

## Chapter 1. Introduction

This manual describes the Cedar language. It is organized into three major parts:

Chapter 2: A description of a much simpler kernel language, in terms of which the Cedar language is explained. This description includes a precise definition (§ 32.1) and a less formal explanation of the ideas of the kernel and the restrictions imposed by Cedar (§§ 32.3-2.9). § 32.1 contains an overview or glossary, in which the major terms used in the kernel are briefly defined.

Chapter 3: The syntax and semantics of the current Cedar language. The semantics is given precisely by a desugaring into the kernel. It is also given more informally by English. This chapter also contains a number of examples to illustrate the syntax.

Chapter 4: The primitive types and procedures of Cedar. For each one, its type is given as well as an English definition of its meaning. This chapter is organized according to the hierarchy of the primitive types (§ 34.1).

In addition, there is a one-page grammar for the full language, a shorter grammar for the kernel language, and a two-page language summary which includes the grammar, the desugaring, and some examples from § 3,

The document you are reading is nearly complete. A few missing sections consist of a few paragraphs in the style of this one.

## Chapter 2. The kernel language

This document describes the Cedar language in terms of a much smaller language, which we usually call the kernel or the Cedar kernel. Cedar differs from the kernel in two ways:

- It has a more elaborate syntax (¶ 3). The meaning of each construct in Cedar is expressed by giving an equivalent kernel program.

Often the kernel program is longer or less readable; the Cedar construct can be thought of as an idiom which conveniently expresses a common operation. Sometimes the Cedar construct has no real advantage, and the difference is the result of backward compatibility with the ten-year history of Mesa and Cedar.

- It has a large number of built-in types and procedures (¶ 4). In the kernel language these could in principle be programmed by the user, though in fact most are provided as special code in the Cedar compiler. In general, you can view these built-in facilities like a library, selecting the ones most useful for your work and ignoring the others.

Unfortunately, the current Cedar language is not a superset of the kernel language. Many objects (notably types, declarations and bindings) which are ordinary values in the kernel cannot be freely passed as arguments or bound to variables, are subject to various restrictions; they can only be written in literal form, cannot be arguments or results of procedures, etc. The long-term goal for evolution of the Cedar language is to make it a superset of the kernel as defined here. In the meantime, however, you should view the kernel as a concise and hopefully clear way of describing the meaning of Cedar programs.

To help in keeping the kernel and current Cedar separate, reserved words and primitives in the kernel which are not available in current Cedar are written in SANS-SERIF SMALL CAPITALS, rather than the SERIF SMALL CAPITALS used for these symbols in current Cedar. Operator symbols of the kernel which are not in current Cedar are not on the keyboard.

The kernel is a distillation of the essential properties of the Cedar language, not an invention. Most Cedar constructs have simple translations into the kernel. Those which do not (some of the features of OPEN) are considered to be mistakes, and should be avoided in new programs.

¶ 2.2 defines the syntax and semantics of the Cedar kernel language, the former with a grammar and the latter by explaining how to take a program and deduce the function it computes and the state changes it causes. The remainder of the chapter is a commentary which explains the reasons behind the kernel. It also gives the restrictions imposed by the current Cedar language. The generality described here; for more on this subject, see ¶ 3. The meaning of the various primitives is given in ¶ 4. ¶ 2.9 describes the incompatibilities between the kernel language and current Cedar, i.e., the constructs in Cedar which would have a different meaning in a kernel program. For the most part, these are bits of syntax which do not have consistent meanings in current Cedar; future evolution of the language will replace them with their kernel equivalents.

Usually, terms are defined and explained before they are used, but some circularity seems unavoidable. ¶ 2.1 gives a brief summary of each major idea which may be helpful as a reference. Both this and the explanations in ¶¶ 2.3-2.7 are given under five major headings, as follows:

Values and computations: Application Value Variable Group Binding Argument  
 The type system: Type Type-checking Mark Cluster Declaration Class  
 Programs: Name Expression Scope Constructors Recursion  
 Conveniences: Coercion Exception Finalization Safety Process  
 Miscellaneous: Allocation Static Pragma

The kernel definition in ¶ 2.2 follows the ordering of the kernel grammar given there.

## 2.1 Overview

This section gives a brief summary of the essential concepts on which the Cedar language. The explanations are informal and incomplete. For more precise but more formal definitions see ¶ 2.2; for more explanation, see ¶ 2.3-2.8.

### 2.1.1 Values and computations

**Application:** The basic mechanism for computing in Cedar is applying a procedure (proc) to arguments. When the proc is finished, it returns some results, which can be discarded or used as arguments to other procs. The application may also change the values of some variables. The denotation of a program an application is denoted by (the denotation of) the proc followed by square brackets enclosing (the denotation of) the arguments:  $f$  [first~ $x$ , last~ $y$ ]. There are special ways of applying many kinds of application:  $x+1$ ,  $person.salary$ ,  $IF\ x<3\ THEN\ red\ ELSE\ green$ ,  $x: INT\_7$ .

**Value:** An entity which takes part in the computation (i.e., acts as a proc, argument or result) is called a value. Values are immutable: they are not changed by the computation. Examples: "Hello",  $1\ x\ IN\ x+3$ ; actually these are all expressions which denote values in an obvious way.

**Variable:** Certain values, called variables, can contain other values. The value contained in a variable (usually called the value of the variable) can change when a new value is assigned to the variable. In addition to its results, a proc may have side-effects by changing the value of a variable. Nearly every type  $T$  has a corresponding variable type  $VAR\ T$ ; values of type  $VAR\ T$  contain a value of type  $T$ . Every  $VAR$  type has a  $NEW$  proc which creates a variable of the type. A variable is represented by a single block of storage; the bits in this block hold the representation of the value.

**Group:** A group is an ordered set of values, often denoted by a constructor like this:  $[3\ "Hello"]$ . Like everything else, a group is itself a value.

**Binding:** A binding is an ordered set of [name, value] pairs, often denoted by a constructor like this:  $[x: INT~3, y: BOOL~TRUE]$ , or simply  $[x~3, y~TRUE]$ . If  $b$  is a binding,  $b.n$  denotes the value of name  $n$  in  $b$ .

**Argument:** A group or binding constructor written explicitly after an expression denotes the value  $P$  denoted by the expression to the value  $a$  denoted by the constructor, called the argument.  $P$  is usually a proc, and  $a$  is a group or binding, which is bound to its domain  $D$  to get the argument which is passed. In making this binding  $a$  is coerced, if necessary, to the declaration:

If it is a group, the names from  $D$  are attached to its elements to turn it into a binding.  
 If a name in  $D$  is missing from  $a$ , a default value is supplied.  
 If a value in  $a$  doesn't have the type required by  $D$ , it is coerced into another value with the same type.

### 2.1.2 The type system

**Type:** A type defines a set of values by specifying certain properties of each value in the set (e.g., integer between 0 and 10); these properties are so simple that the compiler can make sure that all arguments have the specified properties. A value may have many types; i.e., it may be in many of these sets. A type also collects together some procs for computing with the value (e.g.,  $+$ ,  $*$ , multiply).

More precisely, a type is a value which is a pair:

( $T$ ,  $P$ )  
 Its predicate, a function from values to the distinguished type  $BOOL$ . A value has type  $T$  if  $P$ 's predicate returns  $TRUE$  when applied to the value.

Its cluster, a binding in which each value is usually a proc taking one argument of the type of the cluster. The expression `e.f` denotes the result of looking up `f` in the cluster of `e`'s syntactic type and applying the resulting proc to the value of `e`.

A proc's type depends on the types of its domain and range; a proc with domain (argument type) `D` and range (result type) `R` has the type `D_R`. Every expression `e` has a syntactic type denoted by `T(e)`; e.g., the range declared for its outermost proc; in general this may depend on the argument type of `e`; the value of `e` always has this type (satisfies this predicate); of course it may have other types.

Mark: Every value carries a set of marks (e.g., `INT` or `ARRAY`; think of them as little flags on the top of the value). The predicate `HASMARK` tests for a mark on a value; it is normally used with type predicates. The set of all possible marks is partially ordered.

The set of marks carried by a value must have a largest member `m`, and it must include every mark smaller than `m`. Hence all the marks on a value can be represented by the single mark `m`; we can say that `m` is the mark on the value. This does not imply a total ordering on the marks.

Type-checking: The purpose of type-checking is to ensure that the arguments of a proc satisfy the predicate of the domain type; this is a special kind of pre-condition for executing the proc body can then rely on the fact that the parameters satisfy their type predicates. It is used to establish that the results satisfy the predicate of the range type; this is a special kind of post-condition which holds after executing the proc. Finally, the caller can rely on the fact that the results satisfy their type predicate. In summary:

Caller	establish pre-condition: arguments have the domain type;
	rely on post-condition: results have the range type.
Body	rely on pre-condition: parameters have the domain type;
	establish post-condition: returns have the range type.

Declaration: A declaration is an ordered set of `[name, type]` pairs, often denoted like `[x: INT, y: BOOL]`. If `d` is a declaration, a binding `b` has type `d` if it has the same set of names as `d`; each name `n` the value `b.n` has the type `d.n`. A binding `b` matches `d` if the values of `b` can be used to yield a binding `b'` (which has type `d`).

A declaration can be instantiated (e.g., on block entry) to produce a binding in which each name is bound to a variable of the proper type; instantiating the previous example yields

```
[x: VAR INT~(VAR INT).NEW, y: VAR BOOL~(VAR BOOL).NEW].
```

Class: A class is a declaration for the cluster of a type. For instance, the class `Ordered` is `PROC[T, T]_[BOOL], . . .`. `C` is a subclass of `D` if (loosely) `C` includes at least all the `[name, type]` pairs in `D`.

### 2.1.3 Programs

Name: A name (sometimes called an identifier) appearing in a program denotes the value bound to the name in the scope that the name appears in (unless the name is in a pattern before a (declaration) or tilde (binding), or after a dot or `$`). An atom is a value that can be used as a name; a literal atom is written like this: `$alpha`.

Expression: In a program a value is denoted by an expression, which is one of:

- a literal value (`3` or `"Hello"`);
- a name (`x` or `salary`);
- an application of a proc to other values (`Sin[90];`  
`GetProperties[directory, ReadFileName[input]];`
- a l-expression, which yields a proc value (`l [x: INT] IN (IF x<0 THEN -x ELSE x) );`
- a constructor for a declaration or binding (`[x: INT~3, y: REAL~3.14]`).

If a value is given for each free name in an expression, then it can be evaluated to produce a value. Thus an expression is a rule for computing a value.

Scope: A scope is a region of the program in which the value bound to a name does not change (although the value might be a variable, whose contents can change). For each scope there is a binding called ENV (for environment) which determines these values. A new scope is introduced (in the kernel) by IN (after LET or l) or by a REC [...] constructor for a declaration or binding.

```
LET x~3 IN x+5;
LET fact~1 [n: INT] IN IF n=0 THEN 1 ELSE n*fact[n 1].
```

Constructors: Brackets delimit explicit constructors for group, declaration or binding and all have the form  $[x_1, x_2, \dots]$ , and are distinguished by the form of the  $x_i$ :

```
an expression for a group;
n: e for a declaration;
n~e or n: e~e for a binding.
```

Recursion: When names are introduced in a constructor in Cedar, this is done recursively.

If  $v$  is bound to  $n$  in a binding constructor, then in expressions in the constructor  $n$  has the value  $v$  rather than its value in the enclosing scope. Exception: argument bindings are non-recursive.

If  $n$  is declared in a declaration constructor, then it may not be used in the constructor body. Exception: an ordering of the declarations in the constructor such that a name is used only by later declarations.

Exception: declared names may be used in the bodies of l-expressions in the constructor.

¶ 3.3.4)

In the kernel, however, constructors are non-recursive unless preceded by REC.

Dot notation: The form  $e.n$  looks up  $n$  in some binding associated with  $e$ , and does something with the result. There are three cases:

If  $e$  is a binding,  $e.n$  is just the value paired with  $n$  in  $e$ .

If  $e$  is a type,  $e.n$  is  $e.Cluster.n$ .

Otherwise,  $e.n$  is  $(De.n)[e]$ , and  $e.n[\text{more args}]$  is usually  $(De.n)[e, \text{more args}]$ .

In all cases you are supposed to think of  $n$  as some property or behavior associated with  $e$ .  $e.n$  denotes that property or evokes that behavior.

#### 2.1.4 Conveniences

Coercion: Each type cluster contains To and From procs for converting between values of one type and values of other types (e.g., `Float: PROC[INT]_[REAL]`; this would be a To proc in REAL and a From proc in INT). One of these procs is applied automatically if necessary to convert an argument value to the domain type of a proc; this application is a coercion. Each coercion is associated with an atom called its tag (e.g., `$widen` for INT\_REAL or `$output` for INT\_ROPE); several coercions may be composed into a single one if they have the same tag. The tags thus serve to prevent unexpected composition of coercions.

Exception: There is a set of exception values. An expression  $e$  denotes a value which is of type  $De$  or is an exception. Whenever an exception value turns up in evaluating an expression, the expression immediately becomes the value of the whole expression, unless (in the kernel) the expression is of the form  $e \text{ BUT } \{\dots\}$ . The  $\{\dots\}$  tests for exception values and can supply an ordinary value or another exception, as the value of the BUT expression. An exception value may contain an ordinary value, called the argument of the exception, so that arbitrary information can be passed with an exception.

Finalization: When a variable is no longer accessible, the storage it occupies is freed (in the safe language). Before this is done, a finalization proc in the cluster of the variable is called to do any other appropriate resource deallocation. The local variables of a proc body may also be finalized (using UNWIND).

Safe: The safety invariant says that all references are legal, i.e., each REF T value is a variable of type T. A proc is safe if it maintains the safety invariant whenever it is called with arguments of the proper types. If a proc body (l-expression) is

checked, the compiler guarantees that the proc value is safe;  
 trusted, the programmer asserts that it is safe (but the compiler makes no checks), and  
 is treated as safe;  
 unchecked, the compiler makes no checks and the proc value is unsafe.  
 It is best to write checked code whenever possible. However, checked code cannot call un-  
 (since the compiler then cannot guarantee safety).

Process: Concurrency is obtained by creating a number of processes. Each process executes  
 sequential computation, one step at a time. They all share the same address space. Shared  
 (touched by more than one process) can be protected by a monitor; only one process can enter  
 within the procs of the monitor at a time. So that each process can know what to rely on  
 necessary to establish an invariant for the monitored data which is established whenever  
 proc returns or waits. A process can wait on a condition variable within a monitor; other  
 can then enter the monitor. The waiting process runs again when the condition is notified  
 timeout.

### 2.1.5 Miscellaneous

Allocation: Cedar has standard facilities for allocating new variables of any type (the  
 primitive); related variables can be allocated in the same zone. Normally, variables are  
 automatically by the garbage collector when they can no longer be referenced; such variables  
 only be referred to by REFS. It is also possible to have variables which are deallocated  
 FREE, but this is unsafe.

Static: An expression whose value is computed without executing the program is called static.  
 Literals are static, as are names bound to literals, and any expression with static operators  
 bodies are never static unless they are inline

Pragma: Some language constructs do not affect the meaning of the program (except possibly  
 make a legal program illegal), but only its time and space costs; these are called pragmas.  
 are INLINE for proc bodies and PACKED for arrays.

## 2.2 Kernel definition

This section gives the syntax and semantics of the Cedar kernel language. Motivation, an  
 explanation of the relation between the kernel and the current Cedar language, can be found in  
 ¶¶ 2.3-2.7. The kernel is subdivided into

- A rather austere core; everything can be desugared into this, but it isn't very readable.
- A set of conveniences; with these, readable programs can be written.
- Imperative constructs: statements and loops.
- Exception handling.

The format of this section interleaves grammar rules which give the syntax of the language  
 which gives the meaning. The meaning of the core is given in English. For other parts of the  
 kernel, it is given by desugaring rules which show how to rewrite each construct in terms of  
 if rewriting is done repeatedly, the result is a core program, which may invoke some primitive  
 meaning of these is also given in English. There is also some English explanation of the  
 but this is only a commentary and does not have the force of law.

See ¶ 3.1 for the notation used in the grammar and desugaring.

## 2.2.1 The core

In the core, there is syntax only for names, literals, application, l-expressions, a basic recursive binding construction, and syntactic type; everything else is done with primitives. We write anything in the core, however, except to show the desugaring of a kernel construct. The reader need not struggle with programs in the ugly core syntax.

Syntax	Syntactic type	Meaning
Expression ::=		
n	Dn	ENV[\$n].val
literal	Dliteral	
e <sub>1</sub> Z e <sub>2</sub>	(De <sub>1</sub> .RANGE)[e <sub>2</sub> ]	-- Standard application.
l d <sub>1</sub> => d <sub>2</sub> IN e	d <sub>1</sub> -d <sub>2</sub>	-- Standard proc constructor.
L d <sub>1</sub> => d <sub>2</sub> IN e	d <sub>1</sub> -d <sub>2</sub>	-- Unchecked standard proc constructor.
[(n~e), ...]	[(n: De), ...]	-- Vanilla binding constructor.
FIX d ~ e	d	-- Recursive binding constructor.
D e	TYPE	-- Syntactic type.
type ::= e	De --gTYPE--	-- A type is syntactically just an expression.
decl ::= e	De --gDECL--	-- A decl is syntactically just an expression.
name ::=		
letter (letter   digit)n..	(Ddigit)n..	-- Appears as an e or in a pattern.
literal ::=		
\$ n	ATOM	-- ATOM literal.
primitive	Dprimitive	
primitive ::=		
ARROW	[d: DECL, p: (d_DECL)]_[a: --arrow--TYPE]	
DOMAIN   RANGE	*[a: --arrow--TYPE]_[t: TYPE]	
MKPAIR	[t <sub>1</sub> :: TYPE, first: t <sub>1</sub> , t <sub>2</sub> :: TYPE, rest: t <sub>2</sub> ]_[v: t <sub>1</sub> Xt <sub>2</sub> ]	
GROUP	[t <sub>1</sub> : TYPE]_[t: TYPE] --tgTYPE	
MKCROSS	[g: GROUP[TYPE]]_[c: --CROSS--TYPE]	
CDOTG	*[t: --cross--TYPE]_[g: GROUP[TYPE]]	
MKBINDD	[d: DECL, v: d.T]_[b: d]	
BDOTD   BDOTV	[b: BINDING]_[d: DECL]   [d:: DECL, b: d]_[v: d.T ]	
MKBINDP	[p: PATTERN, t:: TYPE, v: t]_[b: MKDECL[p, t]] --=MKBINDD[d~MKDECL[p, t], v]	
LOOKUP	[d:: DECL, b: d, n: ATOM]_[v: DTOB[d].n]	
THEN	[d <sub>1</sub> :: DECL, b <sub>1</sub> : d <sub>1</sub> , d <sub>2</sub> :: DECL, b <sub>2</sub> : d <sub>2</sub> ]_[v: d <sub>1</sub> THEND d <sub>2</sub> ]	
ENV	*BINDING	
MKDECL	*[p: PATTERN, t: TYPE]_[d: DECL]	
DDOTP	*[d: DECL]_[p: PATTERN]	
DDOTT	*[d: DECL]_[t: TYPE]	
DTOB	*[d: DECL]_[b: BINDING]=MKBINDP[p~d.P, v~d.T.G ]	
BTOD	*[b: BINDING]_[d: DECL]=MKDECL[p~b.D.P, t~MKCROSS[b.v] ]	
THEND	[d <sub>1</sub> : DECL, d <sub>2</sub> : DECL]_[v: DECL] --=BTOD[DTOB[d <sub>1</sub> ] THEN DTOB[d <sub>2</sub> ]]	
BOOL   ATOM	TYPE	
TRUE   FALSE	BOOL	
TYPE   DECL   BINDING	TYPE	-- DECLgTYPE, BINDINGgTYPE
PATTERN	TYPE	--=GROUP[ATOM]
ANY	TYPE	-- TgANY for any type T

In the kernel we dress up the primitives as follows:

A primitive denoted `xDOTy` is in the cluster of the type of its argument under the name `x`.

A parameter to a primitive declared with `::` is the type of some other argument; the name of the argument for this parameter may be omitted in an application of the primitive.

A name not in a literal (or pattern, in the kernel) denotes the value to which it is bound in the current environment ENV (A below). An ATOM literal is a value which stands for a name in the kernel. Primitives which deal with declarations and bindings.

A literal denotes a value according to a rule which depends on its syntax. The core has numeric and ATOM literals, and the primitives enumerated above.

An expression denotes a value according to a rule which depends on its syntax. If the expression is a name or literal, the value is the value of the name or literal. The remaining cases are defined in the following sub-sections. Most of these cases define the value of the expression in terms of the value of its sub-expressions. The sub-expressions may be evaluated in any order.

#### A. The current environment ENV

The current environment ENV is a binding. The value of the expression  $n$  is  $ENV.n$ . ENV for an expression is the same as ENV for its containing expression, except that:

For the  $b$  of a closure being applied, ENV is computed according to B below.

For the  $e$  of a FIX, ENV is computed according to E below.

Thus applying a closure and evaluating a FIX are the only ways to change ENV.

#### B. Application

The value of a standard application is obtained by evaluating  $e_1$  and  $e_2$  to obtain  $v_1$  and  $v_2$ , and applying  $v_1$  to  $v_2$ . There are two cases for application:

$v_1$  is a primitive. The value of the application is a function of  $v_2$  given in the definition of the primitive.

$v_1$  is a closure  $c$  (C below), with domain declaration  $d$ , body  $b$  and environment  $E$ . The value of the application is the value of the expression  $b$  in the environment  $MKBINDD[d, v_2]$  THEN  $E$  (D below). Note that if the closure was made with  $L$ , the body must be type-checked when it is applied; a closure made with  $l$  was type-checked when it was made (see below).

An application type-checks if  $De_2$  implies  $De_1.DOMAIN$  (G below).

#### C. Lambda

The value of a  $l$ -expression is a closure, which has three parts:

A domain declaration  $d$ , equal to the value of  $d_1$ .

A body  $b$ , which is the expression  $e$  (not the value of  $e$ ).

An environment  $E$ , equal to the current environment ENV.

A  $l$ -expression type-checks if

$d_1$  evaluates to a declaration  $d$ .

For any  $x$  of type  $d.T$ ,  $De$  implies  $d_2.T$  in the environment  $MKBINDD[d, x]$  THEN  $E$ .

A  $L$ -expression type-checks if  $d_1$  evaluates to a declaration; type-checking of the body is deferred until the closure is applied.



#### D. Pairs, groups and cross types

A pair is the basic structuring mechanism. `MKPAIR[x, y]` yields the pair `<x, y>`. Bigger structures are made, as in Lisp, by making pairs of pairs. When we are interested in the leaves of such a structure we call it a group and call the leaves its elements. A group has type `GROUP[T]` if all its elements have type `T` or are `NIL`. A flat group is a pair in which first is not a group, and rest is `NIL`.

The type of a pair is a cross type: `MKPAIR[x, y]` has type `TXU` iff `x` has type `T` and `y` has type `U`. Cross types are made with `MKCROSS`, which turns a `GROUP[TYPE]` into a cross type in the obvious way:

```
MKCROSS[NIL]=???
```

```
MKCROSS[T]=T if T is a type.
```

```
MKCROSS[ MKPAIR[x, y] ]=MKCROSS[x]XMKCROSS[y]
```

Note that `MKCROSS` of a flat group is flat. `CDOTG` goes the other way, turning a cross type into a `GROUP[TYPE]` in which no element is a cross type. Thus `MKCROSS` is the inverse of `CDOTG`, but not necessarily the other way around.

#### E. Bindings

A binding is either `NIL`, or an `<atom, value>` tuple, or a `<binding, binding>` tuple. The primitive `MKBINDD` constructs a binding from a declaration `d` and a matching value `v`, i.e. (as the type `MKBINDD` indicates), one with the type `d.T`. The resulting binding has type `d`, and consists of the names from `d` paired with the corresponding values from `v`. Example:

```
MKBINDD[ [x: INT, b: BOOL], [3, TRUE] ] = [x~3, b~TRUE]
      = < <$x, 3>, < <$b, TRUE>, NIL > >
```

In this example, `d.T` is `INTXBOOL`.

The declaration and group in this example is written using the syntax of ¶ 2.2.2; in the core they would be `MKDECL[p~[$x, $b], t~MKCROSS[[INT, BOOL]] ]` and `MKPAIR[first~3, rest~MKPAIR[first~TRUE, rest~NIL]]`, where we have written the arguments of these primitives in the kernel syntax.

The primitives `BTOD` and `BTOV` return the arguments of the `MKBINDD` primitive that made the binding. `MKBINDP` is redundant; it is like `MKBINDD`, but takes a type instead of a declaration and hence accepts any `v` with the right structure.

`LOOKUP` returns the value of the name `n` in the binding. `THEN` combines two bindings, giving priority to the first one in case of duplicate names. It works only for flat bindings, i.e. each element of each `<binding, binding>` tuple is an `<atom, value>` tuple, and the second element is another `<binding, binding>` tuple or `NIL`. The value of `b1 THEN b2` is another flat binding, obtained by first replacing any tuple `<<a, v>, b>` in `b2` where `a` is equal to an atom in `b1` by `b`, and then replacing this binding to replace the final `NIL` in `b1`.

The binding constructor `[p~[e], [(..)]` has the value `MKBINDP[p~[n, ...], v~[e, ...] ]`.

`FIX` makes a recursive binding: the value of `FIX d1~e` is `MKBINDD[d, v]`, where `d` is the value of `d1` in the environment `ENV` and `v` is the value of `e` in the environment `(LET FIX d~e IN d~e) THEN ENV`. Of course in general this computation may not terminate; normally the names in `d` occur in `e` only with 1-expressions and in this case it does terminate. Is this really right? The latter checks if `De` in the latter environment implies `DTOT[d]`.

## F. Declarations

A declaration is either NIL, or an <atom, type> tuple, or a <declaration, declaration> tuple. The primitive MKDECL constructs a decl from a pattern p and a value t of type GROUP[TYPE]. A pattern p is a GROUP[ATOM], i.e., either NIL, or an atom, or a pair of patterns; the ATOM elements must be different. An application of MKDECL typechecks if t matches p, i.e., if

both p and t are NIL, or

p is an atom and t has type TYPE, or

p is a pair [p<sub>1</sub>, p<sub>2</sub>] and t is a cross type t<sub>1</sub>Xt<sub>2</sub> and p<sub>1</sub> matches t<sub>1</sub> and p<sub>2</sub> matches t<sub>2</sub>.

The resulting declaration consists of the names from p paired with matching type values.

The primitives DDOTP and DDOTT return the arguments of the MKDECL primitive that made the declaration. Thus

```
DDOTT[NIL]=???
```

```
DDOTT[<$n, T>]=T;
```

```
DDOTT[<d1, d2>]=DDOTT[d1]XDDOTT[d2]
```

DTOB is redundant; it converts a declaration to a binding in which each name has the corresponding type as its value. Thus DTOB[[x: INT, y: REAL]]=[x~INT, y~REAL]. The inverse is BTOD, also redundant; it is defined only if all the values in the binding are types. THEND combines declarations just as THEN combines two bindings: D(b<sub>1</sub> THEN b<sub>2</sub>)=Db<sub>1</sub> THEND Db<sub>2</sub>

## G. Types and type-checking

A type is a value consisting of a pair:

the predicate, a function from values to BOOL.

the cluster, a binding.

A value v has type T if T's predicate applied to v is TRUE.

T implies U iff (Ax) T.Predicate[x]gU.Predicate[x].

Typechecking consists of ensuring that the argument of an application has the type specified by the domain of the proc (B above). The body of a l-expression can then be type-checked (or the implementation of a primitive constructed) independently, assuming that the parameter set has the specified domain predicate. Symmetrically, the result of an application can be assumed to have the type specified by the range of the proc.

To complete the induction, it is also necessary to check that the value of the body of a l-expression has the range type (C above).

The primitive types in the kernel are:

BOOL, with two values TRUE and FALSE.

ATOM, with values denoted by literals of the form \$n.

TYPE, a predicate satisfied by any type value.

ANY, a predicate satisfied by any value.

DECL, the type of a declaration (F above).

BINDING, the type of any binding.

Arrow types, the types of procs (C above). An arrow type has a domain type and a range type.

Cross types, the types of pairs (D above).

GROUP[T], the type of any pair in which all the elements have type T.

Declarations, the types of bindings (E and F above).

There are no non-trivial implications among any of these types, except as follows:

DECLgTYPE; BINDINGgTYPE; GROUP[T]gTYPE.

ANYgT for any type T.

$T_1 \times T_2 \text{g} U_1 \times U_2$  iff  $T_1 \text{g} U_1$  and  $T_2 \text{g} U_2$ .

GROUP[T]gGROUP[U] iff TgU.

$T_1 \text{---} T_2 \text{g} U_1 \text{---} U_2$  iff  $U_1 \text{g} T_1$  and  $(\forall x: U_1) (l T_1 \text{ IN } T_2)[x] \text{g} (l U_1 \text{ IN } U_2)[x]$ . Note the reversal of the domains.

$d_1 \text{g} d_2$  for declarations iff  $d_1.P = d_2.P$  and  $\text{DFOB}[d_1].\text{ngDFOB}[d_2].n$  for each  $n$  in  $d_1.P$ .

## 2.2.2 Conveniences

```

expression ::= coreExpression |
  d1 _ d2 |
  l ( | e1) ( | => e2) IN e3 |
  LET e1 IN e2 |
  LET b, ... IN e |
  IF e1 THEN e2 ELSE e3 |
  e . n |

  e1 [ b, ... ] | e1 [ e2, ... ] |
  e1 infixOp e2 |
  e1 AND e2 | e1 OR e2 |
  [ ] | [ e1 ( | e2, !.. ) ] |
  PATT p |
  [ b, ... ] |
  REC [ (p : t ~ e), ... ] |
  [ d, ... ] |
  XX |
  statements | simpleLoop |
  but

infixOp ::=
  X
  PLUS
  THEN

literal ::= coreLiteral |
  digit digit ... |

declaration ::=
  p: t |
  [(p: t), ... ]

binding ::=
  p ~ e |
  d ~ e |

pattern ::=
  n |
  [p1, ...]

primitive ::= corePrimitive |
  LOOKUP | LOOKUPC |
  PLUS |
  IFPROC |
  ARROW Z [d1, l d1=>DECL IN d2]
  -- The domain defaults to [], the range to x: De3 |
  ( l De1 IN e2 ) e1.v -- e1 a binding |
  LET [b, ...] IN e
  ( IFPROC[De2, e1, l IN e2, _ [IN] e3] )
  IF DegBINDING THEN LOOKUP Z [De, $n]
  ELSE IF DegTYPE THEN LOOKUP Z [De.cluster, $n]
  ELSE_(LOOKUPC Z [De.cluster, Z$1]) |
  e1 . APPLY Z [b, ...] | e1 . APPLY Z MKBINDD[De1.DOMAIN
  e1 . infixOp |
  IF e1 THEN e2 ELSE FALSE | IF e1 THEN TRUE ELSE e2 |
  NIL | MKPAIR[e1, [ ( | -e2 Group) constructor. |
  -- Pattern constructor; see the rule for p below. |
  b PLUS ... PLUS NIL |
  FIX [p, ...] : MKCROSS[[t, ...]]~[e, ...] |
  d PLUS ... PLUS NIL |
  xxxxxx | --Also recursive d maps into this?
  -- See ¶ 2.2.3
  -- See ¶ 2.2.4.

MKCROSS

INT -- Numeric literal, giving the decimal rep
-- A d is not an e; a d must be before ~ or after LE
MKDECL[ PATT p, t ] |
[p, ...]: MKCROSS[[t ,-.to]separate names and types
-- Only the [...] form is an e; a b must be written
MKBINDD[PATT p, e] |
MKBINDD[d, e]
-- Note: a pattern is not an e; it can appear only l
or after PATT in the kernel.
-- PATT n=$n
-- PATT [p1, ...]=[PATT p1, ...]

-- Fill in types

```

The precedence of operators in  $e$  is: (highest)  $[\ ]$ ,  $Z$ ,  $\text{infixOps}$  (all the same),  $\text{BUT}$ ,  $\text{IN}$  (l left associative).

## 2.2.2.1 Expression syntax

Most of this is straightforward sugar.  $\text{LET}$  adds the binding  $e_1$  to  $\text{ENV}$  in evaluating  $e_2$ . The case for  $b, \dots$  simply allows the  $[\ ]$  which normally enclose a binding constructor to be a case; see ¶ 2.2.2.2.  $\text{IF}$  wraps  $e_2$  and  $e_3$  in  $l$ 's so that they don't get evaluated; the  $\text{IFPROC}$  chooses the one to evaluate and applies it.

The dot notation has three cases.

For a binding it just looks up  $n$  in the binding.

For a type it looks up  $n$  in the type's cluster.

For anything else, it looks up  $n$  in the cluster of  $De$  and applies the result to  $e$ . `LOOPUPC` primitive does something special if it finds a `proc` which takes more than one argument: it splits the `proc` into one which takes the first argument and returns a `proc` taking the remaining arguments. This ensures that if  $De.n$  is such a `proc`  $P$ , the expression  $e.n[a, b]$  will desugar into something equivalent to  $P[e, a, b]$ .

The usual syntax for application is a `proc`  $e_1$  followed by an explicit binding constructor. The syntax of application may depend on the type of  $e_1$ , via the `APPLY` element of its type; for a `proc` the standard `apply` operator  $Z$ , `APPLY` is the identity. If  $e_1$  is followed by a group or a binding constructor, the argument is obtained by binding the group to the declaration whose domain.

Infix operators desugar straightforwardly into application; note that the choice of `proc` is by the type of the first operand only. `AND` and `OR` are not ordinary infix operators, since they evaluate no more than necessary; this is expressed by the desugaring into `IF`.

The remaining expression syntax is various constructors, described in the next section, imperative and exception features described in later sections.

#### 2.2.2.2 Group, binding and declaration constructors

A bracketted sequence of expressions (e.g., `[1, 2, 3]`) denotes a flat group with its elements in the same order (e.g., `MKPAIR[1, MKPAIR[2, MKPAIR[3, NIL]]]`). Thus a group constructor is just like the `LIST` function in Lisp. A pattern is a similar construct, except that it contains names instead of literals for the corresponding `ATOM` literals; `PATT` yields the group obtained by replacing each name with the literal `$n`. After desugaring a pattern always appears after `PATT` and hence is always enclosed into an `atom` or a `GROUP[ATOM]`.

Brackets are also used to delimit binding and declaration constructors. They are distinguished from each other, and from group constructors, by the presence of `~` in each element of a binding constructor, and `:` in each element of a declaration constructor. The elements of a binding constructor are sugar for applications of the `MKDECL`, `MKBINDP` and `MKBINDD` primitives. The constructor itself strings the resulting declarations into a big one using the `PLUS` operator which is just like `THEN` except that it does not allow duplicate atoms; the motivation for this is to allow the names and corresponding types or values to be written together, instead of factoring out primitives require. As a result, values made from constructors are always flat.

Note that these constructors do not nest, except that a `d` can be `[(p: t), ...]`. This is the `d~e` form of binding; e.g., if `DivRem` returns two `INT`s, you can write `[d: INT, r, INT]~DivRem` instead of `[d, r]: INTXINT~DivRem[...]`.

The `REC` binding constructor is sugar for `FIX` which exactly parallels the non-recursive one.

#### 2.2.3 Imperatives

These constructs are generally used together with non-functional procs.

```
statements ::= { e; ... }           IF (ISVOID[e]) AND ... THEN [] ELSE ERROR
                                         -- Ordering by non-prompt evaluation.
simpleLoop ::= SIMPLELOOP statements LET REC [loop_1~N { statements; loop_1(N) loop_1(N)}]
                                         -- Only an exception (such as EXIT) will terminate t
```

Each *e* in the statements must evaluate to `VOID`; this is to catch mistakes like writing `x+` statement. The definition of `AND` ensures that the *e*'s are evaluated left-to-right.

The `simpleLoop` is the standard way to express a loop in terms of recursion. You are supposed to use an exception to get out of this loop; Cedar provides a number of convenient ways to do this, such as `EXIT` and `RETURN`.

#### 2.2.4 Exceptions

An exception is treated as a special value returned from an application. The exception value contains an exception code and an `args` value which may be of any type. When an application returns an exception value, it immediately abandons the application and returns the exception value. The application is strict. There has to be some way to stop this, or the first exception would be the end of the program. The `HIDE` primitive takes any value and returns a variant record of type `HEX`. It turns:

a normal value into the normal variant, with the value in its `v` field;

an exception into the exception variant, with the code in its `code` field and the `args` in its `args` field.

`UNHIDE` takes a `HEX` value and returns the original unhidden value.

An exception code has the type `EXCEPTION[T]`, where *T* is a declaration which is the type of the `args`; it is the domain of the exception, and `(DEXCEPTION[T]).DOMAIN=T`. An exception value is constructed by the primitive

```
RAISE: [T:: TYPE, code: EXCEPTION[T], args: T]
```

Thus the `args` always has the type demanded by the code.

This is dressed up with the following syntax.

```
but ::= e BUT { butChoice; ... }      LET v(~HIDE[e]_IN (
                                     IF ISTYPE[v(, HEX.normal] THEN UNHIDE[v()]
                                     ELSE IF ISTYPE[v(, HEX.exception] THEN
                                         LET h(~NARROW[v(, HEX.exception] ] IN
                                             LET selector(~h(.code) ENSURE butChoice UNHIDE[v()]
                                         ELSE ERROR )
butChoice ::= e1 => e2 |           IF selector(=e1) THEN LET MKBINDD[De1.DOMAIN, h(.args] IN
e1 , e1, !.. => e2 |             IF (selector(=e1) OR ... THEN e2 |
ANY => e2                          IF TRUE THEN e2
```

A `BUT` expression evaluates *e*. If it is a normal value, that is the value of the `BUT`. If it is an exception, each `butChoice` in turn gets a look at it. If one of them likes it, then it succeeds; otherwise the exception is the value.

The *e*<sub>1</sub> in a `butChoice` must evaluate to an exception code. If there is just one, and it matches in the exception, then `args` in the exception is bound to the domain of the code, and *e*<sub>2</sub> is evaluated in that environment. If there is more than one, then *e*<sub>2</sub> is just evaluated in the current environment. An `ANY` `butChoice` matches any exception, but of course doesn't bind the arguments.

## 2.3 Doing computations

This section describes the basic mechanisms for doing computations, and the kinds of values that can be manipulated by Cedar programs.

### 2.3.1 Application

The basic mechanism for computing in Cedar is applying a proc to argument values. A proc is a mapping

from argument values and the state of the computation,  
to result values, and a new state of the computation.

The state is the values of all the variables.

A proc is implemented in one of two ways:

By a primitive supplied as part of the language (whose inner workings are not open to inspection).

By a closure, which is the value of a l-expression whose body in turn consists of an l-expression, which may contain further applications of procs to arguments, e.g., `l x+3`. When a closure is applied, the parameters declared after the `l` are bound to the arguments, and then the body after `IN` is evaluated in the new environment thus obtained.

In Cedar, each parameter value thus obtained is used to initialize a variable, which is named by the parameter in the body. Thus the body can assign to the parameters. Use of this feature is not recommended.

Note that when a l-expression is evaluated to obtain a closure its body is not evaluated. The body is saved in the closure, to be evaluated when the closure is applied. Some constructs (`IF`, `SELECT`, `OR`) are defined (see ¶ 2.2.2 and ¶ 3.8) by wrapping l-expressions around some arguments, applying them only when certain conditions hold; e.g., `IF b THEN f [x] ELSE g [y]` evaluates to `f [x]` if `b` is `TRUE` and `g [y]` iff `b` is `FALSE`.

Application is denoted in programs by expressions of the form `f [arg, arg, ...]`. If the value of `f` is a closure, this expression is evaluated by evaluating `f` and all the `arg`'s, and then evaluating the body of the closure with the formal parameters bound to the arguments (unless an exception value is returned; see ¶ 2.6.2). Thus to evaluate `(l [x: INT] IN x+3)[4]`:

evaluate the l-expression to obtain a closure;  
evaluate the argument 4 to obtain the number 4;  
evaluate `x+3` with `x` bound to 4 to obtain the number 7.

The first two evaluations can be done in either order.

To evaluate a primitive application such as `x+3`, evaluate the arguments, and then invoke the primitive on those arguments to obtain the result and any state change. With a few exceptions (assignment and dereferencing or following references), primitives are functions and can be represented as tables which enumerate a result value for each possible combination of arguments. A primitive can therefore be viewed as a simple table lookup using the arguments as the table keys.

Actually there may be one more step in an application. If an argument doesn't have the type expected by the proc, the argument is coerced to the proc's domain type if possible. If no coercion can be found, there is a type error. Coercion is discussed further in ¶ 2.6.1 and ¶ 4.13.

Most procs take a binding as argument, in which the various parts of the argument are named. For example, `OpenFile: PROC[name: ROPE, mode: Files.Mode]` takes a binding with two values named `name` and `mode`. It might be applied like this: `P[name~"Budget.memo", mode~$read]`. If the names are missing

is a positional coercion which supplies them left-to-right, see ¶ 2.3.6. There is also a coercion that supplies missing parts of the binding; see ¶ 4.11.

If  $f$  is neither a primitive nor a closure, the meaning of applying it is defined by the type; this case is discussed further in ¶ 4.4.

There are many ways of writing applications other than  $f [x]$ . In fact, many Cedar primitives are the values of expressions, and can only be applied by writing some other construct. The desugaring rules show how large parts of the Cedar syntax denote various special kinds of application. In each case, the meaning is defined by the standard meaning of application plus a specific meaning of the primitives involved; see ¶ 4.1.

This is partly because of history, and partly because specialized syntax makes the program more readable. Future evolution of the language will improve the situation.

### Functions and order of evaluation

An expression is functional if

its value does not depend on the state, but only on the values bound to its free variables;  
and

evaluating it does not change the state.

As a consequence of this definition,

Two identical functional expressions in the same scope will always have the same value.

A `proc` is a function if every application of it is functional. It doesn't matter when or how many times a function is applied; the order of evaluation doesn't matter for functions. Thus functions can be thought of as mathematical functions for many purposes. Note that a constant can be regarded as an application of a function of no arguments.

Non-functional `procs`, on the other hand, are more complicated objects. Cedar makes no formal distinction, either in syntax or in the type system, between functions and `procs`. However, `procs` do not define the order of evaluation in an expression, except that:

all arguments are evaluated before a `proc` is applied;

because of the desugaring of `IF`, `SELECT`, `AND` and `OR` into `l-expression`, the order of evaluation for these expressions is determined by the first rule;

statements separated by semi-colons are evaluated in the order they are written.

As a consequence, two applications of non-functions should not be written in the same statement unless they don't affect each other; if this is done the effect of the program is unpredictable.

An expression is guaranteed to be functional if it only applies functions; thus if  $f$  is a non-functional `proc`, and  $x$  a variable,  $f [3]$  is functional and  $p[3]$  and  $p[x]$  may not be.  $f [x]$  may not be functional, because it is sugar for  $f [x.VALUEOF]$ , and `VALUEOF` is not a function. The value of a `l-expression` is a function if its body is functional. There are more complications in guaranteeing that an expression is functional, just as for any other interesting property.

Because the values of variables constitute the state, it is only the existence of variables that allows non-functional `procs` to exist. In particular, the `VALUEOF` `proc` which returns the value of a variable is non-functional (because its result depends on the state), and the `ASSIGN` `proc` which changes the value of a variable is non-functional (because it changes the state).



### 2.3.2 Values

A Cedar program manipulates values. Anything which can be denoted by a name or expression in the program is a value. Thus numbers, arrays, variables, procedures, interfaces, and type values. In the kernel language, all values are treated uniformly, in the sense that each

passed as an argument,  
bound to a name, or  
returned as a result.

These operations must work on all values so that application can be used as the basis for computation and l-expressions as the basis for program structure. In addition, each part or type of value has its own primitive operations. Some of these (like assignment and equality) are defined for most types. Others (like addition or subscripting) exist only for certain types (numbers or arrays). None of these operations, however, is fundamental to the language. assignment or equality has the same status as any operation on an abstract type supplied by an implementor; thus `INTEGER.ASSIGN` has the same status as `IO.GetInt`. In practice, of course, special syntax is usually used to invoke these operations, and the implementations are not open to inspection by the editor or debugger. A complete description of the primitives of the language can be found in Chapter 4, organized by the type of the main operand. Table 4.5 is an alphabetized index of these descriptions.

**Restrictions:** In current Cedar, however, there are restrictions on values which are type or bindings: they can only be arguments or results of modules, and hence are first-class in the modelling language, and not within a module. Also, declarations and bindings cannot be constructed or bound to identifiers within a module. Unions are also restricted: they can only appear inside records. Nonetheless, it is simplest to emphasize the uniform treatment of values and consider separately the restrictions on types, declarations, bindings and unions. Future versions will improve this situation.

**Restriction:** In current Cedar you can only use dot notation for some operations of built-in modules, the procs which access record fields, and others as noted in Table 4.5. As a substitute, various syntactic forms which are sugar for dot notation: infix, prefix and postfix operators, functions, and funny applications. These desugarings are given in rules 20-24 of the Cedar grammar in ¶ 3.

### 2.3.3 Variables

Certain values, called variables, can contain other values. A variable containing a value of type `T` has type `VAR T`. If the variable doesn't allow the value to be changed, the type is `READONLY T`. This is not the same as `T`, because there may be a `VAR T` value which is the same container. The value contained by a variable (usually called the value of the variable) can be changed by assignment of a value to the variable. The set of all variables accessible from the process array consists of those of the computation; these are all the variables which can be reached from any process, and a variable which cannot be reached cannot affect the computation. Note that a variable value is a container, which like all values is immutable; it may help to think of it as (the address of) storage. The contents of a variable can be changed by assignment. Thus the value of a variable can change, even though the value that is the variable is immutable.

A suitable abstract representation for a `VAR T` is a value of type `[Get: []_T, Set: T_[]]`. This representation is not used in Cedar, but it clarifies the way in which variables fit into the system: `VAR TgVAR U` only if `T` and `U` have the same predicate, because the `Get` proc requires `TgU` and the `Set` proc requires `UgT`. `READONLY T` corresponds to `[Get: []_T]` and a write-only variable type would be `[Set: T_[]]`.

There is a coercion (an automatically applied conversion; see ¶ 2.6.1) from `VAR T` to `T`, so a variable can be passed without fuss as an argument to a proc which expects a value.

Restriction: In current Cedar, variables generally cannot be passed as arguments or results; the only exception is that an interface can declare a variable (called an exported variable) for which the implementation supplies a value; this is normally written `x: VAR INT` in the interface, but for historical reasons it is also possible to write just `x: INT`. Certain primitives (e.g., `DECL` or `POINTER`) return variables, a variable can (indeed, must) be passed as the first argument to `ASSIGN`, and a variable can be bound to a name by a declaration in a `LET` or block (`LET x: INT` binds a `VAR INT` value to `x`). For the most part, however, a program which wants to handle variables must do so at one remove, through procs or `REFS` (or, unsafely, `POINTERS`).

A variable is often represented by a block of storage; the bits in this block hold the representation of its value. All the built-in `VAR` types are represented in this way. A variable `u` overlaps variable `v` if assigning to `u` can change the value of `v`. The primitive `ASSIGN` procs have the property that

if `r` and `s` are `REFS`, then `r` overlaps `s` iff `r=s`.

For any variables `u` and `v` with the same `VAR` type, `u` overlaps `v` iff `u=v`, provided that no program has given overlapping blocks of storage to the two variables (if `u` and `v` have different types, one might be contained in the other).

The role of variables in non-functional expressions is discussed in ¶ 2.3.1.

#### 2.3.4 Groups

There is a basic mechanism for making a composite value out of several simpler ones. Such a composite value is called a group, and the simpler values are its components or elements. For example, `[x+1, "Hello"]` denotes a group, with components `x+1`, the value of `x+1`, and "Hello". The main use of explicit groups is for passing arguments to procs without naming them (these are sometimes called positional arguments). This is done by binding the group to the declaration which specifies the domain type of the proc; the result is a binding which is the argument the proc expects. For example, `P: [x: INT, y: REAL]_[. . .]`, the application `P [2, 3.14]` is sugar for `P [ [x: INT, y: REAL]_ [2, 3.14]` which is equivalent to `P [x~2, y~3.14]`.

A group has a type which is the cross type of its component types: if `x` has type `T` and `y` has type `U` then `[x, y]` has type `TXU`. Thus for syntactic types, `D[e1, e2, ... ]=De1XDe2X ...`. The `X` type constructor is associative, and type implication (¶ 2.4.2) extends to cross types: if `T1` and `T2` are types, there is a coercion called `MKCROSS` from `[T1, T2, ...]` to `T1X T2X ...`; because of this, explicit cross type is usually not needed.

Restriction: Current Cedar provides no way of making cross types except as domain and range of a proc type (or other transfer type); e.g., `PROC [INT, REAL]_[BOOL, ATOM]`. There are no constructors for taking groups except the group-to-binding coercions. Hence the only thing to do with a group is to pass it to one of the built-in coercion procs by writing it as a proc argument, or to use a type constructor as described in the next section. Current Cedar does not have `X`, but it does have `MKCROSS` to cross type coercion described in the last paragraph and illustrated in the example below.

#### 2.3.5 Bindings

A binding is a group in which each element has a name. Thus, it is an ordered set of [name, value] pairs. There are three main uses for a binding:

- As an argument in an application. Thus, if `P` is a proc with type `PROC[i: INT, b: BOOL]`, the first argument must be a binding such as `[i~3, b~TRUE]`. The application then looks like `P [i~3, b~TRUE]`. A binding argument is sometimes called a keyword argument list. See the next section for details.
- In a `LET` expression, to give names to values in the scope of the `LET`. Thus,

LET i~3, b~TRUE IN (IF b THEN i+5 ELSE 0)  
 has the value 8. Current Cedar doesn't have LET expressions, but a binding at the top of a block has the same effect. See ¶ 2.5.4 on scopes for details.

- As a way of collecting and naming a set of related values. A value can be extracted from the set using dot notation. Thus if b is the binding [i~3, b~TRUE], the value of b is b. In current Cedar this only works for interfaces; see ¶ 3.3.4 and ¶ 4.14 for details.

A binding is usually denoted by a constructor, which takes the form

```
[i~3, b~TRUE]
```

or redundantly (if there are no coercions)

```
[i: INT~3, b: BOOL~TRUE]
```

in which the types are specified explicitly (but you can't write the second form as the application). See ¶ 2.5.5 on constructors for details.

### 2.3.6 Arguments

When a group or binding is bound to a declaration (d~v), there are various conversions or coercions which may be applied to the values. This usually happens when the arguments of an application are bound to the parameter declaration.

First, if v is a group rather than a binding, it is coerced to a binding by attaching the name of the elements of v in order. Thus in

```
[a: INT, b: REAL]~[2, 3.14]
```

the group constructor is coerced to [a~2, b~3.14].

Next, if v is shorter than d, elements of the form n~OMITTED are appended, where n is the corresponding name from the declaration. Thus in

```
[a: INT, b: REAL]~[2]
```

the group constructor is coerced to [a~2, b~OMITTED].

Now the items of the binding are matched by name with the items of the declaration. There is an error unless the names match exactly. The remaining coercions are done on individual items from the declaration and the corresponding n~v from the binding. If v has type t, all items of the declaration must have a coercion to t. Otherwise, if there is a sequence of coercions from the type of v to t, these are applied. If such sequence exists, there is an error. In particular, there is a coercion from OMITTED to t, if any. Thus in

```
[a: INT_0, b: REAL_1.1]~[b~3.14]
```

the group constructor is coerced to [a~0, b~3.14], and in

```
[a: INT_0, b: REAL_1.1]~[]
```

it is coerced to [a~0, b~1.1]. Coercions are discussed in ¶ 2.6.1 and ¶ 4.13, defaulting to the identity coercion.

An important special case is constructors for record and array values. A record type has a constructor proc; e.g.,

```
R: TYPE=RECORD[a: INT, b: REAL_0.]
```

has a proc R.CONST of type PROC[a: INT, b: REAL\_0.]\_[R]. Thus R.CONST[a~2, b~3.1416] constructs a record value. There is also a coercion from BINDING to the particular binding type RB which is the domain type of R.CONST, so that

```
r1: R_[a~2, b~3.1416]
```

is short for

```
r1: R_R.CONST[a~2, b~3.1416].
```

Composing the positional coercion from GROUP to RB with R.CONST makes

```
r1: R_[2, 3.1416]
```

also short for the previous line.

The same scheme works for arrays, but only an array indexed by an enumeration has a corresponding binding which can be written; the elements of an array indexed by numbers have names which can be written in a binding. However, the group constructor still works

## 2.4 The type system

This section describes the way in which types can be used to make assertions about the program which the compiler can verify. It also discusses the role of types in organizing the name space of the program.

### 2.4.1 Types

Types serve two independent but related functions in Cedar:

- A type contains an assertion about some property of a value, e.g., that it is a value between 0 and 10 represented in a single machine word. A value which has the property is said to be of that type, or to have that type.

The assertion part of a type is called its predicate. It is a function which accepts a value (of any type) and returns TRUE iff the value satisfies the assertion.

- A type contains a collection of named procs (and perhaps other values) related in a useful way. Most often, the procs of type T take a value of type T as their first argument. For example, INT has PLUS, TIMES and MINUS procs (usually written as infix or prefix operators) which can be applied to INTs. The dot notation (see ¶ 2.4.4) makes it easy to refer to the procs in a type's collection.

The collection part of a type is called its cluster. It is simply a binding. No rules are enforced about what kind of values are in the binding. However, the idea is that the type is an interface for manipulating values of the type (perhaps the main or even only interface). As with any interface, a tasteful choice of names and values is important.

The predicate and the cluster serve rather different purposes:

The predicate provides the basis for type-checking (¶ 2.4.2). The most important function of type-checking is to guarantee the integrity of abstract data types; this is done with the use of predicates called marks (¶ 2.4.3).

The cluster provides the basis for convenient naming of a large collection of procs and other values (¶ 2.4.4). Clusters are organized into a hierarchy of classes (¶ 2.4.4).

Like everything else which can be named, a type is a value. Hence there is nothing special about binding a type value to a name. If T is a type expression, the binding

```
U: TYPE~T
```

binds T's value to U. In the scope of U, T and U are completely interchangeable (provided U is not rebound). Furthermore, with two exceptions, all type expressions are functional: identical type expressions in the same scope denote the same type value. The exceptions are the record and enumeration type constructors, which make a distinct type each time they are used (by creating a new mark; see ¶ 2.4.3).

Caution: The AMTypes interface does distinguish between T and U as a convenience in debugging but it also provides a procedure GetCanonicalType for obtaining the type value in the sense described.

Restriction: Current Cedar has a number of restrictions on the use of TYPE values, given in the appendix.

### 2.4.2 Type predicates and type-checking

Type predicates provide a way of making assertions in the program which can be checked mechanically. These assertions take the form of declarations for the formal parameters of a procedure. In general the checking must be done during execution. Thus, if the program says

```
a: ARRAY [0..10] OF INT_ALL[0];
i: INT_s.ReadInt;
s.PutF[ a[i] ];
```

there must be a check that  $i > 0$  and  $i < 10$  just before the expression  $a[i]$  is evaluated. This is a bounds check; if it fails there is an exception called `Runtime.BoundsFault`. Where did this come from? Note that  $a[i]$  is short for `Da.APPLY[a, i]`, and `Da.APPLY` is `SUBSCRIPT`, the subscript procedure for `ARRAY [0..10] OF INT`. The type of `SUBSCRIPT` is `PROC[array: ARRAY [0..10] OF INT, index: [0..10]]_[VAR INT]`. So when  $i$  is passed as the index argument, the declaration of `SUBSCRIPT` says it must have the type `[0..10]`. The predicate for this type is

```
1 [x: ANY] IN HASMARK[x, INT] AND LET y: INT~x IN y>=0 AND y<=10.
```

Leaving the `HASMARK` term for later discussion, we see that the rest of the predicate is the bounds check.

The type system is designed, however, so that most assertions can be checked statically by examining the text of the program without running it. Static checking has three obvious advantages:

- It reports any errors after a single examination of the programming, leaving none of that kind) to be discovered later in Peoria.

- It introduces no cost in time or space for run-time checking.

- The compiler can take advantage of the assertions to generate better code.

Of course, there is a corresponding drawback: the assertions made by parameter declarations must be simple enough that the compiler can reliably prove or disprove them.

The proofs done for typechecking have exactly the same form as program correctness proofs on preconditions and postconditions. Consider a proc whose value is the l-expression

```
1 [x: T]=>[y: U] IN e.
```

The domain declaration `[x: T]` is a precondition for the body `e`. This means that any application of the proc must satisfy this condition. As a consequence, the body `e` can be analysed on the assumption that the precondition holds, i.e., that  $x$  has type  $T$ . Similarly, the range declaration `[y: U]` is a postcondition for the body. This means that given the precondition, any evaluation of the body will produce a value  $y$  which has type  $U$ . In summary, for the body we assume the precondition and must establish the postcondition.

To make this hang together, each application must establish the precondition; this means the argument must have the domain type. In return, the application can assume the postcondition; this means that the result of the application has the range type. Thus we have a linkage:

```
argumentgdomaingrangeresult
```

The result in turn will be the argument of another application. In this way the proof is passed on to larger and larger expressions, and finally to the whole program. In summary:

Application	establish pre-condition: arguments have the domain type;
	rely on post-condition: results have the range type.
Body	rely on pre-condition: parameters have the domain type;
	establish post-condition: returns have the range type.

These proofs require showing that an expression always has a particular type  $T$ . This is done by observing that every expression has a unique syntactic type  $U$ , which is the type of every value of that expression; e.g., an application always has the range type of its proc (see below for a detailed discussion of syntactic type). If every value of type  $U$  has type  $T$ , we are done with the usefulness of type implication. One type implies another,  $T \supseteq U$ , iff  $(\forall x) T[x] \supseteq U[x]$ . If two types are equal, each implies the other. However, there are many other useful cases of implication. For instance, `VAR INT` implies `READONLY INT`. The type implications in current Cedar are given in § 4.12.

Of course, not all arguments are applications. The kernel grammar gives the other possible argument expressions, and we enumerate the proof rules for each:

- A literal is like a zero-argument proc: it has a known range (e.g., `3` has type `INT`, `'a'` has type `CHAR`).

A name has the type specified in its declaration or binding.

If there is only a declaration  $n: T$  (e.g.,  $x: INT$ ), it must be the domain dec of a l-expression, and we have already seen how to ensure that the  $n$ 's value type  $T$  when the resulting proc is applied.

If there is a binding  $n: T \sim e$  for the name (e.g.,  $x: INT \sim 3$ ), we must check that type  $T$ .

A l-expression  $l [x: T] \Rightarrow [y: U] \text{ IN } e$  has the type  $[x: T] \_ [y: U]$ . This works for the discussed in the next paragraph.

A binding constructor  $[x \sim e, y \sim f]$  has the type of the corresponding declaration,  $[x: Df]$ .

There is one more link in the chain. An application  $f [x]$  has an arbitrary expression for  $x$ , not necessarily a l-expression. The requirement is that  $f$  must have a proc type, say  $D\_R$ ;  $D$  is domain type and  $R$  the range type. Since the type of  $l D \Rightarrow R \text{ IN } e$  is  $D\_R$ , satisfying the precondition  $D$  for the application is the same as satisfying the precondition  $D$  for the  $l$  and similarly in reverse for the postcondition. The value of  $f$  may be a primitive rather than a closure obtained from a l-expression. In this case, the implementation of the primitive must depend on the precondition and must still establish the postcondition, but since the implementation cannot be examined (within the framework of Cedar) we can say nothing about how this is accomplished. Example: `INT.PLUS`, which is implemented by the machines 32-bit add instruction.

In a proc type  $D\_R$ ,  $D$  and  $R$  may be declarations which provide names for the arguments and results. In general, the expression  $R$  may include names declared in  $D$ . The range type of an application then depends on the argument values.

Restriction: In current Cedar only modules have such types; the type returned by an interface or the interfaces exported by an implementation, may depend on the interface and implementation parameters.

As a by-product of the type-checking proof rules just given, a syntactic type is derived for every expression  $e$  in the program. It is denoted by  $De$ , and computed as follows:

for a name, the declared type;

for a literal, its type;

for an application, the range type (which may depend on the argument);

for a l the obvious proc type;

for a binding constructor, the declaration obtained by pairing the names with the types of the value expressions.

Typechecking ensures that whenever  $e$  is evaluated, the resulting value will have type  $De$  (it may have other types as well, i.e., it may satisfy other predicates). The main use of syntactic types is in connection with dot notation (see ¶ 2.4.4).

In order to carry out the proofs described above, the compiler must either compute the values of type expressions, including those denoted by complex expressions such as `ARRAY [i..j] OF INT`, or it must be able to prove the equality of unevaluated type expressions. For the most part, current Cedar requires the former approach; hence a type expression must have a value which the compiler can compute. Such a value is called static; the rules for static values are given in ¶ 3.9.1.

## 2.4.3 Marks

By this point you may have thought of asking why the assertions provided by type predicates are worth all this fuss. The reason is simple: they are the basis for authenticating values of an abstract type, so the implementation can be sure that it is working on properly formed values. Suppose you are the implementer of an abstraction, e.g., `Table`. You provide operations to `Lookup` a key in the table, to `Insert` a [key, value] pair, and to `Enumerate` the items in the table. A `Table` is represented by a `REF` to a record containing a sorted array `a` of items and an `INT` `n` which gives the number of items. `Lookup` is implemented by binary search. All three operations are programmed on the assumption that elements 0 through `n - 1` of `a` are sorted, and that `n` is smaller than the size of the array. They will not work properly if these assumptions are not satisfied, and indeed they will fail to subscript the array with an out-of-bounds index or to violate other requirements of the abstractions they depend on.

Here is a lower level, but perhaps more dramatic example. The dereferencing operation `^` on a `REAL` returns a `VAR REAL`, which can, for instance, be assigned to, as in the program fragment

```
r: REF REAL_NEW[REAL_1.0];
. . .
r^ _ 3.14159
```

A `REF REAL` is represented by the address of a four-byte block of storage which holds a `REAL`. The assignment to `r^` stores the four bytes which represent 3.14159 into that block. If some other `REF REAL` finds its way into `r`, the assignment will still store four bytes, since it doesn't know any better. But the `REF REAL` points to a two-byte block; the other two bytes that will be modified belong to some unrelated variable, which will be clobbered without warning.

The second example is scarier because the consequences of the bug seem more unpredictable. In both cases, however, the fundamental problem is the same: even if the implementation is correct, the wrong thing happens because it is given an improper value to work on. Or to make the point in different words, the implementation cannot be held responsible for bad results from its operations, if it has no control over the validity of the arguments it receives.

So that the implementation of an abstraction can take responsibility for correct operation, there must be a way to authenticate a value of the abstract type. In Cedar this is done by placing a mark on the value; think of it as a little flag stuck into the value. The mark uniquely identifies the abstract type, and authority to affix it is under the control of the implementation. A correct implementation will mark only values which have the properties needed for a representation of an abstract value, and if no one else can affix the mark, the implementation can be sure that a value with the mark has the desired properties.

A mark can be thought of as an abbreviation for an assertion or type invariant which characterizes a proper abstract value, such as `Table` or `REF REAL`. Such an assertion can be quite complex. In the `Table` example, it would say that the representation is a record of the proper form, that the array size is correct, and that the first `n` array elements are sorted. In the `REF REAL` example, it would say that the address points to a block of storage such that at least the first four bytes do not overlap other blocks. Such assertions are not easy to write down formally, and proving them is beyond the power of any existing program. So the abbreviations are not a mere convenience, but a necessity.

A new mark can be created on demand by the primitive

```
CREATEMARK: PROC[Rep: TYPE, tag: UNIQUEID]_[m: MARK, Affix: [Rep]_[TYPEFROMMARK[m]] ]
```

The primitive `HASMARK` tests a value for the presence of a mark, so `HASMARK[x, m]` tests `x` for the presence of the mark `m`. `Affix` adds the mark to a `Rep` value.

Restriction: `MARK`, `UNIQUEID`, `CREATEMARK`, `HASMARK` and `TYPEFROMMARK` are not accessible in current Cedar. Record and array type constructors provide some access to `CREATEMARK`, as described below. The `ISTYPE` primitive, also described below, is closely related to `HASMARK`.

With these facilities, it is easy to create a new abstract type. Choose its representation, obtain a new mark  $m$ . `TYPEFROMMARK[m]` with an appropriate cluster added is the new abstract type. The implementation must use `Affix` to mark only values which satisfy the properties it demands.

The type returned by `TYPEFROMMARK[m]` has the predicate

```
1 [x: ANY]=>[BOOL] IN HASMARK[x, m]
```

and an empty cluster. Except for subranges and bound unions, all types in current Cedar have a predicate of this form. The built-in types (`INT`, `BOOL` etc.) come with such predicates, and types created in type constructor procs (`ARRAY`, `RECORD` etc.) obtain a mark from `CREATEMARK`. So that two invocations of `ARRAY [0..10] OF INT` will produce the same type, `ARRAY` and most of the other type constructors use a canonical encoding of the constructor and its arguments for the `UNIQUEID` function, hence are functional. `RECORD` and `ENUMERATION` produce a different type each time they are invoked, so they obtain fresh unique identifiers. Since the program cannot invoke `CREATEMARK` directly, we need not explain how to prevent forgery of `UNIQUEIDS`. Future versions of Cedar will solve this problem.

In current Cedar you make a new abstract type by declaring `in` as an opaque type in an interface:

```
T: TYPE[ANY]
```

This generates a new mark, and declares `T` to be a type which has that mark. You get such a type by explicitly painting some other type, normally in an implementation which exports `T` to an interface which declared it:

```
T: PUBLIC TYPE~Interface.T PAINTED RECORD [...].
```

See ¶ 4.3.4 for more details.

The implementation actually stores a mark with each variable allocated by `NEW`. Such a variable is referenced by a `REF`, and in particular by a `REF ANY` value. The type of a `REF ANY` value can be tested at runtime using the primitive

```
ISTYPE: PROC[x: ANY, U: TYPE]_[BOOL]
```

If `De` is `REF ANY` and `RT=REF T`, then the value of `ISTYPE[e, RT]` is `TRUE` iff the predicate for `T` just tests for mark  $m$ , and `x^` has the mark `VAR m`. `ISTYPE` is described in detail in ¶ 4.3.1. The `WITH ... SELECT` construct and the `NARROW` primitive, which are more powerful operations built up from `ISTYPE`.

For other values, there is no mark actually stored; instead, types must be computable strings using the methods described in the last section. The `AMTypes` interface, however, gives a way to refer to any value in a uniform way, and to test its type at runtime.

There is only room for one mark on a variable, and this must encode all the marks that the variable actually carries. We arrange for this by imposing a partial order on the marks, and requiring

The set of marks on a value must have a maximal element.

Every mark smaller than the maximal one must be on the value.

With these rules, a single mark stored on the value is enough to code all the others.

In current Cedar, a value actually has only one mark, since:

The only way to create a new mark is with the record or enumeration type constructor by declaring an opaque type.

When you paint a type `T` with the mark of an opaque type, `T` must be a record or enumeration type, and the opaque type mark replaces the mark it had before.

Note that `VAR T`, `READONLY T` and `T` are different types with different marks, although `VAR T` is a coercion `VALUEOF` from either one to `T`.



#### 2.4.4 Clusters and dot notation

It is convenient to associate with a type the procs supplied by its implementor for deal values of the type. This is done by putting these procs into the type's cluster. The cluster binding which is part of the type value (the predicate is the other part). There are no details about what goes into the cluster. However, there is a special dot notation which makes it easy to populate T's cluster with procs which take a T as their first argument. The usual effect of this: `t.n` is sugar for `Dt.n[t]`, and `t.n[other args]` is sugar for `Dt.n[t, other args]`.

For example, if `t` has type `T`, and a proc `[T, INT]_[BOOL]` is in `T`'s cluster under the name `P`, the proc can be applied by an expression like `t.P [3]`, which is sugar for `Dt.P [t, 3]`. The lookup is done only in `T`'s cluster, not in the current scope. If `Q: [T]_[INT]` is also in the cluster, `t.Q` can be applied with `t.Q`, which is sugar for `Dt.Q [t]`.

The general rule that makes this work is the following: `t.n` is sugar for `LOOKUPC[Dt, $n][t]`. `LOOKUPC[Dt, $n]` is just `Dt.n`, except that if `Dt.n` is a proc that takes several arguments, it is turned into a proc that takes the first argument and returns a proc taking the remaining ones. `LOOKUPC[Dt, $n][t]` will be a proc taking the remaining arguments, and `t.n[other args]=LOOKUPC[Dt, $n][t][other args]` will be the same as `Dt.n[t, other args]`.

Dot notation can also be used to obtain values from a binding or from the cluster of a type. In any application: `T.P` would be the proc named `P` in the previous example. The possible cases of dot notation in current Cedar are described in detail in ¶ 4.14.

Restriction: There is currently no way to explicitly construct clusters. The built-in type constructors have clusters; they are described in detail in ¶ 4. In addition, there is a way to provide a cluster for an opaque or record type in an interface: every proc name in the interface is put into the type's cluster. For a record, the procedures supplied by the record constructor are in the cluster, and they win if there are name conflicts. There is one of these clusters in each imported interface value; if a module imports more than one value of the same interface, however, there are severe restrictions (see ¶ 3.3.3).

#### 2.4.5 Declarations

A declaration is the type of a binding. Thus, the binding `[x~3, y~3.14]` has the type of `[x: INT, y: REAL]`. All the relationships among types, and between types and values, are carried out elementwise to decls and bindings; the elements are matched up by name rather than by position. A decl itself simply has the type `DECL`.

A decl is made up of two parts: the names or pattern, and the types. The basic operation for making decls, `MKDECL`, takes a pattern and a type. Thus `MKDECL[ PATT[x, y], INTXREAL]=[x: INT, y: REAL]`. In general, a pattern is one of `NIL`, a simple name, or a pair of patterns, just like a list expression. Similarly, a type argument to `MKDECL` is one of `NIL`, a type, or a cross type. The type must decompose in a way which matches the pattern. Normally, as in Lisp, we deal only in flat patterns, where the first element of a pattern is always a name. Such flat patterns are denoted by constructors of the form `[x, y, ...]`. The reason for defining things in terms of patterns is that it makes it much simpler to write down precise rules for the semantics, using structural induction on the values.

The main use of a decl is to type-check a binding. The basic binding constructor is `MKBIND`, where `d` is a decl and `e` is a matching group or binding. If `e` is a binding, then its structure must match the structure and names of `d`, and each element of `e` must have the type demanded by the corresponding component of `d`, after a possible coercion. Thus `MKBINDD[[x: INT, y: REAL], [y~3.14]]=[x~3, y~3.14]`. This may seem pointless, but it has two important uses:

Such a binding is used to bind the argument of a proc to the domain declaration. Even though the resulting binding is the same as the argument, the type-checking is essential.

There may be coercions involved, so that the resulting binding is not the same. Coercions on the component values are discussed in ¶ 2.6.1. There are also coercions on the decl itself, which can default missing elements; these are discussed in ¶ 2.3.6.

If  $e$  is a group, it is first coerced to a binding by attaching the names from the decl, ¶ 2.3.6. Thus in `MKBINDD[[x: INT, y: REAL], [3, 3.14]]` the second argument is coerced to `[x~3, y~3.14]`, and things then proceed as before.

Bindings may also be used in LET expressions. Here the types are often redundant, and it is better to use the `MKBINDP` primitive to bind the value directly to a pattern. The syntactic type of the decl is the decl whose type is the syntactic type of the value. Thus `[x~3, y~3.14]` is short for `MKBINDP[PATT[x, y], [3, 3.14]]`; its syntactic type is `MKDECL[[x, y], D[3, 3.14]]=MKDECL[[x, y], INTXREAL]=[x: INT, y: REAL]`.

A decl  $D$  in a block is interpreted somewhat differently. It becomes the argument of the `MKBINDP` primitive, which turns the type of the decl  $D.T$  into the corresponding VAR type  $VT=D.T.MKVAR$ . `MKBINDP` allocates a new value  $v$  of type  $VT$ , and makes the binding `MKBINDP[D.P, v]` over the scope of the block. Thus

```
{x: INT; y: REAL; S}
```

becomes

```
LET [x, y]~[VAR INT, VAR REAL].NEW IN S
```

Here  $D=[x: INT; y: REAL]$ ,  $VT=[VAR INT, VAR REAL]$ , and  $v=[VAR INT, VAR REAL].NEW$ . Note that the types might have defaults, which are used to initialize the values as part of the `NEW` operation.

Actually this is a bit oversimplified, since `NEWFRAME` has to separate the bindings in the decls, construct the variable binding just described from the decl, and then combine the binding from the block. Thus

```
{x: INT; y: REAL; z~TRUE; S}
```

becomes

```
LET [x, y, z]~([VAR INT, VAR REAL].NEW PLUS [TRUE]) IN S
```

or more readably

```
LET x~VAR INT, y~VAR REAL, z~TRUE IN S
```

Anomaly: In Cedar the names in a block are introduced recursively, so that the  $d$ 's and  $b$ 's refer to each other. It is possible for a binding or type to refer to a value which has not yet been initialized, with undefined results. See ¶ 3.4.1 for a further discussion of this point.

#### 2.4.6 Classes

Another important use of a declaration is to characterize the cluster of a type. Since a type is just a binding, it is characterized by its type, which is a decl. When used for this purpose, the decl is called a class. See ¶ 4.1 for further discussion of classes, and an enumeration of the types of decls in Cedar.

### 2.5 Programs

This section describes how meaning is assigned to kernel programs.

#### 2.5.1 Structure of programs

A kernel program is an expression, which is either atomic (a name or literal), or is an expression which involves sub-expressions: the `proc` being applied, and the arguments. The concrete kernel treats certain kinds of expressions specially: modules, blocks (which introduce new variables and return no value), and statements (which return no value). All desugar into simple expressions, however, and are treated identically in the kernel.

### 2.5.2 Names

A name is a part of a program which usually serves to denote a value. There are two contexts in which the occurrence of a name  $n$  denotes a value:

It may occur as an expression. Then  $n$  denotes the value bound to it in the scope in which the expression appears (see ¶ 2.5.4 for details).

It may occur after a dot, as in  $e.n$ . Then the expression  $e.n$  denotes the binding for  $n$  supplied by  $e$  (see ¶ 2.4.4 and ¶ 4.14 for details):

the value bound to  $n$  in  $e$ , if  $e$  is a binding;  
 the value bound to  $n$  in the cluster of  $e$ , if  $e$  is a TYPE;  
 roughly  $(De).n[e]$  otherwise.

There are also two defining contexts for a name  $n$  (see ¶ 2.5.5 for details):

It may occur before a  $\sim$  in a binding constructor, as in  $n\sim e$ . The value of  $e$  is the value bound to  $n$  in the binding denoted by the constructor (see ¶ 2.3.5 for details).

It may occur before a  $:$  in a declaration constructor, as in  $n: t$ . The value of  $t$  is the value bound to  $n$  in the declaration denoted by the constructor (see ¶ 2.4.5 for details).

These constructors are usually recursive in Cedar; that is, the expression  $n$  elsewhere in the constructor denotes the value bound to  $n$  in that constructor; see ¶ 2.5.6 for details. In the kernel they are non-recursive unless preceded by REC.

A name is not a value, but there are values of type ATOM which are related to names. An atom is a print name which is a rope (an immutable sequence of CHARS). A name following a  $\$$  is an atom literal;  $\$n$  denotes the atom with print name  $n$ . Other properties of atoms are described in § 4.7.1.

Caution: Current Cedar has several complications in its treatment of names:

- In an `argBinding`<sup>27</sup>,  $n: e$  may be written instead of  $n\sim e$ . The syntactic context distinguishes this from a declaration, but this usage is not recommended.

An `argBinding` is not recursive: in  $\{a\sim 1; f[a\sim 3, b\sim a+1]\}$   $b$  is bound to 2, not to 4.

The declaration in an `import` list is non-recursive: `IMPORT M` is short for `IMPORT M: M` and the second `M` denotes its binding in the surrounding scope (i.e., the binding supplied by the `DIRECTORY`). Inside the body of the module, of course, `M` denotes the imported parameter.

Names which appear in an `enumerationTC`<sup>54</sup> are treated specially; see ¶ 4.7.1.1 for details.

### 2.5.4 Scope

A scope is a region of the program in which all names retain the same meanings (note that some names denote variables, which can change their values in the same scope, but each name can denote the same variable). In the kernel there are only three constructs which introduce a scope, `l`, `LET` and `REC`. In current Cedar, these are sugared in a variety of ways: modules, lists, `proc` bindings, blocks, `exit` labels, `open`, iterators, `safeSelects` and `withSelects`. In the kernel, straightforward desugarings, however.

### 2.5.5 Constructors

The kernel has constructors, denoted  $[...]$ , to make expressions which denote groups, declarations, and binding values more readable. There is one flavor of constructor for each class:

A binding constructor is a list of binding elements ( $b$  in the kernel syntax) of the form  $b\sim e$  or  $d\sim e$ . The presence of the  $\sim$  distinguishes it from the others. Here  $d$  is a declaration.

(not a declaration), and  $p$  is a pattern, in which the names are being defined rather than evaluated.

A decl constructor is a list of decl elements ( $d$  in the syntax) of the form  $p: t$ . The tilde  $\sim$  of the  $:$  without any  $\sim$  distinguishes it from the others. Again,  $p$  is a pattern.

A group constructor is a list of expressions. Note that decl and binding elements are expressions, although constructors are expressions.

Constructors are useful for making decls and bindings where the names are literal. This is the normal case, and in fact the only case in current Cedar. If you want to make them out of decls, for instance to bind an expression to a decl which is the value of a name  $dn$ , you can use the constructor;  $[dn\sim e]$  would bind the value of  $e$  to the name  $dn$ , not to the decl which is its value. You have to write the decl-constructing primitive directly: `MKDECL[d, e]`.

The only kinds of constructor you can write in current Cedar are:

Decl constructors for proc domains and ranges, and for records and unions (fields in the syntax).

Binding constructors for arguments in an application, or as an expression alone if a scalar or array value is needed (`argBinding`<sup>27</sup> in the syntax).

### 2.5.6 Recursion

In the kernel, you get recursive definition of names only if you write `REC` (or the unsugared `FIX`) explicitly. In Cedar, on the other hand, decls and bindings are normally recursive, as are `argBindings` and `import` lists.

The recursion is legal in a block or interface body (although anomalies are possible in blocks when names are used before they are defined; see ¶ 3.4.1). In fields it is illegal.

## 2.6 Conveniences

### 2.6.1 Coercion

A coercion is a proc which is automatically applied under some circumstances to map a value of one type  $T$  (called the source) to a value of another type  $U$  (called the dest), e.g. from  $T$  to  $U$ . Coercions are obtained from the clusters of the types involved. The coercion mechanism adds a new functionality, since the programmer could always write the applications himself, but it is important in concealing some of the distinctions made by the type system when they are done rather than helpful.

There is exactly one (desugared) context in which a coercion is applied: when an expression of syntactic type  $T$  appears as an argument in an application which expects a value of type  $U$ . This means that there is a binding  $n: U\sim e$ . Since nearly all Cedar constructs are desugared to applications, coercions are widely applicable. The only (desugared) context in which there is no coercion is the first operand of dot, since in that case the cluster of the operand is used to interpret the second operand. Thus in the expression  $e.n$ , it is always  $De$ , the syntactic type of  $e$  is used to look up  $n$ , regardless of the fact that this expression may appear as an argument or parameter of type  $U$ . If  $e$  is not a type or binding, however, then  $e.n$  desugars to  $P[e]$ , where  $P = \text{LOOKUPC}[De.Cluster, \$n]$ , and in the application of  $P$ ,  $e$  does appear as an argument and can be coerced. Usually the cluster for  $T$  is set up with procs which take an argument of type  $T$  and whose domain of  $P$  is  $De$  and no coercion happens. This isn't always true, though; a subrange  $T$  of  $U$  inherits the arithmetic procs of `INT`, for example, and there is a coercion from  $T$  to `INT` which applies.

If  $T \sim U$  it is sometimes natural to think in terms of a coercion from  $T$  to  $U$  that is implemented by the identity function. In fact, implication is stronger than that, since it propagates type constructors, including `PROC`, when coercion does not. Implication is discussed in ¶ 4.12.

There is a rather general rule for finding coercions from the clusters of types, though of much practical importance in current Cedar, since there is no way for the user to define them. The rule goes like this. Each cluster may have a `From` item and a `To` item. `T.From` should consist of pairs with type `[tag: ATOM, proc: T_U]`, and `T.To` of pairs with type `[tag: ATOM, proc: U_T]`. For the moment, consider the binding  $n: U \sim e$ , where  $De = T$ , and  $T \sim U$  is false. For each `proc P` in `T.From` or `U.To` we try  $n: U \sim P[e]$ .

If  $P: T \sim V$  is in `T.From`, it maps  $e$  to a value of type  $V$ , and we have to bind  $n: U \sim P[e]$ . If  $V \sim U$  we are done; otherwise we can recurse on this sub-problem.

If  $P: V \sim U$  is in `U.To`, we have to bind  $m: V \sim e$ . If  $V \sim U$  we are done; otherwise we can recurse on this sub-problem.

The whole process fails if no path of coercion procs takes us from  $T$  to  $U$ . The search can be abandoned when all paths have been explored, and a particular path can be abandoned when a type appears on it for the second time. Since the search is done statically (by the compiler), and since the result of an attempt to coerce  $T$  to  $U$  can be cached, the time required for the search is not a problem.

There are two obvious difficulties with this scheme. First, it may transform erroneous coercions into legal ones, but coercing an argument in ways not intended by the programmer. Second, more than one path of coercion procs may exist, and different paths may give different results. This second difficulty can be avoided, and the first minimized, if every coercion proc  $P$  is correct, that is, if it has a (partial) inverse, and  $P^{-1}[P\{x\}] = x$  for all  $x$  in  $P.DOMAIN$ . This says that a coercion does not lose information, and that different paths give the same answer. Sometimes this is not feasible (for the narrowing coercion from `INT` to `[0..5]`). The following rule gives the builder of coercion procs over proliferating coercions:

If two procs on a coercion path have non-nil tags, they must have the same tag.

In general, coercions that don't lose information can have `NIL` tags, and others should have non-nil tags.

The coercions in current Cedar are described in ¶ 4.13. All have `NIL` tags, and none loses information except the subrange narrowing. Note that coercions extend componentwise to `g` and bindings.

### 2.6.2 Exceptions

The basic idea behind exceptions is to extend the value space, so that it includes not only ordinary values, but also a set of exception values. An exception value has the special property that when it appears in an application, it becomes the value of the application, so that it propagates through the control stack of the program until it finally becomes the value of the whole program. Of course this isn't always what is wanted, so there is a special `HIDE` construct which is like an ordinary application, but takes its argument value, ordinary or exception, and bundles it into a record which is a normal value. Then ordinary code can be used to test for the exception value and take appropriate action. This construct is sugared to give distinctive ways of catching an exception in the kernel with `BUT` (¶ 2.2.4), and in Cedar with `ENABLE`, `EXITS` and `REPEAT` (¶ 3.4.2). Cedar has two kinds of exception: `GOTO` labels and `ERRORS`, which must be raised and caught separately, and which have slightly different semantics.

The main point of this treatment is that it does not require continuations or any other explanation of how control is transferred to catch an exception. The view is that exceptions are simply a convenience feature; the same job could be done by returning a slightly larger value from each proc, with an appropriate status code.

An exception consists of a code and an optional argument value. The type of the code is  $T$  where  $T$  is the type of the argument which does with it. `goto` labels always have empty arguments. The argument is a way of passing some information along in addition to the identity of the exception.

A proper treatment of exceptions in the type system would require that each `proc` range in the exceptions that can emerge from an application of the `proc`. In fact, this is not required possible in current Cedar.

Cedar also has signals, which historically were viewed as a kind of exception but now have a different interpretation, as a way of obtaining dynamic rather than static scoping for names. They are discussed in ¶ 3.4.3.1.

### 2.6.3 Finalization

This subject is discussed in ¶ 3.4.3.1.

### 2.6.5 Concurrency

This subject is discussed in ¶ 4.10, where the Cedar facilities for writing concurrent programs are given. Writing good concurrent programs, or even correct ones, is another matter, which is beyond the scope of this manual to more than hint at. Unfortunately, an adequate reference is lacking.

## 2.7 Miscellaneous

The different kinds of allocation are discussed in ¶ 4.5. Static values are defined in ¶ 4.6.

### 2.7.1 Pragmas

A pragma is a construct that does not change the meaning of the program, except perhaps to make something illegal which was legal without the pragma. Its purpose is to affect the implementation generally by requesting optimization to favor one criterion over others. The pragmas in Cedar are:

`INLINE`, which causes a `proc` body to be expanded inline when it is applied. See ¶ 3.4.3.1 for details.

`PACKED`, which causes array components that fit in 8 or fewer bits to be packed, and incurs the expense of more expensive code to access them.

`CHECKED`, which forbids application of unsafe `procs` in a block, and adds runtime checks for some primitive `procs` which are otherwise unsafe (in particular, narrowing to a type and assigning a `proc`).

`PRIVATE`, which forbids access to items in an interface or instance except to modules that `EXPORT` (or `SHARE`) it.

`MACHINE DEPENDENT`, which allows positions of record fields and representation values for enumeration elements to be specified (strictly, it is the absence of `MACHINE DEPENDENT` that is the pragma)

## 2.8 Relations among groups, types, declarations and bindings

Cedar has are four closely related basic ways of building product values from simple val given precise meanings in ¶ 2.2.1 and ¶ 2.2.2):

a group is simply an n-tuple of values (see ¶ 2.3.4);

a X-type is the type of a group (if  $x: T$  and  $y: U$  then  $[x, y]: TXU$ ) (see ¶ 2.4.5);

a binding is an n-tuple of [name, value] pairs (see ¶ 2.3.5);

a declaration is the type of a binding, an n-tuple of [name, type] pairs (see ¶ 2.

Figure 2 1 illustrates the relations among these kinds of objects. In current Cedar most objects can be constructed and manipulated only as interfaces and instances. In the kern modeller, all of them are first-class citizens. The primitives which go between them are ¶ 2.2.

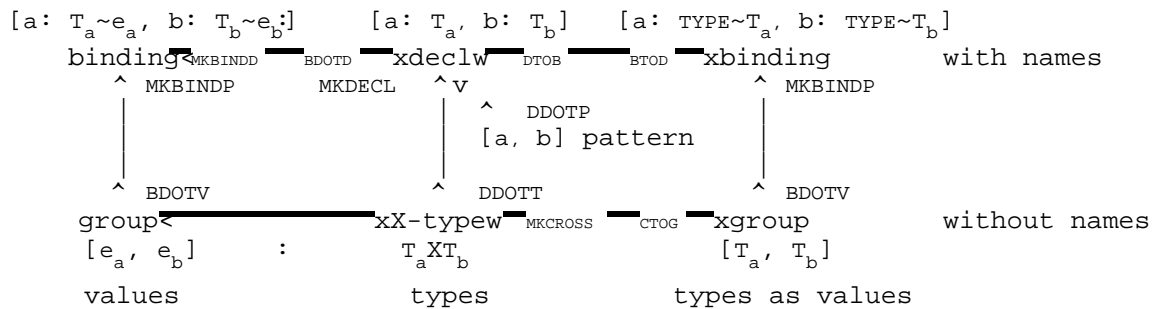


Figure 2 1: Relations among groups, types, bindings and decls

## 2.9 Incompatibilities with current Cedar

Most of the syntax in current Cedar is an extension (or sometimes a restriction) of kernel. There are a few things that have different meanings in the kernel, however, and these are sources of confusion:

Type expressions in Cedar do not have the same syntax as ordinary expressions and appear in the same contexts, for the following reasons:

The use of adjectives for variants (red Node).

The use of `_` for specifying a default value for a type vs its use for assignment.

The use of `{}` for enumeration types vs its use for a block.

In addition to writing `n: t~e` for a binding, you can also write `n: t=e` (in a module or block) and `n: e` (in an `argBinding`). The most unfortunate consequence is that a `argBinding` can look like a kernel `decl` constructor!

Target type overloading for enumeration identifiers (red instead of `Color.red` or `$red`) and union constructors (`[rator~$plus, rands~binary[...]]`) is incompatible with the kernel rules for the meaning of names.

It is now possible to avoid all the conflicting constructs except the relatively harmless defaults, `{}` for enumeration, and union constructors.

## Chapter 3. Syntax and semantics

This chapter gives the concrete syntax for the current Cedar language, together with an explanation of the meaning of each construct, and a precise desugaring of each construct. The kernel language defined in ¶ 2. The desugaring, together with the definitions of the kernel primitives used in it, are the authority for the meaning; the informal explanation is just for reading pleasure. However, paragraphs beginning Anomaly or REstriction document properties of Cedar not captured in the desugaring. The primitive procs and types of Cedar are specified in the appendix.

In addition to the grammar rules and desugaring, there are examples for each construct. These are intended to illustrate the constructs and do not form a meaningful program. The Cedar Manual contains longer examples which do something interesting, and also illustrate the use of the standard packages.

There are several summaries which may be useful as references:

- A two-page summary of all the syntax, desugaring and examples in this chapter (CLRMSumm.press).

- A one-page summary of the full syntax (CLRMFullGram.press).

- A shorter and less cluttered summary of the syntax for the safe language; it also lists a number of constructs which are obsolete or intended only for efficiency hacking (CLRMSafeGram.press).

The chapter begins with a description of the notation (¶ 3.1). The next sections deal systematically with the rules of the grammar, explaining peculiarities of the syntax and giving the semantics.

- ¶ 3.2, rules 56-61: The lexical structure of programs.

- ¶ 3.3, rules 1-5: Modules.

- ¶ 3.4, rules 6-10: Blocks, OPEN, ENABLE, EXITS.

- ¶ 3.5, rules 11-13: Declarations and bindings.

- ¶ 3.6, rules 14-18: Statements.

- ¶ 3.7, rules 19-27: Expressions.

- ¶ 3.8, rules 28-35: Conditional constructs: IF and SELECT.

¶ 3.9 treats various miscellaneous topics. ¶ 4 deals with the syntax and semantics of types.

The order of the grammar rules is:

```

module, block, declaration, statement,
expression, conditional
type,
name, literal

```

and top-down within these.



### 3.1 Notation

This section describes the notation used in the grammar, the desugaring, and the comment this chapter.

#### 3.1.1 Notation for the grammar

The grammar is written in a variant of BNF:

Bold parentheses are for grouping: ( interface | implementation).

Item | item means choose one.

?item means zero or one occurrences of item.

item; ... means zero or more occurrences of item separated by ";". The separator may a

ELSE, IN, or OR, or it may be absent. If the separator is ";", a trailing ";" is option

item; !.. is just like item; ... but there is at least one occurrence.

A terminal is a punctuation character other than bold ()?|, or any character underlined

SMALL CAPS. Note that [] and {} are terminals, and do not denote optional occurrence and repetition as the other variants of BNF.

The rules are numbered sequentially.

Special symbols mark constructs with special properties:

,=unsafe;

•=obsolete;

„=machine-dependent;

μ=efficiency hack.

The grammar is written so that a non-terminal never expands to the empty string. When an of a rule is optional, that is always indicated explicitly by "?" or "...".

The following non-terminals are so basic to the language and so frequently used, that th represented in the grammar by abbreviations:

b=binding<sup>13</sup>

d=declaration<sup>11</sup>

e=expression<sup>19</sup>

n=name<sup>56</sup> (identifier)

s=statement<sup>14</sup>

t=type<sup>36</sup>

I'm afraid this means that you must learn the meaning of these six abbreviations in orde the grammar.

With the exception of these abbreviated non-terminals, each use of a non-terminal is cro referenced with a small superscript number<sup>59</sup>, unless the non-terminal is defined in one c few rules. If a non-terminal (other than e, t or n) is used in more than one rule, then that use it are listed in a comment after its definition.

Except for the entries in Table 3 1, a terminal symbol appears in only one rule. These d do not lead to ambiguity. In most cases they are harmless, since the symbol has essentia meaning in each case, and the rules are separate only for greater readability, to highli use of a construct, or for historical reasons. In some cases, however, the symbol has qu meanings in different rules. These are marked on the left as follows

- The rules marked with • are obsolete and should be avoided.
- 6 In the rules marked with \* the symbol has a different meaning than in the others, confusion is quite possible. The programmer should bear these cases in mind.
- 0 In the rules marked with \* the symbol has a different meaning than in the others, context is sufficiently clear that confusion is unlikely.

A superscript<sup>xn</sup> indicates that the terminal is repeated n times in that rule.

Symbols	Rules	Explanation
0 ( )	19, 25, *51.1, *54	Expr, subrange, *position in record or enumeration
[ ]	19, 25, *37, 43, 51	Constructor/built-in/funnyAppl, subrange, application, *typeName, fields, mdFields
0 { }	2, 6, 8, 13, *54	interface body, block, enable, machine code, *enumerationTC
,	2, 3, 6, 7, 9, 17, 27, 29, 30, 32, 34, 35, 43, 51, 52	see note in ¶ 3.2.
;	6, 8, 10, 17, 27, 30, 33, 35	see note in ¶ 3.2.
0 :	1, 2, 3, 5, *7, 11, 13, 18, *27, 33, *34, *51.1, 53	Introducing names with types, except *51.1=position, *51.1=open, *27=argBinding *34=withSelect
.	19, 37	dot notation for e is repeated for types
0 ..	25 <sup>x4</sup> , *51.1	subrange, *position <sub>0</sub>
* 21, *53	infixOp, *tag	
+	21, 58	infixOp, exponent
• =	20, 21, 58	prefixOp, infixOp, exponent
=>	*13, 22	*binding, infixOp
6 _	6, 9, 17, 31, 33, 35, 52	enable, repeat, select choices <sup>x4</sup> , unionTC
0 ~	14, 16, 18, 21, *55	e_STATE, iterator, e, *defaultTC
~~	2, 3, 13, 20, *22, *27	interface, implementation, b, argBinding, *unaryOp, *relOp
6 ANY	7, 34	open, withSelect
6 CODE	*9, 40, 43	*enable, variableTC, fields
ENDCASE	*13, 23	*new exception, convert t to e
0 ERROR	31, 52	select endChoice, unionTC
IMPORTS	*19, *24, 41.1	*expression, *funnyAppl, transferTC
IN	2, 3	interface and implementation
LONG	18, 22	iterator, relOp
NOT	38 <sup>x2</sup> , 45.1, 48	cardinal/unspecified, pointer, descriptor
• NULL	20, 22	prefixOp, relOp
PACKED	14, *27, *52, *55	statement, *argBinding, *unionTC, *defaultTC
SELECT FROM	44, 45	array, sequence
SHARES	29, 32, 34, 52	select <sup>x3</sup> , unionTC
0 SIGNAL	2, 3	interface and implementation
TRASH	*24, 41.1	*funnyAppl, transferTC
TRUSTED	27 <sup>x2</sup> , 55 <sup>x2</sup>	argBinding, defaultTC
6 USING	6, 13	block and machine code
0 WITH	1, *5	directory, *locks
	*32, 34	*safeSelect, withSelect

Table 3 1: Terminal symbols appearing in more than one rule

## 3.1.2 Notation for desugaring

The right-hand column is desugaring into the Cedar kernel language, or in a few cases in comments describing the meaning in English. This is a purely textual transformation; i.e. on the text of the program, not on the values. The rewriting is done one rule at a time; step of rewriting involves elements from exactly one rule. The desugaring is specified by informal but straightforward rewriting rules, in which:

An occurrence of a non-terminal (written in bold) denotes the text produced by that terminal in the grammar rule.

A | reflects a corresponding alternation in the grammar rule, ? reflects a corresponding optional item in the grammar rule, and (bold parentheses) are for grouping as in a rule. As in grammar rules, literal parentheses are underlined.

Everything else is taken literally.

An underlined non-terminal in the right column means that the desugaring specified for that non-terminal must be done in order to obtain a legal program. Otherwise the transformations are done in any order, yielding a legal program at each step.

Every occurrence of *e* (expression) and *t* (type) in the desugaring is implicitly parenthesized so that the desugared program parses as the rewriting rule indicates. To reduce clutter, these parentheses are not written in the desugaring rules.

For type options like `PACKED`, the desugaring of the construct in which they appear is a call to a built-in type constructor which takes a corresponding `BOOL` argument defaulting to `FALSE`; if the `PACKED` attribute is present, the argument is supplied with the value `TRUE`.

Examples: the following rule for subranges:

```
subrange ::= ( typeName | ) (
  ( [ e1 .. e2 ] | [ e1 .. e2 ] )      (typeName | INT).MKSUBRANGE ( [ e1, ( e2 | e2.PRED ) ] )
  ( _ ( e1 .. e2 ) | _ ( e2 ) )          [ e1.SUCC, ( e2 | e2.PRED ) ] )
```

generates these desugarings

```
Index [ 10 .. 20 ]      Index.MKSUBRANGE[10, 20]
Index [ 10 .. 20 )     Index.MKSUBRANGE[10, 20.PRED ]
( 1 .. 100 )           INT.MKSUBRANGE[1.SUCC, 100.PRED ]
```

Names introduced in the desugaring are written with one or more trailing prime ("') characters. Such names cannot be written in a Cedar program, and hence they are safe from name conflicts. The desugaring is constructed so that the ordinary scope rules prevent multiple uses of the same name from being confused.

### 3.1.3 Notation for the commentary

Each section of the commentary begins with grammar rules, desugaring and examples for part of the language. It continues with text which explains the meaning of the constructs. Generally the meaning is fairly clear from the desugaring, and this text is short. For blocks and especially for modules, however, there are many non-obvious implications of the desugaring, and a number of restrictions; these constructs have a lot of explanatory text.

Some kinds of information are put into specially marked paragraphs, which begin with one of the following italicized words:

*Anomaly*: the meaning of this Cedar construct is not explained by desugaring into the kernel, but by the special rule given here.

*Caution*: here is an implication of the definition which might surprise you.

*Performance*: facts about the time or space required by some construct.

*Representation*: the values of a data type are represented in terms of other types.

*Restriction*: a construct is not fully general, and will cause a static error unless additional conditions stated here are satisfied.

*Style*: advice about good Cedar style.

Symbols written in `SANS-SERIF SMALL CAPITALS` are in the kernel but not in current Cedar. The superscript notation used to cross-reference non-terminals in the grammar is also used in the examples, usually to point to a rule whose example introduces a name.

## 3.2 The lexical structure of programs

```

56 name ::= letter (letter | digit)... -- But not one of the reserved words in Table 3 2.
57 literal ::= num(C|E|D)?num | -- INT literal, decimal if radix omitted or D, octal
digit (digit(C|E|F) ... | H)?num | -- INT literal in hex; must start with digit. |
?num . num ?exponent | -- REAL as a scaled decimal fraction; note no trailing
num exponent | -- With an exponent, the decimal point may be omitted.
' extendedChar | • digit !.. C -- CHAR literal; the C form specifies the code in octal.
" extendedChar .. |" ?•L [ ('extendedChar), ...] -- Rope.ROPE, TEXT, or STRING.
$ n -- ATOM literal.
58 exponent ::= ±e(E(+ | ) num -- Optionally signed decimal exponent.
59 num ::= digit !..
60 extendedChar ::= space | \ extension | anyCharNot'"Or\
61 extension ::= digit1 digit2 digit3 | -- The character with code digit1 digit2 digit3.
| \N | \R | \t | \b | -- CR, '\015 | TAB, '\011 | BACKSPACE, '\010 |
| \f | \l | ' (f" | \ -- FORMFEED, '\014 | LINEFEED, '\012 | ' | " | \

```

## Examples

```

m, x1, x59y, longNameWithSeveralWords: INT;
n: INT~1+12D+2B3+2000B -- = 1+12+1024+1024
+1H+0FFH; -- +1+255
r1: REAL~0.1+.1+1.0E 1 -- = 0.1+0.1+0.1
+1E 1; -- +0.1
a1: ARRAY [0..3] OF CHAR~['x, '\N, '\', '\141];
r2: ROPE~"Hello.\N...\NGoodbye\F";
a2: ATOM~$NameInAnAtomLiteral;

```

The main body of the grammar (rules 1-55) treats a program as a sequence of tokens. Rule give the syntax of most tokens. A token is:

- A literal<sup>57</sup>. More information about literals of type T can be found in ¶ 4, as part of the treatment of type T.
- A name<sup>56</sup>, not one of the reserved words in Table 3 2. Note that case is significant for names.
- A reserved word, which is a string of uppercase letters that appears in the list of reserved words in Table 3 2. A reserved word may not be used as a name, except in an ATOM literal.
- A punctuation symbol: any printing character not a letter or digit, and not part of the two-character sequences below. The legal punctuation symbols in programs are:

! @ # \$ % ~ \* + = | ( ) { } [ ] \_ ^ ; : ' " , . < > /

The following ASCII characters are not legal punctuation symbols (and must not appear in a program except in an extendedChar<sup>60</sup>):

% & \ ?

Note that Cedar uses a variant of ASCII which includes the characters  (instead of the circumflex) and  $\hat{e}$  (instead of the underbar). Note also that the character written here is the ASCII minus character, code 55B, and not any of the various dash or typographer's minus characters with other codes, which are not in the standard ASCII set.

- One of the following two-character symbols (used in the grammar rules indicated):
  - ~= not equal<sup>22</sup>
  - <= less than or equal<sup>22</sup>
  - ~< not less than<sup>22</sup>
  - >= greater than or equal<sup>22</sup>
  - ~> not greater than<sup>22</sup>
  - => chooses<sup>8, 17, 30, 31, 33, 35, 52</sup>
  - .. subrange constructor<sup>25, 51.1</sup>
  - ~~ bind by name<sup>6, 34</sup>

---



---

ABS	EXIT	MONITORED	RETURNS
ALL	EXITS	NARROW	SAFE
AND	EXPORTS	NEW	SELECT
ANY	FINISHED	NIL	SEQUENCE
ARRAY	FIRST	NOT	SHARES
ATOM	FOR	NOTIFY	SIGNAL
BASE	FORK	NULL	SIZE
BEGIN	FRAME	OF	START
BOOL	FREE	OPEN	STATE
BOOLEAN	FROM	OR	STOP
BROADCAST	GO	ORDERED	STRING
CARDINAL	GOTO	OVERLAID	SUCC
CEDAR	IF	PACKED	TEXT
CHAR	IMPORTS	PAINTED	THEN
CHARACTER	IN	POINTER	THROUGH
CHECKED	INLINE	PORT	TO
CODE	INT	PRED	TRANSFER
COMPUTED	INTEGER	PRIVATE	TRASH
CONS	INTERNAL	PROC	TRUSTED
CONTINUE	ISTYPE	PROCEDURE	TYPE
DECREASING	JOIN	PROCESS	UNCHECKED
DEFINITIONS	LAST	PROGRAM	UNCOUNTED
DEPENDENT	LENGTH	PUBLIC	UNTIL
DESCRIPTOR	LIST	READONLY	USING
DIRECTORY	LOCKS	RECORD	WAIT
DO	LONG	REF	WHILE
ELSE	LOOP	REJECT	WITH
ENABLE	LOOPHOLE	RELATIVE	ZONE
END	MACHINE	REPEAT	
ENDCASE	MAX	RESTART	
ENDLOOP	MIN	RESUME	
ENTRY	MOD	RETRY	
ERROR	MONITOR	RETURN	

---



---

Table 3 2: Reserved words and predefined names

The program is parsed into tokens by starting at the beginning and successively taking front the longest sequence of characters which forms a token according to the rules above, discarding any number of initial whitespace characters or comments.

The whitespace characters are space, tab, and carriage return. A Tioga node boundary is also treated as a whitespace character.

A comment is one of:

A sequence of characters beginning with --, not containing -- or a carriage return and ending either with -- or with a carriage return.

A Tioga node with the comment property.

Note that whitespace and comments are not tokens, but may appear before or after any token. Whitespace and comments are token delimiters, and hence cannot appear in the middle of a token. Whitespace and comments thus do not affect the meaning of the program except:

When they delimit a token.

Within a CHAR literal or a ROPE literal, where they are taken literally. Thus ' ' is '\040, and "I am --not--" is equal to "I\Nam --not--" and different from "I\Nam ".

Both reserved words (Table 3 2) and most names with predefined meanings (Table 4 5) are up entirely of upper case letters. They should not be rebound by the program; in some bu cases the compiler forbids their rebinding. All are at least three characters long excep following:

```
DO GO IF IN OF OR TO.
```

A note on lists of items and their separators:

Semi-colons are used to separate declarations, bindings and statements in a body<sup>10</sup>, separate choices in a select statement<sup>29, 32, 34</sup> or in an exits<sup>6, 17</sup> or enable<sup>8, 27.1</sup>.

Commas are used to separate declarations in fields<sup>43, 51</sup> (i.e., in a proc domain or recordTC or a unionTC), bindings in an application<sup>27</sup> or an open<sup>7</sup>, choices in a sele expression<sup>29, 32, 34</sup> or in a unionTC<sup>52</sup>, expressions in a choice<sup>6, 9, 17, 30, 35, 52</sup>, items i exports or shares lists<sup>2, 3</sup>.

In general these sequences may be empty, and an extra separator at the end is harmless w is some kind of closing bracket, except when the sequence is bracketed with [].

The braces which delimit a block<sup>6</sup>, interface body<sup>2</sup>, choices in an enable<sup>8</sup>, or MACHINE CODE may be replaced by BEGIN and END reserved words. BEGIN replaces "{" and END replaces "}". I one brace is replaced, its matching partner must also be replaced. The braces delimiting enumTC<sup>54</sup> may not be replaced by BEGIN and END.

### 3.3 Modules

```
1 module ::= DIRECTORY (nd (: TYPE (nt | [ (nd : ( (TYPE nt | TYPE nd) | TYPE nd), ... ] IN
      ?(USING LETu(nd~RESTRICT)nd, .[$nu; ... ] ] ), ...
      ( interface | implementedinterface | implementation ) )
2 interface ::= nm, !.. : ?CEDAR DEFINITIONSR(~[ nm: INTERFACETYPE[[ nm, ...]] ] IN (imports | l
      ?locks (imports | ) ?*(SHARES ns, ...)- SHARES allows access to PRIVATE names in ns.
      ~ ?*access12 { ?open7 (d | b); !.. } LET REC nm~open [ ?(l(~locks, ) (d | b), ... ] IN MKIN
3 implementation ::= nm : ?CEDAR      LET r(~REC [(ne: ne) , ... , FRAME: TYPE nm ,
      nm: FRAME, CONTROL: PROGRAM]
      ?safety ( PROGRAM ?drType42 |      IN (imports | l=>r() IN
      MONITOR ?drType42                (( || (locks)LOCK:MONITORLOCK IN LET_l(IN(LOCK) IN | ))
      (imports | )                       LET b(~NEWPROGINSTANCE[block].UNCONS IN
      ?(EXPORTS ne, ...)                  [ (ne~BINDDFROM[ne, b[] ), ... , FRAME~MKINTTYPE[bl
      ?*(SHARES ns, ...)                  nm~b( , CONTROL~b(.nm] where the block body is
      ~ ?*access12 block .                [( | ( | l(~locks,)) (d | b), ... , nm: PROGRAM
3.1 imports ::= IMPORTS ( (niv : | ) nit l, [(nit~nit), 3...]=>r( IN LET (niv | nit)~ nit PLUS nit.
4 safety ::= SAFE | UNSAFE --In 3, 41.
5 locks ::= LOCKS e ?( USING nu: t)    l ?( [nu: t] ) IN e
```

## Examples

```

DIRECTORY                                     -- For BufferImpl below.
Rope: TYPE USING [ROPE, Compare],           -- There should always be a USING clause
CIFS: TYPE USING [OpenFile,Error,Open,read,  -- unless most of the interface is used
IO: TYPE IOStream,
Buffer: TYPE;                               -- or it is exported.

Buffer: DEFINITIONS ~
Handle: TYPE~REF BufferObject;
BufferObject: TYPE=Rope.ROPE
New: PROC RETURNS[h: Handle];
Get: PROC[h: Handle] RETURNS[BufferObject];
Put: PROC[h: Handle, o: BufferObject];

BufferImpl: MONITOR [f: CIFS.OpenFile]-- Implementations can have arguments.
LOCKS Buffer.GetLock[h]^                    -- LOCKS only in MONITOR, to specify
  USING h: Buffer.Handle                     -- a non-standard lock.
IMPORTS Files: CIFS, IO, Rope              -- Note the absence of semicolons.
EXPORTS Buffer                             -- EXPORTS in PROGRAM or MONITOR.
~ { -- module body -- } .                 -- Note the final dot.

```

Modules serve a number of functions (which might perhaps better be disentangled, but are

A file of text (BufferImpl.mesa), or its translation into object code (BufferImpl.bcd

The unit handled by the editor, named in DF files and models, and accepted by the compiler, the binder, and the loader.

A set of related structures (types, procedures, variables) which are freely access other, hiding secrets or irrelevant information from other modules.

A procedure which can accept interface types and bindings as arguments, and return interface values as results.

The first two uses are not relevant to the language definition, and are not discussed further. The others are the subject of this section.

There are two kinds of modules: interface modules (written with DEFINITIONS) and implementation modules (written with PROGRAM or MONITOR). They have the same header (except that interface modules have no EXPORTS list); it defines the parameters and results of the module viewed as a procedure and specifies the name  $n_m$  of the module. The bodies (following the ~) are different. Table 3 summarizes the structure of modules and their types; it omits a number of details which are given in rules 1-3 and explained in the text.

Example	Module	Module type	Result	Result type
DIRECTORY Rope, IO; Match: DEFINITIONS~{...}	Interface module	[Rope: TYPE Rope, IO: TYPE IO]_[TYPE Match]	Interface	TYPE Match
DIRECTORY Match, Rope, IO MatchImpl: PROGRAM IMPORTS R: Rope, I: IO EXPORTS Match~{...}	Implementation module	[Match: TYPE Match, IO: TYPE IO, Rope: TYPE Rope]_[Match]	Exported instance	Match

Table 3 3: Interface and implementation modules



The ensuing sub-sections deal in turn with:

- ¶ 3.3.1: Modules as procedures, and the interface or instance values obtained by applying them.
- ¶ 3.3.2: How modules are applied.
- ¶ 3.3.3: Module parameters: the DIRECTORY and IMPORTS lists; USING clauses.
- ¶ 3.3.4: Interface module bodies and interfaces.
- ¶ 3.3.5: Implementation module bodies; the EXPORTS list.
- ¶ 3.3.6: SHARES and access<sup>12</sup>.

The meanings of the other parts of a module header are discussed elsewhere:

CEDAR in ¶ 3.4.4.

MONITOR and LOCKS in ¶ 4.10.

### 3.3.1 Modules and instances

A module is a proc which takes two kinds of arguments:

Interfaces, declared in the DIRECTORY list. These arguments are supplied by the model (on the compiler's command line),

Instances of interfaces, declared in the IMPORTS list. These arguments are also supplied to the model (or in a config file passed to the binder, or implicitly by the loader).

¶ 3.3.3 discusses the types of these arguments and how they are declared. In addition, an implementation may take PROGRAM arguments declared in the drType following PROGRAM or MONITOR. These are ordinary values; they are discussed in ¶ 3.3.2.1.

When a module is applied to its arguments, the resulting value is

For an interface module, an interface.

For an implementation module, a binding whose values are instances:

one interface instance for each interface it exports;

one for the program instance, also called a global frame;

one for the program proc derived from the module body (¶ 3.3.2.1), called CONTROL.

This application cannot be written in the program, only in the model; it is described in

An interface (sometimes called an interface type) is a type, as the latter name suggests. Its declaration (obtained from the declarations which constitute the module body), with an enclosing cluster that includes all the bindings in the module body that don't use declared names. In the example, the Buffer interface (obtained by applying the Buffer module to the argument, in its DIRECTORY) has declarations for New, Get, and Put, and its cluster includes values for New and BufferObject.

An interface instance is a value whose type is an interface; such values are the results of instantiating implementation modules. In the example, BufferImpl returns (exports) an instance of Buffer.

A program instance or a global frame is a frame, as the latter name suggests, i.e., a binding from the bindings and declarations of an implementation (PROGRAM or MONITOR) module body, just like any proc frame (¶ 3.3.5). Normally code outside the module does not deal with frames directly, but only with the exported interface instances. In the example, BufferImpl exports a program instance for the module and a CONTROL proc.

In most cases, there is:

- Exactly one application of each module, and hence exactly one interface or one instance.
- Only one module which exports an interface.
- Only one interface exported by a module.
- Only one argument of the proper type for each module parameter (¶ 3.3.3); hence it is redundant to write the arguments explicitly.

When these conditions hold, there is a close correspondence among the following four objects:

- an interface module;
- the interface it returns (since its arguments need not be written explicitly);
- the implementation module which exports the interface;
- its instance (again, since its arguments need not be written explicitly).

The distinctions made earlier in this section then seem needless; it is sufficient to distinguish the interface and implementation modules, and identify them with the files which hold them. In more complicated situations, however, it is necessary to know what is really going on.

In the example at the start of this section, `BufferImpl` is an implementation module with four parameters:

- Four interface parameters, declared in the `DIRECTORY:` `Rope`, `CIFS`, `IO` and `Buffer`.
- Three instance parameters, declared in the `IMPORTS:` `Files` (of type `CIFS`), `IO` (of type `IO`) and `Rope` (of type `Rope`). Since the instance parameters are declared in an inner scope, only `Rope` is visible in the module body; the interface `Rope` is visible in the header. The same is true for `IO`, but both the interface `CIFS` and the instance `Files` are visible in the body.

When `BufferImpl` is compiled, the four interface parameters must be supplied, in the form of (compiled) interface modules named `Rope`, `CIFS`, `IO` and `Buffer`. When `BufferImpl` is instantiated (normally by loading it), the three instance parameters must be supplied, i.e. there must be instantiated implementation modules which export the `Rope`, `CIFS`, and `IO` interfaces. Normally there will be one of each, and the entire program will consist of eight modules:

- the interface modules `Rope`, `CIFS`, `IO` and `Buffer`;
- implementation modules normally named `RopeImpl`, `CIFSImpl`, `IOImpl` and `BufferImpl`, each exporting an instance of the corresponding interface

The instantiated `BufferImpl` exports an instance of `Buffer`, which can then be used as a parameter to some other module.

### 3.3.2 Applying modules

A module is not applied to all its arguments at once. Instead, the arguments are supplied in two stages:

A module is applied to its interface (`DIRECTORY`) arguments by compiling it; the result is a BCD (represented by a `.bcd` file). The bcd is still a proc, with instance parameters. A proc, a module can be applied to different arguments (i.e., different interfaces) to produce different BCDs.

A BCD is applied to its instance (`IMPORT`) arguments by loading (or binding) it; the result is a program instance, together with any interface instances exported by the module. The BCD can be applied to different arguments (i.e., different interface instances) to produce different instances. Indeed, because an instance may include variables, even two applications to the same arguments will yield different instances.

These two stages are separated for several reasons:

All the type-checking of a module can be (and is) done in the first stage, by the compiler. The only type error possible in the second stage is supplying an unsuitable argument.

Compiling is much slower than loading, and a module needs to be recompiled only when its interface arguments change, not when the interface values change. The latter are in the implementations of the interfaces, and are much more common.

When there are multiple instances of the same module with the same interface parameters, they automatically get the same code.

We've always done it that way.

### 3.3.2.1 Initializing a program instance

The statements in the body of an implementation module form the body of a procedure called `PP`. The function of this procedure is to initialize an instance of the module. An instance `PI` may be uninitialized, because no code in the module is executed when the instance is made. It is the job of the procedure `PP` to initialize `PI`, perhaps using the `PROGRAM` arguments if there are any. Until `PP` has been called, `PI` is not in a good state. It would be better if the `PROGRAM` arguments along with the imported instances, and call `PP` as part of making `PI`, that `PI` is never accessible in its uninitialized state. But it isn't done that way; hence the programmer must ensure that `PP` is called before any use is made of `PI`. To confuse things, `PP` is not an ordinary procedure but a `PROGRAM`, and it must be called using the `START` construct (see § 4.4.1). Note that in addition to the statements of the module body, `PP` also contains specific initialization code for any variables or non-static values in the instance; e.g. the value of `x` will not be 3 until after `PP` has been called.

There is some error detection associated with this kludge. If a procedure in the instance is called before the instance has been initialized by `START`, a start trap occurs. At this point, if `PP` takes arguments it is called automatically, and the original call then proceeds normally; if `PP` takes no arguments, there is a `Runtime.StartFault ERROR`.

Caution: If the module is a monitor, `PP` runs without the monitor lock; if another procedure calls the module while `PP` is running, it will not wait, but will run concurrently with `PP`. This is unlikely to be right. It is unwise to rely on a start trap to initialize a monitor module; initialize it explicitly with `START`.

Caution: If a variable in the instance is referenced before the instance has been initialized, a start trap is detected, and the uninitialized value will be obtained. `PP` can still be called to initialize the instance, and may still be called automatically by a start trap.

The procedure `PP` is bound to the name `CONTROL` in the result of an implementation module if the module type is `PROGRAM[] RETURNS[]` (otherwise the procedure `Runtime.ReportStartFault` is bound to `CONTROL`). This allows the modeller (and binder) to get access to `PP` so as to control the order in which modules are started.

### 3.3.3 Parameters to modules: `DIRECTORY` and `IMPORTS`

The interface parameters of a module are declared in the `DIRECTORY`. An interface `I` has type `T`, where `n` is any one of the names given before `DEFINITIONS` in the header of the interface that produced `I`. The `INTERFACETYPE` primitive in the desugaring takes a list of names and returns a type which implies `TYPE n` for each `n` in the list. The reason for allowing several names is the conversion of an interface from one name to another; both names can continue in use for the same interface. The use of these names provides a clumsy check that the proper interface is supplied as an argument. `DIRECTORY n: TYPE` and `DIRECTORY n` are both short for `DIRECTORY n: TYPE n`.

An interface is a type which can only be used:

Before a dot (¶ 4.14), to obtain a value from its cluster, which simply consists of bindings in the interface module body (¶ 3.3.4).

In an `IMPORTS` list as the type of an instance parameter to a module.

The `USING` clause in the `DIRECTORY`, if present, restricts the cluster of the interface to only items with the names  $n_u$ , ... Thus in the example, only `ROPE` and `Compare` are in the cluster in the `BufferImpl` module. This means that `Rope.ROPE` and `Rope.Compare` are legal, but `Rope.n` for other `n` will be an error. Note that `USING` affects only the cluster of the parameter; it does not affect the clusters of any types or the bodies of any `INLINE` procs obtained from the interface. `Rope`, `Compare` might be bound by

```
Compare: PROC[r1, r2: ROPE] RETURNS [BOOL]~INLINE {
  IF Length[r1]~=Length[r2] THEN ... }
```

A call of `Rope.Compare` in `BufferImpl` is perfectly all right, even though `Rope.Length` in `BufferImpl` would be an error.

In the example, `CIFS`, `IO`, and `Rope` are interfaces. They are the types of three `IMPORTS` parameters named `Files`, `IO`, and `Rope` (if the `IMPORTS` clause gives no name for the parameter, the name of the interface is recycled). An actual argument for an `IMPORT` parameter must be an interface instance, i.e., a value whose type is an interface type. Such a value is obtained from one or more modules which export the interface (¶ 3.3.5). An instance is a binding; in this binding the value declared in the interface is provided by the exporter; the value of a name bound in the instance (e.g., `x~3`) is just the value that the interface binds to the name (in this case, the value has two effects:

The client can ignore the distinction between names bound and declared in the instance since both appear in the instance binding and are referenced uniformly with dot notation. This means that the client is not affected, for example, when a proc is moved from `INLINE` in the interface to an ordinary definition in an implementation.

The client can often ignore the distinction between the interface and the instance since the values in the interface are also in the instance, with the same names. This is the motivation for the shorthand which allows the name of an `IMPORT` parameter to default to the name of the interface; the interface is no longer accessible, but `I.x` has the same meaning (namely 3) whether `I` is the interface or the instance.

Anomaly: Names bound to inline procs in an interface do not appear in the interface binding only in an instance. This somewhat dubious rule ensures that clients won't have to add imports lists if a proc stops being an inline.

Restriction: An interface module may not import more than one instance of a given interface `I`. If an implementing module `P` imports more than one instance of `I`, the principal instance of `I` is the one with no name in the `IMPORTS` list (which is therefore named `I` by default). If `P` imports only one instance of type `I`, then that instance is the principal instance.

Restriction: Often an interface module has no `IMPORTS`, because it only needs access to the static values (types and constants) bound in its interface parameters, and does not need values for any names declared there (procs and variables). If an interface module does have `IMPORTS`, however, and there is more than one instance of any imported interface around, then there is a restriction on the argument values. Suppose that `Int1` imports `Int2`, and that module `P` imports `Int1`. Then `Int1` may only import one instance of `Int2`, and if `P` also imports `Int2`, the principal instance of `Int2` in `P` must be the same as the value of `Int2` imported by the `Int1` imported by `P`. For example, with

```
DIRECTORY Int2; Int1: DEFINITIONS IMPORTS Int2V: Int2 ...
DIRECTORY Int1, Int2; P: PROGRAM IMPORTS Int1V: Int1, Int2V: Int2 ...
```

we must have in `P` that `Int1V.Int2V=Int2V`.

### 3.3.4 Interface module bodies

The body of an interface module `I` is a collection of bindings (e.g., `x: INT~3`) and declarations (e.g., `y: VAR INT` or `P: PROC[a: INT] RETURNS [REAL]`).

Restriction: Only certain things may follow the `~` in one of the bindings<sup>13</sup>:

If it is an expression, it must be static (¶ 3.9.1).

If it is a block (providing the body of a proc), it must be `INLINE` (because there is no place to put the compiled code).

It may not be `CODE`.

The result of applying an interface module is an interface (¶ 3.3.2), which is a type `I`. Applying the primitive `MKINTTYPE` to the `d`'s and `b`'s of the body. This type is simply the `d` obtained by collecting the declarations in the body, with a cluster which is extended to include the bindings of the body. However, `MKINTTYPE` omits any inline proc bindings from the type's cluster, instead leaving the proc declarations in `I`. It puts an extra item `BINDING` in `I`'s cluster, the inline procs in it. When an instance `Inst` of `I` is imported, the binding actually imports `PLUS I.BINDING`. This slightly dubious arrangement ensures that clients don't have to change their lists if a proc stops being inline. This policy is not extended to other items, however, they might change from being bound in the interface to being interface variables.

The interface returned by

```
Red, Blue, Green: DEFINITIONS~...
```

has the types `TYPE Red`, `TYPE Blue`, and `TYPE Green`.

The types and expressions in declarations and bindings may refer to other names in the body, as usual, but they may not refer to names introduced in the declaration, except that:

Any declared name may be used

in the body of an `INLINE`, or

after a `"_"` in a defaultTC<sup>55</sup> in the fields<sup>43</sup> of a transferTC<sup>41</sup> which is the type of the decl in the interface's body.

A declared (opaque) type may be used anywhere.

For example, if an interface contains

```
I: DEFINITIONS~
  x: INT~3;
  y: VAR INT;
  T: TYPE[ANY]
```

then the following may also appear in the interface:

```
xx: INT~x+1;
P: PROC RETURNS[INT]~INLINE {RETURN[x+y]};
Q: PROC [INT_y];
V: TYPE~RECORD[f: REF T, g: U]
```

but the following are illegal:

```
xy: INT~y+1;
U: TYPE~INT_y;
W: TYPE~ARRAY [0..y] OF INT;
```

The values of the bindings can be accessed directly by dot notation in any scope in which the interface value is accessible. Thus if the value of the previous interface module is bound to `J` because `J: TYPE I` appeared in the `DIRECTORY`, then `J.x` is equal to 3. The declarations cannot be accessed directly (`J.y` is an error).

The declarations in an interface module are not quite like ordinary declarations. They are of several kinds, depending on whether the type of a declaration is:

A transfer type; this is just like a declaration of a transfer parameter to an ordinary module, except that it is readonly.

`TYPE[ANY]` or `TYPE[e]`; the type being declared is an opaque type or exported type, discussed in ¶ 4.3.4. The expression `e` must be static. `TYPE[ANY]` or `TYPE[e]` is not allowed in an ordinary declaration; except in an interface, a type name must be bound to a type when it is introduced.

VAR T, OR READONLY T for any type T except TYPE; this is an interface variable; discussed in ¶ 3.3.4.1 below. •In an ordinary declaration in a block you can't write VAR T, but you can write simply T; you can also write simply T here, but this is not recommended

An interface instance II has the interface type I if for each item n: T in the interface, item n~v in the instance, and v has type T. This is the same rule which determines that a variable has the type of a declaration; e.g., that a proc argument has the domain type. In this respect there is nothing special about an interface.

Note that a name can be declared PRIVATE in an interface, even though it must be declared in the exporter (¶ 3.3.6). This can be useful if the name is used in a type constructor in the interface, but its value should not be accessible to the client.

### 3.3.4.1 Interface variables

An interface variable v gives clients of an interface direct access to a variable in a package, namely the variable which is exported to v. This is the only kind of variable parameter in Cedar.

•If you use the obsolete shorthand of T for VAR T in an interface variable declaration, you must declare a transfer type variable as an interface variable, since that already means pass-by-value.

Caution: the variable which is exported to provide the value for an interface variable is not initialized until its module is initialized (¶ 3.3.2.1). However, there is nothing to stop it from being accessed sooner.

Performance: An interface variable can be read and (if not READONLY) set directly, which is significantly faster than Get and Set procs. Of course, the implementor gives up some control; it is not quite as fast as access to an ordinary variable, since there is an extra level of indirection. It costs one or two extra instructions each time. There is also one pointer per interface variable in each module which refers to it. If you use a private interface variable and inline Get and Set procs, you pay nothing in performance, but retain the option of changing the proc definitions later.

•You can get direct access to all the variables of a module by using a POINTER TO FRAME type (¶ 4.5.3).

### 3.3.5 Implementation module bodies

The body of an implementation module Imp is simply a block. This block plays two roles. On one hand, it is an ordinary block, the body of an almost ordinary proc PP( called the PROCPROC, which has parameters and results like any other. PP( is special in one way: it has a type rather than a PROC type. When PP( is applied (using the special construct START; see § 3.3.2), its declarations and bindings are evaluated, its statements are executed, and its result is returned as with any proc. The only difference is that the values bound to the names introduced in the block (i.e., the frame of PP( are retained after the proc returns; in fact, forever (unless RETURN is used to free the frame). Procs local to the block can access these values in the usual way. Names of exported names can also be accessed through interfaces, as explained below; see ¶ 3.3.4.1.

As with any proc (¶ 3.5.1), the frame of PP( includes the parameters and results from Imp as well as the names introduced in the block's d's and b's. It also includes an additional pointer to the frame of the caller.

```
Imp: PROGRAM T~PP(
```

where Imp is the name of the module and T is its drType.

The body of Imp has a second role: to supply values for the names declared in the interfaces exported by Imp. For each interface Ex which Imp exports, an interface value ExI of type Ex is constructed. Each name n in ExI acquires a value as follows:

If  $n: T$  is in  $Ex$  and  $n: \text{PUBLIC } T \sim x$  is in the body of  $\text{Imp}$ , then  $n \sim x$  is in  $ExI$ . This is a slightly peculiar kind of binding; as in an ordinary binding,  $x$  must be coercible (¶ 4.13). Note that  $n$  must have **PUBLIC** access (¶ 3.3.6) in the body.

If  $n$  is declared in  $Ex$  and not bound in the body of  $\text{Imp}$ , then  $n \sim \text{UNBOUND}$  is in  $ExI$ . **UNBOUND** is a special value with the following properties:

- For a proc  $P$ , it causes a `Runtime.UnboundProcedure` signal on any application of  $P$ .
- For a variable  $v$ , it causes a `Runtime.PointerFault` error on any reference to  $v$ .
- For a type  $T$ , it causes no problem.

If  $n \sim x$  in  $Ex$ , then  $n \sim x$  in  $ExI$ . Thus any names bound in the interface are bound the way in any interface value.

Caution: A name can be exported to several interfaces without any warning, if it has a special value. This is unlikely to be what is wanted.

The result of instantiating  $\text{Imp}$  is a binding with:

One item for each exported interface  $Ex$ , namely  $Ex: Ex \sim ExI$ , where  $ExI$  is the interface value constructed above. Here  $Ex$  is the name  $n_d$  given to the interface in the `DIRECTORY`.

One item `CONTROL: PROGRAM[] RETURNS []`, whose value is the program proc `PP` (if that has no arguments, and otherwise `Runtime.ReportStartFault`).

- One item for the type of the module's global frame, namely `FRAME ~ TYPE Imp`.

- One item for  $\text{Imp}$  itself, namely `Imp: FRAME`. The value of this item is the program instance, i.e., the frame of the module's body.

This binding is accessible in a model, where it can be used to get access to the interface, the program proc, the global frame type, and the program instance.

- You can pass `FRAME` as an argument to a `DIRECTORY` parameter `I: TYPE Imp`; like an interface; provides access to constants bound in the module, and allows you to declare an `IMPORTS` parameter whose argument will be a program instance of the module. From `I` you can also obtain a file-like Cedar type `POINTER TO FRAME[I]`; see ¶ 4.3.5. `I`'s cluster includes a coercion from `I` to `POINTER TO FRAME[I]`, and the proc `COPYIMPLINST` (applied by the funnyAppl `NEW`), which is the same as the proc of the same name in cluster of `POINTER TO FRAME[I]`.

- You can import  $\text{Imp}$  into another module (by writing `DIRECTORY Imp ... IMPORTS ImpInst: Imp ..`) and obtain access to all the variables and procs of the program instance.

### 3.3.6 PUBLIC, PRIVATE and SHARES

Cedar has a rather complicated mechanism for controlling access to names. Most uses of it are considered to be obsolete, with the following exceptions:

Names to be exported must be declared **PUBLIC**.

Names included in an interface for use in inline procs etc., but not intended for clients, should be declared **PRIVATE**.

Access to a name is declared by writing **PUBLIC** or **PRIVATE** right after the colon in a declaration of a name:

```
x: PUBLIC T
```

In the Cedar syntax these colons occur in the declarations<sup>11</sup> and bindings<sup>13</sup> in bodies<sup>10</sup>, in interface modules<sup>2</sup>, and in the tag<sup>53</sup> of a unionTC. You can set a default access for a name in a module<sup>2, 3</sup> or record<sup>50</sup> by writing **PUBLIC** or **PRIVATE** just before the `{` or `RECORD;` that is overridden by an explicit **PUBLIC** or **PRIVATE** inside. By default, an interface is **PUBLIC** and an implementation is **PRIVATE**.

A PRIVATE name defined in module M can only be referenced:

from within M;

from a module which EXPORTS M.

•from a module which SHARES M; avoid this feature.

This does not mean that the name is invisible, but rather that it is an error to use it OPENED. Thus in

```
x: INT; {OPEN M; f [x]}
```

if x is bound in M (and not hidden by a USING clause), the call of f is equivalent to f [M] regardless of whether x is PUBLIC or PRIVATE. It is illegal if x is PRIVATE, but it never redeclares x declared by the x: INT.

Furthermore, if a record has any PRIVATE components, a constructor or extractor for the record is legal only in a module where use of the PRIVATE names is legal (even if the private components are not mentioned and have defaults).

### 3.4 Block OPEN and ENABLE

```
6 block ::= ?(CHECKED | UNCHECKED | TRUSTED)
  { ?open ?enable ?body
    ?(EXITS (n, !..=>s); ...) }
  --In 3, 13, 15.
7 open ::= OPEN ( n ~~ e | e ), !.. ; ( LET n~lopen IN e.DEREF |
  In 2, 5, 17. •The ~~ may be written as :.
8 enable ::= ENABLE ( enChoice |
  In 5, 17.
9 enChoice ::= ( e, !.. | ANY ) => s ( e | ANY ), ... => { s; REJECT; EXITS
  In 7, 27.1.
10 body ::= ( d | b ); !.. ; s; ... | SET NEWFRAME[ REC [(d | b), ...] ].UNCONS IN { s; ... }
  In 5, 17.
```

### Examples

```
CHECKED {
  OPEN Buffer, Rope;
  ENABLE Buffer.Overflow=>GOTO HandleOvfl;
  stream: IO.Stream~IO.CreateFileStream[...];
  x: INT_7;
  {OPEN b~~buffer;
    ENABLE {
      Files.Error--[error, file]--=>{
        stream.Put[IO.ropes[error]];
        ERROR Buffer.Error["Help"] };
      ANY=>{ x_12; GOTO AfterQuit } };
  y: INT_9; ... };
  x_stream.GetInt; ...
EXITS
  AfterQuit=>{...};
  HandleOvfl=>{...} };
```

-- Unnamed OPEN OK for exported interface or one with a USING clause.  
 A single choice needn't be in {}.  
 a binding if a name's value is fixed.  
 -- Better to initialize declared names.  
 -- A statement may be a nested block.  
 -- Multiple enable choices must be in {}.  
 -- ERRORS can have parameters.  
 -- Choices are separated by semicolons.  
 -- ANY must be last. ENABLE ends with ;.  
 -- Other bindings, decls and statements.  
 -- Other statements in the outer block.  
 -- Multiple EXIT choices are not in {}.  
 -- AfterQuit, HandleOvfl declared here,  
 -- legal only in a GOTO in the block.

The main function of a block is to establish a new scope (¶ 2.3.4) and to allow for the variables declared in the block, as in Algol or Pascal. A Cedar block has four other fea



attributes: CHECKED, UNCHECKED and TRUSTED are treated in ¶ 3.4.4 on safety.

open<sup>7</sup>: a combination of sugar for LET and call by name; see ¶ 3.4.2.

enable<sup>8</sup>: catches signal and error exceptions in the body; see ¶ 3.4.3.1.

EXITS: catches GOTO exceptions in the body or enable; see ¶ 3.4.3.2.

Note that the braces around a block may be replaced by BEGIN and END (¶ 3.2).

The statements in a block are evaluated in the order they are written. The initialization in the d's and b's are also evaluated in the order they are written; this may be important if they have side effects, although that should be avoided.

### 3.4.1 Scope of names and initialization

The names introduced in the block body's d's and b's (i.e., appearing before a : or ~) are bound to the body with the values supplied by the d's and b's, except in inner scopes where they are reintroduced; they are not known elsewhere in the block. The frame of the block is a binding of a value for each such name.

Actually, the frame is a value of an opaque type which has a coercion (called UNCONS) to this binding. As the decl for body indicates, the frame is constructed (by NEWFRAME), and then a LET makes the names in the binding known in the statements of the body.

Anomaly: A name introduced by a binding, n: T~e, has the value of e throughout the body if e is static. If e is not static, it is evaluated after all preceding d's and b's, but before n's binding. This means that n.VALUEOF is trash in all the d's and b's before its binding. Symmetrically, if n is introduced by a following decl or non-static binding, it will get a trash value. Compiling with the "u" switch will yield a warning in this case. Note that only attempts to get the value of n get trash; n may appear anywhere in a l-expression, and all will be well as long as the l-expression is not applied before n's binding is evaluated.

A name introduced by a declaration, n: T, is bound to a new VAR T. The variable bound to n is allocated, and its value is set to nil before anything in the block is executed (this is done by the NEWFRAME and the desugaring).

Anomaly: However, the INIT proc is executed (to set a REF or transfer value to NIL), and a defaultTC<sup>55</sup> in T is done at the same time that a non-static binding would be evaluated. As with a binding, n.VALUEOF is trash before this time. Furthermore, a (unwise) assignment to n before this time will be overridden by the defaultTC.

Caution: The failure to initialize RC variables is a safety loophole, since the trash can be dereferenced and used as an address.

Style: The expression in a binding or defaultTC should be functional, or at least it should have benign side-effects. There is no enforcement of this recommendation, unfortunately. In Cedar such an expression is evaluated exactly once, at the time described above. This may change in the future, however.

The variables created by a declaration are deallocated when execution of the block is complete, unless the block's frame is retained. Currently only an implementation's block<sup>3</sup> has its frame retained. There are two ways to hang on to a variable v after execution of the block is complete:

Obtain a pointer to v with @; this pointer value can survive the block.

Obtain a proc value for a local procedure which refers to v; this proc value can survive the block.

In the checked language both these dangling references are impossible: the @ operator, `block unsafe`, is forbidden, and `ASSIGN` for proc values gives an error unless the proc is local to the block instance (which has a retained frame).

Caution: An unchecked program can get into trouble with dangling references to frames, `block unsafe`.

Performance: There is no overhead associated with block entry or exit, even if the block contains `open`, `enable` or `EXITS`. The only cost is for initializing the variables bound to its names. The recommended style to use blocks freely to limit the scope of names.

### 3.4.2 OPEN

There are two forms of `open`. The first, `n~e`, binds the name `n` to `lopen IN e.DEREF`. This is like `l IN e.DEREF`, except that there is a coercion from `n` to `n[]`. In other words, every time `n` is used, its value is obtained by evaluating `e.DEREF`. The effect is exactly like call by name in Algol 68. To remind you that this is not ordinary value binding. The value of `e.DEREF` is

```
e if the cluster of De does not include DEREFERENCE or UNWRAP;
e^.DEREF if it includes DEREFERENCE;
e.UNWRAP.DEREF if it includes UNWRAP.
```

In other words, a reference value is dereferenced and a single-component record or binding is replaced by the component, repeatedly if necessary, to obtain a non-reference value. In `e.DEREF` must be a record, interface or instance.

The second, nameless, form of `open` gives an expression without binding it to a name: `{ open e }`. `e.DEREF` must evaluate to a binding `b`:

```
A record value has a corresponding binding (returned by UNCONS in the desugaring) which
has the names of the record fields bound to the field values (or variables, for a
binding).
An application returns a binding, though the call-by-name feature makes it unwise
to use this in an application in an open.
```

```
An interface or instance value is a binding (¶ 3.4.2).
```

The nameless `open` converts `b` into another binding `bp` in which each value is a `lopen proc` (see above), and introduces `bp`'s names in the block with a `LET`. Thus in the program

```
R: TYPE~RECORD [a: INT_3, b: REAL_3.4]; r: R;
{ OPEN r; ... }
```

the names `a` and `b` are known in the body of the block, and have exactly the same meaning as in `r.b`.

Style: Nameless `open` should be used with discretion, with the smallest practicable scope. It is useful if the value being opened is very familiar, or heavily used, or both. Nameless `open` can cause confusion, since it is not obvious from the text of the program where to find the bindings. The names it makes known. It should never be used when evaluation of `e` has a side-effect.

The scope of an `open` is all the rest of the block, including any `enable` and any `EXITS`. A block may have several bindings or expressions. These are applied sequentially, so that those bound by earlier ones are known to the later ones as well as to the rest of the block.

## 3.4.3 ENABLE and EXITS

The `ENABLE` and `EXITS` constructs are two forms of sugar for exception handling (¶ 2.2.4). `ENABLE` catches signals and errors raised in the body (but not the `open`, `enable`, or `exits`). `EXITS` catches `GOTOS` in the body or `enable` (but not the `open` or `exits`). Both are in the scope of the `open`. Neither is in the scope of any names introduced in the body.

## 3.4.3.1 ENABLE

An `enable` has a chance to catch any signal or error raised in the block (and not caught at a higher level). A nearly identical construct can appear in an application<sup>26</sup>; the following explains both cases.

Each `enable` choice (`enChoice`<sup>9</sup>) has a list of expressions with exception values, `•` or `ANY`, `=>`. If `ANY` appears, it must be the last `enChoice`. If the exception is equal to one of the values or if `ANY` appears, the statement after the `=>` is executed. Control leaves this statement in the following ways:

A `REJECT` statement causes the exception to be the value of the block; it will then be propagated within the enclosing block, or if the block is a `proc` body it will be propagated to the application.

A `GOTO` statement sends control to the matching choice in the `EXITS`. There are three special cases<sup>16</sup>:

A `RETURN` is not allowed in an `enChoice`.

A `CONTINUE` statement ends execution of the current statement (in this case the block); execution continues with the next statement following. If the block is a `proc` body, the effect is the same as `RETURN`. You cannot write `CONTINUE` in a body's `d`'s or `b`'s.

•A `RETRY` statement begins execution of the current statement (in this case the block) over again at the beginning. You cannot write `RETRY` in a body's `d`'s or `b`'s.

The semantics of `CONTINUE` and `RETRY` follow from the desugaring of statement<sup>14</sup>.

A `RESUME` statement (signals only) is discussed below.

•If the statement finishes normally, a `REJECT` statement is then executed.

If a single expression `e` appears before the `=>`, then within the `enChoice` statement the names `De.DOMAIN` are declared and initialized to the arguments of the exception. With multiple expressions, or `ANY`, the arguments are inaccessible. •The use of `ANY` is not recommended.

Note that an error is caught by an `enChoice` with a matching exception value, not by one of the matching names. Normally an error `E` will be declared in some interface, its value will be set by a binding of the form `E: PUBLIC ERROR ... ~ CODE`, and both the signaller and the `enChoice` refer to this value by the name `E`. In this case, it is natural to think of the binding as a name. However, it is possible to have a different name for this exception value, e.g. by `ERROR ... ~ E`. It is also possible to bind some other exception value to `E` in a scope which is some `enChoice` examined when the signal is raised. Thus in the silly program

```
E: ERROR~CODE;
F: ERROR~E;
{ENABLE E=>{--Handler 1--...}};
  E: ERROR~CODE;
  {ENABLE E=>{--Handler 2--...}};
  IF switch THEN ERROR F ELSE ERROR E;
```

if `switch` is true handler 1 will be used, and if it is false handler 2 will be used.

## Finalization

You are supposed to think of an `ERROR` as an unusual value `ev` which can be returned from an application; this value immediately stops the evaluation of the containing application, likewise returns `ev` as its value. This propagation is stopped only by an `enable` choice with the `ERROR`. As each application is stopped, it is finalized. Aside from invisible housekeeping, finalization confusingly consists of executing an `enChoice` which catches the `ERROR UNWIND`. A programmer can write any cleanup actions he likes in this statement.

**Caution:** If the finalization raises another `ERROR` which it does not catch, it will itself fail with very confusing consequences. It isn't very useful to know exactly what happens then in this situation.

**Anomaly:** In fact, things are a bit more complicated. When a signal or error is propagated, an `enChoice` statement is called as a `proc` from the `SIGNAL` or `ERROR` which raises the exception. When control leaves the statement by a `GOTO` (including `EXIT`, `CONTINUE`, `RETRY` or `LOOP`, but not `RETURN`, which is forbidden in an `enChoice`), the finalization is done. This means that the statement is executed before any finalization. This is useful for signals, which often raise exceptions. In some cases, however, notably if finalization would release monitor locks, this can cause trouble. A programmer can avoid this problem by exiting from the `enChoice` immediately with a `GOTO`.

**Caution:** An `enChoice` can raise a second exception `ex2` and fail to catch it. This will produce a state of confusion, and should be avoided. If it happens, `ex2` is propagated just like the first exception `ex1`; all the `enChoices` which saw `ex1` will see `ex2`. This is because the `enChoice` statement was called as a `proc`. Unless `ex2` is a signal which is resumed, the `enChoice` which caught `ex1` will be finalized and abandoned.

**Caution:** `ANY` unfortunately catches `UNWIND`, and hence its statement will be taken as the finalization. It is better not to use `ANY`. Also, it is possible to raise `UNWIND` explicitly.

## Signals

Conceptually, a signal is quite different from an error; in fact, it is very much like a function call. The only differences are:

The `proc` to be called is an `enChoice` which is found exactly as though the signal were an error. The effect of this is that `SIGNAL P[args]` binds the `proc` name `P` to the `proc` body `P` dynamically, by searching up the call stack for a binding of `P`. This is just the way `bind` binds free variables, except that a binding for `P` can only be found in an `enChoice` or the frame of a `proc`.

Actually this is not quite right. Like an error handler, the signal `proc` is not found by matching names but by matching exception values. This point is discussed in detail above.

The `enChoice` can be terminated by a `GOTO` out of its body, unlike an ordinary `proc`. A `GOTO` exception is treated exactly like a `GOTO` out of an `enChoice` for an error; it causes the intervening frames to be finalized.

The implementation, however, treats errors and signals in a very similar way; the only difference is that you cannot resume an error (return from the `enChoice`). In fact, you can invoke a `RESUME ERROR`, which prevents it from being resumed; avoid this feature. In the future, however, the distinction between signals and errors will be reflected more clearly in the implementation.

**Anomaly:** The desugaring gives no explanation of how `RESUME` works, since it does not turn an `enChoice` for a signal into a `proc` at all. This is a defect.

## 3.4.3.2 EXITS

An EXITS construct (confusingly called REPEAT in a loop) declares one or more exceptions local to its block, and also catches them. The syntax is just like an enable. However, labels appear before the => rather than expressions, and the EXITS introduces these names scope which includes the block body and any enable, but not an open and not the statement EXITS itself. A label may only be used in a GOTO statement.

Anomaly: Actually labels have their own name space, disjoint from the other names known block. Hence it is possible to declare a label n and still to refer to another n in the feature.

Like the raising of any exception, a GOTO n stops execution of the current statement. The associated with n is executed. If it finishes normally, execution continues after the block was declared. If it raises an exception, that exception becomes the value of the block.

Anomaly: A GOTO skips any UNWIND enChoices that intervene between the GOTO and its matching EXITS. This is the only way to escape from a block without executing the UNWIND. You can avoid this anomaly by not nesting UNWIND enChoices within blocks that have EXITS.

## 3.4.4 Safety

A SAFE proc has the property that if the safety invariants hold before it is called, they afterwards. Roughly, these invariants ensure that the value of every expression has the type of the expression, and that addresses refer only to storage of the proper type (¶ 4.5.1). A proc may lack this property. Hence a safe proc type implies the corresponding unsafe one.

We want to have confidence that the safety invariants hold. To this end, we want to have as few unsafe procs as possible;

a mechanical guarantee that a proc is safe, if possible.

Clearly, a proc whose body calls only safe procs will be safe.

Applying this observation, Cedar provides three attributes which can be applied to a block:

CHECKED: the compiler allows only safe procs to be applied; hence the block is automatically safe, and any proc with the block as its body is safe.

UNCHECKED: there are no restrictions on the block, and it is unsafe.

TRUSTED: there are no restrictions on the block, but the programmer guarantees that it preserves the safety invariants; the compiler assumes that the block is safe. This is a restricted form of LOOPHOLE.

These attributes are defaulted as follows.

A block is checked if its enclosing block is checked; otherwise it is unchecked.

If CEDAR appears in the module header, the outermost block is checked, and a transfer type constructor anywhere in the module defaults the SAFE option to TRUE. Hence the resulting type will be safe, and its initialization must be safe or there is a type error.

Otherwise, the outermost block is unchecked, and a transfer type constructor anywhere in the module defaults the SAFE option to FALSE. Hence the resulting type will be unsafe, and there is no safety restriction on its initialization.

Of course you can override these defaults by writing CHECKED, UNCHECKED or TRUSTED on any block, and SAFE or UNSAFE on any transferTC (except ERROR, which is automatically safe). The defaults are provided to make it convenient to:

write new programs in the safe language;  
 continue to use old, unsafe programs without massive editing.

An unsafe proc value never has a safe type, and hence cannot be bound to a name declared safe type. This applies to enable choices for signals as well as to procs. In both cases be checked or trusted if the type is safe. ERRORS are treated differently, however, because view that an ERROR is a value returned from an application, unlike a signal which calls the enChoice expression. Hence the enChoice for an ERROR is treated just like any statement in enclosing block, and is not considered to be bound to a proc when the ERROR is raised.

The following primitive procs are unsafe:

@, DESCRIPTOR and BASE.

^ or FREE applied to a pointer, and all pointer arithmetic.

APPLY of

a descriptor (because it involves dereferencing a pointer);  
 a computed sequence;  
 a record containing a computed sequence;  
 a base pointer.

APPLY for process and port types (JOIN and port calls).

withSelect<sup>34</sup>.

The fields of an OVERLAID union.

ASSIGN of:

An unspecified type to anything other than the same unspecified type (¶ 4.9)  
 A union or variant record.

LOOPHOLE which produces a RC value (¶ 4.5.1).

### 3.5 Declaration and binding

11 **declaration** ::= n, !.. : ?access<sup>12</sup> varTC<sup>40</sup> varTC ), ...

In 2, 10, 43. VAR, READONLY only for interface var.

12 **access** ::= PUBLIC | PRIVATE

In 2, 3, 11, 13, 50, 51, 53.

13 **binding** ::= n, !.. : ?access<sup>12</sup> t ~ ( n, ... ~LET x( : t ~ (

e	e
t <sub>2</sub> -- if t=TYPE	-2t -- Same as e except for conflicting syntax.
CODE	NEWEXCEPTIONCODE[] --tgsIGNAL or ERROR
?INLINE ?(ENTRY   INTERNAL) block <sup>6</sup>	l [d(: t.DOMAIN) IN LET r(~NEWFRAME[t.RANGE].UNCONS
	INET(r( IN {t.DOMAIN~d(; block; RETURN}
	BUT {Return( ( =>r())
,, ?TRUSTED MACHINE CODE {(e, ...); ...}	MACHINECODE[ (BYTESTOINSTRUCTION[e, ...]), ...]
)	) IN x( -- e is evaluated only once.

In 2, 10. •The ~ may be written as =.

Block or MACHINE CODE only for proc types.

•ENTRY and INTERNAL can also be before t.

## Examples

```

HistValue: TYPE[ANY];           -- Interface: An exported type.
Histogram: TYPE~REF HistValue; -- A type binding.
baseHist: READONLY Histogram;  -- An exported variable .
AddHists: PROC[x, y: Histogram] -- An exported proc.
  RETURNS [Histogram];
LabelValue: PRIVATE TYPE~RECORD[ -- PRIVATE only for secret
  first,last:INT,s:ROPE,x:REAL,f,g:INT,r:REF ANY];  stuff in an interface.
Label: TYPE~REF LabelValue;
Next: PROC[l: Label] RETURNS[Label]~ -- An inline proc binding.
  INLINE { RETURN [NARROW[l.r]] };

H: TYPE~Histogram11; Size: INT~10; -- Implementation binds a TYPE and INT.
HistValue: PUBLIC TYPE~HV40.1; -- PUBLIC for exports.
baseHist: PUBLIC H_NEW[HistValue_ALL[17]]; -- An exported variable
x, y: HistValue_[ 20, 18, 16, 14, 12,-10, 8, 6, 4, 2,with initialization.
FatalError: ERROR[reason: ROPE]~CODE; -- Binds an error.
Setup: PROC [h: Handle3, a: INT]~ENTRY {-..}; -- Binds an entry proc.
i,j,k: INT_0; p,q: BOOL; lb: Label; main: Handle;

```

Declarations are explained in ¶ 2.2.1F and ¶ 2.4.5. Their peculiarities in the different they can appear are explained elsewhere:

```

  interfaces in ¶ 3.3.4;
  blocks in ¶ 3.4.1;
  fields in:
    domains and ranges in ¶ 4.4;
    records and unions in ¶ 4.6;

```

Access is explained in ¶ 3.3.6.

Bindings are explained in ¶ 2.3.5. There are several special forms of binding given in r however, which are defined here. See also ¶ 3.7 on argument bindings. Note that the e in is evaluated just once, even if several names are bound.

A TYPE binding is the only way in which a type value can be bound to a name, since cannot be passed as parameters. Unlike other bindings, this one expects a type<sup>36</sup> ra an expression<sup>19</sup> after the ~.

A name with a signal or error type can be bound to CODE; this use of CODE is not all anywhere else. See ¶ 4.4.1 for details on the meaning of this.

„A MACHINE CODE construct can be bound to a name with a proc type. This construct allows machine instructions to be assembled into a proc value. The instructions ar separated by semicolons. Each instruction is assembled from a list of expressions by commas. An expression in the list is evaluated to yield a [0..256) static value forms one byte of the instruction; successive expressions form successive bytes.

A l-expression derived from a block can be bound to a name with a proc type. The complicated semantics of this construction are explained in the following subsecti

## 3.5.1 PROC bindings

A binding of the form n: T~{...} is the only way to construct a proc value and bind it to since you cannot write a l-expression in current Cedar.

There are other ways to construct proc values:

The expression in a defaultTC<sup>55</sup> is turned into a parameterless proc which is bound to Default in the type cluster (¶ 4.11).

The expression following ~ in an open or WITH ... SELECT is turned into a parameterless proc with a deprecating coercion (¶ 3.4.2).

The statement in an enable choice for an exception is turned into a proc with domain and range given by exception type (¶ 3.4.3.1).

The expression following LOCKS in a module heading is turned into a proc according to a peculiar rule (¶

The l-expression is constructed from the block in the following way. Its domain and range domain and range of the proc type T. Its body implicitly declares a variable for each item in the domain and range; these variables have the names of the domain and range items, and their type is the entire block, not just the block body. The domain variables are initialized to the value of the domain item and the range variables in the usual way according to their types. Then the block, with the domain and range variables tacked on the end, is evaluated. A RETURN exception in the block is caught, and the current value of the range variables are the result of the l-expression. The only other way out of the block is to raise an ERROR.

A RETURN in the block is sugar for GOTO Return(, which is caught as described. RETURN e assigns the value e to the range variables and then does a GOTO Return(.

Anomaly: It is an error to introduce the same name twice in the domain, range or block

Performance: A proc call and return is about 30% faster if the proc is local, i.e., denoted by a name which was bound to a proc body in the same module as the call. A proc which is local to the module, rather than bound in the body of an implementation, is about 20% slower to call. Inline procs introduce some overhead when its parent proc is called, and its access to non-static names is slower than access to other names. A call and return for an inline proc introduced in its parent proc is slower than access to other names. A call and return for an ordinary, non-local proc takes about 10 times as long as the statement x\_y+z.

The attributes ENTRY and INTERNAL can be used only in a MONITOR; they are discussed in ¶ 4.1

The attribute INLINE has no effect on the meaning of the program, but it causes the proc body to be expanded inline whenever it is applied. This saves the cost of a proc call and return, and sometimes the cost of argument passing, and it may allow constant arguments to participate in the evaluation within the proc.

Restrictions: An INLINE proc may not be:

- Recursive.

- Exported.

- Used as a proc value except in an application. Thus you cannot, for example, assign a proc to a proc variable.

- The argument of FORK.

- Accessed from the cluster of a POINTER TO FRAME type.

Anomaly: An inline proc binding in an interface is not accessible from a DIRECTORY argument; the interface must import the interface.

Performance: Excessive application of inline procs will result in much larger compiled code. Excessive definition of inline procs will result in much larger data structures in the compiler, hence in larger symbol table files, and a greater chance of overflowing the compiler's code. The following cases are efficient:

- An inline proc in an implementation which is called zero or one times.

- An inline proc which has a simple body, no locals, no named results, and no access to global formals after potential side effects.



## 3.6 Statements

```

14 statement ::= sS                                { SIMPLELOOP {sS; GOTO Cont((; EXITS Retry(=>NULL};
    In 6, 10, 17, 19.                               EXITS Cont(=>NULL }
15 sS ::= e1-e2 | e | block6 | escape | loop11 | VOID | e --must yield VOID-- | --all four yield VOID
16 escape ::= GOTO n | GO TO n |                   HEX[exception[code~ n((, args~NIL]] |
    EXIT | CONTINUE | •LOOP | •RETRY |             GOTO ( Exit(17 | Cont(9 | Loop(17 | Retry(9) |
    (RETURN | RESUME) ?e |                          { ?(r(13_e;) GOTO (Return(13 | Resume(13) } |
    •REJECT | ,,e _ STATE                           THISEXCEPTION[] | DUMPSTATE[e]

17 loop ::= (iterator | )                          { ( iterator done(~FALSE; Next(: PROC~{;} ; )
    (WHILE e | UNTIL e | )                          { Test(~1 IN (NOT e | e | FALSE);
    DO ?•open7 ?•enable8 ?body10                { open SIMPLELOOP {
                                                    IF Test([] OR done( THEN GOTO FINISHED;
                                                    { enable body EXITS Loop(=>NULL }; Next
    ?(REPEAT (n, !..=>s); ...) ENDLLOOP            EXITS Exit(gNULL; (n, !..gs); ...; FINISHEDgNULL)}}}

18 iterator ::= THROUGH e |                        FOR x(: e IN|e
    FOR (n : t | μn) ( n: t; | )                     ( ( | DECR Range|: IN yEE |e; done(: BOOL_Range(.ISEMPTY;
    ( ( | DECR Range|: IN yEE |e; done(: BOOL_Range(.ISEMPTY;
    Next(: PROC~{ IF n ( >Range(.LAST | <Range(.FIRST )
    THEN done(_TRUE ELSE n_n.(SUCC | PRED) );
    n_range(. (FIRST | LAST); |
    _ e1 done2(0): BOOL~FALSE; Next(: PROC~{n_e2}; n_e1 ) ;
    e is a subrange. In FOR n: t ... , n is readonly except for the assignment in the iterator's desugaring.

```

## Examples

```

x_AddHists[baseHist, baseHist]^;    -- A statement can be an assignment,
Setup[bh~main, a~3];                -- or an application without results,
{ENABLE FatalError=>RETURN[0]; []_f[3];--.or} a block,
IF i>3 THEN RETURN[25] ELSE GOTO NotPresent; or an IF or an escape statement,

```

```

FOR t:INT DECREASING IN [0..5) UNTIL f[t]>3 -- or a loop. Try to declare t in the FOR
u: INT_0; ... ; u_t+4; ...            -- as shown. Avoid OPEN or ENABLE
REPEAT Out=>{...}; FINISHED=>{...} ENDLLOOP; after DO (use a block). FINISHED
-- must be last.

```

```

THROUGH [1..5) DO i_i*i ENDLLOOP;    -- Raises i to the 16th power.
FOR i: INT_1, i+2 WHILE i<8 DO j_j+i ...- Accumulates odd numbers in [1..8).
FOR l: Label_lb, l.Next WHILE l#NIL DO --. Sequences through a list of Labels.

```

Cedar makes a distinction between expressions and statements. This distinction is most easily defined in terms of a special type called VOID, which is equivalent to the empty declaration. The range type of a PROC [...]<sub>1</sub>[], and it is also the result type of a block, control statement. An expression whose value is a VOID can be used as a statement, and cannot be an ordinary value in a binding (since it wouldn't have the right type). If you want to control which returns values as a statement, you must assign the results to an empty group:

```

[]_f [...]
```

Assignment is a special case; an assignment can be used as a statement even though its value is the value of the right operand. This is explained in the desugaring<sup>15</sup> using a special proc TO cluster of every assignable type; it takes a value of the type and returns a VOID. Note that the grammar is ambiguous here, since there are two parsings of e<sub>1</sub>-e<sub>2</sub> as a statement; the one using the rule for statement is preferred.

Anomaly: In a `select`<sup>29</sup> which is a statement (i.e., returns `VOID`), the choices are separated by semicolons; in an ordinary `select` expression they are separated by commas.

Anomaly: •If you write an expression whose value is a `proc` taking no arguments as a statement, the `proc` gets applied. Thus

```
P;
```

is the same as

```
P[];
```

This is the only situation in which an ordinary `proc` gets applied by coercion (but see ¶ open procs).

A `statement`<sup>14</sup> is actually a rather complicated construct, as the desugaring shows. This is the case for the `CONTINUE` and `RETRY` statements, which respectively terminate and repeat the statement containing the `enable`<sup>9</sup> in which they appear. The desugaring shows exactly what this means in various obscure cases. `CONTINUE` and `RETRY` are legal only in an `enable` choice (¶ 3.4.2), and may not appear in a declaration at all. •`RETRY` should be avoided everywhere, since it interrupts the loop into the program in a distinctly non-obvious way.

`Escape`<sup>16</sup> consists mainly of the various flavors of `GOTO` (including `EXIT`, `CONTINUE`, `LOOP`, `RETRY`, `RETURN` and `RESUME`) which raise a local exception bound in an `EXITS`; this is explained in ¶ 3.4.3.2. `REJECT` is explained in ¶ 3.4.3.1.

Anomaly: You cannot use a `GOTO` to escape from a `proc` body, even though the body is within the scope of the label. Only normal completion, or a `RETURN` or `ERROR` exception (or a `SIGNAL` which is not resumed) can terminate the execution of a `proc` body.

A `loop`<sup>17</sup> is repeated indefinitely until stopped by an exception, or by the `iterator`<sup>18</sup> or `UNTIL` test. It has a body, bracketted by `DO` and `ENDLOOP`, which is almost like a block, but with some confusing differences:

You catch `GOTO` exceptions with `REPEAT`, which is exactly like `EXITS` in a block immediately around the loop, except for the different delimiting reserved word. Note that the labels does not include the `iterator` or the `test`, even though these are evaluated repeatedly during execution of the loop. This feature is best avoided, but unfortunately necessary if you want to catch the `FINISHED` exception explained below.

•You can write an `open` or `enable`. This is also best avoided, since the scope is confusing. It is better to write a block explicitly inside the `DO` if you need these facilities.

There are three special exceptions associated with loops:

`EXIT` is equivalent to `GOTO Exit(`, where `Exit(` is a label automatically declared in the body of every loop. Its `enable` choice does nothing. Thus `EXIT` simply terminates the smallest block that encloses it.

`FINISHED` is raised when the `iterator` or the `WHILE/UNTIL` test terminates the loop. It is declared in the `REPEAT` like any label, but it must come last. If it is not declared, an `enable` choice is supplied for it.

•`LOOP` causes the next repetition of the loop to start immediately.

Anomaly: You cannot write `GOTO FINISHED`.

An `iterator`<sup>18</sup> declares a control variable `v` which is initialized by the `iterator` and updated at the end of every iteration; the scope of `v` is the entire loop, and it is read-only in the loop. If the loop is terminated by the `iterator` (i.e., in the `FINISHED` clause), the value of `v` is undefined. If you omit the declaration and simply name an already declared variable, it will be used as the control variable, and will not be read-only; it will still be undefined after the loop is terminated. Avoid this feature.

There are three flavors of iterator:

THROUGH, which has no explicit control variable; THROUGH [0..k) is convenient when you just want to loop k times.

FOR v: T IN [first, last] ...; v is initialized to first, and set to succ[v] after each iteration finishes the loop after a repetition which leaves v>last. The > case can occur for v IN ..., when an out-of-range value is assigned to v in the loop body. DECREASING reverses the order in which the elements of the subrange are used. The subrange need not be static. Note that the subrange is evaluated only once, before execution of the loop begins.

FOR v: T\_first, next ...; v is initialized to first, and set to next after each repetition; the iterator never finishes the loop. Note that the expression next is reevaluated each time around the loop. The usual application is something like

```
FOR v: List_header, v.next UNTIL v=NIL.
```

Note that the WHILE or UNTIL test is made with v equal to its value during the next repetition; that both tests are made before the first repetition, so that zero repetitions are possible.

### 3.7 Expressions

19 expression ::= n | literal<sup>26</sup> | application<sup>26</sup> |  
 (e | typeName<sup>37</sup>) . (9) n |  
 prefixOp e | e<sub>1</sub> infixOp e<sub>2</sub> | e . prefixOp e<sub>1</sub> . infixOp e<sub>2</sub> |  
 e<sub>1</sub> relOp (4) e<sub>2</sub> | ( l [x(: De<sub>1</sub>, y( :De<sub>2</sub>)] ) in relOp |  
 e<sub>1</sub> AND (2) e<sub>2</sub> | e<sub>1</sub> OR (1) e<sub>2</sub> | IF e<sub>1</sub> THEN e<sub>2</sub> ELSE FALSE | IF e<sub>1</sub> THEN TRUE ELSE e<sub>2</sub> |  
 e ^ (9) | •STOP | ERROR | e . DEREFERENCE | STOP[] | ERROR NAMELESSERROR |  
 builtIn [ e<sub>1</sub> ?( , e<sub>2</sub>, !.. ) ?applEn<sup>27</sup> ] . builtIn ?( [e<sub>2</sub>, ... ?applEn ] ) |  
 funnyAppl e ?( [?argBinding<sup>27</sup> ?applEn<sup>27</sup>] funnyAppl ?( [argBinding ] ) |  
 [ argBinding<sup>27</sup> ] | --Binding must coerce to a record, array, or •local  
 s | subrange<sup>25</sup> | if<sup>28</sup> | select<sup>29</sup> | safeSelect<sup>32</sup> | •withSelect<sup>34</sup>  
 Precedence is in bold in rules 19-21. All operators associate to the left except \_, which associates to the right. Application has highest precedence. Subrange only after IN or THROUGH. s only in if<sup>28</sup> and sel

20 prefixOp ::= @ (8) | (7) | (~ | NOT | NOTPOINTER | UMINUS | NOT  
 21 infixOp ::= \* | / | MOD (6) | + | TIMES | DIVIDE | REM | PLUS | MINUS | ASSIGN  
 22 relOp ::= ?NOT ( ?~ (= | < | > ) | #NOT ( ?NOT x(.EQUAL | LESS | GREATER)[y() | x(~=y( |  
 x(=y( &R>=0 (<IN)>) y( | x(>=y(x(=y(AND&R>=0 ) BUT {Bound

--In 19, 30.  
 23 builtIn ::= -- These are enumerated in Table 4 5.  
 24 funnyAppl ::= FORK | JOIN | WAIT | NOTIFY | BROADCAST |  
 SIGNAL | ERROR | RETURN WITH ERROR |  
 •NEW | •START | •RESTART | „,TRANSFER WITH | „,RETURN WITH

25 subrange ::= (typeName<sup>37</sup> | ) LET t(~(typeName | INT) , first(~( e<sub>1</sub> | e<sub>1</sub>.SUCC ) IN  
 - ) e<sub>1</sub> .. e<sub>2</sub>)( [ ] || ( ) t(.MKSUBRANGE[first(, (e<sub>2</sub> | e<sub>2</sub>.PRED )] BUT  
 --In 19, 39, 48. {BoundsFault=>t(.MKEMPTYSUBRANGE[e<sub>1</sub>])

26 application ::= e [?argBinding ?applEn] (~e, a(~[argBinding](. APPLY Za(?applEn  
 27 argBinding ::= (n ~ (e | | μTRASH ))(n!~.(e | OMITTED | TRASH) ), !.. |  
 (e | | μTRASH ), OMITTED | TRASH ), ...  
 In 19, 26. •TRASH may be written as NULL, ~ as :.  
 27.1applEn ::= ! enChoice<sup>9</sup>; ...-- In 19, B06. {enChoice... }

## Examples

```
lv: LabelValue13[ i, 3, "Hello", 31.4E A, const] constructor with some sample
  g[x]+lb.f+j.PRED, NIL ]; -- expressions.
p1: PROCESS RETURNS [INT]_FORK f[i, j]; -- FunnyAppls take one unbracketted
ERROR NoSpace; WAIT bufferFilled; -- arg; many return no result, so
RT: RTBasic.Type_CODE[LabelValue13]; -- must be statements.
h[ 3, NOT(i>j), i*j, i_3, i NOT >j, p OR An application with sample expressions.
lv19[first~0,last~5,x~3.2,g~2,f~5,r~NIL,Short]for lv_LabelValue13[...].
[first~i, last~j]_lv19; -- Assignment to VAR binding
-- (extractor).
```

```
b: BOOL_i IN [1..10]; FOR x: INT IN (0..1)Subrange only in types or with IN.
b_( c IN Color54(red..green) OR x IN INT{0The0}is redundant.
```

```
fh_Files.Open[name~lb.s, mode~Files.readKeywords are best for multiple args.
! AccessDenied=>{...}; FatalError=>{-.Semicolons separate choices.
(GetProcs[j].ReadProc)[k]; -- The proc can be computed.
file.Read[buffer~b, count~k]; -- WFile.Read[file, b, k] (object notation).
f[i~3, j~ , k~TRASH]; f[i~3, k~TRASH]; -- j and k may be trash (see defaultTC55).
f[3, , TRASH]; -- Likewise, if i, j, and k are in that order.
```

Most of the forms of expression are straightforward sugar for application: prefix, infix operators, explicit application of a primitive proc<sup>23</sup>, or the funnyApp<sup>24</sup> in which the file follows the proc name without any brackets. All of these constructs desugar into dot notation (¶ 2.4.4, ¶ 4.14); this means that the procs come from the cluster of the first argument. Exceptions to this rule are ALL, CONS for variant records and lists, LIST, and the single-argument forms of LOOPHOLE and NARROW, and VAL; all of these get the proc from the target type of the expression (¶ 4.2.3). All the primitive procs are described in ¶ 4.

Note that AND and OR are not simply sugar for application. Rather, they are sugar for an application, since the second operand is evaluated only if the first one is TRUE or FALSE or

The order of evaluation for arguments of an application, and therefore for operands in an expression is not defined (unless the operator is AND or OR). However, the arguments are evaluated one at a time, and all arguments are evaluated before the proc is applied. In particular assignment which executes completely behaves as though both left and right operands are completely evaluated before any assignments are done, even if the left side is a binding [a~x, b~y.f].

Rules 19-21 give the precedence for operators: ^ and . are highest (bind most tightly) and \_ are lowest. All are left-associative except \_, which is right-associative. Application has the lowest precedence.

Style: The precedence rules are sufficiently complex that it is wise to parenthesize expressions that depend on subtle differences in precedence.

The first operand of assign can be an argBinding<sup>27</sup> whose value is a variable group or binding whose elements are variables; this is sometimes called an extractor. The second argument is typechecked if it is a group or binding with corresponding elements which can be assigned to the variables. Usually the second argument is either an application which returns more than one value or a record-valued expression. You can omit elements of the left argBinding to discard the corresponding values; however, you can't write TRASH in the left operand. Note that the right operand is fully evaluated before any variables are changed by the assignment.

The expression ERROR is short for raising a nameless ERROR exception. You should think of it as a call to the debugger, appropriate for a state which "can't occur".

A funnyAppl which takes more than one argument has the extra arguments written inside brackets in the usual way; e.g., `START P [3, "Help"]. RETURN WITH ERROR` is explained in ¶ 4.10.

Anomaly: The funnyAppl `NEW e` actually stands for `e.COPYIMPLINST`. See ¶ 4.4.1 and ¶ 4.5.3.

Anomaly: Enable choices are legal only for the following funnyAppls: `FORK JOIN RESTART START STOP WAIT`. You can write empty brackets if necessary to get a place for the enable choices.

A subrange<sup>25</sup> denotes a subrange type; see ¶ 4.7.3. Standard mathematical notation for open/closed intervals is used to indicate whether the endpoints are included in the subrange. `..` can also be used after `IN` in an expression or iterator; in these contexts it need not be

You can write enable choices<sup>9</sup> after a `!` inside the brackets of an application<sup>26</sup>, built-in funnyAppl<sup>24</sup>. See ¶ 3.3.2 for the semantics of this. Note that only an exception returned from an application is caught by these choices, not one resulting from evaluating the proc or arg

An argBinding<sup>27</sup> denotes a binding for the arguments of an application. You can omit a `[name, value]` pair `n~e` in the binding if the corresponding type has a default, or you can omit the name without the value expression (e.g., `n~`) with the same meaning. You can also write `!~` (`!~` or `NULL`) for the value; this supplies a trash value for the argument (¶ 4.11).

### 3.8 IF and SELECT

```

28 if ::= IF e1 THEN e2 (ELSE e3 | )      IF e1 THEN e2 ELSE (e3 | NULL)
29 select ::= SELECT e FROM                    LET selector(~e IN
        choice; ... endChoice                choice ... endChoice
The ";" is " ," in an expression; also in 32- and 34. is a separator for repetitions of the choice.
30 choice ::= ( ( | relOp22 ) e1 ), !..IF>e2(selector( (= | relOp ) e1) OR ... ) THEN e2
31 endChoice ::= ENDCASE ( | => e3)         ELSE (NULL | e3)
In 29, 32, 34.

32 safeSelect ::= WITH e SELECT FROM          LET v(~e IN
        safeChoice; ... endChoice31        safeChoice ... endChoice
33 safeChoice ::= n : t => e2                IF ISTYPE[v(, t)] THEN LET n : t_NARROW[v(, t)] IN e2
34 •withSelect ::= WITH (n1 ~~ e1 | • e1 OPEN v(~e1 IN LET n(~($n1 | NIL), type(~Dv(,
        SELECT ( | ,e11) FROM                selector(~(e1.TAG | e11) IN withChoice endChoice
        withChoice; ... endChoice31        -- e11 must be defaulted except for a COMPUTED variant.
•The ~~ may be written as :.
35 •withChoice ::= n2 => e2 |              IF selector(=$n2 THEN OPEN
        n2, n2, !.. => e2                (BINDP[n(, LOOPHOLE[v(, type(.n2] ] | BINDP[n(, v(
```

## Examples

```

i_(IF j<3 THEN 6 ELSE 8);           -- An IF with results must have an ELSE.
IF k NOT IN Range THEN RETURN[7];
SELECT f[j] FROM                   -- SELECT expressions are also possible.
  <7=>{...};                         -- Wt:INT~f [j]; IF t<7 THEN {...} ELSE ...
  IN [7..8]=>{...};                 -- 7, 8=> or =7, =8=>{...} is the same.
  NOT <=8=>{...};                   -- ENDCASE=>{...} is the same here.
  ENDCASE=>ERROR;                   -- Redundant: choices are exhaustive.

WITH r SELECT FROM                 -- Assume r: REF ANY in this example.
  rInt: REF INT=>RETURN[Gcd[rInt^, 17]];- rInt is declared in this choice only.
  rReal: REF REAL=>RETURN[Floor[Sin[rReal^]]];
  ENDCASE=>RETURN[IF r=NIL THEN 0 ELSE 1] -- Only the REF ANY r is known here.

nr: REF Node52~...; WITH dn~~nr SELECT FROM See rule 52 for the variant record Node.
  binary=>{nr_dn.b};                -- dn is a Node.binary in this choice only.
  unary=>{nr_dn.a};                 -- dn is a Node.unary in this choice only.
  ENDCASE=>{nr_NIL};               -- dn is just a Node here.

```

The kernel construct `if28` evaluates the expression  $e_1$  to a `BOOL` value `test`, and then evaluates `test=TRUE`, or  $e_3$  if `test=FALSE`. In the expression

```
IF test1 THEN IF test2 THEN ifTrue2 ELSE ifFalse2
```

the grammar is ambiguous about which `IF` the `ELSE` belongs to. It belongs to the second one.

A `select29` is a sugared form of `if` which is convenient when one of several cases is chosen to yield a single value. The selector expression  $e$  is evaluated once to yield a value `selector(, a`, and then each choice is tested in turn. Within each choice, each expression  $e_1$  preceding the `=>` is evaluated in turn with `selector(`; the comparison is `selector( relop e1 if e1 is preceded by a relop; otherwise, selector(=e1. If any comparison succeeds, the expression  $e_2$  following the => is evaluated to yield the value of the select. If no comparison succeeds, the next choice is tried. If no choice succeeds, the expression  $e_3$  following the ENDCASE is evaluated to yield the value of the select;  $e_3$  defaults to NULL, and hence must be present when the select is not a statement to prevent a type error.`

**Style:** It is good practice to arrange the tests so that they are disjoint and exhaust the possible values of the selector. `ENDCASE` should be used to mean "in all other cases"; often the application of `ENDCASE` to  $e_2$  raises an error. Don't use `ENDCASE` to mean another specific selector value which you do not wish to bother to mention. Another acceptable form is `SELECT TRUE FROM ...`, which selects the first choice that succeeds, and is sometimes easier to read than a long sequence of `ELSE IF`'s.

**Performance:** If the  $e_2$  are static and select subsets of the selector values, the average number of choices tested is not too large, and the density of unselected values is not too high, a `select` can be implemented as an indexed jump, which executes in a time independent of the number of choices.

A `safeSelect32` is a special form for discriminating cases of unions or `ANY`. The selector expression  $e$  is evaluated to yield a value for which `ISTYPE` can be evaluated dynamically (¶ 4.3.1): `REF ANY`, `PROC ANY_T`, `PROC ANY_T ANY`, `V`, `REF V`, or `(LONG) POINTER TO V`, where  $V$  is a variant record. Each choice specifies a possible type that the selector might have, and declares a name which is initialized to the value of the selector if it has that type. Thus, the example tests for  $r$  having the types `REF INT` and `REF REAL`. If  $r$  has `REF INT`, the first choice's  $e$  is evaluated; within  $e$ , `rInt` is a variable initialized to the value of `r` and has type `REF INT`. Likewise for `REF REAL` and the second choice. As with an ordinary `select`, the `ENDCASE` expression is evaluated (with no new names known) if none of the other choices succeeds. Note that `safeSelect` does ordinary binding by value, not the binding by name done in `open` and `withSelect`.

•,A withSelect<sup>34</sup> is an unsafe and rather tricky construction for discriminating cases of use should be avoided unless a safeSelect can't do the job; this is the case for a COMPUTED if the call by name feature of withSelect is required.

It incorporates an open (¶ 3.4.2) of the  $e_1$  being discriminated. This means that  $e_1$  is dereferenced to yield a variant record value. It also means that this value is not constant, hence it can change its type during execution of a choice, either by assignment to a variant part of a variant record (itself an unsafe operation), or by a change in the value of  $e_1$ .

If the union has a COMPUTED tag, the selector value to be used for the discrimination must be given as  $e_{11}$  in the withSelect. It is entirely up to the programmer to supply a value. If the tag is not COMPUTED,  $e_{11}$  must be omitted and the selector value is  $e_1$ .

The  $n_2$  preceding => in a choice are literals of the (enumerated) type (¶ 4.7.1.1) with the tag type of the union (¶ 4.6.3). They are compared with the selector, and if one matches, the  $e_2$  following => is evaluated as with an ordinary select. If exactly one is given, the  $e_2$  following => is in the scope of

```
OPEN n1~~LOOPHOLE[e1.DEREF, V.n2];
```

or simply

```
OPEN LOOPHOLE[e1.DEREF, V.n2]
```

if no  $n_1$ ~~ followed the WITH. If several  $n_2$  are given, then there is no discrimination; the  $e_2$  following => is in the scope of

```
OPEN n1~~e1.DEREF      or      OPEN e1.DEREF
```

### 3.9 Miscellaneous

This section deals with various topics that are not naturally associated with particular grammar rules.

#### 3.9.1 Static values

An expression has a static value if the compiler can compute the value. In Cedar, an expression has a static value (is static for short) if it is:

a literal;

a name bound to a static value;

an application to static arguments of

a proc declared INLINE with a static body, or

a primitive which is not a loop, a REAL primitive (except unary minus, ABS or INTTOREAL), ASSIGN, @ or NEW. Note that IF and SELECT are evaluated.

Performance: The compiler evaluates all static expressions, not just type expressions. This is important for efficiency.

#### 3.9.2 Size restrictions

Current Cedar has the following restrictions on the sizes of values:

- A record type T must have  $T.SIZE < 2^{16}$ .
- A row type T must have  $T.SIZE < 2^{28}$  and  $T.RANGE.SIZE < 2^{16}$ .

- A type T with  $T.SIZE > 2^{16}$  lacks the following procs:
  - ALL
  - ASSIGN
  - CONS
  - DESCRIPTOR
  - INIT
  - NEW
- A subrange type T must have
  - $0 < T.LAST \leq T.FIRST < 2^{16}$
  - $2^{15} < T.FIRST < 2^{16}$
  - $T.LAST < (IF T.FIRST THEN 2^{15} + T.FIRST ELSE 2^{16})$

### 2.9.3 Checking

Possible errors arising from certain primitive operations are checked, and cause `ERROR` exceptions if they occur, in a `CHECKED` block, or if the compiler's "u" switch is on:

Dereferencing `NIL`.

Narrowing an out-of-range value to a subrange type.

Assigning a local proc to a proc variable (in `CHECKED` blocks only).