

CHAPTER 7

VARIABLE BINDINGS AND THE INTERLISP STACK

A number of schemes have been used in different implementations of LISP for storing the values of variables. These include:

- (1) Storing values on an association list paired with the variable names.
- (2) Storing values on the property list of the atom which is the name of the variable.
- (3) Storing values in a special value cell associated with the atom name, putting old values on a pushdown list, and restoring these values when exiting from a function.
- (4) Storing values on a pushdown list.

Interlisp-10 uses the third scheme, so called “shallow binding”. When a function is entered, the value of each variable bound by the function (function argument) is stored in a value cell associated with that variable name. The value that was in the value cell is stored in a block of storage called the basic frame for this function call. In addition, on exit from the function each variable must be individually unbound; that is, the old value saved in the basic frame must be restored to the value cell. Thus there is a higher cost for binding and unbinding a variable than in the fourth scheme, “deep binding”. However, to find the current value of any variable, it is only necessary to access the variable’s value cell, thus making variable reference considerably cheaper under shallow binding than under deep binding, especially for free variables. However, the shallow binding scheme used does require an additional overhead in switching contexts when doing “spaghetti stack” operations.

Interlisp-D uses the fourth scheme, “deep binding.” Every time a function is entered, a basic frame containing the new variables is put on top of the stack. Therefore, any variable reference requires searching the stack for the first instance of that variable, which makes free variable use somewhat more expensive than in a shallow binding scheme. On the other hand, spaghetti stack operations are considerably faster. Some other tricks involving copying freely-referenced variables to higher frames on the stack are also used to speed up the search.

The basic frames are allocated on a stack or pushdown list; for most user purposes, these frames should be thought of as containing the variable names associated with the function call, and the *current* values for that frame. The descriptions of the stack functions in below are presented from this viewpoint. Both interpreted and compiled functions store both the names and values of variables so that interpreted and compiled functions are compatible and can be freely intermixed, i.e., free variables can be used with no `SPECVAR` declarations necessary. However, it is possible to *suppress* storing of names in compiled functions, either for efficiency or to avoid a clash, via a `LOCALVAR` declaration (see page 12.4). The names are also very useful in debugging, for they make possible a complete symbolic backtrace in case of error.

In addition to the binding information, additional information is associated with each function call: access information indicating the path to search the basic frames for variable bindings, control information, and temporary results are also stored on the stack in a block called the frame extension. The interpreter also stores information about partially evaluated expressions as described on page 7.10.

The Spaghetti Stack

7.1 THE SPAGHETTI STACK

The Bobrow/Wegbreit paper, "A Model and Stack Implementation for Multiple Environments",¹ describes an access and control mechanism more general than the simple pushdown stack. The access and control mechanism used by Interlisp is a slightly modified version of the one proposed by Bobrow and Wegbreit. This mechanism is called the "spaghetti stack."

The spaghetti system presents the access and control stack as a data structure composed of "frames." The functions described below operate on this structure. These primitives allow user functions to manipulate the stack in a machine independent way. Backtracking, coroutines, and more sophisticated control schemes can be easily implemented with these primitives.

The evaluation of a function requires the allocation of storage to hold the values of its local variables during the computation. In addition to variable bindings, an activation of a function requires a return link (indicating where control is to go after the completion of the computation) and room for temporaries needed during the computation. In the spaghetti system, one "stack" is used for storing all this information, but it is best to view this stack as a tree of linked objects called frame extensions (or simply frames).

A frame extension is a variable sized block of storage containing a frame name, a pointer to some variable bindings (the BLINK), and two pointers to other frame extensions (the ALINK and CLINK). In addition to these components, a frame extension contains other information (such as temporaries and reference counts) that does not interest us here.

The block of storage holding the variable bindings is called a basic frame. A basic frame is essentially an array of pairs, each of which contains a variable name and its value. The reason frame extensions point to basic frames (rather than just having them "built in") is so that two frame extensions can share a common basic frame. This allows two processes to communicate via shared variable bindings.

The chain of frame extensions which can be reached via the successive ALINKs from a given frame is called the "access chain" of the frame. The first frame in the access chain is the starting frame. The chain through successive CLINKs is called the "control chain".

A frame extension completely specifies the variable bindings and control information necessary for the evaluation of a function. Whenever a function (or in fact, any form which generally binds local variables) is evaluated, it is associated with some frame extension.

In the beginning there is precisely one frame extension in existence. This is the frame in which the top-level call to the interpreter is being run. This frame is called the "top-level" frame.

Since precisely one function is being executed at any instant, exactly one frame is distinguished as having the "control bubble" in it. This frame is called the active frame. Initially, the top-level frame is the active frame. If the computation in the active frame invokes another function, a new basic frame and frame extension are built. The frame name of this basic frame will be the name of the function being called. The ALINK, BLINK, and CLINK of the new frame all depend on precisely how the function is invoked. The new function is then run in this new frame by passing control to that frame, i.e., it is made the active frame.

¹*Communications of the ACM*, Vol. 16, 10, October 1973.

VARIABLE BINDINGS AND THE INTERLISP STACK

Once the active computation has been completed, control normally returns to the frame pointed to by the CLINK of the active frame. That is, the frame in the CLINK becomes the active frame.

In most cases, the storage associated with the basic frame and frame extension just abandoned can be reclaimed. However, it is possible to obtain a pointer to a frame extension and to “hold on” to this frame even after it has been exited. This pointer can be used later to run another computation in that environment, or even “continue” the exited computation.

A separate data type, called a stack pointer, is used for this purpose. A stack pointer is just a cell that literally points to a frame extension. Stack pointers print as #ADR /FRAMENAME, e.g., #1,13636/COND. Stack pointers are returned by many of the stack manipulating functions described below. Except for certain abbreviations (such as “the frame with such-and-such a name”), stack pointers are the only way the user can reference a frame extension. As long as the user has a stack pointer which references a frame extension, that frame extension (and all those that can be reached from it) will not be garbage collected.

Note that two stack pointers referencing the same frame extension are *not* necessarily EQ, i.e., (EQ (STKPOS 'FOO) (STKPOS 'FOO)) = NIL. However, EQP can be used to test if two different stack pointers reference the same frame extension (page 2.3).

It is possible to evaluate a form with respect to an access chain other than the current one by using a stack pointer to refer to the head of the access chain desired. Note, however, that this can be very expensive when using a shallow binding scheme such as that in Interlisp-10. When evaluating the form, since all references to variables under the shallow binding scheme go through the variable's value cell, the values in the value cells must be adjusted to reflect the values appropriate to the desired access chain. This is done by changing all the bindings on the current access chain (all the name-value pairs) so that they contain the value current at the time of the call. Then along the new access path, all bindings are made to contain the previous value of the variable, and the current value is placed in the value cell. For that part of the access path which is shared by the old and new chain, no work has to be done. The context switching time, i.e. the overhead in switching from the current, active, access chain to another one, is directly proportional to the size of the two branches that are not shared between the access contexts. This cost should be remembered in using generators and coroutines (page 7.13).

7.2 STACK FUNCTIONS

In the descriptions of the stack functions below, when we refer to an argument as a stack descriptor, we mean that it is either a stack pointer or one of the following abbreviations:

NIL means the active frame; that is, the frame of the stack function itself.

T means the top-level frame.

Any other literal atom is equivalent to (STKPOS ATOM -1).

A number is equivalent to (STKNTH NUMBER).

In the stack functions described below, the following errors can occur: The error ILLEGAL STACK ARG occurs when a stack descriptor is expected and the supplied argument is either not a legal stack descriptor (i.e., not a stack pointer, litatom, or number), or is a litatom or number for which there is no corresponding stack frame, e.g., (STKNTH -1 'FOO) where there is no frame named FOO

Stack Functions

in the active control chain or (STKNTH -10 'EVALQT). The error STACK POINTER HAS BEEN RELEASED occurs whenever a released stack pointer is supplied as a stack descriptor argument for any purpose other than as a stack pointer to re-use.

Note: The creation of a single stack pointer can result in the retention of a large amount of stack space. Therefore, one should try to release stack pointers when they are no longer needed. See page 7.10.

(STKPOS NAME N POS OLDPOS) [Function]
Returns a stack pointer to the Nth frame with frame name NAME . The search begins with (and includes) the frame specified by the stack descriptor POS . The search proceeds along the control chain from POS if N is negative, or along the access chain if N is positive. If N is NIL, -1 is used. Returns a stack pointer to the frame if such a frame exists, otherwise returns NIL. If OLDPOS is supplied and is a stack pointer, it is reused. If OLDPOS is supplied and is a stack pointer and STKPOS returns NIL, OLDPOS is released. If OLDPOS is not a stack pointer it is ignored.

Note: (STKPOS 'STKPOS) causes an error, ILLEGAL STACK ARG; it is not permissible to create a stack pointer to the active frame.

(STKNTH N POS OLDPOS) [Function]
Returns a stack pointer to the Nth frame back from the frame specified by the stack descriptor POS . If N is negative, the control chain from POS is followed. If N is positive the access chain is followed. If N equals 0, STKNTH returns a stack pointer to POS (this provides a way to copy a stack pointer). Returns NIL if there are fewer than N frames in the appropriate chain. If OLDPOS is supplied and is a stack pointer, it is reused. If OLDPOS is not a stack pointer it is ignored.

Note: (STKNTH 0) causes an error, ILLEGAL STACK ARG; it is not possible to create a stack pointer to the active frame.

(STKNAME POS) [Function]
Returns the frame name of the frame specified by the stack descriptor POS .

(SETSTKNAME POS NAME) [Function]
Changes the frame name of the frame specified by POS to be NAME . Returns NAME .

(STKNTHNAME N POS) [Function]
Returns the frame name of the Nth frame back from POS . Equivalent to (STKNAME (STKNTH N POS)) but avoids creation of a stack pointer.

In summary, STKPOS converts function names to stack pointers, STKNTH converts numbers to stack pointers, STKNAME converts stack pointers to function names, and STKNTHNAME converts numbers to function names.

(DUMMYFRAMEP POS) [Function]
Returns T if the user never wrote a call to the function at POS , e.g. in Interlisp-10, DUMMYFRAMEP is T for *PROG*LAM, *ENV*, and FOOBLOCK frames (see block compiler, page 12.13).

REALFRAMEP and REALSTKNTH can be used to write functions which manipulate the stack and work on either interpreted or compiled code:

VARIABLE BINDINGS AND THE INTERLISP STACK

(REALFRAMEP POS INTERPFL G) [Function]
Returns POS if POS is a “real” frame, i.e. if POS is not a dummy frame and POS is a frame that does not disappear when compiled (such as COND); otherwise NIL. If INTERPFL G = T, returns POS if POS is not a dummy frame. For example, if (STKNAME POS) = COND, (REALFRAMEP POS) is NIL, but (REALFRAMEP POS T) is POS.

(REALSTKNTH N POS INTERPFL G OLDPOS) [Function]
Returns a stack pointer to the Nth (or -Nth) frames for which (REALFRAMEP POS INTERPFL G) is POS.

The following functions are used for accessing and changing bindings. Some of functions take an argument, N, which specifies a particular binding in the basic frame. If N is a literal atom, it is assumed to be the name of a variable bound in the basic frame. If N is a number, it is assumed to reference the Nth binding in the basic frame. The first binding is 1. If the basic frame contains no binding with the given name or if the number is too large or too small, the error ILLEGAL ARG occurs.

(STKSCAN VAR IPOS OPOS) [Function]
Searches beginning at IPOS for a frame in which a variable named VAR is bound. The search follows the access chain. Returns a stack pointer to the frame if found, otherwise returns NIL. If OPOS is a stack pointer it is reused, otherwise it is ignored.

(FRAMESCAN ATOM POS) [Function]
Returns the relative position of the binding of ATOM in the basic frame of POS. Returns NIL if ATOM is not found.

(STKARG N POS _) [Function]
Returns the value of the binding specified by N in the basic frame of the frame specified by the stack descriptor POS. N can be a literal atom or number.

(STKARGNAME N POS) [Function]
Returns the name of the binding specified by N, in the basic frame of the frame specified by the stack descriptor POS. N can be a literal atom or number.

(SETSTKARG N POS VALUE) [Function]
Sets the value of the binding specified by N in the basic frame of the frame specified by the stack descriptor POS. N can be a literal atom or a number. Returns value.

(SETSTKARGNAME N POS NAME) [Function]
Sets the NAME of the binding specified by N in the basic frame of the frame specified by the stack descriptor POS. N can be a literal atom or a number. Returns NAME.

(STKNARGS POS _) [Function]
Returns the number of arguments bound in the basic frame of the frame specified by the stack descriptor POS.

(VARIABLES POS) [Function]
Returns a list of the variables bound at POS.

As an example of the use of STKNARGS and STKARGNAME, VARIABLES could be defined by:

Stack Functions

```
(VARIABLES
  [LAMBDA (POS)
    (for N from 1 to (STKNARGS POS)
      collect (STKARGNAME N POS]))
```

(STKARGS POS _) [Function]
 Returns a list of the *values* of variables bound at POS .

The following functions are used to evaluate an expression in a different environment, and/or to alter the flow of control.

```
(ENVEVAL FORM APOS CPOS AFLG CFLG ) [Function]
```

Evaluates FORM in the environment specified by APOS and CPOS . That is, a new active frame is created with the frame specified by the stack descriptor APOS as its ALINK , and the frame specified by the stack descriptor CPOS as its CLINK . Then FORM is evaluated. If AFLG is not NIL, and APOS is a stack pointer, then APOS will be released. Similarly, if CFLG is not NIL, and CPOS is a stack pointer, then CPOS will be released.

```
(ENVAPPLY FN ARGS APOS CPOS AFLG CFLG ) [Function]
```

APPLYs FN to ARGS in the environment specified by APOS and CPOS . AFLG and CFLG have the same interpretation as with ENVEVAL.

```
(STKEVAL POS FORM FLG _ ) [Function]
```

Evaluates FORM in the access environment of the frame specified by the stack descriptor POS . If FLG is not NIL and POS is a stack pointer, releases POS . The definition of STKEVAL is (ENVEVAL FORM POS NIL FLG) .

```
(STKAPPLY POS FN ARGS FLG _ ) [Function]
```

Similar to STKEVAL but applies FN to ARGS .

```
(RETEVAL POS FORM FLG _ ) [Function]
```

Evaluates FORM in the access environment of the frame specified by the stack descriptor POS , and then returns from POS with that value. If FLG is not NIL and POS is a stack pointer, then POS is released. The definition of RETEVAL is equivalent to (ENVEVAL FORM POS (STKNTH -1 POS) FLG T) , except that RETEVAL does not create a stack pointer.

```
(RETAPPLY POS FN ARGS FLG _ ) [Function]
```

Similar to RETEVAL except applies FN to ARGS .

```
(RETFROM POS VAL FLG ) [Function]
```

Return from the frame specified by the stack descriptor POS , with the value VAL . If FLG is not NIL, and POS is a stack pointer, then POS is released. An attempt to RETFROM the top level (e.g., (RETFROM T)) causes an error, ILLEGAL STACK ARG. RETFROM can be written in terms of ENVEVAL as follows:

```
(RETFROM
  (LAMBDA (POS VAL FLG)
    (ENVEVAL (LIST 'QUOTE VAL)
      NIL
      (if (STKNTH -1 POS (if FLG then POS))
```

VARIABLE BINDINGS AND THE INTERLISP STACK

```
else (ERRORX (LIST 19 POS)))  
NIL  
T)))
```

(RETTO POS VAL FLG) [Function]
Like RETFROM, except returns *to* the frame specified by POS.

(EVALV VAR POS) [Function]
Evaluates VAR, where VAR is assumed to be a litatom, in the access environment specified by the stack descriptor POS. If VAR is unbound, EVALV returns NOBIND and does not generate an error. While EVALV could be defined as (ENVEVAL VAR POS) it is in fact a SUBR which is somewhat faster. EVALV compiles open when POS = NIL.

The following functions and variables are used to manipulate stack pointers.

(STACKP X) [Function]
Returns X if X is a stack pointer, otherwise returns NIL.

(RELSTK POS) [Function]
Release the stack pointer POS (see page 7.10). If POS is not a stack pointer, does nothing. Returns POS.

(RELSTKP X) [Function]
Returns T if X is a released stack pointer, NIL otherwise.

(CLEARSTK FLG) [Function]
If FLG is NIL, releases all active stack pointers, and returns NIL. If FLG is T, returns a list of all the active (unreleased) stack pointers.

CLEARSTKLST [Variable]
A variable used by top-level EVALQT. Every time EVALQT is re-entered (e.g., following errors, or control-D), CLEARSTKLST is checked. If its value is T, all active stack pointers are released using CLEARSTK. If its value is a list, then all stack pointers on that list are released. If its value is NIL, nothing is released. CLEARSTKLST is initially T.

NOCLEARSTKLST [Variable]
A variable used by top-level EVALQT. If CLEARSTKLST is T (see above) all active stack pointers *except* those on NOCLEARSTKLST are released. NOCLEARSTKLST is initially NIL.

Thus if one wishes to use multiple environments that survive through control-D, either CLEARSTKLST should be set to NIL, or else those stack pointers to be retained should be explicitly added to NOCLEARSTKLST.

(COPYSTK POS1 POS2) [Function]
(Interlisp-10) Copies the stack, including basic frames, from the frame specified by the stack descriptor POS1 to the frame specified by the stack descriptor POS2. That is, copies the frame extensions and basic frames in the access chain from POS2 to POS1 (inclusive). POS1 must be in the access chain of POS2, i.e., "above" POS2. Returns the new POS2. This provides a way to save an entire environment

Stack Functions

including variable bindings.

(MAPDL MAPDLFN MAPDLPOS) [Function]
Starts at MAPDLPOS and applies MAPDLFN, a function of two arguments, to the function *name* at each frame, and the frame (stack pointer) itself, until the top of the stack is reached. Returns NIL. For example,

```
[MAPDL (FUNCTION (LAMBDA (X POS)
                (if (IGREATERP (STKNARGS POS) 2)
                    then (PRINT X)))]
```

will print all functions of more than two arguments.

(SEARCHPDL SRCHFN SRCHPOS) [Function]
Similar to MAPDL, except searches the pushdown list starting at position SRCHPOS until it finds a frame for which SRCHFN, a function of two arguments applied to the *name* of the frame and the frame itself, is not NIL. Returns (NAME . FRAME) if such a frame is found, otherwise NIL.

(BACKTRACE IPOS EPOS FLAGS FILE PRINTFN) [Function]
Performs a backtrace beginning at the frame specified by the stack descriptor IPOS, and ending with the frame specified by the stack descriptor EPOS. FLAGS is a number in which the options of the BACKTRACE are encoded. If a bit is set, the corresponding information is included in the backtrace.

bit 0 - print arguments of non-SUBRs.

bit 1 - print temporaries of the interpreter.

bit 2 - print SUBR arguments and local variables.

bit 3 - omit printing of UNTRACE: and function names.

bit 4 - follow access chain instead of control chain.

bit 5 - print temporaries, i.e. the blips.

For example: if FLAGS = 47Q, everything is printed; if FLAGS = 21Q, follows the access chain, prints arguments.

FILE is the file that the backtrace is printed to. FILE must be open. PRINTFN is used when printing the values of variables, temporaries, blips, etc. PRINTFN = NIL defaults to PRINT.

(BACKTRACE IPOS EPOS SKIPFNS FLAGS FILE) [Function]
Prints a backtrace from IPOS to EPOS onto FILE. FLAGS specifies the options of the backtrace, e.g., do/don't print arguments, do/don't print temporaries of the interpreter, etc., and is the same as for BACKTRACE.²

²BACKTRACE calls BACKTRACE with a PRINTFN of SHOWPRINT (page 6.17), so that if SYSPRETTYFLG = T, the values will be prettyprinted.

VARIABLE BINDINGS AND THE INTERLISP STACK

SKIPFNS is a list of functions. As BAKTRACE scans down the stack, the stack name of each frame is passed to each function in SKIPFNS, and if any of them return non-NIL, POS is skipped (including all variables).

BAKTRACE collapses the sequence of several function calls corresponding to a call to a system package into a single "function" using BAKTRACELST as described below. For example, any call to the editor is printed as **EDITOR**, a break is printed as **BREAK**, etc.

BAKTRACE is used by the BT, BTV, BTV+, BTV*, and BTV! commands, with FLAG = 0, 1, 5, 7, and 47Q respectively.

BAKTRACELST

[Variable]

Used for telling BAKTRACE (therefore, the BT, BTV, etc. commands) to abbreviate various sequences of function calls on the stack by a single key, e.g. **BREAK**, **EDITOR**, etc.

The operation of BAKTRACE and format of BAKTRACELST is described so that the user can add his own entries to BAKTRACELST. Each entry on BAKTRACELST is a list of the form (FRAMENAME KEY . PATTERN) or (FRAMENAME (KEY₁ . PATTERN₁) (KEY_N . PATTERN_N)), where a pattern is a list of elements that are either atoms, which match a single frame, or lists, which are interpreted as a list of alternative patterns, e.g. (PROGN **BREAK** EVAL ((ERRORSET BREAK1A BREAK1) (BREAK1)))

BAKTRACE operates by scanning up the stack and, at each point, comparing the current frame name, with the frame names on BAKTRACELST, i.e. it does an ASSOC. If the frame name does appear, BAKTRACE attempts to match the stack as of that point with (one of) the patterns. If the match is successful, BAKTRACE prints the corresponding key, and continues with where the match left off. If the frame name does not appear, or the match fails, BAKTRACE simply prints the frame name and continues with the next higher frame (unless the SKIPFNS applied to the frame name are non-NIL as described above).

Matching is performed by comparing atoms in the pattern with the current frame name, and matching lists as patterns, i.e. sequences of function calls, always working up the stack. For example, either of the sequence of function calls " BREAK1 BREAK1A ERRORSET EVAL PROGN " or " BREAK1 EVAL PROGN " would match with the sample entry given above, causing **BREAK** to be printed.

Special features:

The litatom & can be used to match any frame.

The pattern "-" can be used to match nothing. - is useful for specifying an optional match, e.g. the example above could also have been written as (PROGN **BREAK** EVAL ((ERRORSET BREAK1A) -) BREAK1).

It is not necessary to provide in the pattern for matching dummy frames, i.e. frames for which DUMMYFRAMEP (see page 7.4) is true, e.g. in Interlisp-10, *PROG*LAM, *ENV*, NOLINKDEF1, etc. When working on a match, the matcher automatically skips over these frames when they do not match.

If a match succeeds and the KEY is NIL, nothing is printed. For example, (*PROG*LAM NIL EVALA *ENV). This sequence will occur following an error which then causes a break if some of the function's

Releasing and Reusing Stack Pointers

arguments are LOCALVARS.

7.3 RELEASING AND REUSING STACK POINTERS

The creation of a single stack pointer can result in the retention of a large amount of stack space. Furthermore, this space will not be freed until the next garbage collection, *even if the stack pointer is no longer being used*, unless the stack pointer is explicitly released or reused. If there is sufficient amount of stack space tied up in this fashion, a STACK OVERFLOW condition can occur, even in the simplest of computations. For this reason, the user should consider releasing a stack pointer when the environment referenced by the stack pointer is no longer needed.

The effects of releasing a stack pointer are:

- (1) The link between the stack pointer and the stack is broken by setting the contents of the stack pointer to the “released mark” (currently unboxed 0). A released stack pointer prints as #ADR /#0.
- (2) If this stack pointer was the last remaining reference to a frame extension; that is, if no other stack pointer references the frame extension and the extension is not contained in the active control or access chain, then the extension may be reclaimed, and is reclaimed immediately. The process repeats for the access and control chains of the reclaimed extension so that all stack space that was reachable only from the released stack pointer is reclaimed.

A stack pointer may be released using the function RELSTK, but there are some cases for which RELSTK is not sufficient. For example, if a function contains a call to RETFROM in which a stack pointer was used to specify where to return to, it would not be possible to simultaneously release the stack pointer. (A RELSTK appearing in the function following the call to RETFROM would not be executed!) To permit release of a stack pointer in this situation, the stack functions that relinquish control have optional flags arguments to denote whether or not a stack pointer is to be released (AFLG and CFLG). Note that in this case releasing the stack pointer will *not* cause the stack space to be reclaimed immediately because the frame referenced by the stack pointer will have become part of the active environment.

Another way of avoiding creating new stack pointers is to *reuse* stack pointers that are no longer needed. The stack functions that create stack pointers (STKPOS, STKNTH, and STKSCAN) have an optional argument which is a stack pointer to reuse. When a stack pointer is reused, two things happen. First the stack pointer is released (see above). Then the pointer to the new frame extension is deposited in the stack pointer. The old stack pointer (with its new contents) is the value of the function. Note that the reused stack pointer will be released even if the function does not find the specified frame.

Note that even if stack pointers are explicitly being released, *creation* of many stack pointers can cause a garbage collection of stack pointer space. Thus, if the user’s application requires creating many stack pointers, he definitely should take advantage of reusing stack pointers.

7.4 THE PUSH-DOWN LIST AND THE INTERPRETER

In addition to the names and values of arguments for functions, information regarding partially-evaluated expressions is kept on the push-down list. For example, consider the following definition of the function

VARIABLE BINDINGS AND THE INTERLISP STACK

FACT (intentionally faulty):

```
(FACT
  [LAMBDA (N)
    (COND
      ((ZEROP N)
       L)
      (T (ITIMES N (FACT (SUB1 N)))))
```

In evaluating the form (FACT 1), as soon as FACT is entered, the interpreter begins evaluating the implicit PROG N following the LAMBDA. The first function entered in this process is COND. COND begins to process its list of clauses. After calling ZEROP and getting a NIL value, COND proceeds to the next clause and evaluates T. Since T is true, the evaluation of the implicit PROG N that is the consequent of the T clause is begun. This requires calling the function ITIMES. However before ITIMES can be called, its arguments must be evaluated. The first argument is evaluated by retrieving the current binding of N from its value cell; the second involves a recursive call to FACT, and another implicit PROG N, etc.

Note that at each stage of this process, some portion of an expression has been evaluated, and another is awaiting evaluation. The output below (from Interlisp-10) illustrates this by showing the state of the push-down list at the point in the computation of (FACT 1) when the unbound atom L is reached.

```
_FACT(1)
u.b.a. L {in FACT} in ((ZEROP N) L)
(L broken)
:BTV!

*TAIL* (L)

*ARG1 (((ZEROP N) L) (T (ITIMES N (FACT (SUB1 N)))))
COND

*FORM* (COND ((ZEROP N) L) (T (ITIMES N (FACT (SUB1 N)))))
*TAIL* ((COND ((ZEROP N) L) (T (ITIMES N (FACT (SUB1 N)))))

N 0
FACT

*FORM* (FACT (SUB1 N))
*FN* ITIMES
*TAIL* ((FACT (SUB1 N)))
*ARGVAL* 1
*FORM* (ITIMES N (FACT (SUB1 N)))
*TAIL* ((ITIMES N (FACT (SUB1 N))))

*ARG1 (((ZEROP N) L) (T (ITIMES N (FACT (SUB1 N)))))
COND

*FORM* (COND ((ZEROP N) L) (T (ITIMES N (FACT (SUB1 N)))))
*TAIL* ((COND ((ZEROP N) L) (T (ITIMES N (FACT (SUB1 N)))))
```


VARIABLE BINDINGS AND THE INTERLISP STACK

7.5 GENERATORS AND COROUTINES

This section describes an application of the spaghetti stack facility to provide mechanisms for creating and using simple generators, generalized coroutines, and Conniver style possibility lists.

7.5.1 Generators

A *generator* is like a subroutine except that it retains information about previous times it has been called. Some of this state may be data (for example, the seed in a random number generator), and some may be in program state (as in a recursive generator which finds all the atoms in a list structure). For example, if LISTGEN is defined as:

```
(LISTGEN (L)
  (IF L THEN (PRODUCE (CAR L))
            (LISTGEN (CDR L))))
```

we can use the function GENERATOR (described below) to create a generator that uses LISTGEN to produce the elements of a list one at a time, e.g.,

```
(SETQ GR (GENERATOR (LISTGEN '(A B C))))
```

creates a generator, which can be called by

```
(GENERATE GR)
```

to produce as values on successive calls, A, B, C. When GENERATE (not GENERATOR) is called the first time, it simply starts evaluating (LISTGEN '(A B C)). PRODUCE gets called from LISTGEN, and pops back up to GENERATE with the indicated value after saving the state. When GENERATE gets called again, it continues from where the last PRODUCE left off. This process continues until finally LISTGEN completes and returns a value (it doesn't matter what it is). GENERATE then returns GR itself as its value, so that the program that called GENERATE can tell that it is finished, i.e., there are no more values to be generated.

```
(GENERATOR FORM  $\alpha\alpha$  COMV AR  $\alpha\alpha$ ) [NLambda Function]
```

An nlambda function that creates a generator which uses FORM $\alpha\alpha$ to compute values. GENERATOR returns a *generator handle* which is represented by a dotted pair of stack pointers.

COMV AR $\alpha\alpha$ is optional. If its value (EVAL of) is a generator handle, the list structure and stack pointers will be reused. Otherwise, a new generator handle will be constructed.

GENERATOR compiles open.

```
(PRODUCE VAL) [Function]
```

Used from within (below) a generator to return VAL as the value of the corresponding call to GENERATE.

```
(GENERATE HANDLE VAL) [Function]
```

Restarts the generator represented by HANDLE. VAL is returned as the value of

Coroutines

the `PRODUCE` which last suspended the operation of the generator. When the generator runs out of values, `GENERATE` returns `HANDLE` itself.

Examples:

The following function will go down recursively through a list structure and produce the atoms in the list structure one at a time.

```
[LEAVESG (L)
  (if (ATOM L)
    then (PRODUCE L)
    else (LEAVESG (CAR L))
    (if (CDR L)
      then (LEAVESG (CDR L))])]
```

The following function prints each of these atoms as it appears. It illustrates how a loop can be set up to use a generator.

```
(PLEAVESG1 (L)
  (PROG (X LHANDLE)
    (SETQ LHANDLE (GENERATOR (LEAVESG L)))
    LP (SETQ X (GENERATE LHANDLE))
      (if (EQ X LHANDLE)
        then (RETURN NIL))
      (PRINT X)
      (GO LP)))
```

Note that the loop terminates when the value of the generator is `EQ` to the dotted pair which is the value produced by the call to `GENERATOR`. A CLISP iterative operator, `OUTOF`, is provided which makes it much easier to write the loop in `PLEAVESG1`. `OUTOF` (or `outof`) can precede a form which is to be used as a generator. On each iteration, the iteration variable will be set to successive values returned by the generator; the loop will be terminated automatically when the generator runs out. Therefore, the following is equivalent to the above program `PLEAVESG1`:

```
(PLEAVESG2 (L)
  (for X outof (LEAVESG L) do (PRINT x)))
```

Here is another example; the following form will print the first `N` atoms.

```
(for X outof (MAPATOMS (FUNCTION PRODUCE))
  as I from 1 to N do (PRINT X))
```

7.5.2 Coroutines

This package provides facilities for the creation and use of fully general coroutine structures. It uses a stack pointer to preserve the state of a coroutine, and allows arbitrary switching between N different coroutines, rather than just a call to a generator and return. This package is slightly more efficient than the generator package described above, and allows more flexibility on specification of what to do when a coroutine terminates.

VARIABLE BINDINGS AND THE INTERLISP STACK

```
(COROUTINE CALLPTR qq COR OUTPTR qq COR OUTFORM qq ENDFORM qq)
```

[NLambda Function]

This lambda function is used to create a coroutine and initialize the linkage. CALLPTR qq and COR OUTPTR qq are the names of two variables, which will be set to appropriate stack pointers. If the values of CALLPTR qq or COR OUTPTR qq are already stack pointers, the stack pointers will be reused. COR OUTFORM qq is the form which is evaluated to start the coroutine; ENDFORM qq is a form to be evaluated if COR OUTFORM qq actually returns when it runs out of values.

COROUTINE compiles open.

```
(RESUME FROMPTR TOPTR VAL)
```

[Function]

Used to transfer control from one coroutine to another. FROMPTR should be the stack pointer for the current coroutine, which will be smashed to preserve the current state. TOPTR should be the stack pointer which has preserved the state of the coroutine to be transferred to, and VAL is the value that is to be returned to the latter coroutine as the value of the RESUME which suspended the operation of that coroutine.

For example, the following is the way one might write the LEAVES program using the coroutine package:

```
(LEAVESC (L COROUTPTR CALLPTR)
  (if (ATOM L)
    then (RESUME COROUTPTR CALLPTR L)
    else (LEAVESC (CAR L) COROUTPTR CALLPTR)
        (if (CDR L) then (LEAVESC (CDR L) COROUTPTR CALLPTR))))
```

A function PLEAVESC which uses LEAVESC can be defined as follows:

```
(PLEAVESC (L)
  (bind PLHANDLE LHANDLE
    first (COROUTINE PLHANDLE LHANDLE
          (LEAVESC L LHANDLE PLHANDLE)
          (RETFROM 'PLEAVESC))
    do (PRINT (RESUME PLHANDLE LHANDLE))))
```

By RESUMEing LEAVESC repeatedly, this function will print all the leaves of list L and then return out of PLEAVESC via the RETFROM. The RETFROM is necessary to break out of the non-terminating do-loop. This was done to illustrate the additional flexibility allowed through the use of ENDFORM qq.

We use two coroutines working on two trees in the example EQLEAVES, defined below. EQLEAVES tests to see whether two trees have the same leaf set in the same order, e.g., (EQLEAVES '(A B C) '(A B C)) is true.

```
(EQLEAVES (L1 L2)
  (bind LHANDLE1 LHANDLE2 PE EL1 EL2
    first (COROUTINE PE LHANDLE1 (LEAVESC L1 LHANDLE1 PE) 'NO-MORE)
          (COROUTINE PE LHANDLE2 (LEAVESC L2 LHANDLE2 PE) 'NO-MORE)
    do (SETQ EL1 (RESUME PE LHANDLE1))
        (SETQ EL2 (RESUME PE LHANDLE2))
        (if (NEQ EL1 EL2)
          then (RETURN NIL))
```

Possibilities Lists

```
repeatuntil (EQ EL1 'NO-MORE)
finally (RETURN T))
```

7.5.3 Possibilities Lists

A possibilities list is the interface between a generator and a consumer. The possibilities list is initialized by a call to `POSSIBILITIES`, and elements are obtained from it by using `TRYNEXT`. By using the spaghetti stack to maintain separate environments, this package allows a regime in which a generator can put a few items in a possibilities list, suspend itself until they have been consumed, and be subsequently aroused and generate some more.

(`POSSIBILITIES` `FORM` `q q`) [NLambda Function]

This nlambda function is used for the initial creation of a possibilities list. `FORM` `q q` will be evaluated to create the list. It should use the functions `NOTE` and `AU-REVOIR` described below to generate possibilities. Normally, one would set some variable to the possibilities list which is returned, so it can be used later, e.g.:

```
(SETQ PLIST (POSSIBILITIES (GENERFN V1 V2))).
```

`POSSIBILITIES` compiles open.

(`NOTE` `VAL` `LSTFL G`) [Function]

Used within a generator to put items on the possibilities list being generated. If `LSTFL G` is equal to `NIL`, `VAL` is treated as a single item. If `LSTFL G` is non-`NIL`, then the list `VAL` is `NCONC`d on the end of the possibilities list. Note that it is perfectly reasonable to create a possibilities list using a second generator, and `NOTE` that list as possibilities for the current generator with `LSTFL G` equal to `T`. The lower generator will be resumed at the appropriate point.

(`AU-REVOIR` `VAL` `q q`) [NoSpread Function]

Puts `VAL` `q q` on the possibilities list if it is given, and then suspends the generator and returns to the consumer in such a fashion that control will return to the generator at the `AU-REVOIR` if the consumer exhausts the possibilities list.

Note: `NIL` is not put on the possibilities list unless it is explicitly given as an argument to `AU-REVOIR`, i.e., (`AU-REVOIR`) and (`AU-REVOIR` `NIL`) are *not* the same. `AU-REVOIR` and `ADIEU` are lambda nospreads to enable them to distinguish these two cases.

(`ADIEU` `VAL` `q q`) [NoSpread Function]

Like `AU-REVOIR` except releases the generator instead of suspending it.

(`TRYNEXT` `PLST` `q q` `ENDF ORM` `q q` `VAL` `q q`) [NLambda Function]

This nlambda function allows a consumer to use a possibilities list. It removes the `rst` item from the possibilities list named by `PLST` `q q` (i.e. `PLST` `q q` must be an atom whose value is a possibilities list), and returns that item, provided it is not a generator handle. If a generator handle is encountered, the generator is reawakened. When it returns a possibilities list, this list is added to the front of the current list. When a call to `TRYNEXT` causes a generator to be awakened, `VAL` `q q` is returned as the value of the `AU-REVOIR` which put that generator to sleep. If `PLST` `q q` is empty, it evaluates `ENDF ORM` `q q` in the caller's environment.

VARIABLE BINDINGS AND THE INTERLISP STACK

TRYNEXT compiles open.

```
(CLEANPOSLST PLST )
```

[Function]
This function is provided to release any stack pointers which may be left in the PLST which was not used to exhaustion.

For example, FIB is a generator for fibonacci numbers. It starts out by NOTEing its two arguments, then suspends itself. Thereafter, on being re-awakened, it will NOTE two more terms in the series and suspends again. PRINTFIB uses FIB to print the first N fibonacci numbers.

```
[FIB (F1 F2)
  (do (NOTE F1)
      (NOTE F2)
      (SETQ F1 (IPLUS F1 F2))
      (SETQ F2 (IPLUS F1 F2))
      (AU-REVOIR)]
```

Note that this AU-REVOIR just suspends the generator and adds nothing to the possibilities list except the generator.

```
[PRINTFIB (N)
  (PROG ((FL (POSSIBILITIES (FIB 0 1))))
        (RPTQ N (PRINT (TRYNEXT FL)))
        (CLEANPOSLST FL)]
```

Note that FIB itself will never terminate.

Possibilities Lists