

The Scala Language Specification

Version 2.0

DRAFT
October 24, 2006

Martin Odersky

PROGRAMMING METHODS LABORATORY
EPFL
SWITZERLAND

Contents

1	Lexical Syntax	3
1.1	Identifiers	4
1.2	Newline Characters	5
1.3	Literals	6
1.3.1	Integer Literals	6
1.3.2	Floating Point Literals	7
1.3.3	Boolean Literals	7
1.3.4	Character Literals	7
1.3.5	String Literals	8
1.3.6	Escape Sequences	8
1.3.7	Symbol literals	9
1.3.8	Primitive Types Defined in <i>Predef</i>	9
1.4	Whitespace and Comments	9
1.5	XML mode	10
2	Identifiers, Names and Scopes	11
3	Types	13
3.1	Paths	13
3.2	Value Types	14
3.2.1	Singleton Types	14
3.2.2	Type Projection	14
3.2.3	Type Designators	15
3.2.4	Parameterized Types	15
3.2.5	Compound Types	16
3.2.6	Function Types	16
3.3	Non-Value Types	17
3.3.1	Method Types	17

3.3.2	Polymorphic Method Types	18
3.4	Base Types and Member Definitions	18
3.5	Relations between types	20
3.5.1	Type Equivalence	20
3.5.2	Conformance	21
3.6	Type Erasure	22
4	Basic Declarations and Definitions	23
4.1	Value Declarations and Definitions	23
4.2	Variable Declarations and Definitions	25
4.3	Type Declarations and Type Aliases	26
4.4	Type Parameters	28
4.5	Variance Annotations	29
4.6	Function Declarations and Definitions	30
4.7	Import Clauses	32
5	Classes and Objects	35
5.1	Templates	35
5.1.1	Constructor Invocations	37
5.1.2	Class Linearization	37
5.1.3	Class Members	38
5.1.4	Overriding	39
5.2	Modifiers	40
5.3	Class Definitions	43
5.3.1	Constructor Definitions	44
5.3.2	Case Classes	45
5.3.3	Traits	47
5.4	Object Definitions	49
6	Expressions	51
6.1	Literals	52
6.2	The <i>Null</i> Value	52
6.3	Designators	53
6.4	This and Super	53

6.5	Function Applications	54
6.6	Type Applications	55
6.7	Instance Creation Expressions	55
6.8	Blocks	56
6.9	Prefix, Infix, and Postfix Operations	57
6.10	Typed Expressions	58
6.11	Method closures	58
6.12	Assignments	59
6.13	Conditional Expressions	60
6.14	While Loop Expressions	61
6.15	Do Loop Expressions	61
6.16	For-Comprehensions	61
6.17	Return Expressions	64
6.18	Throw Expressions	64
6.19	Try Expressions	64
6.20	Anonymous Functions	65
6.21	Statements	66
6.22	Implicit Conversions	66
6.22.1	Value Conversions	67
6.22.2	Method Conversions	67
6.22.3	Overloading Resolution	68
6.22.4	Local Type Inference	69
6.22.5	Eta Expansion	72
7	Implicit Parameters and Views	73
7.1	The Implicit Modifier	73
7.2	Implicit Parameters	74
7.3	Views	76
7.4	View Bounds	77
8	Pattern Matching	79
8.1	Patterns	79
8.1.1	Regular Pattern Matching	81
8.2	Pattern Matching Expressions	81

8.3	Pattern Matching Anonymous Functions	83
9	Top-Level Definitions	85
9.1	Compilation Units	85
9.2	Packagings	85
9.3	Programs	86
10	XML expressions and patterns	89
10.1	XML expressions	89
10.2	XML patterns	91
11	Attributes	93
12	The Scala Standard Library	95
12.1	Root Classes	95
12.2	Value Classes	97
12.2.1	Numeric Value Types	97
12.2.2	Class Boolean	100
12.2.3	Class Unit	100
12.3	Standard Reference Classes	101
12.3.1	Class String	101
12.3.2	The Tuple classes	101
12.3.3	The Function Classes	101
12.3.4	Class Array	102
12.4	Class Node	104
12.5	The Predef Object	106
A	Scala Syntax Summary	113
B	Changes between Scala version 1.0 and 2.0	119
B.1	New Keywords	119
B.2	Newlines as Statement Separators	119
B.3	Syntax Restrictions	119
B.4	Selftype Annotations	120
B.5	For-comprehensions	121

B.6	Conversions	121
B.7	Class Parameters	122
B.8	Private Qualifiers	122
B.9	Changes in the Mixin Model	123
B.10	Implicit Parameters	123
B.11	Flexible Typing of Pattern Matching	123

Preface

Scala is a Java-like programming language which unifies object-oriented and functional programming. It is a pure object-oriented language in the sense that every value is an object. Types and behavior of objects are described by classes. Classes can be composed using mixin composition. Scala is designed to work seamlessly with two less pure but mainstream object-oriented languages – Java and C#.

Scala is a functional language in the sense that every function is a value. Nesting of function definitions and higher-order functions are naturally supported. Scala also supports a general notion of pattern matching which can model the algebraic types used in many functional languages.

Scala has been designed to interoperate seamlessly with Java (an alternative implementation of Scala also works for .NET). Scala classes can call Java methods, create Java objects, inherit from Java classes and implement Java interfaces. None of this requires interface definitions or glue code.

Scala has been developed from 2001 in the programming methods laboratory at EPFL. Version 1.0 was released in November 2003. This document describes the second version of the language, which was released in March 2006. It acts a reference for the language definition and some core library modules. It is not intended to teach Scala or its concepts; for this there are other documents [Oa04, Ode06, OZ05b, OCRZ03, OZ05a].

Scala has been a collective effort of many people. The design and the implementation of version 1.0 was completed by Philippe Altherr, Vincent Cremet, Gilles Dubochet, Burak Emir, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, Matthias Zenger, and the author. Iulian Dragos, Gilles Dubochet, Sean McDirmid and Lex Spoon joined in the effort to develop the second version of the language and tools. Gilad Bracha, Craig Chambers, Erik Ernst, Matthias Felleisen, Shriram Krishnamurti, Gary Leavens, Sebastian Maneth, Erik Mejer, Klaus Ostermann, Didier Rémy, Mads Torgersen, and Philip Wadler have shaped the design of the language through lively and inspiring discussions and comments on previous versions of this document. The contributors to the Scala mailing list have also given very useful feedback that helped us improve the language and its tools.

Chapter 1

Lexical Syntax

Scala programs are written using the Unicode character set. This chapter defines the two modes of Scala's lexical syntax, the Scala mode and the XML mode. If not otherwise mentioned, the following descriptions of Scala tokens refer to Scala mode, and literal characters 'c' refer to the ASCII fragment \u0000-\u007F.

In Scala mode, *Unicode escapes* are replaced by the corresponding Unicode character with the given hexadecimal code.

```
UnicodeEscape ::= {\{\}\u{u} hexDigit hexDigit hexDigit hexDigit
hexDigit      ::= '0' | ... | '9' | 'A' | ... | 'F' | 'a' | ... | 'f' |
```

To construct tokens, characters are distinguished according to the following classes (Unicode general category given in parentheses):

1. Whitespace characters. \u0020 | \u0009 | \u000D | \u000A
2. Letters, which include lower case letters(Ll), upper case letters(Lu), title-case letters(Lt), other letters(Lo), letter numerals(Nl) and the two characters \u0024 '\$' and \u005F '_', which both count as upper case letters
3. Digits '0' | ... | '9'.
4. Parentheses '(' | ')' | '[' | ']' | '{' | '}'.
5. Delimiter characters '' | '' | "" | '.' | ';' | ','.
6. Operator characters. These consist of all printable ASCII characters \u0020-\u007F. which are in none of the sets above, mathematical symbols(Sm) and other symbols(So).

1.1 Identifiers

Syntax:

```

op      ::= special {special}
varid   ::= lower idrest
plainid ::= upper idrest
        | varid
        | op
id      ::= plainid
        | '\\'string chars '\\'
idrest  ::= {letter | digit} ['_' op]

```

There are three ways to form an identifier. First, an identifier can start with a letter which can be followed by an arbitrary sequence of letters and digits. This may be followed by underscore ‘_’ characters and other string composed of either letters and digits or of special characters. Second, an identifier can start with a special character followed by an arbitrary sequence of special characters. The preceding two forms are called *plain* identifiers. Finally, an identifier may also be formed by an arbitrary string between back-quotes (host systems may impose some restrictions on which strings are legal for identifiers). The identifier then is composed of all characters excluding the backquotes themselves.

As usual, a longest match rule applies. For instance, the string

```
big_bob++='def'
```

decomposes into the three identifiers `big_bob`, `++=`, and `def`. The rules for pattern matching further distinguish between *variable identifiers*, which start with a lower case letter, and *constant identifiers*, which do not.

The ‘\$’ character is reserved for compiler-synthesized identifiers. User programs should not define identifiers which contain ‘\$’ characters.

The following names are reserved words instead of being members of the syntactic class `id` of lexical identifiers.

abstract	case	catch	class	def
do	else	extends	false	final
finally	for	if	implicit	import
match	new	null	object	override
package	private	protected	requires	return
sealed	super	this	throw	trait
try	true	type	val	var
while	with	yield		
_	:	=	=>	<-
				<:
			<%	>:
			#	@

The Unicode operator \u21D2 ‘⇒’, which has the ASCII equivalent ‘=>’, is also re-

served.

Example 1.1.1 Here are examples of identifiers:

x	Object	maxIndex	p2p	empty_?
+	+_field	'yield'	αρετη	

1.2 Newline Characters

Syntax:

```
StatementSeparator ::= NewLine | ';' ;
```

Scala is a line-oriented language where statements may be terminated by semi-colons or newlines. A newline in a Scala source text is treated as the special `NewLine` token if the three following criteria are satisfied:

1. The token immediately preceding the newline can terminate a statement.
2. The token immediately following the newline can begin a statement.
3. The token appears in a region where multiple statements are allowed.

The tokens that can terminate a statement are: literals, identifiers and the following delimiters and reserved words:

```
this    null    true    false    return    _    )    ]    }
```

The tokens that can begin a statement are all Scala tokens *except* the following delimiters and reserved words:

```
catch    else    extends    finally    match    requires    with
yield    ,    .    ;    :    _    =    =>    <-    <:    <%    >:
#    @    )    ]    }
```

A **case** token can begin a statement only if followed by a **class** or **object** token.

Multiple statements are allowed in:

1. all of a Scala source file, except for nested regions where newlines are suppressed, and
2. the interval between matching { and } brace tokens, except for nested regions where newlines are suppressed.

Multiple statements are disabled in:

1. the interval between matching (and) parenthesis tokens, except for nested regions where newlines are enabled, and
2. the interval between matching [and] bracket tokens, except for nested regions where newlines are enabled.
3. The interval between a **case** token and its matching => token, except for nested regions where newlines are enabled.
4. Any regions analyzed in XML mode (§1.5).

Note that the brace characters of { . . . } escapes in XML and string literals are not tokens, and therefore do not enclose a region where newlines are enabled.

1.3 Literals

There are literals for integer numbers floating point numbers, characters, booleans, symbols, strings. The syntax of these literals is in each case as in Java.

Syntax:

```

Literal      ::= integerLiteral
              | floatingPointLiteral
              | booleanLiteral
              | characterLiteral
              | stringLiteral
              | symbolLiteral

```

1.3.1 Integer Literals

Syntax:

```

integerLiteral ::= (decimalNumeral | hexadecimal | octalNumeral) ['L' | 'l']
decimalNumeral ::= '0' | nonZeroDigit {digit}
hexadecimal    ::= '0' 'x' hexDigit {hexDigit}
octalNumeral   ::= '0' octalDigit {octalDigit}
digit          ::= '0' | nonZeroDigit
nonZeroDigit   ::= '1' | ... | '9'
octalDigit     ::= '0' | ... | '7'

```

Integer literals are usually of type `int`, or of type `long` when followed by a `L` or `l` suffix. Values of type `int` are all integer numbers between -2^{31} and $2^{31} - 1$, inclusive. Values of type `long` are all integer numbers between -2^{63} and $2^{63} - 1$, inclusive. A compile-time error occurs if an integer literal denotes a number outside these ranges.

However, if the expected type *pt* (§6) of a literal in an expression is either `byte`, `short`, or `char` and the integer number fits in the numeric range defined by the type,

then the number is converted to type *pt* and the literal's type is *pt*. The numeric ranges given by these types are:

byte	-2^7 to $2^7 - 1$
short	-2^{15} to $2^{15} - 1$
char	0 to $2^{16} - 1$

Example 1.3.1 Here are some integer literals:

0 -21 0xFFFFFFFF 0777L

1.3.2 Floating Point Literals

Syntax:

```
floatingPointLiteral ::= digit {digit} '.' {digit} [exponentPart] [floatType]
                      |   '.' digit {digit} [exponentPart] [floatType]
                      |   digit {digit} exponentPart [floatType]
                      |   digit {digit} [exponentPart] floatType
exponentPart         ::= ('E' | 'e') ['+' | '-'] digit {digit}
floatType             ::= 'F' | 'f' | 'D' | 'd'
```

Floating point literals are of type float when followed by a floating point type suffix F or f, and are of type double otherwise. The type float consists of all IEEE 754 32-bit single-precision binary floating point values, whereas the type double consists of all IEEE 754 64-bit double-precision binary floating point values.

Example 1.3.2 Here are some floating point literals:

0.0 1e30f 3.14159f 1.0e-100 .1

1.3.3 Boolean Literals

Syntax:

```
booleanLiteral      ::= true | false
```

The boolean literals **true** and **false** are members of type boolean.

1.3.4 Character Literals

Syntax:

```
characterLiteral    ::= '\ ' char '\ '
                      |   '\ ' charEscapeSeq '\ '
```

A character literal is a single character enclosed in quotes. The character is either a printable unicode character or is described by an escape sequence (§1.3.6).

Example 1.3.3 Here are some character literals:

```
'a'      '\u0041'      '\n'      '\t'
```

Note that '\u000A' is *not* a valid character literal because Unicode conversion is done before literal parsing and the Unicode character \u000A (line feed), and is not a printable character. One can use instead the escape sequence '\n' or the octal escape '\12' (§1.3.6).

1.3.5 String Literals

Syntax:

```
stringLiteral      ::= '\'' {stringElement} '\''
stringElement      ::= charNoDoubleQuote | charEscapeSeq
```

A string literal is a sequence of characters in double quotes. The characters are either printable unicode character or are described by escape sequences (§1.3.6). If the string literal contains a double quote character, it must be escaped, i.e. \". The value of a string literal is an instance of class `String`. T

Example 1.3.4 Here are some string literals:

```
"Hello,\nWorld!"
"This string contains a \" character."
```

1.3.6 Escape Sequences

The following escape sequences are recognized in character and string literals.

\b	\u0008: backspace BS
\t	\u0009: horizontal tab HT
\n	\u000a: linefeed LF
\f	\u000c: form feed FF
\r	\u000d: carriage return CR
\"	\u0022: double quote "
\'	\u0027: single quote '
\\	\u005c: backslash \

A character with Unicode between 0 and 255 may also be represented by an octal escape, i.e. a backslash '\' followed by a sequence of up to three octal characters.

It is a compile time error if a backslash character in a character or string literal does not start a valid escape sequence.

1.3.7 Symbol literals

Syntax:

```
symbolLiteral ::= ' ' plainid
```

A symbol literal `'x'` is a shorthand for the expression `scala.Symbol("x")`. `Symbol` is a case class (§5.3.2), which is defined as follows.

```
package scala
final case class Symbol(name: String) {
  override def toString(): String = "'" + name
}
```

1.3.8 Primitive Types Defined in *Predef*

The object `Predef` is imported implicitly into every Scala program. It contains type definitions which establish the primitive types mentioned above as aliases of class types. Numeric and boolean types are equated with standard Scala classes. The `String` type is equated with the string class of the underlying host system. In a Java environment, `Predef` contains the following bindings, among others:

```
type byte    = scala.Byte
type short   = scala.Short
type char     = scala.Char
type int      = scala.Int
type long     = scala.Long
type float    = scala.Float
type double   = scala.Double
type boolean  = scala.Boolean
type String   = java.lang.String
```

1.4 Whitespace and Comments

Tokens may be separated by whitespace characters and/or comments. Comments come in two forms:

A single-line comment is a sequence of characters which starts with `//` and extends to the end of the line.

A multi-line comment is a sequence of characters between `/*` and `*/`. Multi-line comments may be nested.

1.5 XML mode

In order to allow literal inclusion of XML fragments, lexical analysis switches from Scala mode to XML mode when encountering an opening angle bracket '`<`' in the following circumstance: The '`<`' must be preceded either by whitespace, an opening parenthesis or an opening brace and immediately followed by a character starting an XML name.

Syntax:

```
( whitespace | '(' | '{' ) '<' (XNameStart | '!' | '?')
```

```
XNameStart ::= '_' | BaseChar | Ideographic (as in W3C XML, but without ':'
```

The scanner switches from XML mode to Scala mode if either

- the XML expression or the XML pattern started by the initial '`<`' has been successfully parsed, or if
- the parser encounters an embedded Scala expression or pattern and forces the Scanner back to normal mode, until the Scala expression or pattern is successfully parsed. In this case, since code and XML fragments can be nested, the parser has to maintain a stack that reflects the nesting of XML and Scala expressions adequately.

Note that no Scala tokens are constructed in XML mode, and that comments are interpreted as text.

Example 1.5.1 The following value definition uses an XML literal with two embedded Scala expressions

```
val b = <book>
  <title>The Scala Language Specification</title>
  <version>{scalaBook.version}</version>
  <authors>{scalaBook.authors.mkList("", " ", " ")}</authors>
</book>
```

Chapter 2

Identifiers, Names and Scopes

Names in Scala identify types, values, methods, and classes which are collectively called *entities*. Names are introduced by local definitions and declarations (§4), inheritance (srefsec:members), import clauses (§4.7), or package clauses (§9.2) which are collectively called *bindings*.

Bindings of different kinds have a precedence defined on them: Definitions (local or inherited) have highest precedence, followed by explicit imports, followed by wild-card imports, followed by package members, which have lowest precedence.

There are two different name spaces, one for types (§3) and one for terms (§6). The same name may designate a type and a term, depending on the context where the name is used.

A binding has a *scope* in which the entity defined by a single name can be accessed using a simple name. Scopes are nested. A binding in some inner scope *shadows* bindings of lower precedence in the same scope as well as bindings of the same or lower precedence in outer scopes.

Note that shadowing is only a partial order. In a situation like

```
val x = 1;
{ import p.x;
  x }
```

neither binding of `x` shadows the other. Consequently, the reference to `x` in the third line above would be ambiguous.

A reference to an unqualified (type- or term-) identifier *x* is bound by the unique binding, which

- defines an entity with name *x* in the same namespace as the identifier, and
- shadows all other bindings that define entities with name *x* in that namespace.

It is an error if no such binding exists. If x is bound by an import clause, then the simple name x is taken to be equivalent to the qualified name to which x is mapped by the import clause. If x is bound by a definition or declaration, then x refers to the entity introduced by that binding. In that case, the type of x is the type of the referenced entity.

Example 2.0.2 Assume the following two definitions of a objects named X in packages P and Q .

```
package P {
  object X { val x = 1; val y = 2 }
}

package Q {
  object X { val x = true; val y = "" }
}
```

The following program illustrates different kinds of bindings and precedences between them.

```
package P {                                     // 'X' bound by package clause
import Console._                               // 'println' bound by wildcard import
object A {
  println("L4: "+X)                             // 'X' refers to 'P.X' here
  object B {
    import Q._                                   // 'X' bound by wildcard import
    println("L7: "+X)                             // 'X' refers to 'Q.X' here
    import X._                                   // 'x' and 'y' bound by wildcard import
    println("L8: "+x)                             // 'x' refers to 'Q.X.x' here
    object C {
      val x = 3                                   // 'x' bound by local definition
      println("L12: "+x)                         // 'x' refers to constant '3' here
      { import Q.X._                             // 'x' and 'y' bound by wildcard import
//      println("L14: "+x)                       // reference to 'x' is ambiguous here
      import X.y                                 // 'y' bound by explicit import
      println("L16: "+y)                         // 'y' refers to 'Q.X.y' here
      { val x = "abc"                             // 'x' bound by local definition
      import P.X._                               // 'x' and 'y' bound by wildcard import
//      println("L19: "+y)                       // reference to 'y' is ambiguous here
      println("L20: "+x)                         // 'x' refers to string "abc" here
      }}}}}
}}}}}}
```

A reference to a qualified (type- or term-) identifier $e.x$ refers to the member of the type T of e which has the name x in the same namespace as the identifier. It is an error if T is not a value type (§3.2). The type of $e.x$ is the member type of the referenced entity in T .

Chapter 3

Types

Syntax:

```
Type      ::= Type1 '=>' Type
           | '(' [Types] ')' '=>' Type
           | Type1
Type1      ::= SimpleType {with SimpleType} [Refinement]
SimpleType ::= StableId
           | SimpleType '#' id
           | Path '.' type
           | SimpleType TypeArgs
           | '(' Type ')'
Types      ::= Type {' , ' Type}
```

We distinguish between first-order types and type constructors, which take type parameters and yield types. A subset of first-order types called *value types* represents sets of (first-class) values. Value types are either *concrete* or *abstract*. Every concrete value type can be represented as a *class type*, i.e. a type designator (§3.2.3) that refers to a class¹ (§5.3), or as a *compound type* (§3.2.5) representing an intersection of types, possibly with a refinement (§3.2.5) that further constrains the types of its members. Abstract value types are introduced by type parameters and abstract type bindings (§4.3). Parentheses in types are used for grouping.

Non-value types capture properties of identifiers that are not values (§3.3). There is no syntax to express these types directly in Scala.

3.1 Paths

Syntax:

¹We assume that objects and packages also implicitly define a class (of the same name as the object or package, but inaccessible to user programs).

```

Path      ::= StableId
           | [id '.'] this
StableId   ::= id
           | Path '.' id
           | [id '.'] super [ClassQualifier] '.' id
ClassQualifier ::= '[' id ']'

```

Paths are not types themselves, but they can be a part of named types and in that function form a central role in Scala's type system.

A path is one of the following.

- The empty path ϵ (which cannot be written explicitly in user programs).
- $C.\mathbf{this}$, where C references a class. The path **this** is taken as a shorthand for $C.\mathbf{this}$ where C is the name of the class directly enclosing the reference.
- $p.x$ where p is a path and x is a stable member of p . *Stable members* are members introduced by value or object definitions, as well as packages.
- $C.\mathbf{super}.x$ or $C.\mathbf{super}[M].x$ where C references a class and x references a stable member of the super class or designated parent class M of C . The prefix **super** is taken as a shorthand for $C.\mathbf{super}$ where C is the name of the class directly enclosing the reference.

A *stable identifier* is a path which ends in an identifier.

3.2 Value Types

Every value in Scala has a type which is of one of the following forms.

3.2.1 Singleton Types

Syntax:

```
SimpleType ::= Path '.' type
```

A singleton type is of the form $p.\mathbf{type}$, where p is a path pointing to a value expected to conform (§6) to `scala.AnyRef`. The type denotes the set of values consisting of `null` and the value denoted by p .

3.2.2 Type Projection

Syntax:

```
SimpleType ::= SimpleType '#' id
```

A type projection $T\#x$ references the type member named x of type T . If x references an abstract type member, then T must be a singleton type.

3.2.3 Type Designators

Syntax:

`SimpleType ::= StableId`

A type designator refers to a named value type. It can be simple or qualified. All such type designators are shorthands for type projections.

Specifically, the unqualified type name t where t is bound in some class, object, or package C is taken as a shorthand for $C.\mathbf{this.type}\#t$. If t is not bound in a class, object, or package, then t is taken as a shorthand for $\epsilon.\mathbf{type}\#t$.

A qualified type designator has the form $p.t$ where p is a path (§3.1) and t is a type name. Such a type designator is equivalent to the type projection $p.\mathbf{type}\#x$.

Example 3.2.1 Some type designators and their expansions are listed below. We assume a local type parameter t , a value maintable with a type member `Node` and the standard class `scala.Int`,

<code>t</code>	<code>$\epsilon.\mathbf{type}\#t$</code>
<code>Int</code>	<code><code>scala</code>.$\mathbf{type}\#\mathbf{Int}$</code>
<code>scala.Int</code>	<code><code>scala</code>.$\mathbf{type}\#\mathbf{Int}$</code>
<code>data.maintable.Node</code>	<code><code>data.maintable</code>.$\mathbf{type}\#\mathbf{Node}$</code>

3.2.4 Parameterized Types

Syntax:

`SimpleType ::= SimpleType TypeArgs`
`TypeArgs ::= '[' Types '[']'`

A parameterized type $T[U_1, \dots, U_n]$ consists of a type designator T and type parameters U_1, \dots, U_n where $n \geq 1$. T must refer to a type constructor which takes n type parameters a_1, \dots, a_n .

Say the type parameters have lower bounds L_1, \dots, L_n and upper bounds U_1, \dots, U_n . The parameterized type is well-formed if each actual type parameter *conforms to its bounds*, i.e. $\sigma L_i <: T_i <: \sigma U_i$ where σ is the substitution $[a_1 := T_1, \dots, a_n := T_n]$.

Example 3.2.2 Given the partial type definitions:

```
class TreeMap[a <: Comparable[a], b] { ... }
class List[a] { ... }
class I extends Comparable[I] { ... }
```

the following parameterized types are well formed:

```
TreeMap[I, String]
List[I]
List[List[Boolean]]
```

Example 3.2.3 Given the type definitions of §Example 3.2.2, the following types are ill-formed:

```
TreeMap[I] // illegal: wrong number of parameters
TreeMap[List[I], Boolean] // illegal: type parameter not within bound
```

3.2.5 Compound Types

Syntax:

```
Type          ::= SimpleType {with SimpleType} [Refinement]
Refinement     ::= '{' [RefineStat {StatementSeparator RefineStat}] '}'
RefineStat     ::= Decl
                |  type TypeDef
                |
```

A compound type T_1 **with** ... **with** T_n $\{R\}$ represents objects with members as given in the component types T_1, \dots, T_n and the refinement $\{R\}$. A refinement $\{R\}$ contains declarations and type definitions. Each declaration or definition in a refinement must override a declaration or definition in one of the component types T_1, \dots, T_n . The usual rules for overriding (§5.1.4) apply. If no refinement is given, the empty refinement is implicitly added, i.e. T_1 **with** ... **with** T_n is a shorthand for T_1 **with** ... **with** T_n $\{\}$.

3.2.6 Function Types

Syntax:

```
SimpleType     ::= Type1 '=>' Type
                |  '(' => Type ')' '=>' Type
                |  '(' [Types] ')' '=>' Type
```

The type $(T_1, \dots, T_n) \Rightarrow U$ represents the set of function values that take arguments of types T_1, \dots, T_n and yield results of type U . In the case of exactly one argument type $T \Rightarrow U$ is a shorthand for $(T) \Rightarrow U$. The type $(\Rightarrow T) \Rightarrow U$ represents functions with a call-by-name parameter (§4.6) of type T which yield a result of type U . Function types associate to the right, e.g. $(S) \Rightarrow (T) \Rightarrow U$ is the same as $(S) \Rightarrow ((T) \Rightarrow U)$.

Function types are shorthands for class types that define apply functions. Specifically, the n -ary function type $(T_1, \dots, T_n) \Rightarrow U$ is a shorthand for the class type

$\text{Function}_n[T_1, \dots, T_n, U]$. Such class types are defined in the Scala library for n between 0 and 9 as follows.

```
package scala
trait Function_n[-T_1, ..., -T_n, +R] {
  def apply(x_1: T_1, ..., x_n: T_n): R
  override def toString() = "<function>"
}
```

Hence, function types are covariant (§4.5) in their result type and contravariant in their argument types.

A call-by-name function type $(\Rightarrow T) \Rightarrow U$ is a shorthand for the class type $\text{ByNameFunction}[T, U]$, which is defined as follows.

```
package scala
trait ByNameFunction[-T, +R] {
  def apply(x: => T): R
  override def toString() = "<function>"
}
```

3.3 Non-Value Types

The types explained in the following do not denote sets of values, nor do they appear explicitly in programs. They are introduced in this report as the internal types of defined identifiers.

3.3.1 Method Types

A method type is denoted internally as $(Ts)U$, where (Ts) is a sequence of types (T_1, \dots, T_n) for some $n \geq 0$ and U is a (value or method) type. This type represents named methods that take arguments of types T_1, \dots, T_n and that return a result of type U .

We let method types associate to the right: $(Ts_1)(Ts_2)U$ is treated as $(Ts_1)((Ts_2)U)$.

A special case are types of methods without any parameters. They are written here $\Rightarrow T$. Parameterless methods name expressions that are re-evaluated each time the parameterless method name is referenced.

Method types do not exist as types of values. If a method name is used as a value, its type is implicitly converted to a corresponding function type (§6.22).

Example 3.3.1 The declarations

```
def a: Int
def b (x: Int): Boolean
```

```
def c (x: Int) (y: String, z: String): String
```

produce the typings

```
a: => Int
b: (Int) Boolean
c: (Int) (String, String) String
```

3.3.2 Polymorphic Method Types

A polymorphic method type is denoted internally as $[tps]T$ where $[tps]$ is a type parameter section $[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]$ for some $n \geq 0$ and T is a (value or method) type. This type represents named methods that take type arguments S_1, \dots, S_n which conform (§3.2.4) to the lower bounds L_1, \dots, L_n and the upper bounds U_1, \dots, U_n and that yield results of type T .

Example 3.3.2 The declarations

```
def empty[a]: List[a]
def union[a <: Comparable[a]] (x: Set[a], xs: Set[a]): Set[a]
```

produce the typings

```
empty : [a >: Bottom <: Any] List[a]
union : [a >: Bottom <: Comparable[a]] (x: Set[a], xs: Set[a]) Set[a] .
```

3.4 Base Types and Member Definitions

Types of class members depend on the way the members are referenced. Central here are three notions, namely:

1. the notion of the set of base types of a type T ,
2. the notion of a type T in some class C seen from some prefix type S ,
3. the notion of the set of member bindings of some type T .

These notions are defined mutually recursively as follows.

1. The set of *base types* of a type is a set of class types, given as follows.

- The base types of a class type C with parents T_1, \dots, T_n are C itself, as well as the base types of the compound type T_1 **with** ... **with** T_n $\{R\}$.
- The base types of an aliased type are the base types of its alias.
- The base types of an abstract type are the base types of its upper bound.

- The base types of a parameterized type $C[T_1, \dots, T_n]$ are the base types of type C , where every occurrence of a type parameter a_i of C has been replaced by the corresponding parameter type T_i .
- The base types of a singleton type $p.\mathbf{type}$ are the base types of the type of p .
- The base types of a compound type $T_1 \mathbf{with} \dots \mathbf{with} T_n \{R\}$ are the *reduced union* of the base classes of all T_i 's. This means: Let the multi-set \mathcal{S} be the multi-set-union of the base types of all T_i 's. If \mathcal{S} contains several type instances of the same class, say $S^i \# C[T_1^i, \dots, T_n^i]$ ($i \in I$), then all those instances are replaced by one of them which conforms to all others. It is an error if no such instance exists. It follows that the reduced union, if it exists, produces a set of class types, where different types are instances of different classes.
- The base types of a type selection $S\#T$ are determined as follows. If T is an alias or abstract type, the previous clauses apply. Otherwise, T must be a (possibly parameterized) class type, which is defined in some class B . Then the base types of $S\#T$ are the base types of T in B seen from the prefix type S .

2. The notion of a type T in class C seen from some prefix type S makes sense only if the prefix type S has a type instance of class C as a base type, say $S' \# C[T_1, \dots, T_n]$. Then we define as follows.

- If $S = \epsilon.\mathbf{type}$, then T in C seen from S is T itself.
- Otherwise, if T is the i 'th type parameter of some class D , then
 - If S has a base type $D[U_1, \dots, U_n]$, for some type parameters $[U_1, \dots, U_n]$, then T in C seen from S is U_i .
 - Otherwise, if C is defined in a class C' , then T in C seen from S is the same as T in C' seen from S' .
 - Otherwise, if C is not defined in another class, then T in C seen from S is T itself.
- Otherwise, if T is the singleton type $D.\mathbf{this.type}$ for some class D then
 - If D is a subclass of C and S has a type instance of class D among its base types, then T in C seen from S is S .
 - Otherwise, if C is defined in a class C' , then T in C seen from S is the same as T in C' seen from S' .
 - Otherwise, if C is not defined in another class, then T in C seen from S is T itself.
- If T is some other type, then the described mapping is performed to all its type components.

If T is a possibly parameterized class type, where T 's class is defined in some other class D , and S is some prefix type, then we use “ T seen from S ” as a shorthand for “ T in D seen from S ”.

3. The *member bindings* of a type T are all bindings d such that there exists a type instance of some class C among the base types of T and there exists a definition or declaration d' in C such that d results from d' by replacing every type T' in d' by T' in C seen from T .

The *definition* of a type projection $S\#t$ is the member binding d_t of the type t in S . In that case, we also say that $S\#t$ is *defined by* d_t .

3.5 Relations between types

We define two relations between types.

<i>Type equivalence</i>	$T \equiv U$	T and U are interchangeable in all contexts.
<i>Conformance</i>	$T <: U$	Type T conforms to type U .

3.5.1 Type Equivalence

Equivalence (\equiv) between types is the smallest congruence² such that the following holds:

- If t is defined by a type alias **type** $t = T$, then t is equivalent to T .
- If a path p has a singleton type $q.\mathbf{type}$, then $p.\mathbf{type} \equiv q.\mathbf{type}$.
- If O is defined by an object definition, and p is a path consisting only of package or object selectors and ending in O , then $O.\mathbf{this.type} \equiv p.\mathbf{type}$.
- Two compound types are equivalent if the sequences of their component are pairwise equivalent, and occur in the same order, and their refinements are equivalent. Two refinements are equivalent if they bind the same names and the modifiers, types and bounds of every declared entity are equivalent in both refinements.
- Two method types are equivalent if they have equivalent result types, both have the same number of parameters, and corresponding parameters have equivalent types. Note that the names of parameters do not matter for method type equivalence.
- Two polymorphic types are equivalent if they have the same number of type parameters, and, after renaming one set of type parameters by another, the result types as well as lower and upper bounds of corresponding type parameters are equivalent.

² A congruence is an equivalence relation which is closed under formation of contexts

3.5.2 Conformance

The conformance relation ($<:$) is the smallest transitive relation that satisfies the following conditions.

- Conformance includes equivalence. If $T \equiv U$ then $T <: U$.
- For every value type T , $\text{scala.Bottom} <: T <: \text{scala.Any}$.
- For every class type $T <: \text{scala.AnyRef}$ one has $\text{scala.Null} <: T$.
- A type variable or abstract type t conforms to its upper bound and its lower bound conforms to t .
- A class type or parameterized type conforms to any of its base-types.
- A type projection $T\#t$ conforms to $U\#t$ if T conforms to U .
- A parameterized type $T[T_1, \dots, T_n]$ conforms to $T[U_1, \dots, U_n]$ if the following three conditions hold for $i = 1, \dots, n$.
 - If the i 'th type parameter of T is declared covariant, then $T_i <: U_i$.
 - If the i 'th type parameter of T is declared contravariant, then $U_i <: T_i$.
 - If the i 'th type parameter of T is declared neither covariant nor contravariant, then $U_i \equiv T_i$.
- A compound type $T_1 \text{ with } \dots \text{ with } T_n \{R\}$ conforms to each of its component types T_i .
- If $T <: U_i$ for $i = 1, \dots, n$ and for every binding d of a type or value x in R there exists a member binding of x in T which subsumes d , then T conforms to the compound type $U_1 \text{ with } \dots \text{ with } U_n \{R\}$.
- If $T_i \equiv T'_i$ for $i = 1, \dots, n$ and U conforms to U' then the method type $(T_1, \dots, T_n)U$ conforms to $(T'_1, \dots, T'_n)U'$.
- The polymorphic type $[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]T$ conforms to the polymorphic type $[a_1 >: L'_1 <: U'_1, \dots, a_n >: L'_n <: U'_n]T'$ if, assuming $L'_1 <: a_1 <: U'_1, \dots, L'_n <: a_n <: U'_n$ one has $T <: T'$ and $L_i <: L'_i$ and $U'_i <: U_i$ for $i = 1, \dots, n$.

A declaration or definition in some compound type of class type C *subsumes* another declaration of the same name in some compound type or class type C' , if one of the following holds.

- A value or declaration or definition that defines a name x with type T subsumes a value or method declaration that defines x with type T' , provided $T <: T'$.
- A method declaration or definition that defines a name x with type T subsumes a method declaration that defines x with type T' , provided $T <: T'$.

- A type alias **type** $t = T$ subsumes a type alias **type** $t = T'$ if $T \equiv T'$.
- A type declaration **type** $t >: L <: U$ subsumes a type declaration **type** $t >: L' <: U'$ if $L' <: L$ and $U <: U'$.
- A type or class definition that binds a type name t subsumes an abstract type declaration **type** $t >: L <: U$ if $L <: t <: U$.

The ($<:$) relation forms pre-order between types, i.e. it is transitive and reflexive. *least upper bounds* and *greatest lower bounds* of a set of types are understood to be relative to that order.

Note. The least upper bound or greatest lower bound of a set of types does not always exist. For instance, consider the class definitions

```
class A[+t] {}
class B extends A[B]
class C extends A[C]
```

Then the types $A[\text{Any}]$, $A[A[\text{Any}]]$, $A[A[A[\text{Any}]]]$, \dots form a descending sequence of upper bounds for B and C. The least upper bound would be the infinite limit of that sequence, which does not exist as a Scala type. Since cases like this are in general impossible to detect, a Scala compiler is free to reject a term which has a type specified as a least upper or greatest lower bound, and that bound would be more complex than some compiler-set limit³.

The least upper bound or greatest lower bound might also not be unique. For instance $A \text{ with } B$ and $B \text{ with } A$ are both least upper bounds of A and B. If there are several least upper bounds or greatest lower bounds, the Scala compiler is free to pick any one of them.

3.6 Type Erasure

A type is called *generic* if it contains type arguments or type variables. *Type erasure* is a mapping from (possibly generic) types to non-generic types. We write $|T|$ for the erasure of type T . The erasure mapping is defined as follows.

- The erasure of a type variable is the erasure of its upper bound.
- The erasure of a parameterized type $T[T_1, \dots, T_n]$ is $|T|$.
- The erasure of a singleton type $p.\text{type}$ is the erasure of the type of p .
- The erasure of a type projection $T\#x$ is $|T|\#x$.
- The erasure of a compound type $T_1 \text{ with } \dots \text{ with } T_n \{R\}$ is $|T_1|$.
- The erasure of every other type is the type itself.

³The current Scala compiler limits the nesting level of parameterization in such bounds to 10.

Chapter 4

Basic Declarations and Definitions

Syntax:

```
Dcl          ::=  val ValDcl
               |  var VarDcl
               |  def FunDcl
               |  type TypeDcl
Def          ::=  val PatDef
               |  var VarDef
               |  def FunDef
               |  type TypeDef
               |  TmplDef
```

A *declaration* introduces names and assigns them types. It can form part of a class definition (§5.1) or of a refinement in a compound type (§3.2.5).

A *definition* introduces names that denote terms or types. It can form part of an object or class definition or it can be local to a block. Both declarations and definitions produce *bindings* that associate type names with type definitions or bounds, and that associate term names with types.

The scope of a name introduced by a declaration or definition is the whole statement sequence containing the binding. However, there is a restriction on forward references in blocks: In a statement sequence $s_1 \dots s_n$ making up a block, if a simple name in s_i refers to an entity defined by s_j where $j \geq i$, then none of the definitions between and including s_i and s_j may be a value or variable definition.

4.1 Value Declarations and Definitions

Syntax:

```
Dcl          ::=  val ValDcl
```

```

ValDcl      ::= ids ':' Type
Def         ::= val PatDef
PatDef      ::= Pattern2 {',', Pattern2} [':' Type] '=' Expr
ids         ::= id {'', id}

```

A value declaration **val** $x: T$ introduces x as a name of a value of type T .

A value definition **val** $x: T = e$ defines x as a name of the value that results from the evaluation of e . The type T may be omitted, in which case the type of expression e is assumed. If a type T is given, then e is expected to conform to it (§6).

Evaluation of the value definition implies evaluation of its right-hand side e . The effect of the value definition is to bind x to the value of e converted to type T .

Value definitions can alternatively have a pattern (§8.1) as left-hand side. If p is some pattern other than a simple name or a name followed by a colon and a type, then the value definition **val** $p = e$ is expanded as follows:

1. If the pattern p has bound variables x_1, \dots, x_n , where $n > 1$:

```

val $x = e match { case p => scala.Tuplen( $x_1, \dots, x_n$ ) }
val  $x_1 = $x._1$ 
...
val  $x_n = $x._n$  .

```

Here, $$x$ is a fresh name. The class `Tuplen` is defined for $n = 2, \dots, 9$ in package `scala`.

2. If p has a unique bound variable x :

```

val  $x = e$  match { case p =>  $x$  }

```

3. If p has no bound variables:

```

 $e$  match { case p => () }

```

Example 4.1.1 The following are examples of value definitions

```

val pi = 3.1415
val pi: double = 3.1415    // equivalent to first definition
val Some(x) = f()         // a pattern definition
val x :: xs = mylist      // an infix pattern definition

```

The last two definitions have the following expansions.

```

val x = f() match { case Some(x) => x }

val x$ = mylist match { case x :: xs => scala.Tuple2(x, xs) }
val x = x$._1
val xs = x$._2

```


A value declaration **val** $x_1, \dots, x_n: T$ is a shorthand for the sequence of value declarations **val** $x_1: T$; ...; **val** $x_n: T$. A value definition **val** $p_1, \dots, p_n = e$ is a shorthand for the sequence of value definitions **val** $p_1 = e$; ...; **val** $p_n = e$. A value definition **val** $p_1, \dots, p_n: T = e$ is a shorthand for the sequence of value definitions **val** $p_1: T = e$; ...; **val** $p_n: T = e$.

4.2 Variable Declarations and Definitions

Syntax:

```

Dcl      ::= var VarDcl
Def      ::= var VarDef
VarDcl   ::= ids ':' Type
VarDef   ::= ids [':' Type] '=' Expr
           | ids ':' Type '=' '_'

```

A variable declaration **var** $x: T$ is equivalent to declarations of a *getter function* x and a *setter function* $x_=$, defined as follows:

```

def x: T
def x_= (y: T): unit

```

An implementation of a class containing variable declarations may define these variables using variable definitions, or it may define setter and getter functions directly.

A variable definition **var** $x: T = e$ introduces a mutable variable with type T and initial value as given by the expression e . The type T can be omitted, in which case the type of e is assumed. If T is given, then e is expected to conform to it (§6).

A variable definition **var** $x: T = _$ introduces a mutable variable with type T and a default initial value. The default value depends on the type T as follows:

```

0          if  $T$  is int or one of its subrange types,
0L         if  $T$  is long,
0.0f       if  $T$  is float,
0.0d       if  $T$  is double,
false    if  $T$  is boolean,
()         if  $T$  is unit,
null     for all other types  $T$ .

```

When they occur as members of a template, both forms of variable definition also introduce a getter function x which returns the value currently assigned to the variable, as well as a setter function $x_=$ which changes the value currently assigned to the variable. The functions have the same signatures as for a variable declaration.

The template then has the these getter and setter functions as members, whereas the original variable cannot be accessed directly as a template member.

Example 4.2.1 The following example shows how *properties* can be simulated in Scala. It defines a class `TimeOfDayVar` of time values with updatable integer fields representing hours, minutes, and seconds. Its implementation contains tests that allow only legal values to be assigned to these fields. The user code, on the other hand, accesses these fields just like normal variables.

```
class TimeOfDayVar {
  private var h: int = 0
  private var m: int = 0
  private var s: int = 0

  def hours          = h
  def hours_= (h: int) = if (0 <= h && h < 24) this.h = h
                        else throw new DateError()

  def minutes        = m
  def minutes_= (m: int) = if (0 <= m && m < 60) this.m = m
                        else throw new DateError()

  def seconds        = s
  def seconds_= (s: int) = if (0 <= s && s < 60) this.s = s
                        else throw new DateError()
}
val d = new TimeOfDayVar
d.hours = 8; d.minutes = 30; d.seconds = 0
d.hours = 25 // throws a DateError exception
```

A variable declaration `var $x_1, \dots, x_n: T$` is a shorthand for the sequence of variable declarations `var $x_1: T$; ...; var $x_n: T$` . A variable definition `var $x_1, \dots, x_n = e$` is a shorthand for the sequence of variable definitions `var $x_1 = e$; ...; var $x_n = e$` . A variable definition `var $x_1, \dots, x_n: T = e$` is a shorthand for the sequence of variable definitions `var $x_1: T = e$; ...; var $x_n: T = e$` .

4.3 Type Declarations and Type Aliases

Syntax:

```
Dcl ::= type TypeDcl
TypeDcl ::= id [>: Type] [<: Type]
Def ::= type TypeDef
TypeDef ::= id [TypeParamClause] '=' Type
```

A *type declaration* **type** $t >: L <: U$ declares t to be an abstract type with lower bound type L and upper bound type U . If such a declaration appears as a member declaration of a type, implementations of the type may implement t with any type T for which $L <: T <: U$. It is a compile-time error if L does not conform to U . Either or both bounds may be omitted. If the lower bound L is absent, the bottom type `scala.Bottom` is assumed. If the upper bound U is absent, the top type `scala.Any` is assumed.

A *type alias* **type** $t = T$ defines t to be an alias name for the type T . The left hand side of a type alias may have a type parameter clause, e.g. **type** $t[tps] = T$. The scope of a type parameter extends over the right hand side T and the type parameter clause tps itself.

The scope rules for definitions (§4) and type parameters (§4.6) make it possible that a type name appears in its own bound or in its right-hand side. However, it is a static error if a type alias refers recursively to the defined type constructor itself. That is, the type T in a type alias **type** $t[tps] = T$ may not refer directly or indirectly to the name t . It is also an error if an abstract type is directly or indirectly its own upper or lower bound.

Example 4.3.1 The following are legal type declarations and definitions:

```
type IntList = List[Integer]
type T <: Comparable[T]
type Two[a] = Tuple2[a, a]
```

The following are illegal:

```
type Abs = Comparable[Abs]           // recursive type alias

type S <: T                          // S, T are bounded by themselves.
type T <: S

type T >: Comparable[T.That]        // Cannot select from T.
                                   // T is a type, not a value
```

If a type alias **type** $t[tps] = S$ refers to a class type S , the name t can also be used as a constructor for objects of type S .

Example 4.3.2 The `Predef` object contains a definition which establishes `Pair` as an alias of the parameterized class `Tuple2`:

```
type Pair[+a, +b] = Tuple2[a, b]
```

As a consequence, for any two types S and T , the type `Pair[S, T]` is equivalent to the type `Tuple2[S, T]`. `Pair` can also be used as a constructor instead of `Tuple2`. Furthermore, because `Tuple2` is a case class (§5.3.2), `Pair2` is also an alias

4.5 Variance Annotations

Variance annotations indicate how instances of parameterized types vary with respect to subtyping (§3.5.2). A ‘+’ variance indicates a covariant dependency, a ‘-’ variance indicates a contravariant dependency, and a missing variance indication indicates an invariant dependency.

A variance annotation constrains the way the annotated type variable may appear in the type or class which binds the type parameter. In a type definition **type** $t[tps] = S$, type parameters labeled ‘+’ must only appear in covariant position in S whereas type parameters labeled ‘-’ must only appear in contravariant position. Analogously, for a class definition **class** $c[tps](ps)$ **requires** s **extends** t , type parameters labeled ‘+’ must only appear in covariant position in the self type s and the template t , whereas type parameters labeled ‘-’ must only appear in contravariant position.

The variance position of a type parameter in a type or template is defined as follows. Let the opposite of covariance be contravariance, and the opposite of invariance be itself. The top-level of the type or template is always in covariant position. The variance position changes at the following constructs.

- The variance position of a method parameter is the opposite of the variance position of the enclosing parameter clause.
- The variance position of a type parameter is the opposite of the variance position of the enclosing type parameter clause.
- The variance position of the lower bound of a type declaration or type parameter is the opposite of the variance position of the type declaration or parameter.
- The right hand side S of a type alias **type** $t[tps] = S$ is always in invariant position.
- The type of a mutable variable is always in invariant position.
- The prefix S of a type selection $S\#T$ is always in invariant position.
- For a type argument T of a type $S[\dots T \dots]$: If the corresponding type parameter is invariant, then T is in invariant position. If the corresponding type parameter is contravariant, the variance position of T is the opposite of the variance position of the enclosing type $S[\dots T \dots]$.

Example 4.5.1 The following variance annotation is legal.

```
abstract class P[+a, +b] {
  def fst: a; def snd: b
}
```

With this variance annotation, elements of type P subtype covariantly with respect to their arguments. For instance,

```
P[IOException, String] <: P[Throwable, AnyRef] .
```

If we make the elements of P mutable, the variance annotation becomes illegal.

```
abstract class Q[+a, +b] {
  var fst: a           // **** error: illegal variance:
  var snd: b           // 'a', 'b' occur in invariant position.
}
```

Example 4.5.2 The following variance annotation is illegal, since a appears in contravariant position in the parameter of `append`:

```
abstract class Vector[+a] {
  def append(x: Vector[a]): Vector[a]
                        // **** error: illegal variance:
                        // 'a' occurs in contravariant position.
}
```

The problem can be avoided by generalizing the type of `append` by means of a lower bound:

```
abstract class Vector[+a] {
  def append[b >: a](x: Vector[b]): Vector[b]
}
```

Example 4.5.3 Here is a case where a contravariant type parameter is useful.

```
abstract class OutputChannel[-a] {
  def write(x: a): unit
}
```

With that annotation, we have that `OutputChannel[AnyRef]` conforms to `OutputChannel[String]`. That is, a channel on which one can write any object can substitute for a channel on which one can write only strings.

4.6 Function Declarations and Definitions

Syntax:

```
Dcl           ::= def FunDcl
FunDcl        ::= FunSig ':' Type
Def           ::= def FunDef
```

```

FunDef      ::= FunSig [':' Type] '=' Expr
FunSig      ::= id [FunTypeParamClause] ParamClauses
FunTypeParamClause ::= [NewLine] [' TypeParam {' , ' TypeParam} ' ]
ParamClauses ::= {ParamClause} [[NewLine] '(' implicit Params ')']
ParamClause ::= [NewLine] '(' [Params] ')'
Params      ::= Param {' , ' Param}
Param       ::= id ':' ParamType
ParamType   ::= ['=>'] Type ['*']

```

A function declaration has the form **def** *fpsig*: *T*, where *f* is the function's name, *psig* is its parameter signature and *T* is its result type. A function definition *fpsig*: *T* = *e* also includes a *function body* *e*, i.e. an expression which defines the function's result. A parameter signature consists of an optional type parameter clause [*tps*], followed by zero or more value parameter clauses (*ps*₁)...(*ps*_{*n*}). Such a declaration or definition introduces a value with a (possibly polymorphic) method type whose parameter types and result type are as given.

The type of the function body must conform to the function's declared result type, if one is given. If the function definition is not recursive nor overloaded, the result type may be omitted, in which case it is determined from the type of the function body.

A type parameter clause *tps* consists of one or more type declarations (§4.3), which introduce type parameters, possibly with bounds. The scope of a type parameter includes the whole signature, including any of the type parameter bounds as well as the function body, if it is present.

A value parameter clause *ps* consists of zero or more formal parameter bindings such as *x*: *T*, which bind value parameters and associate them with their types. The scope of a formal value parameter name *x* is the function body, if one is given. Both type parameter names and value parameter names must be pairwise distinct.

The type of a value parameter may be prefixed by =>, e.g. *x*: => *T*. The type of such a parameter is then the parameterless method type => *T*. This indicates that the corresponding argument is not evaluated at the point of function application, but instead is evaluated at each use within the function. That is, the argument is evaluated using *call-by-name*.

Example 4.6.1 The declaration

```
def whileLoop (cond: => Boolean) (stat: => Unit): Unit
```

indicates that both parameters of `whileLoop` are evaluated using call-by-name.

The last value parameter of a parameter section may be suffixed by “*”, e.g. (... , *x*:*T**). The type of such a *repeated* parameter inside the method is then the sequence type `scala.Seq[T]`. Methods with repeated parameters *T** take a variable number of arguments of type *T*. That is, if a method *m* with type

$(T_1, \dots, T_n, S^*)U$ is applied to arguments (e_1, \dots, e_k) where $k \geq n$, then m is taken in that application to have type $(T_1, \dots, T_n, S, \dots, S)U$, with $k - n$ occurrences of type S . The only exception to this rule is if the last argument is marked to be a *sequence argument* via a $_*$ type annotation. If m above is applied to arguments $(e_1, \dots, e_n, e': _*)$, then the type of m in that application is taken to be $(T_1, \dots, T_n, \text{scala.Seq}[S])$.

Example 4.6.2 The following method definition computes the sum of a variable number of integer arguments.

```
def sum(args: int*) {
  var result = 0
  for (val arg <- args.elements) result = result + arg
  result
}
```

The following applications of this method yield 0, 1, 6, in that order.

```
sum()
sum(1)
sum(1, 2, 3)
```

Furthermore, assume the definition:

```
val xs = List(1, 2, 3)
```

The following applications method sum is ill-formed:

```
sum(xs)          // ***** error: expected: int, found: List[int]
```

By contrast, the following application is well formed and yields again the result 6:

```
sum(xs: \_*)
```

4.7 Import Clauses

Syntax:

```
Import          ::= import ImportExpr {',' ImportExpr}
ImportExpr      ::= StableId '.' (id | '\_' | ImportSelectors)
ImportSelectors ::= '{' {ImportSelector '\_','}
                  (ImportSelector | '\_') '}'
ImportSelector  ::= id ['=>' id | '<=>' '\_']
```

An import clause has the form **import** $p.I$ where p is a stable identifier (§3.1) and I is an import expression. The import expression determines a set of names

of members of p which are made available without qualification. The most general form of an import expression is a list of *import selectors*

$$\{ x_1 \Rightarrow y_1, \dots, x_n \Rightarrow y_n, _ \} .$$

for $n \geq 0$, where the final wildcard ‘ $_$ ’ may be absent. It makes available each member $p.x_i$ under the unqualified name y_i . I.e. every import selector $x_i \Rightarrow y_i$ renames $p.x_i$ to y_i . If a final wildcard is present, all members z of p other than x_1, \dots, x_n are also made available under their own unqualified names.

Import selectors work in the same way for type and term members. For instance, an import clause **import** $p.\{x \Rightarrow y\}$ renames the term name $p.x$ to the term name y and the type name $p.x$ to the type name y . At least one of these two names must reference a member of p .

If the target in an import selector is a wildcard, the import selector hides access to the source member. For instance, the import selector $x \Rightarrow _$ “renames” x to the wildcard symbol (which is inaccessible as a name in user programs), and thereby effectively prevents unqualified access to x . This is useful if there is a final wildcard in the same import selector list, which imports all members not mentioned in previous import selectors.

The scope of a binding introduced by an import-clause starts immediately after the import clause and extends to the end of the enclosing block, template, package clause, or compilation unit, whichever comes first.

Several shorthands exist. An import selector may be just a simple name x . In this case, x is imported without renaming, so the import selector is equivalent to $x \Rightarrow x$. Furthermore, it is possible to replace the whole import selector list by a single identifier or wildcard. The import clause **import** $p.x$ is equivalent to **import** $p.\{x\}$, i.e. it makes available without qualification the member x of p . The import clause **import** $p._$ is equivalent to **import** $p.\{_\}$, i.e. it makes available without qualification all members of p (this is analogous to **import** $p.*$ in Java).

An import clause with multiple import expressions **import** $p_1.I_1, \dots, p_n.I_n$ is interpreted as a sequence of import clauses **import** $p_1.I_1$; ...; **import** $p_n.I_n$.

Example 4.7.1 Consider the object definition:

```
object M {
  def z = 0, one = 1
  def add(x: Int, y: Int): Int = x + y
}
```

Then the block

```
{ import M.{one, z => zero, \_}; add(zero, one) }
```

is equivalent to the block

`{ M.add(M.z, M.one) } .`

Chapter 5

Classes and Objects

Syntax:

```
TplDef ::= [case] class ClassDef
        | [case] object ObjectDef
        | trait TraitDef
```

Classes (§5.3) and objects (§5.4) are both defined in terms of *templates*.

5.1 Templates

Syntax:

```
Template      ::= TemplateParents [TemplateBody]
TemplateParents ::= SimpleType {'(' [Exprs] ')'} {with SimpleType}
TemplateBody  ::= '{' [TemplateStat {StatementSeparator TemplateStat}] '}'
```

A template defines the type signature, behavior and initial state of a class of objects or of a single object. Templates form part of instance creation expressions, class definitions, and object definitions. A template *sc with* mt_1 **with** ... **with** mt_n {*stats*} consists of a constructor invocation *sc* which defines the template's *superclass*, trait references mt_1, \dots, mt_n ($n \geq 0$), which define the template's *traits*, and a statement sequence *stats* which contains initialization code and additional member definitions for the template.

Each trait reference mt_i must denote a trait (§5.3.3). By contrast, the superclass constructor *sc* normally refers to a class which is not a trait. It is possible to write a list of parents that starts with a trait reference, e.g. mt_1 **with** ... **with** mt_n . In that case the list of parents is implicitly extended to include the supertype of mt_1 as first parent type. The new supertype must have at least one constructor that does not take parameters. In the following, we will always assume that this implicit

extension has been performed, so that the first parent class of a template is a regular superclass constructor, not a trait reference.

The list of parents of every class is also always implicitly extended by a reference to the `scala.ScalaObject` trait as last mixin. E.g.

```
sc with mt1 with ... with mtn {stats}
```

becomes

```
mt1 with ... with mtn {stats} with ScalaObject {stats} .
```

The list of parents of a template must be well-formed. This means that the class denoted by the superclass constructor `sc` must be a subclass of the superclasses of all the traits mt_1, \dots, mt_n . In other words, the non-trait classes inherited by a template form a chain in the inheritance hierarchy which starts with the template's superclass.

The *least proper supertype* of a template is the class type or compound type (§3.2.5) consisting of all its parent class types.

The statement sequence `stats` contains member definitions that define new members or overwrite members in the parent classes. If the template forms part of a class definition, the statement part `stats` may also contain declarations of abstract members. Furthermore, `stats` may contain expressions that are executed in the order they are given as part of the initialization of a template.

Example 5.1.1 Consider the following class definitions:

```
class Base extends Object {}
trait Mixin extends Base {}
object O extends Mixin {}
```

In this case, the definition of `O` is expanded to:

```
object O extends Base with Mixin {}
```

Inheriting from Java Types. A template may have a Java class as its superclass and Java interfaces as its mixins.

Template Evaluation. Consider a template `sc with mt1 with mtn {stats}`.

If this is the template of a trait (§5.3.3) then its *mixin-evaluation* consists of an evaluation of the statement sequence `stats`.

If this is not a template of a trait, then its *evaluation* consists of the following steps.

- First, the superclass constructor `sc` is evaluated (§5.1.1).

- Then, all base classes in the template's linearization (§5.1.2) up to the template's superclass denoted by sc are mixin-evaluated. Mixin-evaluation happens in reverse order of occurrence in the linearization, i.e. the class immediately preceding sc is evaluated first.
- Finally the statement sequence $stats$ is evaluated.

5.1.1 Constructor Invocations

Syntax:

$\text{Constr} ::= \text{StableId} [\text{TypeArgs}] \{ ' (' [\text{Exprs}] ') ' \}$

Constructor invocations define the type, members, and initial state of objects created by an instance creation expression, or of parts of an object's definition which are inherited by a class or object definition. A constructor invocation is a function application $x.c[targs](args_1) \dots (args_n)$, where x is a stable identifier (§3.1), c is a type name which either designates a class or defines an alias type for one, $targs$ is a type argument list, and $args_1, \dots, args_n$ are argument lists, which match the parameters of one of the constructors of that class.

The prefix ' $x.$ ' can be omitted. A type argument list can be given only if the class c takes type parameters. Even then it can be omitted, in which case a type argument list is synthesized using local type inference (§6.22.4). If no explicit arguments are given, an empty list $()$ is implicitly supplied.

An evaluation of a constructor invocation $x.c[targs](args_1) \dots (args_n)$ consists of the following steps:

- First, the prefix x is evaluated.
- Then, the arguments $args_1, \dots, args_n$ are evaluated from left to right.
- Finally, the being constructed is initialized by evaluating the template of the class referred to by c .

5.1.2 Class Linearization

The classes reachable through transitive closure of the direct inheritance relation from a class C are called the *base classes* of C . Because of mixins, the inheritance relationship on base classes forms in general a directed acyclic graph. A linearization of this graph is defined as follows.

Definition 5.1.2 Let C be a class with template C_1 **with** ... **with** C_n { $stats$ }. The *linearization* of C , $\mathcal{L}(C)$ is defined as follows:

$$\mathcal{L}(C) = C, \mathcal{L}(C_n) \vec{+} \dots \vec{+} \mathcal{L}(C_1)$$

Here $\vec{+}$ denotes concatenation where elements of the right operand replace identical elements of the left operand:

$$\begin{aligned} \{a, A\} \vec{+} B &= a, (A \vec{+} B) && \text{if } a \notin B \\ &= A \vec{+} B && \text{if } a \in B \end{aligned}$$

Example 5.1.3 Consider the following class definitions.

```
abstract class AbsIterator extends AnyRef with ScalaObject { ... }
trait RichIterator extends AbsIterator { ... }
class StringIterator extends AbsIterator { ... }
class Iter extends StringIterator with RichIterator { ... }
```

Then the linearization of class Iter is

```
{ Iter, RichIterator, StringIterator, AbsIterator, AnyRef, Any }
```

Note that the linearization of a class refines the inheritance relation: if C is a subclass of D , then C precedes D in any linearization where both C and D occur. Definition 5.1.2 also satisfies the property that a linearization of a class always contains the linearization of its direct superclass as a suffix. For instance, the linearization of StringIterator is

```
{ StringIterator, AbsIterator, AnyRef, Any }
```

which is a suffix of the linearization of its subclass Iter. The same is not true for the linearization of mixins. For instance, the linearization of RichIterator is

```
{ RichIterator, AbsIterator, AnyRef, Any }
```

which is not a suffix of the linearization of Iter.

5.1.3 Class Members

A class C defined by a template C_1 **with** ... **with** C_n { *stats* } can define members in its statement sequence *stats* and can inherit members from all parent classes. Scala adopts Java and C#'s conventions for static overloading of methods. It is thus possible that a class defines and/or inherits several methods with the same name. To decide whether a defined member of a class C overrides a member of a parent class, or whether the two co-exist as overloaded variants in C , Scala uses the following definition of *matching* on members:

Definition 5.1.4 A member definition M *matches* a member definition M' , if M and M' bind the same name, and one of following holds.

1. Neither M nor M' is a method definition.

2. M and M' define both monomorphic methods with equal argument types.
3. M defines a parameterless method and M' defines a method with an empty parameter list $()$ or *vice versa*.
4. M and M' define both polymorphic methods with equal number of argument types \overline{T} , \overline{T}' and equal numbers of type parameters \overline{t} , \overline{t}' , say, and $\overline{T}' = [\overline{t}'/\overline{t}]\overline{T}$.

Member definitions fall into two categories: concrete and abstract. Members of class C are either *directly defined* (i.e. they appear in C 's statement sequence *stats*) or they are *inherited*. There are two rules that determine the set of members of a class, one for each category:

Definition 5.1.5 A *concrete member* of a class C is any concrete definition M in some class $C_i \in \mathcal{L}(C)$, except if there is a preceding class $C_j \in \mathcal{L}(C)$ where $j < i$ which directly defines a concrete member M' matching M .

An *abstract member* of a class C is any abstract definition M in some class $C_i \in \mathcal{L}(C)$, except if C contains already a concrete member M' matching M , or if there is a preceding class $C_j \in \mathcal{L}(C)$ where $j < i$ which directly defines an abstract member M' matching M .

This definition also determines the overriding relationships between matching members of a class C and its parents (§5.1.4). First, a concrete definition always overrides an abstract definition. Second, for definitions M and M' which are both concrete or both abstract, M overrides M' if M appears in a class that precedes (in the linearization of C) the class in which M' is defined.

It is an error if a template directly defines two matching members. It is also an error if a template contains two members (directly defined or inherited) with the same name and the same erased type (§3.6).

Example 5.1.6 Consider the class definitions

```
class A { def f: Int = 1 ; def g: Int = 2 ; def h: Int = 3 }
abstract class B { def f: Int = 4 ; def g: Int }
abstract class C extends A with B { def h: Int }
```

Then class C has a directly defined abstract member h . It inherits member f from class B and member g from class A .

5.1.4 Overriding

A member M of class C that matches a non-private member M' of a base class of C is said to *override* that member. In this case the binding of the overriding member M must subsume (§3.5.2) the binding of the overridden member M' . Furthermore, the following restrictions on modifiers apply to M and M' :

- M' must not be labeled **final**.
- M must not be **private** (§5.2).
- If M is labeled **private**[C] for some enclosing class or package C , then M' must be labeled **private**[C'] for some class or package C' where C' equals C or C' is contained in C .
- If M is labeled **protected**, then M' must also be labeled **protected**.
- If M' is not an abstract member, then M must be labeled **override**.
- If M' is incomplete (§5.2) in C then M must be labeled **abstract override**.

Example 5.1.7 Consider the definitions:

```
trait Root { type T <: Root }
trait A extends Root { type T <: A }
trait B extends Root { type T <: B }
trait C extends A with B
```

Then the class definition C is not well-formed because the binding of T in C is **type** $T <: B$, which fails to subsume the binding **type** $T <: A$ of T in type A . The problem can be solved by adding an overriding definition of type T in class C :

```
class C extends A with B { type T <: C }
```

5.2 Modifiers

Syntax:

```
Modifier      ::= LocalModifier
                | override
                | private [ "[" id "]" ]
                | protected [ "[" id "]" ]
LocalModifier ::= abstract
                | final
                | sealed
                | implicit
```

Member definitions may be preceded by modifiers which affect the accessibility and usage of the identifiers bound by them. If several modifiers are given, their order does not matter, but the same modifier may not occur repeatedly. Modifiers preceding a repeated definition apply to all constituent definitions. The rules governing the validity and meaning of a modifier are as follows.

- The **private** modifier can be used with any definition or declaration in a template. Such members can be accessed only from within the directly enclosing

template and its companion module. They are not inherited by subclasses and they may not override definitions in parent classes.

The modifier can be *qualified* with an identifier *C* (e.g. **private**[*C*]) that must denote a class or package enclosing the definition. Members labeled with such a modifier are accessible respectively only from code inside the package *C* or only from code inside the class *C* and its companion module (§5.4). Qualified private members are inherited only from templates inside *C*.

We say, a member *is private* if it is labeled with a **private** modifier without qualification. Unqualified **private** may not be applied to abstract members, and it may not be combined in one modifier list with **protected**, **final** or **override**.

- The **protected** modifier applies to class member definitions. Protected members can be accessed from within
 - the template of the defining class,
 - all templates that have the defining class as a base class,
 - the companion module of any of those classes.

A **protected** modifier can be qualified with an package identifier *C* (e.g. **protected**[*C*]) that must denote a class or package enclosing the definition. Members labeled with such a modifier are also accessible respectively from all code inside the package *C* or from all code inside the class *C* and its companion module (§5.4).

A protected identifier *x* may be used as a member name in a selection *r.x* only if one of the following applies:

- The access is within the template defining the member, or, if a qualification *C* is given, inside the package *C*, or the class *C*, or its companion module, or
 - *r* is one of the reserved words **this** and **super**, or
 - *r*'s type conforms to a type-instance of the class which contains the access.
- The **override** modifier applies to class member definitions or declarations. It is mandatory for member definitions or declarations that override some other concrete member definition in a parent class. If an **override** modifier is given, there must be at least one overridden member definition or declaration (either concrete or abstract).
 - The **override** modifier has a different significance when combined with the **abstract** modifier. That modifier combination is only allowed for value members of traits. A member labeled **abstract override** must override at least one other member and all members overridden by it must be incomplete.

We call a member *M* of a template *incomplete* if it is either abstract (i.e. defined by a declaration), or it is labeled **abstract** and **override** and every member overridden by *M* is again incomplete.

Note that the **abstract override** modifier combination does not influence the concept whether a member is concrete or abstract. A member is *abstract* if only a declaration is given for it; it is *concrete* if a full definition is given.

- The **abstract** modifier is used in class definitions. It is redundant for traits, and mandatory for all other classes which have incomplete members. Abstract classes cannot be instantiated (§6.7) with a constructor invocation unless followed by mixins and/or a refinement which override all incomplete members of the class.

The **abstract** modifier can also be used in conjunction with **override** for class member definitions. In that case the previous discussion applies.

- The **final** modifier applies to class member definitions and to class definitions. A **final** class member definition may not be overridden in subclasses. A **final** class may not be inherited by a template. **final** is redundant for object definitions. Members of final classes or objects are implicitly also final, so the **final** modifier is redundant for them, too. **final** may not be applied to incomplete members, and it may not be combined in one modifier list with **private** or **sealed**.
- The **sealed** modifier applies to class definitions. A **sealed** class may not be directly inherited, except if the inheriting template is defined the same source file as the inherited class. However, subclasses of a sealed class can be inherited anywhere.

Example 5.2.1 The following code illustrates the use of qualified private:

```
package outerpkg.innerpkg
class Outer {
  class Inner {
    private[Outer] def f()
    private[innerpkg] def g()
    private[outerpkg] def h()
  }
}
```

Here, accesses to the method *f* can appear anywhere within *OuterClass*, but not outside it. Accesses to method *g* can appear anywhere within the package *outerpkg.innerpkg*, as would be the case for package-private methods in Java. Finally, accesses to method *h* can appear anywhere within package *outerpkg*, including packages contained in it.

Example 5.2.2 A useful idiom to prevent clients of a class from constructing new instances of that class is to declare the class **abstract** and **sealed**:

```

object m {
  abstract sealed class C (x: Int) {
    def nextC = new C(x + 1) {}
  }
  val empty = new C(0) {}
}

```

For instance, in the code above clients can create instances of class `m.C` only by calling the `nextC` method of an existing `m.C` object; it is not possible for clients to create objects of class `m.C` directly. Indeed the following two lines are both in error:

```

new m.C(0)    // **** error: C is abstract, so it cannot be instantiated.
new m.C(0) {} // **** error: illegal inheritance from sealed class.

```

5.3 Class Definitions

Syntax:

```

TplDef      ::= class ClassDef
ClassDef    ::= id [TypeParamClause] ClassParamClauses
              [‘requires’ SimpleType] ClassTemplate
ClassTemplate ::= [‘extends’ TemplateParents] [[NewLine] TemplateBody]
ClassParamClauses ::= {ClassParamClause}
                  [[NewLine] ‘(’ implicit ClassParams ‘)’]
ClassParamClause ::= [NewLine] ‘(’ [ClassParams ‘)’]
ClassParams    ::= ClassParam {‘ ’ ClassParam}
ClassParam     ::= [{Modifier} (‘val’ | ‘var’)] Param
TemplateParents ::= Constr {‘with’ SimpleType}

```

The most general form of class definition is

class $c[tps](ps_1)\dots(ps_n)$ **requires** s **extends** t $(n \geq 0)$.

Here,

c is the name of the class to be defined.

tps is a non-empty list of type parameters of the class being defined. The scope of a type parameter is the whole class definition including the type parameter section itself. It is illegal to define two type parameters with the same name. The type parameter section $[tps]$ may be omitted. A class with a type parameter section is called *polymorphic*, otherwise it is called *monomorphic*.

$(ps_1)\dots(ps_n)$ are formal value parameter clauses for the *primary constructor* of the class. The scope of a formal value parameter includes the template t . However, a formal value parameter may not form part of the types of any of

the parent classes or members of the class template t . It is illegal to define two formal value parameters with the same name. If no formal parameter sections are given, an empty parameter section $()$ is assumed.

If a formal parameter declaration $x : T$ is preceded by a **val** or **var** keyword, an accessor (getter) definition (§4.2) for this parameter is implicitly added to the class. The getter introduces a value member x of class c that is defined as an alias of the parameter. If the introducing keyword is **var**, a setter accessor $x_=(e)$ (§4.2) is also implicitly added to the class. In invocation of that setter $x_=(e)$ changes the value of the parameter to the result of evaluating e . The formal parameter declaration may contain modifiers, which then carry over to the accessor definition(s). A formal parameter prefixed by **val** or **var** may not at the same time be a call-by-name parameter (§4.6).

s is the *self type* of the class. Inside the class, the type of **this** is assumed to be s . The self type must conform to the self types of all classes which are inherited by the template t . The self type declaration **requires** s may be omitted, in which case the self type of the class is assumed to be equal to $c[tps]$.

t is a template (§5.1) of the form

$$sc \text{ with } mt_1 \text{ with } \dots \text{ with } mt_m \{ stats \} \quad (m \geq 0)$$

which defines the base classes, behavior and initial state of objects of the class. The extends clause **extends** $sc \text{ with } mt_1 \text{ with } \dots \text{ with } mt_m$ can be omitted, in which case **extends** `scala.AnyRef` is assumed. The class body $\{stats\}$ may also be omitted, in which case the empty body $\{\}$ is assumed.

This class definition defines a type $c[tps]$ and a constructor which when applied to parameters conforming to types ps initializes instances of type $c[tps]$ by evaluating the template t .

Example 5.3.1 The following example illustrates **val** and **var** parameters of a class C :

```
class C(x: int, val y: String, var z: List[String])
val c = new C(1, "abc", List())
c.z = c.y :: c.z
```

5.3.1 Constructor Definitions

Syntax:

```
FunDef      ::= this ParamClause ParamClauses '=' ConstrExpr
ConstrExpr  ::= SelfInvocation
              | '{' SelfInvocation {StatementSeparator BlockStat} '}'
SelfInvocation ::= this ArgumentExprs {ArgumentExprs}
```

A class may have additional constructors besides the primary constructor. These are defined by constructor definitions of the form **def this**(ps_1)...(ps_n) = e . Such a definition introduces an additional constructor for the enclosing class, with parameters as given in the formal parameter lists ps_1, \dots, ps_n , and whose evaluation is defined by the constructor expression e . The scope of each formal parameter is the constructor expression e . A constructor expression is either a self constructor invocation **this**($args_1$)...($args_n$) or a block which begins with a self constructor invocation. The self constructor invocation must construct a generic instance of the class. I.e. if the class in question has name C and type parameters $[tps]$, then a self constructor invocation must generate an instance of $C[tps]$; it is not permitted to instantiate formal type parameters.

The signature and the self constructor invocation of a constructor definition are type-checked and evaluated in the scope which is in effect at the point of the enclosing class definition, augmented by any type parameters of the enclosing class. The rest of the constructor expression is type-checked and evaluated as a function body in the current class.

If there are auxiliary constructors of a class C , they form together with C 's primary constructor (§5.3) an overloaded constructor definition. The usual rules for overloading resolution (§6.22.3) apply for constructor invocations of C , including for the self constructor invocations in the constructor expressions themselves. However, unlike other methods, constructors are never inherited. To prevent infinite cycles of constructor invocations, there is the restriction that every self constructor invocation must refer to a constructor definition which precedes it (i.e. it must refer to either a preceding auxiliary constructor or the primary constructor of the class).

Example 5.3.2 Consider the class definition

```
class LinkedList[a]() {
  var head = _
  var tail = null
  def isEmpty = tail != null
  def this(head: a) = { this(); this.head = head }
  def this(head: a, tail: List[a]) = { this(head); this.tail = tail }
}
```

This defines a class `LinkedList` with three constructors. The second constructor constructs an singleton list, while the third one constructs a list with a given head and tail.

5.3.2 Case Classes

Syntax:

```
TplDef ::= case class ClassDef
```

If a class definition is prefixed with **case**, the class is said to be a *case class*.

The formal parameters in the first parameter section of a case class are called *elements*; they are treated specially. First, the value of such a parameter can be extracted as a field of a constructor pattern. Second, a **val** prefix is implicitly added to such a parameter, unless the parameter carries already a **val** or **var** modifier. Hence, an accessor definition for the parameter is generated (§5.3).

A case class definition of $c[tps](ps_1) \dots (ps_n)$ with type parameters tps and value parameters ps implicitly generates a function definition for a *case class factory* together with the class definition itself:

```
def c[tps](ps1)...(psn): s = new c[tps](xs1)...(xsn)
```

(Here, s is the self type of class c and each xs_i denotes the parameters of ps_i . If a type parameter section is missing in the class, it is also missing in the factory definition).

Every case class implicitly overrides some method definitions of class `scala.AnyRef` (§12.1) unless a definition of the same method is already given in the case class itself or a concrete definition of the same method is given in some base class of the case class different from `AnyRef`. In particular:

Method `equals: (Any)boolean` is structural equality, where two instances are equal if they belong to the same class and have equal (with respect to `equals`) elements.

Method `hashCode: ()int` computes a hash-code depending on the data structure in a way which maps equal (with respect to `equals`) values to equal hash-codes.

Method `toString: ()String` returns a string representation which contains the name of the class and its elements.

Example 5.3.3 Here is the definition of abstract syntax for lambda calculus:

```
class Expr
case class Var (x: String) extends Expr
case class Apply (f: Expr, e: Expr) extends Expr
case class Lambda(x: String, e: Expr) extends Expr
```

This defines a class `Expr` with case classes `Var`, `Apply` and `Lambda`. A call-by-value evaluator for lambda expressions could then be written as follows.

```
type Env = String => Value
case class Value(e: Expr, env: Env)

def eval(e: Expr, env: Env): Value = e match {
  case Var (x) =>
    env(x)
```

```

case Apply(f, g) =>
  val Value(Lambda (x, e1), env1) = eval(f, env)
  val v = eval(g, env)
  eval (e1, (y => if (y == x) v else env1(y)))
case Lambda(_, _) =>
  Value(e, env)
}

```

It is possible to define further case classes that extend type Expr in other parts of the program, for instance

```
case class Number(x: Int) extends Expr
```

This form of extensibility can be excluded by declaring the base class Expr **sealed**; in this case, all classes that directly extend Expr must be in the same source file as Expr.

5.3.3 Traits

Syntax:

```

TplDef      ::= trait TraitDef
TraitDef     ::= id [TypeParamClause]
               ['requires' SimpleType] [extends MixinParents]
               [[NewLine] TemplateBody]
MixinParents ::= SimpleType {'with' SimpleType}

```

A trait is a class that is meant to be added to some other class as a mixin. Unlike normal classes, traits cannot have constructor parameters. Furthermore, no constructor arguments are passed to its superclass. This is not necessary as traits are initialized after the superclass is initialized.

Assume a trait D defines some aspect of an instance x of type C (i.e. D is a base class of C). Then the *actual supertype* of D in x is the compound type consisting of all the base classes in $\mathcal{L}(C)$ that succeed D . The actual supertype gives the context for resolving a **super** reference in a trait (§6.4). Note that the actual supertype depends on the type to which the trait is added in a mixin composition; it is not statically known at the time the trait is defined.

If D is not a trait, then its actual supertype is simply its least proper supertype (which is statically known).

Example 5.3.4 The following trait defines the property of being comparable to objects of some type. It contains an abstract method `<` and default implementations of the other comparison operators `<=`, `>`, and `>=`.

```

trait Comparable[t <: Comparable[t]] requires t {
  def < (that: t): boolean

```

```

def <=(that: t): boolean = this < that || this == that
def > (that: t): boolean = that < this
def >=(that: t): boolean = that <= this
}

```

Example 5.3.5 Consider an abstract class `Table` that implements maps from a type of keys `A` to a type of values `B`. The class has a method `set` to enter a new key / value pair into the table, and a method `get` that returns an optional value matching a given key. Finally, there is a method `apply` which is like `get`, except that it returns a given default value if the table is undefined for the given key. This class is implemented as follows.

```

abstract class Table[A, B](defaultValue: B) {
  def get(key: A): Option[B]
  def set(key: A, value: B)
  def apply(key: A) = get(key) match {
    case Some(value) => value
    case None => defaultValue
  }
}

```

Here is a concrete implementation of the `Table` class.

```

class ListTable[A, B](defaultValue: B) extends Table[A, B](defaultValue) {
  private var elems: List[Pair[A, B]]
  def get(key: A) = elems.find(_.1==(key)).map(_.2)
  def set(key: A, value: B) = { elems = Pair(key, value) :: elems }
}

```

Here is a trait that prevents concurrent access to the `get` and `set` operations of its parent class:

```

trait SynchronizedTable[A, B] extends Table[A, B] {
  abstract override def get(key: A): B =
    synchronized { super.get(key) }
  abstract override def set((key: A, value: B) =
    synchronized { super.set(key, value) }
}

```

Note that `SynchronizedTable` does not pass an argument to its superclass, `Table`, even though `Table` is defined with a formal parameter. Note also that the **super** calls in `SynchronizedTable`'s `get` and `set` methods statically refer to abstract methods in class `Table`. This is legal, as long as the calling method is labeled **abstract override** (§5.2).

Finally, the following mixin composition creates a synchronized list table with strings as keys and integers as values and with a default value 0:


```
object MyTable extends ListTable[String, int](0) with SynchronizedTable
```

The object `MyTable` inherits its `get` and `set` method from `SynchronizedTable`. The **super** calls in these methods are re-bound to refer to the corresponding implementations in `ListTable`, which is the actual supertype of `SynchronizedTable` in `MyTable`.

5.4 Object Definitions

Syntax:

```
ObjectDef ::= id ClassTemplate
```

An object definition defines a single object of a new class. Its most general form is **object** *m* **extends** *t*. Here, *m* is the name of the object to be defined, and *t* is a template (§5.1) of the form

```
sc with mt1 with ... with mtn { stats }
```

which defines the base classes, behavior and initial state of *m*. The **extends** clause **extends** *sc with mt₁ with ... with mt_n* can be omitted, in which case **extends** `scala.AnyRef` is assumed. The class body `{stats}` may also be omitted, in which case the empty body `{}` is assumed.

The object definition defines a single object (or: *module*) conforming to the template *t*. It is roughly equivalent to the following three definitions, which together define a class and create a single object of that class on demand:

```
final class m$cls extends t
private var m$instance = null
final def m = {
  if (m$instance == null) m$instance = new m$cls
  m$instance
}
```

Here, the **final** modifiers are omitted if the definition occurs as part of a block. The names `m$cls` and `m$instance` are inaccessible for user programs.

Note that the value defined by an object definition is instantiated lazily. The **new** `m$cls` constructor is evaluated not at the point of the object definition, but is instead evaluated the first time *m* is dereferenced during execution of the program (which might be never at all). An attempt to dereference *m* again in the course of evaluation of the constructor leads to a infinite loop or run-time error.

However, the expansion given above is not accurate for top-level objects. It cannot be because variable and method definition cannot appear on the top-level. Instead, top-level objects are translated to static fields.

Example 5.4.1 Classes in Scala do not have static members; however, an equivalent effect can be achieved by an accompanying object definition E.g.

```
abstract class Point {  
  val x: Double  
  val y: Double  
  def isOrigin = (x == 0.0 && y == 0.0)  
}  
object Point {  
  val origin = new Point() { val x = 0.0; val y = 0.0 }  
}
```

This defines a class `Point` and an object `Point` which contains `origin` as a member. Note that the double use of the name `Point` is legal, since the class definition defines the name `Point` in the type name space, whereas the object definition defines a name in the term namespace.

This technique is applied by the Scala compiler when interpreting a Java class with static members. Such a class *C* is conceptually seen as a pair of a Scala class that contains all instance members of *C* and a Scala object that contains all static members of *C*.

Generally, a *companion module* of a class is an object which has the same name as the class and is defined in the same scope and compilation unit.

Chapter 6

Expressions

Syntax:

```
Expr      ::= [(Bindings | Id) '=>'] Expr
           | Expr1
Expr1     ::= if '(' Expr ')' [NewLine] Expr [[';'] else Expr]
           | try '{' block '}' [catch '{' CaseClauses '}']
             [finally Expr]
           | while '(' Expr ')' [NewLine] Expr
           | do Expr [StatementSeparator] while '(' Expr ')'
           | for '(' Enumerators ')' | '{' Enumerators '}'
             [NewLine] [yield] Expr
           | return [Expr]
           | throw Expr
           | SimpleExpr '.' id '=' Expr
           | SimpleExpr ArgumentExprs '=' Expr
           | PostfixExpr [':' Type1]
           | PostfixExpr match '{' CaseClauses '}'
           | MethodClosure
PostfixExpr ::= InfixExpr [id [NewLine]]
InfixExpr  ::= PrefixExpr
           | InfixExpr id [NewLine] PrefixExpr
PrefixExpr ::= ['- ' | '+ ' | '~ ' | '! ' | '&'] SimpleExpr
SimpleExpr ::= Literal
           | null
           | Path
           | '(' [Expr] ')'
           | BlockExpr
           | new Template
           | SimpleExpr '.' id
           | SimpleExpr TypeArgs
           | SimpleExpr ArgumentExprs
```

```

      | XmlExpr
ArgumentExprs ::= '(' [Exprs] ')'
      | BlockExpr
MethodClosure ::= '.' id {'.' id | TypeArgs | ArgumentExprs}
BlockExpr    ::= '{' CaseClauses '}'
      | '{' Block '}'
Block        ::= {BlockStat StatementSeparator} [ResultExpr]
ResultExpr   ::= Expr1
      | (Bindings | Id ':' Type1) '=>' Block
Exprs        ::= Expr {',' Expr}

```

Expressions are composed of operators and operands. Expression forms are discussed subsequently in decreasing order of precedence.

The typing of expressions is often relative to some *expected type* (which might be undefined). When we write “expression e is expected to conform to type T ”, we mean: (1) the expected type of e is T , and (2) the type of expression e must conform to T .

6.1 Literals

Syntax:

```
SimpleExpr ::= Literal
```

Typing of literals is as described in §1.3; their evaluation is immediate.

6.2 The *Null* Value

The **null** value is of type `scala.Null`, and is thus compatible with every reference type. It denotes a reference value which refers to a special “**null**” object. This object implements the methods in class `scala.AnyRef` as follows:

- `eq(x)`, `==(x)`, `equals(x)` return **true** iff their argument x is also the “null” object.
- `isInstanceOf[T]` always returns **false**.
- `asInstanceOf[T]` returns the “null” object itself if T conforms to `scala.AnyRef`, and throws a `NullPointerException` otherwise.
- `toString()` returns the string “null”.

A reference to any other member of the “null” object causes a `NullPointerException` to be thrown.

6.3 Designators

Syntax:

```
Designator ::= Path
            | SimpleExpr '.' id
```

A designator refers to a named term. It can be a *simple name* or a *selection*. If r is a stable identifier (§3.1) of type T , the selection $r.x$ refers statically to a term member m of r that is identified in T by the name x .

For other expressions e , $e.x$ is typed as if it was $\{ \text{val } y = e; y.x \}$, for some fresh name y . The typing rules for blocks implies that in that case x 's type may not refer to any abstract type member of e .

The expected type of a designator's prefix is always undefined. The type of a designator is the type of the entity it refers to, with the following exception: The type of a path (§3.1) p which occurs as the prefix of a selection, or which has a singleton type as expected type, is the singleton type $p.\text{type}$.

The selection $e.x$ is evaluated by first evaluating the qualifier expression e , which yields an object r , say. The selection's result is then the member r that is either defined by m or defined by a definition overriding m .

6.4 This and Super

Syntax:

```
SimpleExpr ::= [id '.'] this
            | [id '.'] super '[' id ']' '.' id
```

The expression **this** can appear in the statement part of a template or compound type. It stands for the object being defined by the innermost template or compound type enclosing the reference. If this is a compound type, the type of **this** is that compound type. If it is a template of an instance creation expression, the type of **this** is the type of that template. If it is a template of a class or object definition with simple name C , the type of this is the same as the type of $C.\text{this}$.

The expression $C.\text{this}$ is legal in the statement part of an enclosing class or object definition with simple name C . It stands for the object being defined by the innermost such definition. If the expression's expected type is a singleton type, or $C.\text{this}$ occurs as the prefix of a selection, its type is $C.\text{this.type}$, otherwise it is the self type of class C .

A reference **super**. m refers statically to a member m in the least proper supertype of the innermost template containing the reference. It evaluates to the member m' in the actual supertype of that template which is equal to m or which overrides

m . The statically referenced member m must be concrete, or the template containing the reference must have a member m' which overrides m and which is labeled **abstract override**.

A reference $C.\mathbf{super}.m$ refers statically to a member m in the least proper supertype of the innermost enclosing class or object definition named C which encloses the reference. It evaluates to the member m' in the actual supertype of that class or object which is equal to m or which overrides m . The statically referenced member m must be concrete, or the innermost enclosing class or object definition named C must have a member m' which overrides m and which is labeled **abstract override**.

The **super** prefix may be followed by a class qualifier $[C]$, as in $C.\mathbf{super}[C].x$. This is called a *static super reference*. In this case, the reference is to the member of x in the parent class of C whose simple name is M . That member must be uniquely defined and concrete.

Example 6.4.1 Consider the following class definitions

```
class Root { val x = "Root" }
class A extends Root { override val x = "A" ; val superA = super.x }
trait B extends Root { override val x = "B" ; val superB = super.x }
class C extends Root with B {
  override val x = "C" ; val superC = super.x }
}
class D extends A with B {
  override val x = "D" ; val superD = super.x }
}
```

The linearization of class C is $\{C, B, \text{Root}\}$ and the linearization of class D is $\{D, B, A, \text{Root}\}$. Then we have:

```
(new A).superA == "Root",
(new C).superA == "Root", (new C).superB = "Root", (new C).superC = "B",
(new D).superA == "Root", (new D).superB = "A",      (new D).superD = "B",
```

Note that the `superB` function returns different results depending on whether B is mixed in with class `Root` or `A`.

6.5 Function Applications

Syntax:

```
SimpleExpr ::= SimpleExpr ArgumentExprs
```

An application $f(e_1, \dots, e_n)$ applies the function f to the argument expressions

e_1, \dots, e_n . If f has a method type $(T_1, \dots, T_n)U$, the type of each argument expression e_i must conform to the corresponding parameter type T_i . If f has some value type, the application is taken to be equivalent to $f.\text{apply}(e_1, \dots, e_n)$, i.e. the application of an `apply` method defined by f .

Evaluation of $f(e_1, \dots, e_n)$ usually entails evaluation of f and e_1, \dots, e_n in that order. Each argument expression is converted to the type of its corresponding formal parameter. After that, the application is rewritten to the function's right hand side, with actual arguments substituted for formal parameters. The result of evaluating the rewritten right-hand side is finally converted to the function's declared result type, if one is given.

The case of a formal parameter with a parameterless method type $\Rightarrow T$ is treated specially. In this case, the corresponding actual argument expression is not evaluated before the application. Instead, every use of the formal parameter on the right-hand side of the rewrite rule entails a re-evaluation of the actual argument expression. In other words, the evaluation order for \Rightarrow -parameters is *call-by-name* whereas the evaluation order for normal parameters is *call-by-value*.

6.6 Type Applications

Syntax:

```
SimpleExpr ::= SimpleExpr '[' Types ']'
```

A type application $e[T_1, \dots, T_n]$ instantiates a polymorphic value e of type $[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]S$ with argument types T_1, \dots, T_n . Every argument type T_i must obey the corresponding bounds L_i and U_i . That is, for each $i = 1, \dots, n$, we must have $\sigma L_i <: T_i <: \sigma U_i$, where σ is the substitution $[a_1 := T_1, \dots, a_n := T_n]$. The type of the application is σS .

If the function part e is of some value type, the type application is taken to be equivalent to $e.\text{apply}[T_1, \dots, T_n]$, i.e. the application of an `apply` method defined by e .

Type applications can be omitted if local type inference (§6.22.4) can infer best type parameters for a polymorphic functions from the types of the actual function arguments and the expected result type.

6.7 Instance Creation Expressions

Syntax:

```
SimpleExpr ::= new Template
```

A simple instance creation expression is of the form `new c` where c is a constructor

invocation (§5.1.1). Let T be the type of c . Then T must denote a (a type instance of) a non-abstract subclass of `scala.AnyRef` which conforms to its self type (§5.3). The expression is evaluated by creating a fresh object of type T which is initialized by evaluating c . The type of the expression is T .

A general instance creation expression is of the form `new t` for some template t (§5.1). Such an expression is equivalent to the block

```
{ class a extends t; new a }
```

where a is a fresh name of an *anonymous class*.

6.8 Blocks

Syntax:

```
BlockExpr ::= '{' Block '}'
Block      ::= [{BlockStat StatementSeparator} ResultExpr]
```

A block expression $\{s_1; \dots; s_n; e\}$ is constructed from a sequence of block statements s_1, \dots, s_n and a final expression e . The statement sequence may not contain two definitions or declarations that bind the same name in the same namespace. The final expression can be omitted, in which case the unit value `()` is assumed.

The expected type of the final expression e is the expected type of the block. The expected type of all preceding statements is undefined.

The type of a block $s_1; \dots; s_n; e$ is usually the type of e . That type must be equivalent to a type which does not refer to an entity defined locally in the block. If this condition is violated, there are two other possibilities:

1. If a fully defined expected type is given, the type of the block is instead assumed to be the expected type.
2. Otherwise, if the type of e is an anonymous class a introduced by the expansion of an instance creation expression (§6.7), the type of the block is taken to be the least class type or refinement type which is a proper supertype of the type a .

It is a compile-time error if neither of the previous two clauses applies.

Evaluation of the block entails evaluation of its statement sequence, followed by an evaluation of the final expression e , which defines the result of the block.

Example 6.8.1 Written in isolation, the block

```
{ class C extends B {...} ; new C }
```


is illegal, since its type refers to class *C*, which is defined locally in the block.

However, when used in a definition such as

```
val x: B = { class C extends B {...} ; new C }
```

the block is well-formed, since the problematic type *C* can be replaced by the expected type *B*.

6.9 Prefix, Infix, and Postfix Operations

Syntax:

```
PostfixExpr ::= InfixExpr [id]
InfixExpr   ::= PrefixExpr
              | InfixExpr id PrefixExpr
PrefixExpr  ::= ['- ' | '+ ' | '! ' | '~ ' | '&'] SimpleExpr
```

Expressions can be constructed from operands and operators. A prefix operation *op e* consists of a prefix operator *op*, which must be one of the identifiers '+', '-', '!', '~' or '&'. For every prefix operator *op* other than '&', the expression *op e* is equivalent to the postfix method application *e.op*. The expression *&e* is well-formed if *e* is of method type. It represents *e* converted to a function type by eta expansion (§6.22.5).

Prefix operators are different from normal function applications in that their operand expression need not be atomic. For instance, the input sequence *-sin(x)* is read as *-(sin(x))*, whereas the function application *negate sin(x)* would be parsed as the application of the infix operator *sin* to the operands *negate* and *(x)*.

An infix or postfix operator can be an arbitrary identifier. Infix operators have precedence and associativity defined as follows:

The *precedence* of an infix operator is determined by the operator's first character. Characters are listed below in increasing order of precedence, with characters on the same line having the same precedence.

```
(all letters)
|
^
&
< >
= !
:
+ -
* / %
(all other special characters)
```

That is, operators starting with a letter have lowest precedence, followed by opera-

tors starting with ‘|’, etc.

The *associativity* of an operator is determined by the operator’s last character. Operators ending with a colon ‘:’ are right-associative. All other operators are left-associative.

Precedence and associativity of operators determine the grouping of parts of an expression as follows.

- If there are several infix operations in an expression, then operators with higher precedence bind more closely than operators with lower precedence.
- If there are consecutive infix operations $e_0 \text{ op}_1 e_1 \text{ op}_2 \dots \text{op}_n e_n$ with operators $\text{op}_1, \dots, \text{op}_n$ of the same precedence, then all these operators must have the same associativity. If all operators are left-associative, the sequence is interpreted as $(\dots (e_0 \text{ op}_1 e_1) \text{ op}_2 \dots) \text{op}_n e_n$. Otherwise, if all operators are right-associative, the sequence is interpreted as $e_0 \text{ op}_1 (e_1 \text{ op}_2 (\dots \text{op}_n e_n) \dots)$.
- Postfix operators always have lower precedence than infix operators. E.g. $e_1 \text{ op}_1 e_2 \text{ op}_2$ is always equivalent to $(e_1 \text{ op}_1 e_2) \text{ op}_2$.

A postfix operation $e \text{ op}$ is interpreted as $e.\text{op}$. A left-associative binary operation $e_1 \text{ op } e_2$ is interpreted as $e_1.\text{op}(e_2)$. If op is right-associative, the same operation is interpreted as $(\text{val } x = e_1; e_2.\text{op}(x))$, where x is a fresh name.

6.10 Typed Expressions

Syntax:

$\text{Expr1} \quad ::= \text{PostfixExpr} \text{ [‘:’ Type1]}$

The typed expression $e : T$ has type T . The type of expression e is expected to conform to T . The result of the expression is the value of e converted to type T .

Example 6.10.1 Here are examples of well-typed and illegally typed expressions.

```
1: int           // legal, of type int
1: long          // legal, of type long
// 1: string     // ***** illegal
```

6.11 Method closures

Syntax:

$\text{MethodClosure} \quad ::= \text{‘.’ id} \{ \text{‘.’ id} \mid \text{TypeArgs} \mid \text{ArgumentExprs} \}$

A method closure *.id* starts with a period and an identifier, which may be followed by selections and type- and value-arguments. This expression is equivalent to an anonymous function $x \Rightarrow x.id$ where x is a fresh parameter name. No type for x is given; hence this type needs to be inferable from the context of the expression.

Example 6.11.1 The following method returns the n 'th column of a given list of row-lists *xss*, using methods *map*, *drop* and *head* defined in class *scala.List*.

```
def column[T](xss: List[List[T]], n: Int): List[T] =
  xss.map(_.drop(n-1)).map(_.head)
```

6.12 Assignments

Syntax:

```
Expr1      ::= Designator '=' Expr
              | SimpleExpr ArgumentExprs '=' Expr
```

The interpretation of an assignment to a simple variable $x = e$ depends on the definition of x . If x denotes a mutable variable, then the assignment changes the current value of x to be the result of evaluating the expression e . The type of e is expected to conform to the type of x . If x is a parameterless function defined in some template, and the same template contains a setter function $x_=(e)$ as member, then the assignment $x = e$ is interpreted as the invocation $x_=(e)$ of that setter function. Analogously, an assignment $f.x = e$ to a parameterless function x is interpreted as the invocation $f.x_=(e)$.

An assignment $f(args) = e$ with a function application to the left of the “=” operator is interpreted as $f.update(args, e)$, i.e. the invocation of an update function defined by f .

Example 6.12.1 Here is the usual imperative code for matrix multiplication.

```
def matmul(xss: Array[Array[Double]], yss: Array[Array[Double]]) = {
  val zss: Array[Array[Double]] = new Array(xss.length, yss.length)
  var i = 0
  while (i < xss.length) {
    var j = 0
    while (j < yss(0).length) {
      var acc = 0.0
      var k = 0
      while (k < yss.length) {
        acc = acc + xss(i)(k) * yss(k)(j)
        k = k + 1
      }
    }
  }
```

```

        zss(i)(j) = acc
        j = j + 1
    }
    i = i + 1
}
zss
}

```

Desugaring the array accesses and assignments yields the following expanded version:

```

def matmul(xss: Array[Array[Double]], yss: Array[Array[Double]]) = {
  val zss: Array[Array[Double]] = new Array(xss.length, yss.length)
  var i = 0
  while (i < xss.length) {
    var j = 0
    while (j < yss(0).length) {
      var acc = 0.0
      var k = 0
      while (k < yss.length) {
        acc = acc + xss.apply(i).apply(k) * yss.apply(k).apply(j)
        k = k + 1
      }
      zss.apply(i).update(j, acc)
      j = j + 1
    }
    i = i + 1
  }
  zss
}

```

6.13 Conditional Expressions

Syntax:

```
Expr1 ::= if '(' Expr ')' [NewLine] Expr [[';']] else Expr]
```

The conditional expression **if** (e_1) e_2 **else** e_3 chooses one of the values of e_2 and e_3 , depending on the value of e_1 . The condition e_1 is expected to conform to type `boolean`. The then-part e_2 and the else-part e_3 are both expected to conform to the expected type of the conditional expression. The type of the conditional expression is the least upper bound of the types of e_1 and e_2 . A semicolon preceding the **else** symbol of a conditional expression is ignored.

The conditional expression is evaluated by evaluating first e_1 . If this evaluates to

true, the result of evaluating e_2 is returned, otherwise the result of evaluating e_3 is returned.

A short form of the conditional expression eliminates the else-part. The conditional expression **if** (e_1) e_2 is evaluated as if it was **if** (e_1) e_2 **else** (). The type of this expression is `unit` and the then-part e_2 is also expected to conform to type `unit`.

6.14 While Loop Expressions

Syntax:

```
Expr1      ::=  while '(' Expr ')' Expr
```

The while loop expression **while** (e_1) e_2 is typed and evaluated as if it was an application of `whileLoop` (e_1) (e_2) where the hypothetical function `whileLoop` is defined as follows.

```
def whileLoop(cond: => Boolean)(body: => Unit): Unit =
  if (cond) { body ; whileLoop(cond)(body) } else {}
```

6.15 Do Loop Expressions

Syntax:

```
Expr1      ::=  do Expr [StatementSeparator] while '(' Expr ')'
```

The do loop expression **do** e_1 **while** (e_2) is typed and evaluated as if it was the expression (e_1 ; **while** (e_2) e_1). A semicolon preceding the **while** symbol of a do loop expression is ignored.

6.16 For-Comprehensions

Syntax:

```
Expr1      ::=  for '(' Enumerators ')' [yield] Expr
Enumerators ::=  Generator {StatementSeparator Enumerator}
Enumerator ::=  Generator
               |  val Pattern1 '=' Expr
               |  Expr
Generator   ::=  val Pattern1 '<- ' Expr
```

A comprehension **for** ($enums$) **yield** e evaluates expression e for each binding generated by the enumerators $enums$. An enumerator sequence always starts with

a generator; this can be followed by further generators, value definitions, or filters. A *generator* **val** $p \leftarrow e$ produces bindings from an expression e which is matched in some way against pattern p . A *value definition* binds a value name (or several names in a pattern) to the result of evaluating an expression. A *filter* is an expression which restricts enumerated bindings. The precise meaning of generators and filters is defined by translation to invocations of four methods: `map`, `filter`, `flatMap`, and `foreach`. These methods can be implemented in different ways for different carrier types.

The translation scheme is as follows. In a first step, every generator **val** $p \leftarrow e$, where p is not irrefutable (§8.1) for the type of e , is replaced by

```
val  $p \leftarrow e$ .filter { case  $p \Rightarrow$  true; case  $\_ \Rightarrow$  false }
```

Then, the following rules are applied repeatedly until all comprehensions have been eliminated.

- A for-comprehension **for** (**val** $p \leftarrow e$) **yield** e' is translated to e .map { **case** $p \Rightarrow e'$ }.
- A for-comprehension **for** (**val** $p \leftarrow e$) e' is translated to e .foreach { **case** $p \Rightarrow e'$ }.
- A for-comprehension

```
for (val  $p \leftarrow e$ ; val  $p' \leftarrow e' \dots$ ) yield  $e''$  ,
```

where \dots is a (possibly empty) sequence of generators or filters, is translated to

```
 $e$ .flatMap { case  $p \Rightarrow$  for (val  $p' \leftarrow e' \dots$ ) yield  $e''$  } .
```

- A for-comprehension

```
for (val  $p \leftarrow e$ ; val  $p' \leftarrow e' \dots$ )  $e''$  .
```

where \dots is a (possibly empty) sequence of generators or filters, is translated to

```
 $e$ .foreach { case  $p \Rightarrow$  for (val  $p' \leftarrow e' \dots$ )  $e''$  } .
```

- A generator **val** $p \leftarrow e$ followed by a filter f is translated to a single generator **val** $p \leftarrow e$.filter($(x_1, \dots, x_n) \Rightarrow f$) where x_1, \dots, x_n are the free variables of p .
- A generator **val** $p \leftarrow e$ followed by a value definition **val** $p' = e'$ is translated to the following generator of pairs of values, where x and x' are fresh names:

```

val Pair( $p$ ,  $p'$ ) <-
  for (val  $x@p$  <-  $e$ ) yield { val  $x'@p' = e'$ ; Pair( $x$ ,  $x'$ ) }

```

Example 6.16.1 the following code produces all pairs of numbers between 1 and $n - 1$ whose sums are prime.

```

for { val  $i$  <- range(1,  $n$ )
      val  $j$  <- range(1,  $i$ )
      isPrime( $i+j$ )
    } yield Pair ( $i$ ,  $j$ )

```

The for-comprehension is translated to:

```

range(1,  $n$ )
  .flatMap {
    case  $i \Rightarrow$  range(1,  $i$ )
      .filter {  $j \Rightarrow$  isPrime( $i+j$ ) }
      .map { case  $j \Rightarrow$  Pair( $i$ ,  $j$ ) } }

```

Example 6.16.2 For comprehensions can be used to express vector and matrix algorithms concisely. For instance, here is a function to compute the transpose of a given matrix:

```

def transpose[ $a$ ](xss: Array[Array[ $a$ ]]) {
  for (val  $i$  <- Array.range(0, xss(0).length)) yield
    Array(for (val  $xs$  <- xss) yield  $xs(i)$ )
}

```

Here is a function to compute the scalar product of two vectors:

```

def scalprod(xs: Array[double], ys: Array[double]) {
  var acc = 0.0
  for (val Pair( $x$ ,  $y$ ) <- xs zip ys) acc = acc +  $x * y$ 
  acc
}

```

Finally, here is a function to compute the product of two matrices. Compare with the imperative version of Example 6.12.1.

```

def matmul(xss: Array[Array[double]], yss: Array[Array[double]]) = {
  val ysst = transpose(yss)
  for (val  $xs$  <- xss) yield
    for (val  $yst$  <- ysst) yield
      scalprod( $xs$ ,  $yst$ )
}

```

The code above makes use of the fact that map, flatmap, filter, and foreach are defined for members of class `scala.Array`.

6.17 Return Expressions

Syntax:

Expr1 ::= **return** [Expr]

A return expression **return** *e* must occur inside the body of some enclosing named method or function *f*. This function must have an explicitly declared result type, and the type of *e* must conform to it. The return expression evaluates the expression *e* and returns its value as the result of *f*. The evaluation of any statements or expressions following the return expression is omitted. The type of a return expression is `scala.Bottom`.

6.18 Throw Expressions

Syntax:

Expr1 ::= **throw** Expr

A throw expression **throw** *e* evaluates the expression *e*. The type of this expression must conform to `Throwable`. If *e* evaluates to an exception reference, evaluation is aborted with the thrown exception. If *e* evaluates to **null**, evaluation is instead aborted with a `NullPointerException`. If there is an active **try** expression (§6.19) which handles the thrown exception, evaluation resumes with the handler; otherwise the thread executing the **throw** is aborted. The type of a throw expression is `scala.Bottom`.

6.19 Try Expressions

Syntax:

Expr1 ::= **try** '{' Block '}' [**catch** Expr] [**finally** Expr]

A try expression **try** { *b* } **catch** *e* evaluates the block *b*. If evaluation of *b* does not cause an exception to be thrown, the result of *b* is returned. Otherwise the handler *e* is applied to the thrown exception. Let *pt* be the expected type of the try expression. The block *b* is expected to conform to *pt*. The handler *e* is expected to conform to type `scala.PartialFunction[scala.Throwable, pt]`. The type of the try expression is the least upper bound of the type of *b* and the result type of *e*.

A try expression **try** { *b* } **finally** *e* evaluates the block *b*. If evaluation of *b* does not cause an exception to be thrown, the expression *e* is evaluated. If an exception is thrown during evaluation of *e*, the evaluation of the try expression is aborted with the thrown exception. If no exception is thrown during evaluation of *e*, the


```

f => g => x => f(g(x))           // Curried function composition

(x: Int,y: Int) => x + y         // A summation function

() => { count = count + 1; count } // The function which takes an
                                   // empty parameter list (),
                                   // increments a non-local variable
                                   // 'count' and returns the new value.

```

6.21 Statements

Syntax:

```

BlockStat ::= Import
           | [implicit] Def
           | {LocalModifier} TmplDef
           | Expr1
           |
TemplateStat ::= Import
              | {AttributeClause} {Modifier} Def
              | {AttributeClause} {Modifier} Dcl
              | Expr
              |

```

Statements occur as parts of blocks and templates. A statement can be an import, a definition or an expression, or it can be empty. Statements used in the template of a class definition can also be declarations. An expression that is used as a statement can have an arbitrary value type. An expression statement e is evaluated by evaluating e and discarding the result of the evaluation.

Block statements may be definitions which bind local names in the block. The only modifiers allowed in block-local definitions are modifiers **abstract**, **final**, or **sealed** preceding a class or object definition.

Evaluation of a statement sequence entails evaluation of the statements in the order they are written.

6.22 Implicit Conversions

Implicit conversions can be applied to expressions whose type does not match their expected type, as well as to unapplied methods. The available implicit conversions are given in the next two sub-sections.

We say, a type T is *compatible* to a type U if T conforms to U after applying eta-expansion (§6.22.5) and view applications (§7.3).

6.22.1 Value Conversions

The following five implicit conversions can be applied to an expression e which has some value type T and which is type-checked with some expected type pt .

Overloading Resolution. If an expression denotes several possible members of a class, overloading resolution (§6.22.3) is applied to pick a unique member.

Type Instantiation. An expression e of polymorphic type

$$[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]T$$

which does not appear as the function part of a type application is converted to a type instance of T by determining with local type inference (§6.22.4) instance types T_1, \dots, T_n for the type variables a_1, \dots, a_n and implicitly embedding e in the type application $e[T_1, \dots, T_n]$ (§6.6).

Numeric Literal Narrowing. If the expected type is byte, short or char, and the expression e is an integer literal fitting in the range of that type, it is converted to the same literal in that type.

Value Discarding. If e has some value type and the expected type is unit, e is converted to the expected type by embedding it in the term $\{ e; () \}$.

View Application. If none of the previous conversions applies, and the e 's type does not conform to the expected type pt , it is attempted to convert e to the expected type with a view (§7.3).

6.22.2 Method Conversions

The following four implicit conversions can be applied to methods which are not applied to some argument list.

Evaluation. A parameterless method m of type $\Rightarrow T$ is always converted to type T by evaluating the expression to which m is bound.

Implicit Application. If the method takes only implicit parameters, implicit arguments are passed following the rules of §7.2.

Eta Expansion. Otherwise, if the method is not a constructor, and the expected type pt is a function type $(Ts') \Rightarrow T'$ eta-expansion (§6.22.5) is performed on the expression e .

Empty Application. Otherwise, if e has method type $()T$ where T is compatible with pt , it is implicitly applied to the empty argument list, yielding $e()$.

6.22.3 Overloading Resolution

If an identifier or selection e references several members of a class, the context of the reference is used to identify a unique member. The way this is done depends on whether or not e is used as a function. Let \mathcal{A} be the set of members referenced by e .

Assume first that e appears as a function in an application, as in $e(args)$. If there is precisely one alternative in \mathcal{A} which is a (possibly polymorphic) method type whose arity matches the number of arguments given, that alternative is chosen.

Otherwise, let Ts be the vector of types obtained by typing each argument with an undefined expected type. One determines first the set of applicable alternatives. A method type alternative is *applicable* if each type in Ts is compatible with the corresponding formal parameter type in the alternative, and, if the expected type is defined, the method's result type is compatible to it. A polymorphic method type is applicable if local type inference can determine type arguments so that the instantiated method type is applicable.

Let \mathcal{B} be the set of applicable alternatives. It is an error if \mathcal{B} is empty. Otherwise, one chooses the *most specific* alternative among the alternatives in \mathcal{B} , according to the following definition of being “more specific”.

- A method type $(Ts)U$ is more specific than some other type S if S is applicable to arguments (ps) of types Ts .
- A polymorphic method type $[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]T$ is more specific than some other type S if T is more specific than S under the assumption that for $i = 1, \dots, n$ each a_i is an abstract type name bounded from below by L_i and from above by U_i .
- Any other type is always more specific than a parameterized method type or a polymorphic type.

It is an error if there is no unique alternative in \mathcal{B} which is more specific than all other alternatives in \mathcal{B} .

Assume next that e appears as a function in a type application, as in $e[targs]$. Then we choose all alternatives in \mathcal{A} which take the same number of type parameters as there are type arguments in $targs$. It is an error if no such alternative exists. If there are several such alternatives overloading resolution is applied again to the whole expression $e[targs]$.

Assume finally that e does not appear as a function in either an application or a type application. If an expected type is given, let \mathcal{B} be the set of those alternatives in \mathcal{A} which are compatible (§6.22) to it. Otherwise, let \mathcal{B} be the same as \mathcal{A} . We choose in this case the most specific alternative among all alternatives in \mathcal{B} . It is an error if there is no unique alternative in \mathcal{B} which is more specific than all other alternatives in \mathcal{B} .

Example 6.22.1 Consider the following definitions:

```
class A extends B {}
def f(x: B, y: B) = ...
def f(x: A, y: B) = ...
val a: A
val b: B
```

Then the application $f(b, b)$ refers to the first definition of f whereas the application $f(a, a)$ refers to the second. Assume now we add a third overloaded definition

```
def f(x: B, y: A) = ...
```

Then the application $f(a, a)$ is rejected for being ambiguous, since no most specific applicable signature exists.

6.22.4 Local Type Inference

Local type inference infers type arguments to be passed to expressions of polymorphic type. Say e is of type $[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]T$ and no explicit type parameters are given.

Local type inference converts this expression to a type application $e[T_1, \dots, T_n]$. The choice of the type arguments T_1, \dots, T_n depends on the context in which the expression appears and on the expected type pt . There are three cases.

Case 1: Selections. If the expression appears as the prefix of a selection with a name x , then type inference is *deferred* to the whole expression $e.x$. That is, if $e.x$ has type S , it is now treated as having type $[a_1 >: L_1 <: U_1, \dots, a_n >: L_n <: U_n]S$, and local type inference is applied in turn to infer type arguments for a_1, \dots, a_n , using the context in which $e.x$ appears.

Case 2: Values. If the expression e appears as a value without being applied to value arguments, the type arguments are inferred by solving a constraint system which relates the expression's type T with the expected type pt . Without loss of generality we can assume that T is a value type; if it is a method type we apply eta-expansion (§6.22.5) to convert it to a function type. Solving means finding a substitution σ of types T_i for the type parameters a_i such that

- All type parameter bounds are respected, i.e. $\sigma L_i <: \sigma a_i$ and $\sigma a_i <: \sigma U_i$ for $i = 1, \dots, n$.
- The expression's type conforms to the expected type, i.e. $\sigma T <: \sigma pt$.

It is a compile time error if no such substitution exists. If several substitutions exist, local-type inference will choose for each type variable a_i a minimal or maximal type T_i of the solution space. A *maximal* type T_i will be chosen if the type parameter a_i appears contravariantly (§4.5) in the type T of the expression. A *minimal* type T_i will be chosen in all other situations, i.e. if the variable appears covariantly, non-variantly or not at all in the type T . We call such a substitution an *optimal solution* of the given constraint system for the type T .

Case 3: Methods. The last case applies if the expression e appears in an application $e(d_1, \dots, d_m)$. In that case T is a method type $(R_1, \dots, R_m)T'$. Without loss of generality we can assume that the result type T' is a value type; if it is a method type we apply eta-expansion (§6.22.5) to convert it to a function type. One computes first the types S_j of the argument expressions d_j , using two alternative schemes. Each argument expression d_j is typed first with the expected type R_j , in which the type parameters a_1, \dots, a_n are taken as type constants. If this fails, the argument d_j is typed instead with an expected type R'_j which results from R_j by replacing every type parameter in a_1, \dots, a_n with *undefined*.

In a second step, type arguments are inferred by solving a constraint system which relates the method's type with the expected type pt and the argument types S_1, \dots, S_m . Solving the constraint system means finding a substitution σ of types T_i for the type parameters a_i such that

- All type parameter bounds are respected, i.e. $\sigma L_i <: \sigma a_i$ and $\sigma a_i <: \sigma U_i$ for $i = 1, \dots, n$.
- The method's result type T' conforms to the expected type, i.e. $\sigma T' <: \sigma pt$.
- Each argument type conforms to the corresponding formal parameter type, i.e. $\sigma S_j <: \sigma R_j$ for $j = 1, \dots, m$.

It is a compile time error if no such substitution exists. If several solutions exist, an optimal one for the type T' is chosen.

All or parts of an expected type pt may be undefined. The rules for conformance (§3.5.2) are extended to this case by adding the rule that for any type T the following two statements are always true:

$$\text{undefined} <: T \quad \text{and} \quad T <: \text{undefined}.$$

It is possible that no minimal or maximal solution for a type variable exists, in which case a compile-time error results. Because $<:$ is a pre-order, it is also possible that a

solution set has several optimal solutions for a type. In that case, a Scala compiler is free to pick any one of them.

Example 6.22.2 Consider the two methods:

```
def cons[a](x: a, xs: List[a]): List[a] = x :: xs
def nil[b]: List[b] = Nil
```

and the definition

```
val xs = cons(1, nil) .
```

The application of `cons` is typed with an undefined expected type. This application is completed by local type inference to `cons[int](1, nil)`. Here, one uses the following reasoning to infer the type argument `int` for the type parameter `a`:

First, the argument expressions are typed. The first argument `1` has type `int` whereas the second argument `nil` is itself polymorphic. One tries to type-check `nil` with an expected type `List[a]`. This leads to the constraint system

```
List[b?] <: List[a]
```

where we have labeled `b?` with a question mark to indicate that it is a variable in the constraint system. Because class `List` is covariant, the optimal solution of this constraint is

```
b = scala.Bottom .
```

In a second step, one solves the following constraint system for the type parameter `a` of `cons`:

```
int <: a?
List[scala.Bottom] <: List[a?]
List[a?] <: undefined
```

The optimal solution of this constraint system is

```
a = int ,
```

so `int` is the type inferred for `a`.

Example 6.22.3 Consider now the definition

```
val ys = cons("abc", xs)
```

where `xs` is defined of type `List[int]` as before. In this case local type inference proceeds as follows.

First, the argument expressions are typed. The first argument `"abc"` has type `String`. The second argument `xs` is first tried to be typed with expected type

List[a]. This fails, as List[int] is not a subtype of List[a]. Therefore, the second strategy is tried; xs is now typed with expected type List[undefined]. This succeeds and yields the argument type List[int].

In a second step, one solves the following constraint system for the type parameter a of cons:

```
String <: a?
List[int] <: List[a?]
List[a?] <: undefined
```

The optimal solution of this constraint system is

```
a = scala.Any ,
```

so scala.Any is the type inferred for a.

6.22.5 Eta Expansion

Eta-expansion converts an expression of method type to an equivalent expression of function type. It proceeds in two steps.

First, one identifies the maximal sub-expressions of e ; let's say these are e_1, \dots, e_m . For each of these, one creates a fresh name x_i . Let e' be the expression resulting from replacing every maximal subexpression e_i in e by the corresponding fresh name x_i . Second, one creates a fresh name y_i for every argument type T_i of the method ($i = 1, \dots, n$). The result of eta-conversion is then:

```
{ val x1 = e1;
  ...
  val xm = em;
  (y1: T1, ..., yn: Tn) => e'(y1, ..., yn)
}
```

Eta expansion is not applicable to methods with call-by-name parameters $x: \Rightarrow T$ or repeated parameters $x: T^*$ (§4.6), because its result type would not be expressible as a function type. Hence, methods with such parameters always need to be applied to arguments immediately.

Chapter 7

Implicit Parameters and Views

7.1 The Implicit Modifier

Syntax:

```
Modifier      ::= implicit
ParamClauses ::= {[NewLine] '(' [Param {',' Param}] ')'}
               {[NewLine] '(' implicit Param {',' Param} ')')}
```

Template members and parameters labeled with an **implicit** modifier can be passed to implicit parameters (§7.2) and can be used as implicit conversions called views (§7.3). The **implicit** modifier is illegal for all type members, as well as for top-level (§9.2) objects.

Example 7.1.1 The following code defines an abstract class of monoids and two concrete implementations, `StringMonoid` and `IntMonoid`. The two implementations are marked implicit.

```
abstract class Monoid[a] extends SemiGroup[a] {
  def unit: a
}
object Monoids {
  implicit object StringMonoid extends Monoid[String] {
    def add(x: String, y: String): String = x.concat(y)
    def unit: String = ""
  }
  implicit object IntMonoid extends Monoid[Int] {
    def add(x: Int, y: Int): Int = x + y
    def unit: Int = 0
  }
}
```

7.2 Implicit Parameters

An implicit parameter list (**implicit** p_1, \dots, p_n) marks the parameters p_1, \dots, p_n as implicit. A method or constructor can have only one implicit parameter list, and it must be the last parameter list given.

A method with implicit parameters can be applied to arguments just like a normal method. In this case the **implicit** label has no effect. However, if such a method misses arguments for its implicit parameters, such arguments will be automatically provided.

The actual arguments that are eligible to be passed to an implicit parameter of type T fall into two categories. First, eligible are all identifiers x that can be accessed at the point of the method call without a prefix and that denote an implicit definition (§7.1) or an implicit parameter. An eligible identifier may thus be a local name, or a member of an enclosing template, or it may have been made accessible without a prefix through an import clause (§4.7). Second, eligible are also all **implicit** members of some object that belongs to the implicit scope of the implicit parameter's type, T .

The *implicit scope* of a type T consists of all companion modules (§5.4) of classes that are associated with the implicit parameter's type. Here, we say a class C is *associated* with a type T , if it is a base class (§5.1.2) of some part of T . The *parts* of a type T are:

- if T is a compound type T_1 **with** ... **with** T_n , the union of the parts of T_1, \dots, T_n , as well as T itself,
- if T is a parameterized type $S[T_1, \dots, T_n]$, the union of the parts of S and T_1, \dots, T_n ,
- if T is a singleton type p .**type**, the parts of the type of p ,
- if T is a type projection $S\#U$, the parts of S as well as T itself,
- in all other cases, just T itself.

If there are several eligible arguments which match the implicit parameter's type, a most specific one will be chosen using the rules of static overloading resolution (§6.22.3).

Example 7.2.1 Assuming the classes from Example 7.1.1, here is a method which computes the sum of a list of elements using the monoid's `add` and `unit` operations.

```
def sum[a](xs: List)(implicit m: Monoid[a]): a =
  if (xs.isEmpty) m.unit
  else m.add(xs.head, sum(xs.tail))
```

The monoid in question is marked as an implicit parameter, and can therefore be inferred based on the type of the list. Consider for instance the call

```
sum(List(1, 2, 3))
```

in a context where `stringMonoid` and `intMonoid` are visible. We know that the formal type parameter `a` of `sum` needs to be instantiated to `Int`. The only eligible object which matches the implicit formal parameter type `Monoid[Int]` is `intMonoid` so this object will be passed as implicit parameter.

This discussion also shows that implicit parameters are inferred after any type arguments are inferred (§6.22.4).

Implicit methods can themselves have implicit parameters. An example is the following method from module `scala.List`, which injects lists into the `scala.Ordered` class, provided the element type of the list is also convertible to this type.

```
implicit def list2ordered[a](x: List[a])
  (implicit elem2ordered: a => Ordered[a]): Ordered[List[a]] =
  ...
```

Assume in addition a method

```
implicit def int2ordered(x: int): Ordered[int]
```

that injects integers into the `Ordered` class. We can now define a `sort` method over ordered lists:

```
sort(xs: List[a])(implicit a2ordered: a => Ordered[a]) = ...
```

We can apply `sort` to a list of lists of integers `yss: List[List[int]]` as follows:

```
sort(yss)
```

The call above will be completed by passing two nested implicit arguments:

```
sort(yss)(xs: List[int] => list2ordered[int](xs)(int2ordered)) .
```

The possibility of passing implicit arguments to implicit arguments raises the possibility of an infinite recursion. For instance, one might try to define the following method, which injects *every* type into the `Ordered` class:

```
def magic[a](x: a)(implicit a2ordered: a => Ordered[a]): Ordered[a] =
  a2ordered(x)
```

Now, if one tried to apply `sort` to an argument `arg` of a type that did not have another injection into the `Ordered` class, one would obtain an infinite expansion:

```
sort(arg)(x => magic(x)(x => magic(x)(x => ... )))
```

To prevent such infinite expansions, we require that every implicit method defini-

tion is contractive.

A method definition is *contractive* if the type of every implicit parameter type is properly contained in the type that is obtained by removing all implicit parameters from the method type and converting the rest to a function type.

A type T is *contained* in a type U if one of the following holds:

- T is the same as some part of U ,
- U is a function type and T is not.
- U and T are both function types, and the arity of U is greater than the arity of T .
- U and T both parameterized types (including function types) with the same type constructor, and each type argument of T is contained in the corresponding type argument of U .

A type T is *properly contained* in a type U if T is contained in U and different from U .

Example 7.2.2 The type of `list2ordered` is

```
(List[a])(implicit a => Ordered[a]): Ordered[List[a]] .
```

This type is contractive, because the type of the implicit parameter, `a => Ordered[a]`, is properly contained in the function type of the method without implicit parameters, `List[a] => Ordered[List[a]]`.

The type of `magic` is

```
(a)(implicit a => Ordered[a]): Ordered[a] .
```

This type is not contractive, because the type of the implicit parameter, `a => Ordered[a]`, is the same as the function type of the method without implicit parameters.

7.3 Views

Implicit parameters and methods can also define implicit conversions called views. A *view* from type S to type T is defined by an implicit value which has function type $S \Rightarrow T$ or $(\Rightarrow S) \Rightarrow T$ or by a method convertible to a value of that type.

Views are applied in two situations.

1. If an expression e is of type T , and T does not conform to the expression's expected type pt . In this case an implicit v is searched which is applicable to e and whose result type conforms to pt . The search proceeds as in the case of

implicit parameters, where the implicit scope is the one of $T \Rightarrow pt$. If such a view is found, the expression e is converted to $v(e)$.

2. In a selection $e.m$ with e of type T , if the selector m does not denote a member of T . In this case, a view v is searched which is applicable to e and whose result contains a member named m . The search proceeds as in the case of implicit parameters, where the implicit scope is the one of T . If such a view is found, the selection $e.m$ is converted to $v(e).m$.

As for implicit parameters, overloading resolution is applied if there are several possible candidates.

Example 7.3.1 Class `scala.Ordered[a]` contains a method

```
def <= [b >: a](that: b)(implicit b2ordered: b => Ordered[b]): boolean .
```

Assume two lists `xs` and `ys` of type `List[int]` and assume that the `list2ordered` and `int2ordered` methods defined in §7.2 are in scope. Then the operation

```
xs <= ys
```

is legal, and is expanded to:

```
list2ordered(xs)(int2ordered).<=
  (ys)
  (xs => list2ordered(xs)(int2ordered))
```

The first application of `list2ordered` converts the list `xs` to an instance of class `Ordered`, whereas the second occurrence is part of an implicit parameter passed to the `<=` method.

7.4 View Bounds

Syntax:

```
TypeParam      ::= id [>: Type] [<: Type] [<% Type]
```

A type parameter a may have a view bound $a <% T$. In this case the type parameter may be instantiated to any type S which is convertible by application of a view to the bound T .

A method or class containing such a type parameter is treated as being equivalent to a method with a view parameter. E.g.

```
def f[a <% T](ps): R = ...
```

is expanded to

```
def f[a](ps)(implicit v: a => T): R = ...
```

where v is a fresh name for the implicit parameter.

Example 7.4.1 The `<=` method mentioned in Example 7.3.1 can be declared more concisely as follows:

```
def <= [b >: a <% Ordered[b]](that: b): boolean
```

Chapter 8

Pattern Matching

8.1 Patterns

Syntax:

```
Pattern      ::= Pattern1 { '|' Pattern1 }
Pattern1     ::= varid ':' Type1
              | '_' ':' Type1
              | Pattern2
Pattern2     ::= varid ['@' Pattern3]
              | Pattern3
Pattern3     ::= SimplePattern
              | SimplePattern {id SimplePattern}
SimplePattern ::= '_'
              | varid
              | Literal
              | null
              | StableId [ '(' [Patterns] ')' ]
              | '(' [Pattern] ')'
              | XmlPattern
Patterns     ::= Pattern [',' Patterns]
              | '_' '*'
```

A pattern is built from constants, constructors, variables and type tests. Pattern matching tests whether a given value (or sequence of values) has the shape defined by a pattern, and, if it does, binds the variables in the pattern to the corresponding components of the value (or sequence of values). The same variable name may not be bound more than once in a pattern.

Pattern matching is always done in a context which supplies an expected type of the pattern. We distinguish the following kinds of patterns.

A *variable pattern* x is a simple identifier which starts with a lower case letter. It matches any value, and binds the variable name to that value. The type of x is the expected type of the pattern as given from outside. A special case is the wild-card pattern `_` which is treated as if it was a fresh variable.

A *typed pattern* $x : T$ consists of a pattern variable x and a type T . The type T may be a class type or a compound type; it may not contain a refinement (§3.2.5). This pattern matches any value which is an instance of the erasure (§3.6) of T and binds the variable name to that value. T must conform to the pattern's expected type. The type of x is T . A special case is if type T is a singleton type, say $p.type$. In this case the pattern matches only the single value designated by p (i.e. the pattern match amounts to a test whether the selector value is equal to p using method `eq` in class `AnyRef`).

A *pattern literal* l matches any value that is equal (in terms of `==`) to it. Its type must conform to the expected type of the pattern.

A *named pattern constant* r is a stable identifier (§3.1). To resolve the syntactic overlap with a variable pattern, a named pattern constant may not be a simple name starting with a lower-case letter. The type of r must conform to the expected type of the pattern. The pattern matches any value v such that $r == v$ (§12.1).

A *constructor pattern* $c(p_1, \dots, p_n)$ where $n \geq 0$ consists of an identifier c , followed by element patterns p_1, \dots, p_n . The constructor c is a simple or qualified name which denotes a case class (§5.3.2). If the case class is monomorphic, then it must conform to the expected type of the pattern, and the formal parameter types of c 's primary constructor (§5.3) are taken as the expected types of the element patterns p_1, \dots, p_n . If the case class is polymorphic, then its type parameters are instantiated so that the instantiation of c conforms to the expected type of the pattern. The instantiated formal parameter types of c 's primary constructor are then taken as the expected types of the component patterns p_1, \dots, p_n . The pattern matches all objects created from constructor invocations $c(v_1) \dots (v_n)$ where each element pattern p_i matches the corresponding value v_i .

A *sequence pattern* $s(p_1, \dots, p_n)$ where $n \geq 0$ consists of an identifier s , followed by element patterns p_1, \dots, p_n . The last element pattern may also be a *sequence wildcard* `_*`. The constructor s is a simple or qualified name which denotes some subclass of class `scala.Seq`. So as to disambiguate sequence patterns from constructor patterns, s is not permitted to denote a case class. The expected type of the sequence pattern must itself conform to some type instance `scala.Seq[T]` of the sequence class. Each element pattern p_i is type-checked with T as expected type, unless it is a sequence wildcard. If a final sequence wildcard is present, the pattern matches all values v that are sequences of class s which start with elements matching patterns p_1, \dots, p_{n-1} . If no final sequence wildcard is given, the pattern matches all values v that are sequences of class s and of length n which consist of elements matching patterns p_1, \dots, p_n .

An *infix operation pattern* `p id p'` is a shorthand for the constructor pattern

`id_class(p, p')`. The precedence and associativity of operators in patterns is the same as in expressions (§6.9).

A *pattern alternative* $p_1 \mid \dots \mid p_n$ consists of a number of alternative patterns p_i . All alternative patterns are type checked with the expected type of the pattern. They may not bind variables other than wildcards. The alternative pattern matches a value v if at least one of its alternatives matches v .

XML patterns are treated in §10.2.

Example 8.1.1 Some examples of patterns are:

1. The pattern `ex: IOException` matches all instances of class `IOException`, binding variable `ex` to the instance.
2. The pattern `(x, _)` matches pairs of values, binding `x` to the first component of the pair. The second component is matched with a wildcard pattern.
3. The pattern `x :: y :: xs` matches lists of length ≥ 2 , binding `x` to the list's first element, `y` to the list's second element, and `xs` to the remainder.
4. The pattern `1 | 2 | 3` matches the integers between 1 and 3.

A pattern p is *irrefutable* for a type T , if one of the following applies:

1. p is a variable pattern,
2. p is a typed pattern $x : T'$, and $T <: T'$,
3. p is a constructor pattern $c(p_1, \dots, p_n)$, the type T is an instance of class c , the primary constructor (§5.3) of type T has argument types T_1, \dots, T_n , and each p_i is irrefutable for T_i .

8.1.1 Regular Pattern Matching

Regular expression patterns are currently unimplemented in Scala version 2.0.

Scala provides a much simplified version of regular expression patterns that covers most scenarios of non-text sequence processing. A *sequence pattern* is a pattern that stands in a position where either (1) a pattern of a type T which is conforming to `Seq[A]` for some A is expected, or (2) a case class constructor that has an iterated formal parameter A^* . A wildcard star pattern `_*` in the rightmost position stands for arbitrary long sequences. It can be bound to variables, as usual, in which case the variable will have the type `Seq[A]`.

8.2 Pattern Matching Expressions

Syntax:

```

Expr          ::= PostfixExpr match '{' CaseClauses '}'
CaseClauses   ::= CaseClause {CaseClause}
CaseClause    ::= case Pattern [if PostfixExpr] '=>' Block

```

A pattern matching expression

```
e match { case  $p_1 \Rightarrow b_1$  ... case  $p_n \Rightarrow b_n$  }
```

consists of a selector expression e and a number $n > 0$ of cases. Each case consists of a (possibly guarded) pattern p_i and a block b_i . Each p_i might be complemented by a guard **if** e where e is a boolean expression. The scope of the pattern variables in p_i comprises the pattern's guard and the corresponding block b_i .

Let T be the type of the selector expression e and let a_1, \dots, a_m be the type parameters of all methods enclosing the pattern matching expression. For every a_i , let L_i be its lower bound and U_i be its higher bound. Every pattern $p \in \{p_1, \dots, p_n\}$ can be typed in two ways. First, it is attempted to type p with T as its expected type. If this fails, p is instead typed with a modified expected type T' which results from T by replacing every occurrence of a type parameter a_i by *undefined*. If this second step fails also, a compile-time error results. If the second step succeeds, let T_p be the type of pattern p seen as an expression. One then determines minimal bounds L'_1, \dots, L'_m and maximal bounds U'_1, \dots, U'_m such that for all i , $L_i <: L'_i$ and $U'_i <: U_i$ and the following constraint system is satisfied:

$$L_1 <: a_1 <: U_1 \wedge \dots \wedge L_m <: a_m <: U_m \Rightarrow T_p <: T$$

If no such bounds can be found, a compile time error results. If such bounds are found, the pattern matching clause starting with p is then typed under the assumption that each a_i has lower bound L'_i instead of L_i and has upper bound U'_i instead of U_i .

The expected type of every block b_i is the expected type of the whole pattern matching expression. The type of the pattern matching expression is then the least upper bound of the types of all blocks b_i .

When applying a pattern matching expression to a selector value, patterns are tried in sequence until one is found which matches the selector value (§8.1). Say this case is **case** $p_i \Rightarrow b_i$. The result of the whole expression is then the result of evaluating b_i , where all pattern variables of p_i are bound to the corresponding parts of the selector value. If no matching pattern is found, a `scala.MatchError` exception is thrown.

The pattern in a case may also be followed by a guard suffix **if** e with a boolean expression e . The guard expression is evaluated if the preceding pattern in the case matches. If the guard expression evaluates to **true**, the pattern match succeeds as normal. If the guard expression evaluates to **false**, the pattern in the case is considered not to match and the search for a matching pattern continues.

In the interest of efficiency the evaluation of a pattern matching expression may try patterns in some other order than textual sequence. This might affect evaluation

through side effects in guards. However, it is guaranteed that a guard expression is evaluated only if the pattern it guards matches.

Example 8.2.1 Consider the following definitions of arithmetic terms:

```
abstract class Term[T]
case class Lit(x: int) extends Term[int]
case class Succ(t: Term[int]) extends Term[int]
case class IsZero(t: Term[int]) extends Term[boolean]
case class If[T](c: Term[boolean],
                t1: Term[T],
                t2: Term[T]) extends Term[T]
```

There are terms to represent numeric literals, incrementation, a zero test, and a conditional. Every term carries as a type parameter the type of the expression it represents (either int or boolean).

A type-safe evaluator for such terms can be written as follows.

```
def eval[T](t: Term[T]): T = t match {
  case Lit(n)           => n
  case Succ(u)          => eval(u) + 1
  case IsZero(u)        => eval(u) == 0
  case If(c, u1, u2) => eval(if (eval(c)) u1 else u2)
}
```

Note that the evaluator makes crucial use of the fact that type parameters of enclosing methods can acquire new bounds through pattern matching.

For instance, the type of the pattern in the second case, `Succ(u)`, is `int`. It conforms to the selector type `T` only if we assume an upper and lower bound of `int` for `T`. Under the assumption `int <: T <: int` we can also verify that the type right hand side of the second case, `int` conforms to its expected type, `T`.

8.3 Pattern Matching Anonymous Functions

Syntax:

```
BlockExpr ::= '{' CaseClauses '}'
```

An anonymous function can be defined by a sequence of cases

```
{ case  $p_1 \Rightarrow b_1$  ... case  $p_n \Rightarrow b_n$  }
```

which appear as an expression without a prior **match**. The expected type of such an expression must in part be defined. It must be either `scala.Function1[T_p , T_r]` or

`scala.PartialFunction[T_p , T_r]`, where the argument type T_p must be fully determined, but the result type T_r may be undetermined.

If the expected type is `scala.Function1[T_p , T_r]`, the expression is taken to be equivalent to the anonymous function:

```
(x:  $T_p$ ) => x match { case  $p_1$  =>  $b_1$  ... case  $p_n$  =>  $b_n$  }
```

As was shown in §6.20 this anonymous function is in turn equivalent to the following instance creation expression:

```
new scala.Function1[ $T_p$ ,  $T$ ] {  
  def apply( $x$ :  $T_p$ ):  $T$  =  $x$  match {  
    case  $p_1$  =>  $b_1$  ... case  $p_n$  =>  $b_n$   
  }  
}
```

Here, x is a fresh name and T is the least upper bound of the types of all b_i .

If the expected type is `scala.PartialFunction[T_p , T_r]`, the expression is taken to be equivalent to the following instance creation expression:

```
new scala.PartialFunction[ $T_p$ ,  $T_r$ ] {  
  def apply( $x$ :  $T_p$ ):  $T$  =  $x$  match {  
    case  $p_1$  =>  $b_1$  ... case  $p_n$  =>  $b_n$   
  }  
  def isDefinedAt( $x$ :  $T_p$ ):  $T$  = {  
    case  $p_1$  => true ... case  $p_n$  => true  
    case _ => false  
  }  
}
```

Here, x is a fresh name and T is the least upper bound of the types of all b_i . The final default case in the `isDefinedAt` method is omitted if one of the patterns p_1, \dots, p_n is already a variable or wildcard pattern.

Chapter 9

Top-Level Definitions

9.1 Compilation Units

Syntax:

```
CompilationUnit ::= [package QualId StatementSeparator] TopStatSeq
TopStatSeq     ::= TopStat {StatementSeparator TopStat}
TopStat        ::= {AttributeClause} {Modifier} TmplDef
                | Import
                | Packaging
                |
QualId          ::= id {'.' id}
```

A compilation unit consists of a sequence of packagings, import clauses, and class and object definitions, which may be preceded by a package clause.

A compilation unit **package** *p*; *stats* starting with a package clause is equivalent to a compilation unit consisting of a single packaging **package** *p* { *stats* }.

Implicitly imported into every compilation unit are, in that order : the package `java.lang`, the package `scala`, and the object `scala.Predef` (§12.5). Members of a later import in that order hide members of an earlier import.

9.2 Packagings

Syntax:

```
Packaging      ::= package QualId '{' TopStatSeq '}'
```

A package is a special object which defines a set of member classes, objects and packages. Unlike other objects, packages are not introduced by a definition. In-

stead, the set of members of a package is determined by packagings.

A packaging **package** *p* { *ds* } injects all definitions in *ds* as members into the package whose qualified name is *p*. Members of a package are called *top-level* definitions. If a definition in *ds* is labeled **private**, it is visible only for other members in the package.

Selections *p.m* from *p* as well as imports from *p* work as for objects. However, unlike other objects, packages may not be used as values. It is illegal to have a package with the same fully qualified name as a module or a class.

Top-level definitions outside a packaging are assumed to be injected into a special empty package. That package cannot be named and therefore cannot be imported. However, members of the empty package are visible to each other without qualification.

9.3 Programs

A *program* is a top-level object that has a member method `main` of type `(Array[String])unit`. Programs can be executed from a command shell. The program's command arguments are passed to the `main` method as a parameter of type `Array[String]`.

The `main` method of a program can be directly defined in the object, or it can be inherited. The `scala` library defines a class `scala.Application` that defines an empty inherited `main` method. An object *m* inheriting from this class is thus a program, which executes the initialization code of the object *m*.

Example 9.3.1 The following example will create a hello world program by defining a method `main` in module `test.HelloWorld`.

```
package test

object HelloWorld {
  def main(args: Array[String]) = System.out.println("hello world")
}
```

This program can be started by either of the commands

```
scala test.HelloWorld
```

In a Java environment, the command

```
java test.HelloWorld
```

would work as well.

`HelloWorld` can also be defined without a `main` method by inheriting from

Application instead:

```
package test
object HelloWorld extends Application {
    System.out.println("hello world")
}
```


Chapter 10

XML expressions and patterns

By Burak Emir

This chapter describes the syntactic structure of XML expressions and patterns. It follows as close as possible the XML 1.0 specification [W3C], changes being mandated by the possibility of embedding Scala code fragments.

10.1 XML expressions

XML expressions are expressions generated by the following production, where the opening bracket ‘<’ of the first element must be in a position to start the lexical XML mode (§1.5).

Syntax:

```
XmlExpr ::= XmlContent {Element}
```

Well-formedness constraints of the XML specification apply, which means for instance that start tags and end tags must match, and attributes may only be defined once, with the exception of constraints related to entity resolution.

The following productions describe Scala’s extensible markup language, designed as close as possible to the W3C extensible markup language standard. Only the productions for attribute values and character data are changed. Scala does not support neither declarations, CDATA sections nor processing instructions. Entity references are not resolved at runtime.

Syntax:

```
Element      ::=      EmptyElemTag  
                |      STag Content ETag
```

```

EmptyElemTag ::= '<' Name {S Attribute} [S] '>'

STag          ::= '<' Name {S Attribute} [S] '>'
ETag          ::= '</' Name [S] '>'
Content       ::= [CharData] {Content1 [CharData]}
Content1      ::= XmlContent
                | Reference
                | ScalaExpr
XmlContent    ::= Element
                | CDsect
                | PI
                | Comment

```

If an XML expression is a single element, its value is a runtime representation of an XML node (an instance of a subclass of `scala.xml.Node`). If the XML expression consists of more than one element, then its value is a runtime representation of a sequence of XML nodes (an instance of a subclass of `scala.Seq[scala.xml.Node]`).

If an XML expression is an entity reference, CDATA section, processing instructions or a comments, it is represented by an instance of the corresponding Scala runtime class.

By default, beginning and trailing whitespace in element content is removed, and consecutive occurrences of whitespace are replaced by a single space character `\u0020`. This behavior can be changed to preserve all whitespace with a compiler option. **Syntax:**

```

Attribute      ::= Name Eq AttValue

AttValue       ::= '"' {CharQ | CharRef} '"'
                | "'" {CharA | CharRef} "'"
                | ScalaExp

ScalaExpr      ::= '{' expr '}'

CharData       ::= { CharNoRef } without {CharNoRef}{'CharB {CharNoRef}
                                     and without {CharNoRef}']>'{CharNoRef}

```

XML expressions may contain Scala expressions as attribute values or within nodes. In the latter case, these are embedded using a single opening brace `{` and ended by a closing brace `}`. To express a single opening braces within XML text as generated by `CharData`, it must be doubled. Thus, `{{` represents the XML text `"` and does not introduce an embedded Scala expression.

Syntax:

BaseChar, Char, Comment, CombiningChar, Ideographic, NameChar, S, Reference

```

::= "as in W3C XML"

Char1      ::= Char without '<' | '&'
CharQ      ::= Char1 without '"'
CharA      ::= Char1 without "'"
CharB      ::= Char1 without '{'

Name        ::= XNameStart {NameChar}

XNameStart  ::= '_' | BaseChar | Ideographic
              (as in W3C XML, but without ':')

```

10.2 XML patterns

XML patterns are patterns generated by the following production, where the opening bracket '`<`' of the element patterns must be in a position to start the lexical XML mode (§1.5).

Syntax:

```
XmlPattern ::= ElementPattern {ElementPattern}
```

Well-formedness constraints of the XML specification apply.

If an XML pattern is a single element pattern, it expects the type of runtime representation of an XML tree, and matches exactly one instance of this type that has the same structure as described by the pattern. If an XML pattern consists of more than one element, then it expects the type of sequences of runtime representations of XML trees, and matches every sequence whose elements match the sequence described by the pattern.

XML patterns may contain Scala patterns (§8.2).

Whitespace is treated the same way as in XML expressions. Patterns that are entity references, CDATA sections, processing instructions and comments match runtime representations which are the the same.

By default, beginning and trailing whitespace in element content is removed, and consecutive occurrences of whitespace are replaced by a single space character `\u0020`. This behavior can be changed to preserve all whitespace with a compiler option.

Syntax:

```

ElemPattern  ::=      EmptyElemTagP
                   |      STagP ContentP ETagP

EmptyElemTagP ::=      '<' Name [S] '/>'

```

```
STagP      ::= '<' Name [S] '>'
ETagP      ::= '</' Name [S] '>'
ContentP    ::= [CharData] {(ElemPattern|ScalaPatterns) [CharData]}
ContentP1   ::= ElemPattern
              | Reference
              | CD Sect
              | PI
              | Comment
              | ScalaPatterns
ScalaPatterns ::= '{' patterns '}'
```

Chapter 11

Attributes

Syntax:

```
AttributeClause ::= '[' Attribute {',' Attribute} ']' [NewLine]
Attribute       ::= Constr
```

Attributes associate meta-information with definitions. A simple attribute clause has the form $[C]$ or $[C(a_1, \dots, a_n)]$. Here, c is a constructor of a class C , which must conform to the class `scala.Attribute`. All given constructor arguments a_1, \dots, a_n must be constant expressions. An attribute clause applies to the first definition or declaration following it. More than one attribute clause may precede a definition and declaration. The order in which these clauses are given does not matter. It is also possible to combine several attributes separated by commas in one clause. Such a combined clause $[A_1, \dots, A_n]$ is equivalent to a set of clauses $[A_1] \dots [A_n]$.

The meaning of attribute clauses is implementation-dependent. On the Java platform, the following attributes have a standard meaning.

`transient`

Marks a field to be non-persistent; this is equivalent to the `transient` modifier in Java.

`volatile`

Marks a field which can change its value outside the control of the program; this is equivalent to the `volatile` modifier in Java.

`serializable`

Marks a class to be serializable; this is equivalent to inheriting from the `java.io.Serializable` interface in Java.

SerialVersionUID(<longlit>)

Attaches a serial version identifier (a long constant) to a class. This is equivalent to a the following field definition in Java:

```
private final static SerialVersionUID = <longlit>
```

beanProperty

When prefixed to a definition of some variable X, this attribute causes getter and setter methods getX, setX in the Java bean style to be added in the class containing the variable. The first letter of the variable appears capitalized after the get or set. When the attribute is added to the definition of an immutable value definition X, only a setter is generated.

Chapter 12

The Scala Standard Library

The Scala standard library consists of the package `scala` with a number of classes and modules. Some of these classes are described in the following.

12.1 Root Classes

The root of the Scala class hierarchy is formed by class `Any`. Every class in a Scala execution environment inherits directly or indirectly from this class. Class `Any` has two direct subclasses: `AnyRef` and `AnyVal`.

The subclass `AnyRef` represents all values which are represented as objects in the underlying host system. Every user-defined Scala class inherits directly or indirectly from this class. Furthermore, every user-defined Scala class also inherits the trait `scala.ScalaObject`. Classes written in other languages still inherit from `scala.AnyRef`, but not from `scala.ScalaObject`.

The class `AnyVal` has a fixed number subclasses, which describe values which are not implemented as objects in the underlying host system.

Classes `AnyRef` and `AnyVal` are required to provide only the members declared in class `Any`, but implementations may add host-specific methods to these classes (for instance, an implementation may identify class `AnyRef` with its own root class for objects).

The signatures of these root classes are described by the following definitions.

```
package scala
/** The universal root class */
abstract class Any {

    /** Defined equality; abstract here */
    def equals(that: Any): boolean
```

```

/** Semantic equality between values of same type */
final def == (that: Any): boolean = this equals that

/** Semantic inequality between values of same type */
final def != (that: Any): boolean = !(this == that)

/** Hash code; abstract here */
def hashCode(): Int = ...

/** Textual representation; abstract here */
def toString(): String = ...

/** Type test; needs to be inlined to work as given */
def isInstanceOf[a]: Boolean = this match {
  case x: a => true
  case _ => false
}

/** Type cast; needs to be inlined to work as given */ */
def asInstanceOf[a]: a = this match {
  case x: a => x
  case _ => if (this eq null) this
    else throw new ClassCastException()
}

/** The root class of all value types */
final class AnyVal extends Any

/** The root class of all reference types */
class AnyRef extends Any {
  def equals(that: Any): Boolean = this eq that
  final def eq(that: Any): Boolean = ... // reference equality

  def hashCode(): Int = ... // hashCode computed from allocation address
  def toString(): String = ... // toString computed from hashCode and class name
}

/** A mixin class for every user-defined Scala class */
trait ScalaObject extends AnyRef

```

The test `x.asInstanceOf[T]` is treated specially if `T` is a numeric value type (§12.2. In this case the cast will be translated to an application of a conversion method `x.toT` (§12.2.1). For non-numeric values `x` the operation will raise a `ClassCastException`.

12.2 Value Classes

Value classes are classes whose instances are not represented as objects by the underlying host system. All value classes inherit from class `AnyVal`. Scala implementations need to provide the value classes `Unit`, `Boolean`, `Double`, `Float`, `Long`, `Int`, `Char`, `Short`, and `Byte` (but are free to provide others as well). The signatures of these classes are defined in the following.

12.2.1 Numeric Value Types

Classes `Double`, `Float`, `Long`, `Int`, `Char`, `Short`, and `Byte` are together called *numeric value types*. Classes `Byte`, `Short`, or `Char` are called *subrange types*. Subrange types, as well as `Int` and `Long` are called *integer types*, whereas `Float` and `Double` are called *floating point types*.

Numeric value types are ranked in the following partial order:

```

Byte - Short
      \
      Int - Long - Float - Double
      /
Char

```

`Byte` and `Short` are the lowest-ranked types in this order, whereas `Double` is the highest-ranked. Ranking does *not* imply a conformance (§3.5.2) relationship; for instance `Int` is not a subtype of `Long`. However, object `Predef` (§12.5) defines views (§7.3) from every numeric value type to all higher-ranked numeric value types. Therefore, lower-ranked types are implicitly converted to higher-ranked types when required by the context (§6.22).

Given two numeric value types S and T , the *operation type* of S and T is defined as follows: If both S and T are subrange types then the operation type of S and T is `Int`. Otherwise the operation type of S and T is the larger of the two types wrt ranking. Given two numeric values v and w the operation type of v and w is the operation type of their run-time types.

Any numeric value type T supports the following methods.

- Comparison methods for equals (`==`), not-equals (`!=`), less-than (`<`), greater-than (`>`), less-than-or-equals (`<=`), greater-than-or-equals (`>=`), which each exist in 7 overloaded alternatives. Each alternative takes a parameter of some numeric value type. Its result type is type `Boolean`. The operation is evaluated by converting the receiver and its argument to their operation type and performing the given comparison operation of that type.
- Arithmetic methods addition (`+`), subtraction (`-`), multiplication (`*`), division (`/`), and remainder (`%`), which each exist in 7 overloaded alternatives. Each alternative takes a parameter of some numeric value type U . Its result type is

the operation type of T and U . The operation is evaluated by converting the receiver and its argument to their operation type and performing the given arithmetic operation of that type.

- Parameterless arithmetic methods identity (+) and negation (-), with result type T . The first of these returns the receiver unchanged, whereas the second returns its negation.
- Conversion methods `toByte`, `toShort`, `toChar`, `toInt`, `toLong`, `toFloat`, `toDouble` which convert the receiver object to the target type, using the rules of Java's numeric type cast operation. The conversion might truncate the numeric value (as when going from `Long` to `Int` or from `Int` to `Byte`) or it might lose precision (as when going from `Double` to `Float` or when converting between `Long` and `Float`).

Integer numeric value types support in addition the following operations:

- Bit manipulation methods bitwise-and (&), bitwise-or |, and bitwise-exclusive-or (^), which each exist in 5 overloaded alternatives. Each alternative takes a parameter of some integer numeric value type. Its result type is the operation type of T and U . The operation is evaluated by converting the receiver and its argument to their operation type and performing the given bitwise operation of that type.
- A parameterless bit-negation method (~). Its result type is the receiver type T or `Int`, whichever is larger. The operation is evaluated by converting the receiver to the result type and negating every bit in its value.
- Bit-shift methods left-shift (<<), arithmetic right-shift (>>), and unsigned right-shift (>>>). Each of these methods has two overloaded alternatives, which take a parameter n of type `Int`, respectively `Long`. The result type of the operation is the receiver type T , or `Int`, whichever is larger. The operation is evaluated by converting the receiver to the result type and performing the specified shift by n bits.

Numeric value types also implement operations `equals`, `hashCode`, and `toString` from class `Any`.

The `equals` method tests whether the argument is a numeric value type. If this is true, it will perform the `==` operation which is appropriate for that type. That is, the `equals` method of a numeric value type can be thought of being defined as follows:

```
def equals(other: Any): Boolean = other match {
  case that: Byte   => this == that
  case that: Short  => this == that
  case that: Char    => this == that
  case that: Int     => this == that
  case that: Long    => this == that
```

```

    case that: Float => this == that
    case that: Double => this == that
    case _ => false
  }

```

The hashCode method returns an integer hashCode that maps equal numeric values to equal results. It is guaranteed to be the identity for type Int and for all subrange types.

The toString method displays its receiver as an integer or floating point number.

Example 12.2.1 As an example, here is the signature of the numeric value type Int:

```

package scala
abstract sealed class Int extends AnyVal {
  def == (that: Double): Boolean // double equality
  def == (that: Float): Boolean // float equality
  def == (that: Long): Boolean // long equality
  def == (that: Int): Boolean // int equality
  def == (that: Short): Boolean // int equality
  def == (that: Byte): Boolean // int equality
  def == (that: Char): Boolean // int equality
  /* analogous for !=, <, >, <=, >= */

  def + (that: Double): Double // double addition
  def + (that: Float): Double // float addition
  def + (that: Long): Long // long addition
  def + (that: Int): Int // int addition
  def + (that: Short): Int // int addition
  def + (that: Byte): Int // int addition
  def + (that: Char): Int // int addition
  /* analogous for -, *, /, % */

  def & (that: Long): Long // long bitwise and
  def & (that: Int): Int // int bitwise and
  def & (that: Short): Int // int bitwise and
  def & (that: Byte): Int // int bitwise and
  def & (that: Char): Int // int bitwise and
  /* analogous for |, ^ */

  def << (cnt: Int): Int // int left shift
  def << (cnt: Long): Int // long left shift
  /* analogous for >>, >>> */

  def + : Int // int identity
  def - : Int // int negation
  def ~ : Int // int bitwise negation

```

```
def toByte: Byte           // convert to Byte
def toShort: Short         // convert to Short
def toChar: Char           // convert to Char
def toInt: Int             // convert to Int
def toLong: Long           // convert to Long
def toFloat: Float         // convert to Float
def toDouble: Double       // convert to Double
}
```

12.2.2 Class Boolean

Class `Boolean` has only two values: **true** and **false**. It implements operations as given in the following signature:

```
package scala
abstract sealed class Boolean extends AnyVal {
  def && (p: => Boolean): Boolean // boolean and
  def || (p: => Boolean): Boolean // boolean or
  def & (x: Boolean): Boolean    // boolean strict and
  def | (x: Boolean): Boolean    // boolean strict or

  def == (x: Boolean): Boolean   // boolean equality
  def != (x: Boolean): Boolean   // boolean inequality

  def !: Boolean                // boolean negation
}
```

The class also implements operations `equals`, `hashCode`, and `toString` from class `Any`.

The `equals` method returns **true** if the argument is the same boolean value as the receiver, **false** otherwise. The `hashCode` method returns 1 when invoked on **true**, and 0 when invoked on **false**. The `toString` method returns the receiver converted to a string, i.e. either **"true"** or **"false"**.

12.2.3 Class Unit

Class `Unit` has only one value: `()`. It implements only the three methods `equals`, `hashCode`, and `toString` from class `Any`.

The `equals` method returns **true** if the argument is the unit value `()`, **false** otherwise. The `hashCode` method returns a fixed, implementation-specific hash-code. The `toString` method returns **"()"**.

12.3 Standard Reference Classes

This section presents some standard Scala reference classes which are treated in a special way in Scala compiler – either Scala provides syntactic sugar for them, or the Scala compiler generates special code for their operations. Other classes in the standard Scala library are documented in the Scala library documentation by HTML pages.

12.3.1 Class String

Scala's `String` class is usually derived from the standard `String` class of the underlying host system (and may be identified with it). For Scala clients the class is taken to support in each case a method

```
def + (that: Any): String
```

which concatenates its left operand with the textual representation of its right operand.

12.3.2 The Tuple classes

Scala defines tuple classes `Tuple n` for $n = 2, \dots, 9$. These are defined as follows.

```
package scala
case class Tuple $n$ [+a1, ..., +a $n$ ](_1: a1, ..., _ $n$ : a $n$ ) {
  def toString = "(" ++ _1 ++ "," ++ ... ++ "," ++ _ $n$  ++ ")"
}
```

The implicitly imported `Predef` object (§12.5) defines the names `Pair` as an alias of `Tuple2` and `Triple` as an alias for `Tuple3`.

12.3.3 The Function Classes

Scala defines function classes `Function n` for $n = 1, \dots, 9$. These are defined as follows.

```
package scala
trait Function $n$ [-a1, ..., -a $n$ , +b] {
  def apply(x1: a1, ..., x $n$ : a $n$ ): b
  def toString = "<function>"
}
```

A subclass of `Function1` represents partial functions, which are undefined on some points in their domain. In addition to the `apply` method of functions, partial functions also have a `isDefined` method, which tells whether the function is defined at the given argument:

```

class PartialFunction[-a,+b] extends Function1[a, b] {
  def isDefinedAt(x: a): Boolean
}

```

The implicitly imported Predef object (§12.5) defines the name `Function` as an alias of `Function1`.

12.3.4 Class Array

The class of generic arrays is given as follows.

```

final class Array[A](len: Int) extends Seq[A] {
  def length: Int = len
  def apply(i: Int): A = ...
  def update(i: Int, x: A): Unit = ...
  def elements: Iterator[A] = ...
  def subArray(from: Int, end: Int): Array[a] = ...
  def filter(p: a => Boolean): Array[a] = ...
  def map[b](f: a => b): Array[b] = ...
  def flatMap[b](f: a => Array[b]): Array[b] = ...
}

```

If T is not a type parameter or abstract type, the type `Array[T]` is represented as the native array type `[]T` in the underlying host system. In that case `length` returns the length of the array, `apply` means subscribing, and `update` means element update. Because of the syntactic sugar for `apply` and `update` operations (§6.22, we have the following correspondences between Scala and Java/C# code for operations on an array `xs`:

<i>Scala</i>	<i>Java/C#</i>
<code>xs.length</code>	<code>xs.length</code>
<code>xs(i)</code>	<code>xs[i]</code>
<code>xs(i) = e</code>	<code>xs[i] = e</code>

Arrays also implement the sequence trait `scala.Seq` by defining an `elements` method which returns all elements of the array in an `Iterator`.

Because of the tension between parametrized types in Scala and the ad-hoc implementation of arrays in the host-languages, some subtle points need to be taken into account when dealing with arrays. These are explained in the following.

First, unlike arrays in Java or C#, arrays in Scala are *not* co-variant; That is, $S <: T$ does not imply `Array[S] <: Array[T]` in Scala. However, it is possible to cast an array of S to an array of T if such a cast is permitted in the host environment.

For instance `Array[String]` does not conform to `Array[Object]`, even though `String` conforms to `Object`. However, it is possible to cast an expression of type `Array[String]` to `Array[Object]`, and this cast will succeed without raising a

ClassCastException. Example:

```
val xs = new Array[String](2)
// val ys: Array[Object] = xs    // **** error: incompatible types
val ys: Array[Object] = xs.asInstanceOf[Array[Object]] // OK
```

Second, for *polymorphic arrays*, that have a type parameter or abstract type T as their element type, a representation different from `[]T` might be used. However, it is guaranteed that `isInstanceOf` and `asInstanceOf` still work as if the array used the standard representation of monomorphic arrays:

```
val ss = new Array[String](2)

def f[T](xs: Array[T]): Array[String] =
  if (xs.isInstanceOf[Array[String]]) xs.asInstanceOf[Array[String]]
  else throw new Error("not an instance")

f(ss)                                // returns ss
```

The representation chosen for polymorphic arrays also guarantees that polymorphic array creations work as expected. An example is the following implementation of method `mkArray`, which creates an array of an arbitrary type T , given a sequence of T 's which defines its elements.

```
def mkArray[T](elems: Seq[T]): Array[T] = {
  val result = new Array[T](elems.length)
  var i = 0
  for (val elem <- elems) {
    result(i) = elem
    i = i + 1
  }
}
```

Note that under Java's erasure model of arrays the method above would not work as expected – in fact it would always return an array of `Object`.

Third, in a Java environment there is a method `System.arraycopy` which takes two objects as parameters together with start indices and a length argument, and copies elements from one object to the other, provided the objects are arrays of compatible element types. `System.arraycopy` will not work for Scala's polymorphic arrays because of their different representation. One should instead use method `Array.copy`, defined as follows:

```
package scala
object Array {
  def copy(src: AnyRef, srcPos: Int,
           dest: AnyRef, destPos: Int,
```

```
length: Int): Unit = ...
```

Example 12.3.1 The following method duplicates a given argument array and returns a pair consisting of the original and the duplicate:

```
def duplicate[T](xs: Array[T]) = {
  val ys = new Array[T](xs.length)
  Array.copy(xs, 0, ys, 0, xs.length)
  Pair(xs, ys)
}
```

12.4 Class Node

```
package scala.xml
```

```
trait Node {
```

```
  /** the label of this node */
```

```
  def label: String
```

```
  /** attribute axis */
```

```
  def attribute: Map[String, String]
```

```
  /** child axis (all children of this node) */
```

```
  def child: Seq[Node]
```

```
  /** descendant axis (all descendants of this node) */
```

```
  def descendant: Seq[Node] = child.toList.flatMap {
    x => x::x.descendant.asInstanceOf[List[Node]]
  }
```

```
  /** descendant axis (all descendants of this node) */
```

```
  def descendant_or_self: Seq[Node] = this::child.toList.flatMap {
    x => x::x.descendant.asInstanceOf[List[Node]]
  }
```

```
  override def equals(x: Any): boolean = x match {
```

```
    case that:Node =>
```

```
      that.label == this.label &&
```

```
      that.attribute.sameElements(this.attribute) &&
```

```
      that.child.sameElements(this.child)
```

```
    case _ => false
```

```
}
```



```

/** XPath style projection function. Returns all children of this node
 * that are labeled with 'that'. The document order is preserved.
 */
def \ (that: Symbol): NodeSeq = {
  new NodeSeq({
    that.name match {
      case "_" => child.toList
      case _ =>
        var res:List[Node] = Nil
        for (val x <- child.elements x.label == that.name) {
          res = x::res
        }
        res.reverse
    }
  })
}

/** XPath style projection function. Returns all nodes labeled with the
 * name 'that' from the descendant_or_self axis. Document order is preserved.
 */
def \\ (that: Symbol): NodeSeq = {
  new NodeSeq(
    that.name match {
      case "_" => this.descendant_or_self
      case _ => this.descendant_or_self.asInstanceOf[List[Node]].
        filter(x => x.label == that.name)
    })
}

/** hashCode for this XML node */
override def hashCode() =
  Utility.hashCode(label, attribute.toList.hashCode(), child)

/** string representation of this node */
override def toString() = Utility.toXML(this)
}

```

12.5 The Predef Object

The Predef object defines standard functions and type aliases for Scala programs. It is always implicitly imported, so that all its defined members are available without qualification. Its definition for the JVM environment conforms to the following signature:

```
package scala
object Predef {

  // Standard type aliases -----

  type byte = scala.Byte
  type short = scala.Short
  type char = scala.Char
  type int = scala.Int
  type long = scala.Long
  type float = scala.Float
  type double = scala.Double
  type boolean = scala.Boolean
  type unit = scala.Unit

  type String = java.lang.String
  type NullPointerException = java.lang.NullPointerException
  type Throwable = java.lang.Throwable

  type Pair[+p, +q] = Tuple2[p, q]
  type Triple[+a, +b, +c] = Tuple3[a, b, c]

  type Function[-a, +b] = Function1[a, b]

  // Factory methods -----

  def Pair[a, b](x: a, y: b) = Tuple2(x, y)
  def Triple[a, b, c](x: a, y: b, z: c) = Tuple3(x, y, z)

  def Tuple[a1, a2](x1: a1, x2: a2) = Tuple2(x1, x2)
  def Tuple[a1, a2, a3](x1: a1, x2: a2, x3: a3) = Tuple3(x1, x2, x3)

  // analogous for tuples of length 4-9:
  ...
}
```

```

def Array[A <: AnyRef](xs: A*): Array[A] = {
  val array = new Array[A](xs.length);
  var i = 0
  for (val x <- xs.elements) { array(i) = x; i = i + 1; }
  array
}

// analogous to above:
def Array(xs: boolean*): Array[boolean] = ...
def Array(xs: byte*) : Array[byte]      = ...
def Array(xs: short*) : Array[short]    = ...
def Array(xs: char*)  : Array[char]     = ...
def Array(xs: int*)   : Array[int]      = ...
def Array(xs: long*)  : Array[long]     = ...
def Array(xs: float*) : Array[float]    = ...
def Array(xs: double*) : Array[double]  = ...
def Array(xs: unit*)  : Array[unit]     = ...

// The “catch-all” view -----

implicit def identity[a](x: a): a = x

// Views into class Ordered

implicit def int2ordered(x: int): Ordered[int] = new Ordered[int] with Proxy {
  def self: Any = x
  def compare [b >: int <% Ordered[b]](y: b): int = y match {
    case y1: int =>
      if (x < y1) -1
      else if (x > y1) 1
      else 0
    case _ => -(y compare x)
  }
}

// The implementations of following methods are analogous to the last one:

implicit def char2ordered(x: char): Ordered[char] = ...
implicit def long2ordered(x: long): Ordered[long] = ...
implicit def float2ordered(x: float): Ordered[float] = ...
implicit def double2ordered(x: double): Ordered[double] = ...
implicit def boolean2ordered(x: boolean): Ordered[boolean] = ...

```

```

implicit def seq2ordered[A <% Ordered[A]](xs: Array[A]): Ordered[Seq[A]] =
  new Ordered[Seq[A]] with Proxy {
    def compare[B >: Seq[A] <% Ordered[B]](that: B): Int = that match {
      case that: Seq[A] =>
        var res = 0
        val these = this.elements
        val those = that.elements
        while (res == 0 && these.hasNext)
          res = if (!those.hasNext) 1 else these.next compare those.next
        case _ => - (that compare xs)
    }
  }

implicit def string2ordered(x: String): Ordered[String] =
  new Ordered[String] with Proxy {
    def self: Any = x
    def compare [b >: String <% Ordered[b]](y: b): int = y match {
      case y1: String => x compare y1
      case _ => -(y compare x)
    }
  }

implicit def tuple2ordered[a1 <% Ordered[a1], a2 <% Ordered[a2]]
  (x: Tuple2[a1, a2]): Ordered[Tuple2[a1, a2]] =
  new Ordered[Tuple2[a1, a2]] with Proxy {
    def self: Any = x
    def compare[T >: Tuple2[a1, a2] <% Ordered[T]](y: T): Int = y match {
      case y: Tuple2[a1, a2] =>
        val res = x._1 compare y._1
        if (res == 0) x._2 compare y._2
        else res
      case _ => -(y compare x)
    }
  }

// Analogous for Tuple3 to Tuple9

// Views into class Seq

implicit def string2seq(str: String): Seq[Char] = new Seq[Char] {
  def length = str.length()
  def elements = Iterator.fromString(str)
  def apply(n: Int) = str.charAt(n)
  override def hashCode(): Int = str.hashCode()
  override def equals(y: Any): Boolean = (str == y)
  override protected def stringPrefix: String = "String"
}

```

// Views from primitive types to Java's boxed types

```
implicit def byte2Byte(x: byte) = new java.lang.Byte(x)
implicit def short2Short(x: short) = new java.lang.Short(x)
implicit def char2Character(x: char) = new java.lang.Character(x)
implicit def int2Integer(x: int) = new java.lang.Integer(x)
implicit def long2Long(x: long) = new java.lang.Long(x)
implicit def float2Float(x: float) = new java.lang.Float(x)
implicit def double2Double(x: double) = new java.lang.Double(x)
implicit def boolean2Boolean(x: boolean) = new java.lang.Boolean(x)
```

// Numeric conversion views

```
implicit def byte2short(x: byte): short = x.toShort
implicit def byte2int(x: byte): int = x.toInt
implicit def byte2long(x: byte): long = x.toLong
implicit def byte2float(x: byte): float = x.toFloat
implicit def byte2double(x: byte): double = x.toDouble

implicit def short2int(x: short): int = x.toInt
implicit def short2long(x: short): long = x.toLong
implicit def short2float(x: short): float = x.toFloat
implicit def short2double(x: short): double = x.toDouble

implicit def char2int(x: char): int = x.toInt
implicit def char2long(x: char): long = x.toLong
implicit def char2float(x: char): float = x.toFloat
implicit def char2double(x: char): double = x.toDouble

implicit def int2long(x: int): long = x.toLong
implicit def int2float(x: int): float = x.toFloat
implicit def int2double(x: int): double = x.toDouble

implicit def long2float(x: long): float = x.toFloat
implicit def long2double(x: long): double = x.toDouble

implicit def float2double(x: float): double = x.toDouble
```

```
// Errors and asserts -----  
  
def error(message: String): Bottom = throw new Error(message)  
  
def exit(): Bottom = exit(0)  
def exit(status: Int): Bottom = {  
  java.lang.System.exit(status)  
  throw new Throwable()  
}  
  
def assert(assertion: Boolean): Unit =  
  if (!assertion)  
    throw new Error("assertion failed")  
  
def assert(assertion: Boolean, message: Any): Unit =  
  if (!assertion)  
    throw new Error("assertion failed: " + message)  
  
def assume(assumption: Boolean): Unit =  
  if (!assumption)  
    throw new Error("assumption failed")  
  
def assume(assumption: Boolean, message: Any): Unit =  
  if (!assumption)  
    throw new Error("assumption failed: " + message)  
}
```

Bibliography

- [Oa04] Martin Odersky and al. An Overview of the Scala Programming Language. Technical Report IC/2004/64, EPFL Lausanne, Switzerland, 2004.
- [OCRZ03] Martin Odersky, Vincent Cremet, Christine Röckl, and Matthias Zenger. A Nominal Theory of Objects with Dependent Types. In *Proc. ECOOP'03*, Springer LNCS, July 2003.
- [Ode06] Martin Odersky. The Scala Experiment – Can We Provide Better Language Support for Component Systems? In *Proc. ACM Symposium on Principles of Programming Languages*, 2006.
- [OZ05a] Martin Odersky and Matthias Zenger. Independently Extensible Solutions to the Expression Problem. In *Proc. FOOL 12*, January 2005. <http://homepages.inf.ed.ac.uk/wadler/fool>.
- [OZ05b] Martin Odersky and Matthias Zenger. Scalable Component Abstractions. In *Proc. OOPSLA*, 2005.
- [W3C] W3C. Extensible Markup Language (XML). <http://www.w3.org/TR/REC-xml>.

Chapter A

Scala Syntax Summary

(this needs to be updated)

The lexical syntax of Scala is given by the following grammar in EBNF form.

```
upper      ::= 'A' | ... | 'Z' | '$' | '_' and Unicode Lu
lower      ::= 'a' | ... | 'z' and Unicode Ll
letter     ::= upper | lower and Unicode categories Lo, Lt, Nl
digit      ::= '0' | ... | '9'
special    ::= "all other characters in \u0020-007F and Unicode categories
               Sm, So except parentheses ([]) and periods"

op         ::= special {special}
varid      ::= lower idrest
plainid    ::= upper idrest
           | varid
           | op
id         ::= plainid
           | '\\' stringLit '\\'
idrest     ::= {letter | digit} ['_' op]

integerLiteral ::= (decimalNumeral | hexNumeral | octalNumeral) ['L' | 'l']
decimalNumeral ::= '0' | nonZeroDigit {digit}
hexNumeral     ::= '0' 'x' hexDigit {hexDigit}
octalNumeral   ::= '0' octalDigit {octalDigit}
digit          ::= '0' | nonZeroDigit
nonZeroDigit   ::= '1' | ... | '9'
octalDigit     ::= '0' | ... | '7'

floatingPointLiteral
    ::= digit {digit} '.' {digit} [exponentPart] [floatType]
    | '.' digit {digit} [exponentPart] [floatType]
```

```

| digit {digit} exponentPart [floatType]
| digit {digit} [exponentPart] floatType
exponentPart ::= ('E' | 'e') ['+' | '-'] digit {digit}
floatType    ::= 'F' | 'f' | 'D' | 'd'

booleanLiteral ::= true | false

characterLiteral ::= '\'' char '\''
| '\'' charEscapeSeq '\''

stringLiteral ::= "" {stringElement} ""
stringElement ::= charNoDoubleQuote
| charEscapeSeq

symbolLiteral ::= ''' plainid

comment ::= '/*' "any sequence of characters" '*/'
| '//' "any sequence of characters up to end of line"

```

The context-free syntax of Scala is given by the following EBNF grammar.

```

Literal      ::= integerLiteral
| floatingPointLiteral
| booleanLiteral
| characterLiteral
| stringLiteral
| symbolLiteral
| null

StableId     ::= id
| Path '.' id
Path         ::= StableId
| [id '.'] this
| [id '.''] super '[' id ']' '.' id

Type         ::= Type1 '=>' Type
| '(' [Types] ')' '=>' Type
| Type1

Type1        ::= SimpleType {with SimpleType} [Refinement]
SimpleType   ::= SimpleType TypeArgs
| SimpleType '#' id
| StableId
| Path '.' type
| '(' Type ')'

TypeArgs     ::= '[' Types ']'
Types        ::= Type {',' Type}

```

```

Refinement      ::= '{' [RefineStat {StatementSeparator RefineStat}] '}'
RefineStat      ::= Dcl
                  | type TypeDef
                  |

Exprs           ::= Expr {' Expr} [':' '_' '*']
Expr            ::= (Bindings | Id) '=>' Expr
                  | Expr1
Expr1           ::= if '(' Expr1 ')' [NewLine] Expr [[';'] else Expr]
                  | try '{' Block '}' [catch '{' CaseClauses '}']
                    [finally Expr]
                  | while '(' Expr ')' [NewLine] Expr
                  | do Expr [StatementSeparator] while '(' Expr ')'
                  | for '(' Enumerators ')' | '{' Enumerators '}'
                    [NewLine] [yield] Expr
                  | throw Expr
                  | return [Expr]
                  | SimpleExpr '.' id '=' Expr
                  | SimpleExpr ArgumentExprs '=' Expr
                  | PostfixExpr [':' Type1]
                  | PostfixExpr match '{' CaseClauses '}'
                  | MethodClosure
PostfixExpr      ::= InfixExpr [id [NewLine]]
InfixExpr        ::= PrefixExpr
                  | InfixExpr id [NewLine] PrefixExpr
PrefixExpr       ::= ['- ' | '+ ' | '~ ' | '! ' | '&'] SimpleExpr
SimpleExpr       ::= Literal
                  | Path
                  | '(' [Expr] ')'
                  | BlockExpr
                  | new Template
                  | SimpleExpr '.' id
                  | SimpleExpr TypeArgs
                  | SimpleExpr ArgumentExprs
                  | XmlExpr
ArgumentExprs    ::= '(' [Exprs] ')'
                  | BlockExpr
MethodClosure    ::= '.' Id {'.' Id | TypeArgs | ArgumentExprs}
BlockExpr        ::= '{' CaseClauses '}'
                  | '{' Block '}'
Block            ::= {BlockStat StatementSeparator} [ResultExpr]
BlockStat        ::= Import
                  | [implicit] Def
                  | {LocalModifier} TmplDef
                  | Expr1
                  |

```

```

ResultExpr      ::= Expr1
                  | (Bindings | Id [':' Type1]) '=>' Block

Enumerators     ::= Generator {StatementSeparator Enumerator}
Enumerator      ::= Generator
                  | val Pattern1 '=' Expr
                  | Expr
Generator       ::= val Pattern1 '<-' Expr

CaseClauses     ::= CaseClause { CaseClause }
CaseClause      ::= case Pattern ['if' PostfixExpr] '=>' Block

Constr          ::= StableId [TypeArgs] {'(' [Exprs] ')'}

Pattern         ::= Pattern1 { '|' Pattern1 }
Pattern1        ::= varid ':' Type
                  | '_' ':' Type
                  | Pattern2
Pattern2        ::= varid ['@' Pattern3]
                  | Pattern3
Pattern3        ::= SimplePattern [ '*' | '?' | '+' ]
                  | SimplePattern { id SimplePattern }
SimplePattern   ::= '_'
                  | varid
                  | Literal
                  | StableId [ '(' [Patterns] ')' ]
                  | '(' [Patterns] ')'
                  | XmlPattern
Patterns        ::= Pattern {',' Pattern}

TypeParamClause ::= [NewLine] '[' VariantTypeParam {',' VariantTypeParam} ']'
FunTypeParamClause ::= [NewLine] '[' TypeParam {',' TypeParam} ']'
VariantTypeParam ::= ['+' | '-'] TypeParam
TypeParam       ::= id [>: Type] [<: Type] [<% Type]
ParamClauses    ::= {ParamClause} [[NewLine] '(' implicit Params ')']
ParamClause     ::= [NewLine] '(' [Params] ')'}
Params          ::= Param {',' Param}
Param           ::= id ':' ParamType
ParamType       ::= ['=>'] Type ['*']
ClassParamClauses ::= {ClassParamClause} [[NewLine]
                                     '(' implicit ClassParams ')']
ClassParamClause ::= [NewLine] '(' [ClassParams] ')'
ClassParams     ::= ClassParam {'' ClassParam}
ClassParam      ::= [{Modifier} ('val' | 'var')] Param
Bindings        ::= '(' Binding {',' Binding ')'}
Binding         ::= id [':' Type]

```

```

Modifier      ::= LocalModifier
                | override
                | private [ "[" id "]" ]
                | protected [ "[" id "]" ]

LocalModifier ::= abstract
                | final
                | sealed
                | implicit

AttributeClause ::= '[' Attribute {',' Attribute} ']' [NewLine]
Attribute       ::= Constr

Template        ::= TemplateParents [TemplateBody]
TemplateParents ::= Constr {with SimpleType}
TemplateBody    ::= '{' [TemplateStat {StatementSeparator TemplateStat}] '}'
TemplateStat    ::= Import
                | {AttributeClause} {Modifier} Def
                | {AttributeClause} {Modifier} Dcl
                | Expr
                |

Import          ::= import ImportExpr {',' ImportExpr}
ImportExpr     ::= StableId '.' (id | '_' | ImportSelectors)
ImportSelectors ::= '{' {ImportSelector ','}
                (ImportSelector | '_') '}'
ImportSelector ::= id ['=>' id | '=>' '_']

Dcl            ::= val ValDcl
                | var VarDcl
                | def FunDcl
                | type TypeDcl

ValDcl         ::= ids ':' Type
VarDcl         ::= ids ':' Type
FunDcl         ::= FunSig ':' Type
FunSig         ::= id [FunTypeParamClause] {ParamClause}
TypeDcl        ::= id [>: Type] [<: Type]

Def            ::= val PatDef
                | var VarDef
                | def FunDef
                | type TypeDef
                | TmplDef
PatDef         ::= Pattern2 {',' Pattern2} [':' Type] '=' Expr
VarDef         ::= ids [':' Type] '=' Expr

```

```

FunDef      ::= | ids ':' Type '=' '_'
               ::= FunSig ':' Type '=' Expr
               | this ParamClause ParamClauses '=' ConstrExpr
TypeDef     ::= id [TypeParamClause] '=' Type

TplDef     ::= ([case] class ClassDef
               | [case] object ObjectDef
               | trait TraitDef
ClassDef    ::= id [TypeParamClause] ClassParamClauses
               ['requires' SimpleType] ClassTemplate
ClassTemplate ::= [extends TemplateParents] [[NewLine] TemplateBody]
TraitDef    ::= id [TypeParamClause]
               ['requires' SimpleType] TraitTemplate
TraitTemplate ::= [extends MixinParents] [[NewLine] TemplateBody]
MixinParents ::= SimpleType {'with' SimpleType}

ObjectDef   ::= id ClassTemplate

ConstrExpr  ::= SelfInvocation
               | '{' SelfInvocation {StatementSeparator BlockStat} '}'
SelfInvocation ::= this ArgumentExprs {ArgumentExprs}

CompilationUnit ::= [package QualId StatementSeparator] TopStatSeq
TopStatSeq    ::= TopStat {StatementSeparator TopStat}
TopStat       ::= {AttributeClause} {Modifier} TmplDef
               | Import
               | Packaging
               |
Packaging     ::= package QualId '{' TopStatSeq '}'

QualId        ::= id {'.' id}
ids           ::= id {'.', ' id}

```

Chapter B

Changes between Scala version 1.0 and 2.0

Scala in its second version is different in some details from the first version of the language. There have been several additions and some old idioms are no longer supported. This appendix summarizes the main changes.

B.1 New Keywords

The following three words are now reserved; they cannot be used as identifiers (§1.1)

`implicit` `match` `requires`

B.2 Newlines as Statement Separators

Newlines can now be used as statement separators in place of semicolons (§1.2)

B.3 Syntax Restrictions

There are some other situations where old constructs no longer work:

Pattern matching expressions. The `match` keyword now appears only as infix operator between a selector expression and a number of cases, as in:

```
expr match {  
  case Some(x) => ...  
  case None => ...
```

```
}
```

Variants such as `expr.match {...}` or just `match {...}` are no longer supported.

“With” in extends clauses. The idiom

```
class C with M { ... }
```

is no longer supported. A **with** connective is only allowed following an **extends** clause. For instance, the line above would have to be written

```
class C extends AnyRef with M { ... } .
```

However, assuming M is a trait (see 5.3.3), it is also legal to write

```
class C extends M { ... }
```

The latter expression is treated as equivalent to

```
class C extends S with M { ... }
```

where S is the superclass of M.

Regular Expression Patterns. The only form of regular expression pattern that is currently supported is a sequence pattern, which might end in a sequence wildcard `_*`. Example:

```
case List(1, 2, _*) => ... // will match all lists starting with \code{1,2}.
```

It is at current not clear whether this is a permanent restriction. We are evaluating the possibility of re-introducing full regular expression patterns in Scala.

B.4 Selftype Annotations

The recommended syntax of selftype annotations has changed.

```
class C: T extends B { ... }
```

becomes

```
class C requires T extends B { ... }
```

That is, selftypes are now indicated by the new **requires** keyword. The old syntax is still available but is considered deprecated. Conversions

B.5 For-comprehensions

For-comprehensions (§6.16) now admit value and pattern definitions. Example:

```
for {  
  val x <- List.range(1, 100)  
  val y <- List.range(1, x)  
  val z = x + y  
  isPrime(z)  
} yield Pair(x, y)
```

Note the definition `val z = x + y` as the third item in the for-comprehension.

B.6 Conversions

The rules for implicit conversions of methods to functions (§6.22) have been tightened. Previously, a parameterized method used as a value was always implicitly converted to a function. This could lead to unexpected results when method arguments were forgotten. Consider for instance the statement below:

```
show(x.toString)
```

where `show` is defined as follows:

```
def show(x: String) = Console.println(x) .
```

Most likely, the programmer forgot to supply an empty argument list `()` to `toString`. The previous Scala version would treat this code as a partially applied method, and expand it to:

```
show(() => x.toString())
```

As a result, the address of a closure would be printed instead of the value of `s`.

Scala version 2.0 will apply a conversion from partially applied method to function value only if the expected type of the expression is indeed a function type. For instance, the conversion would not be applied in the code above because the expected type of `show`'s parameter is `String`, not a function type.

The new convention disallows some previously legal code. Example:

```
def sum(f: int => double)(a: int, b: int): double =  
  if (a > b) 0 else f(a) + sum(f)(a + 1, b)  
  
val sumInts = sum(x => x) // error: missing arguments
```

The partial application of `sum` in the last line of the code above will not be converted to a function type. Instead, the compiler will produce an error message which states that arguments for method `sum` are missing. The problem can be fixed by providing an expected type for the partial application, for instance by annotating the definition of `sumInts` with its type:

```
val sumInts: (int, int) => double = sum(x => x) // OK
```

On the other hand, Scala version 2.0 now automatically applies methods with empty parameter lists to `()` argument lists when necessary. For instance, the `show` expression above will now be expanded to

```
show(x.toString()) .
```

Scala version 2.0 also relaxes the rules of overriding with respect to empty parameter lists. The revised definition of *matching members* (§5.1.3) makes it now possible to override a method with an explicit, but empty parameter list `()` with a parameterless method, and *vice versa*. For instance, the following class definition is now legal:

```
class C {  
  override def toString: String = ...  
}
```

Previously this definition would have been rejected, because the `toString` method as inherited from `java.lang.Object` takes an empty parameter list.

B.7 Class Parameters

A class parameter may now be prefixed by **val** or **var** (§5.3).

B.8 Private Qualifiers

Previously, Scala had three levels of visibility: *private*, *protected* and *public*. There was no way to restrict accesses to members of the current package, as in Java. Scala 2 now defines access qualifiers that let one express this level of visibility, among others. In the definition

```
private[C] def f(...)
```

access to `f` is restricted to all code within the class or package `C` (which must contain the definition of `f`) (§5.2)

B.9 Changes in the Mixin Model

The model which details mixin composition of classes has changed significantly. The main differences are:

1. We now distinguish between *traits* that are used as mixin classes and normal classes. The syntax of traits has been generalized from version 1.0, in that traits are now allowed to have mutable fields. However, as in version 1.0, traits may still do not have constructor parameters.
2. Member resolution and super accesses are now both defined in terms of a *class linearization*.
3. Scala's notion of method overloading has been generalized; in particular, it is now possible to have overloaded variants of the same method in a subclass and in a superclass, or in several different mixins. This makes method overloading in Scala conceptually the same as in Java.

The new mixin model is explained in more detail in §5.

B.10 Implicit Parameters

Views in Scala 1.0 have been replaced by the more general concept of implicit parameters (§7)

B.11 Flexible Typing of Pattern Matching

The new version of Scala implements more flexible typing rules when it comes to pattern matching over heterogeneous class hierarchies (§8.2. A *heterogeneous class hierarchy* is one where subclasses inherit a common superclass with different parameter types. With the new rules in Scala version 2.0 one can perform pattern matches over such hierarchies with more precise typings that keep track of the information gained by comparing the types of a selector and a matching pattern (§Example 8.2.1). This gives Scala capabilities analogous to guarded algebraic data types.