The Caml Light system, release 0.6
Documentation and user's manual

Xavier Leroy

September 1993

Contents

I Getting started **7**

1 Installation instructions **9**
1.1 The Unix version 9
1.2 The Macintosh version 9
1.3 The PC versions 10

II The Caml Light language reference manual **17**

2 The Caml Light language reference manual **19**
2.1 Lexical conventions 20
2.2 Global names 22
2.3 Values 23
2.4 Type expressions 24
2.5 Constants 25
2.6 Patterns 26
2.7 Expressions 28
2.8 Global definitions 36
2.9 Directives 37
2.10 Module implementations 37
2.11 Module interfaces 38

3 Language extensions **41**
3.1 Streams, parsers, and printers 41
3.2 Range patterns 42
3.3 Recursive definitions of values 42
3.4 Mutable variant types 43
3.5 Directives 43

III The Caml Light library **45**

4 The core library **47**
4.1 `bool`: boolean operations 47
4.2 `builtin`: base types and constructors 48
4.3 `char`: character operations 48
4.4 `eq`: equality functions 49

4.5 `exc`: exceptions . . . 49
4.6 `fchar`: character operations, without sanity checks . . . 50
4.7 `float`: operations on floating-point numbers . . . 50
4.8 `fstring`: string operations, without sanity checks . . . 52
4.9 `fvect`: operations on vectors, without sanity checks . . . 52
4.10 `int`: operations on integers . . . 52
4.11 `io`: buffered input and output . . . 54
4.12 `list`: operations on lists . . . 60
4.13 `pair`: operations on pairs . . . 62
4.14 `ref`: operations on references . . . 63
4.15 `stream`: operations on streams . . . 63
4.16 `string`: string operations . . . 64
4.17 `vect`: operations on vectors . . . 65

5 The standard library **69**
5.1 `arg`: parsing of command line arguments . . . 69
5.2 `baltree`: basic balanced binary trees . . . 70
5.3 `filename`: operations on file names . . . 71
5.4 `genlex`: a generic lexical analyzer . . . 72
5.5 `hashtbl`: hash tables and hash functions . . . 73
5.6 `lexing`: the run-time library for lexers generated by camllex . . . 74
5.7 `map`: association tables over ordered types . . . 75
5.8 `parsing`: the run-time library for parsers generated by mlyacc . . . 76
5.9 `printexc`: a catch-all exception handler . . . 76
5.10 `printf`: formatting printing functions . . . 77
5.11 `queue`: queues . . . 77
5.12 `random`: pseudo-random number generator . . . 78
5.13 `set`: sets over ordered types . . . 79
5.14 `sort`: sorting and merging lists . . . 80
5.15 `stack`: stacks . . . 80
5.16 `sys`: system interface . . . 81

6 The graphics library **83**
6.1 `graphics`: machine-independent graphics primitives . . . 85

IV The Caml Light commands **91**

7 Batch compilation (camlc) **93**
7.1 Overview of the compiler . . . 93
7.2 Options . . . 94
7.3 Modules and the file system . . . 96
7.4 Common errors . . . 96

8 The toplevel system (camllight) **99**
8.1 Options . . . 100
8.2 Toplevel control functions . . . 101
8.3 The toplevel and the module system . . . 102

8.4 Common errors . . . 104
8.5 Building custom toplevel systems: `camlmktop` . . . 105
8.6 Options . . . 105

9 The runtime system (camlrun) 107
9.1 Overview . . . 107
9.2 Options . . . 107
9.3 Common errors . . . 108

10 The librarian (camllibr) 109
10.1 Overview . . . 109
10.2 Options . . . 110
10.3 Turning code into a library . . . 110

11 Lexer and parser generators (camllex, camlyacc) 113
11.1 Overview of `camllex` . . . 113
11.2 Syntax of lexer definitions . . . 114
11.3 Overview of `camlyacc` . . . 116
11.4 Syntax of grammar definitions . . . 116
11.5 Options . . . 118
11.6 A complete example . . . 118

12 Interfacing C with Caml Light 121
12.1 Overview and compilation information . . . 121
12.2 The `value` type . . . 123
12.3 Representation of Caml Light data types . . . 124
12.4 Operations on values . . . 125
12.5 Living in harmony with the garbage collector . . . 127
12.6 A complete example . . . 129

V Appendix 133

13 Further reading 135
13.1 Programming in ML . . . 135
13.2 Descriptions of ML dialects . . . 136
13.3 Implementing functional programming languages . . . 137
13.4 Applications of ML . . . 137

Foreword

This manual documents the release 0.6 of the Caml Light system. It is organized as follows.

- Part I, "Getting started", explains how to install Caml Light on your machine.
- Part II, "The Caml Light language reference manual", is the reference description of the Caml Light language.
- Part III, "The Caml Light library", describes the modules provided in the standard library.
- Part IV, "The Caml Light commands", documents the Caml Light compiler, toplevel system, and programming utilities.
- Part V, "Appendix", contains a short bibliography, an index of all identifiers defined in the standard library, and an index of Caml Light keywords.

Conventions

The Caml Light system comes in several versions: for Unix machines, for Macintoshes, and for PCs. The parts of this manual that are specific to one version are presented as shown below:

Unix: This is material specific to the Unix version.

Mac: This is material specific to the Macintosh version.

86: This is material specific to the 8086 (and 80286) PC version.

386: This is material specific to the 80386 PC version.

PC: This is material specific to the two PC versions.

License

The Caml Light system is copyright © 1989, 1990. 1991, 1992, 1993 Institut National de Recherche en Informatique et en Automatique (INRIA). INRIA holds all ownership rights to the Caml Light system. See the file `COPYRIGHT` in the distribution for the copyright notice.

The Caml Light system can be freely copied, but not sold. More precisely, INRIA grants any user of the Caml Light system the right to reproduce it, provided that the copies are distributed free of charge and under the conditions given in the `COPYRIGHT` file.

Availability by FTP

The complete Caml Light distribution resides on the machine `ftp.inria.fr`. The distribution files can be transferred by anonymous FTP:

Host:	`ftp.inria.fr` (Internet address `192.93.2.54`)
Login name:	`anonymous`
Password:	your e-mail address
Directory:	`lang/caml-light`
Files:	see the index in file `README`

Part I

Getting started

Chapter 1

Installation instructions

This chapter explains how to install Caml Light on your machine.

1.1 The Unix version

Requirements. Any machine that runs under one of the various flavors of the Unix operating system, and that has a flat, non-segmented, 32-bit or 64-bit address space. 4M of RAM, 2M of free disk space. The graphics library requires X11 release 4 or 5.

Installation. The Unix version is distributed in source format, as a compressed `tar` file named `cl6unix.tar.Z`. To extract, move to the directory where you want the source files to reside, transfer `cl6unix.tar.Z` to that directory, and execute

```
zcat cl6unix.tar.Z | tar xBf -
```

This extracts the source files in the current directory. The file `INSTALL` contains complete instructions on how to configure, compile and install Caml Light. Read it and follow the instructions.

Troubleshooting. See the file `INSTALL`.

1.2 The Macintosh version

Requirements. Any Macintosh with at least 1M of RAM (2M is recommended), running System 6 or 7. About 850K of free space on the disk. The parts of the Caml Light system that support batch compilation currently require the Macintosh Programmer's Workshop (MPW) version 3.2. MPW is Apple's development environment, and it is distributed by APDA, Apple's Programmers and Developers Association. See the file `READ ME` in the distribution for APDA's address.

Installation. Create the folder where the Caml Light files will reside. Double-click on the file `cl6macbin.sea` from the distribution. This displays a file dialog box. Open the folder where the Caml Light files will reside, and click on the `Extract` button. This will re-create all files from the distribution in the Caml Light folder.

To test the installation, double-click on the application `Caml Light`. The "Caml Light output" window should display something like

```
>       Caml Light version 0.6

#
```

In the "Caml Light input" window, enter `1+2;;` and press the `Return` key. The "Caml Light output" window should display:

```
>       Caml Light version 0.6

#1+2;;
- : int = 3
#
```

Select "Quit" from the "File" menu to return to the Finder.

If you have MPW, you can install the batch compilation tools as follows. The tools and scripts from the `tools` folder must reside in a place where MPW will find them as commands. There are two ways to achieve this result: either copy the files in the `tools` folder to the `Tools` or the `Scripts` folder in your MPW folder; or keep the files in the `tools` folder and add the following line to your `UserStartup` file (assuming Caml Light resides in folder `Caml Light` on the disk named `My HD`):

```
Set Commands "{Commands},My HD:Caml Light:tools:"
```

In either case, you now have to edit the `camlc` script, and replace the string

```
Macintosh HD:Caml Light:lib:
```

(in the first line) with the actual pathname of the `lib` folder. For example, if you put Caml Light in folder `Caml Light` on the disk named `My HD`, the first line of `camlc` should read:

```
Set stdlib "My HD:Caml Light:lib:"
```

Troubleshooting. Here is one commonly encountered problem.

`Cannot find file stream.zi`
: (Displayed in the "Caml Light output" window, with an alert box telling you that Caml Light has terminated abnormally.) This is an installation error. The folder named `lib` in the distribution must always be in the same folder as the `Caml Light` application. It's OK to move the application to another folder; but remember to move the `lib` directory to the same folder. (To return to the Finder, first select "Quit" from the "File" menu.)

1.3 The PC versions

Two versions are distributed for the PC. The first one, dubbed "8086 PC" here, runs on all PCs, whatever their processor is (8088, 8086, 80286, 80386 DX or SX, 80486 DX or SX, and so on). The second one, dubbed "80386 PC", runs in 32-bit protected mode, and is therefore faster and takes advantage of memory above the standard 640K, but runs only on PCs with 80386, 80486 or Pentium processors. The binaries for these two versions are distributed as two `zip` archive files `cl6pc86bin.zip` and `cl6pc386bin.zip`.

In the following, we assume that the distribution files resides in drive `A:`, and that the hard disk on which you are installing Caml Light is drive `C:`. If this is not the case, replace `A:` and `C:` by the appropriate drives.

1.3.1 The 8086 PC version

Requirements. A PC running MS-DOS version 3.1 or later. 640K of RAM. About 800K of free space on the disk. The graphics primitives require a CGA, EGA, or VGA video card.

Installation. Create a directory on the hard disk where Caml Light will reside. In the following, we assume that this directory is named `C:\caml86`. If you choose a different directory, replace `C:\caml86` in the following by the appropriate absolute path name. Then, execute the following commands:

```
cd C:\caml86
A:pkunzip -d A:cl6pc86
```

(Be careful not to omit the `-d` option to `pkunzip`.) This extracts all files in the `C:\caml86` directory. Then, edit the `C:\autoexec.bat` file, in order to:

- add `C:\caml86\bin` to the `PATH` variable; that is, transform the line that reads

  ```
  SET PATH=C:\dos;...
  ```

  into

  ```
  SET PATH=C:\dos;...;C:\caml86\bin
  ```

- insert the following line

  ```
  SET CAMLLIB=C:\caml86\lib
  ```

Then, save the `autoexec.bat` file and restart the machine. To test the installation, execute:

```
camlc -v
```

The `camlc` command should print something like:

```
The Caml Light system for the 8086 PC, version 0.6
  (standard library from C:\caml86\lib)
The Caml Light runtime system, version 0.6
The Caml Light compiler, version 0.6
The Caml Light linker, version 0.6
```

Then, execute:

```
caml
```

The `caml` command should print something like:

```
>       Caml Light version 0.6

#
```

In response to the `#` prompt, type:

```
quit();;
```

This should get you back to the DOS command interpreter.

If you wish to use the graphics primitives, you must install a BGI driver corresponding to your video card in the `C:\caml86\lib` directory. You will need the BGI driver `CGA.BGI` for a CGA card, `EGAVGA.BGI` for a VGA or EGA card. You may also use other BGI drivers, but they must work in 16-color mode: 256-color modes are not supported. The standard BGI drivers are distributed with Borland's Turbo C/C++ and Turbo Pascal. Just copy the appropriate driver to the `C:\caml86\lib` directory.

Troubleshooting. Here are some commonly encountered problems.

`Cannot find the bytecode file` **or** `camlrun.exe: No such file`
: The installation has been performed incorrectly. Double-check the `autoexec.bat` file for errors in setting the `PATH` and `CAMLLIB` variables.

`Out of memory`
: Caml Light barely fits into the ridiculously small memory space provided by MS-DOS. Hence, you must make sure that as much of the standard 640K of RAM as possible are free before running Caml Light. As a rule of thumb, 500K of available RAM is the bare minimum; 600K is still far from sufficient. To free memory, remove as many device drivers (by suppressing lines in the `config.sys` file) and TSR programs (by suppressing lines in the `autoexec.bat` file) as possible, and restart the machine. Alternatively, try to load your drivers and TSRs outside of the standard 640K zone, with the help of memory managers such as QEMM or the tools in MS-DOS 5 and later. For more details, consult the MS-DOS documentation or your favorite PC magazine.

1.3.2 The 80386 PC version

Requirements. A PC equipped with a 80386 or 80486 processor, running MS-DOS version 3.3 or later. 2M of RAM. About 1.2M of free space on the disk. The graphics primitives require a VGA or SuperVGA video card.

Installation. Create a directory on the hard disk where Caml Light will reside. In the following, we assume that this directory is named `C:\caml386`. If you choose a different directory, replace `C:\caml386` in the following by the appropriate absolute path name. Then, execute the following commands:

```
cd C:\caml386
A:pkunzip -d A:cl6pc386
```

(Be careful not to omit the `-d` option to `pkunzip`.) This extracts all files in the `C:\caml386` directory.

Select or create a directory where Caml Light will put its temporary files. Many machines already have a `C:\tmp` directory for that purpose. If it does not exist, create it.

For the remainder of the configuration process, you will have to determine two things:

- does your machine contain floating-point hardware — that is, a 387 coprocessor, a 486DX (or higher) processor, or a 487SX (co-)processor? (If you really don't know, assume no floating-point hardware.)

- what kind of SuperVGA card do you have? Caml Light has graphics primitives that work on any VGA card in 320x200 pixels, 256 colors, but it can take advantage of the extra possibilities of various SuperVGA cards to work at higher resolution. To do so, you must determine which chipset is used in your SuperVGA card. Re-read the documentation for the card, then look at the files with extension `.GRD` or `.GRN` (the graphics drivers) in directory `C:\caml386\dev`, and find one whose name seems to match the name of the chipset. If you can't determine which graphics driver to use, don't worry: you'll stick with the default VGA graphics, that's all.

Now, edit the `C:\autoexec.bat` file, in order to:

- Add `C:\caml386\bin` to the `PATH` variable; that is, transform the line that reads

  ```
  SET PATH=C:\dos;...
  ```

  into

  ```
  SET PATH=C:\dos;...;C:\caml386\bin
  ```

- Insert the following lines

  ```
  SET CAMLLIB=C:\caml386\lib
  SET GO32TMP=C:\tmp
  ```

- If your machine has floating-point hardware, insert the following line:

  ```
  SET GO32=driver C:\caml386\dev\graph.grd gw 640 gh 480
  ```

  where `graph.grd` stands for the name of the graphics driver for your SuperVGA card, as determined above. The `640` and `480` specify the default graphics resolution to use; you can put `800` and `600`, or `1024` and `768` instead, depending on your taste and the capabilities of your card.

  If you were unable to determine the correct graphics driver, do not insert anything, leaving the `GO32` variable undefined.

- If your machine has no floating-point hardware, insert the following line:

  ```
  SET GO32=emu C:\caml386\dev\emu387
           driver C:\caml386\dev\graph.grd gw 640 gh 480
  ```

  (on a single line; we had to fold the line above for typesetting) where `graph.grd` stands for the name of the graphics driver for your SuperVGA card, as determined above. As explained in the previous item, you can choose another default graphics resolution instead of `640` and `480`.

  If you were unable to determine the correct graphics driver, insert the following line instead:

  ```
  SET GO32=emu C:\caml386\dev\emu387
  ```

Then, save the `autoexec.bat` file and restart the machine. To test the installation, execute:

```
camlc -v
```

The `camlc` command should print something like:

```
The Caml Light system for the 80386 PC, version 0.6
  (standard library from C:\caml386\lib)
The Caml Light runtime system, version 0.6
The Caml Light compiler, version 0.6
The Caml Light linker, version 0.6
```

Then, execute:

```
caml
```

The `caml` command should print something like:

```
>       Caml Light version 0.6

#
```

In response to the `#` prompt, type:

```
quit();;
```

This should get you back to the DOS command interpreter.

Troubleshooting. Here are some commonly encountered problems.

`Cannot find the bytecode file` **or** `camlrun.exe: No such file`
: The installation has been performed incorrectly. Double-check the `autoexec.bat` file for errors in setting the `PATH` and `CAMLLIB` variables.

`CPU must be a 386 to run this program`
: Self-explanatory. You'll have to content yourself with the 8086 PC version.

`CPU must be in REAL mode (not V86 mode) to run this program`
: Ah. That's a tricky one. A number of utility programs switch the 80386 processor to a particular mode, called "Virtual 86" or "V86" mode, that prevents 32-bit protected-mode applications like the 80386 PC version of Caml Light from starting. These programs include:

  - device drivers that provide memory management services
  - device drivers that provide enhanced debugging possibilities, such as `tdh386.sys` from Turbo Debugger.

  The 80386 PC version cannot start when any of these programs is active. To solve the problem, remove the guilty device drivers from your `config.sys` file.

  On the other hand, the 80386 PC version knows how to cohabit with DPMI-compliant or VCPI-compliant environments and memory managers. These include Windows 3, Desqview, and the EMM386, QEMM386 and 386MAX memory managers. In the case of `EMM386.EXE` from the MS-DOS distribution, EMS must be enabled (do not give it the `NOEMS` option). If you run the 80386 PC under a VCPI-compliant memory manager, configure the memory manager so that it allocates at least 1M of EMS, and preferably 2M or more.

Caml Light runs slowly and does a lot of disk accesses

When Caml Light cannot allocate the RAM it requires, it starts paging to a disk file, which considerably slows down execution. To avoid this, make sure that at least 1M or memory is available to Caml Light, and preferably 2M. Caml Light uses XMS memory if you are not running under a VCPI-compliant memory manager, and EMS memory if you are running under a VCPI-compliant memory manager. In the latter case, make sure to configure the memory manager so that it can allocate enough EMS.

Part II

The Caml Light language reference manual

Chapter 2

The Caml Light language reference manual

Foreword

This document is intended as a reference manual for the Caml Light language. It lists all language constructs, and gives their precise syntax and informal semantics. It is by no means a tutorial introduction to the language: there is not a single example. A good working knowledge of the language, as provided by the companion tutorial *Functional programming using Caml Light*, is assumed.

No attempt has been made at mathematical rigor: words are employed with their intuitive meaning, without further definition. As a consequence, the typing rules have been left out, by lack of the mathematical framework required to express them, while they are definitely part of a full formal definition of the language. The reader interested in truly formal descriptions of languages from the ML family is referred to *The definition of Standard ML* and *Commentary on Standard ML*, by Milner, Tofte and Harper, MIT Press.

Warning

Even though there is currently only one implementation of the Caml Light language, this document carefully distinguishes the language and its implementation(s). Implementations can provide extra language constructs; moreover, all points left unspecified in this reference manual can be interpreted as the language implementor wishes. All these implementation-dependent features are of course subject to change in future releases; only the features specified in this reference manual can be relied upon.

Notations

The syntax of the language is given in BNF-like notation. Terminal symbols are set in typewriter font (`like this`). Non-terminal symbols are set in italic font (*like that*). Square brackets [...] denote optional components. Curly brackets {...} denotes zero, one or several repetitions of the enclosed components. Curly bracket with a trailing plus sign $\{\ldots\}^+$ denote one or several repetitions of the enclosed components. Parentheses (...) denote grouping.

2.1 Lexical conventions

Blanks

The following characters are considered as blanks: space, newline, horizontal tabulation, carriage return, line feed and form feed. Blanks are ignored, but they separate adjacent identifiers, literals and keywords that would otherwise be confused as one single identifier, literal or keyword.

Comments

Comments are introduced by the two characters `(*`, with no intervening blanks, and terminated by the characters `*)`, with no intervening blanks. Comments are treated as blank characters. Comments do not occur inside string or character literals. Nested comments are correctly handled.

Identifiers

$$\begin{array}{rcl} \textit{ident} & ::= & \textit{letter}\ \{\textit{letter} \mid \texttt{0}\ldots\texttt{9} \mid \texttt{_}\} \\ \textit{letter} & ::= & \texttt{A}\ldots\texttt{Z} \mid \texttt{a}\ldots\texttt{z} \end{array}$$

Identifiers are sequences of letters, digits and `_` (the underscore character), starting with a letter. Letters contain at least the 52 lowercase and uppercase letters from the ASCII set. Implementations can recognize as letters other characters from the extended ASCII set. Identifiers cannot contain two adjacent underscore characters (`__`). Implementation may limit the number of characters of an identifier, but this limit must be above 256 characters. All characters in an identifier are meaningful.

Integer literals

$$\begin{array}{rcl} \textit{integer-literal} & ::= & [\texttt{-}]\ \{\texttt{0}\ldots\texttt{9}\}^{+} \\ & \mid & [\texttt{-}]\ (\texttt{0x} \mid \texttt{0X})\ \{\texttt{0}\ldots\texttt{9} \mid \texttt{A}\ldots\texttt{F} \mid \texttt{a}\ldots\texttt{f}\}^{+} \\ & \mid & [\texttt{-}]\ (\texttt{0o} \mid \texttt{0O})\ \{\texttt{0}\ldots\texttt{7}\}^{+} \\ & \mid & [\texttt{-}]\ (\texttt{0b} \mid \texttt{0B})\ \{\texttt{0}\ldots\texttt{1}\}^{+} \end{array}$$

An integer literal is a sequence of one or more digits, optionally preceded by a minus sign. By default, integer literals are in decimal (radix 10). The following prefixes select a different radix:

`0x`, `0X`	hexadecimal (radix 16)
`0o`, `0O`	octal (radix 8)
`0b`, `0B`	binary (radix 2).

(The initial `0` is the digit zero; the `O` for octal is the letter O.)

Floating-point literals

$$\textit{float-literal} \quad ::= \quad [\texttt{-}]\ \{\texttt{0}\ldots\texttt{9}\}^{+}\ [\texttt{.}\ \{\texttt{0}\ldots\texttt{9}\}]\ [(\texttt{e} \mid \texttt{E})\ [\texttt{+} \mid \texttt{-}]\ \{\texttt{0}\ldots\texttt{9}\}^{+}]$$

Floating-point decimals consist in an integer part, a decimal part and an exponent part. The integer part is a sequence of one or more digits, optionally preceded by a minus sign. The decimal part is a decimal point followed by zero, one or more digits. The exponent part is the character `e` or `E` followed by an optional `+` or `-` sign, followed by one or more digits. The decimal part or the exponent part can be omitted, but not both to avoid ambiguity with integer literals.

Character literals

char-literal ::= ‘ *regular-char* ‘
| ‘ \ (\ | ‘ | n | t | b | r) ‘
| ‘ \ (0...9) (0...9) (0...9) ‘

Character literals are delimited by ‘ (backquote) characters. The two backquotes enclose either one character different from ‘ and \, or one of the escape sequences below:

Sequence	Character denoted
\\	backslash (\)
\‘	backquote (‘)
\n	newline (LF)
\r	return (CR)
\t	horizontal tabulation (TAB)
\b	backspace (BS)
ddd	the character with ASCII code *ddd* in decimal

String literals

string-literal ::= " {*string-character*} "

string-character ::= *regular-char*
| \ (\ | " | n | t | b | r)
| \ (0...9) (0...9) (0...9)

String literals are delimited by " (double quote) characters. The two double quotes enclose a sequence of either characters different from " and \, or escape sequences from the table below:

Sequence	Character denoted
\\	backslash (\)
\"	double quote (")
\n	newline (LF)
\r	return (CR)
\t	horizontal tabulation (TAB)
\b	backspace (BS)
ddd	the character with ASCII code *ddd* in decimal

Implementations must support string literals up to $2^{16} - 1$ characters in length (65535 characters).

Keywords

The identifiers below are reserved as keywords, and cannot be employed otherwise:

```
and     as      begin       do      done      downto
else    end     exception   for     fun       function
if      in      let         match   mutable   not
of      or      prefix      rec     then      to
try     type    value       where   while     with
```

The following character sequences are also keywords:

```
#    !    !=   &    (    )    *    *.   +    +.
,    -    -.   ->   .    .(   /    /.   :    ::
:=   ;    ;;   <    <.   <-   <=   <=.  <>   <>.
=    =.   ==   >    >.   >=   >=.  @    [    [|
]    ^    _    __   {    |    |]   }    '
```

Ambiguities

Lexical ambiguities are resolved according to the "longest match" rule: when a character sequence can be decomposed into two tokens in several different ways, the decomposition retained is the one with the longest first token.

2.2 Global names

Global names are used to denote value variables, value constructors (constant or non-constant), type constructors, and record labels. Internally, a global name consists of two parts: the name of the defining module (the module name), and the name of the global inside that module (the local name). The two parts of the name must be valid identifiers. Externally, global names have the following syntax:

global-name ::= *ident*
| *ident* `__` *ident*

The form *ident* `__` *ident* is called a qualified name. The first identifier is the module name, the second identifier is the local name. The form *ident* is called an unqualified name. The identifier is the local name; the module name is omitted. The compiler infers this module name following the completion rules given below, therefore transforming the unqualified name into a full global name.

To complete an unqualified identifier, the compiler checks a list of modules, the opened modules, to see if they define a global with the same local name as the unqualified identifier. When one is found, the identifier is completed into the full name of that global. That is, the compiler takes as module name the name of an opened module that defines a global with the same local name as the unqualified identifier. If several modules satisfy this condition, the one that comes first in the list of opened modules is selected.

The list of opened modules always includes the module currently being compiled (checked first). (In the case of a toplevel-based implementation, this is the module where all toplevel definitions are entered.) It also includes a number of standard library modules that provide the initial environment (checked last). In addition, the `#open` and `#close` directives can be used to add or remove modules from that list. The modules added with `#open` are checked after the module currently being compiled, but before the initial standard library modules.

variable	::=	*global-name*
	|	`prefix` *operator-name*
operator-name	::=	`+` | `-` | `*` | `/` | `mod` | `+.` | `-.` | `*.` | `/.`
	|	`@` | `^` | `!` | `:=` | `=` | `<>` | `==` | `!=` | `!`
	|	`<` | `<=` | `>` | `<=` | `<.` | `<=.` | `>.` | `<=.`
cconstr	::=	*global-name*
	|	`[]`
	|	`()`
ncconstr	::=	*global-name*
	|	`prefix ::`
typeconstr	::=	*global-name*
label	::=	*global-name*

Depending on the context, global names can stand for global variables (*variable*), constant value constructors (*cconstr*), non-constant value constructors (*ncconst*), type constructors (*typeconstr*), or record labels (*label*). For variables and value constructors, special names built with `prefix` and an operator name are recognized. The tokens `[]` and `()` are also recognized as built-in constant constructors (the empty list and the unit value).

The syntax of the language restricts labels and type constructors to appear in certain positions, where no other kind of global names are accepted. Hence labels and type constructors have their own name spaces. Value constructors and value variables live in the same name space: a global name in value position is interpreted as a value constructor if it appears in the scope of a type declaration defining that constructor; otherwise, the global name is taken to be a value variable. For value constructors, the type declaration determines whether a constructor is constant or not.

2.3 Values

This section describes the kinds of values that are manipulated by Caml Light programs.

2.3.1 Base values

Integer numbers

Integer values are integer numbers from -2^{30} to $2^{30} - 1$, that is -1073741824 to 1073741823. Implementations may support a wider range of integer values.

Floating-point numbers

Floating-point values are numbers in floating-point representation. Everything about floating-point values is implementation-dependent, including the range of representable numbers, the number of significant digits, and the way floating-point results are rounded.

Characters

Character values are represented as 8-bit integers between 0 and 255. Character codes between 0 and 127 are interpreted following the ASCII standard. The interpretation of character codes

between 128 and 255 is implementation-dependent.

Character strings

String values are finite sequences of characters. Implementations must support strings up to $2^{16}-1$ characters in length (65535 characters). Implementations may support longer strings.

2.3.2 Tuples

Tuples of values are written $(v_1, \ldots, v_n)$, standing for the n-tuple of values v_1 to v_n. Tuples of up to $2^{14}-1$ elements (16383 elements) must be supported, though implementations may support tuples with more elements.

2.3.3 Records

Record values are labeled tuples of values. The record value written $\{label_1 = v_1; \ldots; label_n = v_n\}$ associates the value v_i to the record label $label_i$, for $i = 1 \ldots n$. Records with up to $2^{14}-1$ fields (16383 fields) must be supported, though implementations may support records with more fields.

2.3.4 Arrays

Arrays are finite, variable-sized sequences of values of the same type. Arrays of length up to $2^{14}-1$ (16383 elements) must be supported, though implementations may support larger arrays.

2.3.5 Variant values

Variant values are either a constant constructor, or a pair of a non-constant constructor and a value. The former case is written *cconstr*; the latter case is written *ncconstr*(v), where v is said to be the argument of the non-constant constructor *ncconstr*.

The following constants are treated like built-in constant constructors:

`false`	the boolean false
`true`	the boolean true
`()`	the "unit" value
`[]`	the empty list

2.3.6 Functions

Functional values are mappings from values to values.

2.4 Type expressions

typexpr ::= `'` *ident*
| `(` *typexpr* `)`
| *typexpr* `->` *typexpr*
| *typexpr* {`*` *typexpr*}$^+$
| *typeconstr*
| *typexpr typeconstr*
| `(` *typexpr* {`,` *typexpr*} `)` *typeconstr*

The table below shows the relative precedences and associativity of operators and non-closed type constructions. The constructions with higher precedences come first.

Operator	Associativity
Type constructor application	–
`*`	–
`->`	right

Type expressions denote types in definitions of data types as well as in type constraints over patterns and expressions.

Type variables

The type expression `'` *ident* stands for the type variable named *ident*. In data type definitions, type variables are names for the data type parameters. In type constraints, they represent unspecified types that can be instantiated by any type to satisfy the type constraint.

Parenthesized types

The type expression `(` *typexpr* `)` denotes the same type as *typexpr*.

Function types

The type expression $typexpr_1$ `->` $typexpr_2$ denotes the type of functions mapping arguments of type $typexpr_1$ to results of type $typexpr_2$.

Tuple types

The type expression $typexpr_1$ `*` ... `*` $typexpr_n$ denotes the type of tuples whose elements belong to types $typexpr_1, \ldots typexpr_n$ respectively.

Constructed types

Type constructors with no parameter, as in *typeconstr*, are type expressions.

The type expression *typexpr typeconstr*, where *typeconstr* is a type constructor with one parameter, denotes the application of the unary type constructor *typeconstr* to the type *typexpr*.

The type expression $(typexpr_1, \ldots, typexpr_n)$ *typeconstr*, where *typeconstr* is a type constructor with n parameters, denotes the application of the n-ary type constructor *typeconstr* to the types $typexpr_1$ through $typexpr_n$.

2.5 Constants

constant ::= *integer-literal*
| *float-literal*
| *char-literal*
| *string-literal*
| *cconstr*

The syntactic class of constants comprises literals from the four base types (integers, floating-point numbers, characters, character strings), and constant constructors.

2.6 Patterns

pattern ::= *ident*
| `_`
| *pattern* `as` *ident*
| `(` *pattern* `)`
| `(` *pattern* `:` *typexpr* `)`
| *pattern* `|` *pattern*
| *constant*
| *ncconstr pattern*
| *pattern* `,` *pattern* {`,` *pattern*}
| `{` *label* `=` *pattern* {`;` *label* `=` *pattern*} `}`
| `[]`
| `[` *pattern* {`;` *pattern*} `]`
| *pattern* `::` *pattern*

The table below shows the relative precedences and associativity of operators and non-closed pattern constructions. The constructions with higher precedences come first.

Operator	Associativity
Constructor application	–
`::`	right
`,`	–
`|`	left
`as`	–

Patterns are templates that allow selecting data structures of a given shape, and binding identifiers to components of the data structure. This selection operation is called pattern matching; its outcome is either "this value does not match this pattern", or "this value matches this pattern, resulting in the following bindings of identifiers to values".

Variable patterns

A pattern that consists in an identifier matches any value, binding the identifier to the value. The pattern `_` also matches any value, but does not bind any identifier.

Alias patterns

The pattern *pattern*$_1$ `as` *ident* matches the same values as *pattern*$_1$. If the matching against *pattern*$_1$ is successful, the identifier *ident* is bound to the matched value, in addition to the bindings performed by the matching against *pattern*$_1$.

Parenthesized patterns

The pattern `(` *pattern*$_1$ `)` matches the same values as *pattern*$_1$. A type constraint can appear in a parenthesized patterns, as in `(` *pattern*$_1$ `:` *typexpr* `)`. This constraint forces the type of *pattern*$_1$ to be compatible with *type*.

"Or" patterns

The pattern *pattern*$_1$ `|` *pattern*$_2$ represents the logical "or" of the two patterns *pattern*$_1$ and *pattern*$_2$. A value matches *pattern*$_1$ `|` *pattern*$_2$ either if it matches *pattern*$_1$ or if it matches *pattern*$_2$. The two sub-patterns *pattern*$_1$ and *pattern*$_2$ must contain no identifiers. Hence no bindings are returned by matching against an "or" pattern.

Constant patterns

A pattern consisting in a constant matches the values that are equal to this constant.

Variant patterns

The pattern *ncconstr pattern*$_1$ matches all variants whose constructor is equal to *ncconstr*, and whose argument matches *pattern*$_1$.

The pattern *pattern*$_1$ `::` *pattern*$_2$ matches non-empty lists whose heads match *pattern*$_1$, and whose tails match *pattern*$_2$. This pattern behaves like `prefix :: (` *pattern*$_1$ `,` *pattern*$_2$ `)`.

The pattern `[` *pattern*$_1$ `;` ... `;` *pattern*$_n$ `]` matches lists of length n whose elements match *pattern*$_1$ ... *pattern*$_n$, respectively. This pattern behaves like *pattern*$_1$ `::` ... `::` *pattern*$_n$ `::` `[]`.

Tuple patterns

The pattern *pattern*$_1$ `,` ... `,` *pattern*$_n$ matches n-tuples whose components match the patterns *pattern*$_1$ through *pattern*$_n$. That is, the pattern matches the tuple values $(v_1, \ldots, v_n)$ such that *pattern*$_i$ matches v_i for $i = 1, \ldots, n$.

Record patterns

The pattern `{` *label*$_1$ `=` *pattern*$_1$ `;` ... `;` *label*$_n$ `=` *pattern*$_n$ `}` matches records that define at least the labels *label*$_1$ through *label*$_n$, and such that the value associated to *label*$_i$ match the pattern *pattern*$_i$, for $i = 1, \ldots, n$. The record value can define more labels than *label*$_1$ ... *label*$_n$; the values associated to these extra labels are not taken into account for matching.

2.7 Expressions

expr ::= *ident*
| *variable*
| *constant*
| `(` *expr* `)`
| `begin` *expr* `end`
| `(` *expr* `:` *typexpr* `)`
| *expr* `,` *expr* {`,` *expr*}
| *ncconstr expr*
| *expr* `::` *expr*
| `[` *expr* {`;` *expr*} `]`
| `[|` *expr* {`;` *expr*} `|]`
| `{` *label* `=` *expr* {`;` *label* `=` *expr*} `}`
| *expr expr*
| *prefix-op expr*
| *expr infix-op expr*
| *expr* `.` *label*
| *expr* `.` *label* `<-` *expr*
| *expr* `.(` *expr* `)`
| *expr* `.(` *expr* `)` `<-` *expr*
| *expr* `&` *expr*
| *expr* `or` *expr*
| `if` *expr* `then` *expr* [`else` *expr*]
| `while` *expr* `do` *expr* `done`
| `for` *ident* `=` *expr* (`to` | `downto`) *expr* `do` *expr* `done`
| *expr* `;` *expr*
| `match` *expr* `with` *simple-matching*
| `fun` *multiple-matching*
| `function` *simple-matching*
| `try` *expr* `with` *simple-matching*
| `let` [`rec`] *let-binding* {`and` *let-binding*} `in` *expr*

simple-matching ::= *pattern* `->` *expr* {`|` *pattern* `->` *expr*}

multiple-matching ::= *pattern-list* `->` *expr* {`|` *pattern-list* `->` *expr*}

pattern-list ::= *pattern* {*pattern*}

let-binding ::= *pattern* `=` *expr*
| *variable pattern-list* `=` *expr*

prefix-op ::= `-` | `-.` | `!`

infix-op ::= `+` | `-` | `*` | `/` | `mod` | `+.` | `-.` | `*.` | `/.` | `@` | `^` | `!` | `:=`
| `=` | `<>` | `==` | `!=` | `<` | `<=` | `>` | `<=` | `<.` | `<=.` | `>.` | `<=.`

The table below shows the relative precedences and associativity of operators and non-closed constructions. The constructions with higher precedence come first.

Construction or operator	Associativity
`!`	–
`.  .(`	–
function application	right
constructor application	–
`-  -.` (prefix)	–
`mod`	left
`*  *.  /  /.`	left
`+  +.  -  -.`	left
`::`	right
`@ ^`	right
comparisons (`=  ==  <` etc.)	left
`not`	–
`&`	left
`or`	left
`,`	–
`<-  :=`	right
`if`	–
`;`	right
`let match fun function try`	–

2.7.1 Simple expressions

Constants

Expressions consisting in a constant evaluate to this constant.

Variables

Expressions consisting in a variable evaluate to the value bound to this variable in the current evaluation environment. The variable can be either a qualified identifier or a simple identifier. Qualified identifiers always denote global variables. Simple identifiers denote either a local variable, if the identifier is locally bound, or a global variable, whose full name is obtained by qualifying the simple identifier, as described in section 2.2.

Parenthesized expressions

The expressions `(` *expr* `)` and `begin` *expr* `end` have the same value as *expr*. Both constructs are semantically equivalent, but it is good style to use `begin`...`end` inside control structures:

```
if ... then begin ... ; ... end else begin ... ; ... end
```

and `(`...`)` for the other grouping situations.

Parenthesized expressions can contain a type constraint, as in `(` *expr* `:` *type* `)`. This constraint forces the type of *expr* to be compatible with *type*.

Function abstraction

The most general form of function abstraction is:

$$\begin{array}{lll} \texttt{fun} & pattern_1^1 \ \ldots \ pattern_1^m & \texttt{->} \ \ expr_1 \\ \texttt{|} & \ldots & \\ \texttt{|} & pattern_n^1 \ \ldots \ pattern_n^m & \texttt{->} \ \ expr_n \end{array}$$

This expression evaluates to a functional value with m curried arguments. When this function is applied to m values $v_1 \ldots v_m$, the values are matched against each pattern row $pattern_i^1 \ \ldots \ pattern_i^m$, for i from 1 to n. If one of these matchings succeeds, that is if the value v_j matches the pattern $pattern_i^j$ for all $j = 1 \ldots m$, then the expression $expr_i$ associated to the selected pattern row is evaluated, and its value becomes the value of the function application. The evaluation of $expr_i$ takes place in an environment enriched by the bindings performed during the matching.

If several pattern rows match the arguments, the one that occurs first in the function definition is selected. If none of the pattern rows matches the argument, the exception `Match_failure` is raised.

If the function above is applied to less than m arguments, a functional value is returned, that represents the partial application of the function to the provided arguments. This partial application is a function that, when applied to the remaining arguments, matches all arguments against the pattern rows as described above. Matching does not start until all m arguments have been provided to the function; hence, partial applications of the function to less than m arguments never raise `Match_failure`.

All pattern rows in the function body must contain the same number of patterns. A variable must not be bound more than once in one pattern row.

Functions with only one argument can be defined with the `function` keyword instead of `fun`:

$$\begin{array}{lll} \texttt{function} & pattern_1 & \texttt{->} \ \ expr_1 \\ \texttt{|} & \ldots & \\ \texttt{|} & pattern_n & \texttt{->} \ \ expr_n \end{array}$$

The function thus defined behaves exactly as described above. The only difference between the two forms of function definition is how a parsing ambiguity is resolved. The two forms *cconstr pattern* (two patterns in a row) and *ncconstr pattern* (one pattern) cannot be distinguished syntactically. Function definitions introduced by `fun` resolve the ambiguity to the former form; function definitions introduced by `function` resolve it to the latter form (the former form makes no sense in this case).

Function application

Function application is denoted by juxtaposition of expressions. The expression $expr_1 \ \ expr_2 \ldots expr_n$ evaluates the expressions $expr_1$ to $expr_n$. The expression $expr_1$ must evaluate to a functional value, which is then applied to the values of $expr_2, \ldots, expr_n$. The order in which the expressions $expr_1, \ldots, expr_n$ are evaluated is not specified.

Local definitions

The `let` and `let rec` constructs bind variables locally. The construct

$$\texttt{let} \ pattern_1 \ \texttt{=} \ expr_1 \ \texttt{and} \ldots \texttt{and} \ pattern_n \ \texttt{=} \ expr_n \ \texttt{in} \ expr$$

evaluates $expr_1 \ldots expr_n$ in some unspecified order, then matches their values against the patterns $pattern_1 \ldots pattern_n$. If the matchings succeed, *expr* is evaluated in the environment enriched by the bindings performed during matching, and the value of *expr* is returned as the value of the whole `let` expression. If one of the matchings fails, the exception `Match_failure` is raised.

An alternate syntax is provided to bind variables to functional values: instead of writing

$$ident \texttt{ = fun } pattern_1 \ldots pattern_m \texttt{ -> } expr$$

in a `let` expression, one may instead write

$$ident\ pattern_1 \ldots pattern_m \texttt{ = } expr$$

Both forms bind *ident* to the curried function with m arguments and only one case,

$$pattern_1 \ldots pattern_m \texttt{ -> } expr.$$

Recursive definitions of variables are introduced by `let rec`:

$$\texttt{let rec } pattern_1 \texttt{ = } expr_1 \texttt{ and} \ldots \texttt{and } pattern_n \texttt{ = } expr_n \texttt{ in } expr$$

The only difference with the `let` construct described above is that the bindings of variables to values performed by the pattern-matching are considered already performed when the expressions $expr_1$ to $expr_n$ are evaluated. That is, the expressions $expr_1$ to $expr_n$ can reference identifiers that are bound by one of the patterns $pattern_1, \ldots, pattern_n$, and expect them to have the same value as in *expr*, the body of the `let rec` construct.

The recursive definition is guaranteed to behave as described above if the expressions $expr_1$ to $expr_n$ are function definitions (`fun`... or `function`...), and the patterns $pattern_1 \ldots pattern_n$ consist in a single variable, as in:

$$\texttt{let rec } ident_1 \texttt{ = fun} \ldots \texttt{and} \ldots \texttt{and } ident_n \texttt{ = fun} \ldots \texttt{ in } expr$$

This defines $ident_1 \ldots ident_n$ as mutually recursive functions local to *expr*. The behavior of other forms of `let rec` definitions is implementation-dependent.

2.7.2 Control constructs

Sequence

The expression $expr_1$; $expr_2$ evaluates $expr_1$ first, then $expr_2$, and returns the value of $expr_2$.

Conditional

The expression `if` $expr_1$ `then` $expr_2$ `else` $expr_3$ evaluates to the value of $expr_2$ if $expr_1$ evaluates to the boolean `true`, and to the value of $expr_3$ if $expr_1$ evaluates to the boolean `false`.

The `else` $expr_3$ part can be omitted, in which case it defaults to `else ()`.

Case expression

The expression

```
match  expr
 with  pattern1  ->  expr1
     |  ...
     |  patternn  ->  exprn
```

matches the value of *expr* against the patterns $pattern_1$ to $pattern_n$. If the matching against $pattern_i$ succeeds, the associated expression $expr_i$ is evaluated, and its value becomes the value of the whole `match` expression. The evaluation of $expr_i$ takes place in an environment enriched by the bindings performed during matching. If several patterns match the value of *expr*, the one that occurs first in the `match` expression is selected. If none of the patterns match the value of *expr*, the exception `Match_failure` is raised.

Boolean operators

The expression $expr_1$ `&` $expr_2$ evaluates to `true` if both $expr_1$ and $expr_2$ evaluate to `true`; otherwise, it evaluates to `false`. The first component, $expr_1$, is evaluated first. The second component, $expr_2$, is not evaluated if the first component evaluates to `false`. Hence, the expression $expr_1$ `&` $expr_2$ behaves exactly as

`if` $expr_1$ `then` $expr_2$ `else false`.

The expression $expr_1$ `or` $expr_2$ evaluates to `true` if one of $expr_1$ and $expr_2$ evaluates to `true`; otherwise, it evaluates to `false`. The first component, $expr_1$, is evaluated first. The second component, $expr_2$, is not evaluated if the first component evaluates to `true`. Hence, the expression $expr_1$ `or` $expr_2$ behaves exactly as

`if` $expr_1$ `then true else` $expr_2$.

Loops

The expression `while` $expr_1$ `do` $expr_2$ `done` repeatedly evaluates $expr_2$ while $expr_1$ evaluates to `true`. The loop condition $expr_1$ is evaluated and tested at the beginning of each iteration. The whole `while`...`done` expression evaluates to the unit value `()`.

The expression `for` *ident* `=` $expr_1$ `to` $expr_2$ `do` $expr_3$ `done` first evaluates the expressions $expr_1$ and $expr_2$ (the boundaries) into integer values n and p. Then, the loop body $expr_3$ is repeatedly evaluated in an environment where the local variable named *ident* is successively bound to the values n, $n+1$, ..., $p-1$, p. The loop body is never evaluated if $n > p$.

The expression `for` *ident* `=` $expr_1$ `downto` $expr_2$ `do` $expr_3$ `done` first evaluates the expressions $expr_1$ and $expr_2$ (the boundaries) into integer values n and p. Then, the loop body $expr_3$ is repeatedly evaluated in an environment where the local variable named *ident* is successively bound to the values n, $n-1$, ..., $p+1$, p. The loop body is never evaluated if $n < p$.

In both cases, the whole `for` expression evaluates to the unit value `()`.

Exception handling

The expression

```
try   expr
with  pattern1  ->  expr1
    |  ...
    |  patternn  ->  exprn
```

evaluates the expression *expr* and returns its value if the evaluation of *expr* does not raise any exception. If the evaluation of *expr* raises an exception, the exception value is matched against the patterns $pattern_1$ to $pattern_n$. If the matching against $pattern_i$ succeeds, the associated expression $expr_i$ is evaluated, and its value becomes the value of the whole `try` expression. The evaluation of $expr_i$ takes place in an environment enriched by the bindings performed during matching. If several patterns match the value of *expr*, the one that occurs first in the `try` expression is selected. If none of the patterns matches the value of *expr*, the exception value is raised again, thereby transparently "passing through" the `try` construct.

2.7.3 Operations on data structures

Products

The expression $expr_1$ `,` ... `,` $expr_n$ evaluates to the n-tuple of the values of expressions $expr_1$ to $expr_n$. The evaluation order for the subexpressions is not specified.

Variants

The expression *ncconstr expr* evaluates to the variant value whose constructor is *ncconstr*, and whose argument is the value of *expr*.

For lists, some syntactic sugar is provided. The expression $expr_1$ `::` $expr_2$ stands for the constructor `prefix ::` applied to the argument `(` $expr_1$ `,` $expr_2$ `)`, and therefore evaluates to the list whose head is the value of $expr_1$ and whose tail is the value of $expr_2$. The expression `[` $expr_1$ `;` ... `;` $expr_n$ `]` is equivalent to $expr_1$ `::` ... `::` $expr_n$ `::` `[]`, and therefore evaluates to the list whose elements are the values of $expr_1$ to $expr_n$.

Records

The expression `{` $label_1$ `=` $expr_1$ `;` ... `;` $label_n$ `=` $expr_n$ `}` evaluates to the record value $\{label_1 = v_1; \ldots; label_n = v_n\}$, where v_i is the value of $expr_i$ for $i = 1, \ldots, n$. The labels $label_1$ to $label_n$ must all belong to the same record types; all labels belonging to this record type must appear exactly once in the record expression, though they can appear in any order. The order in which $expr_1$ to $expr_n$ are evaluated is not specified.

The expression $expr_1$ `.` *label* evaluates $expr_1$ to a record value, and returns the value associated to *label* in this record value.

The expression $expr_1$ `.` *label* `<-` $expr_2$ evaluates $expr_1$ to a record value, which is then modified in-place by replacing the value associated to *label* in this record by the value of $expr_2$. This operation is permitted only if *label* has been declared `mutable` in the definition of the record type. The whole expression $expr_1$ `.` *label* `<-` $expr_2$ evaluates to the unit value `()`.

Arrays

The expression `[|` $expr_1$ `;` ... `;` $expr_n$ `|]` evaluates to a n-element array, whose elements are initialized with the values of $expr_1$ to $expr_n$ respectively. The order in which these expressions are evaluated is unspecified.

The expression $expr_1$ `.(` $expr_2$ `)` is equivalent to the application `vect_item` $expr_1$ $expr_2$. In the initial environment, the identifier `vect_item` resolves to a built-in function that returns the value of element number $expr_2$ in the array denoted by $expr_1$. The first element has number 0; the last element has number $n - 1$, where n is the size of the array. The exception `Invalid_argument` is raised if the access is out of bounds.

The expression $expr_1$ `.(` $expr_2$ `)` `<-` $expr_3$ is equivalent to `vect_assign` $expr_1$ $expr_2$ $expr_3$. In the initial environment, the identifier `vect_assign` resolves to a built-in function that modifies in-place the array denoted by $expr_1$, replacing element number $expr_2$ by the value of $expr_3$. The exception `Invalid_argument` is raised if the access is out of bounds. The built-in function returns `()`. Hence, the whole expression $expr_1$ `.(` $expr_2$ `)` `<-` $expr_3$ evaluates to the unit value `()`.

This behavior of the two constructs $expr_1$ `.(` $expr_2$ `)` and $expr_1$ `.(` $expr_2$ `)` `<-` $expr_3$ may change if the meaning of the identifiers `vect_item` and `vect_assign` is changed, either by redefinition or by modification of the list of opened modules. See the discussion below on operators.

2.7.4 Operators

The operators written `infix-op` in the grammar above can appear in infix position (between two expressions). The operators written `prefix-op` in the grammar above can appear in prefix position (in front of an expression).

The expression *prefix-op expr* is interpreted as the application *ident expr*, where *ident* is the identifier associated to the operator *prefix-op* in the table below. Similarly, the expression $expr_1$ *infix-op* $expr_2$ is interpreted as the application *ident* $expr_1$ $expr_2$, where *ident* is the identifier associated to the operator *infix-op* in the table below. The identifiers written *ident* above are then evaluated following the rules in section 2.7.1. In the initial environment, they evaluate to built-in functions whose behavior is described in the table. The behavior of the constructions *prefix-op expr* and $expr_1$ *infix-op* $expr_2$ may change if the meaning of the identifiers associated to *prefix-op* or *infix-op* is changed, either by redefinition of the identifiers, or by modification of the list of opened modules, through the `#open` and `#close` directives.

Operator	Associated identifier	Behavior in the default environment
`+`	`prefix +`	Integer addition.
`-` (infix)	`prefix -`	Integer subtraction.
`-` (prefix)	`minus`	Integer negation.
`*`	`prefix *`	Integer multiplication.
`/`	`prefix /`	Integer division. Raise `Division_by_zero` if second argument is null. The result is unspecified if either argument is negative.
`mod`	`prefix mod`	Integer modulus. Raise `Division_by_zero` if second argument is null. The result is unspecified if either argument is negative.
`+.`	`prefix +.`	Floating-point addition.
`-.` (infix)	`prefix -.`	Floating-point subtraction.
`-.` (prefix)	`minus_float`	Floating-point negation.
`*.`	`prefix *.`	Floating-point multiplication.
`/.`	`prefix /.`	Floating-point division. The result is unspecified if second argument is null.
`@`	`prefix @`	List concatenation.
`^`	`prefix ^`	String concatenation.
`!`	`prefix !`	Dereferencing (return the current contents of a reference).
`:=`	`prefix :=`	Reference assignment (update the reference given as first argument with the value of the second argument).
`=`	`prefix =`	Structural equality test.
`<>`	`prefix <>`	Structural inequality test.
`==`	`prefix ==`	Physical equality test.
`!=`	`prefix !=`	Physical inequality test.
`<`	`prefix <`	Test "less than" on integers.
`<=`	`prefix <=`	Test "less than or equal " on integers.
`>`	`prefix >`	Test "greater than" on integers.
`>=`	`prefix >=`	Test "greater than or equal" on integers.
`<.`	`prefix <.`	Test "less than" on floating-point numbers.
`<=.`	`prefix <=.`	Test "less than or equal " on floating-point numbers.
`>.`	`prefix >.`	Test "greater than" on floating-point numbers.
`>=.`	`prefix >=.`	Test "greater than or equal" on floating-point numbers.

The behavior of the `+`, `-`, `*`, `/`, `mod`, `+.`, `-.`, `*.` or `/.` operators is unspecified if the result falls outside of the range of representable integers or floating-point numbers, respectively. See chapter 4 for a more precise description of the behavior of the operators above.

2.8 Global definitions

This section describes the constructs that bind global identifiers (value variables, value constructors, type constructors, record labels).

2.8.1 Type definitions

type-definition ::= `type` *typedef* {`and` *typedef*}

typedef ::= *type-params ident* `=` *constr-decl* {`|` *constr-decl*}
| *type-params ident* `=` `{` *label-decl* {`;` *label-decl*} `}`
| *type-params ident* `==` *typexpr*
| *type-params ident*

type-params ::= *nothing*
| `'` *ident*
| `(` `'` *ident* {`,` `'` *ident*} `)`

constr-decl ::= *ident*
| *ident* `of` *typexpr*

label-decl ::= *ident* `:` *typexpr*
| `mutable` *ident* `:` *typexpr*

Type definitions bind type constructors to data types: either variant types, record types, type abbreviations, or abstract data types.

Type definitions are introduced by the `type` keyword, and consist in one or several simple definitions, possibly mutually recursive, separated by the `and` keyword. Each simple definition defines one type constructor.

A simple definition consists in an identifier, possibly preceded by one or several type parameters, and followed by a data type description. The identifier is the local name of the type constructor being defined. (The module name for this type constructor is the name of the module being compiled.) The optional type parameters are either one type variable `'` *ident*, for type constructors with one parameter, or a list of type variables (`'` $ident_1, \ldots,$ `'` $ident_n$), for type constructors with several parameters. These type parameters can appear in the type expressions of the right-hand side of the definition.

Variant types

The type definition *typeparams ident* `=` $constr\text{-}decl_1$ `|` ... `|` $constr\text{-}decl_n$ defines a variant type. The constructor declarations $constr\text{-}decl_1, \ldots, constr\text{-}decl_n$ describe the constructors associated to this variant type. The constructor declaration *ident* `of` *typexpr* declares the local name *ident* (in the module being compiled) as a non-constant constructor, whose argument has type *typexpr*. The constructor declaration *ident* declares the local name *ident* (in the module being compiled) as a constant constructor.

Record types

The type definition *typeparams ident* `=` `{` $label\text{-}decl_1$ `;` ... `;` $label\text{-}decl_n$ `}` defines a record type. The label declarations $label\text{-}decl_1, \ldots, label\text{-}decl_n$ describe the labels associated to this record type. The label declaration *ident* `:` *typexpr* declares the local name *ident* in the module being compiled as a

label, whose argument has type *typexpr*. The label declaration `mutable` *ident* : *typexpr* behaves similarly; in addition, it allows physical modification over the argument to this label.

Type abbreviations

The type definition *typeparams ident* `==` *typexpr* defines the type constructor *ident* as an abbreviation for the type expression *typexpr*.

Abstract types

The type definition *typeparams ident* defines *ident* as an abstract type. When appearing in a module interface, this definition allows exporting a type constructor while hiding how it is represented in the module implementation.

2.8.2 Exception definitions

exception-definition ::= `exception` *constr-decl* {`and` *constr-decl*}

Exception definitions add new constructors to the built-in variant type `exn` of exception values. The constructors are declared as for a definition of a variant type.

2.9 Directives

directive ::= `#` `open` *string*
| `#` `close` *string*
| `#` *ident string*

Directives control the behavior of the compiler. They apply to the remainder of the current compilation unit.

The two directives `#open` and `#close` modify the list of opened modules, that the compiler uses to complete unqualified identifiers, as described in section 2.2. The directive `#open` *string* adds the module whose name is given by the string constant *string* to the list of opened modules, in first position. The directive `#close` *string* removes the first occurrence of the module whose name is given by the string constant *string* from the list of opened modules.

Implementations can provide other directives, provided they follow the syntax `#` *ident string*, where *ident* is the name of the directive, and the string constant *string* is the argument to the directive. The behavior of these additional directives is implementation-dependent.

2.10 Module implementations

implementation ::= {*impl-phrase* `;;`}

impl-phrase ::= *expr*
| *value-definition*
| *type-definition*
| *exception-definition*
| *directive*

value-definition ::= `let` [`rec`] *let-binding* {`and` *let-binding*}

A module implementation consists in a sequence of implementation phrases, terminated by double semicolons. An implementation phrase is either an expression, a value definition, a type or exception definition, or a directive. At run-time, implementation phrases are evaluated sequentially, in the order in which they appear in the module implementation.

Implementation phrases consisting in an expression are evaluated for their side-effects.

Value definitions bind global value variables in the same way as a `let...in...` expression binds local variables. The expressions are evaluated, and their values are matched against the left-hand sides of the `=` sides, as explained in section 2.7.1. If the matching succeeds, the bindings of identifiers to values performed during matching are interpreted as bindings to the global value variables whose local name is the identifier, and whose module name is the name of the module. If the matching fails, the exception `Match_failure` is raised. The scope of these bindings is the phrases that follow the value definition in the module implementation.

Type and exception definitions introduce type constructors, variant constructors and record labels as described in sections 2.8.1 and 2.8.2. The scope of these definitions is the phrases that follow the value definition in the module implementation. The evaluation of an implementation phrase consisting in a type or exception definition produces no effect at run-time.

Directives modify the behavior of the compiler on the subsequent phrases of the module implementation, as described in section 2.9. The evaluation of an implementation phrase consisting in a directive produces no effect at run-time. Directives apply only to the module currently being compiled; in particular, they have no effect on other modules that refer to globals exported by the module being compiled.

2.11 Module interfaces

interface ::= {*intf-phrase* ;;}

intf-phrase ::= *value-declaration*
| *type-definition*
| *exception-definition*
| *directive*

value-declaration ::= `value` *ident* : *typexpr* {`and` *ident* : *typexpr*}

Module interfaces declare the global objects (value variables, type constructors, variant constructors, record labels) that a module exports, that is, makes available to other modules. Other modules can refer to these globals using qualified identifiers or the `#open` directive, as explained in section 2.2.

A module interface consists in a sequence of interface phrases, terminated by double semicolons. An interface phrase is either a value declaration, a type definition, an exception definition, or a directive.

Value declarations declare global value variables that are exported by the module implementation, and the types with which they are exported. The module implementation must define these variables, with types at least as general as the types declared in the interface. The scope of the bindings for these global variables extends from the module implementation itself to all modules that refer to those variables.

Type or exception definitions introduce type constructors, variant constructors and record labels as described in sections 2.8.1 and 2.8.2. Exception definitions and type definitions that are not

abstract type declarations also take effect in the module implementation; that is, the type constructors, variant constructors and record labels they define are considered bound on entrance to the module implementation, and can be referred to by the implementation phrases. Type definitions that are not abstract type declarations must not be redefined in the module implementation. In contrast, the type constructors that are declared abstract in a module interface must be defined in the module implementation, with the same names.

Directives modify the behavior of the compiler on the subsequent phrases of the module interface, as described in section 2.9. Directives apply only to the interface currently being compiled; in particular, they have no effect on other modules that refer to globals exported by the interface being compiled.

Chapter 3

Language extensions

This chapter describes the language features that are implemented in Caml Light, but not described in the Caml Light reference manual. In contrast with the fairly stable kernel language that is described in the reference manual, the extensions presented here are still experimental, and may be removed or changed in the future.

3.1 Streams, parsers, and printers

Caml Light comprises a built-in type for *streams* (possibly infinite sequences of elements, that are evaluated on demand), and associated stream expressions, to build streams, and stream patterns, to destructure streams. Streams and stream patterns provide a natural approach to the writing of recursive-descent parsers.

Streams are presented by the following extensions to the syntactic classes of expressions:

expr	::=	...
	|	`[<` `>]`
	|	`[<` *stream-component* {`;` *stream-component*} `>]`
	|	`function` *stream-pattern* `->` *expr* {`|` *stream-pattern* `->` *expr*}
	|	`match` *expr* `with` *stream-pattern* `->` *expr* {`|` *stream-pattern* `->` *expr*}
stream-component	::=	`'` *expr*
	|	*expr*
stream-pattern	::=	`[<` `>]`
	|	`[<` *stream-comp-pat* {`;` *stream-comp-pat*} `>]`
stream-comp-pat	::=	`'` *pattern*
	|	*expr pattern*
	|	*ident*

Stream expressions are bracketed by `[<` and `>]`. They represent the concatenation of their components. The component `'` *expr* represents the one-element stream whose element is the value of *expr*. The component *expr* represents a sub-stream. For instance, if both `s` and `t` are streams of integers, then `[<'1; s; t; '2>]` is a stream of integers containing the element `1`, then the elements of `s`, then those of `t`, and finally `2`. The empty stream is denoted by `[< >]`.

Unlike any other kind of expressions in the language, stream expressions are submitted to lazy evaluation: the components are not evaluated when the stream is built, but only when they are accessed during stream matching. The components are evaluated once, the first time they are accessed; the following accesses reuse the value computed the first time.

Stream patterns, also bracketed by `[<` and `>]`, describe initial segments of streams. In particular, the stream pattern `[< >]` matches all streams. Stream pattern components are matched against the corresponding elements of a stream. The component `'` *pattern* matches the corresponding stream element against the pattern. The component *expr pattern* applies the function denoted by *expr* to the current stream, then matches the result of the function against *pattern*. Finally, the component *ident* simply binds the identifier to the stream being matched. (The current implementation limits *ident* to appear last in a stream pattern.)

Stream matching proceeds destructively: once a component has been matched, it is discarded from the stream (by in-place modification).

Stream matching proceeds in two steps: first, a pattern is selected by matching the stream against the first components of the stream patterns; then, the following components of the selected pattern are checked against the stream. If the following components do not match, the exception `Parse_error` is raised. There is no backtracking here: stream matching commits to the pattern selected according to the first element. If none of the first components of the stream patterns match, the exception `Parse_failure` is raised. The `Parse_failure` exception causes the next alternative to be tried, if it occurs during the matching of the first element of a stream, before matching has committed to one pattern.

See *Functional programming using Caml Light* for a more gentle introductions to streams, and for some examples of their use in writing parsers. A more formal presentation of streams, and a discussion of alternate semantics, can be found in *Parsers in ML* by Michel Mauny and Daniel de Rauglaudre, in the proceedings of the 1992 ACM conference on Lisp and Functional Programming.

3.2 Range patterns

In patterns, Caml Light recognizes the form ‘ c ‘ `..` ‘ d ‘ (two character constants separated by `..`) as a shorthand for the pattern

$$‘\ c\ ‘\ |\ ‘\ c_1\ ‘\ |\ ‘\ c_2\ ‘\ |\ \ldots |\ ‘\ c_n\ ‘\ |\ ‘\ d\ ‘$$

where $c_1, c_2, \ldots, c_n$ are the characters that occur between c and d in the ASCII character set. For instance, the pattern `'0'..'9'` matches all characters that are digits.

3.3 Recursive definitions of values

Besides `let rec` definitions of functional values, as described in the reference manual, Caml Light supports a certain class of recursive definitions of non-functional values. For instance, the following definition is accepted:

```
let rec x = 1 :: y and y = 2 :: x;;
```

and correctly binds `x` to the cyclic list `1 :: 2 :: 1 :: 2 ::...`, and `y` to the cyclic list `2 :: 1 :: 2 :: 1 ::...` Informally, the class of accepted definitions consists of those definitions where the defined variables occur only inside function bodies or as a field of a data structure. Moreover, the patterns in the left-hand sides must be identifiers, nothing more complex.

3.4 Mutable variant types

The argument of a value constructor can be declared "mutable" when the variant type is defined:

```
type foo = A of mutable int
         | B of mutable int * int
         | ...
```

This allows in-place modification of the argument part of a constructed value. Modification is performed by a new kind of expressions, written *ident* `<-` *expr*, where *ident* is an identifier bound by pattern-matching to the argument of a mutable constructor, and *expr* denotes the value that must be stored in place of that argument. Continuing the example above:

```
let x = A 1 in
  begin match x with A y -> y <- 2 | _ -> () end;
  x
```

returns the value `A 2`. The notation *ident* `<-` *expr* works also if *ident* is an identifier bound by pattern-matching to the value of a mutable field in a record. For instance,

```
type bar = {mutable lbl : int};;
let x = {lbl = 1} in
  begin match x with {lbl = y} -> y <- 2 end;
  x
```

returns the value `{lbl = 2}`.

3.5 Directives

In addition to the standard `#open` and `#close` directives, Caml Light provides three additional directives.

`#infix` " *id* "
: Change the lexical status of the identifier *id*: in the remainder of the compilation unit, *id* is recognized as an infix operator, just like `+`. The notation `prefix` *id* can be used to refer to the identifier *id* itself. Expressions of the form $expr_1$ *id* $expr_2$ are parsed as the application `prefix` *id* $expr_1$ $expr_2$. The argument to the `#infix` directive must be an identifier, that is, a sequence of letters, digits and underscores starting with a letter; otherwise, the `#infix` declaration has no effect. Example:

```
#infix "union";;
let prefix union = fun x y -> ... ;;
[1,2] union [3,4];;
```

`#uninfix` " *id* "
: Remove the infix status attached to the identifier *id* by a previous `#infix` " *id* " directive.

`#directory` " *dir-name* "
: Add the named directory to the path of directories searched for compiled module interface files. This is equivalent to the `-I` command-line option to the batch compiler and the toplevel system.

Part III

The Caml Light library

Chapter 4

The core library

This chapter describes the functions provided by the Caml Light core library. This library is special in two ways:

- It is automatically linked with the user's object code files by the `camlc` command (chapter 7). Hence, the globals defined by these libraries can be used in standalone programs without having to add any `.zo` file on the command line for the linking phase. Similarly, in interactive use, these globals can be used in toplevel phrases without having to load any `.zo` file in memory.
- The interfaces for the modules below are automatically "opened" when a compilation starts, or when the toplevel system is launched. Hence, it is possible to use unqualified identifiers to refer to the functions provided by these modules, without adding `#open` directives. Actually, the list of automatically opened modules depend on the `-O` option given to the compiler or to the toplevel system:

`-O` option	Opened modules (reverse opening order)
`-O cautious` (default)	`io`, `eq`, `int`, `float`, `ref`, `pair`, `list`, `vect`, `char`, `string`, `bool`, `exc`, `stream`
`-O fast`	`io`, `eq`, `int`, `float`, `ref`, `pair`, `list`, `fvect`, `fchar`, `fstring`, `bool`, `exc`, `stream`
`-O none`	none

Conventions

For easy reference, the modules are listed below in alphabetical order of module names. For each module, the declarations from its interface file are printed one by one in typewriter font, followed by a short comment. All modules and the identifiers they export are indexed at the end of this report.

4.1 `bool`: boolean operations

The boolean and is written `e1 & e2`. The boolean or is written `e1 or e2`. Both constructs are sequential, left-to-right: `e2` is evaluated only if needed. Actually, `e1 & e2` is equivalent to `if e1 then e2 else false`, and `e1 or e2` is equivalent to `if e1 then true else e2`.

```
value prefix not : bool -> bool
```

The boolean negation.

4.2 builtin: base types and constructors

This module defines some types and exceptions for which the language provides special syntax, and therefore wich are treated specially by the compiler.

```
type int
type float
type string
type char
```

The types of integers, floating-point numbers, character strings, and characters, respectively.

```
type exn
```

The type of exception values.

```
type bool = false | true
```

The type of boolean values.

```
type 'a vect
```

The type of arrays whose elements have type `'a`.

```
type unit = ()
```

The type of the unit value.

```
type 'a list = [] | prefix :: of 'a * 'a list
```

The type of lists.

```
exception Match_failure of string * int * int
```

The exception raised when a pattern-matching fails. The argument indicates the position in the source code of the pattern-matching (source file name, position of the first character of the matching, position of the last character.

4.3 char: character operations

```
value int_of_char : char -> int
```

Return the ASCII code of the argument.

```
value char_of_int : int -> char
```

Return the character with the given ASCII code. Raise `Invalid_argument "char_of_int"` if the argument is outside the range 0–255.

```
value char_for_read : char -> string
```

Return a string representing the given character, with special characters escaped following the lexical conventions of Caml Light.

4.4 eq: equality functions

`value prefix = : 'a -> 'a -> bool`

`e1 = e2` tests for structural equality of `e1` and `e2`. Mutable structures (e.g. references) are equal if and only if their current contents are structurally equal, even if the two mutable objects are not the same physical object. Equality between functional values raises `Invalid_argument`. Equality between cyclic data structures may not terminate.

`value prefix <> : 'a -> 'a -> bool`

Negation of `prefix =`.

`value prefix == : 'a -> 'a -> bool`

`e1 == e2` tests for physical equality of `e1` and `e2`. On integers and characters, it is the same as structural equality. On mutable structures, `e1 == e2` is true if and only if physical modification of `e1` also affects `e2`. On non-mutable structures, the behavior of `prefix ==` is implementation-dependent, except that `e1 == e2` implies `e1 = e2`.

`value prefix != : 'a -> 'a -> bool`

Negation of `prefix ==`.

4.5 exc: exceptions

`value raise : exn -> 'a`

Raise the given exception value.

A few general-purpose predefined exceptions.

`exception Out_of_memory`

Raised by the garbage collector, when there is insufficient memory to complete the computation.

`exception Invalid_argument of string`

Raised by some library functions, to signal that the given arguments do not make sense.

`exception Failure of string`

Raised by some library functions, to signal that they are undefined on the given arguments.

`exception Not_found`

Raised by search library functions, when the required object could not be found.

`exception Exit`

This exception is not raised by any library function. It is provided for use in your programs.

`value failwith : string -> 'a`

Raise exception `Failure` with the given string.

`value invalid_arg : string -> 'a`

Raise exception `Invalid_argument` with the given string.

4.6 `fchar`: character operations, without sanity checks

This module implements the same functions as the `char` module, but does not perform bound checks on the arguments of the functions. The functions are therefore faster than those in the `char` module, but calling these functions with incorrect parameters (that is, parameters that would cause the `Invalid_argument` exception to be raised by the corresponding functions in the `char` module) can crash the program.

4.7 `float`: operations on floating-point numbers

```
value int_of_float : float -> int
```

Truncate the given float to an integer value. The result is unspecified if it falls outside the range of representable integers.

```
value float_of_int : int -> float
```

Convert an integer to floating-point.

```
value minus : float -> float
value minus_float : float -> float
```

Unary negation.

```
value prefix + : float -> float -> float
value prefix +. : float -> float -> float
value add_float : float -> float -> float
```

Addition.

```
value prefix - : float -> float -> float
value prefix -. : float -> float -> float
value sub_float : float -> float -> float
```

Subtraction.

```
value prefix * : float -> float -> float
value prefix *. : float -> float -> float
value mult_float : float -> float -> float
```

Product.

```
value prefix / : float -> float -> float
value prefix /. : float -> float -> float
value div_float : float -> float -> float
```

Division. The result is unpredictable if the dividend is 0.0.

```
value eq_float : float -> float -> bool
value prefix =. : float -> float -> bool
```

Floating-point equality. Equivalent to generic equality, just faster.

```
value neq_float : float -> float -> bool
value prefix <>. : float -> float -> bool
```

Negation of `eq_float`.

```
value prefix < : float -> float -> bool
value prefix <. : float -> float -> bool
value lt_float : float -> float -> bool
value prefix > : float -> float -> bool
value prefix >. : float -> float -> bool
value gt_float : float -> float -> bool
value prefix <= : float -> float -> bool
value prefix <=. : float -> float -> bool
value le_float : float -> float -> bool
value prefix >= : float -> float -> bool
value prefix >=. : float -> float -> bool
value ge_float : float -> float -> bool
```

Usual comparisons between floating-point numbers.

```
value exp : float -> float
value log : float -> float
value sqrt : float -> float
value power : float -> float -> float
value sin : float -> float
value cos : float -> float
value tan : float -> float
value asin : float -> float
value acos : float -> float
value atan : float -> float
value atan2 : float -> float -> float
```

Usual transcendental functions on floating-point numbers.

```
value abs_float : float -> float
```

Return the absolute value of the argument.

```
value string_of_float : float -> string
```

Convert the given float to its decimal representation.

```
value float_of_string : string -> float
```

Convert the given string to a float, in decimal. The result is unspecified if the given string is not a valid representation of a float.

4.8 `fstring`: string operations, without sanity checks

This module implements the same functions as the `string` module, but does not perform bound checks on the arguments of the functions. The functions are therefore faster than those in the `string` module, but calling these functions with incorrect parameters (that is, parameters that would cause the `Invalid_argument` exception to be raised by the corresponding functions in the `string` module) can crash the program.

4.9 `fvect`: operations on vectors, without sanity checks

This module implements the same functions as the `vect` module, but does not perform bound checks on the arguments of the functions. The functions are therefore faster than those in the `vect` module, but calling these functions with incorrect parameters (that is, parameters that would cause the `Invalid_argument` exception to be raised by the corresponding functions in the `vect` module) can crash the program.

4.10 `int`: operations on integers

Integers are 31 bits wide. All operations are taken modulo 2^{31}. They do not fail on overflow.

```
exception Division_by_zero
value minus : int -> int
value minus_int : int -> int
```

Unary negation. You can write `-e` instead of `minus e`.

```
value succ : int -> int
```

`succ x` is `x+1`.

```
value pred : int -> int
```

`pred x` is `x-1`.

```
value prefix + : int -> int -> int
value add_int : int -> int -> int
```

Addition.

```
value prefix - : int -> int -> int
value sub_int : int -> int -> int
```

Subtraction.

```
value prefix * : int -> int -> int
value mult_int : int -> int -> int
```

Multiplication.

```
value prefix / : int -> int -> int
value div_int : int -> int -> int
value prefix quo : int -> int -> int
```

Integer division. Raise `Division_by_zero` if the second argument is 0. Give unpredictable results if either argument is negative.

```
value prefix mod : int -> int -> int
```

Remainder. Raise `Division_by_zero` if the second argument is 0. Give unpredictable results if either argument is negative.

```
value eq_int : int -> int -> bool
```

Integer equality. Equivalent to generic equality, just faster.

```
value neq_int : int -> int -> bool
```

Negation of `eq_int`.

```
value prefix < : int -> int -> bool
value lt_int : int -> int -> bool
value prefix > : int -> int -> bool
value gt_int : int -> int -> bool
value prefix <= : int -> int -> bool
value le_int : int -> int -> bool
value prefix >= : int -> int -> bool
value ge_int : int -> int -> bool
```

Usual comparisons between integers.

```
value min : int -> int -> int
```

Return the smaller of the arguments.

```
value max : int -> int -> int
```

Return the greater of the arguments.

```
value abs : int -> int
```

Return the absolute value of the argument.

Bitwise operations

```
value prefix land : int -> int -> int
```

Bitwise logical and.

```
value prefix lor : int -> int -> int
```

Bitwise logical or.

```
value prefix lxor : int -> int -> int
```

Bitwise logical exclusive or.

```
value lnot : int -> int
```

Bitwise complement

```
value prefix lsl : int -> int -> int
value lshift_left : int -> int -> int
```

`n lsl m`, or equivalently `lshift_left n m`, shifts `n` to the left by `m` bits.

```
value prefix lsr : int -> int -> int
```

`n lsr m` shifts `n` to the right by `m` bits. This is a logical shift: zeroes are inserted regardless of sign.

```
value prefix asr : int -> int -> int
value lshift_right : int -> int -> int
```

`n asr m`, or equivalently `lshift_right n m`, shifts `n` to the right by `m` bits. This is an arithmetic shift: the sign bit is replicated.

Conversion functions

```
value string_of_int : int -> string
```

Convert the given integer to its decimal representation.

```
value int_of_string : string -> int
```

Convert the given string to an integer, in decimal (by default) or in hexadecimal, octal or binary if the string begins with `0x`, `0o` or `0b`. Raise `Failure "int_of_string"` if the given string is not a valid representation of an integer.

4.11 `io`: buffered input and output

```
type in_channel
type out_channel
```

The abstract types of input channels and output channels.

```
exception End_of_file
```

Raised when an operation cannot complete, because the end of the file has been reached.

```
value stdin : in_channel
value std_in : in_channel
value stdout : out_channel
value std_out : out_channel
value stderr : out_channel
value std_err : out_channel
```

The standard input, standard output, and standard error output for the process. `std_in`, `std_out` and `std_err` are respectively synonymous with `stdin`, `stdout` and `stderr`.

```
value exit : int -> 'a
```

Flush all pending writes on `std_out` and `std_err`, and terminate the process, returning the given status code to the operating system (usually 0 to indicate no errors, and a small positive integer to indicate failure.) This function should be called at the end of all standalone programs that output results on `std_out` or `std_err`; otherwise, the program may appear to produce no output, or its output may be truncated.

Output functions on standard output

```
value print_char : char -> unit
```

Print the character on standard output.

```
value print_string : string -> unit
```

Print the string on standard output.

```
value print_int : int -> unit
```

Print the integer, in decimal, on standard output.

```
value print_float : float -> unit
```

Print the floating-point number, in decimal, on standard output.

```
value print_endline : string -> unit
```

Print the string, followed by a newline character, on standard output.

```
value print_newline : unit -> unit
```

Print a newline character on standard output, and flush standard output. This can be used to simulate line buffering of standard output.

Output functions on standard error

`value prerr_char : char -> unit`

Print the character on standard error.

`value prerr_string : string -> unit`

Print the string on standard error.

`value prerr_int : int -> unit`

Print the integer, in decimal, on standard error.

`value prerr_float : float -> unit`

Print the floating-point number, in decimal, on standard error.

`value prerr_endline : string -> unit`

Print the string, followed by a newline character on standard error and flush standard error.

Input functions on standard input

`value read_line : unit -> string`

Flush standard output, then read characters from standard input until a newline character is encountered. Return the string of all characters read, without the newline character at the end.

`value read_int : unit -> int`

Flush standard output, then read one line from standard input and convert it to an integer. Raise `Failure "int_of_string"` if the line read is not a valid representation of an integer.

`value read_float : unit -> float`

Flush standard output, then read one line from standard input and convert it to a floating-point number. The result is unspecified if the line read is not a valid representation of a floating-point number.

General output functions

`value open_out : string -> out_channel`

Open the named file for writing, and return a new output channel on that file, positionned at the beginning of the file. The file is truncated to zero length if it already exists. It is created if it does not already exists. Raise `sys__Sys_error` if the file could not be opened.

`value open_out_bin : string -> out_channel`

Same as `open_out`, but the file is opened in binary mode, so that no translation takes place during writes. On operating systems that do not distinguish between text mode and binary mode, this function behaves like `open_out`.

```
value open_out_gen : sys__open_flag list -> int -> string -> out_channel
```

`open_out_gen mode rights filename` opens the file named `filename` for writing, as above. The extra argument `mode` specify the opening mode (see `sys__open`). The extra argument `rights` specifies the file permissions, in case the file must be created (see `sys__open`). `open_out` and `open_out_bin` are special cases of this function.

```
value open_descriptor_out : int -> out_channel
```

`open_descriptor_out fd` returns a buffered output channel writing to the file descriptor `fd`. The file descriptor `fd` must have been previously opened for writing, else the behavior is undefined.

```
value flush : out_channel -> unit
```

Flush the buffer associated with the given output channel, performing all pending writes on that channel. Interactive programs must be careful about flushing `std_out` at the right times.

```
value output_char : out_channel -> char -> unit
```

Write the character on the given output channel.

```
value output_string : out_channel -> string -> unit
```

Write the string on the given output channel.

```
value output : out_channel -> string -> int -> int -> unit
```

`output chan buff ofs len` writes `len` characters from string `buff`, starting at offset `ofs`, to the output channel `chan`. Raise `Invalid_argument "output"` if `ofs` and `len` do not designate a valid substring of `buff`.

```
value output_byte : out_channel -> int -> unit
```

Write one 8-bit integer (as the single character with that code) on the given output channel. The given integer is taken modulo 256.

```
value output_binary_int : out_channel -> int -> unit
```

Write one integer in binary format on the given output channel. The only reliable way to read it back is through the `input_binary_int` function. The format is compatible across all machines for a given version of Caml Light.

```
value output_value : out_channel -> 'a -> unit
```

Write the representation of a structured value of any type to a channel. Circularities and sharing inside the value are detected and preserved. The object can be read back, by the function `input_value`. The format is compatible across all machines for a given version of Caml Light.

```
value seek_out : out_channel -> int -> unit
```

`seek_out chan pos` sets the current writing position to `pos` for channel `chan`. This works only for regular files. On files of other kinds (such as terminals, pipes and sockets,) the behavior is unspecified.

```
value pos_out : out_channel -> int
```

Return the current writing position for the given channel.

```
value out_channel_length : out_channel -> int
```

Return the total length (number of characters) of the given channel.

```
value close_out : out_channel -> unit
```

Close the given channel, flushing all buffered write operations. The behavior is unspecified if any of the above functions is called on a closed channel.

General input functions

```
value open_in : string -> in_channel
```

Open the named file for reading, and return a new input channel on that file, positionned at the beginning of the file. Raise `sys__Sys_error` if the file could not be opened.

```
value open_in_bin : string -> in_channel
```

Same as `open_in`, but the file is opened in binary mode, so that no translation takes place during reads. On operating systems that do not distinguish between text mode and binary mode, this function behaves like `open_in`.

```
value open_in_gen : sys__open_flag list -> int -> string -> in_channel
```

`open_in_gen mode rights filename` opens the file named `filename` for reading, as above. The extra arguments `mode` and `rights` specify the opening mode and file permissions (see `sys__open`). `open_in` and `open_in_bin` are special cases of this function.

```
value open_descriptor_in : int -> in_channel
```

`open_descriptor_in fd` returns a buffered input channel reading from the file descriptor `fd`. The file descriptor `fd` must have been previously opened for reading, else the behavior is undefined.

```
value input_char : in_channel -> char
```

Read one character from the given input channel. Raise `End_of_file` if there are no more characters to read.

```
value input_line : in_channel -> string
```

Read characters from the given input channel, until a newline character is encountered. Return the string of all characters read, without the newline character at the end. Raise `End_of_file` if the end of the file is reached before the line is complete.

`value input : in_channel -> string -> int -> int -> int`

`input chan buff ofs len` attempts to read `len` characters from channel `chan`, storing them in string `buff`, starting at character number `ofs`. It returns the actual number of characters read, between 0 and `len` (inclusive). A return value of 0 means that the end of file was reached. A return value between 0 and `len` exclusive means that no more characters were available at that time; `input` must be called again to read the remaining characters, if desired. Exception `Invalid_argument "input"` is raised if `ofs` and `len` do not designate a valid substring of `buff`.

`value really_input : in_channel -> string -> int -> int -> unit`

`input chan buff ofs len` reads `len` characters from channel `chan`, storing them in string `buff`, starting at character number `ofs`. Raise `End_of_file` if the end of file is reached before `len` characters have been read. Raise `Invalid_argument "really_input"` if `ofs` and `len` do not designate a valid substring of `buff`.

`value input_byte : in_channel -> int`

Same as `input_char`, but return the 8-bit integer representing the character. Raise `End_of_file` if an end of file was reached.

`value input_binary_int : in_channel -> int`

Read an integer encoded in binary format from the given input channel. See `output_binary_int`. Raise `End_of_file` if an end of file was reached while reading the integer.

`value input_value : in_channel -> 'a`

Read the representation of a structured value, as produced by `output_value`, and return the corresponding value. This is not type-safe. The type of the returned object is not `'a` properly speaking: the returned object has one unique type, which cannot be determined at compile-time. The programmer should explicitly give the expected type of the returned value, using the following syntax: `(input_value chan : type)`. The behavior is unspecified if the object in the file does not belong to the given type.

`value seek_in : in_channel -> int -> unit`

`seek_in chan pos` sets the current reading position to `pos` for channel `chan`.

`value pos_in : in_channel -> int`

Return the current reading position for the given channel.

`value in_channel_length : in_channel -> int`

Return the total length (number of characters) of the given channel. This works only for regular files. On files of other kinds, the result is meaningless.

`value close_in : in_channel -> unit`

Close the given channel. Anything can happen if any of the above functions is called on a closed channel.

4.12 `list`: operations on lists

`value list_length : 'a list -> int`

Return the length (number of elements) of the given list.

`value prefix @ : 'a list -> 'a list -> 'a list`

List concatenation.

`value hd : 'a list -> 'a`

Return the first element of the given list. Raise `Failure "hd"` if the list is empty.

`value tl : 'a list -> 'a list`

Return the given list without its first element. Raise `Failure "tl"` if the list is empty.

`value rev : 'a list -> 'a list`

List reversal.

`value map : ('a -> 'b) -> 'a list -> 'b list`

`map f [a1; ...; an]` applies function `f` to `a1, ..., an`, and builds the list `[f a1; ...; f an]` with the results returned by `f`.

`value do_list : ('a -> 'b) -> 'a list -> unit`

`do_list f [a1; ...; an]` applies function `f` in turn to `a1; ...; an`, discarding all the results. It is equivalent to `begin f a1; f a2; ...; f an; () end`.

`value it_list : ('a -> 'b -> 'a) -> 'a -> 'b list -> 'a`

`it_list f a [b1; ...; bn]` is `f (... (f (f a b1) b2) ...) bn`.

`value list_it : ('a -> 'b -> 'b) -> 'a list -> 'b -> 'b`

`list_it f [a1; ...; an] b` is `f a1 (f a2 (... (f an b) ...))`.

`value map2 : ('a -> 'b -> 'c) -> 'a list -> 'b list -> 'c list`

`map2 f [a1; ...; an] [b1; ...; bn]` is `[f a1 b1; ...; f an bn]`. Raise `Invalid_argument "map2"` if the two lists have different lengths.

`value do_list2 : ('a -> 'b -> 'c) -> 'a list -> 'b list -> unit`

`do_list2 f [a1; ...; an] [b1; ...; bn]` calls in turn `f a1 b1; ...; f an bn`, discarding the results. Raise `Invalid_argument "do_list2"` if the two lists have different lengths.

`value it_list2 : ('a -> 'b -> 'c -> 'a) -> 'a -> 'b list -> 'c list -> 'a`

`it_list2 f a [b1; ...; bn] [c1; ...; cn]` is `f (... (f (f a b1 c1) b2 c2) ...) bn cn`. Raise `Invalid_argument "it_list2"` if the two lists have different lengths.

```
value list_it2 : ('a -> 'b -> 'c -> 'c) -> 'a list -> 'b list -> 'c -> 'c
```

`list_it2 f [a1; ...; an] [b1; ...; bn] c` is `f a1 b1 (f a2 b2 (... (f an bn c) ...))`. Raise `Invalid_argument "list_it2"` if the two lists have different lengths.

```
value flat_map : ('a -> 'b list) -> 'a list -> 'b list
```

`flat_map f [l1; ...; ln]` is `(f l1) @ (f l2) @ ... @ (f ln)`.

```
value for_all : ('a -> bool) -> 'a list -> bool
```

`for_all p [a1; ...; an]` is `(p a1) & (p a2) & ... & (p an)`.

```
value exists : ('a -> bool) -> 'a list -> bool
```

`exists p [a1; ...; an]` is `(p a1) or (p a2) or ... or (p an)`.

```
value mem : 'a -> 'a list -> bool
```

`mem a l` is true if and only if `a` is structurally equal (see module `eq`) to an element of `l`.

```
value memq : 'a -> 'a list -> bool
```

`memq a l` is true if and only if `a` is physically equal (see module `eq`) to an element of `l`.

```
value except : 'a -> 'a list -> 'a list
```

`except a l` returns the list `l` where the first element structurally equal to `a` has been removed. The list `l` is returned unchanged if it does not contain `a`.

```
value exceptq : 'a -> 'a list -> 'a list
```

Same as `except`, with physical equality instead of structural equality.

```
value subtract : 'a list -> 'a list -> 'a list
```

`subtract l1 l2` returns the list `l1` where all elements structurally equal to one of the elements of `l2` have been removed.

```
value union : 'a list -> 'a list -> 'a list
```

`union l1 l2` appends before list `l2` all the elements of list `l1` that are not structurally equal to an element of `l2`.

```
value intersect : 'a list -> 'a list -> 'a list
```

`intersect l1 l2` returns the list of the elements of `l1` that are structurally equal to an element of `l2`.

```
value index : 'a -> 'a list -> int
```

`index a l` returns the position of the first element of list `l` that is structurally equal to `a`. The head of the list has position 0. Raise `Not_found` if `a` is not present in `l`.

```
value assoc : 'a -> ('a * 'b) list -> 'b
```

`assoc a l` returns the value associated with key `a` in the list of pairs `l`. That is, `assoc a [ ...; (a,b); ...] = b` if `(a,b)` is the leftmost binding of `a` in list `l`. Raise `Not_found` if there is no value associated with `a` in the list `l`.

```
value assq :  'a -> ('a * 'b) list -> 'b
```

Same as `assoc`, but use physical equality instead of structural equality to compare keys.

```
value mem_assoc : 'a -> ('a * 'b) list -> bool
```

Same as `assoc`, but simply return true if a binding exists, and false if no bindings exist for the given key.

4.13 pair: operations on pairs

```
value fst : 'a * 'b -> 'a
```

Return the first component of a pair.

```
value snd : 'a * 'b -> 'b
```

Return the second component of a pair.

```
value split : ('a * 'b) list -> 'a list * 'b list
```

Transform a list of pairs into a pair of lists: `split [(a1,b1); ...; (an,bn)]` is `([a1; ...; an], [b1; ...; bn])`

```
value combine : 'a list * 'b list -> ('a * 'b) list
```

Transform a pair of lists into a list of pairs: `combine ([a1; ...; an], [b1; ...; bn])` is `[(a1,b1); ...; (an,bn)]`. Raise `Invalid_argument "combine"` if the two lists have different lengths.

```
value map_combine : ('a * 'b -> 'c) -> 'a list * 'b list -> 'c list
```

`map_combine f ([a1; ...; an], [b1; ...; bn])` is `[f (a1, b1); ...; f (an, bn)]`. Raise `invalid_argument "map_combine"` if the two lists have different lengths.

```
value do_list_combine : ('a * 'b -> 'c) -> 'a list * 'b list -> unit
```

`do_list_combine f ([a1; ...; an], [b1; ...; bn])` calls in turn `f (a1, b1); ...; f (an, bn)`, discarding the results. Raise `Invalid_argument "do_list_combine"` if the two lists have different lengths.

4.14 `ref`: operations on references

```
type 'a ref = ref of mutable 'a
value prefix ! : 'a ref -> 'a
```

`!r` returns the current contents of reference `r`. Could be defined as `fun (ref x) -> x`.

```
value prefix := : 'a ref -> 'a -> unit
```

`r := a` stores the value of `a` in reference `r`.

```
value incr : int ref -> unit
```

Increment the integer contained in the given reference. Could be defined as `fun r -> r := succ !r`.

```
value decr : int ref -> unit
```

Decrement the integer contained in the given reference. Could be defined as `fun r -> r := pred !r`.

4.15 `stream`: operations on streams

```
type 'a stream
exception Parse_failure
```

Raised by parsers when none of the first component of the stream patterns is accepted

```
exception Parse_error
```

Raised by parsers when the first component of a stream pattern is accepted, but one of the following components is rejected

```
value stream_next : 'a stream -> 'a
```

`stream_next s` returns the first element of stream `s`, and removes it from the stream. Raise `Parse_failure` if the stream is empty.

```
value stream_from : (unit -> 'a) -> 'a stream
```

`stream_from f` returns the stream which fetches its terminals using the function `f`. This function could be defined as:

```
let rec stream_from f = [< 'f(); stream_from f >]
```

but is implemented more efficiently.

```
value stream_of_string : string -> char stream
```

`stream_of_string s` returns the stream of the characters in string `s`.

```
value stream_of_channel : in_channel -> char stream
```

`stream_of_channel ic` returns the stream of characters read on channel `ic`.

```
value do_stream : ('a -> 'b) -> 'a stream -> unit
```

do_stream f s scans the whole stream s, applying the function f in turn to each terminal encountered

```
value stream_check : ('a -> bool) -> 'a stream -> 'a
```

stream_check p returns the parser which returns the first terminal of the stream if the predicate p returns true on this terminal, and raises Parse_failure otherwise.

```
value end_of_stream : 'a stream -> unit
```

Return () iff the stream is empty, and raise Parse_failure otherwise.

```
value stream_get : 'a stream -> 'a * 'a stream
```

stream_get s return the first element of the stream s, and a stream containing the remaining elements of s. Raise Parse_failure if the stream is empty. The stream s is not modified. This function makes it possible to access a stream non-destructively.

4.16 string: string operations

```
value string_length : string -> int
```

Return the length (number of characters) of the given string.

```
value nth_char : string -> int -> char
```

nth_char s n returns character number n in string s. The first character is character number 0. The last character is character number string_length s - 1. Raise Invalid_argument "nth_char" if n is ouside the range 0 – (string_length s - 1).

```
value set_nth_char : string -> int -> char -> unit
```

set_nth_char s n c modifies string s in place, replacing the character number n by c. Raise Invalid_argument "set_nth_char" if n is ouside the range 0 to (string_length s - 1).

```
value prefix ^ : string -> string -> string
```

s1 ^ s2 returns a fresh string containing the concatenation of the strings s1 and s2.

```
value sub_string : string -> int -> int -> string
```

sub_string s start len returns a fresh string of length len, containing the characters number start to start + len - 1 of string s. Raise Invalid_argument "sub_string" if start and len do not designate a valid substring of s; that is, if start < 0, or len < 0, or start + len > string_length s.

```
value create_string : int -> string
```

create_string n returns a fresh string of length n. The string initially contains arbitrary characters.

```
value make_string : int -> char -> string
```

`make_string n c` returns a fresh string of length `n`, filled with the character `c`.

```
value fill_string : string -> int -> int -> char -> unit
```

`fill_string s start len c` modifies string `s` in place, replacing the characters number `start` to `start + len - 1` by `c`. Raise `Invalid_argument "fill_string"` if `start` and `len` do not designate a valid substring of `s`.

```
value blit_string : string -> int -> string -> int -> int -> unit
```

`blit_string s1 o1 s2 o2 len` copies `len` characters from string `s1`, starting at character number `o1`, to string `s2`, starting at character number `o2`. It works correctly even if `s1` and `s2` are the same string, and the source and destination chunks overlap. Raise `Invalid_argument "blit_string"` if `o1` and `len` do not designate a valid substring of `s1`, or if `o2` and `len` do not designate a valid substring of `s2`.

```
value replace_string : string -> string -> int -> unit
```

`replace_string dest src start` copies all characters from the string `src` into the string `dst`, starting at character number `start` in `dst`. Raise `Invalid_argument "replace_string"` if copying would overflow string `dest`.

```
value eq_string : string -> string -> bool
value neq_string : string -> string -> bool
value le_string : string -> string -> bool
value lt_string : string -> string -> bool
value ge_string : string -> string -> bool
value gt_string : string -> string -> bool
```

Comparison functions (lexicographic ordering) between strings.

```
value compare_strings : string -> string -> int
```

General comparison between strings. `compare_strings s1 s2` returns 0 if `s1` and `s2` are equal, or else -2 if `s1` is a prefix of `s2`, or 2 if `s2` is a prefix of `s1`, or else -1 if `s1` is lexicographically before `s2`, or 1 if `s2` is lexicographically before `s1`.

```
value string_for_read : string -> string
```

Return a copy of the argument, with special characters represented by escape sequences, following the lexical conventions of Caml Light.

4.17 `vect`: operations on vectors

```
value vect_length : 'a vect -> int
```

Return the length (number of elements) of the given vector.

```
value vect_item : 'a vect -> int -> 'a
```

vect_item v n returns the element number n of vector v. The first element has number 0. The last element has number vect_length v - 1. Raise Invalid_argument "vect_item" if n is outside the range 0 – (vect_length v - 1). You can also write v.(n) instead of vect_item v n.

```
value vect_assign : 'a vect -> int -> 'a -> unit
```

vect_assign v n x modifies vector v in place, replacing element number n with x. Raise Invalid_argument "vect_assign" if n is outside the range 0 – vect_length v - 1. You can also write v.(n) <- x instead of vect_assign v n x.

```
value make_vect : int -> 'a -> 'a vect
```

make_vect n x returns a fresh vector of length n, initialized with x. All the elements of this new vector are initially eq to x.

```
value make_matrix : int -> int -> 'a -> 'a vect vect
```

make_matrix dimx dimy e returns a two-dimensional array (a vector of vectors) with first dimension dimx and second dimension dimy. All the elements of this new matrix are initially eq to e. The element (x,y) of a matrix m is accessed with the notation m.(x).(y).

```
value concat_vect : 'a vect -> 'a vect -> 'a vect
```

concat_vect v1 v2 returns a fresh vector containing the concatenation of vectors v1 and v2.

```
value sub_vect : 'a vect -> int -> int -> 'a vect
```

sub_vect v start len returns a fresh vector of length len, containing the elements number start to start + len - 1 of vector v. Raise Invalid_argument "sub_vect" if start and len do not designate a valid subvector of v; that is, if start < 0, or len < 0, or start + len > vect_length v.

```
value copy_vect : 'a vect -> 'a vect
```

copy_vect v returns a copy of v, that is, a fresh vector containing the same elements as v.

```
value fill_vect : 'a vect -> int -> int -> 'a -> unit
```

fill_vect v ofs len x modifies vector v in place, storing x in elements number ofs to ofs + len - 1. Raise Invalid_argument "fill_vect" if ofs and len do not designate a valid subvector of v.

```
value blit_vect : 'a vect -> int -> 'a vect -> int -> int -> unit
```

blit_vect v1 o1 v2 o2 len copies len elements from vector v1, starting at element number o1, to vector v2, starting at element number o2. It works correctly even if v1 and v2 are the same vector, and the source and destination chunks overlap. Raise Invalid_argument "blit_vect" if o1 and len do not designate a valid subvector of v1, or if o2 and len do not designate a valid subvector of v2.

```
value list_of_vect : 'a vect -> 'a list
```

`list_of_vect v` returns the list of all the elements of `v`, that is: `[v.(0); v.(1); ...; v.(vect_length v - 1)]`.

```
value vect_of_list : 'a list -> 'a vect
```

`vect_of_list l` returns a fresh vector containing the elements of `l`.

```
value map_vect : ('a -> 'b) -> 'a vect -> 'b vect
```

`map_vect f v` applies function `f` to all the elements of `v`, and builds a vector with the results returned by `f`: `[| f v.(0); f v.(1); ...; f v.(vect_length v - 1) |]`.

```
value map_vect_list : ('a -> 'b) -> 'a vect -> 'b list
```

`map_vect_list f v` applies function `f` to all the elements of `v`, and builds a list with the results returned by `f`: `[ f v.(0); f v.(1); ...; f v.(vect_length v - 1) ]`.

```
value do_vect : ('a -> 'b) -> 'a vect -> unit
```

`do_vect f v` applies function `f` in turn to all the elements of `v`, discarding all the results: `f v.(0); f v.(1); ...; f v.(vect_length v - 1); ()`.

Chapter 5

The standard library

This chapter describes the functions provided by the Caml Light standard library. Just as the modules from the core library, the modules from the standard library are automatically linked with the user's object code files by the `camlc` command. Hence, the globals defined by these libraries can be used in standalone programs without having to add any `.zo` file on the command line for the linking phase. Similarly, in interactive use, these globals can be used in toplevel phrases without having to load any `.zo` file in memory.

Unlike the modules from the core library, the modules from the standard library are not automatically "opened" when a compilation starts, or when the toplevel system is launched. Hence it is necessary to use qualified identifiers to refer to the functions provided by these modules, or to add `#open` directives.

Conventions

For easy reference, the modules are listed below in alphabetical order of module names. For each module, the declarations from its interface file are printed one by one in typewriter font, followed by a short comment. All modules and the identifiers they export are indexed at the end of this report.

5.1 `arg`: parsing of command line arguments

This module provides a general mechanism for extracting options and arguments from the command line to the program.

Syntax of command lines. A keyword is a character string starting with a `-`. An option is a keyword alone or followed by an argument. There are 4 types of keywords: Unit, String, Int, and Float. Unit keywords do not take an argument. String, Int, and Float keywords take the following word on the command line as an argument. Arguments not preceded by a keyword are called anonymous arguments. Examples (`foo` is assumed to be the command name):

`foo -flag`	(a unit option)
`foo -int 1`	(an int option with argument `1`)
`foo -string foobar`	(a string option with argument `"foobar"`)
`foo -float 12.34`	(a float option with argument `12.34`)

```
foo 1 2 3            (three anonymous arguments: "1", "2", and "3")
foo 1 2 -flag 3 -string bar 4
                     (four anonymous arguments, a unit option, and
                      a string option with argument "bar")
```

```
type spec =
  String of (string -> unit)
| Int of (int -> unit)
| Unit of (unit -> unit)
| Float of (float -> unit)
```

The concrete type describing the behavior associated with a keyword.

```
value parse : (string * spec) list -> (string -> unit) -> unit
```

`parse speclist anonfun` parses the command line, calling the functions in `speclist` whenever appropriate, and `anonfun` on anonymous arguments. The functions are called in the same order as they appear on the command line. The strings in the `(string * spec) list` are keywords and must start with a `-`, else they are ignored. For the user to be able to specify anonymous arguments starting with a `-`, include for example `("--", String anonfun)` in `speclist`.

```
exception Bad of string
```

Functions in `speclist` or `anonfun` can raise `Bad message` to reject invalid arguments.

5.2 `baltree`: basic balanced binary trees

This module implements balanced ordered binary trees. All operations over binary trees are applicative (no side-effects). The `set` and `map` modules are based on this module. This modules gives a more direct access to the internals of the binary tree implementation than the `set` and `map` abstractions, but is more delicate to use and not as safe. For advanced users only.

```
type 'a t = Empty | Node of 'a t * 'a * 'a t * int
```

The type of trees containing elements of type `'a`. `Empty` is the empty tree (containing no elements).

```
type 'a contents = Nothing | Something of 'a
```

Used with the functions `modify` and `split`, to represent the presence or the absence of an element in a tree.

```
value add: ('a -> int) -> 'a -> 'a t -> 'a t
```

`add f x t` inserts the element `x` into the tree `t`. `f` is an ordering function: `f y` must return `0` if `x` and `y` are equal (or equivalent), a negative integer if `x` is smaller than `y`, and a positive integer if `x` is greater than `y`. The tree `t` is returned unchanged if it already contains an element equivalent to `x` (that is, an element `y` such that `f y` is `0`). The ordering `f` must be consistent with the orderings used to build `t` with `add`, `remove`, `modify` or `split` operations.

`value contains: ('a -> int) -> 'a t -> bool`

> `contains f t` checks whether `t` contains an element satisfying `f`, that is, an element `x` such that `f x` is `0`. `f` is an ordering function with the same constraints as for `add`. It can be coarser (identify more elements) than the orderings used to build `t`, but must be consistent with them.

`value find: ('a -> int) -> 'a t -> 'a`

> Same as `contains`, except that `find f t` returns the element `x` such that `f x` is `0`, or raises `Not_found` if none has been found.

`value remove: ('a -> int) -> 'a t -> 'a t`

> `remove f t` removes one element `x` of `t` such that `f x` is `0`. `f` is an ordering function with the same constraints as for `add`. `t` is returned unchanged if it does not contain any element satisfying `f`. If several elements of `t` satisfy `f`, only one is removed.

`value modify: ('a -> int) -> ('a contents -> 'a contents) -> 'a t -> 'a t`

> General insertion/modification/deletion function. `modify f g t` searchs `t` for an element `x` satisfying the ordering function `f`. If one is found, `g` is applied to `Something x`; if `g` returns `Nothing`, the element `x` is removed; if `g` returns `Something y`, the element `y` replaces `x` in the tree. (It is assumed that `x` and `y` are equivalent, in particular, that `f y` is `0`.) If the tree does not contain any `x` satisfying `f`, `g` is applied to `Nothing`; if it returns `Nothing`, the tree is returned unchanged; if it returns `Something x`, the element `x` is inserted in the tree. (It is assumed that `f x` is `0`.) The functions `add` and `remove` are special cases of `modify`, slightly more efficient.

`value split: ('a -> int) -> 'a t -> 'a t * 'a contents * 'a t`

> `split f t` returns a triple `(less, elt, greater)` where `less` is a tree containing all elements `x` of `t` such that `f x` is negative, `greater` is a tree containing all elements `x` of `t` such that `f x` is positive, and `elt` is `Something x` if `t` contains an element `x` such that `f x` is `0`, and `Nothing` otherwise.

`value compare: ('a -> 'a -> int) -> 'a t -> 'a t -> int`

> Compare two trees. The first argument `f` is a comparison function over the tree elements: `f e1 e2` is zero if the elements `e1` and `e2` are equal, negative if `e1` is smaller than `e2`, and positive if `e1` is greater than `e2`. `compare f t1 t2` compares the fringes of `t1` and `t2` by lexicographic extension of `f`.

5.3 `filename`: operations on file names

`value current_dir_name : string`

> The conventional name for the current directory (e.g. . in Unix).

`value concat : string -> string -> string`

> `concat dir file` returns a file name that designates file `file` in directory `dir`.

```
value is_absolute : string -> bool
```

Return `true` if the file name is absolute or starts with an explicit reference to the current directory (`./` or `../` in Unix), and `false` if it is relative to the current directory.

```
value check_suffix : string -> string -> bool
```

`check_suffix name suff` returns `true` if the filename `name` ends with the suffix `suff`.

```
value chop_suffix : string -> string -> string
```

`chop_suffix name suff` removes the suffix `suff` from the filename `name`. The behavior is undefined if `name` does not end up with the suffix `suff`.

```
value basename : string -> string
value dirname : string -> string
```

Split a file name into directory name / base file name. `concat (dirname name) (basename name)` returns a file name which is equivalent to `name`. Moreover, after setting the current directory to `dirname name` (with `sys__chdir`), references to `basename name` (which is a relative file name) designate the same file as `name` before the call to `chdir`.

5.4 genlex: a generic lexical analyzer

This module implements a simple "standard" lexical analyzer, presented as a function from character streams to token streams. It implements roughly the lexical conventions of Caml, but is parameterized by the set of keywords of your language.

```
type token =
    Kwd of string
  | Ident of string
  | Int of int
  | Float of float
  | String of string
  | Char of char
```

The type of tokens. The lexical classes are: `Int` and `Float` for integer and floating-point numbers; `String` for string literals, enclosed in double quotes; `Char` for character literals, enclosed in backquotes; `Ident` for identifiers (either sequences of letters, digits, underscores and quotes, or sequences of "operator characters" such as `+`, `*`, etc); and `Kwd` for keywords (either identifiers or single "special characters" such as `(`, `}`, etc).

```
value make_lexer: string list -> (char stream -> token stream)
```

Construct the lexer function. The first argument is the list of keywords. An identifier `s` is returned as `Kwd s` if `s` belongs to this list, and as `Ident s` otherwise. A special character `s` is returned as `Kwd s` if `s` belongs to this list, and cause a lexical error (exception `Parse_error`) otherwise. Blanks and newlines are skipped. Comments delimited by `(*` and `*)` are skipped as well, and can be nested. Example: a lexer suitable for a desk calculator is obtained by

```
let lexer = make_lexer ["+";"-";"*";"/";"let";"="; "("; ")"]
```

The associated parser would be a function from `token stream` to, for instance, `int`, and would have rules such as:

```
let parse_expr = function
       [< 'Int n >] -> n
     | [< 'Kwd "("; parse_expr n; 'Kwd ")" >] -> n
     | [< parse_expr n1; (parse_end n1) n2 >] -> n2
and parse_end n1 = function
       [< 'Kwd "+"; parse_expr n2 >] -> n1+n2
     | ...
```

5.5 `hashtbl`: hash tables and hash functions

Hash tables are hashed association tables, with in-place modification.

`type ('a, 'b) t mutable`

The type of hash tables from type `'a` to type `'b`.

`value new : int -> ('a,'b) t`

`new n` creates a new, empty hash table, with initial size `n`. The table grows as needed, so `n` is just an initial guess. Better results are said to be achieved when `n` is a prime number.

`value clear : ('a, 'b) t -> unit`

Empty a hash table.

`value add : ('a, 'b) t -> 'a -> 'b -> unit`

`add tbl x y` adds a binding of `x` to `y` in table `tbl`. Previous bindings for `x` are not removed, but simply hidden. That is, after performing `remove tbl x`, the previous binding for `x`, if any, is restored. (This is the semantics of association lists.)

`value find : ('a, 'b) t -> 'a -> 'b`

`find tbl x` returns the current binding of `x` in `tbl`, or raises `Not_found` if no such binding exists.

`value find_all : ('a, 'b) t -> 'a -> 'b list`

`find_all tbl x` returns the list of all data associated with `x` in `tbl`. The current binding is returned first, then the previous bindings, in reverse order of introduction in the table.

`value remove : ('a, 'b) t -> 'a -> unit`

`remove tbl x` removes the current binding of `x` in `tbl`, restoring the previous binding if it exists. It does nothing if `x` is not bound in `tbl`.

`value do_table : ('a -> 'b -> 'c) -> ('a, 'b) t -> unit`

`do_table f tbl` applies `f` to all bindings in table `tbl`, discarding all the results. `f` receives the key as first argument, and the associated value as second argument. The order in which the bindings are passed to `f` is unpredictable. The same binding can be presented several times to `f`.

The polymorphic hash primitive

```
value hash : 'a -> int
```

`hash x` associates a positive integer to any value of any type. It is guaranteed that if `x = y`, then `hash x = hash y`. Moreover, `hash` always terminates, even on cyclic structures.

```
value hash_param : int -> int -> 'a -> int
```

`hash_param n m x` computes a hash value for `x`, with the same properties as for `hash`. The two extra parameters `n` and `m` give more precise control over hashing. Hashing performs a depth-first, right-to-left traversal of the structure `x`, stopping after `n` meaningful nodes were encountered, or `m` nodes, meaningful or not, were encountered. Meaningful nodes are: integers; floating-point numbers; strings; characters; booleans; and constant constructors. Larger values of `m` and `n` means that more nodes are taken into account to compute the final hash value, and therefore collisions are less likely to happen. However, hashing takes longer. The parameters `m` and `n` govern the tradeoff between accuracy and speed.

5.6 `lexing`: the run-time library for lexers generated by camllex

Lexer buffers

```
type lexbuf =
  { refill_buff : lexbuf -> unit;
    lex_buffer : string;
    mutable lex_abs_pos : int;
    mutable lex_start_pos : int;
    mutable lex_curr_pos : int;
    mutable lex_last_pos : int;
    mutable lex_last_action : lexbuf -> obj }
```

The type of lexer buffers. A lexer buffer is the argument passed to the scanning functions defined by the generated scanners. The lexer buffer holds the current state of the scanner, plus a function to refill the buffer from the input.

```
value create_lexer_channel : in_channel -> lexbuf
```

Create a lexer buffer on the given input channel. `create_lexer_channel inchan` returns a lexer buffer which reads from the input channel `inchan`, at the current reading position.

```
value create_lexer_string : string -> lexbuf
```

Create a lexer buffer which reads from the given string. Reading starts from the first character in the string. An end-of-input condition is generated when the end of the string is reached.

```
value create_lexer : (string -> int -> int) -> lexbuf
```

Create a lexer buffer with the given function as its reading method. When the scanner needs more characters, it will call the given function, giving it a character string `s` and a character count `n`. The function should put `n` characters or less in `s`, starting at character number 0, and return the number of characters provided. A return value of 0 means end of input.

Functions for lexer semantic actions

The following functions can be called from the semantic actions of lexer definitions (the ML code enclosed in braces that computes the value returned by lexing functions). They give access to the character string matched by the regular expression associated with the semantic action. These functions must be applied to the argument `lexbuf`, which, in the code generated by camllex, is bound to the lexer buffer passed to the parsing function.

```
value get_lexeme : lexbuf -> string
```

`get_lexeme lexbuf` returns the string matched by the regular expression.

```
value get_lexeme_char : lexbuf -> int -> char
```

`get_lexeme_char lexbuf i` returns character number `i` in the matched string.

```
value get_lexeme_start : lexbuf -> int
```

`get_lexeme_start lexbuf` returns the position in the input stream of the first character of the matched string. The first character of the stream has position 0.

```
value get_lexeme_end : lexbuf -> int
```

`get_lexeme_start lexbuf` returns the position in the input stream of the character following the last character of the matched string. The first character of the stream has position 0.

5.7 `map`: association tables over ordered types

This module implements applicative association tables, also known as finite maps or dictionaries, given a total ordering function over the keys. All operations over maps are purely applicative (no side-effects). The implementation uses balanced binary trees, and therefore searching and insertion take time logarithmic in the size of the map.

```
type ('a, 'b) t
```

The type of maps from type `'a` to type `'b`.

```
value empty: ('a -> 'a -> int) -> ('a, 'b) t
```

The empty map. The argument is a total ordering function over the set elements. This is a two-argument function `f` such that `f e1 e2` is zero if the elements `e1` and `e2` are equal, `f e1 e2` is strictly negative if `e1` is smaller than `e2`, and `f e1 e2` is strictly positive if `e1` is greater than `e2`. Examples: a suitable ordering function for type `int` is `prefix -`. For type `string`, you could use `compare_strings`.

```
value add: 'a -> 'b -> ('a, 'b) t -> ('a, 'b) t
```

`add x y m` returns a map containing the same bindings as `m`, plus a binding of `x` to `y`. Previous bindings for `x` in `m` are not removed, but simply hidden: they reappear after performing a `remove` operation. (This is the semantics of association lists.)

```
value find:'a -> ('a, 'b) t -> 'b
```

`find x m` returns the current binding of `x` in `m`, or raises `Not_found` if no such binding exists.

```
value remove: 'a -> ('a, 'b) t -> ('a, 'b) t
```

`remove x m` returns a map containing the same bindings as `m` except the current binding for `x`. The previous binding for `x` is restored if it exists. `m` is returned unchanged if `x` is not bound in `m`.

```
value iter: ('a -> 'b -> 'c) -> ('a, 'b) t -> unit
```

`iter f m` applies `f` to all bindings in map `m`, discarding the results. `f` receives the key as first argument, and the associated value as second argument. The order in which the bindings are passed to `f` is unspecified. Only current bindings are presented to `f`: bindings hidden by more recent bindings are not passed to `f`.

5.8 `parsing`: the run-time library for parsers generated by mly-acc

The following functions can be called by user code.

```
value symbol_start : unit -> int
value symbol_end : unit -> int
```

`symbol_start` and `symbol_end` are to be called in the action part of a grammar rule only. They return the position of the string that matches the left-hand side of the rule: `symbol_start()` returns the position of the first character; `symbol_end()` returns the position of the last character, plus one. The first character in a file is at position 0.

```
value rhs_start: int -> int
value rhs_end: int -> int
```

Same as `symbol_start` and `symbol_end` above, but return then position of the string matching the `n`th item on the right-hand side of the rule, where `n` is the integer parameter to `lhs_start` and `lhs_end`. `n` is 1 for the leftmost item.

```
value clear_parser : unit -> unit
```

Empty the parser stack. Call it just after a parsing function has returned, to remove all pointers from the parser stack to structures that were built by semantic actions during parsing. This is optional, but lowers the memory requirements of the programs.

5.9 `printexc`: a catch-all exception handler

```
value f: ('a -> 'b) -> 'a -> 'b
```

`f fn x` applies `fn` to `x` and returns the result. If the evaluation of `fn x` raises any exception, the name of the exception is printed on standard error output, and the programs aborts with exit code 2. Typical use is `f main ()`, where `main`, with type `unit->unit`, is the entry point of a standalone program, to catch and print stray exceptions. For `f` to work properly, the program must be linked with the `-g` option.

5.10 printf: formatting printing functions

`value fprintf: out_channel -> string -> 'a`

`fprintf outchan format arg1 ... argN` formats the arguments `arg1` to `argN` according to the format string `format`, and outputs the resulting string on the channel `outchan`. The format is a character string which contains two types of objects: plain characters, which are simply copied to the output channel, and conversion specifications, each of which causes conversion and printing of one argument. Conversion specifications consist in the `%` character, followed by optional flags and field widths, followed by one conversion character. The conversion characters and their meanings are:

`d` or `i`: convert an integer argument to signed decimal

`u`: convert an integer argument to unsigned decimal

`x`: convert an integer argument to unsigned hexadecimal, using lowercase letters.

`X`: convert an integer argument to unsigned hexadecimal, using uppercase letters.

`s`: insert a string argument

`c`: insert a character argument

`f`: convert a floating-point argument to decimal notation, in the style `dddd.ddd`

`e` or `E`: convert a floating-point argument to decimal notation, in the style `d.ddd e+-dd` (mantissa and exponent)

`g` or `G`: convert a floating-point argument to decimal notation, in style `f` or `e`, `E` (whichever is more compact)

`b`: convert a boolean argument to the string `true` or `false`

Refer to the C library `printf` function for the meaning of flags and field width specifiers. The exception `Invalid_argument` is raised if the types of the provided arguments do not match the format. The exception is also raised if too many arguments are provided. If too few arguments are provided, printing stops just before converting the first missing argument.

`value printf: string -> 'a`

Same as `fprintf`, but output on `std_out`.

`value fprint: out_channel -> string -> unit`

Print the given string on the given output channel, without any formatting.

`value print: string -> unit`

Print the given string on `std_out`, without any formatting. This is the same function as `print_string` of module `io`.

5.11 queue: queues

This module implements queues (FIFOs), with in-place modification.

`type 'a t mutable`

The type of queues containing elements of type `'a`.

`exception Empty`

Raised when `take` is applied to an empty queue.

`value new: unit -> 'a t`

Return a new queue, initially empty.

`value add: 'a -> 'a t -> unit`

`add x q` adds the element `x` at the end of the queue `q`.

`value take: 'a t -> 'a`

`take q` removes and returns the first element in queue `q`, or raises `Empty` if the queue is empty.

`value peek: 'a t -> 'a`

`peek q` returns the first element in queue `q`, without removing it from the queue, or raises `Empty` if the queue is empty.

`value clear : 'a t -> unit`

Discard all elements from a queue.

`value length: 'a t -> int`

Return the number of elements in a queue.

`value iter: ('a -> 'b) -> 'a t -> unit`

`iter f q` applies `f` in turn to all elements of `q`, from the least recently entered to the most recently entered. The queue itself is unchanged.

5.12 `random`: pseudo-random number generator

`value init : int -> unit`

Initialize the generator, using the argument as a seed. The same seed will always yield the same sequence of numbers.

`value int : int -> int`

`random__int bound` returns a random number between 0 (inclusive) and `bound` (exclusive). `bound` must be smaller than 2^{30}.

`value float : float -> float`

`random__float` returns a random number between 0 (inclusive) and `bound` (exclusive).

5.13 `set`: sets over ordered types

This module implements the set data structure, given a total ordering function over the set elements. All operations over sets are purely applicative (no side-effects). The implementation uses balanced binary trees, and is therefore reasonably efficient: insertion and membership take time logarithmic in the size of the set, for instance.

```
type 'a t
```

The type of sets containing elements of type `'a`.

```
value empty: ('a -> 'a -> int) -> 'a t
```

The empty set. The argument is a total ordering function over the set elements. This is a two-argument function `f` such that `f e1 e2` is zero if the elements `e1` and `e2` are equal, `f e1 e2` is strictly negative if `e1` is smaller than `e2`, and `f e1 e2` is strictly positive if `e1` is greater than `e2`. Examples: a suitable ordering function for type `int` is `prefix -`. For type `string`, you could use `compare_strings`.

```
value is_empty: 'a t -> bool
```

Test whether a set is empty or not.

```
value mem: 'a -> 'a t -> bool
```

`mem x s` tests whether `x` belongs to the set `s`.

```
value add: 'a -> 'a t -> 'a t
```

`add x s` returns a set containing all elements of `s`, plus `x`. If `x` was already in `s`, `s` is returned unchanged.

```
value remove: 'a -> 'a t -> 'a t
```

`remove x s` returns a set containing all elements of `s`, except `x`. If `x` was not in `s`, `s` is returned unchanged.

```
value union: 'a t -> 'a t -> 'a t
value inter: 'a t -> 'a t -> 'a t
value diff: 'a t -> 'a t -> 'a t
```

Union, intersection and set difference.

```
value equal: 'a t -> 'a t -> bool
```

`equal s1 s2` tests whether the sets `s1` and `s2` are equal, that is, contain the same elements.

```
value compare: 'a t -> 'a t -> int
```

Total ordering between sets. Can be used as the ordering function for doing sets of sets.

```
value elements: 'a t -> 'a list
```

Return the list of all elements of the given set. The elements appear in the list in some non-specified order.

`value iter: ('a -> 'b) -> 'a t -> unit`

`iter f s` applies `f` in turn to all elements of `s`, and discards the results. The elements of `s` are presented to `f` in a non-specified order.

`value fold: ('a -> 'b -> 'b) -> 'a t -> 'b -> 'b`

`fold f s a` computes `(f xN ... (f x2 (f x1 a))...)`, where `x1 ... xN` are the elements of `s`. The order in which elements of `s` are presented to `f` is not specified.

5.14 `sort`: sorting and merging lists

`value sort : ('a -> 'a -> bool) -> 'a list -> 'a list`

Sort a list in increasing order according to an ordering predicate. The predicate should return `true` if its first argument is less than or equal to its second argument.

`value merge : ('a -> 'a -> bool) -> 'a list -> 'a list -> 'a list`

Merge two lists according to the given predicate. Assuming the two argument lists are sorted according to the predicate, `merge` returns a sorted list containing the elements from the two lists. The behavior is undefined if the two argument lists were not sorted.

5.15 `stack`: stacks

This module implements stacks (LIFOs), with in-place modification.

`type 'a t mutable`

The type of stacks containing elements of type `'a`.

`exception Empty`

Raised when `pop` is applied to an empty stack.

`value new: unit -> 'a t`

Return a new stack, initially empty.

`value push: 'a -> 'a t -> unit`

`push x s` adds the element `x` at the top of stack `s`.

`value pop: 'a t -> 'a`

`pop s` removes and returns the topmost element in stack `s`, or raises `Empty` if the stack is empty.

`value clear : 'a t -> unit`

Discard all elements from a stack.

`value length: 'a t -> int`

Return the number of elements in a stack.

`value iter: ('a -> 'b) -> 'a t -> unit`

`iter f s` applies `f` in turn to all elements of `s`, from the element at the top of the stack to the element at the bottom of the stack. The stack itself is unchanged.

5.16 sys: system interface

This module provides a simple interface to the operating system.

```
exception Sys_error of string
```

Raised by some functions in the `sys` and `io` modules, when the underlying system calls fail. The argument to `Sys_error` is a string describing the error. The texts of the error messages are implementation-dependent, and should not be relied upon to catch specific system errors.

```
value command_line : string vect
```

The command line arguments given to the process. The first element is the command name used to invoke the program.

```
value interactive: bool
```

True if we're running under the toplevel system. False if we're running as a standalone program.

```
type file_perm == int
value s_irusr : file_perm
value s_iwusr : file_perm
value s_ixusr : file_perm
value s_irgrp : file_perm
value s_iwgrp : file_perm
value s_ixgrp : file_perm
value s_iroth : file_perm
value s_iwoth : file_perm
value s_ixoth : file_perm
value s_isuid : file_perm
value s_isgid : file_perm
value s_irall : file_perm
value s_iwall : file_perm
value s_ixall : file_perm
```

Access permissions for files. `r` is reading permission, `w` is writing permission, `x` is execution permission. `usr` means permissions for the user owning the file, `grp` for the group owning the file, `oth` for others. `isuid` and `isgid` are for set-user-id and set-group-id files, respectively. The remaining are combinations of the permissions above.

```
type open_flag =
    O_RDONLY                    (* open read-only *)
  | O_WRONLY                    (* open write-only *)
  | O_RDWR                      (* open for reading and writing *)
  | O_APPEND                    (* open for appending *)
  | O_CREAT                     (* create the file if nonexistent *)
  | O_TRUNC                     (* truncate the file to 0 if it exists *)
  | O_EXCL                      (* fails if the file exists *)
  | O_BINARY                    (* open in binary mode *)
  | O_TEXT                      (* open in text mode *)
```

The commands for `open`.

`value exit : int -> 'a`

Terminate the program and return the given status code to the operating system. In contrast with the function `exit` from module `io`, this `exit` function does not flush the standard output and standard error channels.

`value open : string -> open_flag list -> file_perm -> int`

Open a file. The second argument is the opening mode. The third argument is the permissions to use if the file must be created. The result is a file descriptor opened on the file.

`value close : int -> unit`

Close a file descriptor.

`value remove : string -> unit`

Remove the given file name from the file system.

`value rename : string -> string -> unit`

Rename a file. The first argument is the old name and the second is the new name.

`value getenv : string -> string`

Return the value associated to a variable in the process environment. Raise `Not_found` if the variable is unbound.

`value chdir : string -> unit`

Change the current working directory of the process. Note that there is no easy way of getting the current working directory from the operating system.

`exception Break`

Exception `Break` is raised on user interrupt if `catch_break` is on.

`value catch_break : bool -> unit`

`catch_break` governs whether user interrupt terminates the program or raises the `Break` exception. Call `catch_break true` to enable raising `Break`, and `catch_break false` to let the system terminate the program on user interrupt.

Chapter 6

The graphics library

This chapter describes the portable graphics primitives that come standard in the implementation of Caml Light on micro-computers.

Unix: On Unix workstations running the X11 windows system, an implementation of the graphics primitives is available in the directory `contrib/libgraph` in the distribution. See the file `README` in this directory for information on building and using `camlgraph`, a toplevel system that includes the graphics primitives, and linking standalone programs with the library. Drawing takes place in a separate window that is created when `open_graph` is called.

Mac: The graphics primitive are available from the standalone application that runs the toplevel system. They are not available from programs compiled by `camlc` and run under the MPW shell. Drawing takes place in a separate window, that can be made visible with the "Show graphics window" menu entry.

PC: The graphics primitive are available both at toplevel and in separately compiled programs. The `open_graph` primitive switches the whole screen to graphics mode. Under the toplevel system, the user's phrases and the system responses are printed on the graphics screen too, possibly overwriting parts of the drawings, and causing scrolling when the bottom of the screen is reached.

The screen coordinates are interpreted as shown in the figure below. Notice that the coordinate system used is the same as in mathematics: y increases from the bottom of the screen to the top of the screen, and angles are measured counterclockwise (in degrees). Drawing is clipped to the screen.

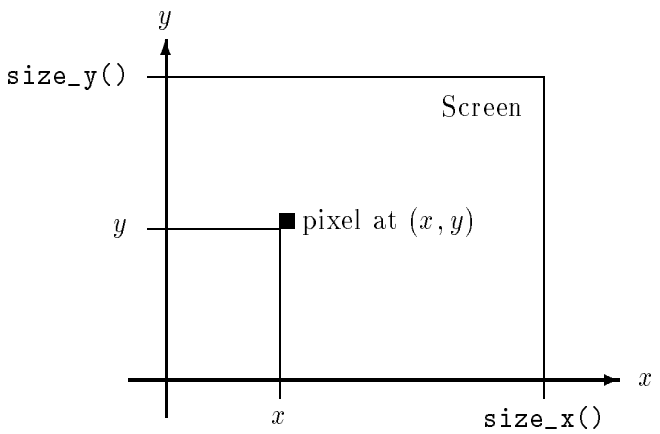

Here are the graphics mode specifications supported by `open_graph` on the various implementations of this library.

Unix: The argument to `open_graph` has the format "*display-name geometry*", where *display-name* is the name of the X-windows display to connect to, and *geometry* is a standard X-windows geometry specification. The two components are separated by a space. Either can be omitted, or both. Examples:

`open_graph "foo:0"`
: connects to the display `foo:0` and creates a window with the default geometry

`open_graph "foo:0 300x100+50-0"`
: connects to the display `foo:0` and creates a window 300 pixels wide by 100 pixels tall, at location $(50, 0)$

`open_graph " 300x100+50-0"`
: connects to the default display and creates a window 300 pixels wide by 100 pixels tall, at location $(50, 0)$

`open_graph ""`
: connects to the default display and creates a window with the default geometry.

Mac: The argument to `open_graph` is ignored.

86: The 86 PC version supports the CGA, EGA and VGA graphics cards. The argument to `open_graph` is interpreted as follows:

Argument to `open_graph`	Graphic mode
`"cga"`	320x200, 4 colors
`"ega_low"`	640x200, 16 colors
`"ega_high"`	640x350, 16 colors
`"vga_low"`	640x200, 16 colors
`"vga_med"`	640x350, 16 colors
`"vga_high"`	640x480, 16 colors
anything else	highest possible resolution

386: The 386 PC version supports VGA graphics cards, and a number of SuperVGA cards, provided the correct driver has been installed. (See the installation instructions.) Only 256 color modes are supported. This means a resolution of 320x200 on standard VGA, and from 640x480 to 1024x768 on SuperVGA.

The graphics primitive do not work if Caml Light runs under Windows 3 or any DPMI-compliant memory manager.

The argument to `open_graph` is interpreted as follows:

Argument to `open_graph`	Graphic mode
`"highest"`	highest possible resolution
`"noninterlaced"`	highest possible resolution, non-interlaced mode
`"320x200"`	standard 320x200 VGA mode
"*w*`x`*h*"	at least w by h pixels, if possible
anything else	default resolution, as set in the `GO32` variable (see the installation instructions).

6.1 graphics: machine-independent graphics primitives

```
exception Graphic_failure of string
```

Raised by the functions below when they encounter an error.

Initializations

```
value open_graph: string -> unit
```

Show the graphics window or switch the screen to graphic mode. The graphics window is cleared. The string argument is used to pass optional information on the desired graphics mode, the graphics window size, and so on. Its interpretation is implementation-dependent. If the empty string is given, a sensible default is selected.

```
value close_graph: unit -> unit
```

Delete the graphics window or switch the screen back to text mode.

```
value clear_graph : unit -> unit
```

Erase the graphics window.

```
value size_x : unit -> int
value size_y : unit -> int
```

Return the size of the graphics window. Coordinates of the screen pixels range over `0 .. size_x()-1` and `0 .. size_y()-1`. Drawings outside of this rectangle are clipped, without causing an error. The origin (0,0) is at the lower left corner.

Colors

```
type color == int
```

A color is specified by its R, G, B components. Each component is in the range `0..255`. The three components are packed in an `int: 0xRRGGBB`, where `RR` are the two hexadecimal digits for the red component, `GG` for the green component, `BB` for the blue component.

```
value rgb: int -> int -> int -> int
```

`rgb r g b` returns the integer encoding the color with red component `r`, green component `g`, and blue component `b`. `r`, `g` and `b` are in the range `0..255`.

```
value set_color : color -> unit
```

Set the current drawing color.

```
value black : color
value white : color
value red : color
value green : color
value blue : color
value yellow : color
value cyan : color
value magenta : color
```

Some predefined colors.

```
value background: color
value foreground: color
```

Default background and foreground colors (usually, either black foreground on a white background or white foreground on a black background). `clear_graph` fills the screen with the `background` color. The initial drawing color is `foreground`.

Point and line drawing

```
value plot : int -> int -> unit
```

Plot the given point with the current drawing color.

```
value point_color : int -> int -> color
```

Return the color of the given point.

```
value moveto : int -> int -> unit
```

Position the current point.

```
value current_point : unit -> int * int
```

Return the position of the current point.

```
value lineto : int -> int -> unit
```

Draw a line with endpoints the current point and the given point, and move the current point to the given point.

```
value draw_arc : int -> int -> int -> int -> int -> int -> unit
```

`draw_arc x y rx ry a1 a2` draws an elliptical arc with center `x,y`, horizontal radius `rx`, vertical radius `ry`, from angle `a1` to angle `a2` (in degrees). The current point is unchanged.

```
value draw_ellipse : int -> int -> int -> int -> unit
```

`draw_ellipse x y rx ry` draws an ellipse with center `x,y`, horizontal radius `rx` and vertical radius `ry`. The current point is unchanged.

```
value draw_circle : int -> int -> int -> unit
```

`draw_circle x y r` draws a circle with center `x,y` and radius `r`. The current point is unchanged.

```
value set_line_width : int -> unit
```

Set the width of points and lines drawn with the functions above.

Text drawing

```
value draw_char : char -> unit
value draw_string : string -> unit
```

Draw a character or a character string with lower left corner at current position. After drawing, the current position is set to the lower right corner of the text drawn.

```
value set_font : string -> unit
value set_text_size : int -> unit
```

Set the font and character size used for drawing text. The interpretation of the arguments to `set_font` and `set_text_size` is implementation-dependent.

```
value text_size : string -> int * int
```

Return the dimensions of the given text, if it were drawn with the current font and size.

Filling

```
value fill_rect : int -> int -> int -> int -> unit
```

`fill_rect x y w h` fills the rectangle with lower left corner at `x,y`, width `w` and heigth `h`, with the current color.

```
value fill_poly : (int * int) vect -> unit
```

Fill the given polygon with the current color. The array contains the coordinates of the vertices of the polygon.

```
value fill_arc : int -> int -> int -> int -> int -> int -> unit
```

Fill an elliptical pie slice with the current color. The parameters are the same as for `draw_arc`.

```
value fill_ellipse : int -> int -> int -> int -> unit
```

Fill an ellipse with the current color. The parameters are the same as for `draw_ellipse`.

```
value fill_circle : int -> int -> int -> unit
```

Fill a circle with the current color. The parameters are the same as for `draw_circle`.

Images

```
type image
```

The abstract type for images, in internal representation. Externally, images are represented as matrices of colors.

```
value transp : color
```

In matrices of colors, this color represent a "transparent" point: when drawing the corresponding image, all pixels on the screen corresponding to a transparent pixel in the image will not be modified, while other points will be set to the color of the corresponding point in the image. This allows superimposing an image over an existing background.

```
value make_image : color vect vect -> image
```

Convert the given color matrix to an image. Each sub-array represents one horizontal line. All sub-arrays must have the same length; otherwise, exception `Graphic_failure` is raised.

```
value dump_image : image -> color vect vect
```

Convert an image to a color matrix.

```
value draw_image : image -> int -> int -> unit
```

Draw the given image with lower left corner at the given point.

```
value get_image : int -> int -> int -> int -> image
```

Capture the contents of a rectangle on the screen as an image. The parameters are the same as for `fill_rect`.

```
value create_image : int -> int -> image
```

`create_image w h` returns a new image `w` pixels wide and `h` pixels tall, to be used in conjunction with `blit_image`. The initial image contents are random.

```
value blit_image : image -> int -> int -> unit
```

`blit_image img x y` copies screen pixels into the image `img`, modifying `img` in-place. The pixels copied are those inside the rectangle with lower left corner at `x,y`, and width and height equal to those of the image.

Mouse and keyboard events

```
type status =
  { mouse_x : int;              (* X coordinate of the mouse *)
    mouse_y : int;              (* Y coordinate of the mouse *)
    button : bool;              (* true if a mouse button is pressed *)
    keypressed : bool;          (* true if a key has been pressed *)
    key : char }                (* the character for the key pressed *)
```

To report events.

```
type event =
    Button_down                 (* A mouse button is pressed *)
  | Button_up                   (* A mouse button is released *)
  | Key_pressed                 (* A key is pressed *)
  | Mouse_motion                (* The mouse is moved *)
  | Poll                        (* Don't wait; return immediately *)
```

To specify events to wait for.

```
value wait_next_event : event list -> status
```

Wait until one of the events specified in the given event list occurs, and return the status of the mouse and keyboard at that time. If `Poll` is given in the event list, return immediately with the current status. If the mouse cursor is outside of the graphics window, the `mouse_x` and `mouse_y` fields of the event are outside the range `0..size_x()-1, 0..size_y()-1`. Keypresses are queued, and dequeued one by one when the `Key_pressed` event is specified.

Mouse and keyboard polling

```
value mouse_pos : unit -> int * int
```

Return the position of the mouse cursor, relative to the graphics window. If the mouse cursor is outside of the graphics window, `mouse_pos()` returns a point outside of the range `0..size_x()-1, 0..size_y()-1`.

```
value button_down : unit -> bool
```

Return `true` if the mouse button is pressed, `false` otherwise.

```
value read_key : unit -> char
```

Wait for a key to be pressed, and return the corresponding character. Keypresses are queued.

```
value key_pressed : unit -> bool
```

Return `true` if a keypress is available; that is, if `read_key` would not block.

Sound

```
value sound : int -> int -> unit
```

`sound freq dur` plays a sound at frequency `freq` (in hertz) for a duration `dur` (in milliseconds). On the Macintosh, for obscure technical reasons, the frequency is rounded to the nearest note in the equal-tempered scale.

Part IV

The Caml Light commands

Chapter 7

Batch compilation (camlc)

This chapter describes how Caml Light programs can be compiled non-interactively, and turned into standalone executable files. This is achieved by the command `camlc`, which compiles and links Caml Light source files.

Mac: This command is not a standalone Macintosh application. To run `camlc`, you need the Macintosh Programmer's Workshop (MPW) programming environment. The programs generated by `camlc` are also MPW tools, not standalone Macintosh applications.

7.1 Overview of the compiler

The `camlc` command has a command-line interface similar to the one of most C compilers. It accepts several types of arguments: source files for module implementations; source files for module interfaces; and compiled module implementations.

- Arguments ending in `.mli` are taken to be source files for module interfaces. Module interfaces declare exported global identifiers, define public data types, and so on. From the file *x*`.mli`, the `camlc` compiler produces a compiled interface in the file *x*`.zi`.
- Arguments ending in `.ml` are taken to be source files for module implementation. Module implementations bind global identifiers to values, define private data types, and contain expressions to be evaluated for their side-effects. From the file *x*`.ml`, the `camlc` compiler produces compiled object code in the file *x*`.zo`. If the interface file *x*`.mli` exists, the module implementation *x*`.ml` is checked against the corresponding compiled interface *x*`.zi`, which is assumed to exist. If no interface *x*`.mli` is provided, the compilation of *x*`.ml` produces a compiled interface file *x*`.zi` in addition to the compiled object code file *x*`.zo`. The produced file *x*`.zi` produced corresponds to an interface that exports everything that is defined in the implementation *x*`.ml`.
- Arguments ending in `.zo` are taken to be compiled object code. These files are linked together, along with the object code files obtained by compiling `.ml` arguments (if any), and the Caml Light standard library, to produce a standalone executable program. The order in which `.zo` and `.ml` arguments are presented on the command line is relevant: global identifiers are initialized in that order at run-time, and it is a link-time error to use a global identifier before having initialized it. Hence, a given *x*`.zo` file must come before all `.zo` files that refer to identifiers defined in the file *x*`.zo`.

The output of the linking phase is a file containing compiled code that can be executed by the Caml Light runtime system: the command named `camlrun`. If `caml.out` is the name of the file produced by the linking phase, the command

`camlrun caml.out` $arg_1\ arg_2\ \ldots\ arg_n$

executes the compiled code contained in `caml.out`, passing it as arguments the character strings arg_1 to arg_n. (See the chapter on `camlrun` for more details.)

Unix: On most Unix systems, the file produced by the linking phase can be run directly, as in:

`./caml.out` $arg_1\ arg_2\ \ldots\ arg_n$

The produced file has the executable bit set, and it manages to launch the bytecode interpreter by itself.

PC: The output file produced by the linking phase is directly executable, provided it is given extension `.EXE`. Hence, if the output file is named `caml_out.exe`, you can execute it with the command

`caml_out` $arg_1\ arg_2\ \ldots\ arg_n$

Actually, the produced file `caml_out.exe` is a tiny executable file prepended to the bytecode file. The executable simply runs the `camlrun` runtime system on the remainder of the file. (As a consequence, this is not a standalone executable: it still requires `camlrun.exe` to reside in one of the directories in the path.)

7.2 Options

The following command-line options are recognized by `camlc`.

`-c` Compile only. Suppress the linking phase of the compilation. Source code files are turned into compiled files, but no executable file is produced. This option is useful to compile modules separately.

`-custom`
Link in "custom runtime" mode. In the default linking mode, the linker produces bytecode that is intended to be executed with the shared runtime system, `camlrun`. In the custom runtime mode, the linker produces an output file that contains both the runtime system and the bytecode for the program. The resulting file is considerably larger, but it can be executed directly, even if the `camlrun` command is not installed. Moreover, the "custom runtime" mode enables linking Caml Light code with user-defined C functions, as described in chapter 12.

Unix: Never strip an executable produced with the `-custom` option.

PC: This option is not implemented.

Mac: This option is not implemented.

`-files` *response-file*
: Process the files whose names are listed in file *response-file*, just as if these names appeared on the command line. File names in *response-file* are separated by blanks (spaces, tabs, newlines). This option allows to overcome silly limitations on the length of the command line.

`-g` Cause the compiler to produce additional debugging information. During the linking phase, this option add information at the end of the executable bytecode file produced. This information is required in particular by the catch-all exceptions handlers provided by the module `printexc` in the standard library.

During the compilation of an implementation file (`.ml` file), the compiler writes a `.zix` file when the `-g` option is set. This `.zix` file describes the full interface of the `.ml` file, that is, all types and values defined in the `.ml` file, including those that are local to the `.ml` file (i.e. not declared in the `.mli` interface file). Used in conjunction with the `-g` option to the toplevel system (chapter **??**), the `.zix` file gives access to the local values of the module, making it possible to print or "trace" them. The `.zix` file is not produced if the implementation file has no explicit interface, since, in this case, the module has no local values.

`-i` Cause the compiler to print the declared types, exceptions, and global variables (with their inferred types) when compiling an implementation (`.ml` file). This can be useful to check the types inferred by the compiler. Also, since the output follows the syntax of module interfaces, it can help in writing an explicit interface (`.mli` file) for a file: just redirect the standard output of the compiler to a `.mli` file, and edit that file to remove all declarations of unexported globals.

`-I` *directory*
: Add the given directory to the list of directories searched for compiled interface files (`.zi`) and compiled object code files (`.zo`). By default, the current directory is searched first, then the standard library directory. Directories added with `-I` are searched after the current directory, but before the standard library directory. When several directories are added with several `-I` options on the command line, these directories are searched from right to left (the rightmost directory is searched first, the leftmost is searched last). (Directories can also be added to the search path from inside the programs with the `#directory` directive; see chapter 3.)

`-o` *exec-file*
: Specify the name of the output file produced by the linker.

Unix: The default output name is `a.out`, in keeping with the tradition.

PC: The default output name is `CAML_OUT.EXE`.

Mac: The default output name is `Caml.Out`.

`-O` *module-set*
: Specify which set of standard modules is to be implicitly "opened" at the beginning of a compilation. There are three module sets currently available:

cautious
: provides the standard operations on integers, floating-point numbers, characters, strings,

arrays, ..., as well as exception handling, basic input/output, etc. Operations from the `cautious` set perform range and bound checking on string and array operations, as well as various sanity checks on their arguments.

fast
: provides the same operations as the `cautious` set, but without sanity checks on their arguments. Programs compiled with `-O fast` are therefore slightly faster, but unsafe.

none
: suppresses all automatic opening of modules. Compilation starts in an almost empty environment. This option is not of general use, except to compile the standard library itself.

The default compilation mode is `-O cautious`. See chapter 4 for a complete listing of the modules in the `cautious` and `fast` sets.

`-v` Print the version number the various passes of the compiler.

7.3 Modules and the file system

This short section is intended to clarify the relationship between the names of the modules and the names of the files that contain their compiled interface and compiled implementation.

The compiler always derives the name of the compiled module by taking the base name of the source file (`.ml` or `.mli` file). That is, it strips the leading directory name, if any, as well as the `.ml` or `.mli` suffix. The produced `.zi` and `.zo` files have the same base name as the source file; hence, the compiled files produced by the compiler always have their base name equal to the name of the module they describe (for `.zi` files) or implement (for `.zo` files).

For compiled interface files (`.zi` files), this invariant must be preserved at all times, since the compiler relies on it to load the compiled interface file for the modules that are used from the module being compiled. Hence, it is risky and generally incorrect to rename `.zi` files. It is admissible to move them to another directory, if their base name is preserved, and the correct `-I` options are given to the compiler.

Compiled bytecode files (`.zo` files), on the other hand, can be freely renamed once created. That's because 1- `.zo` files contain the true name of the module they define, so there is no need to derive that name from the file name; 2- the linker never attempts to find by itself the `.zo` file that implements a module of a given name: it relies on the user providing the list of `.zo` files by hand.

7.4 Common errors

This section describes and explains the most frequently encountered error messages.

Cannot find file *filename*
: The named file could not be found in the current directory, nor in the directories of the search path. The *filename* is either a compiled interface file (`.zi` file), or a compiled bytecode file (`.zo` file). If *filename* has the format *mod*`.zi`, this means you are trying to compile a file that references identifiers from module *mod*, but you have not yet compiled an interface for module *mod*. Fix: compile *mod*`.mli` or *mod*`.ml` first, to create the compiled interface *mod*`.zi`.

  If *filename* has the format *mod*`.zo`, this means you are trying to link a bytecode object file that does not exist yet. Fix: compile *mod*`.ml` first.

If your program spans several directories, this error can also appear because you haven't specified the directories to look into. Fix: add the correct `-I` options to the command line.

Corrupted compiled interface file *filename*
The compiler produces this error when it tries to read a compiled interface file (`.zi` file) that has the wrong structure. This means something went wrong when this `.zi` file was written: the disk was full, the compiler was interrupted in the middle of the file creation, and so on. This error can also appear if a `.zi` file is modified after its creation by the compiler. Fix: remove the corrupted `.zi` file, and rebuild it.

Expression of type t_1 cannot be used with type t_2
This is by far the most common type error in programs. Type t_1 is the type inferred for the expression (the part of the program that is displayed in the error message), by looking at the expression itself. Type t_2 is the type expected by the context of the expression; it is deduced by looking at how the value of this expression is used in the rest of the program. If the two types t_1 and t_2 are not compatible, then the error above is produced.

In some cases, it is hard to understand why the two types t_1 and t_2 are incompatible. For instance, the compiler can report that "expression of type `foo` cannot be used with type `foo`", and it really seems that the two types `foo` are compatible. This is not always true. Two type constructors can have the same name, but actually represent different types. This can happen if a type constructor is redefined. Example:

```
type foo = A | B;;
let f = function A -> 0 | B -> 1;;
type foo = C | D;;
f C;;
```

This result in the error message "expression `C` of type `foo` cannot be used with type `foo`".

Incompatible types with the same names can also appear when a module is changed and recompiled, but some of its clients are not recompiled. That's because type constructors in `.zi` files are not represented by their name (that would not suffice to identify them, because of type redefinitions), but by unique stamps that are assigned when the type declaration is compiled. Consider the three modules:

```
mod1.ml:    type t = A | B;;
            let f = function A -> 0 | B -> 1;;

mod2.ml:    let g x = 1 + mod1__f(x);;

mod3.ml:    mod2__g mod1__A;;
```

Now, assume `mod1.ml` is changed and recompiled, but `mod2.ml` is not recompiled. The recompilation of `mod1.ml` can change the stamp assigned to type `t`. But the interface `mod2.zi` will still use the old stamp for `mod1__t` in the type of `mod2__g`. Hence, when compiling `mod3.ml`, the system complains that the argument type of `mod2__g` (that is, `mod1__t` with the old stamp) is not compatible with the type of `mod1__A` (that is, `mod1__t` with the new stamp). Fix: use `make` or a similar tool to ensure that all clients of a module *mod* are recompiled when the interface *mod*`.zi` changes. To check that the `Makefile` contains the right dependencies,

remove all `.zi` files and rebuild the whole program; if no "Cannot find file" error appears, you're all set.

mod__name **is referenced before being defined**

This error appears when trying to link an incomplete or incorrectly ordered set of files. Either you have forgotten to provide an implementation for the module named *mod* on the command line (typically, the file named *mod*`.zo`, or a library containing that file). Fix: add the missing `.ml` or `.zo` file to the command line. Or, you have provided an implementation for the module named *mod*, but it comes too late on the command line: the implementation of *mod* must come before all bytecode object files that reference one of the global variables defined in module *mod*. Fix: change the order of `.ml` and `.zo` files on the command line.

Of course, you will always encounter this error if you have mutually recursive functions across modules. That is, function `mod1__f` calls function `mod2__g`, and function `mod2__g` calls function `mod1__f`. In this case, no matter what permutations you perform on the command line, the program will be rejected at link-time. Fixes:

- Put `f` and `g` in the same module.
- Parameterize one function by the other. That is, instead of having

```
mod1.ml:    let f x = ... mod2__g ... ;;
mod2.ml:    let g y = ... mod1__f ... ;;
```

  define

```
mod1.ml:    let f g x = ... g ... ;;
mod2.ml:    let rec g y = ... mod1__f g ... ;;
```

  and link `mod1` before `mod2`.

- Use a reference to hold one of the two functions, as in :

```
mod1.ml:    let forward_g =
                ref((fun x -> failwith "forward_g") : <type>);;
            let f x = ... !forward_g ... ;;
mod2.ml:    let g y = ... mod1__f ... ;;
            mod1__forward_g := g;;
```

Unavailable C primitive *f*

This error appears when trying to link code that calls external functions written in C in "default runtime" mode. As explained in chapter 12, such code must be linked in "custom runtime" mode. Fix: add the `-custom` option, as well as the (native code) libraries and (native code) object files that implement the required external functions.

Chapter 8

The toplevel system (camllight)

This chapter describes the toplevel system for Caml Light, that permits interactive use of the Caml Light system, through a read-eval-print loop. In this mode, the system repeatedly reads Caml Light phrases from the input, then typechecks, compile and evaluate them, then prints the inferred type and result value, if any. The system prints a `#` (sharp) prompt before reading each phrase. A phrase can extends several lines; phrases are delimited by `;;` (the final double-semicolon).

From the standpoint of the module system, all phrases entered at toplevel are treated as the implementation of a module named `top`. Hence, all toplevel definitions are entered in the module `top`.

Unix: The toplevel system is started by the command `camllight`. Phrases are read on standard input, results are printed on standard output, errors on standard error. End-of-file on standard input terminates `camllight` (see also the `quit` system function below).

The toplevel system does not perform line editing, but it can easily be used in conjunction with an external line editor such as `fep`; just run `fep -emacs camllight` or `fep -vi camllight`.

At any point, the parsing, compilation or evaluation of the current phrase can be interrupted by pressing `ctrl-C` (or, more precisely, by sending the `intr` signal to the `camllight` process). This goes back to the `#` prompt.

Mac: The toplevel system is presented as the standalone Macintosh application `Caml Light`. This application does not require the Macintosh Programmer's Workshop to run.

Once launched from the Finder, the application opens two windows, "Caml Light Input" and "Caml Light Output". Phrases are entered in the "Caml Light Input" window. The "Caml Light Output" window displays a copy of the input phrases as they are processed by the Caml Light toplevel, interspersed with the toplevel responses. The "Return" key sends the contents of the Input window to the Caml Light toplevel. The "Enter" key inserts a newline without sending the contents of the Input window. (This can be configured with the "Preferences" menu item.)

The contents of the input window can be edited at all times, with the standard Macintosh interface. An history of previously entered phrases is maintained, and can be accessed with the "Previous entry" (command-P) and "Next entry" (command-N) menu items.

To quit the `Caml Light` application, either select "Quit" from the "Files" menu, or use the `quit` function described below.

At any point, the parsing, compilation or evaluation of the current phrase can be interrupted by pressing "command-dot", or by selecting the item "Interrupt Caml Light" in the "Caml Light" menu. This goes back to the `#` prompt.

PC: The toplevel system is started by the command `caml`. (The full name `camllight` had to be truncated because of the well-known limitations of MS-DOS.) Phrases are read on standard input, results are printed on standard output, errors on standard error. The `quit` system function terminates `caml`.

A number of TSR line editors that provide line editing and history facilities for the `command.com` command-line interpreter also work in conjunction with `caml`. Some TSR's that work: `dosedit`, `ced`, `toddy`. Some TSR's that don't work: `doskey` from the MS-DOS 5.0 distribution.

With the 8086 version, the parsing, compilation or evaluation of the current phrase can be interrupted by pressing `ctrl-break`. This goes back to the `#` prompt. Pressing `ctrl-C` interrupts input, but does not stop a phrase that performs no input-output.

With the 80386 version, pressing `ctrl-C` or `ctrl-break` interrupts the system at any point: during parsing as well as during compilation or evaluation.

8.1 Options

The following command-line options are recognized by the `caml` or `camllight` commands.

`-g` Start the toplevel system in debugging mode. This mode gives access to values and types that are local to a module, that is, not exported by the interface of the module. When debugging mode is off, these local objects are not accessible (attempts to access them produce an "Unbound identifier" error). When debugging mode is on, these objects become visible, just like the objects that are exported in the module interface. In particular, values of abstract types are printed using their concrete representations, and the functions local to a module can be "traced" (see the `trace` function in section ??). This applies only to the modules that have been compiled in debugging mode (either by the batch compiler with the `-g` option, or by the toplevel system in debugging mode), that is, those modules that have an associated `.zix` file.

`-I` *directory*
Add the given directory to the list of directories searched for compiled interface files (`.zi`) and compiled object code files (`.zo`). By default, the current directory is searched first, then the standard library directory. Directories added with `-I` are searched after the current directory, but before the standard library directory. When several directories are added with several `-I` options on the command line, these directories are searched from right to left (the rightmost directory is searched first, the leftmost is searched last). Directories can also be added to the search path once the toplevel is running with the `#directory` directive; see chapter 3.

`-O` *module-set*
Specify which set of standard modules is to be implicitly "opened" when the toplevel starts. There are three module sets currently available:

cautious
provides the standard operations on integers, floating-point numbers, characters, strings,

arrays, ..., as well as exception handling, basic input/output, ...Operations from the `cautious` set perform range and bound checking on string and vector operations, as well as various sanity checks on their arguments.

fast
: provides the same operations as the `cautious` set, but without sanity checks on their arguments. Programs compiled with `-O fast` are therefore slightly faster, but unsafe.

none
: suppresses all automatic opening of modules. Compilation starts in an almost empty environment. This option is not of general use.

The default compilation mode is `-O cautious`. See chapter 4 for a complete listing of the modules in the `cautious` and `fast` sets.

8.2 Toplevel control functions

The standard library module `toplevel`, opened by default when the toplevel system is launched, provides a number of functions that control the toplevel behavior, load files in memory, and trace program execution.

`quit ()`
: Exit the toplevel loop and terminate the Caml Light command.

`include "`*filename*`"`
: Read, compile and execute source phrases from the file *filename*. The `.ml` extension is automatically added to *filename*, if not present. This is textual inclusion: phrases are processed just as if they were typed on standard input. In particular, global identifiers defined by these phrases are entered in the module named `top`, not in a new module.

`load "`*filename*`"`
: Load in memory the source code for a module implementation. Read, compile and execute source phrases from file *filename*. The `.ml` extension is automatically added to *filename*, if not present. The `load` function behaves much like `include`, except that a new module is created, with name the base name of *filename*. Global identifiers defined in file *filename* are entered in this module, instead of the `top` module as in the case of `include`. For instance, assuming file `foo.ml` contains the single phrase

```
let bar = 1;;
```

executing `load "foo"` defines the identifier `foo__bar` with value 1.

Caveat: the loaded module is not automatically opened: the identifier `bar` does not automatically complete to `foo__bar`. To achieve this, you must execute the directive `#open "foo"` afterwards.

`compile "`*filename*`"`
: Compile the source code for a module implementation or interface (`.ml` or `.mli` file). Compilation proceeds as with the batch compiler, and produces the same results as `camlc -c` *filename*. If the toplevel system is in debugging mode (option `-g` or function `debug_mode` below), the compilation is also performed in debugging mode, as when giving the `-g` option

to the batch compiler. The result of the compilation is left in files (`.zo`, `.zi`, `.zix`). The compiled code is not loaded in memory.

`load_object "`*filename*`"`
: Load in memory the compiled bytecode contained in file *filename*. The `.zo` extension is automatically added to *filename*, if not present. The bytecode file has been produced by the standalone compiler `camlc`. Global identifiers defined in file *filename* are entered in their own module, not in the `top` module, just as with the `load` function.

`trace "`*function-name*`"`
: After the execution of `trace "foo"`, all calls to the global function named `foo` will be "traced". That is, the argument and the result are displayed for each call, as well as the exceptions escaping out of `foo`, either raised by `foo` itself, or raised by one of the functions called from `foo`. If `foo` is a curried function, each argument is printed as it is passed to the function.

`untrace "`*function-name*`"`
: Executing `untrace "foo"` stops all tracing over the global function named `foo`.

`debug_mode` *boolean-flag*
: Set the debugging mode on (`debug_mode true`) or off (`debug_mode false`). Executing `debug_mode true` is equivalent to starting the toplevel system with the `-g` option. In debugging mode, values and types local to a module (not exported by its interface) are visible, just like the exported values and types. See the description of the `-g` option in section ??.

`verbose_mode` *boolean-flag*
: After executing `verbose_mode true`, the compilations performed by `compile` print the objects defined by the module, along with their types, just as with the `-i` option to the batch compiler. To revert to silent compilations, execute `verbose_mode false`.

`cd "`*directory-name*`"`
: Set the current working directory.

`gc ()`
: Finish the current garbage collection cycle, and return the amount of free space in the heap, in bytes. If you want to perform a complete GC cycle, call this function twice.

8.3 The toplevel and the module system

Toplevel phrases can refer to identifiers defined in modules other than the `top` module with the same mechanisms as for separately compiled modules: either by using qualified identifiers (`modulename__localname`), or by using unqualified identifiers that are automatically completed by searching the list of opened modules. (See section 2 of the reference manual.) The modules opened at startup are given in the documentation for the standard library. Other modules can be opened with the `#open` directive.

However, before referencing a global variable from a module other than the `top` module, a definition of that global variable must have been entered in memory. At start-up, the toplevel system contains the definitions for all the identifiers in the standard library. The definitions for user modules can be entered with the `load` or `load_object` functions described above. Referencing

a global variable for which no definition has been provided by `load` or `load_object` results in the error "`Identifier foo__bar is referenced before being defined`". Since this is a tricky point, let us consider some examples.

1. The library function `sub_string` is defined in module `string`. This module is part of the standard library, and is one of the modules automatically opened at start-up. Hence, both phrases

   ```
   sub_string "qwerty" 1 3;;
   string__sub_string "qwerty" 1 3;;
   ```

   are correct, without having to use `#open`, `load`, or `load_object`.

2. The library function `printf` is defined in module `printf`. This module is part of the standard library, but it is not automatically opened at start-up. Hence, the phrase

   ```
   printf__printf "%s %s" "hello" "world";;
   ```

   is correctly executed, while

   ```
   printf "%s %s" "hello" "world";;
   ```

   causes the error "`Variable printf is unbound`", since none of the currently opened modules define a global with local name `printf`. However,

   ```
   #open "printf";;
   printf "%s %s" "hello" "world";;
   ```

   executes correctly.

3. Assume the file `foo.ml` resides in the current directory, and contains the single phrase

   ```
   let x = 1;;
   ```

   When the toplevel starts, references to `x` will cause the error "`Variable x is unbound`". References to `foo__x` will cause the error "`Cannot find file foo.zi`", since the typechecker is attempting to load the compiled interface for module `foo` in order to find the type of `x`. To load in memory the module `foo`, just do:

   ```
   load "foo";;
   ```

   Then, references to `foo__x` typecheck and evaluate correctly. Since `load` does not open the module it loads, references to `x` will still fail with the error "`Variable x is unbound`". You will have to do

   ```
   #open "foo";;
   ```

   explicitly, for `x` to complete automatically into `foo__x`.

4. Finally, assume the file `foo.ml` above has been previously compiled with the `camlc -c` command. The current directory therefore contains a compiled interface `foo.zi`, claiming that `foo__x` is a global variable with type `int`, and a compiled bytecode file `foo.zo`, defining `foo__x` to have the value `1`. When the toplevel starts, references to `foo__x` will cause the error "`foo__x is referenced before being defined`". In contrast with case 3 above, the typechecker has succeeded in finding the compiled interface for module `foo`. But execution cannot proceed, because no definition for `foo__x` has been entered in memory. To do so, execute:

```
load_object "foo";;
```

   This loads the file `foo.zo` in memory, therefore defining `foo__x`. Then, references to `foo__x` evaluate correctly. As in case 3 above, references to `x` still fail, because `load_object` does not open the module it loads. Again, you will have to do

```
#open "foo";;
```

   explicitly, for `x` to complete automatically into `foo__x`.

8.4 Common errors

This section describes and explains the most frequently encountered error messages.

Cannot find file *filename*
: The named file could not be found in the current directory, nor in the directories of the search path.

  If *filename* has the format *mod*`.zi`, this means the current phrase references identifiers from module *mod*, but you have not yet compiled an interface for module *mod*. Fix: either `load` the file *mod*`.ml`, which will also create in memory the compiled interface for module *mod*; or use `camlc` to compile *mod*`.mli` or *mod*`.ml`, creating the compiled interface *mod*`.zi`, before you start the toplevel.

  If *filename* has the format *mod*`.zo`, this means you are trying to load with `load_object` a bytecode object file that does not exist yet. Fix: compile *mod*`.ml` with `camlc` before you start the toplevel. Or, use `load` instead of `load_object` to load the source code instead of a compiled object file.

  If *filename* has the format *mod*`.ml`, this means `load` or `include` could not find the specified source file. Fix: check the spelling of the file name, or write it if it does not exist.

mod__name **is referenced before being defined**
: You have neglected to load in memory an implementation for a module, with `load` or `load_object`. This is explained in full detail in section 8.3 above.

Corrupted compiled interface file *filename*
: See section 7.4.

Expression of type t_1 **cannot be used with type** t_2
: See section 7.4.

8.5 Building custom toplevel systems: `camlmktop`

The `camlmktop` command builds Caml Light toplevels that contain user code preloaded at start-up.

Mac: This command is not available in the Macintosh version.

PC: This command is not available in the PC versions.

The `camlmktop` command takes as argument a set of `.zo` files, and links them with the object files that implement the Caml Light toplevel. The typical use is:

```
camlmktop -o mytoplevel foo.zo bar.zo gee.zo
```

This creates the bytecode file `mytoplevel`, containing the Caml Light toplevel system, plus the code from the three `.zo` files. To run this toplevel, give it as argument to the `camllight` command:

```
camllight mytoplevel
```

This starts a regular toplevel loop, except that the code from `foo.zo`, `bar.zo` and `gee.zo` is already loaded in memory, just as if you had typed:

```
load_object "foo";;
load_object "bar";;
load_object "gee";;
```

on entrance to the toplevel. The modules `foo`, `bar` and `gee` are not opened, though; you still have to do

```
#open "foo";;
```

yourself, if this is what you wish.

8.6 Options

The following command-line options are recognized by `camlmktop`.

`-custom`
: Link in "custom runtime" mode. See the corresponding option for `camlc`, in chapter 7.

`-I` *directory*
: Add the given directory to the list of directories searched for compiled object code files (`.zo`).

`-o` *exec-file*
: Specify the name of the toplevel file produced by the linker. The default is `camltop.out`.

Chapter 9

The runtime system (camlrun)

The `camlrun` command executes bytecode files produced by the linking phase of the `camlc` command.

Mac: This command is a MPW tool, not a standalone Macintosh application.

9.1 Overview

The `camlrun` command comprises three main parts: the bytecode interpreter, that actually executes bytecode files; the memory allocator and garbage collector; and a set of C functions that implement primitive operations such as input/output.

The usage for `camlrun` is:

`camlrun` *options bytecode-executable* arg_1 arg_2 ... arg_n

The first non-option argument is taken to be the name of the file containing the executable bytecode. (That file is searched in the executable path as well as in the current directory.) The remaining arguments are passed to the Caml Light program, in the string array `sys__command_line`. Element 0 of this array is the name of the bytecode executable file; elements 1 to n are the remaining arguments arg_1 to arg_n.

As mentioned in chapter 7, in most cases, the bytecode executable files produced by the `camlc` command are self-executable, and manage to launch the `camlrun` command on themselves automatically. That is, assuming `caml.out` is a bytecode executable file,

`caml.out` arg_1 arg_2 ... arg_n

works exactly as

`camlrun caml.out` arg_1 arg_2 ... arg_n

Notice that it is not possible to pass options to `camlrun` when invoking directly `caml.out`.

9.2 Options

The following command-line option is recognized by `camlrun`. (There are additional options to control the behavior of the garbage collector, but they are not for the casual user.)

`-v` Cause `camlrun` to print various diagnostic messages relative to memory allocation and garbage collection. Useful to debug memory-related problems.

9.3 Common errors

This section describes and explains the most frequently encountered error messages.

filename: `no such file or directory`
: If *filename* is the name of a self-executable bytecode file, this means that either that file does not exist, or that it failed to run the `camlrun` bytecode interpreter on itself. The second possibility indicates that Caml Light has not been properly installed on your system.

`Cannot exec camlrun`
: (When launching a self-executable bytecode file.) The `camlrun` command could not be found in the executable path. Check that Caml Light has been properly installed on your system.

`Cannot find the bytecode file`
: The file that `camlrun` is trying to execute (e.g. the file given as first non-option argument to `camlrun`) either does not exist, or is not a valid executable bytecode file.

`Truncated bytecode file`
: The file that `camlrun` is trying to execute is not a valid executable bytecode file. Probably it has been truncated or mangled since created. Erase and rebuild it.

`Uncaught exception`
: The program being executed contains a "stray" exception. That is, it raises an exception at some point, and this exception is never caught. This causes immediate termination of the program. If you wish to know which exception thus escapes, use the `printexc__f` function from the standard library (and don't forget to link your program with the `-g` option).

`Out of memory`
: The program being executed requires more memory than available. Either the program builds too large data structures; or the program contains too many nested function calls, and the stack overflows. In some cases, your program is perfectly correct, it just requires more memory than your machine provides. (This happens quite frequently on small microcomputers, but is unlikely on Unix machines.) In other cases, the "out of memory" message reveals an error in your program: non-terminating recursive function, allocation of an excessively large array or string, attempts to build an infinite list or other data structure, ...

  To help you diagnose this error, run your program with the `-v` option to `camlrun`. If it displays lots of "`Growing stack...`" messages, this is probably a looping recursive function. If it displays lots of "`Growing heap...`" messages, with the heap size growing slowly, this is probably an attempt to construct a data structure with too many (infinitely many?) cells. If it displays few "`Growing heap...`" messages, but with a huge increment in the heap size, this is probably an attempt to build an excessively large array or string.

Chapter 10

The librarian (camllibr)

Mac: This command is a MPW tool, not a standalone Macintosh application.

10.1 Overview

The `camllibr` program packs in one single file a set of bytecode object files (`.zo` files). The resulting file is also a bytecode object file and also has the `.zo` extension. It can be passed to the link phase of the `camlc` compiler in replacement of the original set of bytecode object files. That is, after running

```
camllibr -o library.zo mod1.zo mod2.zo mod3.zi mod4.zo
```

all calls to the linker with the form

```
camlc ... library.zo ...
```

are exactly equivalent to

```
camlc ... mod1.zo mod2.zo mod3.zi mod4.zo ...
```

The typical use of `camllibr` is to build a library composed of several modules: this way, users of the library have only one `.zo` file to specify on the command line to `camlc`, instead of a bunch of `.zo` files, one per module contained in the library.

The linking phase of `camlc` is clever enough to discard the code corresponding to useless phrases: in particular, definitions for global variables that are never used after their definitions. Hence, there is no problem with putting many modules, even rarely used ones, into one single library: this will not result in bigger executables.

The usage for `camllibr` is:

`camllibr` *options* *file*$_1$`.zo` ... *file*$_n$`.zo`

where *file*$_1$`.zo` through *file*$_n$`.zo` are the object files to pack together. The order in which these file names are presented on the command line is relevant: the compiled phrases contained in the library will be executed in that order. (Remember that it is a link-time error to refer to a global variable that has not yet been defined.)

10.2 Options

The following command-line options are recognized by `camllibr`.

`-files` *response-file*
: Process the files whose names are listed in file *response-file*, just as if these names appeared on the command line. File names in *response-file* are separated by blanks (spaces, tabs, newlines). This option allows to overcome silly limitations on the length of the command line.

`-I` *directory*
: Add the given directory to the list of directories searched for the input `.zo` files. By default, the current directory is searched first, then the standard library directory. Directories added with `-I` are searched after the current directory, but before the standard library directory. When several directories are added with several `-I` options on the command line, these directories are searched from right to left (the rightmost directory is searched first, the leftmost is searched last).

`-o` *library-name*
: Specify the name of the output file. The default is `library.zo`.

10.3 Turning code into a library

To develop a library, it is usually more convenient to split it into several modules, that reflect the internal structure of the library. From the standpoint of the library users, however, it is preferable to view the library as a single module, with only one interface file (`.zi` file) and one implementation file (`.zo` file): linking is easier, and there is no need to put a bunch of `#open` directives, nor to have to remember the internal structure of the library.

The `camllibr` command allows having a single `.zo` file for the whole library. Here is how the Caml Light module system can be used (some say "abused") to have a single `.zi` file for the whole library. To be more concrete, assume that the library comprises three modules, `windows`, `images` and `buttons`. The idea is to add a fourth module, `mylib`, that re-exports the public parts of `windows`, `images` and `buttons`. The interface `mylib.mli` contains definitions for those types that are public (exported with their definitions), declarations for those types that are abstract (exported without their definitions), and declarations for the functions that can be called from the user's code:

```
(* File mylib.mli *)
type 'a option = None | Some of 'a;;     (* a public type *)
type window and image and button;;      (* three abstract types *)
value new_window : int -> int -> window (* the public functions *)
  and draw_image : image -> window -> int -> int -> unit
  and ...
```

The implementation of the `mylib` module simply equates the abstract types and the public functions to the corresponding types and functions in the modules `windows`, `images` and `buttons`:

```
(* File mylib.ml *)
type window == windows__win
 and image  == images__pixmap
```

```
 and button == buttons__t;;
let new_window = windows__open_window
and draw_image = images__draw
and ...
```

The files `windows.ml`, `images.ml` and `buttons.ml` can open the `mylib` module, to access the public types defined in the interface `mylib.mli`, such as the `option` type. Of course, these modules must not reference the abstract types nor the public functions, to avoid circularities.

Types such as `windows__win` in the example above can be exported by the `windows` module either abstractly or concretely (with their definition). Often, it is necessary to export them concretely, because the other modules in the library (`images`, `buttons`) need to build or destructure directly values of that type. Even if `windows__win` is exported concretely by the `windows` module, that type will remain abstract to the library user, since it is abstracted by the public interface `mylib`.

The actual building of the library `mylib` proceeds as follows:

```
camlc -c mylib.mli              # create mylib.zi
camlc -c windows.mli windows.ml images.mli images.ml
camlc -c buttons.mli buttons.ml
camlc -c mylib.ml               # create mylib.zo
mv mylib.zo tmplib.zo           # renaming to avoid overwriting mylib.zo
camllibr -o mylib.zo windows.zo images.zo buttons.zo tmplib.zo
```

Then, copy `mylib.zi` and `mylib.zo` to a place accessible to the library users. The other `.zi` and `.zo` files need not be copied.

Chapter 11

Lexer and parser generators (camllex, camlyacc)

This chapter describes two program generators: `camllex`, that produces a lexical analyzer from a set of regular expressions with associated semantic actions, and `camlyacc`, that produces a parser from a grammar with associated semantic actions.

These program generators are very close to the well-known `lex` and `yacc` commands that can be found in most C programming environments. This chapter assumes a working knowledge of `lex` and `yacc`: while it describes the input syntax for `camllex` and `camlyacc` and the main differences with `lex` and `yacc`, it does not explain the basics of writing a lexer or parser description in `lex` and `yacc`. Readers unfamiliar with `lex` and `yacc` are referred to "Compilers: principles, techniques, and tools" by Aho, Sethi and Ullman (Addison-Wesley, 1986), "Compiler design in C" by Holub (Prentice-Hall, 1990), or "Lex & Yacc", by Mason and Brown (O'Reilly, 1990).

Streams and stream matching, as described in chapter **??**, provide an alternative way to write lexers and parsers. The stream matching technique is more powerful than the combination of `camllex` and `camlyacc` in some cases (higher-order parsers), but less powerful in other cases (precedences). Choose whichever approach is more adapted to your parsing problem.

Mac: These commands are MPW tool, not standalone Macintosh applications.

86: These commands are not available in the 8086 PC version.

11.1 Overview of `camllex`

The `camllex` command produces a lexical analyzer from a set of regular expressions with attached semantic actions, in the style of `lex`. Assuming the input file is *lexer*`.mll`, executing

`camllex` *lexer*`.mll`

produces Caml Light code for a lexical analyzer in file *lexer*`.ml`. This file defines one lexing function per entry point in the lexer definition. These functions have the same names as the entry points. Lexing functions take as argument a lexer buffer, and return the semantic attribute of the corresponding entry point.

Lexer buffers are an abstract data type implemented in the standard library module `lexing`. The functions `create_lexer_channel`, `create_lexer_string` and `create_lexer` from module

`lexing` create lexer buffers that read from an input channel, a character string, or any reading function, respectively. (See the description of module `lexing` in chapter 4.)

When used in conjunction with a parser generated by `camlyacc`, the semantic actions compute a value belonging to the type `token` defined by the generated parsing module. (See the description of `camlyacc` below.)

11.2 Syntax of lexer definitions

The format of lexer definitions is as follows:

```
{ header }
rule entrypoint =
  parse regexp { action }
      | ...
      | regexp { action }
and entrypoint =
  parse ...
and ...
;;
```

Comments are delimited by `(*` and `*)`, as in Caml Light.

11.2.1 Header

The *header* section is arbitrary Caml Light text enclosed in curly braces. It can be omitted. If it is present, the enclosed text is copied as is at the beginning of the output file. Typically, the header section contains the `#open` directives required by the actions, and possibly some auxiliary functions used in the actions.

11.2.2 Entry points

The names of the entry points must be valid Caml Light identifiers.

11.2.3 Regular expressions

The regular expressions are in the style of `lex`, with a more Caml-like syntax.

‘ *char* ‘
: A character constant, with the same syntax as Caml Light character constants. Match the denoted character.

`_`
: Match any character.

`eof`
: Match the end of the lexer input.

" *string* "
: A string constant, with the same syntax as Caml Light string constants. Match the corresponding sequence of characters.

`[` *character-set* `]`
: Match any single character belonging to the given character set. Valid character sets are: single character constants ‘ *c* ‘; ranges of characters ‘ c_1 ‘ `-` ‘ c_2 ‘ (all characters between c_1 and c_2, inclusive); and the union of two or more character sets, denoted by concatenation.

`[^` *character-set* `]`
: Match any single character not belonging to the given character set.

regexp `*`
: (Repetition.) Match the concatenation of zero or more strings that match *regexp*.

regexp `+`
: (Strict repetition.) Match the concatenation of one or more strings that match *regexp*.

regexp `?`
: (Option.) Match either the empty string, or a string matching *regexp*.

$regexp_1$ `|` $regexp_2$
: (Alternative.) Match any string that matches either $regexp_1$ or $regexp_2$

$regexp_1$ $regexp_2$
: (Concatenation.) Match the concatenation of two strings, the first matching $regexp_1$, the second matching $regexp_2$.

`(` *regexp* `)`
: Match the same strings as *regexp*.

Concerning the precedences of operators, `*` and `+` have highest precedence, followed by `?`, then concatenation, then `|` (alternation).

11.2.4 Actions

The actions are arbitrary Caml Light expressions. They are evaluated in a context where the identifier `lexbuf` is bound to the current lexer buffer. Some typical uses for `lexbuf`, in conjunction with the operations on lexer buffers provided by the `lexing` standard library module, are listed below.

`lexing__get_lexeme lexbuf`
: Return the matched string.

`lexing__get_lexeme_char lexbuf` n
: Return the n^{th} character in the matched string. The first character corresponds to $n = 0$.

`lexing__get_lexeme_start lexbuf`
: Return the absolute position in the input text of the beginning of the matched string. The first character read from the input text has position 0.

`lexing__get_lexeme_end lexbuf`
: Return the absolute position in the input text of the end of the matched string. The first character read from the input text has position 0.

entrypoint `lexbuf`
: (Where *entrypoint* is the name of another entry point in the same lexer definition.) Recursively call the lexer on the given entry point. Useful for lexing nested comments, for example.

11.3 Overview of `camlyacc`

The `camlyacc` command produces a parser from a context-free grammar specification with attached semantic actions, in the style of `yacc`. Assuming the input file is *grammar*`.mly`, executing

`camlyacc` *options grammar*`.mly`

produces Caml Light code for a parser in the file *grammar*`.ml`, and its interface in file *grammar*`.mli`.

The generated module defines one parsing function per entry point in the grammar. These functions have the same names as the entry points. Parsing functions take as arguments a lexical analyzer (a function from lexer buffers to tokens) and a lexer buffer, and return the semantic attribute of the corresponding entry point. Lexical analyzer functions are usually generated from a lexer specification by the `camllex` program. Lexer buffers are an abstract data type implemented in the standard library module `lexing`. Tokens are values from the concrete type `token`, defined in the interface file *grammar*`.mli` produced by `camlyacc`.

11.4 Syntax of grammar definitions

Grammar definitions have the following format:

```
%{
  header
%}
  declarations
%%
  rules
%%
  trailer
```

Comments are enclosed between `/*` and `*/` pairs, as in C, except in the header and trailer sections.

11.4.1 Header and trailer

The header and the trailer sections are Caml Light code that is copied as is into file *grammar*`.ml`. Both sections are optional. The header goes at the beginning of the output file; it usually contains `#open` directives required by the semantic actions of the rules. The trailer goes at the end of the output file.

11.4.2 Declarations

Declarations are given one per line. They all start with a `%` sign.

`%token` *symbol ... symbol*

Declare the given symbols as tokens (terminal symbols). These symbols are added as constant constructors for the `token` concrete type.

`%token <` *type* `>` *symbol ... symbol*

Declare the given symbols as tokens with an attached attribute of the given type. These symbols are added as constructors with arguments of the given type for the `token` concrete type.

The *type* part is an arbitrary Caml Light type expression, except that all type constructor names must be fully qualified (e.g. `modname__typename`) for all types except standard built-in types, even if the proper `#open` directives (e.g. `#open "modname"`) were given in the header section. That's because the header is copied only to the `.ml` output file, but not to the `.mli` output file, while the *type* part of a `%token` declaration is copied to both.

`%start` *symbol ... symbol*

Declare the given symbols as entry points for the grammar. For each entry point, a parsing function with the same name is defined in the output module. Non-terminals that are not declared as entry points have no such parsing function. Start symbols must be given a type with the `%type` directive below.

`%type <` *type* `>` *symbol ... symbol*

Specify the type of the semantic attributes for the given symbols. This is mandatory for start symbols only. Other nonterminal symbols need not be given types by hand: these types will be inferred when running the output files through the Caml Light compiler (unless the `-s` option is in effect). The *type* part is an arbitrary Caml Light type expression, except that all type constructor names must be fully qualified (e.g. `modname__typename`) for all types except standard built-in types, even if the proper `#open` directives (e.g. `#open "modname"`) were given in the header section. That's because the header is copied only to the `.ml` output file, but not to the `.mli` output file, while the *type* part of a `%token` declaration is copied to both.

`%left` *symbol ... symbol*

`%right` *symbol ... symbol*

`%nonassoc` *symbol ... symbol*

Associate precedences and associativities to the given symbols. All symbols on the same line are given the same precedence. They have higher precedence than symbols declared before in a `%left`, `%right` or `%nonassoc` line. They have lower precedence than symbols declared after in a `%left`, `%right` or `%nonassoc` line. The symbols are declared to associate to the left (`%left`), to the right (`%right`), or to be non-associative (`%nonassoc`). The symbols are usually tokens. They can also be dummy nonterminals, for use with the `%prec` directive inside the rules.

11.4.3 Rules

The syntax for rules is as usual:

nonterminal :
    symbol ... symbol `{` *semantic-action* `}`
  `|` ...
  `|` *symbol ... symbol* `{` *semantic-action* `}`
;

Rules can also contain the `%prec` *symbol* directive in the right-hand side part, to override the default precedence and associativity of the rule with the precedence and associativity of the given symbol.

Semantic actions are arbitrary Caml Light expressions, that are evaluated to produce the semantic attribute attached to the defined nonterminal. The semantic actions can access the semantic attributes of the symbols in the right-hand side of the rule with the `$` notation: `$1` is the attribute for the first (leftmost) symbol, `$2` is the attribute for the second symbol, etc.

Actions occurring in the middle of rules are not supported. Error recovery is not implemented.

11.5 Options

The `camlyacc` command recognizes the following options:

`-v` Generate a description of the parsing tables and a report on conflicts resulting from ambiguities in the grammar. The description is put in file *grammar*`.output`.

`-s` Generate a *grammar*`.ml` file with smaller phrases. Semantic actions are presented in the *grammar*`.ml` output file as one large vector of functions. By default, this vector is built by a single phrase. When the grammar is large, or contains complicated semantic actions, the resulting phrase may require large amounts of memory to be compiled by Caml Light. With the `-s` option, the vector of actions is constructed incrementally, one phrase per action. This lowers the memory requirements for the compiler, but it is no longer possible to infer the types of nonterminal symbols: typechecking is turned off on symbols that do not have a type specified by a `%type` directive.

`-b`*prefix*
Name the output files *prefix*`.ml`, *prefix*`.mli`, *prefix*`.output`, instead of the default naming convention.

11.6 A complete example

The all-time favorite: a desk calculator. This program reads arithmetic expressions on standard input, one per line, and prints their values. Here is the grammar definition:

```
/* File parser.mly */
%token <int> INT
%token PLUS MINUS TIMES DIV
%token LPAREN RPAREN
%token EOL
%left PLUS MINUS        /* lowest precedence */
%left TIMES DIV         /* medium precedence */
%nonassoc UMINUS        /* highest precedence */
%start Main             /* the entry point */
%type <int> Main
%%
Main:
    Expr EOL                { $1 }
;
```

```
Expr:
    INT                     { $1 }
  | LPAREN Expr RPAREN      { $2 }
  | Expr PLUS Expr          { $1 + $3 }
  | Expr MINUS Expr         { $1 - $3 }
  | Expr TIMES Expr         { $1 * $3 }
  | Expr DIV Expr           { $1 / $3 }
  | MINUS Expr %prec UMINUS { - $2 }
;
```

Here is the definition for the corresponding lexer:

```
(* File lexer.mll *)
{
#open "parser";;        (* The type token is defined in parser.mli *)
exception Eof;;
}
rule Token = parse
    [` ` `\t`]     { Token lexbuf }     (* skip blanks *)
  | [`\n` ]        { EOL }
  | [`0`-`9`]+     { INT(int_of_string (get_lexeme lexbuf)) }
  | `+`            { PLUS }
  | `-`            { MINUS }
  | `*`            { TIMES }
  | `/`            { DIV }
  | `(`            { LPAREN }
  | `)`            { RPAREN }
  | eof            { raise Eof }
;;
```

Here is the main program, that combines the parser with the lexer:

```
(* File calc.ml *)
try
  let lexbuf = lexing__create_lexer_channel std_in in
  while true do
    let result = parser__Main lexer__Token lexbuf in
      print_int result; print_newline(); flush std_out
  done
with Eof ->
  ()
;;
```

To compile everything, execute:

```
camllex lexer.mll       # generates lexer.ml
camlyacc parser.mly     # generates parser.ml and parser.mli
camlc -c parser.mli
camlc -c lexer.ml
camlc -c parser.ml
```

```
camlc -c calc.ml
camlc -o calc lexer.zo parser.zo calc.zo
```

Chapter 12

Interfacing C with Caml Light

This chapter describes how user-defined primitives, written in C, can be added to the Caml Light runtime system and called from Caml Light code.

Mac: This facility is not implemented in the Macintosh version.

PC: This facility is not implemented in the PC versions.

12.1 Overview and compilation information

12.1.1 Declaring primitives

User primitives are declared in a module interface (a `.mli` file), in the same way as a regular ML value, except that the declaration is followed by the `=` sign, the function arity (number of arguments), and the name of the corresponding C function. For instance, here is how the `input` primitive is declared in the interface for the standard library module `io`:

```
value input : in_channel -> string -> int -> int -> int
            = 4 "input"
```

Primitives with several arguments are always curried. The C function does not necessarily have the same name as the ML function.

Values thus declared primitive in a module interface must not be implemented in the module implementation (the `.ml` file). They can be used inside the module implementation.

12.1.2 Implementing primitives

User primitives with arity $n \leq 5$ are implemented by C functions that take n arguments of type `value`, and return a result of type `value`. The type `value` is the type of the representations for Caml Light values. It encodes objects of several base types (integers, floating-point numbers, strings, ...), as well as Caml Light data structures. The type `value` and the associated conversion functions and macros are described in details below. For instance, here is the declaration for the C function implementing the `input` primitive:

```
value input(channel, buffer, offset, length)
        value channel, buffer, offset, length;
{
```

```
  ...
}
```

When the primitive function is applied in a Caml Light program, the C function is called with the values of the expressions to which the primitive is applied as arguments. The value returned by the function is passed back to the Caml Light program as the result of the function application.

User primitives with arity greater than 5 are implemented by C functions that receive two arguments: a pointer to an array of Caml Light values (the values for the arguments), and an integer which is the number of arguments provided:

```
value prim_with_lots_of_args(argv, argn)
        value * argv;
        int argn;
{
  ... argv[0] ...;              /* The first argument */
  ... argv[6] ...;              /* The seventh argument */
}
```

Implementing a user primitive is actually two separate tasks: on the one hand, decoding the arguments to extract C values from the given Caml Light values, and encoding the return value as a Caml Light value; on the other hand, actually computing the result from the arguments. Except for very simple primitives, it is often preferable to have two distinct C functions to implement these two tasks. The first function actually implements the primitive, taking native C values as arguments and returning a native C value. The second function, often called the "stub code", is a simple wrapper around the first function that converts its arguments from Caml Light values to C values, call the first function, and convert the returned C value to Caml Light value. For instance, here is the stub code for the `input` primitive:

```
value input(channel, buffer, offset, length)
        value channel, buffer, offset, length;
{
  return Val_long(getblock((struct channel *) channel,
                           &Byte(buffer, Long_val(offset)),
                           Long_val(length)));
}
```

(Here, `Val_long`, `Long_val` and so on are conversion macros for the type `value`, that will be described later.) The hard work is performed by the function `getblock`, which is declared as:

```
long getblock(channel, p, n)
     struct channel * channel;
     char * p;
     long n;
{
  ...
}
```

To write C code that operates on Caml Light values, the following include files are provided:

`mlvalues.h`	definition of the `value` type, and conversion macros
`alloc.h`	allocation functions (to create structured Caml Light objects)
`memory.h`	miscellaneous memory-related functions (for in-place modification of structures, etc).

These files reside in the Caml Light standard library directory (usually `/usr/local/lib/caml-light`).

12.1.3 Linking C code with Caml Light code

The Caml Light runtime system comprises three main parts: the bytecode interpreter, the memory manager, and a set of C functions that implement the primitive operations. Some bytecode instructions are provided to call these C functions, designated by their offset in a table of functions (the table of primitives).

In the default mode, the Caml Light linker produces bytecode for the standard runtime system, with a standard set of primitives. References to primitives that are not in this standard set result in the "unavailable C primitive" error.

In the "custom runtime" mode, the Caml Light linker scans the bytecode object files (`.zo` files) and determines the set of required primitives. Then, it builds a suitable runtime system, by calling the native code linker with:

- the table of the required primitives
- a library that provides the bytecode interpreter, the memory manager, and the standard primitives
- libraries and object code files (`.o` files) mentioned on the command line for the Caml Light linker, that provide implementations for the user's primitives.

This builds a runtime system with the required primitives. The Caml Light linker generates bytecode for this custom runtime system. The bytecode is appended to the end of the custom runtime system, so that it will be automatically executed when the output file (custom runtime + bytecode) is launched.

To link in "custom runtime" mode, execute the `camlc` command with:

- the `-custom` option
- the names of the desired Caml Light object files (`.zo` files)
- the names of the C object files and libraries (`.o` and `.a` files) that implement the required primitives. (Libraries can also be specified with the usual `-l` syntax.)

12.2 The `value` type

All Caml Light objects are represented by the C type `value`, defined in the include file `mlvalues.h`, along with macros to manipulate values of that type. An object of type `value` is either:

- an unboxed integer
- a pointer to a block inside the heap (such as the blocks allocated through one of the `alloc_*` functions below)
- a pointer to an object outside the heap (e.g., a pointer to a block allocated by `malloc`, or to a C variable).

12.2.1 Integer values

Integer values encode 31-bit signed integers. They are unboxed (unallocated).

12.2.2 Blocks

Blocks in the heap are garbage-collected, and therefore have strict structure constraints. Each block includes a header containing the size of the block (in words), and the tag of the block. The tag governs how the contents of the blocks are structured. A tag lower than `No_scan_tag` indicates a structured block, containing well-formed values, which is recursively traversed by the garbage collector. A tag greater than or equal to `No_scan_tag` indicates a raw block, whose contents are not scanned by the garbage collector. For the benefits of ad-hoc polymorphic primitives such as equality and structured input-output, structured and raw blocks are further classified according to their tags as follows:

Tag	Contents of the block
0 to `No_scan_tag` − 1	A structured block (an array of Caml Light objects). Each field is a `value`.
`Closure_tag`	A closure representing a functional value. The first word is a pointer to a piece of bytecode, the second word is a `value` containing the environment.
`String_tag`	A character string.
`Double_tag`	A double-precision floating-point number.
`Abstract_tag`	A block representing an abstract datatype.
`Final_tag`	A block representing an abstract datatype with a "finalization" function, to be called when the block is deallocated.

12.2.3 Pointers to outside the heap

Any pointer to outside the heap can be safely cast to and from the type `value`. This includes pointers returned by `malloc`, and pointers to C variables obtained with the `&` operator.

12.3 Representation of Caml Light data types

This section describes how Caml Light data types are encoded in the `value` type.

12.3.1 Atomic types

`int`	Unboxed integer values.
`char`	Unboxed integer values (ASCII code).
`float`	Blocks with tag `Double_tag`.
`string`	Blocks with tag `String_tag`.

12.3.2 Product types

Tuples and arrays are represented by pointers to blocks, with tag 0.

Records are also represented by zero-tagged blocks. The ordering of labels in the record type declaration determines the layout of the record fields: the value associated to the label declared first is stored in field 0 of the block, the value associated to the label declared next goes in field 1, and so on.

12.3.3 Concrete types

Constructed terms are represented by blocks whose tag encode the constructor. The constructors for a given concrete type are numbered from 0 to the number of constructors minus one, following the order in which they appear in the concrete type declaration. Constant constructors are represented by zero-sized blocks (atoms), tagged with the constructor number. Non-constant constructors declared with a n-tuple as argument are represented by a block of size n, tagged with the constructor number; the n fields contain the components of its tuple argument. Other non-constant constructors are represented by a block of size 1, tagged with the constructor number; the field 0 contains the value of the constructor argument. Example:

Constructed term	Representation
`()`	Size = 0, tag = 0
`false`	Size = 0, tag = 0
`true`	Size = 0, tag = 1
`[]`	Size = 0, tag = 0
`h::t`	Size = 2, tag = 1, first field = `h`, second field = `t`

12.4 Operations on values

12.4.1 Kind tests

- `Is_int`(v) is true if value v is an immediate integer, false otherwise
- `Is_block`(v) is true if value v is a pointer to a block, and false if it is an immediate integer.

12.4.2 Operations on integers

- `Val_long`(l) returns the value encoding the `long int` l
- `Long_val`(v) returns the `long int` encoded in value v
- `Val_int`(i) returns the value encoding the `int` i
- `Int_val`(v) returns the `int` encoded in value v

12.4.3 Accessing blocks

- `Wosize_val`(v) returns the size of value v, in words, excluding the header.
- `Tag_val`(v) returns the tag of value v.
- `Field`(v, n) returns the value contained in the n^{th} field of the structured block v. Fields are numbered from 0 to `Wosize_val`(v) $-$ 1.
- `Code_val`(v) returns the code part of the closure v.
- `Env_val`(v) returns the environment part of the closure v.
- `string_length`(v) returns the length (number of characters) of the string v.
- `Byte`(v, n) returns the n^{th} character of the string v, with type `char`. Characters are numbered from 0 to `string_length`(v) $-$ 1.

- `Byte_u`(v, n) returns the n^{th} character of the string v, with type `unsigned char`. Characters are numbered from 0 to `string_length`$(v) - 1$.

- `String_val`(v) returns a pointer to the first byte of the string v, with type `char *`. This pointer is a valid C string: there is a null character after the last character in the string. However, Caml Light strings can contain embedded null characters, that will confuse the usual C functions over strings.

- `Double_val`(v) returns the floating-point number contained in value v, with type `double`.

The expressions `Field`(v, n), `Code_val`(v), `Env_val`(v), `Byte`(v, n), `Byte_u`(v, n) and `Double_val`(v) are valid l-values. Hence, they can be assigned to, resulting in an in-place modification of value v. Assigning directly to `Field`(v, n) must be done with care to avoid confusing the garbage collector (see below).

12.4.4 Allocating blocks

From the standpoint of the allocation functions, blocks are divided according to their size as zero-sized blocks, small blocks (with size less than or equal to `Max_young_size`), and large blocks (with size greater than to `Max_young_size`). The constant `Max_young_size` is declared in the include file `mlvalues.h`. It is guaranteed to be at least 64 (words), so that any block with constant size less than or equal to 64 can be assumed to be small. For blocks whose size is computed at run-time, the size must be compared against `Max_young_size` to determine the correct allocation procedure.

- `Atom`(t) returns an "atom" (zero-sized block) with tag t. Zero-sized blocks are preallocated outside of the heap. It is incorrect to try and allocate a zero-sized block using the functions below. For instance, `Atom`(0) represents `()`, `false` and `[]`; `Atom`(1) represents `true`. (As a convenience, `mlvalues.h` defines the macros `Val_unit`, `Val_false` and `Val_true`.)

- `alloc`(n, t) returns a fresh small block of size $n \leq$ `Max_young_size` words, with tag t. If this block is a structured block (i.e. if $t <$ `No_scan_tag`), then the fields of the block (initially containing garbage) must be initialized with legal values (using direct assignment to the fields of the block) before the next allocation.

- `alloc_tuple`(n) returns a fresh small block of size $n \leq$ `Max_young_size` words, with tag 0. The fields of this block must be filled with legal values before the next allocation or modification.

- `alloc_shr`(n, t) returns a fresh block of size n, with tag t. The size of the block can be greater than `Max_young_size`. (It can also be smaller, but in this case it is more efficient to call `alloc` instead of `alloc_shr`.) If this block is a structured block (i.e. if $t <$ `No_scan_tag`), then the fields of the block (initially containing garbage) must be initialized with legal values (using the `initialize` function described below) before the next allocation.

- `alloc_string`(n) returns a string value of length n characters. The string initially contains garbage.

- `copy_string`(s) returns a string value containing a copy of the null-terminated C string s (a `char *`).

- `copy_double`(d) returns a floating-point value initialized with the `double` d.

- `alloc_array`(f, a) allocates an array of values, calling function f over each element of the input array a to transform it into a value. The array a is an array of pointers terminated by the null pointer. The function f receives each pointer as argument, and returns a value. The zero-tagged block returned by `alloc_array`(f, a) is filled with the values returned by the successive calls to f.
- `copy_string_array`(p) allocates an array of strings, copied from the pointer to a string array p (a `char **`).

12.4.5 Raising exceptions

C functions cannot raise arbitrary exceptions. However, two functions are provided to raise two standard exceptions:

- `failwith`(s), where s is a null-terminated C string (with type `char *`), raises exception `Failure` with argument s.
- `invalid_argument`(s), where s is a null-terminated C string (with type `char *`), raises exception `Invalid_argument` with argument s.

12.5 Living in harmony with the garbage collector

Unused blocks in the heap are automatically reclaimed by the garbage collector. This requires some cooperation from C code that manipulates heap-allocated blocks.

Rule 1 *After a structured block (a block with tag less than* `No_scan_tag`*) is allocated, all fields of this block must be filled with well-formed values before the next allocation operation. If the block has been allocated with* `alloc` *or* `alloc_tuple`*, filling is performed by direct assignment to the fields of the block:*

$$\texttt{Field}(v, n) = v_n;$$

If the block has been allocated with `alloc_shr`*, filling is performed through the* `initialize` *function:*

$$\texttt{initialize}(\&\texttt{Field}(v, n), v_n);$$

The next allocation can trigger a garbage collection. The garbage collector assumes that all structured blocks contain well-formed values. Newly created blocks contain random data, which generally do not represent well-formed values.

If you really need to allocate before the fields can receive their final value, first initialize with a constant value (e.g. `Val_long(0)`), then allocate, then modify the fields with the correct value (see rule 3).

Rule 2 *Local variables containing values must be registered with the garbage collector (using the* `Push_roots` *and* `Pop_roots` *macros), if they are to survive a call to an allocation function.*

Registration is performed with the `Push_roots` and `Pop_roots` macros. `Push_roots`(r, n) declares an array r of n values and registers them with the garbage collector. The values contained in $r[0]$ to $r[n-1]$ are treated like roots by the garbage collector. A root value has the following properties: if it points to a heap-allocated block, this block (and its contents) will not be reclaimed; moreover, if this block is relocated by the garbage collector, the root value is updated to point to the

new location for the block. `Push_roots`(r, n) must occur in a C block exactly between the last local variable declaration and the first statement in the block. To un-register the roots, `Pop_roots`() must be called before the C block containing `Push_roots`(r, n) is exited. (Roots are automatically un-registered if a Caml exception is raised.)

Rule 3 *Direct assignment to a field of a block, as in*

$$\texttt{Field}(v, n) = v';$$

is safe only if v is a block newly allocated by `alloc` *or* `alloc_tuple`*; that is, if no allocation took place between the allocation of v and the assignment to the field. In all other cases, never assign directly. If the block has just been allocated by* `alloc_shr`*, use* `initialize` *to assign a value to a field for the first time:*

$$\texttt{initialize}(\&\texttt{Field}(v, n), v');$$

Otherwise, you are updating a field that previously contained a well-formed value; then, call the `modify` *function:*

$$\texttt{modify}(\&\texttt{Field}(v, n), v');$$

To illustrate the rules above, here is a C function that builds and returns a list containing the two integers given as parameters:

```
value alloc_list_int(i1, i2)
        int i1, i2;
{
  value result;
  Push_roots(r, 1);
  r[0] = alloc(2, 1);                    /* Allocate a cons cell */
  Field(r[0], 0) = Val_int(i2);          /* car = the integer i2 */
  Field(r[0], 1) = Atom(0);              /* cdr = the empty list [] */
  result = alloc(2, 1);                  /* Allocate the other cons cell */
  Field(result, 0) = Val_int(i1);        /* car = the integer i1 */
  Field(result, 1) = r[0];               /* cdr = the first cons cell */
  Pop_roots();
  return result;
}
```

The "cons" cell allocated first needs to survive the allocation of the other cons cell; hence, the value returned by the first call to `alloc` must be stored in a registered root. The value returned by the second call to `alloc` can reside in the un-registered local variable `result`, since we won't do any further allocation in this function.

In the example above, the list is built bottom-up. Here is an alternate way, that proceeds top-down. It is less efficient, but illustrates the use of `modify`.

```
value alloc_list_int(i1, i2)
        int i1, i2;
{
  value tail;
  Push_roots(r, 1);
  r[0] = alloc(2, 1);                    /* Allocate a cons cell */
```

```
  Field(r[0], 0) = Val_int(i1);          /* car = the integer i1 */
  Field(r[0], 1) = Val_int(0);           /* A dummy value
  tail = alloc(2, 1);                    /* Allocate the other cons cell */
  Field(tail, 0) = Val_int(i2);          /* car = the integer i2 */
  Field(tail, 1) = Atom(0);              /* cdr = the empty list [] */
  modify(&Field(r[0], 1), tail);         /* cdr of the result = tail */
  Pop_roots();
  return r[0];
}
```

It would be incorrect to perform `Field(r[0], 1) = tail` directly, because the allocation of `tail` has taken place since `r[0]` was allocated.

12.6 A complete example

This section outlines how the functions from the Unix `curses` library can be made available to Caml Light programs. First of all, here is the interface `curses.mli` that declares the `curses` primitives and data types:

```
type window;;                    (* The type "window" remains abstract *)
value initscr: unit -> window = 1 "curses_initscr"
  and endwin: unit -> unit = 1 "curses_endwin"
  and refresh: unit -> unit = 1 "curses_refresh"
  and wrefresh : window -> unit = 1 "curses_wrefresh"
  and newwin: int -> int -> int -> int -> window = 4 "curses_newwin"
  and mvwin: window -> int -> int -> unit = 3 "curses_mvwin"
  and addch: char -> unit = 1 "curses_addch"
  and mvwaddch: window -> int -> int -> char -> unit = 4 "curses_mvwaddch"
  and addstr: string -> unit = 1 "curses_addstr"
  and mvwaddstr: window -> int -> int -> string -> unit = 4 "curses_mvwaddstr"
;; (* lots more omitted *)
```

To compile this interface:

```
        camlc -c curses.mli
```

To implement these functions, we just have to provide the stub code; the core functions are already implemented in the `curses` library. The stub code file, `curses.o`, looks like:

```
#include <curses.h>
#include <mlvalues.h>

value curses_initscr(unit)
        value unit;
{
  return (value) initscr();     /* OK to coerce directly from WINDOW * to value
                                   since that's a block created by malloc() */
}
```

```
value curses_wrefresh(win)
        value win;
{
  wrefresh((value) win);
  return Val_unit;               /* Or Atom(0) */
}

value curses_newwin(nlines, ncols, x0, y0)
        value nlines, ncols, x0, y0;
{
  return (value) newwin(Int_val(nlines), Int_val(ncols),
                        Int_val(x0), Int_val(y0));
}

value curses_addch(c)
        value c;
{
  addch(Int_val(c));            /* Characters are encoded like integers */
  return Val_unit;
}

value curses_addstr(s)
        value s;
{
  addstr(String_val(s));
  return Val_unit;
}

/* This goes on for pages. */
```

(Actually, it would be better to create a library for the stub code, with each stub code function in a separate file, so that linking would pick only those functions from the `curses` library that are actually used.)

The file `curses.c` can be compiled with:

```
cc -c -I/usr/local/lib/caml-light curses.c
```

or, even simpler,

```
camlc -c curses.c
```

(When passed a `.c` file, the `camlc` command simply calls `cc` on that file, with the right `-I` option.)

Now, here is a sample Caml Light program `test.ml` that uses the `curses` module:

```
#open "curses";;
let main_window = initscr () in
let small_window = newwin 10 5 20 10 in
  mvwaddstr main_window 10 2 "Hello";
  mvwaddstr small_window 4 3 "world";
  refresh();
```

```
  for i = 1 to 100000 do () done;
  endwin()
;;
```

To compile this program, run:

```
camlc -c test.ml
```

Finally, to link everything together:

```
camlc -custom -o test test.zo curses.o -lcurses
```

Part V

Appendix

Chapter 13

Further reading

For the interested reader, we list below some references to books and reports related (sometimes loosely) to Caml Light.

13.1 Programming in ML

The books below are programming courses taught in ML. Their main goal is to teach programming, not to describe ML in full details — though most contain fairly good introductions to the ML language. Some of those books use the Standard ML dialect instead of the Caml dialect, so you will have to keep in mind the differences in syntax and in semantics.

- Pierre Weis and Xavier Leroy. *Le langage Caml.* InterÉditions, 1993.

  The natural companion to this manual, provided you read French. This book is a step-by-step introduction to programming in Caml, and presents many realistic examples of Caml programs.

- Lawrence C. Paulson. *ML for the working programmer.* Cambridge University Press, 1991.

  An excellent introduction to programming in Standard ML. Develops a theorem prover as a complete example. Contains a presentation of the module system of Standard ML.

- Chris Reade. *Elements of Functional Programming.* Addison-Wesley, 1989.

  An introduction to Standard ML, with sections on denotational semantics and lambda-calculus.

- Ryan Stansifer. *ML primer.* Prentice-Hall, 1992.

  A short, but nice introduction to programming in Standard ML.

- Thérèse Accart Hardin and Véronique Donzeau-Gouge Viguié. *Concepts et outils de la programmation. Du fonctionnel à l'impératif avec Caml et Ada.* InterÉditions, 1992.

  A first course in programming, that first introduces the main programming notions in Caml, then shows them underlying Ada. Intended for beginners; slow-paced for the others.

- Ake Wikstrom. *Functional programming using Standard ML.* Prentice-Hall, 1987.

  A first course in programming, taught in Standard ML. Intended for absolute beginners; slow-paced for the others.

- Harold Abelson and Gerald Jay Sussman. *Structure and Interpretation of Computer Programs.* The MIT press, 1985. (French translation: *Structure et interprétation des programmes informatiques*, InterÉditions, 1989.)

  An outstanding course on programming, taught in Scheme, the modern dialect of Lisp. Well worth reading, even if you are more interested in ML than in Lisp.

13.2 Descriptions of ML dialects

The books and reports below are descriptions of various programming languages from the ML family. They assume some familiarity with ML.

- Xavier Leroy and Pierre Weis. *Manuel de référence du langage Caml.* InterÉditions, 1993.

  The French edition of the present reference manual and user's manual.

- Robert Harper. *Introduction to Standard ML.* Technical report ECS-LFCS-86-14, University of Edinburgh, 1986.

  An overview of Standard ML, including the module system. Terse, but still readable.

- Robin Milner, Mads Tofte and Robert Harper. *The definition of Standard ML.* The MIT press, 1990.

  A complete formal definition of Standard ML, in the framework of structured operational semantics. This book is probably the most mathematically precise definition of a programming language ever written. It is heavy on formalism and horribly terse, so even readers that are thoroughly familiar with ML will have major difficulties with it.

- Robin Milner and Mads Tofte. *Commentary on Standard ML.* The MIT Press, 1991.

  A commentary on the book above, that attempts to explain the most delicate parts and motivate the design choices. Easier to read than the Definition, but still rather involving.

- Guy Cousineau and Gérard Huet. *The CAML primer.* Technical report 122, INRIA, 1990.

  A short description of the original Caml system, from which Caml Light has evolved. Some familiarity with Lisp is assumed.

- Pierre Weis et al. *The CAML reference manual, version 2.6.1.* Technical report 121, INRIA, 1990.

  The manual for the original Caml system, from which Caml Light has evolved.

- Michael J. Gordon, Arthur J. Milner and Christopher P. Wadsworth. *Edinburgh LCF.* Lecture Notes in Computer Science volume 78, Springer-Verlag, 1979.

  This is the first published description of the ML language, at the time when it was nothing more than the control language for the LCF system, a theorem prover. This book is now obsolete, since the ML language has much evolved since then; but it is still of historical interest.

- Paul Hudak, Simon Peyton-Jones and Philip Wadler. *Report on the programming language Haskell, version 1.1.* Technical report, Yale University, 1991.

Haskell is a purely functional language with lazy semantics that shares many important points with ML (full functionality, polymorphic typing), but has interesting features of its own (dynamic overloading, also called type classes).

13.3 Implementing functional programming languages

The references below are intended for those who are curious to learn how a language like Caml Light is compiled and implemented.

- Xavier Leroy. *The ZINC experiment: an economical implementation of the ML language.* Technical report 117, INRIA, 1990. (Available by anonymous FTP on `ftp.inria.fr`.)

  A description of the ZINC implementation, the prototype ML implementation that has evolved into Caml Light. Large parts of this report still apply to the current Caml Light system, in particular the description of the execution model and abstract machine. Other parts are now obsolete. Yet this report still gives a complete overview of the implementation techniques used in Caml Light.

- Simon Peyton-Jones. *The implementation of functional programming languages.* Prentice-Hall, 1987. (French translation: *Mise en œuvre des langages fonctionnels de programmation*, Masson, 1990.)

  An excellent description of the implementation of purely functional languages with lazy semantics, using the technique known as graph reduction. The part of the book that deals with the transformation from ML to enriched lambda-calculus directly applies to Caml Light. You will find a good description of how pattern-matching is compiled and how types are inferred. The remainder of the book does not apply directly to Caml Light, since Caml Light is not purely functional (it has side-effects), has strict semantics, and does not use graph reduction at all.

- Andrew W. Appel. *Compiling with continuations.* Cambridge University Press, 1992.

  A complete description of an optimizing compiler for Standard ML, based on an intermediate representation called continuation-passing style. Shows how many advanced program optimizations can be applied to ML. Not directly relevant to the Caml Light system, since Caml Light does not use continuation-passing style at all, and makes little attempts at optimizing programs.

13.4 Applications of ML

The following reports show ML at work in various, sometimes unexpected, areas.

- Emmanuel Chailloux and Guy Cousineau. *The MLgraph primer.* Technical report 92-15, École Normale Supérieure, 1992. (Available by anonymous FTP on `ftp.ens.fr`.)

  Describes a Caml Light library that produces Postscript pictures through high-level drawing functions.

- Xavier Leroy. *Programmation du système Unix en Caml Light.* Technical report 147, INRIA, 1992. (Available by anonymous FTP on `ftp.inria.fr`.)

A Unix systems programming course, demonstrating the use of the Caml Light library that gives access to Unix system calls.

- John H. Reppy. *Concurrent programming with events — The concurrent ML manual.* Cornell University, 1990. (Available by anonymous FTP on `research.att.com`.)

  Concurrent ML extends Standard ML of New Jersey with concurrent processes that communicate through channels and events.

- Jeannette M. Wing, Manuel Faehndrich, J. Gregory Morrisett and Scottt Nettles. *Extensions to Standard ML to support transactions.* Technical report CMU-CS-92-132, Carnegie-Mellon University, 1992. (Available by anonymous FTP on `reports.adm.cs.cmu.edu`.)

  How to integrate the basic database operations to Standard ML.

- Emden R. Gansner and John H. Reppy. *eXene.* Bell Labs, 1991. (Available by anonymous FTP on `research.att.com`.)

  An interface between Standard ML of New Jersey and the X Windows windowing system.

Index to the library

! (infix), 63
!= (infix), 49
* (infix), 50, 52
*. (infix), 50
+ (infix), 50, 52
+. (infix), 50
- (infix), 50, 52
-. (infix), 50
/ (infix), 50, 53
/. (infix), 50
< (infix), 51, 53
<. (infix), 51
<= (infix), 51, 53
<=. (infix), 51
<> (infix), 49
<>. (infix), 51
= (infix), 49
=. (infix), 50
== (infix), 49
> (infix), 51, 53
>. (infix), 51
>= (infix), 51, 53
>=. (infix), 51
@ (infix), 60
^ (infix), 64

abs, 53
abs_float, 51
acos, 51
add, 70, 73, 75, 78, 79
add_float, 50
add_int, 52
arg (module), 69
asin, 51
asr (infix), 54
assoc, 62
assq, 62
atan, 51
atan2, 51

background, 86
Bad (exception), 70
baltree (module), 70
basename, 72
black, 86
blit_image, 88
blit_string, 65
blit_vect, 66
blue, 86
bool (module), 47
Break (exception), 82
builtin (module), 48
button_down, 89

catch_break, 82
cd, 102
char (module), 48
char_for_read, 48
char_of_int, 48
chdir, 82
check_suffix, 72
chop_suffix, 72
clear, 73, 78, 80
clear_graph, 85
clear_parser, 76
close, 82
close_graph, 85
close_in, 59
close_out, 58
combine, 62
command_line, 81
compare, 71, 79
compare_strings, 65
compile, 101
concat, 71
concat_vect, 66
contains, 71
copy_vect, 66
cos, 51

create_image, 88
create_lexer, 74
create_lexer_channel, 74
create_lexer_string, 74
create_string, 64
current_dir_name, 71
current_point, 86
cyan, 86

debug_mode, 102
decr, 63
diff, 79
dirname, 72
div_float, 50
div_int, 53
Division_by_zero (exception), 52
do_list, 60
do_list2, 60
do_list_combine, 62
do_stream, 64
do_table, 73
do_vect, 67
draw_arc, 86
draw_char, 87
draw_circle, 86
draw_ellipse, 86
draw_image, 88
draw_string, 87
dump_image, 88

elements, 79
empty, 75, 79
Empty (exception), 78, 80
End_of_file (exception), 54
end_of_stream, 64
eq (module), 49
eq_float, 50
eq_int, 53
eq_string, 65
equal, 79
exc (module), 49
except, 61
exceptq, 61
exists, 61
exit, 55, 82
Exit (exception), 49
exp, 51

Failure (exception), 49
failwith, 49
fchar (module), 50
filename (module), 71
fill_arc, 87
fill_circle, 87
fill_ellipse, 87
fill_poly, 87
fill_rect, 87
fill_string, 65
fill_vect, 66
find, 71, 73, 76
find_all, 73
flat_map, 61
float, 78
float (module), 50
float_of_int, 50
float_of_string, 51
flush, 57
fold, 80
for_all, 61
foreground, 86
fprint, 77
fprintf, 77
fst, 62
fstring (module), 52
fvect (module), 52

gc, 102
ge_float, 51
ge_int, 53
ge_string, 65
genlex (module), 72
get_image, 88
get_lexeme, 75
get_lexeme_char, 75
get_lexeme_end, 75
get_lexeme_start, 75
getenv, 82
Graphic_failure (exception), 85
graphics (module), 85
green, 86
gt_float, 51
gt_int, 53
gt_string, 65

hash, 74
hash_param, 74

`hashtbl` (module), 73
`hd`, 60

`in_channel_length`, 59
`include`, 101
`incr`, 63
`index`, 61
`init`, 78
`input`, 59
`input_binary_int`, 59
`input_byte`, 59
`input_char`, 58
`input_line`, 58
`input_value`, 59
`int`, 78
`int` (module), 52
`int_of_char`, 48
`int_of_float`, 50
`int_of_string`, 54
`inter`, 79
`interactive`, 81
`intersect`, 61
`invalid_arg`, 49
`Invalid_argument` (exception), 49
`io` (module), 54
`is_absolute`, 72
`is_empty`, 79
`it_list`, 60
`it_list2`, 60
`iter`, 76, 78, 80

`key_pressed`, 89

`land` (infix), 53
`le_float`, 51
`le_int`, 53
`le_string`, 65
`length`, 78, 80
`lexing` (module), 74
`lineto`, 86
`list` (module), 60
`list_it`, 60
`list_it2`, 61
`list_length`, 60
`list_of_vect`, 67
`lnot`, 54
`load`, 101
`load_object`, 102
`log`, 51
`lor` (infix), 53
`lshift_left`, 54
`lshift_right`, 54
`lsl` (infix), 54
`lsr` (infix), 54
`lt_float`, 51
`lt_int`, 53
`lt_string`, 65
`lxor` (infix), 54

`magenta`, 86
`make_image`, 88
`make_lexer`, 72
`make_matrix`, 66
`make_string`, 65
`make_vect`, 66
`map`, 60
`map` (module), 75
`map2`, 60
`map_combine`, 62
`map_vect`, 67
`map_vect_list`, 67
`Match_failure` (exception), 30–32, 48
`max`, 53
`mem`, 61, 79
`mem_assoc`, 62
`memq`, 61
`merge`, 80
`min`, 53
`minus`, 50, 52
`minus_float`, 50
`minus_int`, 52
`mod` (infix), 53
`modify`, 71
`mouse_pos`, 89
`moveto`, 86
`mult_float`, 50
`mult_int`, 52

`neq_float`, 51
`neq_int`, 53
`neq_string`, 65
`new`, 73, 78, 80
`not` (infix), 48
`Not_found` (exception), 49
`nth_char`, 64

open, 82
open_descriptor_in, 58
open_descriptor_out, 57
open_graph, 85
open_in, 58
open_in_bin, 58
open_in_gen, 58
open_out, 56
open_out_bin, 56
open_out_gen, 57
out_channel_length, 58
Out_of_memory (exception), 49
output, 57
output_binary_int, 57
output_byte, 57
output_char, 57
output_string, 57
output_value, 57

pair (module), 62
parse, 70
Parse_error (exception), 63
Parse_failure (exception), 63
parsing (module), 76
peek, 78
plot, 86
point_color, 86
pop, 80
pos_in, 59
pos_out, 58
power, 51
pred, 52
prerr_char, 56
prerr_endline, 56
prerr_float, 56
prerr_int, 56
prerr_string, 56
print, 77
print_char, 55
print_endline, 55
print_float, 55
print_int, 55
print_newline, 55
print_string, 55
printexc (module), 76
printf, 77
printf (module), 77
push, 80

queue (module), 77
quit, 101
quo (infix), 53

raise, 49
random (module), 78
read_float, 56
read_int, 56
read_key, 89
read_line, 56
really_input, 59
red, 86
ref (module), 63
remove, 71, 73, 76, 79, 82
rename, 82
replace_string, 65
rev, 60
rgb, 85
rhs_end, 76
rhs_start, 76

s_irall, 81
s_irgrp, 81
s_iroth, 81
s_irusr, 81
s_isgid, 81
s_isuid, 81
s_iwall, 81
s_iwgrp, 81
s_iwoth, 81
s_iwusr, 81
s_ixall, 81
s_ixgrp, 81
s_ixoth, 81
s_ixusr, 81
seek_in, 59
seek_out, 57
set (module), 79
set_color, 85
set_font, 87
set_line_width, 86
set_nth_char, 64
set_text_size, 87
sin, 51
size_x, 85
size_y, 85

`snd`, 62
`sort`, 80
`sort` (module), 80
`sound`, 89
`split`, 62, 71
`sqrt`, 51
`stack` (module), 80
`std_err`, 55
`std_in`, 55
`std_out`, 55
`stderr`, 55
`stdin`, 55
`stdout`, 55
`stream` (module), 63
`stream_check`, 64
`stream_from`, 63
`stream_get`, 64
`stream_next`, 63
`stream_of_channel`, 63
`stream_of_string`, 63
`string` (module), 64
`string_for_read`, 65
`string_length`, 64
`string_of_float`, 51
`string_of_int`, 54
`sub_float`, 50
`sub_int`, 52
`sub_string`, 64
`sub_vect`, 66
`subtract`, 61
`succ`, 52
`symbol_end`, 76
`symbol_start`, 76
`sys` (module), 81
`Sys_error` (exception), 81

`take`, 78
`tan`, 51
`text_size`, 87
`tl`, 60
`trace`, 102
`transp`, 87

`union`, 61, 79
`untrace`, 102

`vect` (module), 65
`vect_assign`, 66
`vect_item`, 66
`vect_length`, 65
`vect_of_list`, 67
`verbose_mode`, 102

`wait_next_event`, 88
`white`, 86

`yellow`, 86

Index of keywords

`and`, *see* `let`, `type`, `exception`, `value`

`begin`, 28, 29

`do`, *see* `while`, `for`
`done`, *see* `while`, `for`
`downto`, *see* `for`

`else`, *see* `if`
`end`, 28, 29
`exception`, *see also* 49, 37, 38

`for`, 28, 32
`fun`, 28
`function`, 28

`if`, 28, 31
`in`, *see* `let`

`let`, 28, 30

`match`, 28, 32
`mutable`, 36, 37

`not`, 28

`of`, *see* `type`, `exception`
`or`, 28, 32

`prefix`, 28, 34

`rec`, *see* `let`

`then`, *see* `if`
`to`, *see* `for`
`try`, 28, 33
`type`, 36, 38

`value`, 38

`while`, 32
`with`, *see* `match`, `try`