

The Standard ML Core Language

Robin Milner

University of Edinburgh

1. Introduction

1.1 How this proposal evolved; 1.2 Design principles; 1.3 An example.

2. The bare language

2.1 Discussion; 2.2 Reserved words; 2.3 Special constants; 2.4 Identifiers;  
2.5 Comments; 2.6 Lexical analysis; 2.7 Delimiters; 2.8 The bare syntax.

3. Evaluation

3.1 Environments and values; 3.2 Environment manipulation;  
3.3 Matching patterns; 3.4 Applying a match;  
3.5 Evaluation of expressions; 3.6 Evaluation of value bindings;  
3.7 Evaluation of type bindings; 3.8 Evaluation of exception bindings;  
3.9 Evaluation of declarations; 3.10 Evaluation of programs.

4. Directives

5. Standard bindings

5.1 Standard type constructors; 5.2 Standard functions and constants;  
5.3 Standard exceptions.

6. Standard derived forms

6.1 Expressions and patterns; 6.2 Bindings and declarations.

7. References and equality

7.1 References and assignment; 7.2 Equality.

8. Exceptions

8.1 Discussion; 8.2 Derived forms; 8.3 An example;  
8.4 Some pathological examples.

9. Type-checking

10. Syntactic restrictions

11. Conclusion

REFERENCES

- APPENDICES: 1. Syntax: Expressions and Patterns  
2. Syntax: Types, Bindings, Declarations and Programs  
3. Predeclared Variables and Constructors

## 1. Introduction

### 1.1 How this proposal evolved

ML is a strongly typed functional programming language, which has been used by a number of people for serious work during the last few years [1]. At the same time HOPE, designed by Rod Burstall and his group, has been similarly used [2]. The original DEC-10 ML was incomplete in some ways, redundant in others. Some of these inadequacies were remedied by Cardelli in his VAX version; others could be put right by importing ideas from HOPE.

In April '83, prompted by Bernard Sufrin, I wrote a tentative proposal to consolidate ML, and while doing so became convinced that this consolidation was possible while still keeping its character. The main strengthening came from generalising the "varstructs" of ML - the patterns of formal parameters - to the patterns of HOPE, which are extendible by the declaration of new data types. Many people immediately discussed the initial proposal. It was extremely lucky that we managed to have several separate discussions, in large and small groups, in the few succeeding months; we could not have chosen a better time to do the job. Also, Luca Cardelli very generously offered to freeze his detailed draft ML manual [3] until this proposal was worked out.

The proposal went through a second draft, on which there were further discussions. The results of these discussions were of two kinds. First, it became clear that two areas were still contentious: input/output and facilities for separate compilation. Second, many points were brought up about the remaining core of the language, and these were almost all questions of fine detail. The conclusion was rather clear; it was obviously better to present at first a definition of a Core language without the two contentious areas. This course is further justified by the fact that the Core language appears to be almost completely unaffected by the choice of input/output primitives and of separate compilation constructs. Also, there are already strong and carefully considered proposals, from Cardelli and MacQueen respectively, on how to design these two vital facilities; together with the Core they will form a complete language definition which can be adopted in its entirety, while still leaving open the possibility of adopting only parts of it. But the strong hope is that the whole will be very widely accepted.

A third draft [4] of the Core language was discussed in detail in a three-day design meeting at Edinburgh in June '84, attended by nine of the people mentioned below; some final points were ironed out, and the present Standard is the outcome. The meeting also looked in detail at the MacQueen Modules proposal and the Cardelli input/output proposal, and agreed on the essentials of these facilities to be embodied soon in a working definition.

The main contributors to the proposed language, through their design work on ML and on HOPE, are:

Rod Burstall, Luca Cardelli, Michael Gordon, David MacQueen,  
Robin Milner, Lockwood Morris, Malcolm Newey, Christopher Wadsworth.

The final proposal also owes much to criticisms and suggestions from many other people: Guy Cousineau, Jim Hook, Gerard Huet, Robert Milne, Kevin Mitchell, Brian Monahan, Peter Mosses, Alan Mycroft, Larry Paulson, David Rydeheard, Don Sannella, David Schmidt, John Scott, Stefan Sokolowski, Bernard Sufrin, Philip Wadler. Most of them have expressed strong support for most of the design; any inadequacies which remain are my fault, but I have tried to represent the consensus.

## 1.2 Design principles

The proposed ML is not intended to be the functional language. There are too many degrees of freedom for such a thing to exist: lazy or eager evaluation, presence or absence of references and assignment, whether and how to handle exceptions, types-as-parameters or polymorphic type-checking, and so on. Nor is the language or its implementation meant to be a commercial product. It aims to be a means for propagating the craft of functional programming and a vehicle for further research into the design of functional languages.

The over-riding design principle is to restrict the Core language to ideas which are simple and well-understood, and also well-tried - either in previous versions of ML or in other functional languages (the main other source being HOPE, mainly for its argument-matching constructs). One effect of this principle has been the omission of polymorphic references and assignment. There is indeed an elegant and sound scheme for polymorphic assignment worked out by Luis Damas; unfortunately it is not yet documented, and we will do better to wait for a clear exposition either from Damas or - as promised - from David MacQueen. In the proposed language much can be done to get the polymorphic effect by passing assignment functions as parameters; it is worthwhile experimenting with this method, and there is further advantage in keeping to the simple polymorphic type-checking discipline which derives from Curry's Combinatory Logic via Hindley.

A second design principle is to generalise well-tried ideas where the generalisation is apparently natural. This has been applied in generalising ML "varstructs" to HOPE patterns, in broadening the structure of declarations (following Cardelli's declaration connectives which go back to Robert Milne's Ph.D. Thesis) and in allowing exceptions which carry values of arbitrary polymorphic type. It should be pointed out here that a difficult decision had to be made concerning HOPE's treatment of data types - present only in embryonic form in the original ML - and the labelled records and variants which Cardelli introduced in his VAX version. The latter have definite advantages which the former lack; on the other hand, the HOPE treatment is well-rounded in its own terms. Though a combination of these features is possible, it seemed (at least to me, but some disagreed!) to entail too rich a language for the present definition. Thus the HOPE treatment is fully adopted here. However, at the design meeting of June '84 it was agreed to experiment with at least two different ways of adding labelled records to the Core as a smooth extension, and to adopt one of these schemes as standard in the near future.

A third principle is to specify the language completely, so that programs will port between correct implementations with minimum fuss. This entails, first, precise concrete syntax (abstract syntax is in some senses more important - but we do not all have structure editors yet, and humans still communicate among themselves in concrete syntax!); second, it entails exact evaluation rules (e.g. we must specify the order of evaluation of two expressions, one applied to the other, just because of the exception mechanism). The present document is not a full language definition; the Core language will only become a full language when the proposals for input/output and for separate compilation are added.

### 1.3 An example

The following declaration illustrates some constructs of the Core Language. A longer expository paper should contain many more examples; here, we hope only to draw attention to some of the less familiar ideas.

The example sets up the abstract type 'a dictionary, in which each entry associates an item (of arbitrary type 'a) with a key (an integer). Besides the null dictionary, the operations provided are for looking up a key, and for adding a new entry which overrides any old entry with the same key. A natural representation is by a list of key-item pairs, ordered by key.

```
abstype 'a dictionary =
  data dict of (int * 'a)list           {dict is the abstraction}
                                         {      constructor.}

with
  val nulldict = dict nil                {The function lookup may}
                                         {  raise an exception.}

  exception lookup : unit                {

  val lookup (key:int)                   {'a is the result type. }
      (dict entrylist) : 'a =
    let val rec search nil = raise lookup {An auxiliary clausal }
      | search ((k,item)::entries) =    { function declaration.}
        if key=k then item
        else if key<k then raise lookup
        else search entries

    in search entrylist
  end

  val enter (newentry as (key,item))    {A layered pattern.   }
      (dict entrylist) : 'a dictionary. =
    let val rec update nil = [ newentry ] {A singleton list.   }
      | update ((entry as (k,_))::entries) =
        if key=k then newentry::entries
        else if key<k then newentry::entry::entries
        else entry::update entries

    in dict(update entrylist)
  end

end                                     {end of dictionary}
```

After the declaration is evaluated, five identifier bindings are reported, and recorded in the top-level environment. They consist of the type binding of dictionary, the exception binding of lookup, and three value bindings with their types:

```
nulldict : 'a dictionary
lookup : int -> 'a dictionary -> 'a
enter : int * 'a -> 'a dictionary -> 'a dictionary
```

The layered pattern construct "as" was first introduced in HOPE, and yields both brevity and efficiency. The discerning reader may be able to find one further use for it in the declaration.

**Note:** the abstype construct is in the Core language for completeness, but is likely to be subsumed by Modules.

## 2. The bare language

### 2.1 Discussion

It is convenient to present the language first in a bare form, containing enough on which to base the semantic description given in Section 3. Things omitted from the bare language description are:

- (1) Derived syntactic forms, whose meaning derives from their equivalent forms in the bare language (Section 6);
- (2) Directives for introducing infix identifier status (Section 4);
- (3) Standard bindings (Section 5);
- (4) References and equality (Section 7);
- (5) Type-checking (Section 9).

The principal syntactic objects are expressions and declarations. The composite expression forms are application, type constraint, tupling, raising and handling exceptions, local declaration (using let) and function abstraction.

Another important syntactic class is the class of patterns; these are essentially expressions containing only variables and value constructors, and are used to create value bindings. Declarations may declare value variables (using value bindings), types with associated constructors or operations (using type bindings), and exceptions (using exception bindings). Apart from this, one declaration may be local to another (using local), and a sequence of declarations is allowed as a single declaration.

An ML program is a series of declarations, called top-level declarations,

```
dec1 ; .. decn ;
```

each terminated by a semicolon (where each `deci` is not itself of the form "`dec ; dec'`"). In evaluating a program, the bindings created by `dec1` are reported before `dec2` is evaluated, and so on. In the complete language, an expression occurring in place of any `deci` is an abbreviated form (see Section 6.2) for a declaration binding the expression value to the variable "it"; such expressions are called top-level expressions.

The bare syntax is in Section 2.8 below; first we consider lexical matters.

### 2.2 Reserved words

The following are the reserved words used in the Core language. They may not (except =) be used as identifiers. In this document the alphabetic reserved words are always underlined>.

```
abstype and andalso as case do data else
end exception fun handle if in infix
infixr let local nonfix of op orelse
raise rec then type val with while

( ) [ ] , : ; | || = => _ ?
```

### 2.3 Special constants

The unique object of type unit is denoted by the special constant ().

An integer constant is any non-empty sequence of digits, possibly preceded by a negation symbol (~).

A real constant is an integer constant, possibly followed by a point (.) and one or more digits, possibly followed by an exponent symbol (E) and an integer constant; at least one of the optional parts must occur, hence no integer constant is a real constant. Examples: 0.7 , ~3.32E5 , 3E~7 . Non-examples: 23 , .3 , 4.E5 , 1E2.0 .

A string constant is a sequence, between quotes ("), of zero or more printable characters, spaces or escape sequences. Each escape sequence is introduced by the escape character \, and stands for a character sequence. The allowed escape sequences are as follows (all other uses of \ being incorrect):

\n	A single character interpreted by the system as end-of-line.
\t	Tab.
~c	The control character c, for any appropriate c.
\ddd	The single character with ASCII code ddd (3 decimal digits).
\"	"
\\	\
\f..f\	This sequence is ignored, where f..f stands for a sequence of one or more formatting characters (a subset of the non-printable characters including at least space, tab, newline, formfeed). This allows one to write long strings on more than one line, by writing \ at the end of one line and at the start of the next.

### 2.4 Identifiers

Identifiers are used to stand for five different syntax classes which, if we had a large enough character set, would be disjoint:

value variables	{var}	type variables	{tyvar}
value constructors	{con}	type constructors	{tycon}
	exception names	{exn}	

An identifier is either alphanumeric: any sequence of letters, digits, primes (') and underbars (\_) starting with a letter or prime, or symbolic: any sequence of the following symbols

! % & \$ + - / : < = > ? @ \ ~ ` ^ | \*

In either case, however, reserved words are excluded. This means that for example ? and | are not identifiers, but ?? and |=| are identifiers. The only exception to this rule is that the symbol =, which is a reserved word, is also allowed as an identifier to stand for the equality predicate (see Section 7.2). The identifier = may not be rebound; this precludes any syntactic ambiguity.

A type variable {tyvar} may be any alphanumeric identifier starting with a prime. The other four classes {var, con, tycon, exn} are represented by identifiers not starting with a prime. Thus type variables are disjoint from the other four classes. Otherwise, the syntax class of an occurrence of identifier id is determined thus:

- (1) In types, `id` is a type constructor, and must be within the scope of the type binding which introduced it.
- (2) Following exception, raise or handle, or in the context "exception `..=id`", `id` is an exception name.
- (3) Elsewhere, `id` is a value constructor if it occurs in the scope of a type binding which introduced it as such, otherwise it is a value variable.

It follows from (3) that no value binding can make a hole in the scope of a value constructor by introducing the same identifier as a variable; this is because, in the scope of the declaration which introduces `id` as a value constructor, any occurrence of `id` in a pattern is interpreted as the constructor and not as the binding occurrence of a new variable.

The syntax-classes `var`, `con`, `tycon` and `exn` all depend on which bindings are in force, but only the classes `var` and `con` are necessarily disjoint. The context determines (as described above) to which class each identifier occurrence belongs.

In the Core language, an identifier may be given infix status by the infix or infixr directive; this status only pertains to its use as a `var` or a `con`. If `id` has infix status, then "`exp1 id exp2`" (resp. "`pat1 id pat2`") may occur wherever the application "`id(exp1,exp2)`" (resp. "`id(pat1,pat2)`") would otherwise occur. On the other hand, non-infix occurrences of `id` must be prefixed by the keyword op. Infix status is cancelled by the nonfix directive.

## 2.5 Comments

A comment is any character sequence within curly brackets `{}` in which curly brackets are properly nested. An unmatched `}` should be detected by the compiler.

## 2.6 Lexical analysis

Each item of lexical analysis is either a reserved word or a special constant or an identifier; comments and formatting characters separate items (except within string constants; see Section 2.3) and are otherwise ignored. At each stage the longest next item is taken.

As a consequence of this simple approach, spaces - or parentheses - are needed sometimes to separate identifiers and reserved words. Two examples are

<code>a := !b</code>	or	<code>a := (!b)</code>	but not	<code>a := !b</code>
				(assigning contents of <code>b</code> to <code>a</code> )
<code>~ :int-&gt;int</code>	or	<code>(~):int-&gt;int</code>	but not	<code>~:int-&gt;int</code>
				(unary minus qualified by its type)

Rules which allow omission of spaces in such examples, such as adopted by Cardelli in VAX ML, also forbid certain symbol sequences as identifiers and - more importantly - are hard to remember; it seems better to keep a simple scheme and tolerate a few extra spaces or parentheses.

## 2.7 Delimiters

Not all constructs have a terminating reserved word; this would be verbose. But a compromise has been adopted; end terminates any construct which declares bindings with local scope. This involves only the let, local and abstype constructs.

## 2.8 The bare syntax

### Conventions

- (1) {...} means optional.
- (2) For any syntax class  $s$ , define  $s\_seq ::= s (s_1, \dots, s_n) (n \geq 1)$
- (3) Alternatives are in order of decreasing precedence.
- (4) L (resp. R) means left (resp. right) association.
- (5) Parentheses may enclose phrases of any named syntax class defined in the table.

<p style="text-align: center;"><u>EXPRESSIONS</u>    exp</p> <pre> aexp ::=   var           (variable)   con           (constructor)   [ exp1 , ... , expn ] (tuple, n ≥ 2)   [ exp ]  exp ::=   aexp           (atomic)   exp aexp       L(application)   exp : ty       L(constraint)   raise exn with exp (raise exc'n)   let dec in exp end (local dec'n)   fun match      (function)   exp handle handler R(handle exc'ns)  match ::=   rule1   ...   rulen (n ≥ 1)  rule ::=   pat =&gt; exp  handler ::=   hrule1    ...    hrulen (n ≥ 1)  hrule ::=   exn with match   ? =&gt; exp </pre>	<p style="text-align: center;"><u>PATTERNS</u>    pat</p> <pre> apat ::=   _             (wildcard)   var          (variable)   con          (constant)   [ pat1 , ... , patn ] (tuple, n ≥ 2)   [ pat ]  pat ::=   apat           (atomic)   con apat       L(construction)   pat : ty       L(constraint)   var {:ty} as pat (layered) </pre>
<p style="text-align: center;"><u>DECLARATIONS</u>    dec</p> <pre> dec ::=   val vb           (values)   type tb         (types)   abstype tb   with dec end    (abs. types)   exception eb    (exceptions)   local dec in dec' end (local dec'n)   dec1 {;} .. decn {;} (sequence, n ≥ 0) </pre>	<p style="text-align: center;"><u>VALUE BINDINGS</u>    vb</p> <pre> vb ::=   pat = exp       (simple)   vb1 and ..and vbn (multiple, n ≥ 2)   rec vb         (recursive) </pre>
<p style="text-align: center;"><u>PROGRAMS</u> :    dec1 ; .. decn ;</p>	<p style="text-align: center;"><u>TYPE BINDINGS</u>    tb</p> <pre> tb ::=   {tyvar_seq} tycon   = data constrs (simple)   {tyvar_seq} tycon   = ty           (simple)   tb1 and ..and tbn (multiple, n ≥ 2)   rec tb        (recursive) </pre>
	<p style="text-align: center;"><u>CONSTRAINTS</u>    constrs</p> <pre> constrs ::=   con1 {of ty1}   ...   conn {of tyn} </pre>
	<p style="text-align: center;"><u>EXCEPTION BINDINGS</u>    eb</p> <pre> eb ::=   exn {:ty} [= exn'] (simple)   eb1 and ..and ebn (multiple, n ≥ 2) </pre>
	<p style="text-align: center;"><u>TYPES</u>    ty</p> <pre> ty ::=   tyvar           (type variable)   {ty_seq} tycon (type constr'n)   ty1 * ..* tyn  (tuple type, n ≥ 2)   ty -&gt; ty'      R(function type) </pre>

The syntax of types binds more tightly than that of expressions, so type constraints should be parenthesized if not followed by a reserved word.

Each iterated construct (match, handler, ..) extends as far right as possible; thus e.g. a match within a match may need to be parenthesised.



### 3. Evaluation

#### 3.1 Environments and Values

Evaluation of phrases takes place in the presence of an ENVIRONMENT and a STORE. An ENVIRONMENT E has two components: a value environment VE associating values to variables and to value constructors, and an exception environment EE associating exceptions to exception names. A STORE S associates values to references, which are themselves values. (A third component of an environment, a type environment TE, is ignored here since it is relevant only to type-checking and compilation, not to evaluation.)

An exception e, associated to an exception name exn in any exception environment, is an object drawn from an infinite set (the nature of e is immaterial, but see Section 3.8). A packet p=(e,v) is an exception e paired with a value v, called the excepted value. Neither exceptions nor packets are values. Besides possibly changing S (by assignment), evaluation of a phrase returns a result as follows:

<u>Phrase</u>	<u>Result</u>
Expression	v or p
Value binding	VE or p
Type binding	VE
Exception binding	EE
Declaration	E or p

For every phrase except a handle expression, whenever its evaluation demands the evaluation of an immediate subphrase which returns a packet p as result, no further evaluation of subphrases occurs and p is also the result of the phrase. This rule should be remembered while reading the evaluation rules below.

A function value f is a partial function which, given a value, may return a value or a packet; it may also change the store as a side-effect. Every other value is either a constant (a nullary constructor), a construction (a constructor with a value), a tuple or a reference.

#### 3.2 Environment manipulation

We may write <(id1,v1) ..(idn,vn)> for a value environment VE (the idi being distinct). Then VE(idi) denotes vi, <> is the empty value environment, and VE+VE' means the value environment in which the associations of VE' supersede those of VE. Similarly for exception environments. If E=(VE,EE) and E'=(VE',EE'), then E+E' means (VE+VE',EE+EE'), E+VE' means E+(VE',<>), etc. This implies that an identifier may be associated both in VE and in EE without conflict.

#### 3.3 Matching patterns

The matching of a pattern pat to a value v either fails or yields a value environment. Failure is distinct from returning a packet, but a packet will be returned when all patterns fail in applying a match to a value (see Section 3.4). In the following rules, if any component pattern fails to match then the whole pattern fails to match.

The following is the effect of matching a pattern `pat` to a value `v`, in each of the cases for `pat`:

`_` : the empty value environment `<>` is returned.  
`var` : the value environment `<(var,v)>` is returned.  
`con{pat}` : if `v = con{v'}` then `pat` is matched to `v'`, else failure.  
`var{:ty} as pat` : `pat` is matched to `v` returning `VE`; then `<(var,v)>+VE` is returned.  
`(pat1, ...,patn)` : if `v=(v1,...,vn)` then `pati` is matched to `vi` returning `VEi`, for each `i`; then `VE1+ ..+VEn` is returned.  
`pat:ty` : `pat` is matched to `v`.

### 3.4 Applying a match

Assume environment `E`. Applying a match `pat1=>exp1! ..!patn=>expn` to value `v` returns a value or packet as follows:

Each `pati` is matched to `v` in turn, from left to right, until one succeeds returning `VEi`; then `expi` is evaluated in `E+VEi`. If none succeeds, then the packet `[ematch,()]` is returned, where `ematch` is the standard exception bound by predeclaration to the exception name "match". But matches which may fail are to be detected by the compiler and flagged with a warning; see Section 10(2).

Thus, for each `E`, a match denotes a function value.

### 3.5 Evaluation of expressions

Assume environment `E=(VE,EE)`. Evaluating an expression `exp` returns a value or packet as follows, in each of the cases for `exp`:

`var` : the value `VE(var)` is returned.  
`con` : the value `VE(con)` is returned.  
`exp aexp` : `exp` is evaluated, returning function value `f`; then `aexp` is evaluated, returning value `v`; then `f(v)` is returned.  
`(exp1, ...,expn)` : the `expi` are evaluated in sequence, from left to right, returning `vi` respectively; then `(v1, ...,vn)` is returned.  
`raise exn with exp` : `exp` is evaluated, returning value `v`; then packet `(e,v)` is returned, where `e = EE(exn)`.  
`exp handle handler` : `exp` is evaluated; if `exp` returns a value `v`, then `v` is returned; if it returns a packet `p = (e,v)` then the handling rules of the handler are scanned from left to right until a rule is found which satisfies one of two conditions:  
 (1) it is of form "`exn with match`" and `e=EE(exn)`,

in which case match is applied to v;  
(2) it is of form "? => exp'", in which case exp'  
is evaluated.

If no such hrule is found, then p is returned.

let dec in exp end : dec is evaluated, returning E'; then exp is  
evaluated in E+E'.

fun match : f is returned, where f is the function of v gained  
by applying match to v in environment E.

exp:ty : exp is evaluated.

### 3.6 Evaluation of value bindings

Assume environment E = (VE,EE). Evaluating a value binding vb returns a  
value environment VE' or a packet as follows, by cases of vb:

pat = exp : exp is evaluated in E, returning value v; then pat is  
matched to v; if this returns VE', then VE' is returned,  
and if it fails then the packet (ebind,()) is returned, where  
ebind is the standard exception bound by predeclaration to  
the exception name "bind".

vb1 and ..and vbn : vb1, ..,vbn are evaluated in E from left to right, returning  
VE1, ..,VEN; then VE1+ ..+VEN is returned.

rec vb : vb is evaluated in E', returning VE', where E' = (VE+VE',EE).  
Because the values bound by "rec vb" must be function values  
(see 10(4)), E' is well defined by "tying knots" (Landin).

### 3.7 Evaluation of type bindings

The components VE and EE of the current environment do not affect the  
evaluation of type bindings (TE affects their type-checking and compilation).  
Evaluating a type binding tb returns a value environment VE' (it cannot return a  
packet) as follows, by cases of tb:

{tyvar\_seq} tycon = data con1 {of ty1} | .. | conn {of tyn} :  
the value environment VE' = <(con1,v1), ..,(conn,vn)> is  
returned, where vi is either the constant value con1 (if  
"of tyi" is absent) or else the function which maps v to  
con1(v). Other effects of this type binding are handled  
by the compiler or type-checker, not by evaluation.

{tyvar\_seq} tycon = ty :  
the value environment VE' = <> is returned. This type  
binding has no effect on evaluation; its purpose, in the  
Core language, is merely to provide an abbreviation for  
a compound type. It may not be qualified by rec.

tb1 and ..and tbn : tb1, ..,tbn are evaluated from left to right, returning  
VE1, ..,VEN; then VE' = VE1+ ..+VEN is returned.

rec tb : tb is evaluated. Note again that the recursion is  
handled by type-checking only..

### 3.8 Evaluation of exception bindings

Assume environment  $E = (VE, EE)$ . The evaluation of an exception binding  $eb$  returns an exception environment  $EE'$  as follows, by cases of  $eb$ :

$exn \{ :ty \} \{ = exn' \}$  :  $EE' = \langle (exn, e) \rangle$  is returned, where

- (1) if  $exn'$  is present then  $e = EE(exn')$ ; this is a non-generative exception binding since it merely re-binds an existing exception to  $exn$ ;
- (2) otherwise  $e$  is a previously unused exception (an object from which the identifier  $exn$  is retrievable, for reporting unhandled exceptions at top-level); this is a generative exception binding.

$eb1$  and .. and  $ebn$  :  $eb1, \dots, ebn$  are evaluated in  $E$  from left to right, returning  $EE1, \dots, EE_n$ ; then  $EE' = EE1 + \dots + EE_n$  is returned.

### 3.9 Evaluation of declarations

Assume environment  $E = (VE, EE)$ . Evaluating a declaration  $dec$  returns an environment  $E'$  or a packet as follows, by cases of  $dec$ :

val  $vb$  :  $vb$  is evaluated, returning  $VE'$ ; then  $E' = (VE', \langle \rangle)$  is returned.

type  $tb$  :  $tb$  is evaluated, returning  $VE'$ ; then  $E' = (VE', \langle \rangle)$  is returned.

abstype  $tb$  with  $dec$  end :  
 $tb$  is evaluated, returning  $VE'$ ; then  $dec$  is evaluated in  $E + VE'$ , returning  $E'$ ; then  $E'$  is returned.

exception  $eb$  :  $eb$  is evaluated, returning  $EE'$ ; then  $E' = (\langle \rangle, EE')$  is returned.

local  $dec1$  in  $dec2$  end :  
 $dec1$  is evaluated, returning  $E1$ , then  $dec2$  is evaluated in  $E + E1$ , returning  $E2$ ; then  $E' = E2$  is returned.

$dec1 \{ ; \} \dots decn \{ ; \}$  :  
each  $dec_i$  is evaluated in  $E + E1 + \dots + E_{i-1}$ , returning  $E_i$ , for  $i = 1, 2, \dots, n$ ; then  $E' = (\langle \rangle, \langle \rangle) + E1 + \dots + E_n$  is returned. Thus when  $n=0$  the empty environment is returned.

Each declaration is defined to return only the new environment which it makes, but the effect of a declaration sequence is to accumulate environments.

### 3.10 Evaluation of programs

The evaluation of a program " $dec1 ; \dots decn ;$ " takes place in the initial presence of the standard top-level environment  $ENV_0$  containing all the standard bindings (see Section 5). For  $i > 0$  the top-level environment  $ENV_i$ , present after the evaluation of  $dec_i$  in the program, is defined recursively as follows:  $dec_i$  is evaluated in  $ENV_{i-1}$  returning environment  $E_i$ , and then  $ENV_i = ENV_{i-1} + E_i$ .

#### 4. Directives

Directives are included in ML as (syntactically) a subclass of declarations. They possess scope, as do all declarations.

There is only one kind of directive in the standard language, namely those concerning the infix status of value variables and constructors. Others, perhaps also concerned with syntactic conventions, may be included in extensions of the language. The directives concerning infix status are:

```
infix{r} {p} id1 ..idn  
nonfix id1 ..idn
```

where  $p$  is a non-negative integer. The infix and infixr directives introduce infix status for each  $id_i$  (as a value variable or constructor), and the nonfix directive cancels it. The integer  $p$  (default 0) determines the precedence, and an infix identifier associates to the left if introduced by infix, to the right if by infixr. Different infix identifiers of equal precedence associate to the left. As indicated in Appendix 2, the precedence of infix application is just weaker than that of application.

While  $id$  has infix status, each occurrence of it (as a value variable or constructor) must be infix or else preceded by op. Note that this includes occurrences of the identifier within patterns, even binding occurrences of variables.

Several standard functions and constructors have infix status (see Appendix 3) with precedence; these are all left associative except "::<".

It may be thought better that the infix status of a variable or constructor should be established in some way within its binding occurrence, rather than by a separate directive. However, the use of directives avoids problems in parsing.

The use of local directives (introduced by let or local) imposes on the parser the burden of determining their textual scope. A quite superficial analysis is enough for this purpose, due to the use of end to delimit local scopes.

## 5. Standard bindings

The bindings of this section form the standard top-level environment ENVO.

### 5.1 Standard type constructors

The bare language provides the function-type constructor,  $\rightarrow$ , and for each  $n \geq 2$  a tuple-type constructor  $*n$ . Type constructors are in general postfix in ML, but  $\rightarrow$  must be infix, and the  $n$ -ary tuple-type constructed from  $ty_1, \dots, ty_n$  must be written " $ty_1 * \dots * ty_n$ ". Besides these type constructors, the following are standard:

```
Type constants (nullary constructors) : unit, bool, int, real, string
Unary type constructors                : list, ref
```

None of the identifiers  $\rightarrow$ ,  $*$ , unit, bool, int, real, string, list, ref may be redeclared as type constructors.

The constructors unit, bool and list are fully defined by the following assumed declaration

```
infixr 30  ::
type unit = data ()
and bool = data true | false
and rec 'a list = data nil | op :: of 'a * 'a list
```

The word "unit" is chosen since the type contains just one value; this is why it is preferred to the word "void" of ALGOL 68. Note that it is also (up to isomorphism) a unit for type tupling, though we do not exploit this isomorphism by allowing a coercion between the types  $ty$  and  $ty * unit$ .

The type constants int, real and string are equipped with special constants as described in Section 2.3. The type constructor ref is for constructing reference types; see Section 7.

### 5.2 Standard functions and constants

All standard functions and constants are listed in Appendix 3. There is not a lavish number; we envisage function libraries provided by each implementation, together with the equivalent ML declaration of each function (though the implementation may be more efficient). In time, some such library functions may accrue to the standard; a likely candidate for this is a group of array-handling functions, grouped in a standard declaration of the unary type constructor "array".

Most of the standard functions and constants are familiar, so we need mention only a few critical points:

- (1) explode yields a list of strings of size 1; implode is iterated string concatenation (^). ord yields the Ascii code number of the first character of a string; chr yields the Ascii character (as a string of size 1) corresponding to an integer.
- (2) ref is a monomorphic function, but in patterns it may be used polymorphically, with type 'a  $\rightarrow$  'a ref.

- (3) The character functions `ord` and `chr`, the arithmetic operators `*`, `/`, `div`, `mod`, `+` and `-`, and the standard functions `floor`, `sqrt`, `exp` and `ln` may raise standard exceptions (see Section 5.3) whose name in each case is the same as that of the function. This occurs for `ord` when the string is empty; for `chr` when the character is undefined; and for the others when the result is undefined or out of range.
- (4) The values  $r = a \bmod d$  and  $q = a \operatorname{div} d$  are determined by the condition  $d \cdot q + r = a$ , where either  $0 \leq r < d$  or  $d < r \leq 0$ . Thus the remainder takes the same sign as the divisor, and has lesser magnitude. The result of `arctan` lies between  $\pm\pi/2$ , and `ln` (the inverse of `exp`) is the natural logarithm. The value `floor(x)` is the largest integer  $\leq x$ ; thus rounding may be done by `floor(x+0.5)`.
- (5) Two multi-typed functions are included as quick debugging aids. The function `print :ty->ty` is an identity function, which as a side-effect prints its argument exactly as it would be printed at top-level. The printing caused by "print(exp)" will depend upon the type ascribed to this particular occurrence of `exp`; thus `print` is not a normal polymorphic function. The function `makestring :ty->string` is similar, but instead of printing it returns as a string what `print` would produce on the screen.

### 5.3 Standard exceptions

All predeclared exception names are of type `unit`. There are three special ones: `match`, `bind` and `interrupt`. These exceptions are raised, respectively, by failures of matching and binding as explained in Sections 3.4 and 3.6, and by an interrupt generated (often by the user) outside the program. Note, however, that `match` and `bind` exceptions cannot occur unless the compiler has given a warning, as detailed in Section 10(2),(3), except in the case of a top-level declaration as indicated in 10(3).

The only other predeclared exception names are

`ord chr * / div mod + - floor sqrt exp ln`

Each name identifies the corresponding standard function, which is ill-defined or out of range for certain arguments, as detailed in Section 5.2. For example, using the derived handle form explained in Section 8.2, the expression

`3 div x handle div => 10000`

will return 10000 when  $x = 0$ .

## 6. Standard Derived Forms

### 6.1 Expressions and Patterns

DERIVED FORM	EQUIVALENT FORM
<u>Expressions</u> :	
<u>raise</u> <i>exn</i>	<u>raise</u> <i>exn with</i> ()
<u>case</u> <i>exp of</i> <i>match</i>	( <u>fun</u> <i>match</i> )( <i>exp</i> )
<u>if</u> <i>exp then</i> <i>exp1 else</i> <i>exp2</i>	<u>case</u> <i>exp of</i> <i>true=&gt;exp1   false=&gt;exp2</i>
<i>exp</i> <u>or</u> <i>else</i> <i>exp'</i>	<u>if</u> <i>exp then</i> <i>true else</i> <i>exp'</i>
<i>exp</i> <u>and</u> <i>also</i> <i>exp'</i>	<u>if</u> <i>exp then</i> <i>exp' else</i> <i>false</i>
<i>exp</i> ; <i>exp'</i>	<u>case</u> <i>exp of</i> ( <u>_</u> ) => <i>exp'</i>
<u>while</u> <i>exp do</i> <i>exp'</i>	<u>let val</u> <i>rec</i> <i>f = fun</i> () => <u>if</u> <i>exp then</i> ( <i>exp'</i> ; <i>f</i> ()) <u>else</u> () <u>in</u> <i>f</i> () <u>end</u>
[ <i>exp1</i> , ... , <i>expn</i> ]	<i>exp1:: ...::expn::nil</i> (n>=0)
<u>Handling rules</u> :	
<i>exn</i> => <i>exp</i>	<i>exn with</i> ( <u>_</u> ) => <i>exp</i>
<u>Patterns</u> :	
[ <i>pat1</i> , ... , <i>patn</i> ]	<i>pat1:: ...::patn::nil</i> (n>=0)

The derived form may be implemented more efficiently than its equivalent form, but must be precisely equivalent to it semantically. The type-checking of each derived form is also defined by that of its equivalent form.

The binding power of all bare and derived forms is shown in Appendix 1. A semicolon, whether used in declaration sequencing or in expression sequencing, always has weakest binding power; also a semicolon always terminates a declaration where this is syntactically possible (thus expression sequencing may need to be parenthesised).

The shortened raise form is only admissible with exceptions of type unit. The shortened form of handling rule is appropriate whenever the excepted value is immaterial, and is therefore (in the full form) matched to the wildcard pattern.



## 6.2 Bindings and Declarations

### DERIVED FORM

### EQUIVALENT FORM

#### Value bindings :

<pre>var apat11 ..apat1n {:ty} = exp1   ..   var apatm1 ..apatmn {:ty} = expm</pre>	<pre>var = <u>fun</u> x1 =&gt; ..<u>fun</u> xn =&gt;       <u>case</u> (x1, .., xn)       <u>of</u> (apat11, .., apat1n) =&gt; exp1 {:ty}         ..         (apatm1, .., apatmn) =&gt; expm {:ty}        { where the xi are new, and m,n≥1 }</pre>
---	---

#### Declarations :

<pre>exp</pre>	<pre><u>val</u> it = exp</pre>
----------------	--------------------------------

The derived value binding allows function definitions, possibly Curried, with several clauses. The derived declaration is only allowed at top-level, for treating top-level expressions as degenerate declarations; "it" is just a normal value variable.

## 7. References and equality

### 7.1 References and assignment

Following Cardelli, references are provided by the type constructor "ref". Since we are sticking to monomorphic references, there are two overloaded functions available at all monotypes mty:

- (1) `ref : mty -> mty ref`, which associates (in the store) a new reference with its argument value. "ref" is a constructor, and may be used polymorphically in patterns, with type `'a -> 'a ref`.
- (2) `op := : mty ref * mty -> unit`, which associates its first (reference) argument with its second (value) argument in the store, and returns () as result.

The polymorphic contents function "!" is provided, and is equivalent to the declaration `"val !(ref x) = x"`.

### 7.2 Equality

The overloaded equality function `op = : ety * ety -> bool` is available at all types ety which admit equality, according to the definition below. The effect of this definition is that equality will only be applied to values which are built up from references (to arbitrary values) by value constructors, including of course constant values. On references, equality means identity; on objects of other types ety, it is defined recursively in the natural way.

The types which admit equality are as follows, assuming that abbreviations introduced by non-generative type bindings have first been expanded out:

- (1) A type ty admits equality iff it is built from arbitrary reference types by type constructors which admit equality.
- (2) The standard type constructors \*n, unit, bool, int, real, string and list all admit equality.

Thus for example, the type `(int * 'a ref)list` admits equality, but `(int * 'a)list` and `(int -> bool)list` do not.

A user-defined type constructor tycon, declared by a generative type binding tb whose form is

```
{tyvar_seq} tycon = data con1 {of ty1} | .. | conn {of tyn}
```

admits equality within its scope (but, if declared by abstype, only within the with part of its declaration) iff it satisfies the following condition:

- (3) Each construction type tyi in this binding is built from arbitrary reference types and type variables, either by type constructors which already admit equality or (if tb is within a rec) by tycon or any other type constructor declared by mutual recursion with tycon, provided these other type constructors also satisfy the present condition.

The first paragraph of this section should be enough for an intuitive understanding of the types which admit equality, but the precise definition is given in a form which is readily incorporated in the type-checking mechanism.

## 8. Exceptions

### 8.1 Discussion

Some discussion of the exception mechanism is needed, as it goes a little beyond what exists in other functional languages. It was proposed by Alan Mycroft, as a means to gain the convenience of dynamic exception trapping without risking violation of the type discipline (and indeed still allowing polymorphic exception-raising expressions). Brian Monahan put forward a similar idea. Don Sannella also contributed, particularly to the nature of the derived forms (Section 8.2); these forms give a pleasant way of treating standard exceptions, as explained in Section 5.3.

The rough and ready rule for understanding how exceptions are handled is as follows. If an exception is raised by a raise expression

raise exn with exp

which lies in the textual scope of a declaration of the exception name exn, then it may be handled by a handling rule

exn with match

in a handler, but only if this handler is in the textual scope of the same declaration. Otherwise it may only be caught by the universal handling rule

? => exp' .

This rule is perfectly adequate for exceptions declared at top level; some examples in Section 8.4 below illustrate what may occur in other cases.

### 8.2 Derived forms

A handler discriminates among exception packets in two ways. First, it handles just those packets (e,v) for which e is the exception bound to the exception name in one of its handling rules; second, the match in this rule may discriminate upon v, the excepted value. Note however that, if a universal handling rule "? => exp'" is activated, then all packets are handled without discrimination. Thus "?" may be considered as a wildcard, matching any packet. It should be used with some care, bearing in mind that it will even handle interrupts.

A case which is likely to be frequent is when discrimination is required upon the exception, but not upon the excepted value; in this case, the derived handling rule

exn => exp'

is appropriate for handling. Further, exceptions of type unit may be raised by the shortened form

raise exn

since the only possible excepted value is ().

### 8.3 An example

To illustrate the generality of exception handling, suppose that we have declared some exceptions as follows:

```
exception oddlist :int list and oddstring :string
```

and that a certain expression `exp:int` may raise either of these exceptions and also runs the risk of dividing by zero. The handler in the following handle expression would deal with these exceptions:

```
exp handle oddlist with [] => 0
      | [x] => 2*x
      | x::y::_ => x·div y
  || oddstring with "" => 0
      | s => size(s)-1
  || div => 10000
```

Note that the whole expression is well-typed because in each handling rule the type of each match-pattern is the same as the exception type, and because the result type of each match is `int`, the same as the type of `exp`. The last handling rule is the shortened form appropriate for exceptions of type `unit`.

Note also that the last handling rule will handle `div` exceptions raised by `exp`, but will not handle the `div` exception which may be raised by "`x div y`" within the first handling rule. Finally, note that a universal handling rule

```
|| ? => 50000
```

at the end would deal with all other exceptions raised by `exp`.

### 8.4 Some pathological examples

We now consider some possible misuses of exception handling, which may arise from the fact that exception declarations have scope, and that each evaluation of a generative exception binding creates a distinct exception. Consider a simple example:

```
exception exn : bool;
val f(x) =
  let exception exn:int in
    if x > 100 then raise exn with x else x+1
  end;
f(200) handle exn with true=>500 | false=>1000;
```

The program is well-typed, but useless. The exception bound to the outer `exn` is distinct from that bound to the inner `exn`; thus the exception raised by `f(200)`, with expected value 200, could only be handled by a handler within the scope of the inner exception declaration - it will not be handled by the handler in the program, which expects a boolean value. So this exception will be reported at top level. This would apply even if the outer exception declaration were also of type `int`; the two exceptions bound to `exn` would still be distinct.

On the other hand, if the last line of the program is changed to

```
f(200) handle ? => 500 ;
```

then the exception will be caught, and the value 500 returned. A universal handling rule (i.e. containing "?") catches any exception packet, even one exported from the scope of the declaration of the associated exception name, but cannot examine the excepted value in the packet, since the type of this value cannot be statically determined.

Even a single textual exception binding - if for example it is declared within a recursively defined function - may bind distinct exceptions to the same identifier. Consider another useless program:

```
val rec f(x) =  
  let exception exn in  
    if p(x) then a(x) else  
    if q(x) then f(b(x)) handle exn with c(x)  
    else raise exn with d(x)  
  end;  
f(v);
```

Now if p(v) is false but q(v) is true, the recursive call will evaluate f(b(v)). Then, if both p(b(v)) and q(b(v)) are false, this evaluation will raise an exn exception with excepted value d(b(v)). But this packet will not be handled, since the exception of the packet is that which is bound to exn by the inner - not outer - evaluation of the exception declaration.

These pathological examples should not leave the impression that exceptions are hard to use or to understand. The rough and ready rule of Section 8.1 will almost always give the correct understanding.

## 9. Type-checking

The type-checking discipline is exactly as in original ML, and therefore need only be described with respect to new phrases.

In a match "pat1=>exp1 | .. | patn=>expn", the types of all pati must be the same (ty say), and if variable var occurs in pati then all free occurrences of var in expi must have the same type as its occurrence in pati. In addition, the types of all the expi must be the same (ty' say). Then ty->ty' is the type of the match. The type of "fun match" is the type of the match.

The type of a handler rule "exn with match" is ty', where exn has type ty and match has type ty->ty'. The type of a universal handling rule "? => exp" is the type of exp. The type of a handler is the type of all its handling rules (which must therefore be the same), and the type of "exp handle handler" is that of both exp and handler. The type of "raise exn with exp" is arbitrary, but exp and exn must have the same type. The type of an exception may be polymorphic; any exn is required to have the same type at all occurrences within the scope of its declaration (and this must be an instance of any type qualifying the declaration).

A type variable is only explicitly bound (in the sense of variable-binding in lambda-calculus) by its occurrence in the tyvar\_seq on the left hand side of a simple type binding "{tyvar\_seq} tycon = ..", and then its scope is the right hand side. (This means for example that bound uses of 'a in both tb1 and tb2 in the type binding "tb1 and tb2" bear no relation to each other.) Otherwise, repeated occurrences of a (free) type variable may serve to link explicit type constraints. The scope of such a type variable is the top-level declaration or expression in which it occurs. In a type-constraint "exp:ty" or "pat:ty" the type-checker must ascribe to exp or to pat a type which is an instance of ty; if this instance is less general than ty then the compiler should issue a warning (but still compile).

The first form of simple type binding "{tyvar\_seq} tycon = data .." is generative, since a new unique type constructor (denoted by tycon) is created by each textual occurrence of such a binding. The second form "{tyvar\_seq} tycon = ty", on the other hand, is non-generative; to take an example, the type binding "'a couple = 'a \* 'a" merely allows the type expression "ty couple" to abbreviate "ty \* ty" (for any ty) within its scope. There is no semantic significance in abbreviation; in the Core language it is purely for brevity, though in the proposed extension of ML to contain Modules non-generative type-bindings are likely to be essential in matching types or Signatures. However, the type-checker should take some advantage of non-local type abbreviations in reporting types at top-level; in doing this, it may need to choose sensibly between different possible abbreviations for the same type.

Some standard function symbols (e.g. =,+) stand for functions of more than one type; in these cases the type-checker should complain if it cannot determine from the context which is intended (an explicit type constraint may be needed). Note that there is no implicit coercion in ML, in particular from int to real; the conversion function real:int->real must be used explicitly.

The type-checker refers to the type environment (TE) component of the environment, and records its findings there. Details of TE are not given in this report; they are compatible with what is done in current ML implementations, except that value constructors (and their types) are associated with the type constructors to which they belong.

## 10. Syntactic restrictions

- (1) No pattern may contain two occurrences of the same variable. No binding may bind the same identifier twice.
- (2) In a match "pat1=>exp1 | ..| patn=>expn", the pattern sequence pat1, .., patn should be irredundant and exhaustive. That is, each patj must match some value (of the right type) which is not matched by pati for any i<j, and every value (of the right type) must be matched by some pati. The compiler must give warning on violation of this restriction, but should still compile the match. Thus the "match" exception (see Section 3.4) will only be raised for a match which has been flagged by the compiler. The restriction is inherited by derived forms; in particular, this means that in the Curried function binding "var apat1 ..apatn { :ty } = exp" (consisting of one clause only), each separate apati should be exhaustive by itself.
- (3) For each value binding "pat = exp" the compiler must issue a report (but still compile) if either pat is not exhaustive or pat contains no variable. This will (on both counts) detect a mistaken declaration like "val nil = exp" in which the user expects to declare a new variable nil (whereas the language dictates that nil is here a constant pattern, so no variable gets declared). Cardelli points out this danger.

However, these warnings should not be given when the binding is a component of a top-level declaration val vb ; e.g. "val x::l = exp1 and y = exp2" is not faulted by the compiler at top level, but may of course generate a "bind" exception (see Section 3.6).

- (4) For each value binding "pat = exp" within rec, exp must be of the form "fun match" (The derived form of value binding given in Section 6.2 necessarily obeys this restriction). Each type binding "{tyvar\_seq} tycon = .." within rec must be generative (i.e. include "data").
- (5) In the left hand side "{tyvar\_seq} tycon" of a simple type binding, the tyvar\_seq must contain no type variable more than once. The right hand side of a simple type binding may contain only the type variables mentioned on the left.
- (6) In "let dec in exp end" and "local dec in dec' end" no type constructor exported by dec may occur in the type of exp or in the type of any variable, value constructor or exception name exported by dec'.
- (7) Every global exception binding - that is, not localised either by let or by local - must be explicitly constrained by a monotype.
- (8) If, within the scope of a type constructor tycon, a type binding tb binds (simultaneously) one or more type constructors tycon1, .., tyconn then: (a) if the identifiers tyconi are all distinct from tycon, then their value constructors must also have identifiers distinct from those of tycon; (b) if any tyconi is the same identifier as tycon, then any value constructor of tycon may be re-bound as a value constructor for one of tycon1, .., tyconn, but is otherwise considered unbound (as a variable or value constructor) within the scope of tb, unless it is bound again therein. This constraint ensures that the scope of a type constructor is identical with the scopes of its associated value constructors, except that in an abstype declaration the scope of the value constructors is restricted to the with part.

## 11. Conclusion

This design has been under discussion for over a year, and the designers are confident in their understanding of it. However, it is only by extensive practice that a language is properly evaluated; there are probably a few infelicities of design from the practical point of view, and we expect these to emerge during the next year or so in the course of experience with implementation and use.

It would be reasonable after such a period to collect reactions and to publish a list of corrections - just those which can be agreed among several seriously concerned implementers and users.

Besides these corrections there will clearly be extensions - design ideas which use the present language as a platform. It will be important to keep these two developments separate as far as possible. Corrections should be few and preferably done at most once; extensions may be many, but need not impair the identity of the present language.

## REFERENCES:

- [1] M.Gordon, R.Milner and C.Wadsworth (1979) Edinburgh LCF. Springer-Verlag, Lecture Notes in Computer Science, Vol 78.
- [2] R.Burstall, D.MacQueen and D.Sannella (1980) HOPE: An Experimental Applicative Language. Report CSR-62-80, Computer Science Dept, Edinburgh University.
- [3] L.Cardelli (1982) ML under UNIX. Bell Laboratories, Murray Hill, New Jersey.
- [4] R.Milner (1983) A Proposal for Standard ML. Report CSR-157-83, Computer Science Dept, Edinburgh University.



## APPENDIX 1

### SYNTAX : EXPRESSIONS and PATTERNS (See Section 2.8 for conventions)

<pre>aexp ::=   {op} var   {op} con   ( exp1 , .. , expn )   [ exp1 , .. , expn ]   ( exp )</pre>	<pre>{variable} {constructor} {tuple, n≥2} {list, n≥0}</pre>
<pre>exp ::=   aexp   exp aexp   exp id exp'   exp : ty   exp <u>andalso</u> exp'   exp <u>orelse</u> exp'   <u>raise</u> exn {<u>with</u> exp}   <u>if</u> exp <u>then</u> exp1 <u>else</u> exp2   <u>while</u> exp <u>do</u> exp'   <u>let</u> dec <u>in</u> exp <u>end</u>   <u>case</u> exp <u>of</u> match   <u>fun</u> match   exp <u>handle</u> handler   exp ; exp'</pre>	<pre>{atomic} L{application}   {infix application} L{constraint}   {conjunction}   {disjunction}   {raise exception}   {conditional}   {iteration}   {local declaration}   {case expression}   {function}   R{handle exception}   {sequence}</pre>
<pre>match ::=   rule1   ..   rulen    (n≥1)</pre>	<pre>handler ::=   hrule1    ..    hrulen    (n≥1)</pre>
<pre>rule ::=   pat =&gt; exp</pre>	<pre>hrule ::=   exn <u>with</u> match   exn =&gt; exp   ? =&gt; exp</pre>
<pre>apat ::=   -   {op} var   con   ( pat1 , .. , patn )   [ pat1 , .. , patn ]   ( pat )</pre>	<pre>{wildcard} {variable} {constant} {tuple, n≥2} {list, n≥0}</pre>
<pre>pat ::=   apat   {op} con apat   pat con pat'   pat : ty   var { : ty } <u>as</u> pat</pre>	<pre>{atomic} L{construction}   {infix construction} L{constraint}   {layered}</pre>

The syntax of types binds more tightly than that of expressions, so type constraints should be parenthesised if not followed by a reserved word.

Each iterated construct (match, handler, ..) extends as far right as possible; thus e.g. a match within a match may need to be parenthesised.

## APPENDIX 2

### SYNTAX : TYPES, BINDINGS, DECLARATIONS and PROGRAMS (See Section 2.8 for conventions)

```

ty ::=
    tyvar                                (type variable)
    {ty_seq} tycon                       (type construction)
    ty1 * ..* tyn                        (tuple type, n≥2)
    ty1 -> ty2                           R(function type)

vb ::=
    pat = exp                            (simple)
    {op} var apat11 ..apat1n { :ty } = exp1 (clausal function) **
    | ..
    | {op} var apatm1 ..apatmn { :ty } = expn (m,n≥1)
    vb1 and ..and vbn                    (multiple, n≥2)
    rec vb                                (recursive)

tb ::=
    {tyvar_seq} tycon = data constrs    (simple, generative)
    {tyvar_seq} tycon = ty                (simple, non-generative)
    tb1 and ..and tbn                    (multiple, n≥2)
    rec tb                                (recursive)

constrs ::=
    con1 {of ty1} | ..| conn {of tyn}    (n≥1)

eb ::=
    exn { :ty } [= exn']                  (simple)
    eb1 and ..and ebn                    (multiple, n≥2)

dec ::=
    val vb                                (value declaration)
    type tb                                (type declaration)
    abstype tb with dec and             (abstract type declaration)
    exception eb                          (exception declaration)
    local dec in dec' and              (local declaration)
    exp                                    (top-level only)
    dir                                    (directive)
    dec1 {;} ..decn {;}                   (declaration sequence, n≥0)

dir ::=
    infix{r} {p} id1 ..idn                (declare infix status, p≥0)
    nonfix id1 ..idn                    (cancel infix status)

```

PROGRAMS: dec1 ; ..decn ;

\*\* If var has infix status then op is required in this form; alternatively -- var may be infix in any clause. Thus at the start of any clause:

op var {apat,apat'}      may be written :      {apat var apat'}

and the parentheses may also be dropped if ":ty" or "=" follows immediately.

## APPENDIX 3

### PREDECLARED VARIABLES and CONSTRUCTORS

In the types of these bindings, `num` stands for either `int` or `real` (the same in each type). Similarly `ty` stands for an arbitrary type, `mty` stands for any monotype, and `ety` (see Section 7.2) stands for any type admitting equality.

<u>nonfix</u>	<u>infix</u>
<code>nil</code> : 'a list	<u>Precedence 50:</u>
<code>map</code> : ('a->'b) -> 'a list	<code>/</code> : real * real -> real
	<code>div</code> : int * int -> int
	<code>mod</code> : " " "
<code>rev</code> : 'a list -> 'a list	<code>*</code> : num * num -> num
	-
<code>true,false</code> : bool	<u>Precedence 40:</u>
<code>not</code> : bool -> bool	<code>+</code> : " " "
	<code>-</code> : " " "
<code>~</code> : num -> num	<code>^</code> : string * string -> string
<code>abs</code> : num -> num	
<code>floor</code> : real -> int	<u>Precedence 30:</u>
<code>real</code> : int -> real	<code>::</code> : 'a * 'a list -> 'a list
<code>sqrt</code> : real -> real	<code>@</code> : 'a list * 'a list
<code>sin,cos,arctan</code> : real -> real	-> 'a list
<code>exp,ln</code> : real -> real	<u>Precedence 20:</u>
	<code>=</code> : ety * ety -> bool
<code>size</code> : string -> int	<code>&lt;&gt;</code> : " " "
<code>chr</code> : int -> string	<code>&lt;</code> : num * num -> bool
<code>ord</code> : string -> int	<code>&gt;</code> : " " "
<code>explode</code> : string -> string list	<code>&lt;=</code> : " " "
<code>implode</code> : string list -> string	<code>&gt;=</code> : " " "
	<u>Precedence 10:</u>
<code>ref</code> : mty -> mty ref	<code>o</code> : ('b->'c) * ('a->'b)
<code>!</code> : 'a ref -> 'a	-> ('a->'c)
<code>print</code> : ty -> ty	<code>:=</code> : mty ref * mty -> unit
<code>makestring</code> : ty -> string	

Special constants: as in Section 2.3.

#### Notes:

(1) The following are constructors, and thus may appear in patterns:

`nil true false ref :: ..` and all special constants.

(2) Infixes of higher precedence bind tighter. `::` associates to the right; otherwise infixes of equal precedence associate to the left.

(3) The meanings of these predeclared bindings are discussed in Section 5.2.