\ e\ \texttt{of}\ \{p_1 match_1; ...; p_n match_n\}$$

where each $match_i$ is of the general form

$$\begin{array}{l} \mid gs_{i1}\ \texttt{->}\ e_{i1} \\ \ldots \\ \mid gs_{im_i}\ \texttt{->}\ e_{im_i} \\ \texttt{where}\ decls_i \end{array}$$

(Notice that in the syntax rule for *guards*, the "`|`" is a terminal symbol, not the syntactic metasymbol for alternation.) Each alternative $p_i match_i$ consists of a pattern p_i and its matches, $match_i$. Each match in turn consists of a sequence of pairs of guards gs_{ij} and bodies e_{ij} (expressions), followed by optional bindings ($decls_i$) that scope over all of the guards and expressions of the alternative.

A *guard* has one of the following forms:

- *pattern guards* are of the form p `<-` e, where p is a pattern (see Section 3.17) of type t and e is an expression type t[2]. They succeed if the expression e matches the pattern p, and introduce the bindings of the pattern to the environment.
- *local bindings* are of the form `let` $decls$. They always succeed, and they introduce the names defined in $decls$ to the environment.
- *boolean guards* are arbitrary expressions of type `Bool`. They succeed if the expression evaluates to `True`, and they do not introduce new names to the environment. A boolean guard, g, is semantically equivalent to the pattern guard `True` `<-` g.

An alternative of the form

$$pat\ \texttt{->}\ exp\ \texttt{where}\ decls$$

is treated as shorthand for:

$$\begin{array}{l} pat \mid \texttt{True}\ \texttt{->}\ exp \\ \texttt{where}\ decls \end{array}$$

A case expression must have at least one alternative and each alternative must have at least one body. Each body must have the same type, and the type of the whole expression is that type.

A case expression is evaluated by pattern matching the expression e against the individual alternatives. The alternatives are tried sequentially, from top to bottom. If e matches the pattern of an alternative, then the guarded expressions for that alternative are tried sequentially from top to bottom in the environment of the case expression extended first by the bindings created during the matching of the pattern, and then by the $decls_i$ in the `where` clause associated with that alternative.

For each guarded expression, the comma-separated guards are tried sequentially from left to right. If all of them succeed, then the corresponding expression is evaluated in the environment extended with the bindings introduced by the guards. That is, the bindings that are introduced by a guard (either by using a let clause or a pattern guard) are in scope in the following guards and the corresponding expression. If any of the guards fail, then this guarded expression fails and the next guarded expression is tried.

[2]Note that the syntax of a pattern guard is the same as that of a generator in a list comprehension. The contextual difference is that, in a list comprehension, a pattern of type t goes with an expression of type $[t]$.

If none of the guarded expressions for a given alternative succeed, then matching continues with the next alternative. If no alternative succeeds, then the result is $\bot$. Pattern matching is described in Section 3.17, with the formal semantics of case expressions in Section 3.17.3.

A note about parsing. The expression

```
case x of { (a,_) | let b = not a in b :: Bool -> a }
```

is tricky to parse correctly. It has a single unambiguous parse, namely

```
case x of { (a,_) | (let b = not a in b :: Bool) -> a }
```

However, the phrase `Bool -> a` is syntactically valid as a type, and parsers with limited lookahead may incorrectly commit to this choice, and hence reject the program. Programmers are advised, therefore, to avoid guards that end with a type signature — indeed that is why a *guard* contains an $infixexp$ not an exp.

3.14 Do Expressions

$lexp$	$\rightarrow$	`do` { $stmts$ }	(do expression)
$stmts$	$\rightarrow$	$stmt_1 ... stmt_n$ exp [;]	($n \geq 0$)
$stmt$	$\rightarrow$	exp ;	
	$\mid$	pat <- exp ;	
	$\mid$	`let` $decls$;	
	$\mid$	;	(empty statement)

A *do expression* provides a more conventional syntax for monadic programming. It allows an expression such as

```
putStr "x: "     >>
getLine          >>= \l ->
return (words l)
```

to be written in a more traditional way as:

```
do putStr "x: "
   l <- getLine
   return (words l)
```

Translation: Do expressions satisfy these identities, which may be used as a translation into the kernel, after eliminating empty *stmts*:

$$
\begin{array}{lcl}
\texttt{do}\ \{e\} & = & e \\
\texttt{do}\ \{e; stmts\} & = & e >> \texttt{do}\ \{stmts\} \\
\texttt{do}\ \{p <\text{-}\ e; stmts\} & = & \texttt{let ok}\ p = \texttt{do}\ \{stmts\} \\
 & & \qquad \texttt{ok \_} = \texttt{fail "..."} \\
 & & \texttt{in}\ e >>= \texttt{ok} \\
\texttt{do}\ \{\texttt{let}\ decls; stmts\} & = & \texttt{let}\ decls\ \texttt{in do}\ \{stmts\}
\end{array}
$$

The ellipsis "..." stands for a compiler-generated error message, passed to `fail`, preferably giving some indication of the location of the pattern-match failure; the functions `>>`, `>>=`, and `fail` are operations in the class `Monad`, as defined in the Prelude; and `ok` is a fresh identifier.

As indicated by the translation of `do`, variables bound by `let` have fully polymorphic types while those defined by `<-` are lambda bound and are thus monomorphic.

3.15 Datatypes with Field Labels

A datatype declaration may optionally define field labels (see Section 4.2.1). These field labels can be used to construct, select from, and update fields in a manner that is independent of the overall structure of the datatype.

Different datatypes cannot share common field labels in the same scope. A field label can be used at most once in a constructor. Within a datatype, however, a field label can be used in more than one constructor provided the field has the same typing in all constructors. To illustrate the last point, consider:

```
data S = S1 { x :: Int } | S2 { x :: Int }    -- OK
data T = T1 { y :: Int } | T2 { y :: Bool }   -- BAD
```

Here `S` is legal but `T` is not, because `y` is given inconsistent typings in the latter.

3.15.1 Field Selection

aexp → *qvar*

Field labels are used as selector functions. When used as a variable, a field label serves as a function that extracts the field from an object. Selectors are top level bindings and so they may be shadowed by local variables but cannot conflict with other top level bindings of the same name. This shadowing only affects selector functions; in record construction (Section 3.15.2) and update (Section 3.15.3), field labels cannot be confused with ordinary variables.

Translation: A field label f introduces a selector function defined as:

$$f x = \texttt{case}\ x\ \texttt{of}\ \{C_1 p_{11} \ldots p_{1k}\ \texttt{->}\ e_1; \ldots; C_n p_{n1} \ldots p_{nk}\ \texttt{->}\ e_n\}$$

where $C_1 \ldots C_n$ are all the constructors of the datatype containing a field labeled with f, p_{ij} is y when f labels the jth component of C_i or _ otherwise, and e_i is y when some field in C_i has a label of f or `undefined` otherwise.

3.15.2 Construction Using Field Labels

$$\begin{array}{lcl} aexp & \rightarrow & qcon\ \{\ fbind_1\ ,\ \ldots\ ,\ fbind_n\ \} \quad \text{(labeled construction, } n \geq 0\text{)} \\ fbind & \rightarrow & qvar = exp \end{array}$$

A constructor with labeled fields may be used to construct a value in which the components are specified by name rather than by position. Unlike the braces used in declaration lists, these are not subject to layout; the `{` and `}` characters must be explicit. (This is also true of field updates and field patterns.) Construction using field labels is subject to the following constraints:

- Only field labels declared with the specified constructor may be mentioned.
- A field label may not be mentioned more than once.
- Fields not mentioned are initialized to $\bot$.
- A compile-time error occurs when any strict fields (fields whose declared types are prefixed by `!`) are omitted during construction. Strict fields are discused in Section 4.2.1.

The expression `F {}`, where `F` is a data constructor, is legal *whether or not `F` was declared with record syntax* (provided `F` has no strict fields — see the fourth bullet above); it denotes $F \bot_1 \ldots \bot_n$, where n is the arity of `F`.

Translation: In the binding $f = v$, the field f labels v.

$$C\{bs\} \quad = \quad C\big(pick_1^C\ bs\ \texttt{undefined}\big) \ldots \big(pick_k^C\ bs\ \texttt{undefined}\big)$$

where k is the arity of C.

The auxiliary function $pick_i^C\ bs\ d$ is defined as follows:

> If the ith component of a constructor C has the field label f, and if $f = v$ appears in the binding list bs , then $pick_i^C\ bs\ d$ is v. Otherwise, $pick_i^C\ bs\ d$ is the default value d.

3.15.3 Updates Using Field Labels

$$aexp \quad \rightarrow \quad aexp_{\langle qcon \rangle}\ \{\ fbind_1\ ,\ \ldots\ ,\ fbind_n\ \} \quad \text{(labeled update, } n \geq 1\text{)}$$

Values belonging to a datatype with field labels may be non-destructively updated. This creates a new value in which the specified field values replace those in the existing value. Updates are restricted in the following ways:

- All labels must be taken from the same datatype.
- At least one constructor must define all of the labels mentioned in the update.
- No label may be mentioned more than once.
- An execution error occurs when the value being updated does not contain all of the specified labels.

Translation: Using the prior definition of *pick*,

$$
\begin{aligned}
e\{bs\} \;=\; & \texttt{case } e \texttt{ of} \\
& \quad C_1 v_1 ... v_{k_1} \texttt{ -> } C_1 \left(pick_1^{C_1}\ bs\ v_1\right) ... \left(pick_{k_1}^{C_1}\ bs\ v_{k_1}\right) \\
& \quad ... \\
& \quad C_j v_1 ... v_{k_j} \texttt{ -> } C_j \left(pick_1^{C_j}\ bs\ v_1\right) ... \left(pick_{k_j}^{C_j}\ bs\ v_{k_j}\right) \\
& \quad \texttt{\_ -> error "Update error"}
\end{aligned}
$$

where $\{C_1, ..., C_j\}$ is the set of constructors containing all labels in bs, and k_i is the arity of C_i.

Here are some examples using labeled fields:

```
data T    = C1 {f1,f2 :: Int}
          | C2 {f1 :: Int,
                f3,f4 :: Char}
```

Expression	Translation
`C1 {f1 = 3}`	`C1 3 undefined`
`C2 {f1 = 1, f4 = 'A', f3 = 'B'}`	`C2 1 'B' 'A'`
`x {f1 = 1}`	`case x of C1 _ f2 -> C1 1 f2` `C2 _ f3 f4 -> C2 1 f3 f4`

The field `f1` is common to both constructors in T. This example translates expressions using constructors in field-label notation into equivalent expressions using the same constructors without field labels. A compile-time error will result if no single constructor defines the set of field labels used in an update, such as `x {f2 = 1, f3 = 'x'}`.

3.16 Expression Type-Signatures

Expression type-signatures have the form $e :: t$, where e is an expression and t is a type (Section 4.1.2); they are used to type an expression explicitly and may be used to resolve ambiguous typings due to overloading (see Section 4.3.4). The value of the expression is just that of *exp*. As with normal type signatures (see Section 4.4.1), the declared type may be more specific than the principal type derivable from *exp*, but it is an error to give a type that is more general than, or not comparable to, the principal type.

Translation:

$$e :: t \;=\; \texttt{let } \{v :: t; v = e\} \texttt{ in } v$$

3.17 Pattern Matching

Patterns appear in lambda abstractions, function definitions, pattern bindings, list comprehensions, do expressions, and case expressions. However, the first five of these ultimately translate into case expressions, so defining the semantics of pattern matching for case expressions is sufficient.

3.17.1 Patterns

Patterns have this syntax:

pat	→	*lpat qconop pat*	(infix constructor)
	|	*lpat*	
lpat	→	*apat*	
	|	- (*integer* | *float*)	(negative literal)
	|	*gcon* $apat_1 ... apat_k$	(arity *gcon* $= k$, $k \geq 1$)
apat	→	*var* [@ *apat*]	(as pattern)
	|	*gcon*	(arity *gcon* $= 0$)
	|	*qcon* { $fpat_1$, ... , $fpat_k$ }	(labeled pattern, $k \geq 0$)
	|	*literal*	
	|	_	(wildcard)
	|	(*pat*)	(parenthesized pattern)
	|	(pat_1 , ... , *pat*)	(tuple pattern, $k \geq 2$)
	|	[pat_1 , ... , *pat*]	(list pattern, $k \geq 1$)
	|	~ *apat*	(irrefutable pattern)
fpat	→	*qvar* = *pat*	

All patterns must be *linear*—no variable may appear more than once. For example, this definition is illegal:

```
f (x,x) = x     -- ILLEGAL; x used twice in pattern
```

Patterns of the form *var*@*pat* are called *as-patterns*, and allow one to use *var* as a name for the value being matched by *pat*. For example,

```
case e of { xs@(x:rest) -> if x==0 then rest else xs }
```

is equivalent to:

```
let { xs = e } in
  case xs of { (x:rest) -> if x==0 then rest else xs }
```

Patterns of the form _ are *wildcards* and are useful when some part of a pattern is not referenced on the right-hand-side. It is as if an identifier not used elsewhere were put in its place. For example,

```
case e of { [x,_,_]  ->  if x==0 then True else False }
```

is equivalent to:

```
case e of { [x,y,z]  ->  if x==0 then True else False }
```

3.17.2 Informal Semantics of Pattern Matching

Patterns are matched against values. Attempting to match a pattern can have one of three results: it may *fail*; it may *succeed*, returning a binding for each variable in the pattern; or it may *diverge* (i.e. return $\bot$). Pattern matching proceeds from left to right, and outside to inside, according to the following rules:

1. Matching the pattern var against a value v always succeeds and binds var to v.
2. Matching the pattern $\sim apat$ against a value v always succeeds. The free variables in $apat$ are bound to the appropriate values if matching $apat$ against v would otherwise succeed, and to $\bot$ if matching $apat$ against v fails or diverges. (Binding does *not* imply evaluation.)

   Operationally, this means that no matching is done on a $\sim apat$ pattern until one of the variables in $apat$ is used. At that point the entire pattern is matched against the value, and if the match fails or diverges, so does the overall computation.
3. Matching the wildcard pattern _ against any value always succeeds, and no binding is done.
4. Matching the pattern $con\ pat$ against a value, where con is a constructor defined by `newtype`, depends on the value:
   - If the value is of the form $con\ v$, then pat is matched against v.
   - If the value is $\bot$, then pat is matched against $\bot$.

   That is, constructors associated with `newtype` serve only to change the type of a value.
5. Matching the pattern $con\ pat_1 ... pat_n$ against a value, where con is a constructor defined by `data`, depends on the value:
   - If the value is of the form $con\ v_1 ... v_n$, sub-patterns are matched left-to-right against the components of the data value; if all matches succeed, the overall match succeeds; the first to fail or diverge causes the overall match to fail or diverge, respectively.
   - If the value is of the form $con' v_1 ... v_m$, where con is a different constructor to con', the match fails.
   - If the value is $\bot$, the match diverges.
6. Matching against a constructor using labeled fields is the same as matching ordinary constructor patterns except that the fields are matched in the order they are named in the field list. All fields listed must be declared by the constructor; fields may not be named more than once. Fields not named by the pattern are ignored (matched against _).
7. Matching a numeric, character, or string literal pattern k against a value v succeeds if $v == k$, where == is overloaded based on the type of the pattern. The match diverges if this test diverges.

   The interpretation of numeric literals is exactly as described in Section 3.2; that is, the overloaded function `fromInteger` or `fromRational` is applied to an `Integer` or `Rational` literal (resp) to convert it to the appropriate type.
8. Matching an as-pattern $var@apat$ against a value v is the result of matching $apat$ against v, augmented with the binding of var to v. If the match of $apat$ against v fails or diverges, then so does the overall match.

Aside from the obvious static type constraints (for example, it is a static error to match a character against a boolean), the following static class constraints hold:

- An integer literal pattern can only be matched against a value in the class `Num`.
- A floating literal pattern can only be matched against a value in the class `Fractional`.

It is sometimes helpful to distinguish two kinds of patterns. Matching an *irrefutable pattern* is non-strict: the pattern matches even if the value to be matched is $\bot$. Matching a *refutable* pattern is strict: if the value to be matched is $\bot$ the match diverges. The irrefutable patterns are as follows: a variable, a wildcard, $N\ apat$ where N is a constructor defined by `newtype` and $apat$ is irrefutable (see

Section 4.2.3), *var*@*apat* where *apat* is irrefutable, or of the form ∼ *apat* (whether or not *apat* is irrefutable). All other patterns are *refutable*.

Here are some examples:

1. If the pattern `['a','b']` is matched against $['x', \bot]$, then `'a'` *fails* to match against `'x'`, and the result is a failed match. But if `['a','b']` is matched against $[\bot,' x']$, then attempting to match `'a'` against $\bot$ causes the match to *diverge*.
2. These examples demonstrate refutable vs. irrefutable matching:

`(\ ~(x,y) -> 0)` $\bot$	$\Rightarrow$	0
`(\  (x,y) -> 0)` $\bot$	$\Rightarrow$	$\bot$
`(\ ~[x] -> 0) []`	$\Rightarrow$	0
`(\ ~[x] -> x) []`	$\Rightarrow$	$\bot$
`(\ ~[x, ~(a,b)] -> x` $[(0,1), \bot]$	$\Rightarrow$	$(0,1)$
`(\ ~[x,  (a,b)] -> x` $[(0,1), \bot]$	$\Rightarrow$	$\bot$
`(\ (x:xs) -> x:x:xs)` $\bot$	$\Rightarrow$	$\bot$
`(\ ~(x:xs) -> x:x:xs)` $\bot$	$\Rightarrow$	$\bot : \bot : \bot$

3. Consider the following declarations:

   ```
   newtype N = N Bool
   data    D = D !Bool
   ```

   These examples illustrate the difference in pattern matching between types defined by `data` and `newtype`:

`(\ (N True) -> True)` $\bot$	$\Rightarrow$	$\bot$
`(\ (D True) -> True)` $\bot$	$\Rightarrow$	$\bot$
`(\ ~(D True) -> True)` $\bot$	$\Rightarrow$	`True`

   Additional examples may be found in Section 4.2.3.

Top level patterns in case expressions and the set of top level patterns in function or pattern bindings may have zero or more associated *guards*. See Section 3.13 for the syntax and semantics of guards.

The guard semantics have an influence on the strictness characteristics of a function or case expression. In particular, an otherwise irrefutable pattern may be evaluated because of a guard. For example, in

```
f :: (Int,Int,Int) -> [Int] -> Int
f ~(x,y,z) [a] | (a == y) = 1
```

both `a` and `y` will be evaluated by `==` in the guard.

3.17.3 Formal Semantics of Pattern Matching

The semantics of all pattern matching constructs other than `case` expressions are defined by giving identities that relate those constructs to `case` expressions. The semantics of `case` expressions themselves are in turn given as a series of identities, in Figure 1 – Figure 2. Any implementation should behave so that these identities hold; it is not expected that it will use them directly, since that would generate rather inefficient code.

In Figure 1 – Figure 2: e, e' and e_i are expressions; g_i and gs_i are guards and sequences of guards respectively; p and p_i are patterns; v, x, and x_i are variables; K and K' are algebraic datatype (`data`) constructors (including tuple constructors); and N is a `newtype` constructor.

Rule (b) matches a general source-language `case` expression, regardless of whether it actually includes guards—if no guards are written, then `True` is substituted for the guards $gs_{i,j}$ in the $match_i$ forms. Subsequent identities manipulate the resulting `case` expression into simpler and simpler forms.

Rule (h) in Listing 3 involves the overloaded operator `==`; it is this rule that defines the meaning of pattern matching against overloaded constants.

These identities all preserve the static semantics. Rules (d), (e), (j), and (q) use a lambda rather than a `let`; this indicates that variables bound by `case` are monomorphically typed (Section 4.1.4).

(a) `case` e `of` { $alts$ } $= (\backslash v$ -> `case` v `of` { $alts$ }) e
where v is a new variable

(b) `case` v `of` { p_1 $match_1$; ... ; p_n $match_n$ }
$=$ `case` v `of` { p_1 $match_1$;
_ -> ... `case` v `of` {
p_n $match_n$;
_ -> `error "No match"`}...}
where each $match_i$ has the form

| $gs_{i,1}$ -> $e_{i,1}$; ... ; gs_{i,m_i} -> e_{i,m_i} `where` {$decls$}

(c) `case` v `of` { p | gs_1 -> e_1 ; ...
| gs_n -> e_n `where` {$decls$}
_ -> e' }
$=$ `case` e' `of` { y ->
`case` v `of` {
p -> `let` {$decls$} `in`
`case` () `of` {
() | gs_1 -> e_1 ;
_ -> ... `case` () `of` {
() | gs_n -> e_n ;
_ -> y } ... }
_ -> y }}
where y is a new variable

(d) `case` v `of` {~ p -> e ; _ -> e'}
$= (\backslash x_1 ... x_n$ -> $e)$(`case` v `of`{p -> x_1})...(`case` v `of` {p -> x_n})
where $x_1, ..., x_n$ are all the variables in p

(e) `case` v `of` { x@p -> e ; _ -> e'}
$=$ `case` v `of` {p -> $(\backslash x$ -> $e)$ v ; _ -> e'}

(f) `case` v `of` { _ -> e ; _ -> e' } $= e$

Figure 1: Semantics of Case Expressions, Part 1

(g) `case` v `of` { $Kp_1...p_n$ -> e _ -> e' }
$=$ `case` v `of` {
$K\ x_1...x_n$ -> `case` x_1 `of` {
p_1 -> ... `case` x_n `of` {p_n -> e ; _ -> e'}...
_ -> e'}
_ -> e'}
at least one of $p_1, ..., p_n$ is not a variable; $x_1, ..., x_n$ are new variables

(h) `case` v `of` { k -> e ; _ -> e' } $=$ `if` $(v == k)$ `then` e `else` e'
where k is a numeric, character, or string literal

(i) `case` v `of` { x -> e ; _ -> e' } $=$ `case` v `of` { x -> e }

(j) `case` v `of` { x -> e } $= (\backslash\ x$ -> $e)\ v$

(k) `case` Nv `of` { Np -> e ; _ -> e' }
$=$ `case` v `of` { p -> e ; _ -> e' }
where N is a `newtype` constructor

(l) `case` $\perp$ `of` { $N\ p$ -> e ; _ -> e' } $=$ `case` $\perp$ `of` { p -> e }
where N is a newtype constructor

(m) `case` v `of` { $K\{f_1 = p_1, f_2 = p_2, ...\}$ -> e ; _ -> e' }
$=$ `case` e' `of` {
y ->
`case` v `of` {
$K\{f_1 = p_1\}$ ->
`case` v `of` {$K\{f_2 = p_2, ...\}$ -> e ; _ -> y} ;
_ -> y}}
where $f_1, f_2, ...$ are fields of constructor K; y is a new variable

(n) `case` v `of` { $K\{f = p\}$ -> e ; _ -> e' }
$=$ `case` v `of` {
$Kp_1...p_n$ -> e ; _ -> e'}
where p_i is p if f labels the ith component of K, _ otherwise

(o) `case` v `of` { K { } -> e ; _ -> e' }
$=$ `case` v `of` {
K _ ... _ -> e ; _ -> e'}

(p) `case` $(K'\ e_1...e_m)$ `of` { $K\ x_1...x_n$ -> e ; _ -> e' } $= e'$
where K and K' are distinct `data` constructors of arity n and m, respectively

(q) `case` $(Ke_1...e_n)$ `of` { $Kx_1...x_n$ -> e ; _ -> e' }
$= (\backslash x_1...x_n$ -> $e)e_1...e_n$
where K is a `data` constructor of arity n

(r) `case` $\perp$ `of` { $Kx_1...x_n$ -> e ; _ -> e' } $= \perp$
where K is a `data` constructor of arity n

Listing 3: Semantics of Case Expressions, Part 2

(s) `case` () `of` { () | g_1, ... ,g_n -> e ; _ -> e' }
= `case` () `of` {
() | g_1 -> ... `case` () `of` {
() | g_n -> e;
_ -> e' } ...
_ -> e' }
where y is a new variable

(t) `case` () `of` { () | p <- e_0 -> e ; _ -> e' }
= `case` e_0 `of` { p -> e ; _ -> e' }

(u) `case` () `of` { () | `let` $decls$ -> e ; _ -> e' }
= `let` $decls$ `in` e

(v) `case` () `of` { () | e_0 -> e ; _ -> e' }
= `if` e_0 `then` e `else` e'

Figure 2: Semantics of Case Expressions, Part 3

Chapter 4

Declarations and Bindings

In this chapter, we describe the syntax and informal semantics of Haskell *declarations*.

module	→	`module` *modid* [*exports*] `where` *body*	
	|	*body*	
body	→	{ *impdecls* ; *topdecls* }	
	|	{ *impdecls* }	
	|	{ *topdecls* }	
topdecls	→	$topdecl_1$; ... ; $topdecl_n$	$(n \geq 1)$
topdecl	→	`type` *simpletype* = *type*	
	|	`data` [*context* =>] *simpletype* [= *constrs*][*deriving*]	
	|	`newtype` [*context* =>] *simpletype* = *newconstr* [*deriving*]	
	|	`class` [*scontext* =>] *tycls tyvar* [`where` *cdecls*]	
	|	`instance` [*scontext* =>] *qtycls inst* [`where` *idecls*]	
	|	`default` ($type_1$, ... , $type_n$)	$(n \geq 0)$
	|	`foreign` *fdecl*	
	|	*decl*	
decls	→	{ $decl_1$; ... ; $decl_n$ }	$(n \geq 0)$
decl	→	*gendecl*	
	|	(*funlhs* | *pat*) *rhs*	
cdecls	→	{ $cdecl_1$; ... ; $cdecl_n$ }	$(n \geq 0)$
cdecl	→	*gendecl*	
	|	(*funlhs* | *var*) *rhs*	
idecls	→	{ $idecl_1$; ... ; $idecl_n$ }	$(n \geq 0)$
idecl	→	(*funlhs* | *var*) *rhs*	
	|		(empty)
gendecl	→	*vars* :: [*context* =>] *type*	(type signature)
	|	*fixity* [*integer*] *ops*	(fixity declaration)
	|		(empty declaration)
ops	→	op_1 , ... , op_n	$(n \geq 1)$
vars	→	var_1 , ... , var_n	$(n \geq 1)$
fixity	→	`infixl` | `infixr` | `infix`	

The declarations in the syntactic category *topdecls* are only allowed at the top level of a Haskell module (see Chapter 5), whereas *decls* may be used either at the top level or in nested scopes (i.e. those within a `let` or `where` construct).

For exposition, we divide the declarations into three groups: user-defined datatypes, consisting of `type`, `newtype`, and `data` declarations (Section 4.2); type classes and overloading, consisting of `class`, `instance`, and `default` declarations (Section 4.3); and nested declarations, consisting of value bindings, type signatures, and fixity declarations (Section 4.4).

Haskell has several primitive datatypes that are "hard-wired" (such as integers and floating-point numbers), but most "built-in" datatypes are defined with normal Haskell code, using normal `type` and `data` declarations. These "built-in" datatypes are described in detail in Section 6.1.

4.1 Overview of Types and Classes

Haskell uses a traditional Hindley-Milner polymorphic type system to provide a static type semantics [4], [5], but the type system has been extended with *type classes* (or just *classes*) that provide a structured way to introduce *overloaded* functions.

A `class` declaration (Section 4.3.1) introduces a new *type class* and the overloaded operations that must be supported by any type that is an instance of that class. An `instance` declaration (Section 4.3.2) declares that a type is an *instance* of a class and includes the definitions of the overloaded operations —called *class methods*—instantiated on the named type.

For example, suppose we wish to overload the operations `(+)` and `negate` on types `Int` and `Float`. We introduce a new type class called `Num`:

```
class Num a  where          -- simplified class declaration for Num
  (+)    :: a -> a -> a     -- (Num is defined in the Prelude)
  negate :: a -> a
```

This declaration may be read "a type `a` is an instance of the class `Num` if there are class methods `(+)` and `negate`, of the given types, defined on it."“

We may then declare `Int` and `Float` to be instances of this class:

```
instance Num Int  where     -- simplified instance of Num Int
  x + y       =  addInt x y
  negate x    =  negateInt x

instance Num Float  where   -- simplified instance of Num Float
  x + y       =  addFloat x y
  negate x    =  negateFloat x
```

where `addInt`, `negateInt`, `addFloat`, and `negateFloat` are assumed in this case to be primitive functions, but in general could be any user-defined function. The first declaration above may be read "`Int` is an instance of the class `Num` as witnessed by these definitions (i.e. class methods) for `(+)` and `negate`."

More examples of type classes can be found in the papers by Jones [6] or Wadler and Blott [7]. The term "type class" was used to describe the original Haskell 1.0 type system; "constructor class" was used to describe an extension to the original type classes. There is no longer any reason to use two different terms: in this report, "type class" includes both the original Haskell type classes and the constructor classes introduced by Jones.

4.1.1 Kinds

To ensure that they are valid, type expressions are classified into different *kinds*, which take one of two possible forms:

- The symbol $*$ represents the kind of all nullary type constructors.
- If κ_1 and κ_2 are kinds, then $\kappa_1 \to \kappa_2$ is the kind of types that take a type of kind κ_1 and return a type of kind κ_2.

Kind inference checks the validity of type expressions in a similar way that type inference checks the validity of value expressions. However, unlike types, kinds are entirely implicit and are not a visible part of the language. Kind inference is discussed in Section 4.6.

4.1.2 Syntax of Types

type	→	*btype* [-> *type*]	(function type)
btype	→	[*btype*] *atype*	(type application)
atype	→	*gtycon*	
	|	*tyvar*	
	|	($type_1$, ... , $type_k$)	(tuple type, $k \geq 2$)
	|	[*type*]	(list type)
	|	(*type*)	(parenthesized constructor)
gtycon	→	*qtycon*	
	|	()	(unit type)
	|	[]	(list constructor)
	|	(->)	(function constructor)
	|	(, {,})	(tupling constructors)

The syntax for Haskell type expressions is given above. Just as data values are built using data constructors, type values are built from *type constructors*. As with data constructors, the names of type constructors start with uppercase letters. Unlike data constructors, infix type constructors are not allowed (other than `(->)`).

The main forms of type expression are as follows:

1. Type variables, written as identifiers beginning with a lowercase letter. The kind of a variable is determined implicitly by the context in which it appears.

2. Type constructors. Most type constructors are written as an identifier beginning with an uppercase letter. For example:
   - `Char`, `Int`, `Integer`, `Float`, `Double` and `Bool` are type constants with kind $*$.
   - `Maybe` and `IO` are unary type constructors, and treated as types with kind $* \to *$.
   - The declarations `data T ...` or `newtype T ...` add the type constructor `T` to the type vocabulary. The kind of `T` is determined by kind inference.

   Special syntax is provided for certain built-in type constructors:
   - The *trivial type* is written as `()` and has kind $*$. It denotes the "nullary tuple" type, and has exactly one value, also written `()` (see Section 3.9 and Section 6.1.5).
   - The *function type* is written as `(->)` and has kind $* \to * \to *$.
   - The *list type* is written as `[]` and has kind $* \to *$.
   - The *tuple types* are written as `(,)`, `(,,)`, and so on. Their kinds are $* \to * \to *$,$* \to * \to * \to *$, and so on.

Use of the `(->)` and `[]` constants is described in more detail below.

3. Type application. If t_1 is a type of kind $\kappa_1 \to \kappa_2$ and t_2 is a type of kind κ_1, then $t_1\ t_2$ is a type expression of kind κ_2.

4. A *parenthesized type*, having form (t), is identical to the type t.

For example, the type expression `IO a` can be understood as the application of a constant, `IO`, to the variable `a`. Since the `IO` type constructor has kind $* \to *$, it follows that both the variable `a` and the whole expression, `IO a`, must have kind $*$. In general, a process of *kind inference* (see Section 4.6) is needed to determine appropriate kinds for user-defined datatypes, type synonyms, and classes.

Special syntax is provided to allow certain type expressions to be written in a more traditional style:

1. A *function type* has the form $t_1 \to t_2$, which is equivalent to the type $(\to)t_1t_2$. Function arrows associate to the right. For example, `Int -> Int -> Float` means `Int -> (Int -> Float)`.
2. A *tuple type* has the form $(t_1, ..., t_k)$, where $k \geq 2$, which is equivalent to the type $(, ...,)t_1...t_k$ where there are $k-1$ commas between the parenthesis. It denotes the type of k-tuples with the first component of type t_1, the second component of type t_2, and so on (see Section 3.8 and Section 6.1.4).
3. A *list type* has the form $[t]$, which is equivalent to the type $[]t$. It denotes the type of lists with elements of type t (see Section 3.7 and Section 6.1.3).

These special syntactic forms always denote the built-in type constructors for functions, tuples, and lists, regardless of what is in scope. In a similar way, the prefix type constructors `(->)`, `[]`, `()`, `(,)`, and so on, always denote the built-in type constructors; they cannot be qualified, nor mentioned in import or export lists (Chapter 5). (Hence the special production, "gtycon", above.)

Although the list and tuple types have special syntax, their semantics is the same as the equivalent user-defined algebraic data types.

Notice that expressions and types have a consistent syntax. If t_i is the type of expression or pattern e_i, then the expressions `(\ e1 -> e2)`, `[e1]`, and `(t1 -> t2)`, `[t1]`, and `(t1, t2)`, respectively.

With one exception (that of the distinguished type variable in a class declaration (Section 4.3.1)), the type variables in a Haskell type expression are all assumed to be universally quantified; there is no explicit syntax for universal quantification [5]. For example, the type expression `a -> a` denotes the type $\forall a.a \to a$. For clarity, however, we often write quantification explicitly when discussing the types of Haskell programs. When we write an explicitly quantified type, the scope of the $\forall$ extends as far to the right as possible; for example, $\forall a.a \to a$ means $\forall a.(a \to a)$.

4.1.3 Syntax of Class Assertions and Contexts

context	→	*class*	
	|	($class_1$, ... , $class_n$)	$(n \geq 0)$
class	→	*qtycls tyvar*	
	|	*qtycls* (*tyvar* $atype_1$...$atype_n$)	$(n \geq 1)$
qtycls	→	[*modid* .] *tycls*	
tycls	→	*conid*	
tyvar	→	*varid*	

A *class assertion* has form *qtycls tyvar*, and indicates the membership of the type *tyvar* in the class *qtycls*. A class identifier begins with an uppercase letter. A *context* consists of zero or more class assertions, and has the general form

$$(C_1 u_1, ..., C_n u_n)$$

where $C_1, ..., C_n$ are class identifiers, and each of the $u_1, ..., u_n$ is either a type variable, or the application of type variable to one or more types. The outer parentheses may be omitted when $n = 1$. In general, we use cx to denote a context and we write $cx \Rightarrow t$ to indicate the type t restricted by the context cx. The context cx must only contain type variables referenced in t. For convenience, we write $cx \Rightarrow t$ even if the context cx is empty, although in this case the concrete syntax contains no `=>`.

4.1.4 Semantics of Types and Classes

In this section, we provide informal details of the type system. (Wadler and Blott [7] and Jones [6] discuss type and constructor classes, respectively, in more detail.)

The Haskell type system attributes a *type* to each expression in the program. In general, a type is of the form $\forall \overline{u}.\ cx \Rightarrow t$, where $\overline{u}$ is a set of type variables $u_1, ..., u_n$. In any such type, any of the universally-quantified type variables u_i that are free in cx must also be free in t. Furthermore, the context cx must be of the form given above in Section 4.1.3. For example, here are some valid types:

```
Eq a => a -> a
(Eq a, Show a, Eq b) => [a] -> [b] -> String
(Eq (f a), Functor f) => (a -> b) -> f a -> f b -> Bool
```

In the third type, the constraint `Eq (f a)` cannot be made simpler because `f` is universally quantified.

The type of an expression e depends on a *type environment* that gives types for the free variables in e, and a *class environment* that declares which types are instances of which classes (a type becomes an instance of a class only via the presence of an `instance` declaration or a `deriving` clause).

Types are related by a generalization preorder (specified below); the most general type, up to the equivalence induced by the generalization preorder, that can be assigned to a particular expression (in a given environment) is called its *principal type*. Haskell's extended Hindley-Milner type system can infer the principal type of all expressions, including the proper use of overloaded class methods (although certain ambiguous overloadings could arise, as described in Section 4.3.4). Therefore, explicit typings (called *type signatures*) are usually optional (see Section 3.16 and Section 4.4.1).

The type $\forall \overline{u}.cx_1 \Rightarrow t_1$ *more general than* the type $\forall \overline{w}.cx_2 \Rightarrow t_2$ if and only if there is a substitution S whose domain is $\overline{u}$ such that:

- t_2 is identical to $S(t_1)$.
- Whenever cx_2 holds in the class environment, $S(cx_1)$ also holds.

A value of type $\forall \overline{u}.cx \Rightarrow t$, may be instantiated at types $\overline{s}$ if and only if the context $cx\left[\frac{\overline{s}}{\overline{u}}\right]$ holds. For example, consider the function `double`:

```
double x = x + x
```

The most general type of `double` is $\forall a.$ `Num` $a \Rightarrow a \rightarrow a$. `double` may be applied to values of type `Int` (instantiating a to `Int`), since `Num Int` holds, because `Int` is an instance of the class `Num`. However, `double` may not normally be applied to values of type `Char`, because `Char` is not normally an instance of class `Num`. The user may choose to declare such an instance, in which case `double` may indeed be applied to a `Char`.

4.2 User-Defined Datatypes

In this section, we describe algebraic datatypes (`data` declarations), renamed datatypes (`newtype` declarations), and type synonyms (`type` declarations). These declarations may only appear at the top level of a module.

4.2.1 Algebraic Datatype Declarations

topdecl	→	`data` [*context* `=>`] *simpletype* [`=` *constrs*][*deriving*]	
simpletype	→	$tycon\ tyvar_1 ... tyvar_k$	$(k \geq 0)$
constrs	→	$constr_1 \mid ... \mid constr_n$	$(n \geq 1)$
constr	→	$con\ [!]atype_1 ... [!]atype_k$	(arity $con = k, k \geq 0$)
	|	(*btype* | `!` *atype*) *conop* (*btype* | `!` *atype*)	(infix *conop*)
	|	$con\ \{\ fielddecl_1\ ,\ ...\ ,\ fielddecl_n\ \}$	$(n \geq 0)$
fielddecl	→	*vars* `::` (*type* | `!` *atype*)	
deriving	→	`deriving` (*dclass* | `(` $dclass_1\ ,\ ...\ ,\ dclass_n$ `)`)	$(n \geq 0)$
dclass	→	*qtycls*	

The precedence for *constr* is the same as that for expressions—normal constructor application has higher precedence than infix constructor application (thus `a : Foo a` parses as `a : (Foo a)`).

An algebraic datatype declaration has the form:

$$\texttt{data}\ cx \Rightarrow Tu_1 ... u_k = K_1 t_{11} ... t_{1k_1} \mid ... \mid K_n t_{n1} ... t_{nk_n}$$

where cx is a context. This declaration introduces a new *type constructor* T with zero or more constituent *data constructors* $K_1, ...K_n$. In this Report, the unqualified term "constructor" always means "data constructor".

The types of the data constructors are given by:

$$K_i :: \forall u_1 ... u_k . cx_i \Rightarrow t_{i1} \to ... \to t_{ik_i} \to (T\ u_1 ... u_k)$$

where cx_i is the largest subset of cx that constrains only those type variables free in the types $t_{i1} ... t_{ik_i}$. The type variables u_1 through u_k must be distinct and may appear in cx and the t_{ij}; it is a static error for any other type variable to appear in cx or on the right-hand-side. The new type constant T has a kind of the form $\kappa_1 \to ... \to \kappa_k \to *$ where the kinds κ_i of the argument variables u_i are determined by kind inference as described in Section 4.6. This means that T may be used in type expressions with anywhere between 0 and k arguments.

For example, the declaration

```
data Eq a => Set a = NilSet | ConsSet a (Set a)
```

introduces a type constructor `Set` of kind $* \to *$, and constructors `NilSet` and `ConsSet` with types

```
NilSet  :: ∀a. Set a
ConsSet :: ∀a. Eq a ⇒ a → Set a → Set a
```

In the example given, the overloaded type for `ConsSet` ensures that `ConsSet` can only be applied to values whose type is an instance of the class `Eq`. Pattern matching against `ConsSet` also gives rise to an `Eq a` constraint. For example:

```
f (ConsSet a s) = a
```

the function `f` has inferred type `Eq a => Set a -> a`. The context in the `data` declaration has no other effect whatsoever.

The visibility of a datatype's constructors (i.e. the "abstractness" of the datatype) outside of the module in which the datatype is defined is controlled by the form of the datatype's name in the export list as described in Section 5.8.

The optional `deriving` part of a `data` declaration has to do with *derived instances*, and is described in Section 4.3.3.

Labelled Fields A data constructor of arity k creates an object with k components. These components are normally accessed positionally as arguments to the constructor in expressions or patterns. For large datatypes it is useful to assign *field labels* to the components of a data object. This allows a specific field to be referenced independently of its location within the constructor.

A constructor definition in a `data` declaration may assign labels to the fields of the constructor, using the record syntax (C { ... }). Constructors using field labels may be freely mixed with constructors without them. A constructor with associated field labels may still be used as an ordinary constructor; features using labels are simply a shorthand for operations using an underlying positional constructor. The arguments to the positional constructor occur in the same order as the labeled fields. For example, the declaration

```
data C = F { f1,f2 :: Int, f3 :: Bool }
```

defines a type and constructor identical to the one produced by

```
data C = F Int Int Bool
```

Operations using field labels are described in Section 3.15. A `data` declaration may use the same field label in multiple constructors as long as the typing of the field is the same in all cases after type synonym expansion. A label cannot be shared by more than one type in scope. Field names share the top level namespace with ordinary variables and class methods and must not conflict with other top level names in scope.

The pattern `F { }` matches any value built with constructor `F`, *whether or not `F` was declared with record syntax.*

Strictness Flags Whenever a data constructor is applied, each argument to the constructor is evaluated if and only if the corresponding type in the algebraic datatype declaration has a strictness flag, denoted by an exclamation point, "`!`". Lexically, "`!`" is an ordinary varsym not a *reservedop*; it has special significance only in the context of the argument types of a data declaration.

Translation: A declaration of the form

$$\texttt{data}\ cx \Rightarrow Tu_1...u_k = ... \mid Ks_1...s_n \mid ...$$

where each s_i is either of the form $!t_i$ or t_i, replaces every occurrence of K in an expression by

$$(\backslash\ x_1...x_n \texttt{ -> } (((Kop_1x_1)op_2x_2)...)op_nx_n)$$

where op_i is the non-strict apply function `$` if s_i is of the form t_i, and op_i is the strict apply function `$!` (see Section 6.2) if s_i is of the form $!t_i$. Pattern matching on K is not affected by strictness flags.

4.2.2 Type Synonym Declarations

topdecl → `type` *simpletype* = *type*
simpletype → *tycon tyvar*$_1$...*tyvar*$_k$ $(k \geq 0)$

A type synonym declaration introduces a new type that is equivalent to an old type. It has the form

$$\texttt{type}\ Tu_1...u_k = t$$

which introduces a new type constructor, T. The type $(Tu_1...u_k)$ is equivalent to the type $t[t_1\ /\ u_1, ..., t_k\ /\ u_k]$. The type variables u_1 through u_k must be distinct and are scoped only over t; it is a static error for any other type variable to appear in t. The kind of the new type constructor T is of the form $\kappa_1 \to ... \to \kappa_k \to \kappa$ where the kinds κ_i of the arguments u_i and κ of the right hand side t are determined by kind inference as described in Section 4.6. For example, the following definition can be used to provide an alternative way of writing the list type constructor:

```
type List = []
```

Type constructor symbols T introduced by type synonym declarations cannot be partially applied; it is a static error to use T without the full number of arguments.

Although recursive and mutually recursive datatypes are allowed, this is not so for type synonyms, *unless an algebraic datatype intervenes*. For example,

```
type Rec a   =  [Circ a]
data Circ a  =  Tag [Rec a]
```

is allowed, whereas

```
type Rec a   =  [Circ a]        -- invalid
type Circ a  =  [Rec a]         -- invalid
```

is not. Similarly, `type Rec a = [Rec a]` is not allowed.

Type synonyms are a convenient, but strictly syntactic, mechanism to make type signatures more readable. A synonym and its definition are completely interchangeable, except in the instance type of an `instance` declaration (Section 4.3.2).

4.2.3 Datatype Renamings

topdecl → `newtype` [*context* =>] *simpletype* = *newconstr* [*deriving*]
newconstr → *con atype*
| *con* { *var* :: *type* }
simpletype → *tycon tyvar*$_1$...*tyvar*$_k$ $(k \geq 0)$

A declaration of the form

$$\texttt{newtype}\ cx \Rightarrow Tu_1 ... u_k = N\ t$$

introduces a new type whose representation is the same as an existing type. The type $(Tu_1 ... u_k)$ renames the datatype t. It differs from a type synonym in that it creates a distinct type that must be explicitly coerced to or from the original type. Also, unlike type synonyms, `newtype` may be used to define recursive types. The constructor N in an expression coerces a value from type `t` to type $(Tu_1 ... u_k)$. Using N in a pattern coerces a value from type $(Tu_1 ... u_k)$ to type t. These coercions may be implemented without execution time overhead; `newtype` does not change the underlying representation of an object.

New instances (see Section 4.3.2) can be defined for a type defined by `newtype` but may not be defined for a type synonym. A type created by `newtype` differs from an algebraic datatype in that the representation of an algebraic datatype has an extra level of indirection. This difference may make access to the representation less efficient. The difference is reflected in different rules for pattern matching (see Section 3.17). Unlike algebraic datatypes, the newtype constructor N is *unlifted*, so that $N\ \bot$ is the same as $\bot$.

The following examples clarify the differences between `data` (algebraic datatypes), `type` (type synonyms), and `newtype` (renaming types.) Given the declarations

```
data D1 = D1 Int
data D2 = D2 !Int
type S = Int
newtype N = N Int
d1 (D1 i) = 42
d2 (D2 i) = 42
s i = 42
n (N i) = 42
```

the expressions (`d1` $\bot$), (`d2` $\bot$) and (`d2` (`D2` $\bot$)) are all equivalent to $\bot$, whereas (`n` $\bot$), (`n` (`N` $\bot$)), (`d1` (`D1` $\bot$)) and (`s` $\bot$) are all equivalent to `42`. In particular, (`N` $\bot$) is equivalent to $\bot$ while (`D1` $\bot$) is not equivalent to $\bot$.

The optional deriving part of a `deriving` declaration is treated in the same way as the deriving component of a `data` declaration; see Section 4.3.3.

A `newtype` declaration may use field-naming syntax, though of course there may only be one field. Thus:

```
newtype Age = Age { unAge :: Int }
```

brings into scope both a constructor and a de-constructor:

```
Age   :: Int -> Age
unAge :: Age -> Int
```

4.3 Type Classes and Overloading

4.3.1 Class Declarations

topdecl → `class` [*scontext* `=>`] *tycls tyvar* [`where` *cdecls*]

scontext	→	*simpleclass*	
	|	($simpleclass_1$, ... , $simpleclass_n$)	$(n \geq 0)$
simpleclass	→	*qtycls tyvar*	
cdecls	→	{ $cdecl_1$; ... ; $cdecl_n$ }	$(n \geq 0)$
cdecl	→	*gendecl*	
	|	(*funlhs* | *var*) *rhs*	

A *class declaration* introduces a new class and the operations (*class methods*) on it. A class declaration has the general form:

$$\texttt{class}\ cx \Rightarrow Cu\ \texttt{where}\ cdecls$$

This introduces a new class name C; the type variable u is scoped only over the class method signatures in the class body. The context cx specifies the superclasses of C, as described below; the only type variable that may be referred to in cx is u.

The superclass relation must not be cyclic; i.e. it must form a directed acyclic graph.

The *cdecls* part of a `class` declaration contains three kinds of declarations:

- The class declaration introduces new *class methods* v_i, whose scope extends outside the `class` declaration. The class methods of a class declaration are precisely the mbox{itv_i} for which there is an explicit type signature

  $$v_i :: cx_i \Rightarrow t_i$$

  in *cdecls*. Class methods share the top level namespace with variable bindings and field names; they must not conflict with other top level bindings in scope. That is, a class method can not have the same name as a top level definition, a field name, or another class method.

  The type of the top-level class method v_i is:

  $$v_i :: \forall u, \overline{w}.(Cu, cx_i) \Rightarrow t_i$$

  The t_i must mention u; it may mention type variables $\overline{w}$ other than u, in which case the type of v_i is polymorphic in both u and $\overline{w}$. The cx_i may constrain only $\overline{w}$; in particular, the cx_i may not constrain u. For example:

  ```
  class Foo a where
    op :: Num b => a -> b -> a
  ```

  Here the type of `op` is $\forall a, b.(\texttt{Foo}\ a, \texttt{Num}\ b) \Rightarrow a \to b \to a$.
- The *cdecls* may also contain a *fixity declaration* for any of the class methods (but for no other values). However, since class methods declare top-level values, the fixity declaration for a class method may alternatively appear at top level, outside the class declaration.
- Lastly, the *cdecls* may contain a *default class method* for any of the v_i. The default class method for v_i is used if no binding for it is given in a particular `instance` declaration (see Section 4.3.2). The default method declaration is a normal value definition, except that the left hand side may only be a variable or function definition. For example:

  ```
  class Foo a where
    op1, op2 :: a -> a
    (op1, op2) = ...
  ```

  is not permitted, because the left hand side of the default declaration is a pattern.

Other than these cases, no other declarations are permitted in *cdecls*.

A `class` declaration with no `where` part may be useful for combining a collection of classes into a larger one that inherits all of the class methods in the original ones. For example:

```
class  (Read a, Show a) => Textual a
```

In such a case, if a type is an instance of all superclasses, it is not *automatically* an instance of the subclass, even though the subclass has no immediate class methods. The `instance` declaration must be given explicitly with no `where` part.

4.3.2 Instance Declarations

topdecl	→	`instance` [*scontext* `=>`] *qtycls inst* [`where` *idecls*]	
inst	→	*gtycon*	
	|	(*gtycon* $tyvar_1 ... tyvar_k$)	($k \geq 0$, *tyvars* distinct)
	|	($tyvar_1$, ... , $tyvar_k$)	($k \geq 2$, *tyvars* distinct)
	|	[*tyvar*]	
	|	($tyvar_1$ -> $tyvar_2$)	($tyvar_1$ and $tyvar_2$ distinct)
idecls	→	{ $idecl_1$; ... ; $idecl_n$ }	($n \geq 0$)
idecl	→	(*funlhs* | *var*) *rhs*	
	|		(empty)

An *instance declaration* introduces an instance of a class. Let

$$\texttt{class}\ cx\ \texttt{=>}\ C\ u\ \texttt{where}\ \{cbody\}$$

be a `class` declaration. The general form of the corresponding instance declaration is:

$$\texttt{instance}\ cx'\ \texttt{=>}\ C\ (T\ u_1 ... u_k)\ \texttt{where}\ \{d\}$$

where $k \geq 0$. The type $(Tu_1...u_k)$ must take the form of a type constructor T applied to simple type variables $u_1, ..., u_k$; furthermore, T must not be a type synonym, and the u_i must all be distinct.

This prohibits instance declarations such as:

```
instance C (a,a) where ...
instance C (Int,a) where ...
instance C [[a]] where ...
```

The declarations d may contain bindings only for the class methods of C. It is illegal to give a binding for a class method that is not in scope, but the name under which it is in scope is immaterial; in particular, it may be a qualified name. (This rule is identical to that used for subordinate names in export lists — Section 5.2.) For example, this is legal, even though `range` is in scope only with the qualified name `Data.Ix.range`.

```
module A where
  import qualified Data.Ix

  instance Data.Ix.Ix T where
    range = ...
```

The declarations may not contain any type signatures or fixity declarations, since these have already been given in the `class` declaration. As in the case of default class methods (Section 4.3.1), the method declarations must take the form of a variable or function definition.

If no binding is given for some class method then the corresponding default class method in the `class` declaration is used (if present); if such a default does not exist then the class method of this instance is bound to `undefined` and no compile-time error results.

An `instance` declaration that makes the type T to be an instance of class C is called a *C-T instance declaration* and is subject to these static restrictions:

- A type may not be declared as an instance of a particular class more than once in the program.
- The class and type must have the same kind; this can be determined using kind inference as described in Section 4.6.
- Assume that the type variables in the instance type $(Tu_1 ... u_k)$ satisfy the constraints in the instance context cx'. Under this assumption, the following two conditions must also be satisfied:
  1. The constraints expressed by the superclass context $cx[(Tu_1 ... u_k)\ /u]$ of C must be satisfied. In other words, T must be an instance of each of C's superclasses and the contexts of all superclass instances must be implied by cx'.
  2. Any constraints on the type variables in the instance type that are required for the class method declarations in d to be well-typed must also be satisfied.

  In fact, except in pathological cases it is possible to infer from the instance declaration the most general instance context cx' satisfying the above two constraints, but it is nevertheless mandatory to write an explicit instance context.

The following example illustrates the restrictions imposed by superclass instances:

```
class Foo a => Bar a where ...

instance (Eq a, Show a) => Foo [a] where ...

instance Num a => Bar [a] where ...
```

This example is valid Haskell. Since `Foo` is a superclass of `Bar`, the second instance declaration is only valid if `[a]` is an instance of `Foo` under the assumption `Num a`. The first instance declaration does indeed say that `[a]` is an instance of `Foo` under this assumption, because `Eq` and `Show` are superclasses of `Num`.

If the two instance declarations instead read like this:

```
instance Num a => Foo [a] where ...

instance (Eq a, Show a) => Bar [a] where ...
```

then the program would be invalid. The second instance declaration is valid only if `[a]` is an instance of `Foo` under the assumptions `(Eq a, Show a)`. But this does not hold, since `[a]` is only an instance of `Foo` under the stronger assumption `Num a`.

Further examples of `instance` declarations may be found in Chapter 9.

4.3.3 Derived Instances

As mentioned in Section 4.2.1, `data` and `newtype` declarations contain an optional `deriving` form. If the form is included, then *derived instance declarations* are automatically generated for the datatype in each of the named classes. These instances are subject to the same restrictions as user-defined instances. When deriving a class C for a type T, instances for all superclasses of C must exist for T, either via an explicit `instance` declaration or by including the superclass in the `deriving` clause.

Derived instances provide convenient commonly-used operations for user-defined datatypes. For example, derived instances for datatypes in the class `Eq` define the operations `==` and `/=`, freeing the programmer from the need to define them.

The only classes in the Prelude for which derived instances are allowed are `Eq`, `Ord`, `Enum`, `Bounded`, `Show`, and `Read`, all mentioned in Figure 3. The precise details of how the derived instances are generated for each of these classes are provided in Chapter 11, including a specification of when such derived instances are possible. Classes defined by the standard libraries may also be derivable.

A static error results if it is not possible to derive an `instance` declaration over a class named in a `deriving` form. For example, not all datatypes can properly support class methods in `Enum`. It is also a static error to give an explicit `instance` declaration for a class that is also derived.

If the `deriving` form is omitted from a `data` or `newtype` declaration, then *no* instance declarations are derived for that datatype; that is, omitting a `deriving` form is equivalent to including an empty deriving form: `deriving ()`.

4.3.4 Ambiguous Types, and Defaults for Overloaded Numeric Operations

$$topdecl \rightarrow \texttt{default}\ (\ type_1\ ,\ \ldots\ ,\ type_n\ ) \quad (n \geq 0)$$

A problem inherent with Haskell-style overloading is the possibility of an *ambiguous type*. For example, using the `read` and `show` functions defined in Chapter 11, and supposing that just `Int` and `Bool` are members of `Read` and `Show`, then the expression

```
let x = read "..." in show x  -- invalid
```

is ambiguous, because the types for `show` and `read`,

$$\texttt{show} :: \forall a.\ \texttt{Show}\ a \Rightarrow a \rightarrow \texttt{String}$$
$$\texttt{read} :: \forall a.\ \texttt{Read}\ a \Rightarrow \texttt{String} \rightarrow a$$

could be satisfied by instantiating `a` as either `Int` in both cases, or `Bool`. Such expressions are considered ill-typed, a static error.

We say that an expression `e` has an *ambiguous type* if, in its type $\forall \overline{u}.cx \Rightarrow t$, there is a type variable u in $\overline{u}$ that occurs in cx but not in t. Such types are invalid.

For example, the earlier expression involving `show` and `read` has an ambiguous type since its type is $\forall a.\ \texttt{Show}\ a, \texttt{Read}\ a \Rightarrow \texttt{String}$.

Ambiguous types can only be circumvented by input from the user. One way is through the use of *expression type-signatures* as described in Section 3.16. For example, for the ambiguous expression given earlier, one could write:

```
let x = read "..." in show (x::Bool)
```

which disambiguates the type.

Occasionally, an otherwise ambiguous expression needs to be made the same type as some variable, rather than being given a fixed type with an expression type-signature. This is the purpose of the function `asTypeOf` (Chapter 9): `x `asTypeOf` y` has the value of x, but x and y are forced to have the same type. For example,

```
approxSqrt x = encodeFloat 1 (exponent x `div` 2) `asTypeOf` x
```

(See Section 6.4.6 for a description of `encodeFloat` and `exponent`.)

Ambiguities in the class `Num` are most common, so Haskell provides another way to resolve them—with a *default declaration*:

$$\texttt{default}\ (t_1, ..., t_n)$$

where $n \geq 0$, and each t_i must be a type for which `Num` t_i holds. In situations where an ambiguous type is discovered, an ambiguous type variable, v, is defaultable if:

- v appears only in constraints of the form $C\ v$, where C is a class, and
- at least one of these classes is a numeric class, (that is, `Num` or a subclass of `Num`), and
- all of these classes are defined in the Prelude or a standard library (Listing 4 – Listing 5 show the numeric classes, and Figure 3 shows the classes defined in the Prelude.)

Each defaultable variable is replaced by the first type in the default list that is an instance of all the ambiguous variable's classes. It is a static error if no such type is found.

Only one default declaration is permitted per module, and its effect is limited to that module. If no default declaration is given in a module then it assumed to be:

```
default (Integer, Double)
```

The empty default declaration, `default ()`, turns off all defaults in a module.

4.4 Nested Declarations

The following declarations may be used in any declaration list, including the top level of a module.

4.4.1 Type Signatures

$$\begin{array}{lcll} gendecl & \rightarrow & vars\ ::\ [context\ =>]\ type & \\ vars & \rightarrow & var_1\ ,\ ...\ ,\ var_n & (n \geq 1) \end{array}$$

A type signature specifies types for variables, possibly with respect to a context. A type signature has the form:

$$v_1, ..., v_n\ ::\ cx\ =>\ t$$

which is equivalent to asserting $v_i\ ::\ cx\ =>\ t$ for each i from 1 to n. Each v_i must have a value binding in the same declaration list that contains the type signature; i.e. it is invalid to give a type signature for a variable bound in an outer scope. Moreover, it is invalid to give more than one type signature for one variable, even if the signatures are identical.

As mentioned in Section 4.1.2, every type variable appearing in a signature is universally quantified over that signature, and hence the scope of a type variable is limited to the type signature that contains it. For example, in the following declarations

```
f :: a -> a
f x = x :: a                    -- invalid
```

the `a`'s in the two type signatures are quite distinct. Indeed, these declarations contain a static error, since `x` does not have type $\forall a.a$. (The type of `x` is dependent on the type of `f`; there is currently no way in Haskell to specify a signature for a variable with a dependent type; this is explained in Section 4.5.4.)

If a given program includes a signature for a variable f, then each use of f is treated as having the declared type. It is a static error if the same type cannot also be inferred for the defining occurrence of f.

If a variable f is defined without providing a corresponding type signature declaration, then each use of f outside its own declaration group (see Section 4.5.1) is treated as having the corresponding inferred, or *principal* type. However, to ensure that type inference is still possible, the defining occurrence, and all uses of f within its declaration group must have the same monomorphic type (from which the principal type is obtained by generalization, as described in Section 4.5.2).

For example, if we define

```
sqr x  =  x*x
```

then the principal type is `sqr` :: $\forall a.$ `Num` $a \Rightarrow a \to a$, which allows applications such as `sqr 5` or `sqr 0.1`. It is also valid to declare a more specific type, such as

```
sqr :: Int -> Int
```

but now applications such as `sqr 0.1` are invalid. Type signatures such as

```
sqr :: (Num a, Num b) => a -> b     -- invalid
sqr :: a -> a                       -- invalid
```

are invalid, as they are more general than the principal type of `sqr`.

Type signatures can also be used to support *polymorphic recursion*. The following definition is pathological, but illustrates how a type signature can be used to specify a type more general than the one that would be inferred:

```
data T a  =  K (T Int) (T a)
f         :: T a -> a
f (K x y) =  if f x == 1 then f y else undefined
```

If we remove the signature declaration, the type of `f` will be inferred as `T Int -> Int` due to the first recursive call for which the argument to `f` is `T Int`. Polymorphic recursion allows the user to supply the more general type signature, `T a -> a`.

4.4.2 Fixity Declarations

gendecl	→	*fixity* [*integer*] *ops*	
fixity	→	`infixl` | `infixr` | `infix`	
ops	→	op_1 , ... , op_n	$(n \geq 1)$
op	→	*varop* | *conop*	

A fixity declaration gives the fixity and binding precedence of one or more operators. The *integer* in a fixity declaration must be in the range 0 to 9. A fixity declaration may appear anywhere that a type signature appears and, like a type signature, declares a property of a particular operator. Also like a type signature, a fixity declaration can only occur in the same sequence of declarations as the declaration of the operator itself, and at most one fixity declaration may be given for any operator. (Class methods are a minor exception; their fixity declarations can occur either in the class declaration itself or at top level.)

There are three kinds of fixity, non-, left- and right-associativity (`infix`, `infixl`, and `infixr`, respectively), and ten precedence levels, 0 to 9 inclusive (level 0 binds least tightly, and level 9 binds most tightly). If the *digit* is omitted, level 9 is assumed. Any operator lacking a fixity declaration is assumed

to be `infixl 9` (See Chapter 3 for more on the use of fixities). Table 1 lists the fixities and precedences of the operators defined in the Prelude.

Precedence	Left associative operators	Non-associative operators	Right associative operators
9	`!!`		`.`
8			`^`,`^^`,`**`
7	`*`, `/`, `div`, `mod`, `rem`, `quot`		
6	`+`,`-`		
5			`:`,`++`
4		`==`, `/=`, `<`,`<=`, `>`, `>=`, `elem`, `notElem`	
3			`&&`
2			`||`
1	`>>`, `>>=`		
0			`$`, `$!`, `seq`

Table 1: Precedences and fixities of prelude operators

Fixity is a property of a particular entity (constructor or variable), just like its type; fixity is not a property of that entity's *name*. For example:

```
module Bar( op ) where
  infixr 7 `op`
  op = ...

module Foo where
  import qualified Bar
  infix 3 `op`

  a `op` b = (a `Bar.op` b) + 1

  f x = let
          p `op` q = (p `Foo.op` q) * 2
        in ...
```

Here, `` `Bar.op` `` is `infixr 7`, `` `Foo.op` `` is `infix 3`, and the nested definition of `op` in `f`'s right-hand side has the default fixity of `infixl 9`. (It would also be possible to give a fixity to the nested definition of `` `op` `` with a nested fixity declaration.)

4.4.3 Function and Pattern Bindings

$$
\begin{array}{lcll}
\mathit{decl} & \to & (\mathit{funlhs} \mid \mathit{pat})\ \mathit{rhs} & \\
\mathit{funlhs} & \to & \mathit{var}\ \mathit{apat}\ \{\mathit{apat}\} & \\
 & \mid & \mathit{pat}\ \mathit{varop}\ \mathit{pat} & \\
 & \mid & (\ \mathit{funlhs}\ )\ \mathit{apat}\ \{\mathit{apat}\} & \\
\mathit{rhs} & \to & = \mathit{exp}\ [\texttt{where}\ \mathit{decls}] & \\
 & \mid & \mathit{gdrhs}\ [\texttt{where}\ \mathit{decls}] & \\
\mathit{gdrhs} & \to & \mathit{guards} = \mathit{exp}\ [\mathit{gdrhs}] & \\
\mathit{guards} & \to & \mid \mathit{guard}_1\ ,\ \ldots\ ,\ \mathit{guard}_n & (n \geq 1) \\
\mathit{guard} & \to & \mathit{pat} <\text{-}\ \mathit{infixexp} & \text{(pattern guard)}
\end{array}
$$

| `let` *decls* (local declaration)
| *infixexp* (boolean guard)

We distinguish two cases within this syntax: a *pattern binding* occurs when the left hand side is a *pat*; otherwise, the binding is called a *function binding*. Either binding may appear at the top-level of a module or within a `where` or `let` construct.

4.4.3.1 Function bindings

A function binding binds a variable to a function value. The general form of a function binding for variable x is:

$$\begin{array}{lll} x & p_{11} \ \ldots \ p_{1k} & match_1 \\ \ldots & & \\ x & p_{n1} \ \ldots \ p_{nk} & match_n \end{array}$$

where each p_{ij} is a pattern, and where each $match_i$ is of the general form:

$$= e_i \ \texttt{where} \ \{ \ decls_i \ \}$$

or

$$\begin{array}{ll} \mid gs_{i1} & = e_{i1} \\ \ldots & \\ \mid gs_{im_i} & = e_{im_i} \\ & \texttt{where} \ \{ \ decls_i \ \} \end{array}$$

and where $n \geq 1$, $1 \leq i \leq n$, $m_i \geq 1$. The former is treated as shorthand for a particular case of the latter, namely:

$$\mid \texttt{True} = e_i \ \texttt{where} \ \{ \ decls_i \ \}$$

Note that all clauses defining a function must be contiguous, and the number of patterns in each clause must be the same. The set of patterns corresponding to each match must be *linear*—no variable is allowed to appear more than once in the entire set.

Alternative syntax is provided for binding functional values to infix operators. For example, these three function definitions are all equivalent:

```
plus x y z = x+y+z
x `plus` y = \ z -> x+y+z
(x `plus` y) z = x+y+z
```

Note that fixity resolution applies to the infix variants of the function binding in the same way as for expressions (Section 10.6). Applying fixity resolution to the left side of the equals in a function binding must leave the *varop* being defined at the top level. For example, if we are defining a new operator `##` with precedence 6, then this definition would be illegal:

```
a ## b : xs = exp
```

because `:` has precedence 5, so the left hand side resolves to `(a ## x) : xs`, and this cannot be a pattern binding because `(a ## x)` is not a valid pattern.

Translation: The general binding form for functions is semantically equivalent to the equation (i.e. simple pattern binding):

$$x = \backslash\ x_1 ... x_k \texttt{ -> case } (x_1, ..., x_k) \texttt{ of } \begin{array}{l} (p_{11}, ..., p_{1k})\ match_1 \\ ... \\ (p_{n1}, ..., p_{nk})\ match_n \end{array}$$

where the x_i are new identifiers.

4.4.3.2 Pattern bindings

A pattern binding binds variables to values. A *simple* pattern binding has form $p = e$. The pattern p is matched "lazily" as an irrefutable pattern, as if there were an implicit ~ in front of it. See the translation in Section Section 3.12.

The *general* form of a pattern binding is $p\ match$, where a $match$ is the same structure as for function bindings above; in other words, a pattern binding is:

$$\begin{array}{ll} p & \mid gs_1 = e_1 \\ & \mid gs_2 = e_2 \\ & ... \\ & \mid gs_m = e_m \\ & \texttt{where } \{\ decls\ \} \end{array}$$

Translation: The pattern binding above is semantically equivalent to this simple pattern binding:

$$\begin{array}{lll} p & = & \texttt{let } decls \texttt{ in} \\ & & \texttt{case } () \texttt{ of} \\ & & \quad () \mid gs_1 \rightarrow e_1 \\ & & \quad\quad \mid gs_2 \rightarrow e_2 \\ & & \quad\quad\quad ... \\ & & \quad\quad \mid gs_m \rightarrow e_m \\ & & \quad \_ \rightarrow \texttt{error "Unmatched pattern"} \end{array}$$

4.5 Static Semantics of Function and Pattern Bindings

The static semantics of the function and pattern bindings of a `let` expression or `where` clause are discussed in this section.

4.5.1 Dependency Analysis

In general the static semantics are given by applying the normal Hindley-Milner inference rules. In order to increase polymorphism, these rules are applied to groups of bindings identified by a *dependency analysis*.

A binding $b1$ *depends* on a binding $b2$ in the same list of declarations if either

1. $b1$ contains a free identifier that has no type signature and is bound by $b2$, or
2. $b1$ depends on a binding that depends on $b2$.

A *declaration group* is a minimal set of mutually dependent bindings. Hindley-Milner type inference is applied to each declaration group in dependency order. The order of declarations in `where`/`let` constructs is irrelevant.

4.5.2 Generalization

The Hindley-Milner type system assigns types to a let-expression in two stages:

1. The declaration groups are considered in dependency order. For each group, a type with no universal quantification is inferred for each variable bound in the group. Then, all type variables that occur in these types are universally quantified unless they are associated with bound variables in the type environment; this is called generalization.
2. Finally, the body of the let-expression is typed.

For example, consider the declaration

```
f x = let g y = (y,y)
      in ...
```

The type of `g`'s definition is $a \to (a, a)$. The generalization step attributes to `g` the polymorphic type $\forall a.a \to (a, a)$, after which the typing of the "`...`" part can proceed.

When typing overloaded definitions, all the overloading constraints from a single declaration group are collected together, to form the context for the type of each variable declared in the group. For example, in the definition:

```
f x = let g1 x y = if x>y then show x else g2 y x
          g2 p q = g1 q p
      in ...
```

The types of the definitions of `g1` and `g2` are both $a \to a \to$ `String`, and the accumulated constraints are `Ord` a (arising from the use of `>`), and `Show` a (arising from the use of `show`). The type variables appearing in this collection of constraints are called the *constrained type variables*.

The generalization step attributes to both `g1` and `g2` the type

$$\forall a.(\texttt{Ord}\ a, \texttt{Show}\ a) \Rightarrow a \to a \to \texttt{String}$$

Notice that `g2` is overloaded in the same way as `g1` even though the occurrences of `>` and `show` are in the definition of `g1`.

If the programmer supplies explicit type signatures for more than one variable in a declaration group, the contexts of these signatures must be identical up to renaming of the type variables.

4.5.3 Context Reduction Errors

As mentioned in Section 4.1.4, the context of a type may constrain only a type variable, or the application of a type variable to one or more types. Hence, types produced by generalization must be expressed in a form in which all context constraints have be reduced to this "head normal form". Consider, for example, the definition:

```
f xs y  =  xs == [y]
```

Its type is given by

```
f :: Eq a => [a] -> a -> Bool
```

and not

```
f :: Eq [a] => [a] -> a -> Bool
```

Even though the equality is taken at the list type, the context must be simplified, using the instance declaration for `Eq` on lists, before generalization. If no such instance is in scope, a static error occurs.

Here is an example that shows the need for a constraint of the form $C(m\ t)$ where m is one of the type variables being generalized; that is, where the class C applies to a type expression that is not a type variable or a type constructor. Consider:

```
f :: (Monad m, Eq (m a)) => a -> m a -> Bool
f x y = return x == y
```

The type of `return` is `Monad m => a -> m a`; the type of `(==)` is `Eq a => a -> a -> Bool`. The type of `f` should be therefore `(Monad m, Eq (m a)) => a -> m a -> Bool`, and the context cannot be simplified further.

The instance declaration derived from a data type `deriving` clause (see Section 4.3.3) must, like any instance declaration, have a *simple* context; that is, all the constraints must be of the form $C\ a$, where a is a type variable. For example, in the type

```
data Apply a b = App (a b)  deriving Show
```

the derived Show instance will produce a context `Show (a b)`, which cannot be reduced and is not simple; thus a static error results.

4.5.4 Monomorphism

Sometimes it is not possible to generalize over all the type variables used in the type of the definition. For example, consider the declaration

```
f x = let g y z = ([x,y], z)
      in ...
```

In an environment where x has type a, the type of g's definition is $a \to b \to ([a], b)$. The generalization step attributes to mbox{tt g} the type $\forall b.a \to b \to ([a], b)$; only b can be universally quantified because a occurs in the type environment. We say that the type of g is *monomorphic in the type variable* a.

The effect of such monomorphism is that the first argument of all applications of g must be of a single type. For example, it would be valid for the "`...`" to be

```
(g True, g False)
```

(which would, incidentally, force x to have type `Bool`) but invalid for it to be

```
(g True, g 'c')
```

In general, a type $\forall \overline{u}.cx \Rightarrow t$ is said to be *monomorphic* in the type variable a if a is free in $\forall \overline{u}.cx \Rightarrow t$.

It is worth noting that the explicit type signatures provided by Haskell are not powerful enough to express types that include monomorphic type variables. For example, we cannot write

```
f x = let
        g :: a -> b -> ([a],b)
        g y z = ([x,y], z)
      in ...
```

because that would claim that g was polymorphic in both a and b (Section 4.4.1). In this program, g can only be given a type signature if its first argument is restricted to a type not involving type variables; for example

```
g :: Int -> b -> ([Int],b)
```

This signature would also cause `x` to have type `Int`.

4.5.5 The Monomorphism Restriction

Haskell places certain extra restrictions on the generalization step, beyond the standard Hindley-Milner restriction described above, which further reduces polymorphism in particular cases.

The monomorphism restriction depends on the binding syntax of a variable. Recall that a variable is bound by either a *function binding* or a *pattern binding*, and that a *simple* pattern binding is a pattern binding in which the pattern consists of only a single variable (Section 4.4.3).

The following two rules define the monomorphism restriction:

The monomorphism restriction

Rule 1. We say that a given declaration group is *unrestricted* if and only if:

- **(a)** every variable in the group is bound by a function binding or a simple pattern binding (Section 4.4.3.2), *and*
- **(b)** an explicit type signature is given for every variable in the group that is bound by simple pattern binding.

The usual Hindley-Milner restriction on polymorphism is that only type variables that do not occur free in the environment may be generalized. In addition, *the constrained type variables of a restricted declaration group may not be generalized* in the generalization step for that group. (Recall that a type variable is constrained if it must belong to some type class; see Section 4.5.2.)

Rule 2. Any monomorphic type variables that remain when type inference for an entire module is complete, are considered *ambiguous*, and are resolved to particular types using the defaulting rules (Section 4.3.4).

Motivation Rule 1 is required for two reasons, both of which are fairly subtle.

- *Rule 1 prevents computations from being unexpectedly repeated.* For example, `genericLength` is a standard function (in library `Data.List`) whose type is given by

  ```
  genericLength :: Num a => [b] -> a
  ```

  Now consider the following expression:

  ```
  let { len = genericLength xs } in (len, len)
  ```

  It looks as if `len` should be computed only once, but without Rule 1 it might be computed twice, once at each of two different overloadings. If the programmer does actually wish the computation to be repeated, an explicit type signature may be added:

  ```
  let { len :: Num a => a; len = genericLength xs } in (len, len)
  ```

- *Rule 1 prevents ambiguity.* For example, consider the declaration group

  ```
  [(n,s)] = reads t
  ```

  Recall that `reads` is a standard function whose type is given by the signature

  ```
  reads :: (Read a) => String -> [(a,String)]
  ```

  Without Rule 1, `n` would be assigned the type $\forall a$. `Read` $a \Rightarrow a \rightarrow a$ and `s` the type $\forall a$. `Read` $a \Rightarrow$ `String`. The latter is an invalid type, because it is inherently ambiguous. It is not possible to determine at what overloading to use `s`, nor can this be solved by adding a type signature for `s`.

Hence, when *non-simple* pattern bindings are used (Section 4.4.3.2), the types inferred are always monomorphic in their constrained type variables, irrespective of whether a type signature is provided. In this case, both `n` and `s` are monomorphic in *a*.

The same constraint applies to pattern-bound functions. For example, in

```
(f,g) = ((+),(-))
```

both `f` and `g` are monomorphic regardless of any type signatures supplied for `f` or `g`.

Rule 2 is required because there is no way to enforce monomorphic use of an *exported* binding, except by performing type inference on modules outside the current module. Rule 2 states that the exact types of all the variables bound in a module must be determined by that module alone, and not by any modules that import it.

```
module M1(len1) where
  default( Int, Double )
  len1 = genericLength "Hello"

module M2 where
  import M1(len1)
  len2 = (2*len1) :: Rational
```

When type inference on module `M1` is complete, `len1` has the monomorphic type `Num a => a` (by Rule 1). Rule 2 now states that the monomorphic type variable `a` is ambiguous, and must be resolved using the defaulting rules of Section 4.3.4. Hence, `len1` gets type `Int`, and its use in `len2` is type-incorrect. (If the above code is actually what is wanted, a type signature on `len1` would solve the problem.)

This issue does not arise for nested bindings, because their entire scope is visible to the compiler.

Consequences The monomorphism rule has a number of consequences for the programmer. Anything defined with function syntax usually generalizes as a function is expected to. Thus in

```
f x y = x+y
```

the function `f` may be used at any overloading in class `Num`. There is no danger of recomputation here. However, the same function defined with pattern syntax:

```
f = \x -> \y -> x+y
```

requires a type signature if `f` is to be fully overloaded. Many functions are most naturally defined using simple pattern bindings; the user must be careful to affix these with type signatures to retain full overloading. The standard prelude contains many examples of this:

```
sum  :: (Num a) => [a] -> a
sum  =  foldl (+) 0
```

Rule 1 applies to both top-level and nested definitions. Consider

```
module M where
  len1 = genericLength "Hello"
  len2 = (2*len1) :: Rational
```

Here, type inference finds that `len1` has the monomorphic type (`Num a => a`); and the type variable `a` is resolved to `Rational` when performing type inference on `len2`.

4.6 Kind Inference

This section describes the rules that are used to perform *kind inference*, i.e. to calculate a suitable kind for each type constructor and class appearing in a given program.

The first step in the kind inference process is to arrange the set of datatype, synonym, and class definitions into dependency groups. This can be achieved in much the same way as the dependency analysis for value declarations that was described in Section 4.5.1. For example, the following program fragment includes the definition of a datatype constructor `D`, a synonym `S` and a class `C`, all of which would be included in the same dependency group:

```
data C a => D a = Foo (S a)
type S a = [D a]
class C a where
    bar :: a -> D a -> Bool
```

The kinds of variables, constructors, and classes within each group are determined using standard techniques of type inference and kind-preserving unification [6]. For example, in the definitions above, the parameter `a` appears as an argument of the function constructor `(->)` in the type of `bar` and hence must have kind $*$. It follows that both `D` and `S` must have kind $* \to *$ and that every instance of class `C` must have kind $*$.

It is possible that some parts of an inferred kind may not be fully determined by the corresponding definitions; in such cases, a default of $*$ is assumed. For example, we could assume an arbitrary kind κ for the `a` parameter in each of the following examples:

```
data App f a = A (f a)
data Tree a  = Leaf | Fork (Tree a) (Tree a)
```

This would give kinds $(\kappa \to *) \to \kappa \to *$ and $\kappa \to *$ for `App` and `Tree`, respectively, for any kind κ, and would require an extension to allow polymorphic kinds. Instead, using the default binding $\kappa = *$, the actual kinds for these two constructors are $(* \to *) \to * \to *$ and $* \to *$, respectively.

Defaults are applied to each dependency group without consideration of the ways in which particular type constructor constants or classes are used in later dependency groups or elsewhere in the program. For example, adding the following definition to those above does not influence the kind inferred for `Tree` (by changing it to $(* \to *) \to *$, for instance), and instead generates a static error because the kind of `[]`, $* \to *$, does not match the kind $*$ that is expected for an argument of `Tree`:

```
type FunnyTree = Tree []     -- invalid
```

This is important because it ensures that each constructor and class are used consistently with the same kind whenever they are in scope.

Chapter 5

Modules

A module defines a collection of values, datatypes, type synonyms, classes, etc. (see Chapter 4), in an environment created by a set of *imports* (resources brought into scope from other modules). It *exports* some of these resources, making them available to other modules. We use the term *entity* to refer to a value, type, or class defined in, imported into, or perhaps exported from a module.

A Haskell *program* is a collection of modules, one of which, by convention, must be called `Main` and must export the value `main`. The *value* of the program is the value of the identifier `main` in module `Main`, which must be a computation of type `IO` τ for some type τ (see Chapter 7). When the program is executed, the computation `main` is performed, and its result (of type τ) is discarded.

Modules may reference other modules via explicit `import` declarations, each giving the name of a module to be imported and specifying its entities to be imported. Modules may be mutually recursive.

Modules are used for name-space control, and are not first class values. A multi-module Haskell program can be converted into a single-module program by giving each entity a unique name, changing all occurrences to refer to the appropriate unique name, and then concatenating all the module bodies[3]. For example, here is a three-module program:

```
module Main where
import A
import B
main = A.f >> B.f

module A where
f = ...

module B where
f = ...
```

It is equivalent to the following single-module program:

```
module Main where
main = af >> bf

af = ...

bf = ...
```

Because they are allowed to be mutually recursive, modules allow a program to be partitioned freely without regard to dependencies.

[3]There are two minor exceptions to this statement. First, `default` declarations scope over a single module (Section 4.3.4). Second, Rule 2 of the monomorphism restriction (Section 4.5.5) is affected by module boundaries.

A module name (lexeme *modid*) is a sequence of one or more identifiers beginning with capital letters, separated by dots, with no intervening spaces. For example, `Data.Bool`, `Main` and `Foreign.Marshal.Alloc` are all valid module names.

$$modid \rightarrow \{conid\ .\}\ conid$$

Module names can be thought of as being arranged in a hierarchy in which appending a new component creates a child of the original module name. For example, the module `Control.Monad.ST` is a child of the `Control.Monad` sub-hierarchy. This is purely a convention, however, and not part of the language definition; in this report a *modid* is treated as a single identifier occupying a flat namespace.

There is one distinguished module, `Prelude`, which is imported into all modules by default (see Chapter 9), plus a set of standard library modules that may be imported as required.

5.1 Module Structure

A module defines a mutually recursive scope containing declarations for value bindings, data types, type synonyms, classes, etc. (see Chapter 4).

$$\begin{array}{lcll}
module & \rightarrow & \texttt{module}\ modid\ [exports]\ \texttt{where}\ body & \\
 & | & body & \\
body & \rightarrow & \{\ impdecls\ ;\ topdecls\ \} & \\
 & | & \{\ impdecls\ \} & \\
 & | & \{\ topdecls\ \} & \\
impdecls & \rightarrow & impdecl_1\ ;\ \ldots\ ;\ impdecl_n & (n \geq 1) \\
topdecls & \rightarrow & topdecl_1\ ;\ \ldots\ ;\ topdecl_n & (n \geq 1)
\end{array}$$

A module begins with a header: the keyword `module`, the module name, and a list of entities (enclosed in round parentheses) to be exported. The header is followed by a possibly-empty list of `import` declarations (*impdecls*, Section 5.3) that specify modules to be imported, optionally restricting the imported bindings. This is followed by a possibly-empty list of top-level declarations (*topdecls*, Chapter 4).

An abbreviated form of module, consisting only of the module body, is permitted. If this is used, the header is assumed to be `module Main(main) where`. If the first lexeme in the abbreviated module is not a `{`, then the layout rule applies for the top level of the module.

5.2 Export Lists

$$\begin{array}{lcll}
exports & \rightarrow & (\ export_1\ ,\ \ldots\ ,\ export_n[,]) & (n \geq 0) \\
export & \rightarrow & qvar & \\
 & | & qtycon\ [(..)\ |\ (\ cname_1\ ,\ \ldots\ ,\ cname_n\ )] & (n \geq 0) \\
 & | & qtycls\ [(..)\ |\ (\ var_1\ ,\ \ldots\ ,\ var_n\ )] & (n \geq 0) \\
 & | & \texttt{module}\ modid &
\end{array}$$

cname → *var* | *con*

An *export list* identifies the entities to be exported by a module declaration. A module implementation may only export an entity that it declares, or that it imports from some other module. If the export list is omitted, all values, types and classes defined in the module are exported, *but not those that are imported.*

Entities in an export list may be named as follows:

1. A value, field name, or class method, whether declared in the module body or imported, may be named by giving the name of the value as a *qvarid*, which must be in scope. Operators should be enclosed in parentheses to turn them into *qvarid*s.
2. An algebraic datatype T declared by a `data` or `newtype` declaration may be named in one of three ways:
   - The form T names the type *but not the constructors or field names.* The ability to export a type without its constructors allows the construction of abstract datatypes (see Section 5.8).
   - The form $T(c_1, ..., c_n)$, names the type and some or all of its constructors and field names.
   - The abbreviated form $T(..)$ names the type and all its constructors and field names that are currently in scope (whether qualified or not).

   In all cases, the (possibly-qualified) type constructor T must be in scope. The constructor and field names c_i in the second form are unqualified; one of these subordinate names is legal if and only if (a) it names a constructor or field of T, and (b) the constructor or field is in scope in the module body *regardless of whether it is in scope under a qualified or unqualified name.* For example, the following is legal

   ```
   module A( Mb.Maybe( Nothing, Just ) ) where
     import qualified Data.Maybe as Mb
   ```

   Data constructors cannot be named in export lists except as subordinate names, because they cannot otherwise be distinguished from type constructors.
3. A type synonym T declared by a `type` declaration may be named by the form T, where T is in scope.
4. A class C with operations $f_1, ..., f_n$ declared in a `class` declaration may be named in one of three ways:
   - The form C names the class *but not the class methods.*
   - The form $C(f_1, ..., f_n)$, names the class and some or all of its methods.
   - The abbreviated form $C(..)$ names the class and all its methods that are in scope (whether qualified or not).

   In all cases, C must be in scope. In the second form, one of the (unqualified) subordinate names f_i is legal if and only if (a) it names a class method of C, and (b) the class method is in scope in the module body regardless of whether it is in scope under a qualified or unqualified name.
5. The form "`module M`" names the set of all entities that are in scope with both an unqualified name "`e`" and a qualified name "`M.e`". This set may be empty. For example:

   ```
   module Queue( module Stack, enqueue, dequeue ) where
   import Stack
   ...
   ```

   Here the module `Queue` uses the module name `Stack` in its export list to abbreviate all the entities imported from `Stack`.

A module can name its own local definitions in its export list using its own name in the "`module M`" syntax, because a local declaration brings into scope both a qualified and unqualified name (Section 5.5.1). For example:

```
module Mod1( module Mod1, module Mod2 ) where
import Mod2
import Mod3
```

Here module `Mod1` exports all local definitions as well as those imported from `Mod2` but not those imported from `Mod3`.

It is an error to use `module M` in an export list unless `M` is the module bearing the export list, or `M` is imported by at least one import declaration (qualified or unqualified).

Exports lists are cumulative: the set of entities exported by an export list is the union of the entities exported by the individual items of the list.

It makes no difference to an importing module how an entity was exported. For example, a field name `f` from data type `T` may be exported individually (`f`, item (1) above); or as an explicitly-named member of its data type (`T(f)`, item (2)); or as an implicitly-named member (`T(..)`, item(2)); or by exporting an entire module (`module M`, item (5)).

The *unqualified* names of the entities exported by a module must all be distinct (within their respective namespace). For example

```
module A ( C.f, C.g, g, module B ) where   -- an invalid module
import B(f)
import qualified C(f,g)
g = f True
```

There are no name clashes within module `A` itself, but there are name clashes in the export list between `C.g` and `g` (assuming `C.g` and `g` are different entities – remember, modules can import each other recursively), and between `module B` and `C.f` (assuming `B.f` and `C.f` are different entities).

5.3 Import Declarations

impdecl	→	`import` [`qualified`] *modid* [`as` *modid*][*impspec*]	
	|		(empty declaration)
impspec	→	($import_1$, … , $import_n$[,])	($n \geq 0$)
	|	`hiding` ($import_1$, … , $import_n$[,])	($n \geq 0$)
import	→	*var*	
	|	*tycon* [(..) | ($cname_1$, … , $cname_n$)]	($n \geq 0$)
	|	*tycls* [(..) | (var_1 , … , var_n)]	($n \geq 0$)
cname	→	*var* | *con*	

The entities exported by a module may be brought into scope in another module with an `import` declaration at the beginning of the module. The `import` declaration names the module to be imported and optionally specifies the entities to be imported. A single module may be imported by more than one `import` declaration. Imported names serve as top level declarations: they scope over the entire body of the module but may be shadowed by local non-top-level bindings.

The effect of multiple `import` declarations is strictly cumulative: an entity is in scope if it is imported by any of the `import` declarations in a module. The ordering of import declarations is irrelevant.

Lexically, the terminal symbols "`as`", "`qualified`" and "`hiding`" are each a *varid* rather than a *reservedid*. They have special significance only in the context of an `import` declaration; they may also be used as variables.

5.3.1 What is imported

Exactly which entities are to be imported can be specified in one of the following three ways:

1. The imported entities can be specified explicitly by listing them in parentheses. Items in the list have the same form as those in export lists, except qualifiers are not permitted and the "`module` *modid*" entity is not permitted. When the `(..)` form of import is used for a type or class, the `(..)` refers to all of the constructors, methods, or field names exported from the module.

   The list must name only entities exported by the imported module. The list may be empty, in which case nothing except the instances is imported.

2. Entities can be excluded by using the form `hiding` $(import_1, ..., import_n)$, which specifies that all entities exported by the named module should be imported except for those named in the list. Data constructors may be named directly in hiding lists without being prefixed by the associated type. Thus, in

   ```
   import M hiding (C)
   ```

   any constructor, class, or type named `C` is excluded. In contrast, using `C` in an import list names only a class or type.

   It is an error to hide an entity that is not, in fact, exported by the imported module.

3. Finally, if *impspec* is omitted then all the entities exported by the specified module are imported.

5.3.2 Qualified import

For each entity imported under the rules of Section 5.3.1, the top-level environment is extended. If the import declaration used the `qualified` keyword, only the *qualified name* of the entity is brought into scope. If the `qualified` keyword is omitted, then *both* the qualified *and* unqualified name of the entity is brought into scope. Section 5.5.1 describes qualified names in more detail.

The qualifier on the imported name is either the name of the imported module, or the local alias given in the `as` clause (Section 5.3.3) on the `import` statement. Hence, *the qualifier is not necessarily the name of the module in which the entity was originally declared.*

The ability to exclude the unqualified names allows full programmer control of the unqualified namespace: a locally defined entity can share the same name as a qualified import:

```
module Ring where
import qualified Prelude     -- All Prelude names must be qualified
import Data.List( nub )

l1 + l2 = l1 Prelude.++ l2  -- This + differs from the one in the Prelude
l1 * l2 = nub (l1 + l2)     -- This * differs from the one in the Prelude

succ = (Prelude.+ 1)
```

5.3.3 Local aliases

Imported modules may be assigned a local alias in the importing module using the `as` clause. For example, in

```
import qualified VeryLongModuleName as C
```

entities must be referenced using "`C.`" as a qualifier instead of "`VeryLongModuleName.`". This also allows a different module to be substituted for `VeryLongModuleName` without changing the qualifiers used for the imported module. It is legal for more than one module in scope to use the same qualifier, provided that all names can still be resolved unambiguously. For example:

```
module M where
import qualified Foo as A
import qualified Baz as A
x = A.f
```

This module is legal provided only that `Foo` and `Baz` do not both export `f`.

An `as` clause may also be used on an un-`qualified` `import` statement:

```
import Foo as A(f)
```

This declaration brings into scope `f` and `A.f`.

5.3.4 Examples

To clarify the above import rules, suppose the module `A` exports `x` and `y`. Then this table shows what names are brought into scope by the specified import statement:

Import declaration	Names brought into scope
`import A`	`x`, `y`, `A.x`, `A.y`
`import A()`	(nothing)
`import A(x)`	`x`, `A.x`
`import qualified A`	`A.x`, `A.y`
`import qualified A()`	(nothing)
`import qualified A(x)`	`A.x`
`import A hiding ()`	`x`, `y`, `A.x`, `A.y`
`import A hiding (x)`	`y`, `A.y`
`import qualified A hiding ()`	`A.x`, `A.y`
`import qualified A hiding (x)`	`A.y`
`import A as B`	`x`, `y`, `B.x`, `B.y`
`import A as B(x)`	`x`,`B.x`
`import qualified A as B`	`B.x`, `B.y`

In all cases, all instance declarations in scope in module `A` are imported (Section 5.4).

5.4 Importing and Exporting Instance Declarations

Instance declarations cannot be explicitly named on import or export lists. All instances in scope within a module are *always* exported and any import brings *all* instances in from the imported module. Thus, an instance declaration is in scope if and only if a chain of `import` declarations leads to the module containing the instance declaration.

For example, `import M()` does not bring any new names in scope from module `M`, but does bring in any instances visible in `M`. A module whose only purpose is to provide instance declarations can have an empty export list. For example

```
module MyInstances() where
instance Show (a -> b) where
  show fn = "<<function>>"
instance Show (IO a) where
  show io = "<<IO action>>"
```

5.5 Name Clashes and Closure

5.5.1 Qualified names

A *qualified name* is written as *modid.name* (Section 2.4). A qualified name is brought into scope:

- *By a top level declaration.* A top-level declaration brings into scope both the unqualified *and* the qualified name of the entity being defined. Thus:

  ```
  module M where
  f x = ...
  g x = M.f x x
  ```

  is legal. The *defining* occurrence must mention the *unqualified* name; therefore, it is illegal to write

  ```
  module M where
  M.f x = ...                  -- ILLEGAL
  g x = let M.y = x+1 in ...   -- ILLEGAL
  ```

- *By an `import` declaration.* An `import` declaration, whether `qualified` or not, always brings into scope the qualified name of the imported entity (Section 5.3). This allows a qualified import to be replaced with an unqualified one without forcing changes in the references to the imported names.

5.5.2 Name clashes

If a module contains a bound occurrence of a name, such as `f` or `A.f`, it must be possible unambiguously to resolve which entity is thereby referred to; that is, there must be only one binding for `f` or `A.f` respectively.

It is *not* an error for there to exist names that cannot be so resolved, provided that the program does not mention those names. For example:

```
module A where
import B
import C
tup = (b, c, d, x)

module B( d, b, x, y ) where
import D
x = ...
y = ...
b = ...

module C( d, c, x, y ) where
import D
x = ...
```

```
y = ...
c = ...

module D( d ) where
d = ...
```

Consider the definition of `tup`.

- The references to `b` and `c` can be unambiguously resolved to `b` declared in `B`, and `c` declared in `C` respectively.
- The reference to `d` is unambiguously resolved to `d` declared in `D`. In this case the same entity is brought into scope by two routes (the import of `B` and the import of `C`), and can be referred to in `A` by the names `d`, `B.d`, and `C.d`.
- The reference to `x` is ambiguous: it could mean `x` declared in `B`, or `x` declared in `C`. The ambiguity could be fixed by replacing the reference to `x` by `B.x` or `C.x`.
- There is no reference to `y`, so it is not erroneous that distinct entities called `y` are exported by both `B` and `C`. An error is only reported if `y` is actually mentioned.

The name occurring in a type signature or fixity declarations is always unqualified, and unambiguously refers to another declaration in the same declaration list (except that the fixity declaration for a class method can occur at top level — Section 4.4.2). For example, the following module is legal:

```
module F where

sin :: Float -> Float
sin x = (x::Float)

f x = Prelude.sin (F.sin x)
```

The local declaration for `sin` is legal, even though the Prelude function `sin` is implicitly in scope. The references to `Prelude.sin` and `F.sin` must both be qualified to make it unambiguous which `sin` is meant. However, the unqualified name `sin` in the type signature in the first line of `F` unambiguously refers to the local declaration for `sin`.

5.5.3 Closure

Every module in a Haskell program must be *closed*. That is, every name explicitly mentioned by the source code must be either defined locally or imported from another module. However, entities that the compiler requires for type checking or other compile time analysis need not be imported if they are not mentioned by name. The Haskell compilation system is responsible for finding any information needed for compilation without the help of the programmer. That is, the import of a variable `x` does not require that the datatypes and classes in the signature of `x` be brought into the module along with `x` unless these entities are referenced by name in the user program. The Haskell system silently imports any information that must accompany an entity for type checking or any other purposes. Such entities need not even be explicitly exported: the following program is valid even though `T` does not escape `M1`:

```
module M1(x) where
data T = T
x = T

module M2 where
import M1(x)
y = x
```

In this example, there is no way to supply an explicit type signature for `y` since `T` is not in scope. Whether or not `T` is explicitly exported, module `M2` knows enough about `T` to correctly type check the program.

The type of an exported entity is unaffected by non-exported type synonyms. For example, in

```
module M(x) where
type T = Int
x :: T
x = 1
```

the type of `x` is both `T` and `Int`; these are interchangeable even when `T` is not in scope. That is, the definition of `T` is available to any module that encounters it whether or not the name `T` is in scope. The only reason to export `T` is to allow other modules to refer it by name; the type checker finds the definition of `T` if needed whether or not it is exported.

5.6 Standard Prelude

Many of the features of Haskell are defined in Haskell itself as a library of standard datatypes, classes, and functions, called the "Standard Prelude." In Haskell, the Prelude is contained in the module `Prelude`. There are also many predefined library modules, which provide less frequently used functions and types. For example, complex numbers, arrays, and most of the input/output are all part of the standard libraries. Separating libraries from the Prelude has the advantage of reducing the size and complexity of the Prelude, allowing it to be more easily assimilated, and increasing the space of useful names available to the programmer.

Prelude and library modules differ from other modules in that their semantics (but not their implementation) are a fixed part of the Haskell language definition. This means, for example, that a compiler may optimize calls to functions in the Prelude without consulting the source code of the Prelude.

5.6.1 The Prelude Module

The `Prelude` module is imported automatically into all modules as if by the statement `import Prelude`, if and only if it is not imported with an explicit `import` declaration. This provision for explicit import allows entities defined in the Prelude to be selectively imported, just like those from any other module.

The semantics of the entities in `Prelude` is specified by a reference implementation of `Prelude` written in Haskell, given in Chapter 9. Some datatypes (such as `Int`) and functions (such as `Int` addition) cannot be specified directly in Haskell. Since the treatment of such entities depends on the implementation, they are not formally defined in Chapter 9. The implementation of `Prelude` is also incomplete in its treatment of tuples: there should be an infinite family of tuples and their instance declarations, but the implementation only gives a scheme.

Chapter 9 defines the module `Prelude` using several other modules: `PreludeList`, `PreludeIO`, and so on. These modules are *not* part of Haskell, and they cannot be imported separately. They are simply there to help explain the structure of the `Prelude` module; they should be considered part of its implementation, not part of the language definition.

5.6.2 Shadowing Prelude Names

The rules about the Prelude have been cast so that it is possible to use Prelude names for nonstandard purposes; however, every module that does so must have an `import` declaration that makes this nonstandard usage explicit. For example:

```
module A( null, nonNull ) where
import Prelude hiding( null )
null, nonNull :: Int -> Bool
null    x = x == 0
nonNull x = not (null x)
```

Module `A` redefines `null`, and contains an unqualified reference to `null` on the right hand side of `nonNull`. The latter would be ambiguous without the `hiding(null)` on the `import Prelude` statement. Every module that imports `A` unqualified, and then makes an unqualified reference to `null` must also resolve the ambiguous use of `null` just as `A` does. Thus there is little danger of accidentally shadowing Prelude names.

It is possible to construct and use a different module to serve in place of the Prelude. Other than the fact that it is implicitly imported, the Prelude is an ordinary Haskell module; it is special only in that some objects in the Prelude are referenced by special syntactic constructs. Redefining names used by the Prelude does not affect the meaning of these special constructs. For example, in

```
module B where
import Prelude()
import MyPrelude
f x = (x,x)
g x = (,) x x
h x = [x] ++ []
```

the explicit `import Prelude()` declaration prevents the automatic import of `Prelude`, while the declaration `import MyPrelude` brings the non-standard prelude into scope. The special syntax for tuples (such as `(x,x)` and `(,)`) and lists (such as `[x]` and `[]`) continues to refer to the tuples and lists defined by the standard `Prelude`; there is no way to redefine the meaning of `[x]`, for example, in terms of a different implementation of lists. On the other hand, the use of `++` is not special syntax, so it refers to `++` imported from `MyPrelude`.

It is not possible, however, to hide `instance` declarations in the `Prelude`. For example, one cannot define a new instance for `Show Char`.

5.7 Separate Compilation

Depending on the Haskell implementation used, separate compilation of mutually recursive modules may require that imported modules contain additional information so that they may be referenced before they are compiled. Explicit type signatures for all exported values may be necessary to deal with mutual recursion. The precise details of separate compilation are not defined by this report.

5.8 Abstract Datatypes

The ability to export a datatype without its constructors allows the construction of abstract datatypes (ADTs). For example, an ADT for stacks could be defined as:

```
module Stack( StkType, push, pop, empty ) where
data StkType a = EmptyStk | Stk a (StkType a)
push x s = Stk x s
```

```
pop (Stk _ s) = s
empty = EmptyStk
```

Modules importing `Stack` cannot construct values of type `StkType` because they do not have access to the constructors of the type. Instead, they must use `push`, `pop`, and `empty` to construct such values.

It is also possible to build an ADT on top of an existing type by using a `newtype` declaration. For example, stacks can be defined with lists:

```
module Stack( StkType, push, pop, empty ) where
newtype StkType a = Stk [a]
push x (Stk s) = Stk (x:s)
pop (Stk (_:s)) = Stk s
empty = Stk []
```

Chapter 6

Predefined Types and Classes

The Haskell Prelude contains predefined classes, types, and functions that are implicitly imported into every Haskell program. In this chapter, we describe the types and classes found in the Prelude. Most functions are not described in detail here as they can easily be understood from their definitions as given in Chapter 9. Other predefined types such as arrays, complex numbers, and rationals are defined in part:libraries[Part]

6.1 Standard Haskell Types

These types are defined by the Haskell Prelude. Numeric types are described in Section 6.4 When appropriate, the Haskell definition of the type is given. Some definitions may not be completely valid on syntactic grounds but they faithfully convey the meaning of the underlying type.

6.1.1 Booleans

```
data  Bool  =  False | True deriving
                              (Read, Show, Eq, Ord, Enum, Bounded)
```

The boolean type `Bool` is an enumeration. The basic boolean functions are `&&` (and), `||` (or), and `not`. The name `otherwise` is defined as `True` to make guarded expressions more readable.

6.1.2 Characters and Strings

The character type `Char` is an enumeration whose values represent Unicode characters [3]. The lexical syntax for characters is defined in Section 2.6; character literals are nullary constructors in the datatype `Char`. Type `Char` is an instance of the classes `Read`, `Show`, `Eq`, `Ord`, `Enum`, and `Bounded`. The `toEnum` and `fromEnum` functions, standard functions from class `Enum`, map characters to and from the `Int` type.

Note that ASCII control characters each have several representations in character literals: numeric escapes, ASCII mnemonic escapes, and the `\^X` notation. In addition, there are the following equivalences: `\a` and `\BEL`, `\b` and `\BS`, `\f` and `\FF`, `\r` and `\CR`, `\t` and `\HT`, `\v` and `\VT`, and `\n` and `\LF`.

A *string* is a list of characters:

```
type  String  =  [Char]
```

Strings may be abbreviated using the lexical syntax described in Section 2.6. For example, `"A string"` abbreviates

```
['A', ' ', 's', 't', 'r', 'i', 'n', 'g']
```

6.1.3 Lists

```
data  [a]  =  [] | a : [a]  deriving (Eq, Ord)
```

Lists are an algebraic datatype of two constructors, although with special syntax, as described in Section 3.7. The first constructor is the null list, written `'[]'` ("nil"), and the second is `':'` ("cons"). The module `PreludeList` (see Section 9.1) defines many standard list functions. Arithmetic sequences and list comprehensions, two convenient syntaxes for special kinds of lists, are described in Section 3.10 and Section 3.11, respectively. Lists are an instance of classes `Read`, `Show`, `Eq`, `Ord`, `Monad`, `Functor`, and `MonadPlus`.

6.1.4 Tuples

Tuples are algebraic datatypes with special syntax, as defined in Section 3.8. Each tuple type has a single constructor. All tuples are instances of `Eq`, `Ord`, `Bounded`, `Read`, and `Show` (provided, of course, that all their component types are).

There is no upper bound on the size of a tuple, but some Haskell implementations may restrict the size of tuples, and limit the instances associated with larger tuples. However, every Haskell implementation must support tuples up to size 15, together with the instances for `Eq`, `Ord`, `Bounded`, `Read`, and `Show`. The Prelude and libraries define tuple functions such as `zip` for tuples up to a size of 7.

The constructor for a tuple is written by omitting the expressions surrounding the commas; thus `(x,y)` and `(,) x y` produce the same value. The same holds for tuple type constructors; thus, `(Int,Bool,Int)` and `(,,) Int Bool Int` denote the same type.

The following functions are defined for pairs (2-tuples): `fst`, `snd`, `curry`, and `uncurry`. Similar functions are not predefined for larger tuples.

6.1.5 The Unit Datatype

```
data  () = () deriving (Eq, Ord, Bounded, Enum, Read, Show)
```

The unit datatype `()` has one non-$\bot$ member, the nullary constructor `()`. See also Section 3.9.

6.1.6 Function Types

Functions are an abstract type: no constructors directly create functional values. The following simple functions are found in the Prelude: `id`, `const`, `(.)`, `flip`, `($)`, and `until`.

6.1.7 The IO and IOError Types

The `IO` type serves as a tag for operations (actions) that interact with the outside world. The `IO` type is abstract: no constructors are visible to the user. `IO` is an instance of the `Monad` and `Functor` classes. Chapter 7 describes I/O operations.

`IOError` is an abstract type representing errors raised by I/O operations. It is an instance of `Show` and `Eq`. Values of this type are constructed by the various I/O functions and are not presented in any further detail in this report. The Prelude contains a few I/O functions (defined in Section 9.3), and part:libraries[Part] contains many more.

6.1.8 Other Types

```
data  Maybe a     =  Nothing | Just a  deriving (Eq, Ord, Read, Show)
data  Either a b  =  Left a | Right b  deriving (Eq, Ord, Read, Show)
data  Ordering    =  LT | EQ | GT deriving
                                 (Eq, Ord, Bounded, Enum, Read, Show)
```

The `Maybe` type is an instance of classes `Functor`, `Monad`, and `MonadPlus`. The `Ordering` type is used by `compare` in the class `Ord`. The functions `maybe` and `either` are found in the Prelude.

6.2 Strict Evaluation

Function application in Haskell is non-strict; that is, a function argument is evaluated only when required. Sometimes it is desirable to force the evaluation of a value, using the `seq` function:

```
seq :: a -> b -> b
```

The function `seq` is defined by the equations:

$$seq \perp b = \perp$$
$$seq\ a\ b = b \quad (\text{if } a \neq \perp)$$

`seq` is usually introduced to improve performance by avoiding unneeded laziness. Strict datatypes (see Section 4.2.1) are defined in terms of the `$!` operator. However, the provision of `seq` has important semantic consequences, because it is available *at every type*. As a consequence, $\perp$ is not the same as `\x -> ⊥`, since `seq` can be used to distinguish them. For the same reason, the existence of `seq` weakens Haskell's parametricity properties.

The operator `$!` is strict (call-by-value) application, and is defined in terms of `seq`. The Prelude also defines the `$` operator to perform non-strict application.

```
infixr 0 $, $!
($), ($!) :: (a -> b) -> a -> b
f $  x  =          f x
f $! x  =  x `seq` f x
```

The non-strict application operator `$` may appear redundant, since ordinary application `(f x)` means the same as `(f $ x)`. However, `$` has low, right-associative binding precedence, so it sometimes allows parentheses to be omitted; for example:

```
f $ g $ h x  =  f (g (h x))
```

It is also useful in higher-order situations, such as `map ($ 0) xs`, or `zipWith ($) fs xs`.

6.3 Standard Haskell Classes

Figure 3 shows the hierarchy of Haskell classes defined in the Prelude and the Prelude types that are instances of these classes.

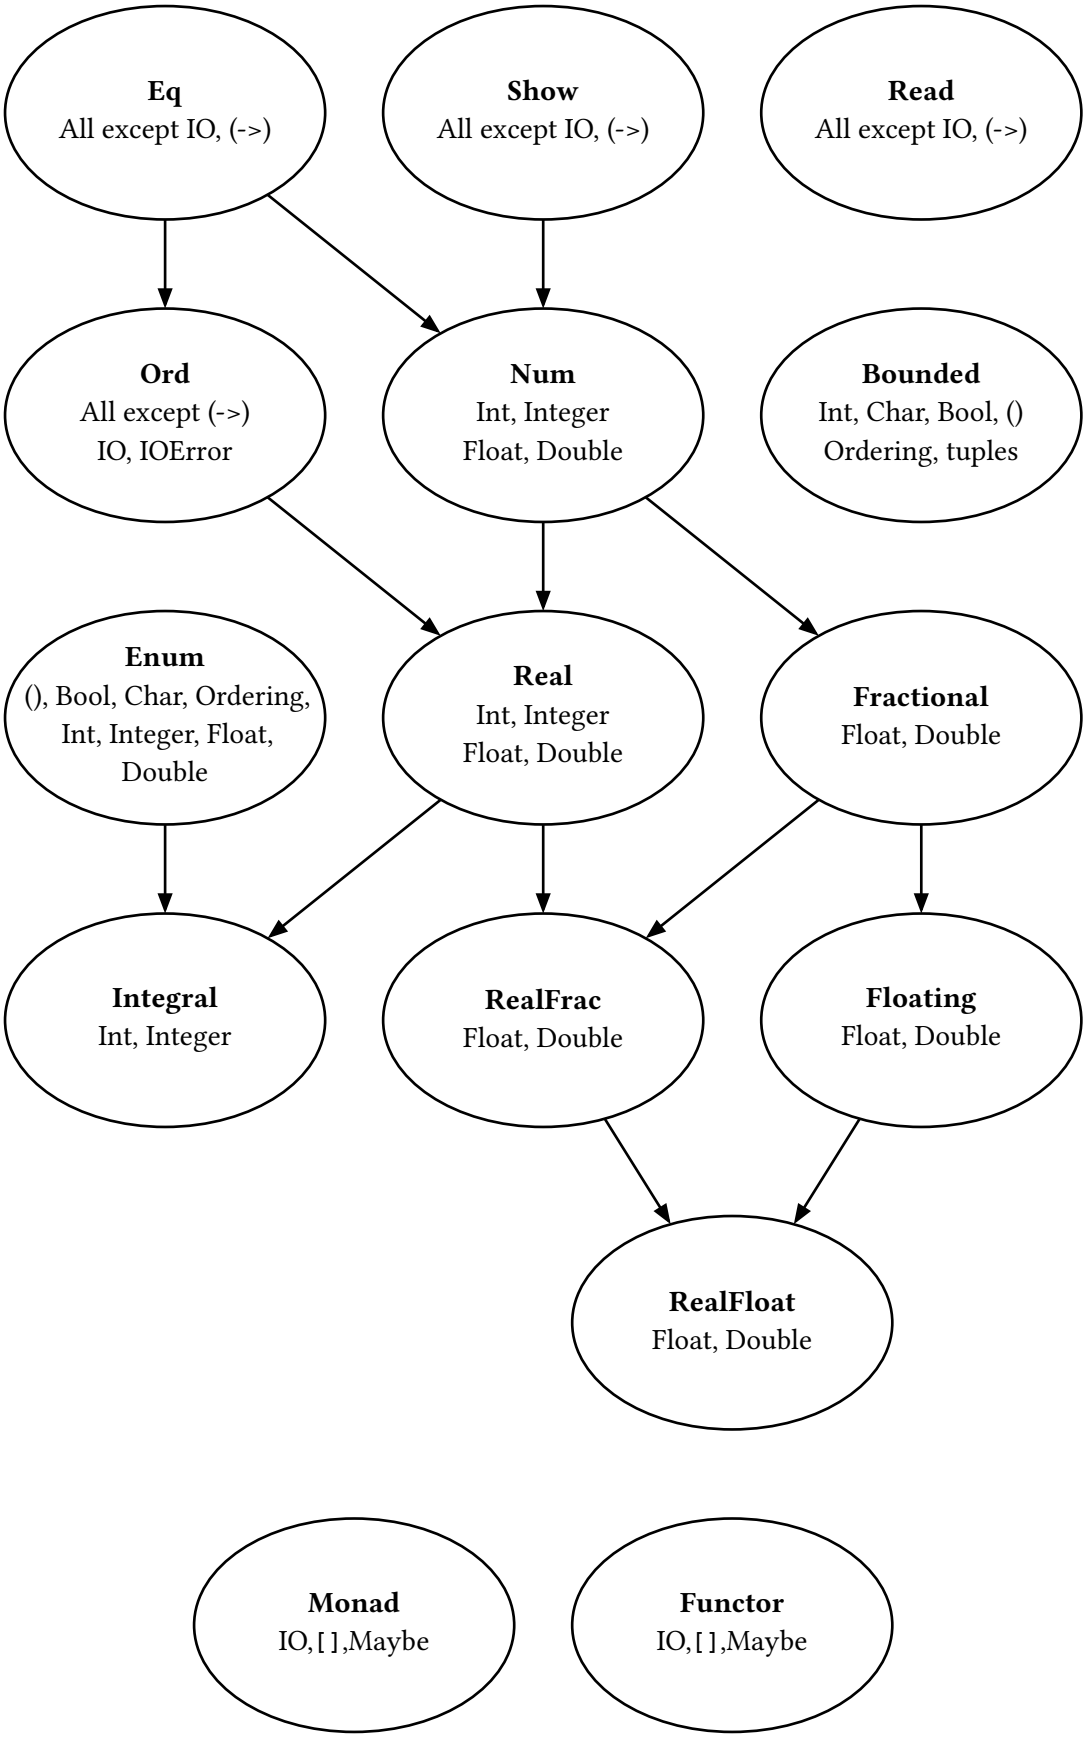

Figure 3: Standard Haskell Classes

Default class method declarations (Section 4.3) are provided for many of the methods in standard classes. A comment with each `class` declaration in Chapter 9 specifies the smallest collection of

method definitions that, together with the default declarations, provide a reasonable definition for all the class methods. If there is no such comment, then all class methods must be given to fully specify an instance.

6.3.1 The Eq Class

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

      x /= y  = not (x == y)
      x == y  = not (x /= y)
```

The `Eq` class provides equality (==) and inequality (/=) methods. All basic datatypes except for functions and `IO` are instances of this class. Instances of `Eq` can be derived for any user-defined datatype whose constituents are also instances of `Eq`.

This declaration gives default method declarations for both `/=` and ==, each being defined in terms of the other. If an instance declaration for `Eq` defines neither == nor `/=`, then both will loop. If one is defined, the default method for the other will make use of the one that is defined. If both are defined, neither default method is used.

6.3.2 The Ord Class

```
class  (Eq a) => Ord a  where
  compare              :: a -> a -> Ordering
  (<), (<=), (>=), (>) :: a -> a -> Bool
  max, min             :: a -> a -> a

  compare x y | x == y    = EQ
              | x <= y    = LT
              | otherwise = GT

  x <= y  = compare x y /= GT
  x <  y  = compare x y == LT
  x >= y  = compare x y /= LT
  x >  y  = compare x y == GT

  -- Note that (min x y, max x y) = (x,y) or (y,x)
  max x y | x <= y    =  y
          | otherwise =  x
  min x y | x <= y    =  x
          | otherwise =  y
```

The `Ord` class is used for totally ordered datatypes. All basic datatypes except for functions, `IO`, and `IOError`, are instances of this class. Instances of `Ord` can be derived for any user-defined datatype whose constituent types are in `Ord`. The declared order of the constructors in the data declaration determines the ordering in derived `Ord` instances. The `Ordering` datatype allows a single comparison to determine the precise ordering of two objects.

The default declarations allow a user to create an `Ord` instance either with a type-specific `compare` function or with type-specific == and <= functions.

6.3.3 The Read and Show Classes

```
type  ReadS a = String -> [(a,String)]
type  ShowS   = String -> String

class  Read a  where
    readsPrec :: Int -> ReadS a
```

```
    readList  :: ReadS [a]
    -- ... default decl for readList given in Prelude

class  Show a  where
    showsPrec :: Int -> a -> ShowS
    show      :: a -> String
    showList  :: [a] -> ShowS

    showsPrec _ x s   = show x ++ s
    show x            = showsPrec 0 x ""
    -- ... default decl for showList given in Prelude
```

The `Read` and `Show` classes are used to convert values to or from strings. The `Int` argument to `showsPrec` and `readsPrec` gives the operator precedence of the enclosing context (see Section 11.4).

`showsPrec` and `showList` return a `String`-to-`String` function, to allow constant-time concatenation of its results using function composition. A specialised variant, `show`, is also provided, which uses precedence context zero, and returns an ordinary `String`. The method `showList` is provided to allow the programmer to give a specialised way of showing lists of values. This is particularly useful for the `Char` type, where values of type `String` should be shown in double quotes, rather than between square brackets.

Derived instances of `Read` and `Show` replicate the style in which a constructor is declared: infix constructors and field names are used on input and output. Strings produced by `showsPrec` are usually readable by `readsPrec`.

All `Prelude` types, except function types and `IO` types, are instances of `Show` and `Read`. (If desired, a programmer can easily make functions and `IO` types into (vacuous) instances of `Show`, by providing an instance declaration.)

For convenience, the Prelude provides the following auxiliary functions:

```
reads   :: (Read a) => ReadS a
reads   =  readsPrec 0

shows   :: (Show a) => a -> ShowS
shows   =  showsPrec 0

read    :: (Read a) => String -> a
read s  =  case [x | (x,t) <- reads s, ("","") <- lex t] of
              [x] -> x
              []  -> error "PreludeText.read: no parse"
              _   -> error "PreludeText.read: ambiguous parse"
```

`shows` and `reads` use a default precedence of 0. The `read` function reads input from a string, which must be completely consumed by the input process.

The function `lex :: ReadS String`, used by `read`, is also part of the Prelude. It reads a single lexeme from the input, discarding initial white space, and returning the characters that constitute the lexeme. If the input string contains only white space, `lex` returns a single successful "lexeme" consisting of the empty string. (Thus `lex "" = [("","")]`.) If there is no legal lexeme at the beginning of the input string, `lex` fails (i.e. returns `[]`).

6.3.4 The Enum Class

```
class  Enum a  where
    succ, pred     :: a -> a
    toEnum         :: Int -> a
```

```
    fromEnum       :: a -> Int
    enumFrom       :: a -> [a]           -- [n..]
    enumFromThen   :: a -> a -> [a]      -- [n,n'..]
    enumFromTo     :: a -> a -> [a]      -- [n..m]
    enumFromThenTo :: a -> a -> a -> [a] -- [n,n'..m]

    -- Default declarations given in Prelude
```

Class `Enum` defines operations on sequentially ordered types. The functions `succ` and `pred` return the successor and predecessor, respectively, of a value. The functions `fromEnum` and `toEnum` map values from a type in `Enum` to and from `Int`. The `enumFrom`... methods are used when translating arithmetic sequences (Section 3.10).

Instances of `Enum` may be derived for any enumeration type (types whose constructors have no fields); see Chapter 11.

For any type that is an instance of class `Bounded` as well as `Enum`, the following should hold:

- The calls `succ maxBound` and `pred minBound` should result in a runtime error.
- `fromEnum` and `toEnum` should give a runtime error if the result value is not representable in the result type. For example, `toEnum 7 :: Bool` is an error.
- `enumFrom` and `enumFromThen` should be defined with an implicit bound, thus:

```
enumFrom     x   = enumFromTo     x maxBound
enumFromThen x y = enumFromThenTo x y bound
  where
    bound | fromEnum y >= fromEnum x = maxBound
          | otherwise                = minBound
```

The following `Prelude` types are instances of `Enum`:

- Enumeration types: `()`, `Bool`, and `Ordering`. The semantics of these instances is given by Chapter 11. For example, `[LT ..]` is the list `[LT,EQ,GT]`.
- `Char`: the instance is given in Chapter 9, based on the primitive functions that convert between a `Char` and an `Int`. For example, `enumFromTo 'a' 'z'` denotes the list of lowercase letters in alphabetical order.
- Numeric types: `Int`, `Integer`, `Float`, `Double`. The semantics of these instances is given next.

For all four numeric types, `succ` adds 1, and `pred` subtracts 1. The conversions `fromEnum` and `toEnum` convert between the type and `Int`. In the case of `Float` and `Double`, the digits after the decimal point may be lost. It is implementation-dependent what `fromEnum` returns when applied to a value that is too large to fit in an `Int`.

For the types `Int` and `Integer`, the enumeration functions have the following meaning:

- The sequence **`enumFrom`** e_1 is the list $[e_1, e_1 + 1, e_1 + 2, ...]$.
- The sequence **`enumFromThen`** $e_1 e_2$ is the list $[e_1, e_1 + i, e_1 + 2i, ...]$, where the increment, i, is $e_2 - e_1$. The increment may be zero or negative. If the increment is zero, all the list elements are the same.
- The sequence **`enumFromTo`** $e_1 e_3$ is the list $[e_1, e_1 + 1, e_1 + 2, ..., e_3]$. The list is empty if $e_1 > e_3$.
- The sequence **`enumFromThenTo`** $e_1 e_2 e_3$ is the list $[e_1, e_1 + i, e_1 + 2i, ..., e_3]$, where the increment, i, is $e_2 - e_1$. If the increment is positive or zero, the list terminates when the next element would be greater than e_3; the list is empty if $e_1 > e_3$. If the increment is negative, the list terminates when the next element would be less than e_3; the list is empty if $e_1 < e_3$.

For `Float` and `Double`, the semantics of the `enumFrom` family is given by the rules for `Int` above, except that the list terminates when the elements become greater than $e_3 + i \ / \ 2$ for positive increment i, or when they become less than $e_3 + i \ / \ 2$ for negative i.

For all four of these Prelude numeric types, all of the `enumFrom` family of functions are strict in all their arguments.

6.3.5 The Functor Class

```
class  Functor f  where
    fmap    :: (a -> b) -> f a -> f b
```

The `Functor` class is used for types that can be mapped over. Lists, `IO`, and `Maybe` are in this class.

Instances of `Functor` should satisfy the following laws:

$$\texttt{fmap id} = \texttt{id}$$

$$\texttt{fmap (f . g)} = \texttt{fmap f . fmap g}$$

All instances of `Functor` defined in the Prelude satisfy these laws.

6.3.6 The Monad Class

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

    m >> k  =  m >>= \_ -> k
    fail s  = error s
```

The `Monad` class defines the basic operations over a *monad*. See Chapter 7 for more information about monads.

"`do`" expressions provide a convenient syntax for writing monadic expressions (see Section 3.14). The `fail` method is invoked on pattern-match failure in a `do` expression.

In the Prelude, lists, `Maybe`, and `IO` are all instances of `Monad`. The `fail` method for lists returns the empty list `[]`, for `Maybe` returns `Nothing`, and for `IO` raises a user exception in the IO monad (see Section 7.3).

Instances of `Monad` should satisfy the following laws:

$$\texttt{return a >>= k} = \texttt{k a}$$

$$\texttt{m >> return} = \texttt{m}$$

$$\texttt{m >>= (\x -> k x >>= h)} = \texttt{(m >>= k) >>= h}$$

Instances of both `Monad` and `Functor` should additionally satisfy the law:

$$\texttt{fmap f xs} = \texttt{xs >>= return . f}$$

All instances of `Monad` defined in the Prelude satisfy these laws.

The Prelude provides the following auxiliary functions:

```
sequence  :: Monad m => [m a] -> m [a]
sequence_ :: Monad m => [m a] -> m ()
mapM      :: Monad m => (a -> m b) -> [a] -> m [b]
mapM_     :: Monad m => (a -> m b) -> [a] -> m ()
(=<<)     :: Monad m => (a -> m b) -> m a -> m b
```

6.3.7 The Bounded Class

```
class  Bounded a  where
    minBound, maxBound :: a
```

The `Bounded` class is used to name the upper and lower limits of a type. `Ord` is not a superclass of `Bounded` since types that are not totally ordered may also have upper and lower bounds. The types `Int`, `Char`, `Bool`, `()`, `Ordering`, and all tuples are instances of `Bounded`. The `Bounded` class may be derived for any enumeration type; `minBound` is the first constructor listed in the `data` declaration and `maxBound` is the last. `Bounded` may also be derived for single-constructor datatypes whose constituent types are in `Bounded`.

6.4 Numbers

Haskell provides several kinds of numbers; the numeric types and the operations upon them have been heavily influenced by Common Lisp and Scheme. Numeric function names and operators are usually overloaded, using several type classes with an inclusion relation shown in Figure 3. The class `Num` of numeric types is a subclass of `Eq`, since all numbers may be compared for equality; its subclass `Real` is also a subclass of `Ord`, since the other comparison operations apply to all but complex numbers (defined in the `Complex` library). The class `Integral` contains integers of both limited and unlimited range; the class `Fractional` contains all non-integral types; and the class `Floating` contains all floating-point types, both real and complex.

The Prelude defines only the most basic numeric types: fixed sized integers `Int`, arbitrary precision integers `Integer`, single precision floating `Float`, and double precision floating `Double`. Other numeric types such as rationals and complex numbers are defined in libraries. In particular, the type `Rational` is a ratio of two `Integern` values, as defined in the `Ratio` library.

Type	Class	Description
`Integer`	`Integral`	Arbitrary-precision integers
`Int`	`Integral`	Rational numbers
`(Integral a) => Ratio a`	`RealFrac`	Rational numbers
`Float`	`RealFloat`	Real floating-point, single precision
`Double`	`RealFloat`	Real floating-point, double precision
`(RealFloat a) => Complex a`	`Floating`	Complex floating-point

Table 2: Standard Numeric Types

The default floating point operations defined by the Haskell Prelude do not conform to current language independent arithmetic (LIA) standards. These standards require considerably more complexity in the numeric structure and have thus been relegated to a library. Some, but not all, aspects of the IEEE floating point standard have been accounted for in Prelude class `RealFloat`.

The standard numeric types are listed in Table 2. The finite-precision integer type `Int` covers at least the range $[-2^{29}, 2^{29} - 1]$. As `Int` is an instance of the `Bounded` class, `maxBound` and `minBound` can be used to determine the exact `Int` range defined by an implementation. `Float` is implementation-defined; it is desirable that this type be at least equal in range and precision to the IEEE single-precision type. Similarly, `Double` should cover IEEE double-precision. The results of exceptional conditions (such as overflow or underflow) on the fixed-precision numeric types are undefined; an implementation may choose error ($\bot$, semantically), a truncated value, or a special value such as infinity, indefinite, etc.

The standard numeric classes and other numeric functions defined in the Prelude are shown in Listing 4 Listing 5. Figure 3 shows the class dependencies and built-in types that are instances of the numeric classes.

```
class  (Eq a, Show a) => Num a  where
    (+), (-), (*)  :: a -> a -> a
    negate         :: a -> a
    abs, signum    :: a -> a
    fromInteger    :: Integer -> a

class  (Num a, Ord a) => Real a  where
    toRational ::  a -> Rational

class  (Real a, Enum a) => Integral a  where
    quot, rem, div, mod :: a -> a -> a
    quotRem, divMod     :: a -> a -> (a,a)
    toInteger           :: a -> Integer

class  (Num a) => Fractional a  where
    (/)          :: a -> a -> a
    recip        :: a -> a
    fromRational :: Rational -> a

class  (Fractional a) => Floating a  where
    pi                  :: a
    exp, log, sqrt      :: a -> a
    (**), logBase       :: a -> a -> a
    sin, cos, tan       :: a -> a
    asin, acos, atan    :: a -> a
    sinh, cosh, tanh    :: a -> a
    asinh, acosh, atanh :: a -> a
```

Listing 4: Standard Numeric Classes and Related Operations, Part 1

```
class  (Real a, Fractional a) => RealFrac a  where
    properFraction   :: (Integral b) => a -> (b,a)
    truncate, round  :: (Integral b) => a -> b
    ceiling, floor   :: (Integral b) => a -> b

class  (RealFrac a, Floating a) => RealFloat a  where
    floatRadix          :: a -> Integer
    floatDigits         :: a -> Int
    floatRange          :: a -> (Int,Int)
    decodeFloat         :: a -> (Integer,Int)
    encodeFloat         :: Integer -> Int -> a
    exponent            :: a -> Int
    significand         :: a -> a
    scaleFloat          :: Int -> a -> a
    isNaN, isInfinite, isDenormalized, isNegativeZero, isIEEE
                        :: a -> Bool
    atan2               :: a -> a -> a

gcd, lcm :: (Integral a) => a -> a-> a
(^)      :: (Num a, Integral b) => a -> b -> a
(^^)     :: (Fractional a, Integral b) => a -> b -> a

fromIntegral :: (Integral a, Num b) => a -> b
realToFrac   :: (Real a, Fractional b) => a -> b
```

Listing 5: Standard Numeric Classes and Related Operations, Part 2

6.4.1 Numeric Literals

The syntax of numeric literals is given in Section 2.5. An integer literal represents the application of the function `fromInteger` to the appropriate value of type `Integer`. Similarly, a floating literal stands for an application of `fromRational` to a value of type `Rational` (that is, `Ratio Integer`). Given the typings:

```
fromInteger  :: (Num a) => Integer -> a
fromRational :: (Fractional a) => Rational -> a
```

integer and floating literals have the typings `(Num a) => a` and `(Fractional a) => a`, respectively. Numeric literals are defined in this indirect way so that they may be interpreted as values of any appropriate numeric type. See Section 4.3.4 for a discussion of overloading ambiguity.

6.4.2 Arithmetic and Number-Theoretic Operations

The infix class methods `(+)`, `(*)`, `(-)`, and the unary function `negate` (which can also be written as a prefix minus sign; see Section 3.4) apply to all numbers. The class methods `quot`, `rem`, `div`, and `mod` apply only to integral numbers, while the class method `(/)` applies only to fractional ones. The `quot`, `rem`, `div`, and `mod` class methods satisfy these laws if `y` is non-zero:

```
(x `quot` y)*y + (x `rem` y) == x
(x `div`  y)*y + (x `mod` y) == x
```

`quot` is integer division truncated toward zero, while the result of `div` is truncated toward negative infinity. The `quotRem` class method takes a dividend and a divisor as arguments and returns a (quotient, remainder) pair; `divMod` is defined similarly:

```
quotRem x y  =  (x `quot` y, x `rem` y)
divMod  x y  =  (x `div`  y, x `mod` y)
```

Also available on integral numbers are the even and odd predicates:

```
even x =  x `rem` 2 == 0
odd    =  not . even
```

Finally, there are the greatest common divisor and least common multiple functions. `gcd` x y is the greatest (positive) integer that divides both x and y; for example `gcd (-3) 6 = 3`, `gcd (-3) (-6) = 3`, `gcd 0 4 = 4`. `gcd 0 0` raises a runtime error.

`lcm` x y is the smallest positive integer that both x and y divide.

6.4.3 Exponentiation and Logarithms

The one-argument exponential function `exp` and the logarithm function `log` act on floating-point numbers and use base e. `logBase` a x returns the logarithm of x in base a. `sqrt` returns the principal square root of a floating-point number. There are three two-argument exponentiation operations: `(^)` raises any number to a nonnegative integer power, `(^^)` raises a fractional number to any integer power, and `(**)` takes two floating-point arguments. The value of `x ^ 0` or `x ^^ 0`} is `1` for any x, including zero; `0**y` is `1` if y is `0`, and `0` otherwise.

6.4.4 Magnitude and Sign

A number has a *magnitude* and a *sign*. The functions `abs` and `signum` apply to any number and satisfy the law:

```
abs x * signum x == x
```

For real numbers, these functions are defined by:

```
abs x    | x >= 0  = x
         | x <  0  = -x

signum x | x >  0  = 1
         | x == 0  = 0
         | x <  0  = -1
```

6.4.5 Trigonometric Functions

Class `Floating` provides the circular and hyperbolic sine, cosine, and tangent functions and their inverses. Default implementations of `tan`, `tanh`, `logBase`, **, and `sqrt` are provided, but implementors are free to provide more accurate implementations.

Class `RealFloat` provides a version of arctangent taking two real floating-point arguments. For real floating x and y, `atan2` y x computes the angle (from the positive x-axis) of the vector from the origin to the point (x, y). `atan2` y x returns a value in the range `[-pi,pi]`. It follows the Common Lisp semantics for the origin when signed zeroes are supported. `atan2` y `1`, with y in a type that is `RealFloat`, should return the same value as `atan` y. A default definition of `atan2` is provided, but implementors can provide a more accurate implementation.

The precise definition of the above functions is as in Common Lisp, which in turn follows Penfield's proposal for APL [8]. See these references for discussions of branch cuts, discontinuities, and implementation.

6.4.6 Coercions and Component Extraction

The `ceiling`, `floor`, `truncate`, and `round` functions each take a real fractional argument and return an integral result. `ceiling` x returns the least integer not less than x, and `floor` x, the greatest integer not greater than x. `truncate` x yields the integer nearest x between 0 and x, inclusive. `round` x returns the nearest integer to x, the even integer if x is equidistant between two integers.

The function `properFraction` takes a real fractional number x and returns a pair (n, f) such that $x = n + f$, and: n is an integral number with the same sign as x; and f is a fraction f with the same type and sign as x, and with absolute value less than 1. The `ceiling`, `floor`, `truncate`, and `round` functions can be defined in terms of `properFraction`.

Two functions convert numbers to type `Rational`: `toRational` returns the rational equivalent of its real argument with full precision; `approxRational` takes two real fractional arguments x and ε and returns the simplest rational number within ε of x, where a rational $p \ / \ q$ in reduced form is *simpler* than another $p' \ / \ q'$ if $|p| \leq |p'|$ and $q \leq q'$. Every real interval contains a unique simplest rational; in particular, note that $0 \ / \ 1$ is the simplest rational of all.

The class methods of class `RealFloat` allow efficient, machine-independent access to the components of a floating-point number. The functions `floatRadix`, `floatDigits`, and `floatRange` give the parameters of a floating-point type: the radix of the representation, the number of digits of this radix in the significand, and the lowest and highest values the exponent may assume, respectively. The function `decodeFloat` applied to a real floating-point number returns the significand expressed as an `Integer` and an appropriately scaled exponent (an `Int`). If `decodeFloat` x yields (m, n), then `x` is equal in value to mb^n, where b is the floating-point radix, and furthermore, either m and n are both zero or else $b^{d-1} \leq |m| < b^d$, where d is the value of `floatDigits x`. `encodeFloat` performs the inverse of this transformation. The functions `significand` and `exponent` together provide the same information as `decodeFloat`, but rather than an `Integer`, `significand` x yields a value of the same type as `x`, scaled to lie in the open interval $(-1, 1)$. `exponent 0` is zero. `scaleFloat` multiplies a floating-point number by an integer power of the radix.

The functions `isNaN`, `isInfinite`, `isDenormalized`, `isNegativeZero`, and `isIEEE` all support numbers represented using the IEEE standard. For non-IEEE floating point numbers, these may all return false.

Also available are the following coercion functions:

```
fromIntegral :: (Integral a, Num b)    => a -> b
realToFrac   :: (Real a, Fractional b) => a -> b
```

Chapter 7

Basic Input/Output

The I/O system in Haskell is purely functional, yet has all of the expressive power found in conventional programming languages. To achieve this, Haskell uses a *monad* to integrate I/O operations into a purely functional context.

The I/O monad used by Haskell mediates between the *values* natural to a functional language and the *actions* that characterize I/O operations and imperative programming in general. The order of evaluation of expressions in Haskell is constrained only by data dependencies; an implementation has a great deal of freedom in choosing this order. Actions, however, must be ordered in a well-defined manner for program execution – and I/O in particular – to be meaningful. Haskell's I/O monad provides the user with a way to specify the sequential chaining of actions, and an implementation is obliged to preserve this order.

The term *monad* comes from a branch of mathematics known as *category theory*. From the perspective of a Haskell programmer, however, it is best to think of a monad as an *abstract datatype*. In the case of the I/O monad, the abstract values are the *actions* mentioned above. Some operations are primitive actions, corresponding to conventional I/O operations. Special operations (methods in the class `Monad`, see Section 6.3.6) sequentially compose actions, corresponding to sequencing operators (such as the semicolon) in imperative languages.

7.1 Standard I/O Functions

Although Haskell provides fairly sophisticated I/O facilities, as defined in the `IO` library, it is possible to write many Haskell programs using only the few simple functions that are exported from the Prelude, and which are described in this section.

All I/O functions defined here are character oriented. The treatment of the newline character will vary on different systems. For example, two characters of input, return and linefeed, may read as a single newline character. These functions cannot be used portably for binary I/O.

In the following, recall that `String` is a synonym for `[Char]` (Section 6.1.2).

Output Functions These functions write to the standard output device (this is normally the user's terminal).

```
putChar  :: Char -> IO ()
putStr   :: String -> IO ()
putStrLn :: String -> IO ()  -- adds a newline
print    :: Show a => a -> IO ()
```

The `print` function outputs a value of any printable type to the standard output device. Printable types are those that are instances of class `Show`; `print` converts values to strings for output using the `show` operation and adds a newline.

For example, a program to print the first 20 integers and their powers of 2 could be written as:

```
main = print ([(n, 2^n) | n <- [0..19]])
```

Input Functions These functions read input from the standard input device (normally the user's terminal).

```
getChar     :: IO Char
getLine     :: IO String
getContents :: IO String
interact    :: (String -> String) -> IO ()
readIO      :: Read a => String -> IO a
readLn      :: Read a => IO a
```

The `getChar` operation raises an exception (Section 7.3) on end-of-file; a predicate `isEOFError` that identifies this exception is defined in the `IO` library. The `getLine` operation raises an exception under the same circumstances as `hGetLine`, defined the `IO` library.

The `getContents` operation returns all user input as a single string, which is read lazily as it is needed. The `interact` function takes a function of type `String->String` as its argument. The entire input from the standard input device is passed to this function as its argument, and the resulting string is output on the standard output device.

Typically, the `read` operation from class `Read` is used to convert the string to a value. The `readIO` function is similar to `read` except that it signals parse failure to the I/O monad instead of terminating the program. The `readLn` function combines `getLine` and `readIO`.

The following program simply removes all non-ASCII characters from its standard input and echoes the result on its standard output. (The `isAscii` function is defined in a library.)

```
main = interact (filter isAscii)
```

Files These functions operate on files of characters. Files are named by strings using some implementation-specific method to resolve strings as file names.

The `writeFile` and `appendFile` functions write or append the string, their second argument, to the file, their first argument. The `readFile` function reads a file and returns the contents of the file as a string. The file is read lazily, on demand, as with `getContents`.

```
type FilePath =  String

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

Note that `writeFile` and `appendFile` write a literal string to a file. To write a value of any printable type, as with `print`, use the `show` function to convert the value to a string first.

```
main = appendFile "squares" (show [(x,x*x) | x <- [0,0.1..2]])
```

7.2 Sequencing I/O Functions

The type constructor `IO` is an instance of the `Monad` class. The two monadic binding functions, methods in the `Monad` class, are used to compose a series of I/O operations. The `>>` function is used where the result of the first operation is uninteresting, for example when it is `()`. The `>>=` operation passes the result of the first operation as an argument to the second operation.

```
(>>=) :: IO a -> (a -> IO b) -> IO b
(>>)  :: IO a -> IO b        -> IO b
```

For example,

```
main = readFile "input-file"                        >>= \ s ->
       writeFile "output-file" (filter isAscii s)   >>
       putStr "Filtering successful\n"
```

is similar to the previous example using `interact`, but takes its input from `"input-file"` and writes its output to `"output-file"`. A message is printed on the standard output before the program completes.

The `do` notation allows programming in a more imperative syntactic style. A slightly more elaborate version of the previous example would be:

```
main = do
        putStr "Input file: "
        ifile <- getLine
        putStr "Output file: "
        ofile <- getLine
        s <- readFile ifile
        writeFile ofile (filter isAscii s)
        putStr "Filtering successful\n"
```

The `return` function is used to define the result of an I/O operation. For example, `getLine` is defined in terms of `getChar`, using `return` to define the result:

```
getLine :: IO String
getLine = do c <- getChar
             if c == '\n' then return ""
                          else do s <- getLine
                                  return (c:s)
```

7.3 Exception Handling in the I/O Monad

The I/O monad includes a simple exception handling system. Any I/O operation may raise an exception instead of returning a result.

Exceptions in the I/O monad are represented by values of type `IOError`. This is an abstract type: its constructors are hidden from the user. The `IO` library defines functions that construct and examine `IOError` values. The only Prelude function that creates an `IOError` value is `userError`. User error values include a string describing the error.

```
userError :: String -> IOError
```

Exceptions are raised and caught using the following functions:

```
ioError :: IOError -> IO a
catch   :: IO a    -> (IOError -> IO a) -> IO a
```

The `ioError` function raises an exception; the `catch` function establishes a handler that receives any exception raised in the action protected by `catch`. An exception is caught by the most recent handler established by `catch`. These handlers are not selective: all exceptions are caught. Exception propagation must be explicitly provided in a handler by re-raising any unwanted exceptions. For example, in

```
f = catch g (\e -> if IO.isEOFError e then return [] else ioError e)
```

the function `f` returns `[]` when an end-of-file exception occurs in `g`; otherwise, the exception is propagated to the next outer handler. The `isEOFError` function is part of `IO` library.

When an exception propagates outside the main program, the Haskell system prints the associated `IOError` value and exits the program.

The `fail` method of the `IO` instance of the `Monad` class (Section 6.3.6) raises a `userError`, thus:

```
instance Monad IO where
  ...bindings for return, (>>=), (>>)

  fail s = ioError (userError s)
```

The exceptions raised by the I/O functions in the Prelude are defined in the module `System.IO.Error`.

Chapter 8

Foreign Function Interface

The Foreign Function Interface (FFI) has two purposes: it enables (1) to describe in Haskell the interface to foreign language functionality and (2) to use from foreign code Haskell routines. More generally, its aim is to support the implementation of programs in a mixture of Haskell and other languages such that the source code is portable across different implementations of Haskell and non-Haskell systems as well as independent of the architecture and operating system.

8.1 Foreign Languages

The Haskell FFI currently only specifies the interaction between Haskell code and foreign code that follows the C calling convention. However, the design of the FFI is such that it enables the modular extension of the present definition to include the calling conventions of other programming languages, such as C++ and Java. A precise definition of the support for those languages is expected to be included in later versions of the language. The second major omission is the definition of the interaction with multithreading in the foreign language and, in particular, the treatment of thread-local state, and so these details are currently implementation-defined.

The core of the present specification is independent of the foreign language that is used in conjunction with Haskell. However, there are two areas where FFI specifications must become language specific: (1) the specification of external names and (2) the marshalling of the basic types of a foreign language. As an example of the former, consider that in C [9] a simple identifier is sufficient to identify an object, while Java [10], in general, requires a qualified name in conjunction with argument and result types to resolve possible overloading. Regarding the second point, consider that many languages do not specify the exact representation of some basic types. For example the type `int` in C may be 16, 32, or 64 bits wide. Similarly, Haskell guarantees only that `Int` covers at least the range $[-2^{29}, 2^{29} - 1]$ (Section 6.4). As a consequence, to reliably represent values of C's `int` in Haskell, we have to introduce a new type `CInt`, which is guaranteed to match the representation of `int`.

The specification of external names, dependent on a calling convention, is described in Section 8.5, whereas the marshalling of the basic types in dependence on a foreign language is described in Section 8.6

8.2 Contexts

For a given Haskell system, we define the *Haskell context* to be the execution context of the abstract machine on which the Haskell system is based. This includes the heap, stacks, and the registers of the

abstract machine and their mapping onto a concrete architecture. We call any other execution context an *external context*. Generally, we cannot assume any compatibility between the data formats and calling conventions between the Haskell context and a given external context, except where Haskell explicitly prescribes a specific data format.

The principal goal of a foreign function interface is to provide a programmable interface between the Haskell context and external contexts. As a result Haskell threads can access data in external contexts and invoke functions that are executed in an external context as well as vice versa. In the rest of this definition, external contexts are usually identified by a calling convention.

8.2.1 Cross Language Type Consistency

Given that many external languages support static types, the question arises whether the consistency of Haskell types with the types of the external language can be enforced for foreign functions. Unfortunately, this is, in general, not possible without a significant investment on the part of the implementor of the Haskell system (i.e., without implementing a dedicated type checker). For example, in the case of the C calling convention, the only other approach would be to generate a C prototype from the Haskell type and leave it to the C compiler to match this prototype with the prototype that is specified in a C header file for the imported function. However, the Haskell type is lacking some information that would be required to pursue this route. In particular, the Haskell type does not contain any information as to when `const` modifiers have to be emitted.

As a consequence, this definition does not require the Haskell system to check consistency with foreign types. Nevertheless, Haskell systems are encouraged to provide any cross language consistency checks that can be implemented with reasonable effort.

8.3 Lexical Structure

The FFI reserves a single keyword `foreign`, and a set of special identifiers. The latter have a special meaning only within foreign declarations, but may be used as ordinary identifiers elsewhere.

The special identifiers `ccall`, `cplusplus`, `dotnet`, `jvm`, and `stdcall` are defined to denote calling conventions. However, a concrete implementation of the FFI is free to support additional, system-specific calling conventions whose name is not explicitly listed here.

To refer to objects of an external C context, we introduce the following phrases:

chname	→	{*chchar*} `.h`	(C header filename)
cid	→	*letter* {*letter* | *ascDigit*}	(C identifier)
chchar	→	*letter* | $ascSymbol_{\langle \& \rangle}$	
letter	→	*ascSmall* | *ascLarge* | _	

The range of lexemes that are admissible for *chname* is a subset of those permitted as arguments to the `#include` directive in C. In particular, a file name *chname* must end in the suffix `.h`. The lexemes produced by *cid* coincide with those allowed as C identifiers, as specified in [9].

8.4 Foreign Declarations

The syntax of foreign declarations is as follows:

topdecl	→	`foreign` *fdecl*	
fdecl	→	`import` *callconv* [*safety*] *impent var* :: *ftype*	(define variable)
	|	`export` *callconv expent var* :: *ftype*	(expose variable)“
callconv	→	`ccall` | `stdcall` | `cplusplus`	(calling convention)
	|	`jvm` | `dotnet`	
	|	**system-specific calling conventions**	
impent	→	[*string*]	
expent	→	[*string*]	
safety	→	`unsafe` | `safe`	

There are two flavours of foreign declarations: import and export declarations. An import declaration makes an *external entity,* i.e., a function or memory location defined in an external context, available in the Haskell context. Conversely, an export declaration defines a function of the Haskell context as an external entity in an external context. Consequently, the two types of declarations differ in that an import declaration defines a new variable, whereas an export declaration uses a variable that is already defined in the Haskell module.

The external context that contains the external entity is determined by the calling convention given in the foreign declaration. Consequently, the exact form of the specification of the external entity is dependent on both the calling convention and on whether it appears in an import declaration (as *impent*) or in an export declaration (as *expent*). To provide syntactic uniformity in the presence of different calling conventions, it is guaranteed that the description of an external entity lexically appears as a Haskell string lexeme. The only exception is where this string would be the empty string (i.e., be of the form `""`); in this case, the string may be omitted in its entirety.

8.4.1 Calling Conventions

The binary interface to an external entity on a given architecture is determined by a calling convention. It often depends on the programming language in which the external entity is implemented, but usually is more dependent on the system for which the external entity has been compiled.

As an example of how the calling convention is dominated by the system rather than the programming language, consider that an entity compiled to byte code for the Java Virtual Machine (JVM) [11] needs to be invoked by the rules of the JVM rather than that of the source language in which it is implemented (the entity might be implemented in Oberon, for example).

Any implementation of the Haskell FFI must at least implement the C calling convention denoted by `ccall`. All other calling conventions are optional. Generally, the set of calling conventions is open, i.e., individual implementations may elect to support additional calling conventions. In addition to `ccall`, Table 3 specifies a range of identifiers for common calling conventions.

Identifier	Represented calling convention
`ccall`	Calling convention of the standard C compiler on a system
`cplusplus`	Calling convention of the standard C++ compiler on a system
`dotnet`	Calling convention of the `.NET` platform
`jvm`	Calling convention of the Java Virtual Machine
`stdcall`	Calling convention of the Win32 API (matches Pascal conventions)

Table 3: Calling conventions

Implementations need not implement all of these conventions, but if any is implemented, it must use the listed name. For any other calling convention, implementations are free to choose a suitable name.

Only the semantics of the calling conventions `ccall` and `stdcall` are defined herein; more calling conventions may be added in future versions of Haskell.

It should be noted that the code generated by a Haskell system to implement a particular calling convention may vary widely with the target code of that system. For example, the calling convention `jvm` will be trivial to implement for a Haskell compiler generating Java code, whereas for a Haskell compiler generating C code, the Java Native Interface (JNI) [12] has to be targeted.

8.4.2 Foreign Types

The following types constitute the set of *basic foreign types*:

- `Char`, `Int`, `Double`, `Float`, and `Bool` as exported by the Haskell `Prelude` as well as
- `Int8`, `Int16`, `Int32`, `Int64`, `Word8`, `Word16`, `Word32`, `Word64`, `Ptr a`, `FunPtr a`, and `StablePtr a`, for any type `a`, as exported by `Foreign` (Section ref{module:Foreign}).

A Haskell system that implements the FFI needs to be able to pass these types between the Haskell and the external context as function arguments and results.

Foreign types are produced according to the following grammar:

$$
\begin{array}{lcl}
\mathit{ftype} & \rightarrow & \mathit{frtype} \\
 & | & \mathit{fatype} \texttt{ -> } \mathit{ftype} \\
\mathit{frtype} & \rightarrow & \mathit{fatype} \\
 & | & () \\
\mathit{fatype} & \rightarrow & \mathit{qtycon}\ \mathit{atype}_1 ... \mathit{atype}_k \quad (k \geq 0)
\end{array}
$$

A foreign type is the Haskell type of an external entity. Only a subset of Haskell's types are permissible as foreign types, as only a restricted set of types can be canonically transferred between the Haskell context and an external context. A foreign type has the form

$$at_1 \rightarrow ... \rightarrow at_n \rightarrow rt$$

where $n \geq 0$. It implies that the arity of the external entity is n.

External functions are strict in all arguments.

Marshallable foreign types. The argument types at_i produced by *fatype* must be *marshallable foreign types;* that is, either

- a basic foreign type,
- a type synonym that expands to a marshallable foreign type,

- a type $T\ t'_1 ... t'_n$ where T is defined by a `newtype` declaration

$$\texttt{newtype}\ T\ a_1 ... a_n = N\ t$$

  and
  - the constructor N is visible where T is used,
  - $t[t'_1\ /\ a_1 ... t'_n\ /\ a_n]$ is a marshallable foreign type

Consequently, in order for a type defined by `newtype` to be used in a `foreign` declaration outside of the module that defines it, the type must not be exported abstractly. The module `Foreign.C.Types` that defines the Haskell equivalents for C types follows this convention.

Marshallable foreign result types. The result type rt produced by *frtype* must be a *marshallable foreign result type;* that is, either

- the type `()`,
- a type matching `Prelude.IO` t, where t is a marshallable foreign type or `()`,
- a basic foreign type,
- a type synonym that expands to marshallable foreign result type,
- a type $T t'_1 ... t'_n$ where T is defined by a `newtype` declaration

$$\texttt{newtype}\ T\ a_1 ... a_n = N\ t$$

  and
  - the constructor N is visible where T is used,
  - $t[t'_1\ /\ a_1 ... t'_n\ /\ a_n]$ is a marshallable foreign result type

8.4.3 Import Declarations

Generally, an import declaration has the form

$$\texttt{foreign import}\ c\ e\ v :: t$$

which declares the variable v of type t to be defined externally. Moreover, it specifies that v is evaluated by executing the external entity identified by the string e using calling convention c. The precise form of e depends on the calling convention and is detailed in Section 8.5 If a variable v is defined by an import declaration, no other top-level declaration for v is allowed in the same module. For example, the declaration

```
foreign import ccall "string.h strlen"
   cstrlen :: Ptr CChar -> IO CSize
```

introduces the function mbox{tt cstrlen}, which invokes the external function `strlen` using the standard C calling convention. Some external entities can be imported as pure functions; for example,

```
foreign import ccall "math.h sin"
   sin :: CDouble -> CDouble.
```

Such a declaration asserts that the external entity is a true function; i.e., when applied to the same argument values, it always produces the same result.

Whether a particular form of external entity places a constraint on the Haskell type with which it can be imported is defined in Section 8.5 Although, some forms of external entities restrict the set of Haskell types that are permissible, the system can generally not guarantee the consistency between the Haskell type given in an import declaration and the argument and result types of the external entity. It is the responsibility of the programmer to ensure this consistency.

Optionally, an import declaration can specify, after the calling convention, the safety level that should be used when invoking an external entity. A `safe` call is less efficient, but guarantees to leave the Haskell system in a state that allows callbacks from the external code. In contrast, an `unsafe` call, while carrying less overhead, must not trigger a callback into the Haskell system. If it does, the system behaviour is undefined. The default for an invocation is to be `safe`. Note that a callback into the Haskell system implies that a garbage collection might be triggered after an external entity was called, but before this call returns. Consequently, objects other than stable pointers (cf.
Section ref{module:Foreign.StablePtr}) may be moved or garbage collected by the storage manager.

8.4.4 Export Declarations

The general form of export declarations is

$$\texttt{foreign export}\ c\ e\ v :: t$$

Such a declaration enables external access to v, which may be a value, field name, or class method that is declared on the top-level of the same module or imported. Moreover, the Haskell system defines the external entity described by the string e, which may be used by external code using the calling convention c; an external invocation of the external entity e is translated into evaluation of v. The type t must be an instance of the type of v. For example, we may have

```
foreign export ccall "addInt"   (+) :: Int   -> Int   -> Int
foreign export ccall "addFloat" (+) :: Float -> Float -> Float
```

If an evaluation triggered by an external invocation of an exported Haskell value returns with an exception, the system behaviour is undefined. Thus, Haskell exceptions have to be caught within Haskell and explicitly marshalled to the foreign code.

8.5 Specification of External Entities

Each foreign declaration has to specify the external entity that is accessed or provided by that declaration. The syntax and semantics of the notation that is required to uniquely determine an external entity depends heavily on the calling convention by which this entity is accessed. For example, for the calling convention `ccall`, a global label is sufficient. However, to uniquely identify a method in the calling convention `jvm`, type information has to be provided. For the latter, there is a choice between the Java source-level syntax of types and the syntax expected by JNI—but, clearly, the syntax of the specification of an external entity depends on the calling convention and may be non-trivial.

Consequently, the FFI does not fix a general syntax for denoting external entities, but requires both *impent* and *expent* to take the form of a Haskell *string* literal. The formation rules for the values of these strings depend on the calling convention and a Haskell system implementing a particular calling convention will have to parse these strings in accordance with the calling convention.

Defining *impent* and *expent* to take the form of a *string* implies that all information that is needed to statically analyse the Haskell program is separated from the information needed to generate the code interacting with the foreign language. This is, in particular, helpful for tools processing Haskell source code. When ignoring the entity information provided by *impent* or *expent*, foreign import and export declarations are still sufficient to infer identifier definition and use information as well as type information.

For more complex calling conventions, there is a choice between the user-level syntax for identifying entities (e.g., Java or C++) and the system-level syntax (e.g., the type syntax of JNI or mangled C++,

respectively). If such a choice exists, the user-level syntax is preferred. Not only because it is more user friendly, but also because the system-level syntax may not be entirely independent of the particular implementation of the foreign language.

The following defines the syntax for specifying external entities and their semantics for the calling conventions `ccall` and `stdcall`. Other calling conventions from Table 3 are expected to be defined in future versions of Haskell.

8.5.1 Standard C Calls

The following defines the structure of external entities for foreign declarations under the `ccall` calling convention for both import and export declarations separately. Afterwards additional constraints on the type of foreign functions are defined.

The FFI covers only access to C functions and global variables. There are no mechanisms to access other entities of C programs. In particular, there is no support for accessing pre-processor symbols from Haskell, which includes `#definedd` constants. Access from Haskell to such entities is the domain of language-specific tools, which provide added convenience over the plain FFI as defined here.

Import Declarations For import declarations, the syntax for the specification of external entities under the `ccall` calling convention is as follows:

$$
\begin{aligned}
\mathit{impent} \rightarrow\ & \texttt{"}\ [\texttt{static}][\mathit{chname}][\&][\mathit{cid}]\ \texttt{"} && \text{(static function or address)} \\
\mid\ & \texttt{"}\ \mathit{dynamic}\ \texttt{"} && \text{(stub factory importing addresses)} \\
\mid\ & \texttt{"}\ \mathit{wrapper}\ \texttt{"} && \text{(stub factory exporting thunks)}
\end{aligned}
$$

The first alternative either imports a static function *cid* or, if & precedes the identifier, a static address. If *cid* is omitted, it defaults to the name of the imported Haskell variable. The optional filename *chname* specifies a C header file, where the intended meaning is that the header file declares the C entity identified by *cid*. In particular, when the Haskell system compiles Haskell to C code, the directive

$$\#\texttt{include}\ \texttt{"}\ \mathit{chname}\ \texttt{"}$$

needs to be placed into any generated C file that refers to the foreign entity before the first occurrence of that entity in the generated C file.

The second and third alternative, identified by the keywords `dynamic` and `wrapper`, respectively, import stub functions that have to be generated by the Haskell system. In the case of `dynamic`, the stub converts C function pointers into Haskell functions; and conversely, in the case of `wrapper`, the stub converts Haskell thunks to C function pointers. If neither of the specifiers `static`, `dynamic`, or `wrapper` is given, `static` is assumed. The specifier `static` is nevertheless needed to import C routines that are named `dynamic` or `wrapper`.

It should be noted that a static foreign declaration that does not import an address (i.e., where & is not used in the specification of the external entity) always refers to a C function, even if the Haskell type is non-functional. For example,

```
foreign import ccall foo :: CInt
```

refers to a pure C function `foo` with no arguments that returns an integer value. Similarly, if the type is `IO CInt`, the declaration refers to an impure nullary function. If a Haskell program needs to access a C variable `bar` of integer type,

```
foreign import ccall "&" bar :: Ptr CInt
```

must be used to obtain a pointer referring to the variable. The variable can be read and updated using the routines provided by the module `Foreign.Storable` (cf.
Section ref{module:Foreign.Storable}).

Export Declarations External entities in *ccall* export declarations are of the form

$$expent \rightarrow \texttt{"}\ cid\ \texttt{"}$$

The optional C identifier *cid* defines the external name by which the exported Haskell variable is accessible in C. If it is omitted, the external name defaults to the name of the exported Haskell variable.

Constraints on Foreign Function Types In the case of import declaration, there are, depending on the kind of import declaration, constraints regarding the admissible Haskell type that the variable defined in the import may have. These constraints are specified in the following.

Static Functions A static function can be of any foreign type; in particular, the result type may or may not be in the IO monad. If a function that is not pure is not imported in the IO monad, the system behaviour is undefined. Generally, no check for consistency with the C type of the imported label is performed.

As an example, consider

```
foreign import ccall "static stdlib.h"
   system :: Ptr CChar -> IO CInt
```

This declaration imports the `system()` function whose prototype is available from `stdlib.h`.

Static addresses The type of an imported address is constrained to be of the form `Ptr` a or `FunPtr` a, where a can be any type.

As an example, consider

```
foreign import ccall "errno.h &errno" errno :: Ptr CInt
```

It imports the address of the variable `errno`, which is of the C type `int`.

Dynamic import The type of a *dynamic* stub has to be of the form (`FunPtr` ft) $\rightarrow$ ft, where ft may be any foreign type.

As an example, consider

```
foreign import ccall "dynamic"
  mkFun :: FunPtr (CInt -> IO ()) -> (CInt -> IO ())
```

The stub factory `mkFun` converts any pointer to a C function that gets an integer value as its only argument and does not have a return value into a corresponding Haskell function.

Dynamic wrapper The type of a *wrapper* stub has to be of the form $ft \rightarrow$ `IO` (`FunPtr` ft), where ft may be any foreign type.

As an example, consider

```
foreign import ccall "wrapper"
  mkCallback :: IO () -> IO (FunPtr (IO ()))
```

The stub factory `mkCallback` turns any Haskell computation of type `IO ()` into a C function pointer that can be passed to C routines, which can call back into the Haskell context by invoking the referenced function.

Specification of Header Files A C header specified in an import declaration is always included by `#include` " *chname* ". There is no explicit support for `#include` < *chname* > style inclu-

sion. The ISO C99 [13] standard guarantees that any search path that would be used for a `#include` < *chname* > is also used for `#include` " *chname* " and it is guaranteed that these paths are searched after all paths that are unique to `#include` " *chname* ". Furthermore, we require that *chname* ends in `.h` to make parsing of the specification of external entities unambiguous.

The specification of include files has been kept to a minimum on purpose. Libraries often require a multitude of include directives, some of which may be system-dependent. Any design that attempts to cover all possible configurations would introduce significant complexity. Moreover, in the current design, a custom include file can be specified that uses the standard C preprocessor features to include all relevant headers.

Header files have no impact on the semantics of a foreign call, and whether an implementation uses the header file or not is implementation-defined. However, as some implementations may require a header file that supplies a correct prototype for external functions in order to generate correct code, portable FFI code must include suitable header files.

C Argument Promotion The argument passing conventions of C are dependent on whether a function prototype for the called functions is in scope at a call site. In particular, if no function prototype is in scope, *default argument promotion* is applied to integral and floating types. In general, it cannot be expected from a Haskell system that it is aware of whether a given C function was compiled with or without a function prototype being in scope. For the sake of portability, we thus require that a Haskell system generally implements calls to C functions as well as C stubs for Haskell functions as if a function prototype for the called function is in scope.

This convention implies that the onus for ensuring the match between C and Haskell code is placed on the FFI user. In particular, when a C function that was compiled without a prototype is called from Haskell, the Haskell signature at the corresponding `foreign import` declaration must use the types *after* argument promotion. For example, consider the following C function definition, which lacks a prototype:

```
void foo (a)
float a;
{
  ...
}
```

The lack of a prototype implies that a C compiler will apply default argument promotion to the parameter `a`, and thus, `foo` will expect to receive a value of type `double`, *not* `float`. Hence, the correct `foreign import` declaration is

```
foreign import ccall foo :: Double -> IO ()
```

In contrast, a C function compiled with the prototype

```
void foo (float a);
```

requires

```
foreign import ccall foo :: Float -> IO ()
```

A similar situation arises in the case of `foreign export` declarations that use types that would be altered under the C default argument promotion rules. When calling such Haskell functions from C, a function prototype matching the signature provided in the `foreign export` declaration must be in scope; otherwise, the C compiler will erroneously apply the promotion rules to all function arguments.

Note that for a C function defined to accept a variable number of arguments, all arguments beyond the explicitly typed arguments suffer argument promotion. However, because C permits the calling

convention to be different for such functions, a Haskell system will, in general, not be able to make use of variable argument functions. Hence, their use is deprecated in portable code.

8.5.2 Win32 API Calls

The specification of external entities under the `stdcall` calling convention is identical to that for standard C calls. The two calling conventions only differ in the generated code.

8.6 Marshalling

In addition to the language extension discussed in previous sections, the FFI includes a set of standard libraries, which ease portable use of foreign functions as well as marshalling of compound structures. Generally, the marshalling of Haskell structures into a foreign representation and vice versa can be implemented in either Haskell or the foreign language. At least where the foreign language is at a significantly lower level, e.g. C, there are good reasons for doing the marshalling in Haskell:

- Haskell's lazy evaluation strategy would require any foreign code that attempts to access Haskell structures to force the evaluation of these structures before accessing them. This would lead to complicated code in the foreign language, but does not need any extra consideration when coding the marshalling in Haskell.
- Despite the fact that marshalling code in Haskell tends to look like C in Haskell syntax, the strong type system still catches many errors that would otherwise lead to difficult-to-debug runtime faults.
- Direct access to Haskell heap structures from a language like C—especially, when marshalling from C to Haskell, i.e., when Haskell structures are created—carries the risk of corrupting the heap, which usually leads to faults that are very hard to debug.

Consequently, the Haskell FFI emphasises Haskell-side marshalling.

The interface to the marshalling libraries is provided by the module `Foreign` (Chapter ref{module:Foreign}) plus a language-dependent module per supported language. In particular, the standard requires the availability of the module `Foreign.C` (Chapter ref{module:Foreign.C}), which simplifies portable interfacing with external C code. Language-dependent modules, such as `Foreign.C`, generally provide Haskell types representing the basic types of the foreign language using a representation that is compatible with the foreign types as implemented by the default implementation of the foreign language on the present architecture. This is especially important for languages where the standard leaves some aspects of the implementation of basic types open. For example, in C, the size of the various integral types is not fixed. Thus, to represent C interfaces faithfully in Haskell, for each integral type in C, we need to have an integral type in Haskell that is guaranteed to have the same size as the corresponding C type.

8.7 The External C Interface

Every Haskell system that implements the FFI needs to provide a C header file named `HsFFI.h` that defines the C symbols listed in Table 4 and Table 5. Table 4 lists symbols that represent types together with the Haskell type that they represent and any constraints that are placed on the concrete C types that implement these symbols. When a C type `HsT` represents a Haskell type `T`, the occurrence of `T` in a foreign function declaration should be matched by `HsT` in the corresponding C function prototype.

Indeed, where the Haskell system translates Haskell to C code that invokes `foreign imported` C routines, such prototypes need to be provided and included via the header that can be specified in external entity strings for foreign C functions (cf. Section 8.5.1); otherwise, the system behaviour is undefined. It is guaranteed that the Haskell value `nullPtr` is mapped to `(HsPtr) NULL` in C and `nullFunPtr` is mapped to `(HsFunPtr) NULL` and vice versa.

C symbol	Haskell symbol	Constraint on concrete C type
`HsChar`	`Char`	integral type
`HsInt`	`Int`	signed integral type, $\geq$ 30 bit
`HsInt8`	`Int8`	signed integral type, 8 bit; `int8_t` if available
`HsInt16`	`Int16`	signed integral type, 16 bit; `int16_t` if available
`HsInt32`	`Int32`	signed integral type, 32 bit; `int32_t` if available
`HsInt64`	`Int64`	signed integral type, 64 bit; `int64_t` if available
`HsWord8`	`Word8`	unsigned integral type, 8 bit; `uint8_t` if available
`HsWord16`	`Word16`	unsigned integral type, 16 bit; `uint16_t` if available
`HsWord32`	`Word32`	unsigned integral type, 32 bit; `uint32_t` if available
`HsWord64`	`Word64`	unsigned integral type, 64 bit; `uint64_t` if available
`HsFloat`	`Float`	floating point type
`HsDouble`	`Double`	floating point type
`HsBool`	`Bool`	`int`
`HsPtr`	`Ptr a`	`(void *)`
`HsFunPtr`	`FunPtr a`	`(void (*) (void))`
`HsStablePtr`	`StablePtr a`	`(void *)`

Table 4: C Interface to Basic Haskell Types

CPP symbol	Haskell value	Description
HS_CHAR_MIN	minBound :: Char	
HS_CHAR_MAX	maxBound :: Char	
HS_INT_MIN	minBound :: Int	
HS_INT_MAX	maxBound :: Int	
HS_INT8_MIN	minBound :: Int8	
HS_INT8_MAX	maxBound :: Int8	
HS_INT16_MIN	minBound :: Int16	
HS_INT16_MAX	maxBound :: Int16	
HS_INT32_MIN	minBound :: Int32	
HS_INT32_MAX	maxBound :: Int32	
HS_INT64_MIN	minBound :: Int64	
HS_INT64_MAX	maxBound :: Int64	
HS_WORD8_MAX	maxBound :: Word8	
HS_WORD16_MAX	maxBound :: Word16	
HS_WORD32_MAX	maxBound :: Word32	
HS_WORD64_MAX	maxBound :: Word64	
HS_FLOAT_RADIX	floatRadix :: Float	
HS_FLOAT_ROUND	n/a	rounding style as per [13]
HS_FLOAT_EPSILON	n/a	difference between 1 and the least value greater than 1 as per [13]
HS_DOUBLE_EPSILON	n/a	(as above)
HS_FLOAT_DIG	n/a	number of decimal digits as per [13]
HS_DOUBLE_DIG	n/a	(as above)
HS_FLOAT_MANT_DIG	floatDigits :: Float	
HS_DOUBLE_MANT_DIG	floatDigits :: Double	
HS_FLOAT_MIN	n/a	minimum floating point number as per [13]
HS_DOUBLE_MIN	n/a	(as above)
HS_FLOAT_MIN_EXP	fst . floatRange :: Float	
HS_DOUBLE_MIN_EXP	fst . floatRange :: Double	
HS_FLOAT_MIN_10_EXP	n/a	minimum decimal exponent as per [13]
HS_DOUBLE_MIN_10_EXP	n/a	(as above)
HS_FLOAT_MAX	n/a	maximum floating point number as per [13]
HS_DOUBLE_MAX	n/a	(as above)
HS_FLOAT_MAX_EXP	snd . floatRange :: Float	
HS_DOUBLE_MAX_EXP	snd . floatRange :: Double	
HS_FLOAT_MAX_10_EXP	n/a	maximum decimal exponent as per [13]
HS_DOUBLE_MAX_10_EXP	n/a	(as above)
HS_BOOL_FALSE	False	
HS_BOOL_TRUE	True	

Table 5: C Interface to Range and Precision of Basic Types

Table 5 contains symbols characterising the range and precision of the types from Table 4. Where available, the table states the corresponding Haskell values. All C symbols, with the exception of `HS_FLOAT_ROUND` are constants that are suitable for use in `#if` preprocessing directives. Note that there is only one rounding style (`HS_FLOAT_ROUND`) and one radix (`HS_FLOAT_RADIX`), as this is all that is supported by ISO C [13].

Moreover, an implementation that does not support 64 bit integral types on the C side should implement `HsInt64` and `HsWord64` as a structure. In this case, the bounds `HS_INT64_MIN`, `HS_INT64_MAX`, and `HS_WORD64_MAX` are undefined.

In addition, to the symbols from Table 4 and Table 5, the header `HsFFI.h` must also contain the following prototypes:

```
void hs_init     (int *argc, char **argv[]);
void hs_exit     (void);
void hs_set_argv (int argc, char *argv[]);

void hs_perform_gc (void);

void hs_free_stable_ptr (HsStablePtr sp);
void hs_free_fun_ptr    (HsFunPtr fp);
```

These routines are useful for mixed language programs, where the main application is implemented in a foreign language that accesses routines implemented in Haskell. The function `hs_init()` initialises the Haskell system and provides it with the available command line arguments. Upon return, the arguments solely intended for the Haskell runtime system are removed (i.e., the values that `argc` and `argv` point to may have changed). This function must be called during program startup before any Haskell function is invoked; otherwise, the system behaviour is undefined. Conversely, the Haskell system is deinitialised by a call to `hs_exit()`. Multiple invocations of `hs_init()` are permitted, provided that they are followed by an equal number of calls to `hs_exit()` and that the first call to `hs_exit()` is after the last call to `hs_init()`. In addition to nested calls to `hs_init()`, the Haskell system may be deinitialised with `hs_exit()` and be re-initialised with `hs_init()` at a later point in time. This ensures that repeated initialisation due to multiple libraries being implemented in Haskell is covered.

The Haskell system will ignore the command line arguments passed to the second and any following calls to `hs_init()`. Moreover, `hs_init()` may be called with `NULL` for both `argc` and `argv`, signalling the absence of command line arguments.

The function `hs_set_argv()` sets the values returned by the functions `getProgName` and `getArgs` of the module `System.Environment` (Section ref{module:System.Environment}). This function may only be invoked after `hs_init()`. Moreover, if `hs_set_argv()` is called at all, this call must precede the first invocation of `getProgName` and `getArgs`. Note that the separation of `hs_init()` and `hs_set_argv()` is essential in cases where in addition to the Haskell system other libraries that process command line arguments during initialisation are used.

The function `hs_perform_gc()` advises the Haskell storage manager to perform a garbage collection, where the storage manager makes an effort to releases all unreachable objects. This function must not be invoked from C functions that are imported `unsafe` into Haskell code nor may it be used from a finalizer.

Finally, `hs_free_stable_ptr()` and `hs_free_fun_ptr()` are the C counterparts of the Haskell functions `freeStablePtr` and `freeHaskellFunPtr`.

Chapter 9

Standard Prelude

In this chapter the entire Haskell Prelude is given. It constitutes a *specification* for the Prelude. Many of the definitions are written with clarity rather than efficiency in mind, and it is not required that the specification be implemented as shown here.

The default method definitions, given with `class` declarations, constitute a specification *only* of the default method. They do not constitute a specification of the meaning of the method in all instances. To take one particular example, the default method for `enumFrom` in class `Enum` will not work properly for types whose range exceeds that of `Int` (because `fromEnum` cannot map all values in the type to distinct `Int` values).

The Prelude shown here is organized into a root module, `Prelude`, and three sub-modules, `PreludeList`, `PreludeText`, and `PreludeIO`. This structure is purely presentational. An implementation is not required to use this organisation for the Prelude, nor are these three modules available for import separately. Only the exports of module `Prelude` are significant.

Some of these modules import Library modules, such as `Data.Char`, `Control.Monad`, `System.IO`, and `Numeric`. These modules are described fully in part:libraries[Part] These imports are not, of course, part of the specification of the `Prelude`. That is, an implementation is free to import more, or less, of the Library modules, as it pleases.

Primitives that are not definable in Haskell, indicated by names starting with "`prim`", are defined in a system dependent manner in module `PreludeBuiltin` and are not shown here. Instance declarations that simply bind primitives to class methods are omitted. Some of the more verbose instances with obvious functionality have been left out for the sake of brevity.

Declarations for special types such as `Integer`, or `()` are included in the Prelude for completeness even though the declaration may be incomplete or syntactically invalid. An ellipsis "`...`" is often used in places where the remainder of a definition cannot be given in Haskell.

To reduce the occurrence of unexpected ambiguity errors, and to improve efficiency, a number of commonly-used functions over lists use the `Int` type rather than using a more general numeric type, such as `Integral a` or `Num a`. These functions are: `take`, `drop`, `!!`, `length`, `splitAt`, and `replicate`. The more general versions are given in the `Data.List` library, with the prefix "`generic`"; for example `genericLength`.

```
module Prelude (
    module PreludeList, module PreludeText, module PreludeIO,
    Bool(False, True),
    Maybe(Nothing, Just),
    Either(Left, Right),
    Ordering(LT, EQ, GT),
    Char, String, Int, Integer, Float, Double, Rational, IO,

--      These built-in types are defined in the Prelude, but
```

```
--      are denoted by built-in syntax, and cannot legally
--      appear in an export list.
--  List type: []((:), [])
--  Tuple types: (,)((,)), (,,)((,,)), etc.
--  Trivial type: ()(())
--  Functions: (->)

    Eq((==), (/=)),
    Ord(compare, (<), (<=), (>=), (>), max, min),
    Enum(succ, pred, toEnum, fromEnum, enumFrom, enumFromThen,
         enumFromTo, enumFromThenTo),
    Bounded(minBound, maxBound),
    Num((+), (-), (*), negate, abs, signum, fromInteger),
    Real(toRational),
    Integral(quot, rem, div, mod, quotRem, divMod, toInteger),
    Fractional((/), recip, fromRational),
    Floating(pi, exp, log, sqrt, (**), logBase, sin, cos, tan,
             asin, acos, atan, sinh, cosh, tanh, asinh, acosh, atanh),
    RealFrac(properFraction, truncate, round, ceiling, floor),
    RealFloat(floatRadix, floatDigits, floatRange, decodeFloat,
              encodeFloat, exponent, significand, scaleFloat, isNaN,
              isInfinite, isDenormalized, isIEEE, isNegativeZero, atan2),
    Monad((>>=), (>>), return, fail),
    Functor(fmap),
    mapM, mapM_, sequence, sequence_, (=<<),
    maybe, either,
    (&&), (||), not, otherwise,
    subtract, even, odd, gcd, lcm, (^), (^^),
    fromIntegral, realToFrac,
    fst, snd, curry, uncurry, id, const, (.), flip, ($), until,
    asTypeOf, error, undefined,
    seq, ($!)
  ) where

import PreludeBuiltin                       -- Contains all `prim' values
import UnicodePrims( primUnicodeMaxChar )   -- Unicode primitives
import PreludeList
import PreludeText
import PreludeIO
import Data.Ratio( Rational )

infixr 9  .
infixr 8  ^, ^^, **
infixl 7  *, /, `quot`, `rem`, `div`, `mod`
infixl 6  +, -

-- The (:) operator is built-in syntax, and cannot legally be given
-- a fixity declaration; but its fixity is given by:
--   infixr 5  :

infix  4  ==, /=, <, <=, >=, >
infixr 3  &&
infixr 2  ||
infixl 1  >>, >>=
infixr 1  =<<
infixr 0  $, $!, `seq`
```

```
-- Standard types, classes, instances and related functions

-- Equality and Ordered classes

class  Eq a  where
    (==), (/=) :: a -> a -> Bool

        -- Minimal complete definition:
        --      (==) or (/=)
    x /= y     =  not (x == y)
    x == y     =  not (x /= y)

class  (Eq a) => Ord a  where
    compare              :: a -> a -> Ordering
    (<), (<=), (>=), (>) :: a -> a -> Bool
    max, min             :: a -> a -> a

        -- Minimal complete definition:
        --      (<=) or compare
        -- Using compare can be more efficient for complex types.
    compare x y
         | x == y    =  EQ
         | x <= y    =  LT
         | otherwise =  GT

    x <= y           =  compare x y /= GT
    x <  y           =  compare x y == LT
    x >= y           =  compare x y /= LT
    x >  y           =  compare x y == GT

-- note that (min x y, max x y) = (x,y) or (y,x)
    max x y
         | x <= y    =  y
         | otherwise =  x
    min x y
         | x <= y    =  x
         | otherwise =  y

-- Enumeration and Bounded classes

class  Enum a  where
    succ, pred       :: a -> a
    toEnum           :: Int -> a
    fromEnum         :: a -> Int
    enumFrom         :: a -> [a]             -- [n..]
    enumFromThen     :: a -> a -> [a]        -- [n,n'..]
    enumFromTo       :: a -> a -> [a]        -- [n..m]
    enumFromThenTo   :: a -> a -> a -> [a]   -- [n,n'..m]

        -- Minimal complete definition:
        --      toEnum, fromEnum
  --
  -- NOTE: these default methods only make sense for types
  --   that map injectively into Int using fromEnum
  --   and toEnum.
```

```
    succ              =  toEnum . (+1) . fromEnum
    pred              =  toEnum . (subtract 1) . fromEnum
    enumFrom x        =  map toEnum [fromEnum x ..]
    enumFromTo x y    =  map toEnum [fromEnum x .. fromEnum y]
    enumFromThen x y  =  map toEnum [fromEnum x, fromEnum y ..]
    enumFromThenTo x y z =
                         map toEnum [fromEnum x, fromEnum y .. fromEnum z]

class  Bounded a  where
    minBound          :: a
    maxBound          :: a

-- Numeric classes

class  (Eq a, Show a) => Num a  where
    (+), (-), (*)     :: a -> a -> a
    negate            :: a -> a
    abs, signum       :: a -> a
    fromInteger       :: Integer -> a

        -- Minimal complete definition:
        --      All, except negate or (-)
    x - y             =  x + negate y
    negate x          =  0 - x

class  (Num a, Ord a) => Real a  where
    toRational        ::  a -> Rational

class  (Real a, Enum a) => Integral a  where
    quot, rem         :: a -> a -> a
    div, mod          :: a -> a -> a
    quotRem, divMod   :: a -> a -> (a,a)
    toInteger         :: a -> Integer

        -- Minimal complete definition:
        --      quotRem, toInteger
    n `quot` d        =  q  where (q,r) = quotRem n d
    n `rem` d         =  r  where (q,r) = quotRem n d
    n `div` d         =  q  where (q,r) = divMod n d
    n `mod` d         =  r  where (q,r) = divMod n d
    divMod n d        =  if signum r == - signum d then (q-1, r+d) else qr
                         where qr@(q,r) = quotRem n d

class  (Num a) => Fractional a  where
    (/)               :: a -> a -> a
    recip             :: a -> a
    fromRational      :: Rational -> a

        -- Minimal complete definition:
        --      fromRational and (recip or (/))
    recip x           =  1 / x
    x / y             =  x * recip y

class  (Fractional a) => Floating a  where
    pi                   :: a
    exp, log, sqrt       :: a -> a
```

```
    (**), logBase       :: a -> a -> a
    sin, cos, tan       :: a -> a
    asin, acos, atan    :: a -> a
    sinh, cosh, tanh    :: a -> a
    asinh, acosh, atanh :: a -> a

        -- Minimal complete definition:
        --      pi, exp, log, sin, cos, sinh, cosh
        --      asin, acos, atan
        --      asinh, acosh, atanh
    x ** y           =  exp (log x * y)
    logBase x y      =  log y / log x
    sqrt x           =  x ** 0.5
    tan  x           =  sin  x / cos  x
    tanh x           =  sinh x / cosh x


class  (Real a, Fractional a) => RealFrac a  where
    properFraction   :: (Integral b) => a -> (b,a)
    truncate, round  :: (Integral b) => a -> b
    ceiling, floor   :: (Integral b) => a -> b

        -- Minimal complete definition:
        --      properFraction
    truncate x       =  m  where (m,_) = properFraction x

    round x          =  let (n,r) = properFraction x
                            m     = if r < 0 then n - 1 else n + 1
                        in case signum (abs r - 0.5) of
                             -1 -> n
                             0  -> if even n then n else m
                             1  -> m

    ceiling x        =  if r > 0 then n + 1 else n
                        where (n,r) = properFraction x

    floor x          =  if r < 0 then n - 1 else n
                        where (n,r) = properFraction x

class  (RealFrac a, Floating a) => RealFloat a  where
    floatRadix       :: a -> Integer
    floatDigits      :: a -> Int
    floatRange       :: a -> (Int,Int)
    decodeFloat      :: a -> (Integer,Int)
    encodeFloat      :: Integer -> Int -> a
    exponent         :: a -> Int
    significand      :: a -> a
    scaleFloat       :: Int -> a -> a
    isNaN, isInfinite, isDenormalized, isNegativeZero, isIEEE
                     :: a -> Bool
    atan2            :: a -> a -> a

        -- Minimal complete definition:
        --      All except exponent, significand,
        --                 scaleFloat, atan2
    exponent x       =  if m == 0 then 0 else n + floatDigits x
```

```
                    where (m,n) = decodeFloat x

    significand x   =  encodeFloat m (- floatDigits x)
                    where (m,_) = decodeFloat x

    scaleFloat k x  =  encodeFloat m (n+k)
                    where (m,n) = decodeFloat x

    atan2 y x
      | x>0           =  atan (y/x)
      | x==0 && y>0   =  pi/2
      | x<0  && y>0   =  pi + atan (y/x)
      |(x<=0 && y<0)  ||
       (x<0 && isNegativeZero y) ||
       (isNegativeZero x && isNegativeZero y)
                      = -atan2 (-y) x
      | y==0 && (x<0 || isNegativeZero x)
                      =  pi    -- must be after the previous test on zero y
      | x==0 && y==0  =  y     -- must be after the other double zero tests
      | otherwise     =  x + y -- x or y is a NaN, return a NaN (via +)

-- Numeric functions

subtract        :: (Num a) => a -> a -> a
subtract        =  flip (-)

even, odd       :: (Integral a) => a -> Bool
even n          =  n `rem` 2 == 0
odd             =  not . even

gcd             :: (Integral a) => a -> a -> a
gcd 0 0         =  error "Prelude.gcd: gcd 0 0 is undefined"
gcd x y         =  gcd' (abs x) (abs y)
                   where gcd' x 0  =  x
                         gcd' x y  =  gcd' y (x `rem` y)

lcm             :: (Integral a) => a -> a -> a
lcm _ 0         =  0
lcm 0 _         =  0
lcm x y         =  abs ((x `quot` (gcd x y)) * y)

(^)             :: (Num a, Integral b) => a -> b -> a
x ^ 0           =  1
x ^ n | n > 0   =  f x (n-1) x
                   where f _ 0 y = y
                         f x n y = g x n  where
                                   g x n | even n  = g (x*x) (n `quot` 2)
                                         | otherwise = f x (n-1) (x*y)
_ ^ _           = error "Prelude.^: negative exponent"

(^^)            :: (Fractional a, Integral b) => a -> b -> a
x ^^ n          =  if n >= 0 then x^n else recip (x^(-n))

fromIntegral    :: (Integral a, Num b) => a -> b
fromIntegral    =  fromInteger . toInteger
```

```
realToFrac     :: (Real a, Fractional b) => a -> b
realToFrac      =  fromRational . toRational

-- Monadic classes

class  Functor f  where
    fmap              :: (a -> b) -> f a -> f b

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

        -- Minimal complete definition:
        --      (>>=), return
    m >> k  =  m >>= \_ -> k
    fail s  = error s

sequence       :: Monad m => [m a] -> m [a]
sequence       =  foldr mcons (return [])
                    where mcons p q = p >>= \x -> q >>= \y -> return (x:y)

sequence_      :: Monad m => [m a] -> m ()
sequence_      =  foldr (>>) (return ())

-- The xxxM functions take list arguments, but lift the function or
-- list element to a monad type
mapM             :: Monad m => (a -> m b) -> [a] -> m [b]
mapM f as        =  sequence (map f as)

mapM_            :: Monad m => (a -> m b) -> [a] -> m ()
mapM_ f as       =  sequence_ (map f as)

(=<<)            :: Monad m => (a -> m b) -> m a -> m b
f =<< x          =  x >>= f


-- Trivial type

data  ()  =  ()  deriving (Eq, Ord, Enum, Bounded)
  -- Not legal Haskell; for illustration only

-- Function type

-- identity function
id               :: a -> a
id x             =  x

-- constant function
const            :: a -> b -> a
const x _        =  x

-- function composition
(.)              :: (b -> c) -> (a -> b) -> a -> c
f . g            =  \ x -> f (g x)
```

```haskell
-- flip f  takes its (first) two arguments in the reverse order of f.
flip             :: (a -> b -> c) -> b -> a -> c
flip f x y       =  f y x

seq :: a -> b -> b
seq = ...       -- Primitive

-- right-associating infix application operators
-- (useful in continuation-passing style)
($), ($!) :: (a -> b) -> a -> b
f $  x    =  f x
f $! x    =  x `seq` f x


-- Boolean type

data  Bool  =  False | True     deriving (Eq, Ord, Enum, Read, Show, Bounded)

-- Boolean functions

(&&), (||)       :: Bool -> Bool -> Bool
True  && x       =  x
False && _       =  False
True  || _       =  True
False || x       =  x

not              :: Bool -> Bool
not True         =  False
not False        =  True

otherwise        :: Bool
otherwise        =  True


-- Character type

data Char = ... 'a' | 'b' ... -- Unicode values

instance  Eq Char  where
    c == c'          =  fromEnum c == fromEnum c'

instance  Ord Char  where
    c <= c'          =  fromEnum c <= fromEnum c'

instance  Enum Char  where
    toEnum            = primIntToChar
    fromEnum          = primCharToInt
    enumFrom c        = map toEnum [fromEnum c .. fromEnum (maxBound::Char)]
    enumFromThen c c' = map toEnum [fromEnum c, fromEnum c' .. fromEnum lastChar]
                        where lastChar :: Char
                              lastChar | c' < c    = minBound
                                       | otherwise = maxBound

instance  Bounded Char  where
    minBound  =  '\0'
```

```
    maxBound  =  primUnicodeMaxChar

type  String = [Char]


-- Maybe type

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

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

instance  Functor Maybe  where
    fmap f Nothing    =  Nothing
    fmap f (Just x)   =  Just (f x)

instance  Monad Maybe  where
    (Just x) >>= k   =  k x
    Nothing  >>= k   =  Nothing
    return           =  Just
    fail s           =  Nothing

-- Either type

data  Either a b  =  Left a | Right b   deriving (Eq, Ord, Read, Show)

either               :: (a -> c) -> (b -> c) -> Either a b -> c
either f g (Left x)  =  f x
either f g (Right y) =  g y

-- IO type

data IO a = ...   -- abstract

instance  Functor IO where
   fmap f x           =  x >>= (return . f)

instance Monad IO where
   (>>=)  = ...
   return = ...
   fail s = ioError (userError s)

-- Ordering type

data  Ordering  =  LT | EQ | GT
          deriving (Eq, Ord, Enum, Read, Show, Bounded)


-- Standard numeric types.  The data declarations for these types cannot
-- be expressed directly in Haskell since the constructor lists would be
-- far too large.

data  Int  =  minBound ... -1 | 0 | 1 ... maxBound
instance  Eq       Int  where ...
instance  Ord      Int  where ...
```

```
instance  Num      Int  where ...
instance  Real     Int  where ...
instance  Integral Int  where ...
instance  Enum     Int  where ...
instance  Bounded  Int  where ...

data  Integer  =  ... -1 | 0 | 1 ...
instance  Eq       Integer  where ...
instance  Ord      Integer  where ...
instance  Num      Integer  where ...
instance  Real     Integer  where ...
instance  Integral Integer  where ...
instance  Enum     Integer  where ...

data  Float
instance  Eq         Float  where ...
instance  Ord        Float  where ...
instance  Num        Float  where ...
instance  Real       Float  where ...
instance  Fractional Float  where ...
instance  Floating   Float  where ...
instance  RealFrac   Float  where ...
instance  RealFloat  Float  where ...

data  Double
instance  Eq         Double  where ...
instance  Ord        Double  where ...
instance  Num        Double  where ...
instance  Real       Double  where ...
instance  Fractional Double  where ...
instance  Floating   Double  where ...
instance  RealFrac   Double  where ...
instance  RealFloat  Double  where ...

-- The Enum instances for Floats and Doubles are slightly unusual.
-- The `toEnum' function truncates numbers to Int.  The definitions
-- of enumFrom and enumFromThen allow floats to be used in arithmetic
-- series: [0,0.1 .. 0.95].  However, roundoff errors make these somewhat
-- dubious.  This example may have either 10 or 11 elements, depending on
-- how 0.1 is represented.

instance  Enum Float  where
    succ x           =  x+1
    pred x           =  x-1
    toEnum           =  fromIntegral
    fromEnum         =  fromInteger . truncate   -- may overflow
    enumFrom         =  numericEnumFrom
    enumFromThen     =  numericEnumFromThen
    enumFromTo       =  numericEnumFromTo
    enumFromThenTo   =  numericEnumFromThenTo

instance  Enum Double  where
    succ x           =  x+1
    pred x           =  x-1
    toEnum           =  fromIntegral
    fromEnum         =  fromInteger . truncate   -- may overflow
```

```
    enumFrom         =  numericEnumFrom
    enumFromThen     =  numericEnumFromThen
    enumFromTo       =  numericEnumFromTo
    enumFromThenTo   =  numericEnumFromThenTo

numericEnumFrom         :: (Fractional a) => a -> [a]
numericEnumFromThen     :: (Fractional a) => a -> a -> [a]
numericEnumFromTo       :: (Fractional a, Ord a) => a -> a -> [a]
numericEnumFromThenTo   :: (Fractional a, Ord a) => a -> a -> a -> [a]
numericEnumFrom         =  iterate (+1)
numericEnumFromThen n m =  iterate (+(m-n)) n
numericEnumFromTo n m   =  takeWhile (<= m+1/2) (numericEnumFrom n)
numericEnumFromThenTo n n' m = takeWhile p (numericEnumFromThen n n')
                           where
                             p | n' >= n   = (<= m + (n'-n)/2)
                               | otherwise = (>= m + (n'-n)/2)

-- Lists

data  [a]  =  [] | a : [a]  deriving (Eq, Ord)
  -- Not legal Haskell; for illustration only

instance Functor [] where
    fmap = map

instance  Monad []  where
    m >>= k          = concat (map k m)
    return x         = [x]
    fail s           = []

-- Tuples

data  (a,b)   =  (a,b)    deriving (Eq, Ord, Bounded)
data  (a,b,c) =  (a,b,c)  deriving (Eq, Ord, Bounded)
  -- Not legal Haskell; for illustration only

-- component projections for pairs:
-- (NB: not provided for triples, quadruples, etc.)
fst             :: (a,b) -> a
fst (x,y)       =  x

snd             :: (a,b) -> b
snd (x,y)       =  y

-- curry converts an uncurried function to a curried function;
-- uncurry converts a curried function to a function on pairs.
curry           :: ((a, b) -> c) -> a -> b -> c
curry f x y     =  f (x, y)

uncurry         :: (a -> b -> c) -> ((a, b) -> c)
uncurry f p     =  f (fst p) (snd p)

-- Misc functions

-- until p f  yields the result of applying f until p holds.
until           :: (a -> Bool) -> (a -> a) -> a -> a
```

```
until p f x
     | p x       =  x
     | otherwise =  until p f (f x)

-- asTypeOf is a type-restricted version of const.  It is usually used
-- as an infix operator, and its typing forces its first argument
-- (which is usually overloaded) to have the same type as the second.
asTypeOf         :: a -> a -> a
asTypeOf         =  const

-- error stops execution and displays an error message

error            :: String -> a
error            =  primError

-- It is expected that compilers will recognize this and insert error
-- messages that are more appropriate to the context in which undefined
-- appears.

undefined        :: a
undefined        =  error "Prelude.undefined"
```

9.1 Prelude PreludeList

```
-- Standard list functions

module PreludeList (
    map, (++), filter, concat, concatMap,
    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,
    zip, zip3, zipWith, zipWith3, unzip, unzip3)
  where

import qualified Data.Char(isSpace)

infixl 9  !!
infixr 5  ++
infix  4  `elem`, `notElem`

-- Map and append
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)
```

```
filter :: (a -> Bool) -> [a] -> [a]
filter p []                 = []
filter p (x:xs) | p x       = x : filter p xs
                | otherwise = filter p xs

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

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

-- head and tail extract the first element and remaining elements,
-- respectively, of a list, which must be non-empty.  last and init
-- are the dual functions working from the end of a finite list,
-- rather than the beginning.

head             :: [a] -> a
head (x:_)       =  x
head []          =  error "Prelude.head: empty list"

tail             :: [a] -> [a]
tail (_:xs)      =  xs
tail []          =  error "Prelude.tail: empty list"

last             :: [a] -> a
last [x]         =  x
last (_:xs)      =  last xs
last []          =  error "Prelude.last: empty list"

init             :: [a] -> [a]
init [x]         =  []
init (x:xs)      =  x : init xs
init []          =  error "Prelude.init: empty list"

null             :: [a] -> Bool
null []          =  True
null (_:_)       =  False

-- length returns the length of a finite list as an Int.
length           :: [a] -> Int
length []        =  0
length (_:l)     =  1 + length l

-- List index (subscript) operator, 0-origin
(!!)                :: [a] -> Int -> a
xs     !! n | n < 0 =  error "Prelude.!!: negative index"
[]     !! _         =  error "Prelude.!!: index too large"
(x:_)  !! 0         =  x
(_:xs) !! n         =  xs !! (n-1)

-- foldl, applied to a binary operator, a starting value (typically the
-- left-identity of the operator), and a list, reduces the list using
-- the binary operator, from left to right:
--  foldl f z [x1, x2, ..., xn] == (...((z `f` x1) `f` x2) `f`...) `f` xn
-- foldl1 is a variant that has no starting value argument, and  thus must
-- be applied to non-empty lists.  scanl is similar to foldl, but returns
```

```haskell
-- a list of successive reduced values from the left:
--      scanl f z [x1, x2, ...] == [z, z `f` x1, (z `f` x1) `f` x2, ...]
-- Note that  last (scanl f z xs) == foldl f z xs.
-- scanl1 is similar, again without the starting element:
--      scanl1 f [x1, x2, ...] == [x1, x1 `f` x2, ...]

foldl            :: (a -> b -> a) -> a -> [b] -> a
foldl f z []     =  z
foldl f z (x:xs) =  foldl f (f z x) xs

foldl1           :: (a -> a -> a) -> [a] -> a
foldl1 f (x:xs)  =  foldl f x xs
foldl1 _ []      =  error "Prelude.foldl1: empty list"

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 _ []      =  []

-- foldr, foldr1, scanr, and scanr1 are the right-to-left duals of the
-- above functions.

foldr            :: (a -> b -> b) -> b -> [a] -> b
foldr f z []     =  z
foldr f z (x:xs) =  f x (foldr f z xs)

foldr1           :: (a -> a -> a) -> [a] -> a
foldr1 f [x]     =  x
foldr1 f (x:xs)  =  f x (foldr1 f xs)
foldr1 _ []      =  error "Prelude.foldr1: empty list"

scanr             :: (a -> b -> b) -> b -> [a] -> [b]
scanr f q0 []     =  [q0]
scanr f q0 (x:xs) =  f x q : qs
                     where qs@(q:_) = scanr f q0 xs

scanr1          :: (a -> a -> a) -> [a] -> [a]
scanr1 f []     =  []
scanr1 f [x]    =  [x]
scanr1 f (x:xs) =  f x q : qs
                   where qs@(q:_) = scanr1 f xs

-- iterate f x returns an infinite list of repeated applications of f to x:
-- iterate f x == [x, f x, f (f x), ...]
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.
repeat           :: a -> [a]
repeat x         =  xs where xs = x:xs

-- replicate n x is a list of length n with x the value of every element
```

```
replicate        :: Int -> a -> [a]
replicate n x    =  take n (repeat x)

-- cycle ties a finite list into a circular one, or equivalently,
-- the infinite repetition of the original list.  It is the identity
-- on infinite lists.

cycle            :: [a] -> [a]
cycle []         =  error "Prelude.cycle: empty list"
cycle xs         =  xs' where xs' = xs ++ xs'

-- take n, applied to a list xs, returns the prefix of xs of length n,
-- or xs itself if n > length xs.  drop n xs returns the suffix of xs
-- after the first n elements, or [] if n > length xs.  splitAt n xs
-- is equivalent to (take n xs, drop n xs).

take                    :: Int -> [a] -> [a]
take n _      | n <= 0 =  []
take _ []               =  []
take n (x:xs)           =  x : take (n-1) xs

drop                    :: Int -> [a] -> [a]
drop n xs     | n <= 0 =  xs
drop _ []               =  []
drop n (_:xs)           =  drop (n-1) xs

splitAt                  :: Int -> [a] -> ([a],[a])
splitAt n xs             =  (take n xs, drop n xs)

-- takeWhile, applied to a predicate p and a list xs, returns the longest
-- prefix (possibly empty) of xs of elements that satisfy p.  dropWhile p xs
-- returns the remaining suffix.  span p xs is equivalent to
-- (takeWhile p xs, dropWhile p xs), while break p uses the negation of p.

takeWhile               :: (a -> Bool) -> [a] -> [a]
takeWhile p []          =  []
takeWhile p (x:xs)
            | p x       =  x : takeWhile p xs
            | otherwise =  []

dropWhile               :: (a -> Bool) -> [a] -> [a]
dropWhile p []          =  []
dropWhile p xs@(x:xs')
            | p x       =  dropWhile p xs'
            | otherwise =  xs

span, break             :: (a -> Bool) -> [a] -> ([a],[a])
span p []            = ([],[])
span p xs@(x:xs')
            | p x       =  (x:ys,zs)
            | otherwise =  ([],xs)
                           where (ys,zs) = span p xs'

break p                 =  span (not . p)

-- lines breaks a string up into a list of strings at newline characters.
```

```
-- The resulting strings do not contain newlines.  Similary, words
-- breaks a string up into a list of words, which were delimited by
-- white space.  unlines and unwords are the inverse operations.
-- unlines joins lines with terminating newlines, and unwords joins
-- words with separating spaces.

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 (\w s -> w ++ ' ':s) ws

-- reverse xs returns the elements of xs in reverse order.  xs must be finite.
reverse         :: [a] -> [a]
reverse         =  foldl (flip (:)) []

-- and returns the conjunction of a Boolean list.  For the result to be
-- True, the list must be finite; False, however, results from a False
-- value at a finite index of a finite or infinite list.  or is the
-- disjunctive dual of and.
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.  Similarly, for all.
any, all        :: (a -> Bool) -> [a] -> Bool
any p           =  or . map p
all p           =  and . map p

-- elem is the list membership predicate, usually written in infix form,
-- e.g., x `elem` xs.  notElem is the negation.
elem, notElem   :: (Eq a) => a -> [a] -> Bool
elem x          =  any (== x)
notElem x       =  all (/= x)

-- lookup key assocs looks up a key in an association list.
lookup          :: (Eq a) => a -> [(a,b)] -> Maybe b
lookup key []   =  Nothing
lookup key ((x,y):xys)
    | key == x  =  Just y
    | otherwise =  lookup key xys
```

```
-- sum and product compute the sum or product of a finite list of numbers.
sum, product     :: (Num a) => [a] -> a
sum              =  foldl (+) 0
product          =  foldl (*) 1

-- maximum and minimum return the maximum or minimum value from a list,
-- which must be non-empty, finite, and of an ordered type.
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

-- 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 in the List library

zip              :: [a] -> [b] -> [(a,b)]
zip              =  zipWith (,)

zip3             :: [a] -> [b] -> [c] -> [(a,b,c)]
zip3             =  zipWith3 (,,)

-- The zipWith family generalises the zip family by zipping with the
-- function given as the first argument, instead of a tupling function.
-- For example, zipWith (+) is applied to two lists to produce the list
-- of corresponding sums.

zipWith          :: (a->b->c) -> [a]->[b]->[c]
zipWith z (a:as) (b:bs)
                 =  z a b : zipWith z as bs
zipWith _ _ _    =  []

zipWith3         :: (a->b->c->d) -> [a]->[b]->[c]->[d]
zipWith3 z (a:as) (b:bs) (c:cs)
                 =  z a b c : zipWith3 z as bs cs
zipWith3 _ _ _ _ =  []


-- unzip transforms a list of pairs into a pair of lists.

unzip            :: [(a,b)] -> ([a],[b])
unzip            =  foldr (\(a,b) ~(as,bs) -> (a:as,b:bs)) ([],[])

unzip3           :: [(a,b,c)] -> ([a],[b],[c])
unzip3           =  foldr (\(a,b,c) ~(as,bs,cs) -> (a:as,b:bs,c:cs))
                          ([],[],[])
```

9.2 Prelude PreludeText

```
module PreludeText (
    ReadS, ShowS,
    Read(readsPrec, readList),
    Show(showsPrec, show, showList),
    reads, shows, read, lex,
    showChar, showString, readParen, showParen ) where

-- The instances of Read and Show for
--      Bool, Maybe, Either, Ordering
-- are done via "deriving" clauses in Prelude.hs

import Data.Char(isSpace, isAlpha, isDigit, isAlphaNum,
                 showLitChar, readLitChar, lexLitChar)

import Numeric(showSigned, showInt, readSigned, readDec, showFloat,
               readFloat, lexDigits)

type  ReadS a  = String -> [(a,String)]
type  ShowS    = String -> String

class  Read a  where
    readsPrec        :: Int -> ReadS a
    readList         :: ReadS [a]

        -- Minimal complete definition:
        --      readsPrec
    readList         = readParen False (\r -> [pr | ("[",s)  <- lex r,
                                                   pr       <- readl s])
                       where readl  s = [([],t)   | ("]",t)  <- lex s] ++
                                        [(x:xs,u) | (x,t)    <- reads s,
                                                    (xs,u)   <- readl' t]
                             readl' s = [([],t)   | ("]",t)  <- lex s] ++
                                        [(x:xs,v) | (",",t)  <- lex s,
                                                    (x,u)    <- reads t,
                                                    (xs,v)   <- readl' u]

class  Show a  where
    showsPrec        :: Int -> a -> ShowS
    show             :: a -> String
    showList         :: [a] -> ShowS

        -- Mimimal complete definition:
        --      show or showsPrec
    showsPrec _ x s   = show x ++ s

    show x            = showsPrec 0 x ""

    showList []       = showString "[]"
    showList (x:xs)   = showChar '[' . shows x . showl xs
                        where showl []     = showChar ']'
                              showl (x:xs) = showChar ',' . shows x .
                                             showl xs
```

```
reads           :: (Read a) => ReadS a
reads           =  readsPrec 0

shows           :: (Show a) => a -> ShowS
shows           =  showsPrec 0

read            :: (Read a) => String -> a
read s          =  case [x | (x,t) <- reads s, ("","") <- lex t] of
                        [x] -> x
                        []  -> error "Prelude.read: no parse"
                        _   -> error "Prelude.read: ambiguous parse"

showChar        :: Char -> ShowS
showChar        =  (:)

showString      :: String -> ShowS
showString      =  (++)

showParen       :: Bool -> ShowS -> ShowS
showParen b p   =  if b then showChar '(' . p . showChar ')' else p

readParen       :: Bool -> ReadS a -> ReadS a
readParen b g   =  if b then mandatory else optional
                   where optional r  = g r ++ mandatory r
                         mandatory r = [(x,u) | ("(",s) <- lex r,
                                                (x,t)   <- optional s,
                                                (")",u) <- lex t    ]

-- This lexer is not completely faithful to the Haskell lexical syntax.
-- Current limitations:
--    Qualified names are not handled properly
--    Octal and hexidecimal numerics are not recognized as a single token
--    Comments are not treated properly

lex             :: ReadS String
lex ""          =  [("","")]
lex (c:s)
   | isSpace c  =  lex (dropWhile isSpace s)
lex ('\'':s)    =  [('\'':ch++"'", t) | (ch,'\'':t)  <- lexLitChar s,
                                        ch /= "'" ]
lex ('"':s)     =  [('"':str, t)      | (str,t) <- lexString s]
                   where
                   lexString ('"':s) = [("\"",s)]
                   lexString s = [(ch++str, u)
                                        | (ch,t)  <- lexStrItem s,
                                          (str,u) <- lexString t  ]

                   lexStrItem ('\\':'&':s) =  [("\\&",s)]
                   lexStrItem ('\\':c:s) | isSpace c
                                         =  [("\\&",t) |
                                             '\\':t <-
                                                 [dropWhile isSpace s]]
                   lexStrItem s          =  lexLitChar s

lex (c:s) | isSingle c = [([c],s)]
          | isSym c    = [(c:sym,t)       | (sym,t) <- [span isSym s]]
```

```
          | isAlpha c  = [(c:nam,t)       | (nam,t) <- [span isIdChar s]]
          | isDigit c  = [(c:ds++fe,t)    | (ds,s)  <- [span isDigit s],
                                            (fe,t)  <- lexFracExp s     ]
          | otherwise  = []    -- bad character
            where
              isSingle c =  c `elem` ",;()[]{}_`"
              isSym c    =  c `elem` "!@#$%&*+./<=>?\\^|:-~"
              isIdChar c =  isAlphaNum c || c `elem` "_'"

              lexFracExp ('.':c:cs) | isDigit c
                          = [('.':ds++e,u) | (ds,t) <- lexDigits (c:cs),
                                             (e,u)  <- lexExp t]
              lexFracExp s  = lexExp s

              lexExp (e:s) | e `elem` "eE"
                       = [(e:c:ds,u) | (c:t)  <- [s], c `elem` "+-",
                                       (ds,u) <- lexDigits t] ++
                         [(e:ds,t)   | (ds,t) <- lexDigits s]
              lexExp s = [("",s)]

instance  Show Int  where
    showsPrec n = showsPrec n . toInteger
        -- Converting to Integer avoids
        -- possible difficulty with minInt

instance  Read Int  where
  readsPrec p r = [(fromInteger i, t) | (i,t) <- readsPrec p r]
        -- Reading at the Integer type avoids
        -- possible difficulty with minInt

instance  Show Integer  where
    showsPrec           = showSigned showInt

instance  Read Integer  where
    readsPrec p         = readSigned readDec

instance  Show Float  where
    showsPrec p         = showFloat

instance  Read Float  where
    readsPrec p         = readSigned readFloat

instance  Show Double  where
    showsPrec p         = showFloat

instance  Read Double  where
    readsPrec p         = readSigned readFloat

instance  Show ()  where
    showsPrec p () = showString "()"

instance Read () where
    readsPrec p    = readParen False
                           (\r -> [((),t) | ("(",s) <- lex r,
                                            (")",t) <- lex s ] )
instance  Show Char  where
```

```
    showsPrec p '\'' = showString "'\\''"
    showsPrec p c    = showChar '\'' . showLitChar c . showChar '\''

    showList cs = showChar '"' . showl cs
                 where showl ""       = showChar '"'
                       showl ('"':cs) = showString "\\\"" . showl cs
                       showl (c:cs)   = showLitChar c . showl cs

instance  Read Char  where
    readsPrec p      = readParen False
                           (\r -> [(c,t) | ('\'':s,t)<- lex r,
                                           (c,"\'")  <- readLitChar s])

    readList = readParen False (\r -> [(l,t) | ('"':s, t) <- lex r,
                                               (l,_)      <- readl s ])
        where readl ('"':s)      = [("",s)]
              readl ('\\':'&':s) = readl s
              readl s            = [(c:cs,u) | (c ,t) <- readLitChar s,
                                               (cs,u) <- readl t       ]

instance  (Show a) => Show [a]  where
    showsPrec p      = showList

instance  (Read a) => Read [a]  where
    readsPrec p      = readList

-- Tuples

instance  (Show a, Show b) => Show (a,b)  where
    showsPrec p (x,y) = showChar '(' . shows x . showChar ',' .
                                       shows y . showChar ')'

instance  (Read a, Read b) => Read (a,b)  where
    readsPrec p       = readParen False
                           (\r -> [((x,y), w) | ("(",s) <- lex r,
                                                (x,t)   <- reads s,
                                                (",",u) <- lex t,
                                                (y,v)   <- reads u,
                                                (")",w) <- lex v ] )

-- Other tuples have similar Read and Show instances
```

9.3 Prelude PreludeIO

```
module PreludeIO (
    FilePath, IOError, ioError, userError, catch,
    putChar, putStr, putStrLn, print,
    getChar, getLine, getContents, interact,
    readFile, writeFile, appendFile, readIO, readLn
  ) where

import PreludeBuiltin
```

```
type  FilePath = String

data IOError    -- The internals of this type are system dependent

instance  Show IOError  where ...
instance  Eq IOError  where ...

ioError    ::  IOError -> IO a
ioError    =   primIOError

userError  ::  String -> IOError
userError  =   primUserError

catch      ::  IO a -> (IOError -> IO a) -> IO a
catch      =   primCatch

putChar    :: Char -> IO ()
putChar    =  primPutChar

putStr     :: String -> IO ()
putStr s   =  mapM_ putChar s

putStrLn   :: String -> IO ()
putStrLn s =  do putStr s
                 putStr "\n"

print      :: Show a => a -> IO ()
print x    =  putStrLn (show x)

getChar    :: IO Char
getChar    =  primGetChar

getLine    :: IO String
getLine    =  do c <- getChar
                 if c == '\n' then return "" else
                    do s <- getLine
                       return (c:s)

getContents :: IO String
getContents =  primGetContents

interact    ::  (String -> String) -> IO ()
-- The hSetBuffering ensures the expected interactive behaviour
interact f  =  do hSetBuffering stdin  NoBuffering
                  hSetBuffering stdout NoBuffering
                  s <- getContents
                  putStr (f s)

readFile   :: FilePath -> IO String
readFile   =  primReadFile

writeFile  :: FilePath -> String -> IO ()
writeFile  =  primWriteFile

appendFile :: FilePath -> String -> IO ()
```

```
appendFile =  primAppendFile

  -- raises an exception instead of an error
readIO   :: Read a => String -> IO a
readIO s =  case [x | (x,t) <- reads s, ("","") <- lex t] of
              [x] -> return x
              []  -> ioError (userError "Prelude.readIO: no parse")
              _   -> ioError (userError "Prelude.readIO: ambiguous parse")

readLn :: Read a => IO a
readLn =  do l <- getLine
             r <- readIO l
             return r
```

Chapter 10

Syntax Reference

10.1 Notational Conventions

These notational conventions are used for presenting syntax:

[*pattern*]	optional
{*pattern*}	zero or more repetitions
(*pattern*)	grouping
$pat_1 \mid pat_2$	choice
$pat_{\langle pat' \rangle}$	difference—elements generated by *pat* except those generated by *pat'*
`fibonacci`	terminal syntax in typewriter font

BNF-like syntax is used throughout, with productions having the form:

$$nonterm \quad \rightarrow \quad alt_1 \mid alt_2 \mid \ldots \mid alt_n$$

In both the lexical and the context-free syntax, there are some ambiguities that are to be resolved by making grammatical phrases as long as possible, proceeding from left to right (in shift-reduce parsing, resolving shift/reduce conflicts by shifting). In the lexical syntax, this is the "maximal munch" rule. In the context-free syntax, this means that conditionals, let-expressions, and lambda abstractions extend to the right as far as possible.

10.2 Lexical Syntax

$$
\begin{array}{lcl}
program & \rightarrow & \{\, lexeme \mid whitespace \,\} \\
lexeme & \rightarrow & qvarid \mid qconid \mid qvarsym \mid qconsym \\
 & \mid & literal \mid special \mid reservedop \mid reservedid \\
literal & \rightarrow & integer \mid float \mid char \mid string \\
special & \rightarrow & (\mid) \mid , \mid ; \mid [\mid] \mid ` \mid \{ \mid \} \\
whitespace & \rightarrow & whitestuff\ \{whitestuff\} \\
whitestuff & \rightarrow & whitechar \mid comment \mid ncomment \\
whitechar & \rightarrow & newline \mid vertab \mid space \mid tab \mid uniWhite \\
newline & \rightarrow & return\ linefeed \mid return \mid linefeed \mid formfeed
\end{array}
$$

return → a carriage return
linefeed → a line feed
vertab → a vertical tab
formfeed → a form feed
space → a space
tab → a horizontal tab
uniWhite → any Unicode character defined as whitespace
comment → *dashes* $\left[any_{\langle symbol \rangle} \{ any \} \right]$ *newline*
dashes → `--` {`-`}
opencom → `{-`
closecom → `-}`
ncomment → *opencom ANYseq* {*ncomment ANYseq*} *closecom*
ANYseq → $\{\ ANY \}_{\langle \{\ ANY \} (opencom \mid closecom) \{\ ANY \} \rangle}$
ANY → *graphic* | *whitechar*
any → *graphic* | *space* | *tab*
graphic → *small* | *large* | *symbol* | *digit* | *special* | `"` | `'`
small → *ascSmall* | *uniSmall* | `_`
ascSmall → `a` | `b` | ... | `z`
uniSmall → any Unicode lowercase letter
large → *ascLarge* | *uniLarge*
ascLarge → `A` | `B` | ... | `Z`
uniLarge → any uppercase or titlecase Unicode letter
symbol → *ascSymbol* | $uniSymbol_{\langle special \mid \_ \mid " \mid ' \rangle}$
ascSymbol → `!` | `#` | `$` | `%` | `&` | `*` | `+` | `.` | `/` | `<` | `=` | `>` | `?` | `@`
| `\` | `^` | `|` | `-` | `~` | `:`
uniSymbol → any Unicode symbol or punctuation
digit → *ascDigit* | *uniDigit*
ascDigit → `0` | `1` | ... | `9`
uniDigit → any Unicode decimal digit
octit → `0` | `1` | ... | `7`
hexit → *digit* | `A` | ... | `F` | `a` | ... | `f`

varid → $(small\ \{small \mid large \mid digit \mid '\})_{\langle reservedid \rangle}$
conid → *large* {*small* | *large* | *digit* | `'`}
reservedid → `case` | `class` | `data` | `default` | `deriving` | `do` | `else`
| `foreign` | `if` | `import` | `in` | `infix` | `infixl`
| `infixr` | `instance` | `let` | `module` | `newtype` | `of`
| `then` | `type` | `where` | `_`

varsym → $\left(symbol_{\langle : \rangle} \{ symbol \} \right)_{\langle reservedop \mid dashes \rangle}$

$$
\begin{array}{lcl}
\mathit{consym} & \to & (\texttt{:}\ \{\mathit{symbol}\})_{\langle \mathit{reservedop} \rangle} \\
\mathit{reservedop} & \to & \texttt{..} \mid \texttt{:} \mid \texttt{::} \mid \texttt{=} \mid \texttt{\textbackslash} \mid \texttt{|} \mid \texttt{<-} \mid \texttt{->} \mid \texttt{@} \mid \texttt{\~{}} \mid \texttt{=>}
\end{array}
$$

$$
\begin{array}{lcll}
\mathit{varid} & & & \text{(variables)} \\
\mathit{conid} & & & \text{(constructors)} \\
\mathit{tyvar} & \to & \mathit{varid} & \text{(type variables)} \\
\mathit{tycon} & \to & \mathit{conid} & \text{(type constructors)} \\
\mathit{tycls} & \to & \mathit{conid} & \text{(type classes)} \\
\mathit{modid} & \to & \{\mathit{conid}\ \texttt{.}\}\ \mathit{conid} & \text{(modules)}
\end{array}
$$

$$
\begin{array}{lcl}
\mathit{qvarid} & \to & [\mathit{modid}\ \texttt{.}]\ \mathit{varid} \\
\mathit{qconid} & \to & [\mathit{modid}\ \texttt{.}]\ \mathit{conid} \\
\mathit{qtycon} & \to & [\mathit{modid}\ \texttt{.}]\ \mathit{tycon} \\
\mathit{qtycls} & \to & [\mathit{modid}\ \texttt{.}]\ \mathit{tycls} \\
\mathit{qvarsym} & \to & [\mathit{modid}\ \texttt{.}]\ \mathit{varsym} \\
\mathit{qconsym} & \to & [\mathit{modid}\ \texttt{.}]\ \mathit{consym}
\end{array}
$$

$$
\begin{array}{lcl}
\mathit{decimal} & \to & \mathit{digit}\{\mathit{digit}\} \\
\mathit{octal} & \to & \mathit{octit}\{\mathit{octit}\} \\
\mathit{hexadecimal} & \to & \mathit{hexit}\{\mathit{hexit}\} \\
\mathit{integer} & \to & \mathit{decimal} \\
 & \mid & \texttt{0o}\ \mathit{octal} \mid \texttt{0O}\ \mathit{octal} \\
 & \mid & \texttt{0x}\ \mathit{hexadecimal} \mid \texttt{0X}\ \mathit{hexadecimal} \\
\mathit{float} & \to & \mathit{decimal}\ \texttt{.}\ \mathit{decimal}\ [\mathit{exponent}] \\
 & \mid & \mathit{decimal}\ \mathit{exponent} \\
\mathit{exponent} & \to & (\texttt{e} \mid \texttt{E})[\texttt{+} \mid \texttt{-}]\ \mathit{decimal}
\end{array}
$$

$$
\begin{array}{lcl}
\mathit{char} & \to & \texttt{'}\ \left(\mathit{graphic}_{\langle \texttt{'} \mid \texttt{\textbackslash} \rangle} \mid \mathit{space} \mid \mathit{escape}_{\langle \texttt{\textbackslash\&} \rangle}\right)\ \texttt{'} \\
\mathit{string} & \to & \texttt{"}\ \left\{\mathit{graphic}_{\langle \texttt{"} \mid \texttt{\textbackslash} \rangle} \mid \mathit{space} \mid \mathit{escape} \mid \mathit{gap}\right\}\ \texttt{"} \\
\mathit{escape} & \to & \texttt{\textbackslash}\ (\mathit{charesc} \mid \mathit{ascii} \mid \mathit{decimal} \mid \texttt{o}\ \mathit{octal} \mid \texttt{x}\ \mathit{hexadecimal}) \\
\mathit{charesc} & \to & \texttt{a} \mid \texttt{b} \mid \texttt{f} \mid \texttt{n} \mid \texttt{r} \mid \texttt{t} \mid \texttt{v} \mid \texttt{\textbackslash} \mid \texttt{"} \mid \texttt{'} \mid \texttt{\&} \\
\mathit{ascii} & \to & \texttt{\^{}}\ \mathit{cntrl} \mid \texttt{NUL} \mid \texttt{SOH} \mid \texttt{STX} \mid \texttt{ETX} \mid \texttt{EOT} \mid \texttt{ENQ} \mid \texttt{ACK} \\
 & \mid & \texttt{BEL} \mid \texttt{BS} \mid \texttt{HT} \mid \texttt{LF} \mid \texttt{VT} \mid \texttt{FF} \mid \texttt{CR} \mid \texttt{SO} \mid \texttt{SI} \mid \texttt{DLE} \\
 & \mid & \texttt{DC1} \mid \texttt{DC2} \mid \texttt{DC3} \mid \texttt{DC4} \mid \texttt{NAK} \mid \texttt{SYN} \mid \texttt{ETB} \mid \texttt{CAN} \\
 & \mid & \texttt{EM} \mid \texttt{SUB} \mid \texttt{ESC} \mid \texttt{FS} \mid \texttt{GS} \mid \texttt{RS} \mid \texttt{US} \mid \texttt{SP} \mid \texttt{DEL} \\
\mathit{cntrl} & \to & \mathit{ascLarge} \mid \texttt{@} \mid \texttt{[} \mid \texttt{\textbackslash} \mid \texttt{]} \mid \texttt{\^{}} \mid \texttt{\_} \\
\mathit{gap} & \to & \texttt{\textbackslash}\ \mathit{whitechar}\ \{\ \mathit{whitechar}\ \}\ \texttt{\textbackslash}
\end{array}
$$

10.3 Layout

Section 2.7 gives an informal discussion of the layout rule. This section defines it more precisely.

The meaning of a Haskell program may depend on its *layout*. The effect of layout on its meaning can be completely described by adding braces and semicolons in places determined by the layout. The meaning of this augmented program is now layout insensitive.

The effect of layout is specified in this section by describing how to add braces and semicolons to a laid-out program. The specification takes the form of a function L that performs the translation. The input to L is:

- A stream of lexemes as specified by the lexical syntax in the Haskell report, with the following additional tokens:
  - If a `let`, `where`, `do`, or `of` keyword is not followed by the lexeme `{`, the token `{n}` is inserted after the keyword, where n is the indentation of the next lexeme if there is one, or 0 if the end of file has been reached.
  - If the first lexeme of a module is not `{` or `module`, then it is preceded by `{n}` where n is the indentation of the lexeme.
  - Where the start of a lexeme is preceded only by white space on the same line, this lexeme is preceded by $<n>$ where n is the indentation of the lexeme, provided that it is not, as a consequence of the first two rules, preceded by `{n}`. (NB: a string literal may span multiple lines – Section 2.6. So in the fragment

    ```
    f = ("Hello \
          \Bill", "Jake")
    ```

    There is no $<n>$ inserted before the `\Bill`, because it is not the beginning of a complete lexeme; nor before the `,`, because it is not preceded only by white space.)
- A stack of "layout contexts", in which each element is either:
  - Zero, indicating that the enclosing context is explicit (i.e. the programmer supplied the opening brace). If the innermost context is 0, then no layout tokens will be inserted until either the enclosing context ends or a new context is pushed.
  - A positive integer, which is the indentation column of the enclosing layout context.

The "indentation" of a lexeme is the column number of the first character of that lexeme; the indentation of a line is the indentation of its leftmost lexeme. To determine the column number, assume a fixed-width font with the following conventions:

- The characters *newline*, *return*, *linefeed*, and *formfeed*, all start a new line.
- The first column is designated column 1, not 0.
- Tab stops are 8 characters apart.
- A tab character causes the insertion of enough spaces to align the current position with the next tab stop.

For the purposes of the layout rule, Unicode characters in a source program are considered to be of the same, fixed, width as an ASCII character. However, to avoid visual confusion, programmers should avoid writing programs in which the meaning of implicit layout depends on the width of non-space characters.

The application

```
L tokens []
```

delivers a layout-insensitive translation of *tokens*, where *tokens* is the result of lexically analysing a module and adding column-number indicators to it as described above. The definition of L is as follows, where we use "`:`" as a stream construction operator, and "`[]`" for the empty stream.

```
L(<n>:ts)(m:ms)    =  ; : (L ts (m : ms))          if m = n
                   =  } : (L(<n>:ts) ms)           if n < m
L(<n>:ts) ms       =  L ts ms
L({n} : ts)(m:ms)  =  { : (L ts (n : m : ms)       if n > m (Note 1)
L({n} : ts)[]      =  { : (L ts [n])               if n > 0 (Note 1)
L({n} : ts) ms     =  { : } : (L (<n> : ts) ms)    (Note 2)
L(} : ts)(0:ms)    =  } : (L ts ms)                (Note 3)
L(} : ts) ms       =  parse-error                  (Note 3)
L({ : ts) ms       =  { : (L ts (0 : ms))          (Note 4)
L(t : ts)(m:ms)    =  } : (L (t:ts) ms)            if m /= 0 and parse-error(t) (Note 5)
L(t : ts) ms       =  t : (L ts ms)
L [] []            =  []
L [] (m : ms)      =  } : (L [] ms)                if m ≠ 0 (Note 6)
```

Note 1. A nested context must be further indented than the enclosing context ($n > m$). If not, L fails, and the compiler should indicate a layout error. An example is:

```
f x = let
          h y = let
  p z = z
                in p
      in h
```

Here, the definition of `p` is indented less than the indentation of the enclosing context, which is set in this case by the definition of `h`.

Note 2. If the first token after a `where` (say) is not indented more than the enclosing layout context, then the block must be empty, so empty braces are inserted. The $\{n\}$ token is replaced by $< n >$, to mimic the situation if the empty braces had been explicit.

Note 3. By matching against 0 for the current layout context, we ensure that an explicit close brace can only match an explicit open brace. A parse error results if an explicit close brace matches an implicit open brace.

Note 4. This clause means that all brace pairs are treated as explicit layout contexts, including labelled construction and update (Section 3.15). This is a difference between this formulation and Haskell 1.4.

Note 5. The side condition *parse-error* (t) is to be interpreted as follows: if the tokens generated so far by L together with the next token t represent an invalid prefix of the Haskell grammar, and the tokens generated so far by L followed by the token "`}`" represent a valid prefix of the Haskell grammar, then *parse-error*(t) is true.

The test `m /= 0` checks that an implicitly-added closing brace would match an implicit open brace.

Note 6. At the end of the input, any pending close-braces are inserted. It is an error at this point to be within a non-layout context (i.e. $m = 0$).

If none of the rules given above matches, then the algorithm fails. It can fail for instance when the end of the input is reached, and a non-layout context is active, since the close brace is missing. Some error conditions are not detected by the algorithm, although they could be: for example `let }`.

Note 1 implements the feature that layout processing can be stopped prematurely by a parse error. For example

```
let x = e; y = x in e'
```

is valid, because it translates to

```
let { x = e; y = x } in e'
```

The close brace is inserted due to the parse error rule above.

10.4 Literate comments

The "literate comment" convention, first developed by Richard Bird and Philip Wadler for Orwell, and inspired in turn by Donald Knuth's "literate programming", is an alternative style for encoding Haskell source code. The literate style encourages comments by making them the default. A line in which "`>`" is the first character is treated as part of the program; all other lines are comments.

The program text is recovered by taking only those lines beginning with "`>`", and replacing the leading "`>`" with a space. Layout and comments apply exactly as described in Chapter 10 in the resulting text.

To capture some cases where one omits an "`>`" by mistake, it is an error for a program line to appear adjacent to a non-blank comment line, where a line is taken as blank if it consists only of whitespace.

By convention, the style of comment is indicated by the file extension, with "`.hs`" indicating a usual Haskell file and "`.lhs`" indicating a literate Haskell file. Using this style, a simple factorial program would be:

```
   This literate program prompts the user for a number
   and prints the factorial of that number:

> main :: IO ()

> main = do putStr "Enter a number: "
>           l <- readLine
>           putStr "n!= "
>           print (fact (read l))

  This is the factorial function.

> fact :: Integer -> Integer
> fact 0 = 1
> fact n = n * fact (n-1)
```

An alternative style of literate programming is particularly suitable for use with the LaTeX text processing system. In this convention, only those parts of the literate program that are entirely enclosed between `\begin{code}...\end{code}` delimiters are treated as program text; all other lines are comments. More precisely:

- Program code begins on the first line following a line that begins `\begin{code}`.

- Program code ends just before a subsequent line that begins `\end{code}` (ignoring string literals, of course).

It is not necessary to insert additional blank lines before or after these delimiters, though it may be stylistically desirable. For example,

```
\documentstyle{article}

\begin{document}

\chapter{Introduction}

This is a trivial program that prints the first 20 factorials.

\begin{code}
main :: IO ()
main =  print [ (n, product [1..n]) | n <- [1..20]]
\end{code}

\end{document}
```

This style uses the same file extension. It is not advisable to mix these two styles in the same file.

10.5 Context-Free Syntax

module	→	`module` *modid* [*exports*] `where` *body*	
	|	*body*	
body	→	{ *impdecls* ; *topdecls* }	
	|	{ *impdecls* }	
	|	{ *topdecls* }	
impdecls	→	$impdecl_1$; ... ; $impdecl_n$	$(n \geq 1)$
exports	→	($export_1$, ... , $export_n$ [,])	$(n \geq 0)$
export	→	*qvar*	
	|	*qtycon* [(..) | ($cname_1$, ... , $cname_n$)]	$(n \geq 0)$
	|	*qtycls* [(..) | (var_1 , ... , var_n)]	$(n \geq 0)$
	|	`module` *modid*	
impdecl	→	`import` [`qualified`] *modid* [`as` *modid*][*impspec*]	
	|		(empty declaration)
impspec	→	($import_1$, ... , $import_n$ [,])	$(n \geq 0)$
	|	`hiding` ($import_1$, ... , $import_n$ [,])	$(n \geq 0)$
import	→	*var*	
	|	*tycon* [(..) | ($cname_1$, ... , $cname_n$)]	$(n \geq 0)$
	|	*tycls* [(..) | (var_1 , ... , var_n)]	$(n \geq 0)$
cname	→	*var* | *con*	
topdecls	→	$topdecl_1$; ... ; $topdecl_n$	$(n \geq 1)$

topdecl	→	`type` *simpletype* = *type*	
	|	`data` [*context* =>] *simpletype* [= *constrs*][*deriving*]	
	|	`newtype` [*context* =>] *simpletype* = *newconstr* [*deriving*]	
	|	`class` [*scontext* =>] *tycls tyvar* [`where` *cdecls*]	
	|	`instance` [*scontext* =>] *qtycls inst* [`where` *idecls*]	
	|	`default` ($type_1$, ... , $type_n$)	$(n \geq 0)$
	|	`foreign` *fdecl*	
	|	*decl*	
decls	→	{ $decl_1$; ... ; $decl_n$ }	$(n \geq 0)$
decl	→	*gendecl*	
	|	(*funlhs* | *pat*) *rhs*	
cdecls	→	{ $cdecl_1$; ... ; $cdecl_n$ }	$(n \geq 0)$
cdecl	→	*gendecl*	
	|	(*funlhs* | *var*) *rhs*	
idecls	→	{ $idecl_1$; ... ; $idecl_n$ }	$(n \geq 0)$
idecl	→	(*funlhs* | *var*) *rhs*	
	|		(empty)
gendecl	→	*vars* :: [*context* =>] *type*	(type signature)
	|	*fixity* [*integer*] *ops*	(fixity declaration)
	|		(empty declaration)
ops	→	op_1 , ... , op_n	$(n \geq 1)$
vars	→	var_1 , ... , var_n	$(n \geq 1)$
fixity	→	`infixl` | `infixr` | `infix`	
type	→	*btype* [-> *type*]	(function type)
btype	→	[*btype*] *atype*	(type application)
atype	→	*gtycon*	
	|	*tyvar*	
	|	($type_1$, ... , $type_k$)	(tuple type, $k \geq 2$)
	|	[*type*]	(list type)
	|	(*type*)	(parenthesized constructor)
gtycon	→	*qtycon*	
	|	()	(unit type)
	|	[]	(list constructor)
	|	(->)	(function constructor)
	|	(, {,})	(tupling constructors)
context	→	*class*	
	|	($class_1$, ... , $class_n$)	$(n \geq 0)$
class	→	*qtycls tyvar*	
	|	*qtycls* (*tyvar* $atype_1$...$atype_n$)	$(n \geq 1)$

scontext	→	*simpleclass*	
	|	(*simpleclass*$_1$, ... , *simpleclass*$_n$)	$(n \geq 0)$
simpleclass	→	*qtycls tyvar*	
simpletype	→	*tycon tyvar*$_1$...*tyvar*$_k$	$(k \geq 0)$
constrs	→	*constr*$_1$ | ... | *constr*$_n$	$(n \geq 1)$
constr	→	*con* [!]*atype*$_1$...[!]*atype*$_k$	(arity *con* $= k, k \geq 0$)
	|	(*btype* | ! *atype*) *conop* (*btype* | ! *atype*)	(infix *conop*)
	|	*con* { *fielddecl*$_1$, ... , *fielddecl*$_n$ }	$(n \geq 0)$
newconstr	→	*con atype*	
	|	*con* { *var* :: *type* }	
fielddecl	→	*vars* :: (*type* | ! *atype*)	
deriving	→	`deriving` (*dclass* | (*dclass*$_1$, ... , *dclass*$_n$))$(n \geq 0)$	
dclass	→	*qtycls*	
inst	→	*gtycon*	
	|	(*gtycon tyvar*$_1$...*tyvar*$_k$)	$(k \geq 0$, *tyvars* distinct)
	|	(*tyvar*$_1$, ... , *tyvar*$_k$)	$(k \geq 2$, *tyvars* distinct)
	|	[*tyvar*]	
	|	(*tyvar*$_1$ -> *tyvar*$_2$)	(*tyvar*$_1$ and *tyvar*$_2$ distinct)
fdecl	→	`import` *callconv* [*safety*] *impent var* :: *ftype*	(define variable)
	|	`export` *callconv expent var* :: *ftype*	(expose variable)
callconv	→	`ccall` | `stdcall` | `cplusplus`	(calling convention)
	|	`jvm` | `dotnet`	
	|	**system-specific calling conventions**	
impent	→	[*string*]	See Section X
expent	→	[*string*]	See Section X
safety	→	`unsafe` | `safe`	
ftype	→	*frtype*	
	|	*fatype* -> *ftype*	
frtype	→	*fatype*	
	|	()	
fatype	→	*qtycon atype*$_1$...*atype*$_k$	$(k \geq 0)$
funlhs	→	*var apat* {*apat*}	
	|	*pat varop pat*	
	|	(*funlhs*) *apat* {*apat*}	
rhs	→	= *exp* [`where` *decls*]	
	|	*gdrhs* [`where` *decls*]	
gdrhs	→	*guards* = *exp* [*gdrhs*]	
guards	→	| *guard*$_1$, ... , *guard*$_n$	$(n \geq 1)$
guard	→	*pat* <- *infixexp*	(pattern guard)

	| `let` *decls*	(local declaration)
	| *infixexp*	
exp	→ *infixexp* `::` [*context* `=>`] *type*	(expression type signature)
	| *infixexp*	
infixexp	→ *lexp qop infixexp*	
	| `-` *infixexp*	(prefix negation)
	| *lexp*	
lexp	→ `\` $apat_1 ... apat_n$ `->` *exp*	(lambda abstraction, $n \geq 1$)
	| `let` *decls* `in` *exp*	(let expression)
	| `if` *exp* [`;`] `then` *exp* [`;`] `else` *exp*	(conditional)
	| `case` *exp* `of` `{` *alts* `}`	(case expression)
	| `do` `{` *stmts* `}`	(do expression)
	| *fexp*	
fexp	→ [*fexp*] *aexp*	(function application)
aexp	→ *qvar*	(variable)
	| *gcon*	(general constructor)
	| *literal*	
	| `(` *exp* `)`	(parenthesized expression)
	| `(` exp_1 `,` ... `,` exp_k `)`	(tuple, $k \geq 2$)
	| `[` exp_1 `,` ... `,` exp_k `]`	(list, $k \geq 1$)
	| `[` exp_1[`,` exp_2] `..` [exp_3] `]`	(arithmetic sequence)
	| `[` *exp* `|` $qual_1$ `,` ... `,` $qual_n$ `]`	(list comprehension, $n \geq 1$)
	| `(` *infixexp qop* `)`	(left section)
	| `(` $qop_{\langle - \rangle}$ *infixexp* `)`	(right section)
	| *qcon* `{` $fbind_1$ `,` ... `,` $fbind_n$ `}`	(labeled construction, $n \geq 0$)
	| $aexp_{\langle qcon \rangle}$ `{` $fbind_1$ `,` ... `,` $fbind_n$ `}`	(labeled update, $n \geq 1$)
qual	→ *pat* `<-` *exp*	(generator)
	| `let` *decls*	(local declaration)
	| *exp*	(guard)
alts	→ alt_1 `;` ... `;` alt_n	($n \geq 1$)
alt	→ *pat* `->` *exp* [`where` *decls*]	
	| *pat gdpat* [`where` *decls*]	
	|	(empty alternative)
gdpat	→ *guards* `->` *exp* [*gdpat*]	
stmts	→ $stmt_1 ... stmt_n$ *exp* [`;`]	($n \geq 0$)
stmt	→ *exp* `;`	
	| *pat* `<-` *exp* `;`	
	| `let` *decls* `;`	
	| `;`	(empty statement)

fbind	→	*qvar* = *exp*	
pat	→	*lpat qconop pat*	(infix constructor)
	|	*lpat*	
lpat	→	*apat*	
	|	- (*integer* | *float*)	(negative literal)
	|	*gcon* $apat_1 ... apat_k$	(arity *gcon* $= k, k \geq 1$)
apat	→	*var* [@ *apat*]	(as pattern)
	|	*gcon*	(arity *gcon* $= 0$)
	|	*qcon* { $fpat_1$, ... , $fpat_k$ }	(labeled pattern, $k \geq 0$)
	|	*literal*	
	|	_	(wildcard)
	|	(*pat*)	(parenthesized pattern)
	|	(pat_1 , ... , *pat*)	(tuple pattern, $k \geq 2$)
	|	[pat_1 , ... , *pat*]	(list pattern, $k \geq 1$)
	|	~ *apat*	(irrefutable pattern)
fpat	→	*qvar* = *pat*	
gcon	→	()	
	|	[]	
	|	(, {,})	
	|	*qcon*	
var	→	*varid* | (*varsym*)	(variable)
qvar	→	*qvarid* | (*qvarsym*)	(qualified variable)
con	→	*conid* | (*consym*)	(constructor)
qcon	→	*qconid* | (*qconsym*)	(qualified constructor)
varop	→	*varsym* | ` *varid* `	(variable operator)
qvarop	→	*qvarsym* | ` *qvarid* `	(qualified variable operator)
conop	→	*consym* | ` *conid* `	(constructor operator)
qconop	→	*gconsym* | ` *qconid* `	(qualified constructor operator)
op	→	*varop* | *conop*	(operator)
qop	→	*qvarop* | *qconop*	(qualified operator)
gconsym	→	: | *qconsym*	

10.6 Fixity Resolution

The following is an example implementation of fixity resolution for Haskell expressions. Fixity resolution also applies to Haskell patterns, but patterns are a subset of expressions so in what follows we consider only expressions for simplicity.

The function `resolve` takes a list in which the elements are expressions or operators, i.e. an instance of the *infixexp* non-terminal in the context-free grammar. It returns either `Just e` where `e` is the

resolved expression, or `Nothing` if the input does not represent a valid expression. In a compiler, of course, it would be better to return more information about the operators involved for the purposes of producing a useful error message, but the `Maybe` type will suffice to illustrate the algorithm here.

```
import Control.Monad

type Prec   = Int
type Var    = String

data Op = Op String Prec Fixity
  deriving (Eq,Show)

data Fixity = Leftfix | Rightfix | Nonfix
  deriving (Eq,Show)

data Exp = Var Var | OpApp Exp Op Exp | Neg Exp
  deriving (Eq,Show)

data Tok = TExp Exp | TOp Op | TNeg
  deriving (Eq,Show)

resolve :: [Tok] -> Maybe Exp
resolve tokens = fmap fst $ parseNeg (Op "" (-1) Nonfix) tokens
  where
    parseNeg :: Op -> [Tok] -> Maybe (Exp,[Tok])
    parseNeg op1 (TExp e1 : rest)
       = parse op1 e1 rest
    parseNeg op1 (TNeg : rest)
       = do guard (prec1 < 6)
            (r, rest') <- parseNeg (Op "-" 6 Leftfix) rest
            parse op1 (Neg r) rest'
       where
          Op _ prec1 fix1 = op1

    parse :: Op -> Exp -> [Tok] -> Maybe (Exp, [Tok])
    parse _   e1 [] = Just (e1, [])
    parse op1 e1 (TOp op2 : rest)
       -- case (1): check for illegal expressions
       | prec1 == prec2 && (fix1 /= fix2 || fix1 == Nonfix)
       = Nothing

       -- case (2): op1 and op2 should associate to the left
       | prec1 > prec2 || (prec1 == prec2 && fix1 == Leftfix)
       = Just (e1, TOp op2 : rest)

       -- case (3): op1 and op2 should associate to the right
       | otherwise
       = do (r,rest') <- parseNeg op2 rest
            parse op1 (OpApp e1 op2 r) rest'
       where
         Op _ prec1 fix1 = op1
         Op _ prec2 fix2 = op2
```

The algorithm works as follows. At each stage we have a call

```
parse op1 E1 (op2 : tokens)
```

which means that we are looking at an expression like

```
E0 `op1` E1 `op2` ...      (1)
```

(the caller holds `E0`). The job of `parse` is to build the expression to the right of `op1`, returning the expression and any remaining input.

There are three cases to consider:

1. if `op1` and `op2` have the same precedence, but they do not have the same associativity, or they are declared to be nonfix, then the expression is illegal.
2. If `op1` has a higher precedence than `op2`, or `op1` and `op2` should left-associate, then we know that the expression to the right of `op1` is `E1`, so we return this to the caller.
3. Otherwise, we know we want to build an expression of the form `E1 `op2` R`. To find `R`, we call `parseNeg op2 tokens` to compute the expression to the right of `op2`, namely `R` (more about `parseNeg` below, but essentially if `tokens` is of the form `(E2 : rest)`, then this is equivalent to `parse op2 E2 rest`). Now, we have

   ```
   E0 `op1` (E1 `op2` R) `op3` ...
   ```

   where `op3` is the next operator in the input. This is an instance of (1) above, so to continue we call parse, with the new `E1 == (E1 `op2` R)`.

To initialise the algorithm, we set `op1` to be an imaginary operator with precedence lower than anything else. Hence `parse` will consume the whole input, and return the resulting expression.

The handling of the prefix negation operator, `-`, complicates matters only slightly. Recall that prefix negation has the same fixity as infix negation: left-associative with precedence 6. The operator to the left of `-`, if there is one, must have precedence lower than 6 for the expression to be legal. The negation operator itself may left-associate with operators of the same fixity (e.g. `+`). So for example `-a + b` is legal and resolves as `(-a) + b`, but `a + -b` is illegal.

The function `parseNeg` handles prefix negation. If we encounter a negation operator, and it is legal in this position (the operator to the left has precedence lower than 6), then we proceed in a similar way to case (3) above: compute the argument to `-` by recursively calling `parseNeg`, and then continue by calling `parse`.

Note that this algorithm is insensitive to the range and resolution of precedences. There is no reason in principle that Haskell should be limited to integral precedences in the range 1 to 10; a larger range, or fractional values, would present no additional difficulties.

Chapter 11

Specification of Derived Instances

A *derived instance* is an instance declaration that is generated automatically in conjunction with a `data` or `newtype` declaration. The body of a derived instance declaration is derived syntactically from the definition of the associated type. Derived instances are possible only for classes known to the compiler: those defined in either the Prelude or a standard library. In this chapter, we describe the derivation of classes defined by the Prelude.

If T is an algebraic datatype declared by:

$$\texttt{data}\ cx \Rightarrow T\ u_1 ... u_k = K_1\ t_{1,1} ... t_{1,k_1} \mid ... \mid K_n\ t_{n,1} ... t_{n,k_n}$$
$$\texttt{deriving}\ (C_1, ..., C_m)$$

(where $m \geq 0$ and the parentheses may be omitted if $m = 1$) then a derived instance declaration is possible for a class C if these conditions hold:

1. C is one of `Eq`, `Ord`, `Enum`, `Bounded`, `Show` or `Read`.
2. There is a context cx' such that $cx' \Rightarrow Ct_{ij}$ holds for each of the constituent types t_{ij}.
3. If C is `Bounded`, the type must be either an enumeration (all constructors must be nullary) or have only one constructor.
4. If C is `Enum`, the type must be an enumeration.
5. There must be no explicit instance declaration elsewhere in the program that makes $Tu_1 ... u_k$ an instance of C.
6. If the data declaration has no constructors (i.e. when $n = 0$), then no classes are derivable (i.e. $m = 0$)

For the purposes of derived instances, a `newtype` declaration is treated as a `data` declaration with a single constructor.

If the `deriving` form is present, an instance declaration is automatically generated for $Tu_1 ... u_k$ over each class C_i. If the derived instance declaration is impossible for any of the C_i then a static error results. If no derived instances are required, the `deriving` form may be omitted or the form `deriving ()` may be used.

Each derived instance declaration will have the form:

$$\texttt{instance}\ (cx, cx') \Rightarrow C_i\ (T\ u_1 ... u_k)\ \texttt{where}\ \{\ d\ \}$$

where d is derived automatically depending on C_i and the data type declaration for T (as will be described in the remainder of this section).

The context cx' is the smallest context satisfying point (2) above. For mutually recusive data types, the compiler may need to perform a fixpoint calculation to compute it.

The remaining details of the derived instances for each of the derivable Prelude classes are now given. Free variables and constructors used in these translations always refer to entities defined by the `Prelude`.

11.1 Derived instances of Eq and Ord

The class methods automatically introduced by derived instances of `Eq` and `Ord` are `(==)`, `(/=)`, `compare` `(<)`, `(<=)`, `(>)`, `(>=)`, `max`, and `min`. The latter seven operators are defined so as to compare their arguments lexicographically with respect to the constructor set given, with earlier constructors in the datatype declaration counting as smaller than later ones. For example, for the `Bool` datatype, we have that `(True > False) == True`.

Derived comparisons always traverse constructors from left to right. These examples illustrate this property:

$$(1,\texttt{undefined}) == (2,\texttt{undefined}) \Rightarrow \texttt{False}$$
$$(\texttt{undefined},1) == (\texttt{undefined},2) \Rightarrow \bot$$

All derived operations of class `Eq` and `Ord` are strict in both arguments. For example, `False <=` $\bot$ is $\bot$, even though `False` is the first constructor of the `Bool` type.

11.2 Derived instances of Enum

Derived instance declarations for the class `Enum` are only possible for enumerations (data types with only nullary constructors).

The nullary constructors are assumed to be numbered left-to-right with the indices 0 through $n-1$. The `succ` and `pred` operators give the successor and predecessor respectively of a value, under this numbering scheme. It is an error to apply `succ` to the maximum element, or `pred` to the minimum element.

The `toEnum` and `fromEnum` operators map enumerated values to and from the `Int` type; `toEnum` raises a runtime error if the `Int` argument is not the index of one of the constructors.

The definitions of the remaining methods are

```
enumFrom x           = enumFromTo x lastCon
enumFromThen x y     = enumFromThenTo x y bound
                     where
                       bound | fromEnum y >= fromEnum x = lastCon
                             | otherwise                = firstCon
enumFromTo x y       = map toEnum [fromEnum x .. fromEnum y]
enumFromThenTo x y z = map toEnum [fromEnum x, fromEnum y .. fromEnum z]
```

where `firstCon` and `lastCon` are respectively the first and last constructors listed in the `data` declaration. For example, given the datatype:

```
data  Color = Red | Orange | Yellow | Green  deriving (Enum)
```

we would have:

```
[Orange ..]         ==  [Orange, Yellow, Green]
fromEnum Yellow     ==  2
```

11.3 Derived instances of Bounded

The `Bounded` class introduces the class methods `minBound` and `maxBound`, which define the minimal and maximal elements of the type. For an enumeration, the first and last constructors listed in the `data` declaration are the bounds. For a type with a single constructor, the constructor is applied to the bounds for the constituent types. For example, the following datatype:

```
data  Pair a b = Pair a b deriving Bounded
```

would generate the following `Bounded` instance:

```
instance (Bounded a,Bounded b) => Bounded (Pair a b) where
  minBound = Pair minBound minBound
  maxBound = Pair maxBound maxBound
```

11.4 Derived instances of Read and Show

The class methods automatically introduced by derived instances of `Read` and `Show` are `showsPrec`, `readsPrec`, `showList`, and `readList`. They are used to coerce values into strings and parse strings into values.

The function `showsPrec d x r` accepts a precedence level `d` (a number from `0` to `11`), a value `x`, and a string `r`. It returns a string representing `x` concatenated to `r`. `showsPrec` satisfies the law:

$$\texttt{showsPrec d x r} +\!\!+ \texttt{s} == \texttt{showsPrec d x (r} +\!\!+ \texttt{s)}$$

The representation will be enclosed in parentheses if the precedence of the top-level constructor in `x` is less than `d`. Thus, if `d` is `0` then the result is never surrounded in parentheses; if `d` is `11` it is always surrounded in parentheses, unless it is an atomic expression (recall that function application has precedence `10`). The extra parameter `r` is essential if tree-like structures are to be printed in linear time rather than time quadratic in the size of the tree.

The function `readsPrec d s` accepts a precedence level `d` (a number from `0` to `10`) and a string `s`, and attempts to parse a value from the front of the string, returning a list of (parsed value, remaining string) pairs. If there is no successful parse, the returned list is empty. Parsing of an un-parenthesised infix operator application succeeds only if the precedence of the operator is greater than or equal to `d`.

It should be the case that

`(x,"")` is an element of `(readsPrec d (showsPrec d x ""))`

That is, `readsPrec` should be able to parse the string produced by `showsPrec`, and should deliver the value that `showsPrec` started with.

`showList` and `readList` allow lists of objects to be represented using non-standard denotations. This is especially useful for strings (lists of `Char`).

`readsPrec` will parse any valid representation of the standard types apart from strings, for which only quoted strings are accepted, and other lists, for which only the bracketed form `[...]` is accepted. See Chapter 9 for full details.

The result of `show` is a syntactically correct Haskell expression containing only constants, given the fixity declarations in force at the point where the type is declared. It contains only the constructor names defined in the data type, parentheses, and spaces. When labelled constructor fields are used, braces, commas, field names, and equal signs are also used. Parentheses are only added where needed, *ignoring associativity*. No line breaks are added. The result of `show` is readable by `read` if all component types are readable. (This is true for all instances defined in the Prelude but may not be true for user-defined instances.)

Derived instances of `Read` make the following assumptions, which derived instances of `Show` obey:

- If the constructor is defined to be an infix operator, then the derived `Read` instance will parse only infix applications of the constructor (not the prefix form).
- Associativity is not used to reduce the occurrence of parentheses, although precedence may be. For example, given

  ```
  infixr 4 :$
  data T = Int :$ T  |  NT
  ```

  then:
  - `show (1 :$ 2 :$ NT)` produces the string `"1 :$ (2 :$ NT)"`.
  - `read "1 :$ (2 :$ NT)"` succeeds, with the obvious result.
  - `read "1 :$ 2 :$ NT"` fails.
- If the constructor is defined using record syntax, the derived `Read` will parse only the record-syntax form, and furthermore, the fields must be given in the same order as the original declaration.
- The derived `Read` instance allows arbitrary Haskell whitespace between tokens of the input string. Extra parentheses are also allowed.

The derived `Read` and `Show` instances may be unsuitable for some uses. Some problems include:

- Circular structures cannot be printed or read by these instances.
- The printer loses shared substructure; the printed representation of an object may be much larger than necessary.
- The parsing techniques used by the reader are very inefficient; reading a large structure may be quite slow.
- There is no user control over the printing of types defined in the Prelude. For example, there is no way to change the formatting of floating point numbers.

11.5 An Example

As a complete example, consider a tree datatype:

```
data Tree a = Leaf a | Tree a :^: Tree a
       deriving (Eq, Ord, Read, Show)
```

Automatic derivation of instance declarations for `Bounded` and `Enum` are not possible, as `Tree` is not an enumeration or single-constructor datatype. The complete instance declarations for `Tree` are shown in Listing 6. Note the implicit use of default class method definitions—for example, only `<=` is defined for `Ord`, with the other class methods (`<`, `>`, `>=`, `max`, and `min`) being defined by the defaults given in the class declaration shown in Figure 3.

```
infixr 5 :^:
data Tree a =  Leaf a  |  Tree a :^: Tree a

instance (Eq a) => Eq (Tree a) where
        Leaf m == Leaf n  =  m==n
        u:^:v  == x:^:y   =  u==x && v==y
             _ == _       =  False

instance (Ord a) => Ord (Tree a) where
        Leaf m <= Leaf n  =  m<=n
        Leaf m <= x:^:y   =  True
        u:^:v  <= Leaf n  =  False
        u:^:v  <= x:^:y   =  u<x || u==x && v<=y

instance (Show a) => Show (Tree a) where

        showsPrec d (Leaf m) = showParen (d > app_prec) showStr
          where
             showStr = showString "Leaf " . showsPrec (app_prec+1) m

        showsPrec d (u :^: v) = showParen (d > up_prec) showStr
          where
             showStr = showsPrec (up_prec+1) u .
                       showString " :^: "      .
                       showsPrec (up_prec+1) v
                -- Note: right-associativity of :^: ignored

instance (Read a) => Read (Tree a) where

        readsPrec d r =  readParen (d > up_prec)
                         (\r -> [(u:^:v,w) |
                                 (u,s) <- readsPrec (up_prec+1) r,
                                 (":^:",t) <- lex s,
                                 (v,w) <- readsPrec (up_prec+1) t]) r

                      ++ readParen (d > app_prec)
                         (\r -> [(Leaf m,t) |
                                 ("Leaf",s) <- lex r,
                                 (m,t) <- readsPrec (app_prec+1) s]) r

up_prec  = 5    -- Precedence of :^:
app_prec = 10   -- Application has precedence one more than
                -- the most tightly-binding operator
```

Listing 6: Example of Derived Instances

Chapter 12

Compiler Pragmas

Some compiler implementations support compiler *pragmas*, which are used to give additional instructions or hints to the compiler, but which do not form part of the Haskell language proper and do not change a program's semantics. This chapter summarizes this existing practice. An implementation is not required to respect any pragma, although pragmas that are not recognised by the implementation should be ignored. Implementations are strongly encouraged to support the LANGUAGE pragma described below as there are many language extensions being used in practice.

Lexically, pragmas appear as comments, except that the enclosing syntax is `{-# #-}`.

12.1 Inlining

$$decl \to \{\text{-\# } \texttt{INLINE} \;\; qvars \; \text{\#-}\}$$
$$decl \to \{\text{-\# } \texttt{NOINLINE} \;\; qvars \; \text{\#-}\}$$

The `INLINE` pragma instructs the compiler to inline the specified variables at their use sites. Compilers will often automatically inline simple expressions. This may be prevented by the `NOINLINE` pragma.

12.2 Specialization

$$decl \to \{\text{-\# } \texttt{SPECIALIZE} \;\; spec_1, ..., spec_k \; \text{\#-}\} \; (k \geq 1)$$
$$spec \to vars :: type$$

Specialization is used to avoid inefficiencies involved in dispatching overloaded functions. For example, in

```
factorial :: Num a => a -> a
factorial 0 = 0
factorial n = n * factorial (n-1)
{-# SPECIALIZE factorial :: Int -> Int,
               factorial :: Integer -> Integer #-}
```

calls to `factorial` in which the compiler can detect that the parameter is either `Int` or `Integer` will use specialized versions of `factorial` which do not involve overloaded numeric operations.

12.3 Language extensions

The `LANGUAGE` pragma is a file-header pragma. A file-header pragma must precede the module keyword in a source file. There can be as many file-header pragmas as you please, and they can be preceded or followed by comments. An individual language pragma begins with the keyword `LANGUAGE` and is followed by a comma-separated list of named language features.

For example, to enable scoped type variables and preprocessing with CPP, if your Haskell implementation supports these extensions:

```
{-# LANGUAGE ScopedTypeVariables, CPP #-}
```

If a Haskell implementation does not recognize or support a particular language feature that a source file requests (or cannot support the combination of language features requested), any attempt to compile or otherwise use that file with that Haskell implementation must fail with an error.

In the interests of portability, multiple attempts to enable the same, supported language features (e.g. via command-line arguments, implementation-specific features dependencies or non-standard pragmas) are specifically permitted. Haskell 2010 implementations that support the `LANGUAGE` pragma are required to support

```
{-# LANGUAGE Haskell2010 #-}
```

Those implementations are also encouraged to support the following named language features:

```
PatternGuards, NoNPlusKPatterns, RelaxedPolyRec,
EmptyDataDecls, ForeignFunctionInterface
```

These are the named language extensions supported by some pre-Haskell 2010 implementations, that have been integrated into this report.

Bibliography

[1] H. B. Curry and R. Feys, *Combinatory Logic*. North-Holland Pub. Co., Amsterdam, 1958.

[2] J. Backus, "Can programming be liberated from the von Neumann style? A functional style and its algebra of programs.," *CACM*, vol. 21, no. 8, pp. 613–641, Aug. 1978.

[3] Unicode Consortium, "Unicode Standard." [Online]. Available: http://unicode.org/standard/standard.html

[4] R. Hindley, "The Principal Type Scheme of an Object in Combinatory Logic," *Transactions of the American Mathematical Society*, vol. 146, pp. 29–60, Dec. 1969.

[5] L. Damas and R. Milner, "Principal Type-Schemes for Functional Programs," in *Conference Record of the 9th Annual ACM Symposium on Principles of Programming Languages*, New York: ACM Press, 1982, pp. 207–212.

[6] M. Jones, "A system of constructor classes: overloading and implicit higher-order polymorphism," *Journal of Functional Programming*, vol. 5, no. 1, pp. 1–36, Jan. 1995.

[7] P. Wadler and S. Blott, "How to Make \em ad hoc Polymorphism Less \em ad hoc," in *Proceedings of 16th ACM Symposium on Principles of Programming Languages*, Austin, Texas, Jan. 1989, pp. 60–76.

[8] P. Penfield Jr., "Principal Values and Branch Cuts in Complex APL," in *APL '81 Conference Proceedings*, San Francisco, Sep. 1981, pp. 248–256.

[9] B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Second. Prentice Hall, 1988.

[10] J. Gosling, B. Joy, and G. Steele, *The Java Language Specification*. in The Java Series. Addison-Wesley, 1997.

[11] T. Lindholm and F. Yellin, *The Java Virtual Machine Specification*. Addison-Wesley, 1996.

[12] S. Liang, *The Java Native Interface: Programmer's Guide and Specification*. Addison Wesley, 1999.

[13] I. S. ISO/IEC, "Programming Languages – C."