

# CARNEGIE-MELLON UNIVERSITY

DEPARTMENT OF COMPUTER SCIENCE

SPICE PROJECT

---

Revised Internal Design of Spice Lisp

Skef Wholey  
Scott F. Fahlman  
Joseph Ginder

20 December 1983

---

DRAFT

Spice Document S026 [Revised]

Keywords and index categories: Lisp

Location of machine-readable file: CMUC::<Wholey.Australia>Revguts.Mss

Copyright © 1983 Carnegie-Mellon University

Supported by the Defense Advanced Research Projects Agency, Department of Defense, ARPA Order 3597, monitored by the Air Force Avionics Laboratory under contract F33615-78-C-1551. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

# Table of Contents

<b>1. Introduction</b>	<b>2</b>
1.1. Scope and Purpose	2
1.2. Notational Conventions	2
<b>2. Data Types and Object Formats</b>	<b>3</b>
2.1. Lisp Objects	3
2.2. Table of Type Codes	3
2.3. Table of Space Codes	4
2.4. Immediate Data Type Descriptions	4
2.5. Pointer-Type Objects and Spaces	5
2.6. Forwarding Pointers	8
2.7. System and Stack Spaces	8
2.8. Vectors and Arrays	9
2.8.1. General Vectors	9
2.8.2. Integer Vectors	10
2.8.3. Arrays	12
2.9. Symbols Known to the Microcode	13
<b>3. Runtime Environment</b>	<b>14</b>
3.1. Control Registers	14
3.2. Function Object Format	15
3.3. Control-Stack Format	16
3.3.1. Call Frames	16
3.3.2. Catch Frames	17
3.4. Binding-Stack Format	17
<b>4. Storage Management</b>	<b>18</b>
4.1. The Transporter	18
4.2. The Scavenger	19
4.3. Purification	19
<b>5. Macro Instruction Set</b>	<b>21</b>
5.1. Macro-Instruction Formats	21
5.2. Instructions	22
5.2.1. Allocation	23
5.2.2. Stack Manipulation	25
5.2.3. List Manipulation	26
5.2.4. Symbol Manipulation	28
5.2.5. Array Manipulation	29
5.2.6. Type Predicates	33
5.2.7. Arithmetic and Logic	35
5.2.8. Branching and Dispatching	39
5.2.9. Function Call and Return	40
5.2.10. Miscellaneous	41
5.2.11. System Hacking	43
<b>6. Control Conventions</b>	<b>46</b>
6.1. Function Calls	46

6.1.1. Starting a Function Call	46
6.1.2. Finishing a Function Call	47
6.1.3. Returning from a Function Call	48
6.1.4. Returning Multiple-Values	48
6.2. Non-Local Exits	49
6.3. Escaping to Macrocode	51
6.4. Errors	53
6.5. Trapping to the Accent Kernel	57
6.6. Interrupts	57
<b>Appendix I. Fasload File Format</b>	<b>59</b>
1.1. General	59
1.2. Strategy	60
1.3. Fasload Language	61
<b>Appendix II. The Opcode Definition File</b>	<b>69</b>
<b>Index</b>	<b>81</b>

### **Acknowledgments**

The following people have been contributors to this and earlier versions of the design of the Spice Lisp instruction set: Guy L. Steele Jr., Gail E. Kaiser, Walter van Roggen, Neal Feinberg, Jim Large, and Rob MacLachlan. The original instruction set was loosely based on the MIT Lisp Machine instruction set.

The FASL file format was designed by Guy L. Steele Jr. and Walter van Roggen, and the appendix on this subject is their document with very few modifications.

# 1. Introduction

## 1.1. Scope and Purpose

NOTE: This document describes a new implementation of Spice Lisp as it is to be implemented on the PERQ, running the Spice operating system, Accent. This new design is undergoing rapid change, and for the present is not guaranteed to accurately describe any past, present, or future implementation of Spice Lisp. All questions and comments on this material should be directed to Skef Wholey (Wholey@CMU-CS-C).

This document specifies the instruction set and virtual memory architecture of the PERQ Spice Lisp system. This is a working document, and it will change frequently as the system is developed and maintained. If some detail of the system does not agree with what is specified here, it is to be considered a bug.

Spice Lisp will be implemented on other microcodable machines, and implementations of Common Lisp based on Spice Lisp exist for conventional machines with fixed instructions sets. These other implementations are very different internally and are described in other documents.

## 1.2. Notational Conventions

Spice Lisp objects are 32 bits long. The low-order bit of each word is numbered 0; the high-order bit is numbered 31. If a word is broken into smaller units, these are packed into the word from right to left. For example, if we break a word into bytes, byte 0 would occupy bits 0-7, byte 1 would occupy 8-15, byte 2 would occupy 16-23, and byte 3 would occupy 24-31. In these conventions we follow the conventions of the VAX; the PDP-10 family follows the opposite convention, packing and numbering left to right.


All Spice Lisp documentation uses decimal as the default radix; other radices will be indicated by a subscript (as in  $77_8$ ) or by a clear statement of what radix is in use.

## 2. Data Types and Object Formats


### 2.1. Lisp Objects

Lisp objects are 32 bits long. They come in 32 basic types, divided into three classes: immediate data types, pointer types, and forwarding pointer types. The storage formats are as follows:

#### Immediate Data Types:



#### Pointer and Forwarding Types:



### 2.2. Table of Type Codes

Code	Type	Class	Explanation
0	Misc	Immediate	Trap object, stacks, system tables
1	Bit-Vector	Pointer	Vector of bits
2	Integer-Vector	Pointer	Vector of integers
3	String	Pointer	Character string
4	Bignum	Pointer	Bignum
5	Long-Float	Pointer	Long float
6	Complex	Pointer	Complex number
7	Ratio	Pointer	Ratio
8	General-Vector	Pointer	Vector of Lisp objects
9	Function	Pointer	Compiled function header
10	Array	Pointer	Array header
11	Symbol	Pointer	Symbol
12	List	Pointer	Cons cell
13-15	Unused		
16	+ Fixnum	Immediate	Fixnum >= 0
17	- Fixnum	Immediate	Fixnum < 0
18	+ Short-Float	Immediate	Short float >= 0
19	- Short-Float	Immediate	Short float < 0
20	Character	Immediate	Character object
21	Values-Marker	Immediate	Multiple values marker
22	Call-Header	Immediate	Control stack call frame header
23	Catch-Header	Immediate	Control stack catch frame header
24	Catch-All	Immediate	Catch-All object
25	GC-Forward	Forward	Object in newspace of same type
26-31	Unused		

### 2.3. Table of Space Codes

Code	Space	Explanation
0	Dynamic-0	Storage normally garbage collected, space 0.
1	Dynamic-1	Storage normally garbage collected, space 1.
2	Static	Permanent objects, never moved or reclaimed.
3	Read-Only	Objects never moved, reclaimed, or altered.

### 2.4. Immediate Data Type Descriptions

**Misc** Reserved for assorted internal values. Bits 25-26 specify a sub-type code.

**0 Trap** Illegal object trap. If you fetch one of these, it's an error except under very specialized conditions. Note that a word of all zeros is of this type, so this is useful for trapping references to uninitialized memory. This value also is used in symbols to flag an undefined value or definition.

**1 Control-Stack-Pointer**

The low 25 bits are a pointer into the control stack. This is a word pointer that points to the proper virtual memory address. Pointers of this form are returned by certain system routines for use by debugging programs.

**2 Binding-Stack-Pointer**


The low 25 bits are a pointer into the binding stack. This is a word pointer that points to the proper virtual memory address. Pointers of this form are returned by certain system routines for use by debugging programs.

**3 System-Table-Pointer**

The low 25 bits are a pointer into an area of memory used for system tables. This is a word pointer into an area of the address space reserved for data sent and received in Accent messages.

**Fixnum** A 28-bit two's complement integer. The sign bit is stored as part of the type code.


**Short-Float** As with fixnums, the sign bit is stored as part of the type code. The format of short floating point number can be viewed as follows:



The sign of the mantissa is moved to the left so that these flonums can be compared just like fixnums. The exponent is the binary two's complement exponent of the number, plus 128; then, if the mantissa is negative, the bits of the exponent field are inverted. The mantissa is a 21-bit two's complement number with the sign moved to bit 27 and the

leading significant bit (which is always the complement of the sign bit and hence carries no information) stripped off. The short flonum representing 0.0 has 0's in bits 0 - 27. It is illegal for the sign bit to be 1 with all the other bits equal to 0. This encoding gives a range of about  $10^{-38}$  to  $10^{+38}$  and about 6 digits of accuracy. Note that long-flonums are available for those wanting more accuracy, but they are less efficient to use because they generate garbage that must be collected later.

**Character** A character object holding a character code, control bits, and font in the following format:



**Values-Marker** Used to mark the presence of multiple values on the stack. The low 16 bits indicate how many values are being returned. Note then, that only 65535 values can be returned from a multiple-values producing form. These are pushed in order, then the Values-Marker is pushed.

**Call-Header** Marks the start of each call frame on the control stack. The low-order 27 bits are used as a place to stash information for certain special kinds of calls.

For a normal function call, as created by the CALL or CALL-0 instruction, the low 27 bits are always 0.

Bit 22, if 1, indicates an "escape to macro" call frame, created when a macro-instruction cannot be completed entirely within the microcode. In this case, bits 16-17 indicate where the result is supposed to go (see section 6.3).

Bit 21, if 1, indicates a call frame that will accept multiple values to be returned. Such frames are created by Call-Multiple, and cause Return to take certain special actions. See section 6.1.3 for details.

Bits 22 and 21 are mutually exclusive. It is undefined what happens when both of these are on at once.

**Catch-Header** Marks a catch frame on the control stack. If bit 21 is on, this indicates that the catching form will accept multiple values. See section 6.2 for details.

**Catch-All** Object used as the catch tag for unwind-protects. Special things happen when a catch frame with this as its tag is encountered during a throw. See section 6.2 for details.

## 2.5. Pointer-Type Objects and Spaces

Each of the pointer-type lisp objects points into a different space in virtual memory. There are separate spaces for Bit-Vectors, Symbols, Lists, and so on. The 5-bit type-code provides the high-order virtual address bits for the object, followed by the 2-bit space code, followed by the 25-bit pointer address. This gives a 31-bit



virtual address to a 32-bit word; since the PERQ is a word-addressed machine, the low-order bit will be 0, and under Accent, the high order bit will be 0 (because the operating system lives in the upper half of the address space). This leaves us with a 30-bit virtual address. In effect we have carved a 30-bit space into a fixed set of 24-bit subspaces, not all of which are used.

The space code divides each of the type spaces into four sub-spaces, as shown in the table above. At any given time, one of the dynamic spaces is considered newspace, while the other is oldspace. The garbage collector continuously moves accessible objects from oldspace into newspace. When oldspace contains no more accessible objects it is considered empty. A "flip" can then be done, turning the old newspace into the new oldspace. All type-spaces are flipped at once. Allocation of new dynamic objects always occurs in newspace.

Optionally, the user (or system functions) may allocate objects in static or read-only space. Such objects are never reclaimed once they are allocated -- they occupy the space in which they were initially allocated for the lifetime of the Lisp process. The advantage of static allocation is that the GC never has to move these objects, thereby saving a significant amount of work, especially if the objects are large. Objects in read-only space are static, in that they are never moved or reclaimed; in addition, they cannot be altered once they are set up. Pointers in read-only space may only point to read-only or static space, never to dynamic space. This saves even more work, since read-only space does not need to be scavenged, and pages of read-only material do not need to be written back onto the disk during paging.

Objects in a particular type-space will contain either pointers to garbage-collectable objects or words of raw non-garbage-collectable bits, but not both. Similarly, a space will contain either fixed-length objects or variable-length objects, but not both. A variable-length object always contains a 24-bit length field right-justified in the first word, with the fixnum type-code in the high-order four bits. The remaining four bits can be used for sub-type information. The length field gives the size of the object in 32-bit words, including the header word. The garbage collector needs this information when the object is moved, and it is also useful for bounds checking.

The format of objects in each space are as follows:


Symbol	Each symbol is represented as a fixed-length block of boxed Lisp cells. The number of cells per symbol is 5, in the following order:
	0 Value cell for shallow binding.
	1 Definition cell: a function or list.
	2 Property list: a list of attribute-value pairs.
	3 Print name: a string.
	4 Package: the obarray holding this symbol.

- List                    A fixed-length block of two boxed Lisp cells, the CAR and the CDR.
  
- General-Vector        Vector of lisp objects, any length. The first word is a fixnum giving the number of words allocated for the vector (up to 24 bits). The highest legal index is this number minus 2. The second word is vector entry 0, and additional entries are allocated contiguously in virtual memory. General vectors are sometimes called G-Vectors. (See section 2.8 for further details.)
  
- Integer-Vector        Vector of integers, any length. The 24 low bits of the first word give the allocated length in 32-bit words. The low-order 28 bits of the second word gives the length of the vector in entries, whatever the length of the individual entries may be. The high-order 4 bits of the second word contain access-type information that yields, among other things, the number of bits per entry. Entry 0 is right-justified in the third word of the vector. Bits per entry will normally be powers of 2, so they will fit neatly into 32-bit words, but if necessary some empty space may be left at the high-order end of each word. Integer vectors are sometimes called I-Vectors. (See section 2.8 for details.)

Bit-Vector            Vector of bits, any length. Bit-Vectors are represented in a form identical to I-Vectors, but live in a different space for efficiency reasons.

Bignum                Bignums are infinite-precision integers, represented in a format identical to I-Vectors. Each bignum is stored as a series of 8-bit bytes, with the low-order byte stored first. The representation is two's complement, but the sign of the number is redundantly encoded in the subtype field of the bignum: positive bignums are sub-type 0, negative bignums sub-type 1. The access-type code is always 8-Bit.

Long-Float            Long floats are stored as two consecutive words of bits, in the following format:



The exponent is biased by 1023. Exponents of 0 and 2047 are reserved. The most significant bit of the fraction is stripped off since it is always the complement of the sign bit, and carries no information.

Ratio                 Ratios are stored as two consecutive words of Lisp objects, which should both be integers.

Complex               Complex numbers are stored as two consecutive words of Lisp objects, which should both be numbers.

Array                 This is actually a header which holds the accessing information and other information about the array. The actual array contents are held in a vector (either an I-Vector or G-Vector) pointed to by an entry in the header. The header is identical in format to a G-Vector. For details on what the array header contains, see section 2.8.3.

String                A vector of bytes. Identical in form to I-Vectors with the access type always 8-Bit.

However, instead of accepting and returning fixnums, string accesses accept and return character objects. Only the 8-bit code field is actually stored, and the returned character object always has bit and font values of 0.

**Function** A compiled Spice Lisp function consists of both lisp objects and raw bits for the code. The Lisp objects are stored in the Function space in a format identical to that used for general vectors, with a 24-bit length field in the first word. This object contains assorted parameters needed by the calling machinery, a pointer to an 8-bit I-Vector containing the compiled byte codes, a number of pointers to symbols used as special variables within the function, and a number of lisp objects used as constants by the function. For details of the code format and definitions of the byte codes, see section 5.1.

## 2.6. Forwarding Pointers

**GC-Forward** When a data structure is transported into newspace, a GC-Forward pointer is left behind in the first word of the oldspace object. This points to the same type-space in which it is found. For example, a GC-Forward in G-Vector space points to a structure in the G-Vector newspace. GC-Forward pointers are only found in oldspace.

## 2.7. System and Stack Spaces

The virtual addresses below  $08000000_{16}$  are not occupied by Lisp objects, since Lisp objects with type code 0 are immediate objects. Some of this space is used for other purposes by Lisp. The current allocations are as follows:

Address (base 16)	Use
-----	---
00000000 - 01FFFFFF	Microcode tables
02000000 - 03FFFFFF	Control Stack
04000000 - 05FFFFFF	Binding Stack
06000000 - 07FFFFFF	System tables

Microcode tables for a given process are never accessed by Lisp-level code from that process (the SAVE function looks at the allocation table of another Lisp process). These tables contain allocation information for the various spaces and pointers to functions that are called when escapes to macrocode are done. There are currently two microcode tables:

Address (base 16)	Use
-----	---
00010000 - 00010100	Allocation table
00020000 - 00020100	Escape function table

The format of the allocation table is described in chapter 4, and the format of the escape function table is described in section 6.3.

The control stack grows upward (toward higher addresses) in memory, and is a framed stack. It contains

only general Lisp objects (with some random things encoded as fixnums or Misc codes). Every object pointed to by an entry on this stack is kept alive. The frame for a function call contains an area for the function's arguments, an area for local variables, a pointer to the caller's frame, and a pointer into the binding stack. The frame for a Catch form contains similar information. The precise stack format can be found in chapter 3.

The special binding stack also grows upward. This stack is used to hold previous values of special variables that have been bound. It grows and shrinks with the depth of the binding environment, as reflected in the control stack. This stack contains symbol-value pairs, with only boxed Lisp objects present.

System table space is used to interface Lisp to the operating system. This is the only part of the address space that contains invalid memory, so all Accent messages received will appear in this space. Since files are sent and received in messages, all files will be mapped into this space. Data in system table space may be accessed and altered by the instructions described in section 5.2.11.

There are significant performance advantages to be gained by aligning all objects on the PERQ's "quad-word" (64-bit) boundaries. This happens automatically for conses, long-floats, complex numbers, and ratios, which are all two Lisp-words long. For all other pointer-type objects, the allocator makes sure that the object starts on a quad-word boundary, wasting a word with a Misc-Trap code if necessary. Thus, every pointer (except possibly for stack and system area pointers) will have its two low-order bits 0. User-level code should never have to notice this distinction.

## 2.8. Vectors and Arrays

Common Lisp arrays can be represented in a few different ways in Spice Lisp -- different representations have different performance advantages. Simple general vectors, simple vectors of integers, and simple strings are basic Spice Lisp data types, and access to these structures is quicker than access to non-simple (or "complex") arrays. However, all multi-dimensional arrays in Spice Lisp are complex arrays, so references to these is always through a header structure.

### 2.8.1. General Vectors

G-Vectors contain Lisp objects. The format is as follows:

```
-----
| Fixnum code (4) | Subtype (4) |   Allocated length (24)   |
-----
| Vector entry 0   (Additional entries in subsequent words) |
-----
```

Note that the subtype field overlaps the type field -- this means that the subtype can change the sign bit of

the fixnum. The first word of the vector is a header indicating its length; the remaining words hold the boxed entries of the vector, one entry per 32-bit word. The header word is of type fixnum. It contains a 3-bit subtype field, which is used to indicate several special types of general vectors. At present, the following subtype codes are defined:

- 0 Normal. Used for assorted things.
- 1 Named structure created by DEFSTRUCT, with type name in entry 0.
- 2 FQ Hash Table, last rehashed in dynamic-0 space.
- 3 FQ Hash Table, last rehashed in dynamic-1 space.
- 4 FQ Hash Table, must be rehashed.

Following the subtype is a 24-bit field indicating how many 32-bit words are allocated for this vector, including the header word. Legal indices into the vector range from zero to the number in the allocated length field minus 2, inclusive. The index is checked on every access to the vector. Entry 0 is stored in the second word of the vector, and subsequent entries follow contiguously in virtual memory.


Once a vector has been allocated, it is possible to reduce its length by using the Shrink-Vector instruction, but never to increase its length, even back to the original size, since the space freed by the reduction may have been reclaimed. This reduction simply stores a new smaller value in the length field of the header word.

It is not an error to create a vector of length 0, though it will always be an out-of-bounds error to access such an object. The maximum possible length for a general vector is  $2^{24}-2$  entries, and that is only possible if no other general vectors are present in the space.

Objects of type Function and Array are identical in format to general vectors, though they have their own spaces. In both cases, only 0 is currently used in the sub-type field of the header word.

### 2.8.2. Integer Vectors

I-Vectors contain unboxed items of data, and their format is more complex. The data items come in a variety of lengths, but are of constant length within a given vector. Data going to and from an I-Vector are passed as Fixnums, right justified. Internally these integers are stored in packed form, filling 32-bit words without any type-codes or other overhead. The format is as follows:



Note that the subtype field overlaps the type field -- this means that the subtype can change the sign bit of the fixnum. The first word of an I-Vector contains the Fixnum type-code in the top 4 bits, a 4-bit subtype code in the next four bits, and the total allocated length of the vector (in 32-bit words) in the low-order 24 bits. At present, the following subtype codes are defined:

- 0            Normal. Used for assorted things.
- 1            Code. This is the code-vector for a function object.

The second word of the vector is the one that is looked at every time the vector is accessed. The low-order 28 bits of this word contain the number of valid entries in the vector, regardless of how long each entry is. The lowest legal index into the vector is always 0; the highest legal index is one less than this number-of-entries field from the second word. These bounds are checked on every access. Once a vector is allocated, it can be reduced in size but not increased. The Shrink-Vector instruction changes both the allocated length field and the number-of-entries field of an integer vector.

The high-order 4 bits of the second word contain an access-type code which indicates how many bits are occupied by each item (and therefore how many items are packed into a 32-bit word). The encoding is as follows:

0	1-Bit	8	Unused
1	2-Bit	9	Unused
2	4-Bit	10	Unused
3	8-Bit	11	Unused
4	16-Bit	12	Unused
5	Unused	13	Unused
6	Unused	14	Unused
7	Unused	15	Unused

In I-Vectors, the data items are packed into the third and subsequent words of the vector. Item 0 is right justified in the third word, item 1 is to its left, and so on until the allocated number of items has been accommodated. All of the currently-defined access types happen to pack neatly into 32-bit words, but if this should not be the case, some unused bits would remain at the left side of each word. No attempt will be made to split items between words to use up these odd bits. When allocated, an I-Vector is initialized to all 0's.

As with G-Vectors, it is not an error to create an I-Vector of length 0, but it will always be an error to access such a vector. The maximum possible length of an I-Vector is  $2^{28}-1$  entries or  $2^{24}-3$  words, whichever is smaller.

Objects of type String are identical in format to I-Vectors, though they have their own space. Strings always have subtype 0 and access-type 3 (8-Bit). Strings differ from normal I-Vectors in that the accessing instructions accept and return objects of type Character rather than Fixnum.

Bignums are also identical in format and operation to I-Vectors, though they may also be operated on directly by microcoded routines. For details of the currently-defined sub-types and their access-codes, see section 2.5.

### 2.8.3. Arrays

An array header is identical in form to a G-Vector. Like any G-Vector, its first word contains a fixnum type-code, a 4-bit subtype code, and a 24-bit total length field (this is the length of the array header, not of the vector that holds the data). At present, the subtype code is always 0. The entries in the header-vector are interpreted as follows:

**0 Data Vector** This is a pointer to the I-Vector, G-Vector, or string that contains the actual data of the array. In a multi-dimensional array, the supplied indices are converted into a single 1-D index which is used to access the data vector in the usual way.

**1 Number of Elements**

This is a fixnum indicating the number of elements for which there is space in the data vector.

**2 Fill Pointer**

This is a fixnum indicating how many elements of the data vector are actually considered to be in use. Normally this is initialized to the same value as the Number of Elements field, but in some array applications it will be given a smaller value. Any access beyond the fill pointer is illegal.

**3 Displacement**

This fixnum value is added to the final code-vector index after the index arithmetic is done but before the access occurs. Used for mapping a portion of one array into another. For most arrays, this is 0.

**4 Range of First Index**

This is the number of index values along the first dimension, or one greater than the largest legal value of this index (since the arrays are always zero-based). A fixnum in the range 0 to  $2^{24}-1$ . If any of the indices has a range of 0, the array is legal but will contain no data and accesses to it will always be out of range. In a 0-dimension array, this entry will not be present.

**5 - N Ranges of Subsequent Dimensions**

The number of dimensions of an array can be determined by looking at the length of the array header. The rank will be this number minus 6. The maximum array rank is 65535 - 6, or 65529.

The ranges of all indices are checked on every access, during the conversion to a single data-vector index. In this conversion, each index is added to the accumulating total, then the total is multiplied by the range of the following dimension, the next index is added in, and so on. In other words, if the data vector is scanned linearly, the last array index is the one that varies most rapidly, then the index before it, and so on.

## 2.9. Symbols Known to the Microcode

A large number of symbols will be pre-defined when a Spice Lisp system is fired up. A few of these are so fundamental to the operation of the system that their addresses have to be assembled into the microcode. These symbols are listed here. All of these symbols are in static space, so they will not be moving around.

**NIL**                    5C000000<sub>16</sub> The value of NIL is always NIL; it is an error to alter it. NIL is unique among symbols in that you can take its CAR and CDR, yielding NIL in either case.

**T**                        5C00000C<sub>16</sub> The value of T is always T; it is an error to alter it.

### %SP-Internal-Apply

5C000018<sub>16</sub> The function stored in the definition cell of this symbol is called by the microcode whenever compiled code calls an interpreted function. See section 6.1.2 for details.

### %SP-Internal-Error

5C000024<sub>16</sub> The function stored in the definition cell of this symbol is called whenever an error is detected during the execution of a byte instruction. See section 6.4 for details.

### %SP-Software-Interrupt-Handler

5C000030<sub>16</sub> The function stored in the definition cell of this symbol is called whenever a software interrupt occurs. See section 6.6 for details.

### %SP-Internal-Throw-Tag

5C00003C<sub>16</sub> This symbol is bound to the tag being thrown when a Catch-All frame is encountered on the stack. See section 6.2 for details.



## 3. Runtime Environment

### 3.1. Control Registers

To describe the instructions in chapter 5 and the complicated control conventions in chapter 6 requires that we talk about a number of "machine registers." All of these registers will be treated as if they contain 32-bit Lisp objects.

#### Control-Stack-Pointer

The stack pointer for the control stack, an object of type Misc-Control-Stack-Pointer. Points to the first unused word in Control-Stack space; this upward growing stack uses a write-increment/decrement-read discipline.

#### TOS

The top entry of the control stack, which is kept in a register for efficiency. References to local variables are faster if they can assume that the local in question is on the stack in main memory and that it has not popped up into the TOS register. To ensure this, the compiler adds an extra local variable to each function, so that none of the locals that are actually used can ever pop into TOS.

#### Binding-Stack-Pointer

The stack pointer for the special variable binding stack, an object of type Misc-Binding-Stack-Pointer. The binding stack follows the same write-increment/decrement-read discipline as the control stack.

#### Active-Frame

An object of type Misc-Control-Stack-Pointer which points to the first word of the call frame for the currently executing function. The virtual address of the start of the arguments and locals area of the active frame is this pointer plus a constant (see section 3.3).

#### Open-Frame

An object of type Misc-Control-Stack-Pointer which points to the first word of the call frame currently being built (i.e. whose arguments are being evaluated).

#### Active-Catch

An object of type Misc-Control-Stack-Pointer which points to the first word of the most recent catch frame built.

#### Active-Function

The compiled function object for the function that is currently being executed. The virtual address of the start of the symbols and constants area of the current function is this pointer plus a constant (see section 3.2).

#### Active-Code

The I-Vector of instructions for the currently executing function.

#### PC

A pointer into I-Vector space that indicates the next quadword from which the instruction buffer will be filled. This and the hardware BPC determine the next instruction to be executed. When a PC is pushed on the stack by a Call or Catch instruction, it is stored in the form of a 16-bit offset from the base of the Active-Code vector and the BPC:

31	27 26	20 19	16 15	0
-----				
Trap type code (5)   Unused (7)   BPC (4)   Offset (16)				
-----				

### 3.2. Function Object Format

Each compiled function is represented in the machine as a Function Object. This is identical in form to a G-Vector of lisp objects, and is treated as such by the garbage collector, but it exists in a special function space. (There is no particular reason for this distinction. We may decide later to store these things in G-Vector space, if we become short on spaces or have some reason to believe that this would improve paging behavior.) Usually, the function objects and code vectors will be kept in read-only space, but nothing should depend on this: some applications may create, compile, and destroy functions often enough to make dynamic allocation of function objects worthwhile.

The function object contains a vector of header information needed by the function-calling mechanism: a pointer to the I-Vector that holds the actual code. Following this is the so-called "symbols and constants" area. The first few entries in this area are fixnums that give the offsets into the code vector for various numbers of supplied arguments. Following this begin the true symbols and constants used by the function. Any symbol used by the code as a special variable or the name of another function will appear here. Fixnum constants in the range of -256 to +255 can be generated within the byte code, and so do not need to be stored in the constants area as full-word constants.

After the one-word G-Vector header, the entries of the function object are as follows:

- 0 A fixnum with bit fields as follows:
  - 0 - 14: Number of symbols/constants in this fn object (0 to 32K-1).  
This number includes the optional-arg offsets.
  - 15 - 26: Not used.
  - 27: 0 => All args eval'd. 1 => This is a FEXPR.
- 1 Pointer to the unboxed Code vector holding the macro-instructions.
- 2 A fixnum with bit fields as follows:
  - 0 - 7: The minimum legal number of args (0 to 255).
  - 8 - 15: The maximum number of args, not counting &rest (0 to 255).
  - 16 - 26: Number of local variables allocated on stack (0 to 2047).
  - 27: 0 => No &rest arg. 1 => One &rest arg.
- 3 Name of this function (a symbol).
- 4 Vector of argument names, in order, for debugging use.
- 5 The symbols and constants area starts here.  
This word is entry 0 of the symbol/constant area.  
The first few entries in this area are fixnums representing  
the code-vector entry points for various numbers of  
optional arguments. See section 6.1.2.

### 3.3. Control-Stack Format

The Spice Lisp control stack is a framed stack. Call frames, which hold information for function calls, are intermixed with catch frames, which hold information used for non-local exits. In addition, the control stack is used as a scratchpad for random computations.

#### 3.3.1. Call Frames

At any given time, the machine contains pointers to the current top of the control stack, the start of the current active frame (in which the current function is executing), and the start of the most recent open frame. In addition, there is a pointer to the current top of the special binding stack. An open frame is one which has been partially built, but which is still having arguments for it computed. When all the arguments have been computed and saved on the frame, the function is then started. This means that the call frame is completed, becomes the current active frame, and the function is executed. At this time, special variables may be bound and the old values are saved on the binding stack. Upon return, the active frame is popped away and the result is either sent as an argument to some previously opened frame or goes to some other destination. The binding stack is popped and old values are restored.

The active frame contains pointers to the previously-active frame, to the most recent open frame, and to the point to which the binding stack will be popped on exit, among other things. Following this is a vector of storage locations for the function's arguments and local variables. Space is allocated for the maximum number of arguments that the function can take, regardless of how many are actually supplied.

In an open frame, the structure is built up to the point where the arguments are stored. Thus, as arguments are computed, they can simply be pushed on the stack. When the function is finally started, the remainder of the frame is built. A call frame looks like this:

- 0 Header word. Type Call-Frame-Header.
- 1 Function object or EXPR for this call.
- 2 Pointer to previous active frame. Type Misc-Control-Stack-Ptr.
- 3 Pointer to previous open frame. Type Misc-Control-Stack-Ptr.
- 4 Pointer to previous binding stack. Type Misc-Binding-Stack-Ptr.
- 5 Saved PC of caller. A fixnum.
- 6 Args-and-locals area starts here. This is entry 0.

The first slot is pointed to by the Active-Frame register if this frame is currently active, and by the Open-Frame register if this frame is the currently opened frame.

### 3.3.2. Catch Frames

Catch frames contain much of the same information that call frames do, and have a very similar format. A catch frame holds the function object for the current function, a stack pointer to the current active and open frames, a pointer to the current top of the binding stack, and a pointer to the previous catch frame. When a Throw occurs, an operation equivalent to returning from this catch frame (as if it were a call frame) is performed, and the stacks are unwound to the proper place for continued execution in the current function. A catch frame looks like this:

```
0  Header word.  Type Catch-Frame-Header.
1  Function object for this call.
2  Pointer to current active frame.
3  Pointer to current open frame.
4  Pointer to current binding stack.
5  Destination PC for a Throw.
6  Tag caught by this catch frame.
7  Pointer to previous catch frame.
```

The conventions used to manipulate call and catch frames are described in chapter 6.

### 3.4. Binding-Stack Format

Each entry of the binding-stack consists of two boxed (32-bit) words. Pushed first is a pointer to the symbol being bound. Pushed second is the symbol's old value (any boxed item) that is to be restored when the binding is popped.

## 4. Storage Management

New objects are allocated from the lowest unused addresses within the specified space. Each allocation call specifies how many words are wanted, and a Free-Storage pointer is incremented by that amount. There is one of these Free-Storage pointers for each space, and it points to the lowest free address in the space. There is also a Clean-Space pointer associated with each space that is used during garbage collection. These pointers are stored in a table in the microcode table area which is indexed by type and space code. The address of the Free-Storage pointer for a given space is

$$(+ \text{ alloc-table-base } (1\text{sh type } 4) (1\text{sh space } 2)).$$

The address of the Clean-Space pointer is

$$(+ \text{ alloc-table-base } (1\text{sh type } 4) (1\text{sh space } 2) 2).$$

PERQ Spice Lisp uses a stop-and-copy garbage collector to reclaim storage. The Collect-Garbage instruction performs a full GC. The algorithm used is a degenerate form of Baker's incremental garbage collection scheme. When the Collect-Garbage instruction is executed, the following happens:

1. The current newspace becomes oldspace, and the current oldspace becomes newspace.
2. The newspace Free-Storage and Clean-Space pointers are initialized to point to the beginning of their spaces.
3. The contents of the "registers inside the barrier" are transported. There are only three such registers: Active-Function, Active-Code, and TOS. However, the PC is stored internally as an absolute address, so it must be recomputed if the code vector in Active-Code is transported. This is easily done by subtracting Active-Code from PC before it is transported, and adding it back in afterwards. Because the Active-Code vector will be transported from a quadword boundary to a quadword boundary, the PERQ's internal BPC needn't be modified.
4. The control stack and binding stack are scavenged.
5. Each static pointer space is scavenged.
6. Each new dynamic space is scavenged. The scavenging of the dynamic spaces is done until an entire pass through all of them does not result in anything being transported. At this point, every live object is in newspace.

A Lisp-level GC function must return the oldspace pages to Accent.

### 4.1. The Transporter

The transporter moves objects from oldspace to newspace. It is given an address  $A$ , which contains the object to be transported,  $B$ . If  $B$  is an immediate object, a pointer into static space, a pointer into read-only space, or a pointer into newspace, the transporter does nothing.

If  $B$  is a pointer into oldspace, the object it points to must be moved. It may, however, already have been moved. Fetch the first word of  $B$ , and call it  $C$ . If  $C$  is a GC-forwarding pointer, we form a new pointer with the type code of  $B$  and the low 27 bits of  $C$ . Write this into  $A$ .

If  $C$  is not a GC-forwarding pointer, we must copy the object the  $B$  points to. Allocate a new object of the same size in newspace, and copy the contents. Replace  $C$  with a GC-forwarding pointer to the new structure, and write the address of the new structure back into  $A$ .

Hash tables maintained with an EQ relation need special treatment by the transporter. Whenever a G-Vector with subtype 2 or 3 is transported to newspace, its subtype code is changed to 4. The Lisp-level hash-table functions will see that the subtype code has changed, and re-hash the entries before any access is made.

## 4.2. The Scavenger

The scavenger looks through an area of pointers for pointers into oldspace, transporting the objects they point to into newspace. The stacks and static spaces need to be scavenged once, but the new dynamic spaces need to be scavenged repeatedly, since new objects will be allocated while garbage collection is in progress. To keep track of how much a dynamic space has been scavenged, a Clean-Space pointer is maintained. The Clean-Space pointer points to the next word to be scavenged. Each call to the scavenger scavenges the area between the Clean-Space pointer and the Free-Storage pointer. The Clean-Space pointer is then set to the Free-Storage pointer. When all Clean-Space pointers are equal to their Free-Storage pointers, GC is complete.

To maintain (and create) locality of list structures, list space is treated specially. Two separate Clean-Space pointers are maintained for list space: one for cars and one for cdrs. The scavenger works on the Clean-Cdr pointer unless it is at the Free-Storage pointer, in which case it works on the Clean-Car pointer. When Clean-Car, Clean-Cdr, and the Free-Storage pointer for list space coincide, list space has been completely scavenged.

## 4.3. Purification

Garbage is created when the files that make up a Spice Lisp system are loaded. Many functions are needed only for initialization and bootstrapping (e.g. the "one-shot" functions produced by the compiler for random forms between function definition), and these can be thrown away once a full system is built. Most of the functions in the system, however, will be used after initialization. Rather than bend over backwards to make the compiler dump some functions in read-only space and others in dynamic space (which involves dumping

their constants in the proper spaces, also), we will dump *everything* into dynamic space, and use the following storage allocation feature to move what we need after initialization into read-only and static space.

The Set-Newspace-For-Type instruction lets us set the free pointer of the next newspace to dynamic or read-only space instead of the corresponding dynamic space. When the next GC happens, objects in newspace will be transported to this other space (static or read-only) instead of dynamic space. Care must be taken, of course, to ensure that objects in read-only space point only to static or read-only space. Probably the following should be used for "purifying" a system:

```
(set-newspace-for-type 1 2) ; bit-vectors to static
(set-newspace-for-type 2 2) ; likewise for i-vectors
(set-newspace-for-type 3 2) ; and strings
(set-newspace-for-type 4 2) ; and bignums
(set-newspace-for-type 5 2) ; and long-floats
(set-newspace-for-type 6 3) ; complexes can be read-only
(set-newspace-for-type 7 3) ; as can ratios
(set-newspace-for-type 8 2) ; g-vectors should be static
(set-newspace-for-type 9 3) ; functions should be read-only
(set-newspace-for-type 10 3) ; arrays can be, also
(set-newspace-for-type 11 2) ; symbols should be static
(set-newspace-for-type 12 2) ; as should conses.
```

## 5. Macro Instruction Set

The intent is that this instruction set should be a very direct mapping from the S-expression source it is derived from. There should therefore never be any temptation for users to write macrocode by hand; all of the system that is not in microcode should be written in Lisp. Since the compiler will run both in Spice Lisp and in MacLisp, we need not hand-code things even for bootstrapping.

### 5.1. Macro-Instruction Formats

There are three ways in which instructions fetch their arguments and store their results.

1. Most instructions pop all of their operands off of the stack and push a result back onto the stack, behaving like little Lisp functions. There are some instructions that will take their last operand from a place other than the stack (an immediate constant, a local variable, etc).
2. Some instructions modify a value in place. This value is sometimes the top of the stack, but could be a local variable, argument, or special variable. In the descriptions of the instructions below, these instructions operate on a pseudo-operand *E*, the effective address, which is specified as part of the opcode.
3. Finally, a few instructions pop the top of the stack but leave no result. The Pop, Branch, and Dispatch instructions do this.

All non-branching Spice Lisp instructions are made up of 1 or 2 opcode bytes, that contain an implicit effective address, and 0 to 2 bytes following the opcode that are used as part of the effective address. Branching instructions have 1 or 2 bytes of opcode followed by a 1 or 2 byte branch offset. The possible effective addresses (and their use of additional effective address bytes) are these:

**Stack**           The operand is taken from the stack. Then the operation takes place, in some cases pushing a result onto the stack. No effective address bytes are fetched. The names of instructions that take all stack operands are not suffixed with an effective address specifier, as others are. These instructions are called "basic" instructions. In most cases, the compiler-writer need concern himself with only these forms of an instruction. The peephole optimizer will replace sequences of stack referencing instructions with instructions with different addressing modes if the resulting sequence is faster.

#### Positive Short Integer Constant

A byte is fetched and is converted to a positive fixnum in the range 0 to 255. This is used as the operand. The "-PSIC" suffix on an instruction name is used for instructions with positive short integer constant operands. Some instructions imply a particular short integer without a second byte. These are suffixed with "-PSIC $n$ " where  $n$  is the short integer. A short integer constant may never be used as a result effective address, of course.

#### Negative Short Integer Constant

A byte is fetched and is converted to a negative fixnum in the range -1 to -256. This is used as the operand. The "-NSIC" suffix on an instruction name is used for instructions with



negative short integer constant operands.

### Arguments & Locals

In most cases, one byte is fetched and used as an unsigned offset (0 - 255) into the arguments and local variables area of the currently active call frame ("AI" suffix). The contents of this cell are used as the operand. For a few instructions, two bytes are fetched to form a 16-bit offset (0 - 65535). In fetching this double offset, the low-order byte comes in first ("LongAI" suffix). Some instructions imply a particular offset without the need for another offset byte. These instructions are those that are suffixed with "AI.n" where *n* is an integer which denotes the implied offset. When used as a result effective address, the result is stored in the specified slot of the call frame.

### Constants

In most cases, one byte is fetched and used as an unsigned offset (0 - 255) into the vector of symbols and constants in the code object of the current function. The constant in this cell is used directly ("C" suffix). For a few instructions, the next two bytes are fetched to form a 16-bit unsigned offset (0 - 65535) ("LongC" suffix). In fetching this double offset, the low-order byte comes in first. Sometimes an instruction implies an offset into the symbols and constants vector without the need of another byte for the offset. In these instances when the offset is implied, the instruction will have the suffix "Cn" where *n* is an integer denoting the offset. Constants may not be used as a result effective address.

### Symbols

In most cases, one byte is fetched and used as an unsigned offset (0 - 255) into the vector of symbols and constants in the code object of the current function. The constant in this cell is supposed to be a symbol pointer, and the operand is fetched from its value cell ("S" suffix). If the value is Misc-Trap, an unbound variable error is signalled. For some instructions, the next two bytes are fetched to form a 16-bit offset ("LongS" suffix). In fetching this double offset, the low-order byte comes in first. Sometimes an instruction implies an offset into the symbols and constants vector without the need of another byte for the offset. In these instances when the offset is implied, the instruction will have the suffix "Sn" where *n* is an integer denoting the offset. When a symbol is used as a result effective address, its value cell is set to the result.

### Ignore

Specified with a "-Ignore" suffix. This may be used only as a result effective address.

In the following listing, the effective address is called "*E*" and its contents are called "*CE*". The opcodes for these instructions are defined in a file read by the microassembler, compiler, error system, and disassembler. This file lives on CMU-CS-C as PRVA:<Slisp.Compiler.New-And-Improved>Instrdefs.Slisp and CMU-Badger as >Slisp>Instrdefs.Lisp. It is included in this document as an appendix.

## 5.2. Instructions

There are 11 classes of instructions: allocation, stack manipulation, list manipulation, symbol manipulation, array manipulation, type predicate, arithmetic and logical, branching and dispatching, function call and return, miscellaneous, and system hacking. A number of the instructions below combine the effect of two or more simpler instructions, and are part of the instruction set for efficiency reasons. It is envisioned that the

compiler will generate code using the stack forms of most instructions, with lots of Push and Pop instructions to get operands and store results. An optimizing assembler will then collapse sequences of these simple instructions into the faster, more compact complex instructions. Each basic instruction is flagged with an asterisk (\*).

### 5.2.1. Allocation

All non-immediate objects are allocated in the "current allocation space," which is dynamic space, static space, or read-only space. The current allocation space is initially dynamic space, but can be changed by using the Set-Allocation-Space instruction below. The current allocation space can be determined by using the Get-Allocation-Space instruction. One usually wants to change the allocation space around some section of code; an unwind protect should be used to insure that the allocation space is restored to some safe value.

Get-Allocation-Space () pushes 0, 2, or 3 if the current allocation space is dynamic, static, or read-only respectively.

Get-Allocation-Space\*

Set-Allocation-Space (*X*) sets the current allocation space to dynamic, static, or read-only if *X* is 0, 2, or 3 respectively. Pushes *X*.

Set-Allocation-Space\*

Alloc-Bit-Vector (*Length*) pushes a new bit-vector *Length* bits long, which is allocated in the current allocation space. *Length* must be a positive fixnum.

Alloc-Bit-Vector\*

Alloc-I-Vector (*Length* *A*) pushes a new I-Vector *Length* bytes long, with the access code specified by *A*. *Length* and *A* must be positive fixnums.

Alloc-I-Vector\*

Alloc-String (*Length*) pushes a new string *Length* characters long. *Length* must be a fixnum.

Alloc-String\*

Alloc-Bignum (*Length*) pushes a new bignum *Length* 8-bit bytes long. *Length* must be a fixnum.

Alloc-Bignum\*

Make-Complex (*Realpart* *Imagpart*) pushes a new complex number with the specified *Realpart* and *Imagpart*. *Realpart* and *Imagpart* should be the same type of non-complex number.

Make-Complex\*

Make-Ratio (*Numerator Denominator*) pushes a new ratio with the specified *Numerator* and *Denominator*. *Numerator* and *Denominator* should be integers.

Make-Ratio\*

Alloc-G-Vector (*Length Initial-Element*) pushes a new G-Vector with *Length* elements initialized to *Initial-Element*. *Length* should be a fixnum.

Alloc-G-Vector\*

Vector (*Elt<sub>0</sub> Elt<sub>1</sub> ... Elt<sub>Length-1</sub> Length*) pushes a new G-Vector containing the specified *Length* elements. *Length* should be a fixnum.

Vector\*

Vector-PSIC

Alloc-Function (*Length*) pushes a new function with *Length* elements. *Length* should be a fixnum.

Alloc-Function\*

Alloc-Array (*Length*) pushes a new array with *Length* elements. *Length* should be a fixnum.

Alloc-Array\*

Alloc-Symbol (*Print-Name*) pushes a new symbol with the print-name as *Print-Name*. The value is initially Misc-Trap, the definition is Misc-Trap, the property list and the package are initially NIL. The symbol is not interned by this operation -- that is done in macrocode. *Print-Name* should be a simple-string.

Alloc-Symbol\*

Cons (*Car Cdr*) pushes a new cons with the specified *Car* and *Cdr*.

Cons\*

Set-LPush (*Car E*) pushes a new cons with the specified *Car* and *CE* as the cdr, and stores the cons back into *E*.

Set-LPush-AL

Set-LPush-S

List (*Elt<sub>0</sub> Elt<sub>1</sub> ... Elt<sub>CE-1</sub> Length*) pushes a new list containing the *Length* elements. *Length* should be fixnum.

List\*

List-PSIC

List\* ( $Elt_0\ Elt_1\ \dots\ Elt_{CF-1}\ Length$ ) pushes a list\* formed by the  $CE$  elements onto the stack.  $Length$  should be a fixnum.

List\*  
List\*-PSIC

### 5.2.2. Stack Manipulation

Push ( $E$ ) pushes  $CE$  onto the stack.

Push-PSIC\*  
Push-PSIC0  
Push-PSIC1  
Push-PSIC2  
Push-PSIC3  
Push-NSIC\*  
Push-AL\*  
Push-AL0  
Push-AL1  
Push-AL2  
Push-AL3  
Push-AL4  
Push-AL5  
Push-AL6  
Push-AL7  
Push-LongAL\*  
Push-C\*  
Push-LongC\*  
Push-S\*  
Push-LongS\*

Pop ( $E$ ) pops the stack into  $E$ .

Pop-AL\*  
Pop-AL0  
Pop-AL1  
Pop-AL2  
Pop-AL3  
Pop-AL4  
Pop-AL5  
Pop-AL6  
Pop-AL7  
Pop-LongAL\*  
Pop-S\*  
Pop-LongS\*  
Pop-Ignore\*

Exchange () exchanges the top two elements of the stack.

Exchange\*

Copy (*E*) copies the top of stack to *E*.

Copy\*  
Copy-AL

NPop (*N*). If *N* is positive, *N* items are popped off of the stack. If *N* is negative, NIL is pushed onto the stack  $-N$  times. *N* must be a fixnum.

NPop-Stack\*  
NPop-PSIC  
NPop-NSIC

Bind-Null (*E*) pushes *CE* (which must be a symbol) and its current value onto the binding stack, and sets the value cell of *CE* to NIL.

Bind-Null\*  
Bind-Null-C

Bind (*Value Symbol*) pushes *Symbol* (which must be a symbol) and its current value onto the binding stack, and sets the value cell of *Symbol* to *Value*.

Bind\*  
Bind-C

Unbind (*N*) undoes the top *N* bindings on the binding stack.

Unbind\*  
Unbind-PSIC

### 5.2.3. List Manipulation

Cxxr (*L*). The cxxr of *L* is pushed onto the stack. *L* should be a list or NIL.

Car\*  
Car-AL  
Cdr\*  
Cdr-AL  
Cadr\*  
Cadr-AL  
Cddr\*  
Cddr-AL  
Cdar\*  
Cdar-AL  
Caar\*

Caar-AL

Set-Cxxr (*E*). The cxxr of *CE* is stored in *E*. *CE* should be either a list or NIL.

Set-Cdr-AL  
Set-Cdr-S  
Set-Cddr-AL  
Set-Cddr-S

Set-l.pop (*E*). The car of *CE* is pushed onto the stack; the cdr of *CE* is stored in *E*. *CE* should be a list or NIL.

Set-l.pop-AL  
Set-l.pop-S

Spread (*E*) pushes the elements of the list *CE* onto the stack in left-to-right order.

Spread\*  
Spread-AL

Replace-Cxr (*List Value*) sets the cxr of the *List* to *Value* and pushes *Value* on the stack.

Replace-Car\*  
Replace-Car-AL  
Replace-Cdr\*  
Replace-Cdr-AL

Endp (*X*) pushes T if *X* is NIL, or NIL if *X* is a cons. Otherwise an error is signalled.

Endp\*  
Endp-AL

Assoc (*List Item*) pushes the first cons in the association-list *List* whose car is EQL to *Item*. If the = part of the EQL comparison bugs out (and it can if the numbers are too complicated), a Lisp-level Assoc function is called with the current cdr of the *List*. Assq pushes the first cons in the association-list *List* whose car is EQ to *Item*.

Assoc\*  
Assq\*

Member (*List Item*) pushes the first cons in the list *List* whose car is EQL to *Item*. If the = part of the EQL comparison bugs out, a Lisp-level Member function is called with the current cdr of the *List*. Memq pushes the first cons in *List* whose car is EQ to the *Item*.

Member\*

Memq\*

GetF (*List Indicator Default*) searches for the *Indicator* in the list *List*, cddring down as the Common Lisp form GetF would. If *Indicator* is found, its associated value is pushed, otherwise *Default* is pushed.

GetF\*

#### 5.2.4. Symbol Manipulation

Get-Value (*Symbol*) pushes the value of *Symbol* (which must be a symbol) onto the stack. An error is signalled if *CE* is unbound.

Get-Value\*

Set-Value (*Symbol Value*) sets the value cell of the symbol *Symbol* to *Value*. *Value* is left on the top of the stack.

Set-Value\*

Get-Definition (*Symbol*) pushes the definition of the symbol *Symbol* onto the stack. If *Symbol* is undefined, an error is signalled.

Get-Definition\*  
Get-Definition-C

Set-Definition (*Symbol Definition*) sets the definition of the symbol *Symbol* to *Definition*. *Definition* is left on the top of the stack.

Set-Definition\*  
Set-Definition-C

Get-Plist (*Symbol*) pushes the property list of the symbol *Symbol* onto the stack.

Get-Plist\*  
Get-Plist-C

Set-Plist (*Symbol Plist*) sets the property list of the symbol *Symbol* to *Plist*. *Plist* is left on the top of the stack.

Set-Plist\*  
Set-Plist-C

Get-Pname (*Symbol*) pushes the print name of the symbol *Symbol* onto the stack.

Get-Pname\*

Get-Package (*Symbol*) pushes the package cell of the symbol *Symbol* onto the stack.

Get-Package\*

Set-Package (*Symbol Package*) sets the package cell of the symbol *Symbol* to *Package*. *Package* is left on the top of the stack.

Set-Package\*

Boundp (*Symbol*) pushes T if the symbol *Symbol* is bound; NIL otherwise.

Boundp\*  
Boundp-C

FBoundp (*Symbol*) pushes T if the symbol *Symbol* is defined; NIL otherwise.

FBoundp\*  
FBoundp-C

### 5.2.5. Array Manipulation

Common Lisp arrays have many manifestations in Spice Lisp. The Spice Lisp data types Bit-Vector, Integer-Vector, String, General-Vector, and Array are used to implement the collection of data types the Common Lisp manual calls "arrays."

In the following instruction descriptions, "simple-array" means an array implemented in Spice Lisp as a Bit-Vector, I-Vector, String, or G-Vector. "Complex-array" means an array implemented as a Spice Lisp Array object. "Complex-bit-vector" means a bit-vector implemented as a Spice Lisp array; similar remarks apply for "complex-string" and so forth.

Vector-Length (*Vector*) pushes the length of the one-dimensional Common Lisp array *Vector*. G-Vector-Length, Simple-String-Length, and Simple-Bit-Vector-Length push the lengths of G-Vectors, Spice Lisp strings, and Spice Lisp Bit-Vectors respectively. *Vector* should be a vector of the appropriate type.

Vector-Length\*  
G-Vector-Length\*  
Simple-String-Length\*  
Simple-Bit-Vector-Length\*

Get-Vector-Subtype (*Vector*) pushes the subtype field of the vector *Vector* as an integer. *Vector* should be a vector of some sort.

Get-Vector-Subtype\*



Set-Vector-Subtype (*Vector A*) sets the subtype field of the vector *Vector* to *A*, which must be a fixnum.

Set-Vector-Subtype\*

Get-Vector-Access-Code (*Vector*) pushes the access code of the I-Vector (or Bit-Vector) *Vector* as an integer.

Get-Vector-Access-Code\*

Shrink-Vector (*Vector Length*) sets the length field and the number-of-entries field of the vector *Vector* to *Length*. If the vector contains Lisp objects, entries beyond the new end are set to Misc-Trap. Pushes the shortened vector. *Length* should be a fixnum. One cannot shrink array headers or function headers.

Shrink-Vector\*

Typed-Vref (*A Vector I*) pushes the *I*th element of the I-Vector *Vector* by indexing into it as if its access-code were *A*. *A* and *I* should be fixnums.

Typed-Vref\*

Typed-Vset (*A Vector I Value*) sets the *I*th element of the I-Vector *Vector* to *Value* indexing into *Vector* as if its access-code were *A*. *A*, *I*, and *Value* should be fixnums. *Value* is pushed onto the stack.

Typed-Vset\*

Header-Length (*Object*) pushes the number of Lisp objects in the header of the function or array *Object*. This is used to find the number of dimensions of an array or the number of constants in a function.

Header-Length\*

Header-Ref (*Object I*) pushes the *I*th element of the function or array header *Object*. *I* must be a fixnum.

Header-Ref\*

Header-Set (*Object I Value*) sets the *I*th element of the function or array header *Object* to *Value*, and pushes *Value*. *I* must be a fixnum.

Header-Set\*

The names of the instructions used to reference and set elements of arrays are somewhat based on the Common Lisp function names. The SVref, SBit, and SChar instructions perform the same operation as their Common Lisp namesakes -- referencing elements of simple-vectors, simple-bit-vectors, and simple-strings respectively. Arefl references any kind of one dimensional array. The names of setting functions are derived by replacing "ref" with "set", "char" with "charset", and "bit" with "bitset."

Aref1 (*Array I*) pushes the *I*th element of the one-dimensional array *Array*. SVref pushes an element of a G-Vector; SChar an element of a string; Sbit an element of a Bit-Vector. *I* should be a fixnum.

Aref1\*  
 Aref1-AL  
 SVref\*  
 SVref-PSIC  
 SVref-AL  
 SVref-PSIC0  
 SVref-PSIC1  
 SVref-PSIC2  
 SVref-PSIC3  
 SVref-PSIC4  
 SVref-PSIC5  
 SChar\*  
 SChar-AL  
 SBit\*

Aset1 (*Array I Value*) sets the *I*th element of the one-dimensional array *Array* to *Value*. SVset sets an element of a G-Vector; SCharset an element of a string; SBitset an element of a Bit-Vector. *I* should be a fixnum and *Value* is pushed on the stack.

Aset1\*  
 Aset1-AL  
 SVset\*  
 SVset-AL  
 SCharset\*  
 SCharset-AL  
 SBitset\*

SVset\* (*Array Value I*) sets the *I*th element of the G-Vector *Array* to *Value*. The operands to the instruction are arranged so that the index can be specified as part of the effective address. This could not be done in general, of course, because order of evaluation must be preserved, but for constant indices (as used in structure accesses) this problem does not come up.

SVset\*-PSIC  
 SVset\*-PSIC0  
 SVset\*-PSIC1  
 SVset\*-PSIC2  
 SVset\*-PSIC3  
 SVset\*-PSIC4  
 SVset\*-PSIC5

CAref2 (*Array I1 I2*) pushes the element (*I1*, *I2*) of the two-dimensional array *Array* onto the stack. *I1* and *I2* should be fixnums. CAref3 pushes the element (*I1*, *I2*, *I3*).

C\ref2  
C\ref3

C\set2 (*Array I1 I2 Value*) sets the element (*I1, I2*) of the two-dimensional array *Array* to *Value* and pushes *Value* on the stack. *I1* and *I2* should be fixnums. C\set3 sets the element (*I1, I2, I3*).

C\set2  
C\set3

Bit-Bash (*V1 V2 V3 Op*). *V1, V2,* and *V3* should be bit-vectors and *Op* should be a fixnum. The elements of the bit vector *V3* are filled with the result of *Op*'ing the corresponding elements of *V1* and *V2*. *Op* should be a Boole-style number (see the Boole instruction in section 5.2.7).

Bit-Bash\*

The rest of the instructions in this section implement special cases of sequence or string operations. Where an operand is referred to as a string, it may actually be an 8-bit I-Vector or system area pointer.

Byte-BIT (*Src-String Src-Start Dst-String Dst-Start Dst-End*) moves bytes from *Src-String* into *Dst-String* between *Dst-Start* (inclusive) and *Dst-End* (exclusive). *Dst-Start - Dst-End* bytes are moved. If the substrings specified overlap, "the right thing happens," i.e. all the characters are moved to the right place. This instruction corresponds to the Common Lisp function REPLACE when the sequences are simple-strings.

Byte-BIT\*

Find-Character (*String Start End Character*) searches *String* for the *Character* from *Start* to *End*. If the character is found, the corresponding index into *String* is returned, otherwise NIL is returned. This instruction corresponds to the Common Lisp function FIND when the sequence is a simple-string.

Find-Character\*

Find-Character-With-Attribute (*String Start End Table Mask*) The codes of the characters of *String* from *Start* to *End* are used as indices into the *Table*, which is an I-Vector of 8-bit bytes. When the number picked up from the table bitwise ANDed with *Mask* is non-zero, the current index into the *String* is returned.

Find-Character-With-Attribute\*

SXHash-Simple-String (*String Length*) Computes the hash code of the first *Length* characters of *String* and pushes it on the stack. This corresponds to the Common Lisp function SXHASH when the object is a simple-string. The *Length* operand can be Nil, in which case the length of the string is calculated in microcode.

SXHash-Simple-String\*

### 5.2.6. Type Predicates

Bit-Vector-P (*Object*) pushes T if *Object* is a Common Lisp bit-vector or NIL, if it is not.

Bit-Vector-P\*

Simple-Bit-Vector-P (*Object*) pushes T if *Object* is a Spice Lisp bit-vector or NIL, if it is not.

Simple-Bit-Vector-P\*

Simple-Integer-Vector-P (*Object*) pushes T if *Object* is a Spice Lisp I-Vector or NIL, if it is not.

Simple-Integer-Vector-P\*

StringP (*Object*) pushes T if *Object* is a Common Lisp string or NIL, if it is not.

StringP\*

Simple-String-P (*Object*) pushes T if *Object* is a Spice Lisp string or NIL, if it is not.

Simple-String-P\*

BignumP (*Object*) pushes T if *Object* is a bignum or NIL, if it is not.

BignumP\*

Long-Float-P (*Object*) pushes T if *Object* is a long-float or NIL, if it is not.

Long-Float-P\*

ComplexP (*Object*) pushes T if *Object* is a complex number or NIL, if it is not.

ComplexP\*

RatioP (*Object*) pushes T if *Object* is a ratio or NIL, if it is not.

RatioP\*

IntegerP (*Object*) pushes T if *Object* is a fixnum or bignum or NIL, if it is not.

IntegerP\*

RationalP (*Object*) pushes T if *Object* is a fixnum, bignum, or ratio.

RationalP\*

FloatP (*Object*) pushes T if *Object* is a short-float or long-float.

FloatP\*

NumberP (*Object*) pushes T if *Object* is a number or NIL, if it is not.  
NumberP\*

General-Vector-P (*Object*) pushes T if *Object* is a Common Lisp general vector or NIL, if it is not.  
General-Vector-P\*

Simple-Vector-P (*Object*) pushes T if *Object* is a Spice Lisp G-Vector or NIL, if it is not.  
Simple-Vector-P\*

Compiled-Function-P (*Object*) pushes T if *Object* is a compiled function or NIL, if it is not.  
Compiled-Function-P\*

ArrayP (*Object*) pushes T if *Object* is a Common Lisp array or NIL if it is not.  
ArrayP\*

VectorP (*Object*) pushes T if *Object* is a Common Lisp vector or NIL, if it is not.  
VectorP\*

Complex-Array-P (*Object*) pushes T if *Object* is a Spice Lisp array or NIL if it is not.  
Complex-Array-P\*

SymbolP (*Object*) pushes T if *Object* is a symbol or NIL, if it is not.  
SymbolP\*

ListP (*Object*) pushes T if *Object* is a cons or NIL, or NIL, if it is not.  
ListP\*

ConsP (*Object*) pushes T if *Object* is a cons or NIL if it is not.  
ConsP\*

FixnumP (*Object*) pushes T if *Object* is a fixnum or NIL, if it is not.  
FixnumP\*

Short-Float-P (*Object*) pushes T if *Object* is a short-float or NIL if it is not.  
Short-Float-P\*

CharacterP (*Object*) pushes T if *Object* is a character or NIL if it is not.

CharacterP\*

### 5.2.7. Arithmetic and Logic

Integer-Length (*Object*) pushes the integer-length (as defined in the Common Lisp manual) of the integer *Object* onto the stack.

Integer-Length\*  
Integer-Length-AL

Float-Short (*Object*) pushes a short-float corresponding to the number *Object*.

Float-Short\*

Float-Long (*Number*) pushes a long float formed by coercing *Number* to a long float. This corresponds to the Common Lisp function Float when given a long float as its second argument.

Float-Long\*

Reapart (*Number*) pushes the reapart of the *Number*.

Reapart\*

Imagpart (*Number*) pushes the imagpart of the *Number*.

Imagpart\*

Numerator (*Number*) pushes the numerator of the rational *Number*.

Numerator\*

Denominator (*Number*) pushes the denominator of the rational *Number*.

Denominator\*

Decode-Float (*Number*) performs the Common Lisp Decode-Float function, leaving 3 values and a Values-Marker on the stack.

Decode-Float\*

Scale-Float (*Number X*) performs the Common Lisp Scale-Float function, pushing the result on the stack.

Scale-Float\*

= (*X Y*) pushes T if *X* and *Y* are numerically equal, or NIL if they are not. If an integer is compared with a flonum, the integer is floated first; if a short flonum is compared with a long flonum, the short one is first extended. Flonums must be exactly identical (after conversion) for a non-null comparison. < and > are

similar.

```

=
=-AL
=-PSIC
<
<-AL
<-PSIC
>
>-AL
>-PSIC
    
```

Truncate ( $N X$ ) performs the Common Lisp TRUNCATE operation. There are 3 cases depending on  $X$ :

- If  $X$  is fixnum 1, push three items: a fixnum or bignum representing the integer part of  $N$  (rounded toward 0), then either 0 if  $N$  was already an integer or the fractional part of  $N$  represented as a flonum or ratio with the same type as  $N$ , then Values-Marker 2 to mimic a multiple return of two values.
- If  $X$  and  $N$  are both fixnums or bignums and  $X$  is not 1, divide  $N$  by  $X$ . Push three items: the integer quotient (a fixnum or bignum), the integer remainder, and Values-Marker 2.
- If either  $X$  or  $N$  is a flonum or ratio, push a fixnum or bignum quotient (the true quotient rounded toward 0), then a flonum or ratio remainder, then push Values-Marker 2. The type of the remainder is determined by the same type-coercion rules as for  $+$ . The value of the remainder is equal to  $N - X * Quotient$ .

If Truncate uses the escape-to-macro mechanism (see section 6.3), it builds a multiple-value frame header rather than an escape header.

```

Truncate
Truncate-AL
Truncate-PSIC
    
```

$+$  ( $N X$ ) pushes  $N + X$ .  $-$ ,  $*$ , and  $/$  are similar.

```

+
+-AL
+-PSIC
+-PSIC1
-
--AL
--PSIC
--PSIC1
*
*-AL
*-PSIC
/
/-AL
    
```

/-PSIC

1+ (*E*) stores  $CE + 1$  into *E*.

1+-AL

1- (*E*) stores  $CE - 1$  into *E*.

1--AL

Negate (*N*) pushes  $-N$ .

Negate\*  
Negate-AL

Abs (*N*) pushes  $|N|$ .

Abs\*  
Abs-AL

Logand (*N X*) pushes the bitwise and of the integers *N* and *X*. Logior and Logxor are analagous.

Logand\*  
Logior\*  
Logxor\*

Lognot (*N*) pushes the bitwise complement of *N*.

Lognot\*

Boole (*Op X Y*) performs the Common Lisp Boole operation *Op* on *X*, and *Y*. The Boole constants for Spice Lisp are these:

boole-clr	0
boole-set	1
boole-1	2
boole-2	3
boole-c1	4
boole-c2	5
boole-and	6
boole-ior	7
boole-xor	8
boole-eqv	9
boole-nand	10
boole-nor	11
boole-andc1	12
boole-andc2	13
boole-orc1	14
boole-orc2	15



Boole\*

Ash ( $N X$ ) performs the Common Lisp ASH operation on  $N$  and  $X$ .

Ash\*  
Ash-PSIC

L.db ( $S P N$ ). All args are integers;  $S$  and  $P$  are non-negative. Performs the Common Lisp LDB operation with  $S$  and  $P$  being the size and position of the byte specifier.

L.db\*

Mask-Field ( $S P N$ ) performs the Common Lisp Mask-Field operation with  $S$  and  $P$  being the size and position of the byte specifier.

Mask-Field\*

Dpb ( $V S P N$ ) performs the Common Lisp DPB operation with  $S$  and  $P$  being the size and position of the byte specifier.

Dpb\*

Deposit-Field ( $V S P N$ ) performs the Common Lisp Deposit-Field operation with  $S$  and  $P$  as the size and position of the byte specifier.

Deposit-Field\*

Lsh ( $N X$ ) pushes a fixnum that is  $N$  shifted left by  $X$  bits, with 0's shifted in on the right. If  $X$  is negative,  $N$  is shifted to the right with 0's coming in on the left. Both  $N$  and  $X$  should be fixnums.

Lsh\*  
Lsh-PSIC

Logldb ( $S P N$ ). All args are fixnums.  $S$  and  $P$  specify a "byte" or bit-field of any length within  $N$ . This is extracted and is pushed right-justified as a fixnum.  $S$  is the length of the field in bits;  $P$  is the number of bits from the right of  $N$  to the beginning of the specified field.  $P = 0$  means that the field starts at bit 0 of  $N$ , and so on. It is an error if the specified field is not entirely within the 28 bits of  $N$ .

Logldb\*

Logdpb ( $V S P N$ ). All args are fixnums. Pushes a number equal to  $N$ , but with the field specified by  $P$  and  $S$  replaced by the  $S$  low-order bits of  $V$ . It is an error if the field does not fit into the 28 bits of  $N$ .

Logdpb\*

### 5.2.8. Branching and Dispatching

Branch instructions add or subtract a 1 or 2 byte relative offset to the PC after the branch instruction and the offset bytes have been fetched. The opcode determines the direction of the branch and the number of offset bytes to be fetched.

Branch-Forward (*Offset*) adds the 1 byte *Offset* to the PC. Long-Branch-Forward adds a 2 byte *Offset*. Branch-Backward and Long-Branch-Backward subtract 1 or 2 byte *Offsets*.

Branch-Forward\*  
 Long-Branch-Forward\*  
 Branch-Backward\*  
 Long-Branch-Backward\*

Branch-Null (*Offset*) pops the top item off the stack and branches if it is NIL; Branch-Not-Null branches if it is not null.

Branch-Null-Forward\*  
 Long-Branch-Null-Forward\*  
 Branch-Not-Null-Forward\*  
 Long-Branch-Not-Null-Forward\*  
 Branch-Null-Backward\*  
 Long-Branch-Null-Backward\*  
 Branch-Not-Null-Backward\*  
 Long-Branch-Not-Null-Backward\*

Branch-Save-Not-Null (*Offset*) looks at the value in TOS. If it is Nil, the stack it is popped off the stack and we fall through. Otherwise the stack is left as is and we take the branch.

Branch-Save-Not-Null-Forward\*  
 Long-Branch-Save-Not-Null-Forward\*  
 Branch-Save-Not-Null-Backward\*  
 Long-Branch-Save-Not-Null-Backward\*

Dispatch (). The top of stack (TOS) is used as an index into a dispatch table located in the current code vector. The next byte in the instruction is a limit. If TOS is not a fixnum, a fixnum less than 0, or a fixnum greater than or equal to the limit, no branch happens and we fall through, continuing with the next instruction. If TOS is within the specified bounds, however, it is added to a 16-bit number formed by fetching the next 1 or 2 bytes from the instruction stream. This result is used as an index into the code vector, and a 16-bit word is fetched from that memory location. The offset into the current code vector is set to this word. The stack is popped whether or not we branch.

Dispatch\*  
 Long-Dispatch\*

### 5.2.9. Function Call and Return

Call (*F*). *F* must be some sort of executable function: a function object, a lambda-expression, or a symbol with one of these stored in its function cell. A call frame for this function is opened. This is explained in more detail in the next chapter.

Call\*  
Call-C  
Call-AL

Call-0 (*F*). *F* must be an executable function, as above, but is a function of 0 arguments. Thus, there is no need to collect arguments. The call frame is opened and activated in a single instruction.

Call-0\*  
Call-0-C  
Call-0-AL

Call-Multiple (*F*). Just like a Call instruction, but it sets bit 21 of the frame header word to indicate that multiple values will be accepted. See section 6.1.4.

Call-Multiple\*  
Call-Multiple-C  
Call-Multiple-AL

Start-Call () closes the currently open call frame, and initiates a function call. See section 6.1.2 for details.

Start-Call\*

Push-Last (*X*) pushes *X* as an argument, closes the currently open call frame, and initiates a function call. See section 6.1.2 for details.

Push-Last-AL  
Push-Last-C  
Push-Last-S  
Push-Last-PSIC

Return (*X*). Return from the current function call. After the current frame is popped off the stack, *X* is pushed as the result being returned. See section 6.1.3 for more details.

Return\*  
Return-C  
Return-AL

Escape-Return (*X*). If the current call frame has an escape frame header, this works like a normal return, but the value *X* is put in the destination indicated in the header rather than just being returned on the stack.

If the current frame is not an escape frame, just return the single value on the stack as a normal return would.

Escape-Return\*

Break-Return (). If the header of the current call frame indicates a break frame, do a Return, but push no return value on the stack. If the current frame is not an escape frame, return NIL.

Break-Return\*

Catch () builds a catch frame. The top of stack should hold the tag caught by this catch frame, and the next entry on the stack should be a saved-format PC (which will come from the constants vector of the function). See section 6.2 for details.

Catch\*

Catch-Multiple () builds a multiple-value catch frame. The top of stack should hold the tag caught by this catch frame, and the next entry on the stack should be a saved-format PC. See section 6.2 for details.

Catch-Multiple\*

Catch-All () builds a catch frame whose tag is the special Catch-All object. The top of stack should hold the saved-format PC, which is the address to branch to if this frame is thrown through. See section 6.2 for details.

Catch-All\*

Throw (X). X is the throw-tag, normally a symbol. The value to be returned, either single or multiple, is on the top of the stack. See section 6.2 for a description of how this instruction works.

Throw\*  
Throw-C

### 5.2.10. Miscellaneous

Eq (X Y) pushes T if X and Y are the same object, NIL otherwise.

Eq\*  
Eq-AL  
Eq-C

Eq1 (X Y) pushes T if X and Y are the same object or if X and Y are numbers of the same type with the same value, NIL otherwise.

Eq1\*  
Eq1-AL

Eq1-C

Set-Null (*E*) sets *CE* to NIL.

Set-Null\*  
Set-Null-AL

Set-T (*E*) sets *CE* to T.

Set-T\*  
Set-T-AL

Set-0 (*E*) sets *CE* to 0.

Set-0\*  
Set-0-AL

Make-Predicate (*X*) pushes NIL if *X* is NIL or T if it is not.

Make-Predicate\*  
Make-Predicate-AL

Not-Predicate (*X*) pushes T if *X* is NIL or NIL if it is not.

Not-Predicate\*  
Not-Predicate-AL

Values-To-N (*V*). *V* must be a Values-Marker. Returns the number of values indicated in the low 24 bits of *V* as a fixnum.

Values-To-N

N-To-Values (*N*). *N* is a fixnum. Returns a Values-Marker with the same low-order 24 bits as *N*.

N-To-Values

Force-Values (). If the top of the stack is a Values-Marker, do nothing; if not, push Values-Marker 1.

Force-Values

Flush-Values (). If the top of the stack is a Values-Marker, remove this marker; if not, do nothing.

Flush-Values

### 5.2.11. System Hacking

Get-Type (*Object*) pushes the five type bits of the *Object* as a fixnum.

Get-Type\*  
Get-Type-Al.

Get-Space (*Object*) pushes the two space bits of *Object* as a fixnum.

Get-Space\*

Make-Immediate-Type (*X A*) pushes an object whose type bits are the integer *A* and whose other bits come from the immediate object or pointer *X*. *A* should be an immediate type code.

Make-Immediate-Type\*

8bit-System-Ref (*X I*). If *X* is an I-Vector, pushes the *I*th byte of *X*, indexing into *X* as an 8-bit I-Vector. If *X* is a system area pointer, pushes the *I*th byte beyond *X* as a fixnum. *I* must be a fixnum.

8bit-System-Ref

8bit-System-Set (*X I V*). If *X* is an I-Vector, sets the *I*th element of *X* to *V*, indexing into *X* as an 8-bit I-Vector. If *X* is a system area pointer, sets the *I*th byte beyond *X* to *V*.

8bit-System-Set

16bit-System-Ref (*X I*). If *X* is an I-Vector, pushes the *I*th 16-bit word of *X*, indexing into *X* as a 16-bit I-Vector. If *X* is a system area pointer, pushes the *I*th word beyond *X* as a fixnum. *I* must be a fixnum.

16bit-System-Ref

16bit-System-Set (*X I V*). If *X* is an I-Vector, sets the *I*th element of *X* to *V*, indexing into *X* as a 16-bit I-Vector. If *X* is a system area pointer, sets the *I*th word beyond *X* to *V*.

16bit-System-Set

Collect-Garbage () causes a stop-and-copy GC to be performed.

Collect-Garbage

Newspace-Bit () pushes 0 if newspace is currently space 0 or 1 if it is 1.

Newspace-Bit

Set-Newspace-For-Type (*type space*) sets the next newspace free pointer for the type corresponding to the *type* number to the space corresponding to the *space* number. There is about one useful thing that you can do with this instruction; see section 4.3. There are a number of not-so-useful but very fun things that you can do

with this instruction that are not documented here.

#### Set-Newspace-For-Type

Kernel-Trap (*Argblock Code*) is for communication with the Accent Kernel. *Code* is the type of trap desired (a fixnum), and *Argblock* is an I-Vector containing assorted argument information. See section 6.5 for details.

#### Kernel-Trap

Halt () stops the execution of Lisp. If continued, T is pushed on the stack.

#### Halt

Arg-In-Frame (*N F*). *N* is a fixnum, *F* is a control stack pointer as returned by the Active-Call-Frame and Open-Call-Frame instructions. Pushes the item in slot *N* of the args-and-locals area of call frame *F*.

#### Arg-In-Frame

Active-Call-Frame () pushes a control-stack pointer to the start of the currently active call frame. This will be of type Misc-Control-Stack-Pointer.

#### Active-Call-Frame

Active-Catch-Frame () pushes the control-stack pointer to the start of the currently active catch frame. This is Nil if there is no active catch.

#### Active-Catch-Frame

Set-Call-Frame (*P*). *P* must be a control stack pointer. This becomes the current active call frame pointer.

#### Set-Call-Frame

Current-Open-Frame () pushes a control-stack pointer to the start of the currently open call frame. This will be of type Misc-Control-Stack-Pointer.

#### Current-Open-Frame

Set-Open-Frame (*P*). *P* must be a control stack pointer. This becomes the current open frame pointer.

#### Set-Open-Frame

Current-Stack-Pointer () pushes the Misc-Control-Stack-Pointer that points to the current top of the stack (before the result of this operation is pushed). Note: by definition, this points to the first unused word of the stack, not to the last thing pushed. The stack manipulation instructions make it appear as if the stack is all in contiguous virtual memory, despite the fact that the TOS register will be holding the top word of the stack.

### Current-Stack-Pointer

Current-Binding-Pointer () pushes a Misc-Binding-Stack-Pointer that points to the first word above the current top of the binding stack.

### Current-Binding-Pointer

Read-Control-Stack (*P*). *P* must be a control stack pointer. Pushes the Lisp object that resides at this location. If the addressed object is totally outside the current stack, this is an error.

### Read-Control-Stack

Write-Control-Stack (*P V*). *P* is a stack pointer, *V* is any Lisp object. Writes *V* into the location addressed. If the addressed cell is totally outside the current stack, this is an error. Obviously, this should only be used by carefully written and debugged system code, since you can destroy the world by using this instruction.

### Write-Control-Stack

Read-Binding-Stack (*B*). *B* must be a binding stack pointer. Reads and returns the Lisp object at this location. An error if the location specified is outside the current binding stack.

### Read-Binding-Stack

Write-Binding-Stack (*B V*). *B* must be a binding stack pointer. Writes *V* into the specified location. An error if the location specified is outside the current binding stack.

### Write-Binding-Stack



## 6. Control Conventions

### 6.1. Function Calls

#### 6.1.1. Starting a Function Call

The Call and Call-Multiple instructions open a call frame on the control stack, but do not transfer control to the called function. The arguments for the function are then evaluated and pushed on the stack, and the call is started by a Push-Last instruction. Call-Multiple sets bit 21, the multiple-values-accepted bit, of the call frame to indicate that it will accept multiple-values. Call-0 opens the call frame and does the equivalent of a Start-Call instruction (see below) to start the called function. All these instructions take the function to be called as *CE*.

If *CE* is a symbol, we fetch the contents of the symbol's definition cell. If it is a Misc-Trap or another symbol, we signal an error. Otherwise, we go on with this definition as the function. We do not allow chains of symbols defined as other symbols. If *CE* is a compiled function, we perform the following steps:

1. Note the current value of the Control-Stack-Pointer register.
2. Push a Call-Frame-Header on control stack (with bit 21 set if this is a Call-Multiple).
3. Push *CE* (the function being called).
4. Push the Active-Frame register.
5. Push the Open-Frame register.
6. Push Binding-Stack-Pointer.
7. Push Fixnum -1 or some other easy-to-generate value. This will later be filled with caller's PC.
8. Open-Frame  $\leftarrow$  = Stack frame pointer saved in step 1.

The open frame is now ready to have arguments pushed.

If *CE* is a list, it is probably a lambda-expression or interpreted lexical closure. The call proceeds as above, with the list stored in the function slot of the new frame. The arguments are pushed normally, and %SP-Internal-Apply will be called when the Push-Last is executed. %SP-Internal-Apply verifies that this function is a lambda or lexical closure.

If *CE* is anything else, an Illegal-Function error is signalled.

### 6.1.2. Finishing a Function Call

Push-Last pushes a final argument  $X$  and starts the function responsible for the current open frame. Start-Call just starts the function. Call-0 opens the frame and performs the equivalent of a Start-Call immediately, since there are no arguments to push.

We look at the function entry of the current open frame. If this contains a compiled function object, proceed as follows:

1. Insert the current PC (points to the NEXT instruction of the caller's code vector) in the PC slot of the open frame.
2. Active-Function  $\leftarrow$  = Called function (from slot 1 of open frame).
3. Active-Code  $\leftarrow$  = Code vector for new active function.
4. Active-Frame  $\leftarrow$  = Open-Frame
5. Note number of args pushed by caller. Let this be  $K$ . We must now compute the proper entry point in the called function's code vector as a function of  $K$  and the number of args the called function wants.
  - a. If number of args  $<$  minimum, signal an error.
  - b. If number of args  $>$  maximum and no &REST arg is allowed, signal an error.
  - c. If number of args  $>$  maximum and a &REST arg is present, pop excess args into a list, push this list back on stack as the &REST arg, and start at offset 0.
  - d. If number of args is between min and max (inclusive), get the starting offset from the appropriate slot of the called function's function object. This is stored as a fixnum in slot  $K - \text{MIN} + 6$  of the function object.
6. Set up the new PC to point at the right place in the code vector and return to the macro-code execution loop to run the new function. This involves setting up PC, the BPC, and refilling the instruction buffer.

If the object in the function entry is a list instead of a function object, we must call %SP-Internal-Apply to interpret the function with the given arguments. We proceed as follows:

1. Note the number of args pushed in the current open frame (call this  $N$ ) and the frame pointer for this frame (call it  $F$ ). Also remember the lambda-expression in this frame (call it  $L$ ).
2. Perform steps 1 - 4 of the sequence above for a normal Start-Call.
3. Perform the equivalent of a Call-Multiple instruction with the symbol %SP-Internal-Apply as  $CE$ . (This symbol is in a fixed location known to the microcode. See section 2.9.)

4. Push L, N, and F in that order as the three arguments to %SP-Internal-Apply.
5. Perform the equivalent of a Push-Last-Stack to start the call.

%SP-Internal-Apply, a function of three arguments, now evaluates the call to the lambda-expression or interpreted lexical closure L, obtaining the arguments from the frame pointed to by F. These arguments are obtained using the Arg-In-Frame instruction. Prior to returning %SP-Internal-Apply sets the Active-Frame register to F, so that it returns from frame F.

### 6.1.3. Returning from a Function Call

Return returns from the current function, popping the stack frame and pushing some number of returned values. If CE is a Values-Marker but bit 21 is not on in the current call frame, only one value is returned. If bit 21 is on, either multiple values or a single value will be returned. The steps are as follows:

1. Pop binding stack back to value saved in slot 5 of the active control frame. For each symbol/value pair popped off the binding stack, restore that value for the symbol.
2. Temp <= = Previous active frame from slot 3 of current frame.
3. Open-Frame <= = Saved value in current frame.
4. PC <= = Saved value in current frame. This requires setting up the internal PC, the BPC, and the instruction buffer.
5. Active-Function <= = Saved value from previous frame. A pointer to this frame is in Temp.
6. Active-Code <= = Code Vector obtained from entry in restored Active-Function object.
7. Pop current frame off stack:
  - Control-Stack-Pointer <= = Active-Frame.
  - Active-Frame <= = Temp.
  - Pop top of stack into TOS register. Since the active frame is inside the barrier, make sure the new top frame has been scavenged, or do it now.
8. Push the return value onto the stack.
9. Resume execution of function popped to.

### 6.1.4. Returning Multiple-Values

If bit 21 is on in the current frame and a Values-Marker indicating N values is on the top of the stack, we proceed as follows:

1. Note the value of the current stack pointer (after CE is popped off if it came from the stack) as OLDSP.
2. Perform steps 1 - 7 of the Return procedure described above.

3. Do a block transfer loop pushing the N words starting at (O.LDSP) - N onto the stack as return values. Then push the original *CV*, which is Values-Marker N.
4. Resume execution of the caller.

To do (MULTIPLE-VALUE-LIST (FOO A B)), we could use this sequence of instructions:

```
(CALL-MULTIPLE (CONSTANT [FOO]))
(PUSH [A])
(PUSH-LAST [B])
(FORCE-VALUES)
(VALUE-TO-N STACK)
(LIST STACK) ;Pop N from stack, then listify N things.
```

To do (MULTIPLE-VALUE-SETQ (X Y Z) (FOO A B)), we could use this code:

```
(CALL-MULTIPLE (CONSTANT [FOO]))
(PUSH [A])
(PUSH-LAST [B])
(FORCE-VALUES)
(VALUE-TO-N STACK)
(- (CONSTANT [3])) ;Get number offered - number wanted.
(NPOP STACK) ;Flush surplus returns or push NILs.
(POP [Z]) ;Now put the three values wherever they
(POP [Y]) ; are supposed to go.
(POP [X])
```

In tail recursive situations, such as in the last form of a PROGN, one function, FOO, may want to call another function, BAR, and return "whatever BAR returns." Call-Multiple is used in this case. If BAR returns multiple values, they will all be passed to FOO. If FOO's caller wants multiple values, the values will be returned. If not, FOO's Return instruction will see that there are multiple values on the stack, but that multiple values will not be accepted by FOO's caller. So Return will return only the first value.

## 6.2. Non-Local Exits

The Catch and Unwind-Protect special forms are implemented using catch frames. Unwind-Protect builds a catch frame whose tag is the Catch-All object. The Catch instruction creates a catch frame for a given tag and PC to branch to in the current instruction. The Throw instruction looks up the stack by following the chain of catch frames until it finds a frame with a matching tag or a frame with the Catch-All object as its tag. If it finds a frame with a matching tag, that frame is "returned from," and that function is resumed. If it finds a frame with the Catch-All object as its tag, that frame is "returned from," and in addition, %SP-Internal-Throw-Tag is set to the tag being searched for. So that interrupted cleanup forms behave correctly, %SP-Internal-Throw-Tag should be bound to the Catch-All object before the Catch-All frame is built. The protected forms are then executed, and if %SP-Internal-Throw-Tag is not the Catch-All object, its value is thrown to. Exactly what we do is this:

1. Put the contents of the Active-Catch register into a register, A. Put NIL into another register, B.
2. If A is NIL, the tag we seek isn't on the stack. Signal an Unseen-Throw-Tag error.
3. Look at the tag for the catch frame in register A. If it's the tag we're looking for, go to step 4. If it's the Catch-All object and B is NIL, copy A to B. Set A to the previous catch frame and go back to step 2.
4. If B is non-NIL, we need to execute some cleanup forms. Return into B's frame and bind %SP-Internal-Throw-Tag to the tag we're searching for. When the cleanup forms are finished executing, they'll throw to this tag again.
5. If B is NIL, return into this frame, pushing the return value (or Binding the multiple values if this frame has bit 21 set and there are multiple values).

If no form inside of a Catch results in a Throw, the catch frame needs to be removed from the stack before execution of the function containing the throw is resumed. For now, the value produced by the forms inside the Catch form are thrown to the tag. Some sort of specialized instruction could be used for this, but right now we'll just go with the throw. The branch PC specified by a Catch instruction is part of the constants area of the function object, much like the function's entry points. To do

```
(catch 'foo
  (baz)
  (bar))
```

we could use this code:

```
(PUSH (CONSTANT [PC-FOR-TAG-1]))
(PUSH (CONSTANT [FOO]))
(CATCH STACK)
(CALL-0 (CONSTANT [BAZ]))
(POP IGNORE)
(CALL-0 (CONSTANT [BAR]))
(PUSH (CONSTANT [FOO]))
(THROW STACK)
TAG-1
```

To do

```
(unwind-protect
  (baz)
  (bar))
```

we could use this code:

```

(PUSH (SYMBOL %CATCH-ALL-OBJECT))
(PUSH (CONSTANT %SP-INTERNAL-THROW-TAG))
(BIND STACK)
(PUSH (CONSTANT [PC-FOR-TAG-1]))
(CATCH-ALL STACK)
(CALL-0 (CONSTANT [BAZ]))
(PUSH (SYMBOL %CATCH-ALL-OBJECT))
(THROW STACK)
TAG-1
(CALL-0 (CONSTANT [BAR]))
(POP IGNORE)
(PUSH (SYMBOL %CATCH-ALL-OBJECT))
(EQ (SYMBOL %SP-INTERNAL-THROW-TAG))
(BRANCH-NOT-NULL TAG-2)
(PUSH (SYMBOL %SP-INTERNAL-THROW-TAG))
(THROW STACK)
TAG-2

```

### 6.3. Escaping to Macrocode

Some instructions can be complex (e.g. \* given a long-float and a bignum), and with limited microstore (and microprogrammer time) on the PERQ, we would like to handle these in Lisp code. Such cases could be handled by a full-scale microcode-to-macrocode subroutine call, which upon a return comes back to the designated return address in the microcode and restores any micro-state that may have been clobbered. This may ultimately be needed if we ever implement a micro-compiler for lisp, but for now we can get by with a simpler scheme. If the microcode for any macro-instruction decides that it has a case too difficult to handle, it can call a macrocoded function that does whatever the original macro-instruction was supposed to do. It does this by opening an escape-type frame on the control stack, pushing an appropriate set of arguments, and then starting the call as though a push-last had been done in macrocode.

When the macrocoded escape function returns (the Escape-Return instruction must be used for this return) the single returned value goes wherever the original macro-instruction was supposed to place its result, and the original instruction stream continues on as if the macrocode instruction had exited normally without an escape.

Instructions can place their return values in any of several destinations. The escape call must set up the frame header word to indicate which of these locations is to get the value returned by the macro-coded escape function. An appropriate effective-address code is stored in bits 16-17:

- |         |  |
|---------|--|
| 0 Stack | The result is pushed onto the stack.   |
| 1 AL    | The result is put into the arguments/locals area of the current call frame. Bits 0-15 contain a 16-bit offset. |

- 2 Symbol            The result is put into the value cell of a symbol in the symbols and constants area of the current function object. Bits 0-15 contain a 16-bit offset.
- 3 Ignore            The result is thrown away.

Given this information in the frame header, Escape-Return will do the right thing to make it appear that the original instruction had exited normally.

Some instructions, notably Truncate, may want to return multiple values from an escape function. These values will always be returned on the stack. In this case, the escape mechanism builds a multiple-value call frame rather than an escape call frame, then escapes in the usual way. The escape routine for Truncate is exited using a normal Return instruction.

A table of pointers to the Lisp-level escape functions is stored in a fixed location in virtual memory, and the address of the start of this table is known to the microcode. This means that microcode routines can select the desired function by means of a table index, and it is not necessary to assemble the addresses of all these functions into the microcode.

The escape mechanism is implemented by a micro-subroutine named ESCAPE, which can be called (or rather, jumped to, since ESCAPE never returns to the caller) by any microcode that wants to escape to macrocode. ESCAPE is passed the index of the macro-function to be called and from 0 to 4 lisp objects as arguments on the PERQ E-Stack. ESCAPE then performs the following steps:

1. It is determined where the currently executing instruction is going to place its result, and an appropriate escape-type call header word is generated.
2. A pointer to the desired function object is fetched from the table of escape functions, as determined by the index that was passed to ESCAPE.
3. The equivalent of a Call instruction is executed for this function object, but the header word determined in step 1 is used instead of the normal header word.
4. The specified arguments, if any, are pushed onto the control stack. The new function is then started by executing the equivalent of a Push-Last instruction.

A second entry point, ESCAPE-MULTIPLE, does the same thing as ESCAPE but creates a multiple-value frame header instead of an escape frame header.

## 6.4. Errors

When an error occurs during the execution of an instruction, a call to %SP-Internal-Error is performed. This call is a break-type call, so if the error is proceeded (with a Break-Return instruction), no value will be pushed on the stack.

%SP-Internal-Error is passed a fixnum error code as its first argument. The second argument is a fixnum offset into the current code vector that points to the location immediately following the instruction that encountered the trouble. From this offset, the Lisp-level error handler can reconstruct the PC of the losing instruction, which is not readily available in the micro-machine. Following the offset, there may be 0 - 2 additional arguments that provide information of possible use to the error handler. For example, an unbound-symbol error will pass the symbol in question as the third arg.

A Lisp-Level error handler may want to provide a result for the instruction. It can find the losing instruction in the way described above, and look at it's opcode to find the destination. The error handler could then store the user-supplied result in the specified place and proceed executing the errorful function at the instruction after the losing instruction.

The following error codes are currently defined. Unless otherwise specified, only the error code and the code-vector offset are passed as arguments.

The following table is pretty bogus. After the microcode is written, and I know what errors are really generated, I'll make a newer table.

1 Control Stack Overflow

The control stack has exceeded the allowable size, currently  $2^{24}$  words.

2 Control Stack Underflow

Can only result from a compiler bug or misuse of an instruction.

3 Binding Stack Overflow

The binding stack has exceeded the allowable size, currently  $2^{24}$  words.

4 Binding Stack Underflow

Can only result from a compiler bug or misuse of an instruction.

5 Virtual Memory Overflow

Some data space has exceeded the maximum size of its segment in virtual memory.

6 Unbound Symbol

Attempted access to the special value of an unbound symbol. Passes the symbol as the third argument to %Sp-Internal-Error.



## 7 Undefined Symbol

Attempted access to the definition cell of an undefined symbol. Passes the symbol as the third argument to %Sp-Internal-Error.

## 8 Unused.

## 9 Altering T or NIL.

Attempt to bind or setq the special value of T or NIL.

## 10 Unused.

## 11 Write Into Read-Only Space

Self-explanatory.

## 12 Object Not Character

The object is passed as the third argument.

## 13 Object Not System Area Pointer

The object is passed as the third argument.

## 14 Object Not Control Stack Pointer

The object is passed as the third argument.

## 15 Object Binding Stack Pointer

The object is passed as the third argument.

## 16 Object Not Values Marker

The object is passed as the third argument.

## 17 Object Not Fixnum

The object is passed as the third argument.

## 18 Object Not Vector-Like

The object is passed as the third argument.

## 19 Object Not Integer-Vector

The object is passed as the third argument.

## 20 Object Not Symbol

The object is passed as the third argument.

## 21 Object Not List

The object is passed as the third argument.

## 22 Object Not List or Nil

The object is passed as the third argument.

- 23 Object Not String  
The object is passed as the third argument.
- 24 Object Not Number  
The object is passed as the third argument.
- 25 Object Not Misc Type  
The object is passed as the third argument.
- 26 Unused.
- 27 Illegal Allocation Space Value  
Self explanatory.
- 28 Illegal Vector Size  
Attempt to allocate a vector with negative size or size too large for vectors of this type.  
Passes the requested size as the third argument.
- 29 Illegal Immediate Type Code  
Passes the code as the third argument.
- 30 Illegal Control Stack Pointer  
Passes the illegal pointer as the third argument.
- 31 Illegal Binding Stack Pointer  
Passes the illegal pointer as the third argument.
- 32 Illegal Instruction  
Must be due to a compiler error or to using obsolete code that does not match the current microcode. No additional args.
- 33 Unused.
- 34 Illegal Divisor  
The divisor is integer or floating 0. Returns the divisor and dividend as the third and fourth args.
- 35 Illegal Vector Access Type  
The specified access type is returned as the third argument.
- 36 Illegal Vector Index  
The specified index is out of bounds for this vector. The bad index is passed as the third argument.
- 37 Illegal Byte Pointer  
Bad S or P value to LDB or related function. Returns S and P as the third and fourth arguments.
- 38 Illegal Function

Bad object being called as a function. The object is passed as the third argument.

39 Too Few Arguments

Attempt to activate the call to a function with too few arguments on the stack. Returns the number of arguments passed as the third argument, the function being called as the fourth.

40 Too Many Arguments

Attempt to activate the call to a function with too few arguments on the stack. Returns the number of arguments passed as the third argument, the function being called as the fourth.

41 Unseen Throw Tag

Returns the tag as the third argument.

42 Null Open Frame

Attempt to activate a function call, but no frame has been opened. No additional args.

43 Undefined Type Code

Can only result from a bug in the micro-machine. Returns the strange object as the third argument.

44 Return From Initial Function

Self-explanatory.

45 GC Forward Not To Newspace

Can only result from internal errors in the micro-machine. No additional args.

46 Attempt To Transport GC Forward

Can only result from internal errors in the micro-machine. No additional args.

47 Object Not Integer

The object is passed as the third argument.

48 Short-float exponent overflow, underflow

No additional args.

49 Long-float exponent overflow, underflow

No additional args.

50 - 63 Unused.

In the Tops-20 virtual machine, the following codes are defined:

64 Illegal File Token

The bad token is passed as the third argument.

65 Illegal I/O Mode Specifier

The bad mode is passed as the third argument.

## 6.5. Trapping to the Accent Kernel

Most of the primitive calls to the Accent kernel are made through a single microcode entry point, SVCall, defined in Accent file process.mic. From Lisp level, these calls are generated by the Kernel-Trap instruction.

Kernel-Trap takes two operands, an argument block and a trap code, in that order. The trap code is a fixnum which specifies the sort of trap call desired. The argument block is an I-Vector which contains the argument information for the trap call. The size and format of the argument block depends on which trap code is called. The return codes and values from the trap are written into elements of the I-Vector by the kernel.

Internally, the trap code and a pointer to the data portion of the I-Vector are passed to Accent on the PERQ E-Stack, as follows:

ETOS	The trap code.
ETOS - 1	The low order 16 bits of the virtual address.
ETOS - 2	The high order 16 bits of the virtual address.

All of the kernel traps called by Lisp-level code use the virtual address as a pointer to an argument block. An argument block is stored at lisp level as an I-Vector of 16-bit quantities. The trap codes are defined in Accent file accenttype.pas, and the arguments to these calls are described in the *Accent Kernel Interface Manual*.

## 6.6. Interrupts

There are three kinds of asynchronous events that the Spice Lisp system must service: hardware interrupts, process breaks, and software interrupts.

Hardware interrupts must be serviced every 70 microinstructions. It is guaranteed that no process registers will be altered and no page faults will occur, so all a microprogrammer need do is check the Intr-Pending condition every now and then, and call the hardware interrupt service routine. Sometimes that routine will set the process break flag, and a process break should occur.

If there are other runnable processes on the machine, a process break will result in the de-scheduling of the Lisp process. Process registers will be saved by the kernel, and restored when the Lisp runs again. After a process break, all cached virtual-to-physical memory translations may be invalid and the instruction buffer will probably be filled with some other process's instructions. The caches must be flushed and the instruction

buffer must be refilled after a process break.

After a process break, it is possible that the Lisp process will have received an "emergency message" from some other process. If so, the software interrupt flag will be set. To service this software interrupt, a break-type call frame is built to %SP-Software-Interrupt-Handler, which should receive the message and figure out what to do with it at Lisp level. The emergency message might, for example, report that an interrupt character has been typed, and the interrupt handler could enter a break loop or throw to the Lisp top level.

## Appendix I Fasload File Format

### I.1. General

The purpose of Fasload files is to allow concise storage and rapid loading of Lisp data, particularly function definitions. The intent is that loading a Fasload file has the same effect as loading the ASCII file from which the Fasload file was compiled, but accomplishes the tasks more efficiently. One noticeable difference, of course, is that function definitions may be in compiled form rather than S-expression form. Another is that Fasload files may specify in what parts of memory the Lisp data should be allocated. For example, constant lists used by compiled code may be regarded as read-only.

In some Lisp implementations, Fasload file formats are designed to allow sharing of code parts of the file, possibly by direct mapping of pages of the file into the address space of a process. This technique produces great performance improvements in a paged time-sharing system. Since the Spice project is to produce a distributed personal-computer network system rather than a time-sharing system, efficiencies of this type are explicitly *not* a goal for the Spice Lisp Fasload file format.

On the other hand, Spice Lisp is intended to be portable, as it will eventually run on a variety of machines. Therefore an explicit goal is that Fasload files shall be transportable among various implementations, to permit efficient distribution of programs in compiled form. The representations of data objects in Fasload files shall be relatively independent of such considerations as word length, number of type bits, and so on. If two implementations interpret the same macrocode (compiled code format), then Fasload files should be completely compatible. If they do not, then files not containing compiled code (so-called "Fasdump" data files) should still be compatible. While this may lead to a format which is not maximally efficient for a particular implementation, the sacrifice of a small amount of performance is deemed a worthwhile price to pay to achieve portability.

The primary assumption about data format compatibility is that all implementations can support I/O on finite streams of eight-bit bytes. By "finite" we mean that a definite end-of-file point can be detected irrespective of the content of the data stream. A Fasload file will be regarded as such a byte stream.

## 1.2. Strategy

A Fasload file may be regarded as a human-readable prefix followed by code in a funny little language. When interpreted, this code will cause the construction of the encoded data structures. The virtual machine which interprets this code has a *stack* and a *table*, both initially empty. The table may be thought of as an expandable register file; it is used to remember quantities which are needed more than once. The elements of both the stack and the table are Lisp data objects. Operators of the funny language may take as operands following bytes of the data stream, or items popped from the stack. Results may be pushed back onto the stack or pushed onto the table. The table is an indexable stack that is never popped; it is indexed relative to the base, not the top, so that an item once pushed always has the same index.

More precisely, a Fasload file has the following macroscopic organization. It is a sequence of zero or more groups concatenated together. End-of-file must occur at the end of the last group. Each group begins with a series of seven-bit ASCII characters terminated by one or more bytes of all ones ( $FF_{16}$ ); this is called the *header*. Following the bytes which terminate the header is the *body*, a stream of bytes in the funny binary language. The body of necessity begins with a byte other than  $FF_{16}$ . The body is terminated by the operation FOP-END-GROUP.

The first nine characters of the header must be "FASL FILE" in upper-case letters. The rest may be any ASCII text, but by convention it is formatted in a certain way. The header is divided into lines, which are grouped into paragraphs. A paragraph begins with a line which does *not* begin with a space or tab character, and contains all lines up to, but not including, the next such line. The first word of a paragraph, defined to be all characters up to but not including the first space, tab, or end-of-line character, is the *name* of the paragraph. A Fasload file header might look something like this:

```
FASL FILE >SteelesPerq>User>Guy>IoHacks>Pretty-Print.Slisp
Package Pretty-Print
Compiled 31-Mar-1988 09:01:32 by some random luser
Compiler Version 1.6, Lisp Version 3.0.
Functions: INITIALIZE DRIVER HACK HACK1 MUNGE MUNGE1 GAZORCH
           MINGLE MUDDLE PERTURB OVERDRIVE GOBBLE-KEYBOARD FRY-USER
           DROP-DEAD HELP CLEAR-MICROCODE %AOS-TRIANGLE
           %HARASS-READTABLE-MAYBE
Macros:   PUSH POP FROB TWIDDLE
<one or more bytes of  $FF_{16}$ >
```

The particular paragraph names and contents shown here are only intended as suggestions.

### 1.3. Fasload Language

Each operation in the binary Fasload language is an eight-bit (one-byte) opcode. Each has a name beginning with "FOP-". In the following descriptions, the name is followed by operand descriptors. Each descriptor denotes operands that follow the opcode in the input stream. A quantity in parentheses indicates the number of bytes of data from the stream making up the operand. Operands which implicitly come from the stack are noted in the text. The notation " $\Rightarrow$  stack" means that the result is pushed onto the stack; " $\Rightarrow$  table" similarly means that the result is added to the table. A construction like " $n(1) \text{ value}(n)$ " means that first a single byte  $n$  is read from the input stream, and this byte specifies how many bytes to read as the operand named *value*. All numeric values are unsigned binary integers unless otherwise specified. Values described as "signed" are in two's-complement form unless otherwise specified. When an integer read from the stream occupies more than one byte, the first byte read is the least significant byte, and the last byte read is the most significant (and contains the sign bit as its high-order bit if the entire integer is signed).

Some of the operations are not necessary, but are rather special cases of or combinations of others. These are included to reduce the size of the file or to speed up important cases. As an example, nearly all strings are less than 256 bytes long, and so a special form of string operation might take a one-byte length rather than a four-byte length. As another example, some implementations may choose to store bits in an array in a left-to-right format within each word, rather than right-to-left. The Fasload file format may support both formats, with one being significantly more efficient than the other for a given implementation. The compiler for any implementation may generate the more efficient form for that implementation, and yet compatibility can be maintained by requiring all implementations to support both formats in Fasload files.

Measurements are to be made to determine which operation codes are worthwhile; little-used operations may be discarded and new ones added. After a point the definition will be "frozen", meaning that existing operations may not be deleted (though new ones may be added; some operations codes will be reserved for that purpose).

- 0 FOP-NOP      No operation. (This is included because it is recognized that some implementations may benefit from alignment of operands to some operations, for example to 32-bit boundaries. This operation can be used to pad the instruction stream to a desired boundary.)
- 1 FOP-POP     $\Rightarrow$  table  
One item is popped from the stack and added to the table.
- 2 FOP-PUSH   *index*(4)  $\Rightarrow$  stack  
Item number *index* of the table is pushed onto the stack. The first element of the table is item number zero.
- 3 FOP-BYTE-PUSH   *index*(1)  $\Rightarrow$  stack



Item number *index* of the table is pushed onto the stack. The first element of the table is item number zero.

4 FOP-EMPTY-LIST  $\Rightarrow$  stack

The empty list (( )) is pushed onto the stack.

5 FOP-TRUTH  $\Rightarrow$  stack

The standard truth value (T) is pushed onto the stack.

6 FOP-SYMBOL-SAVE  $n(4)$  *name(n)*  $\Rightarrow$  stack & table

The four-byte operand *n* specifies the length of the print name of a symbol. The name follows, one character per byte, with the first byte of the print name being the first read. The name is interned in the default package, and the resulting symbol is both pushed onto the stack and added to the table.

7 FOP-SMALL-SYMBOL-SAVE  $n(1)$  *name(n)*  $\Rightarrow$  stack & table

The one-byte operand *n* specifies the length of the print name of a symbol. The name follows, one character per byte, with the first byte of the print name being the first read. The name is interned in the default package, and the resulting symbol is both pushed onto the stack and added to the table.

8 FOP-SYMBOL-IN-PACKAGE-SAVE *index(4)*  $n(4)$  *name(n)*  $\Rightarrow$  stack & table

The four-byte *index* specifies a package stored in the table. The four-byte operand *n* specifies the length of the print name of a symbol. The name follows, one character per byte, with the first byte of the print name being the first read. The name is interned in the specified package, and the resulting symbol is both pushed onto the stack and added to the table.

9 FOP-SMALL-SYMBOL-IN-PACKAGE-SAVE *index(4)*  $n(1)$  *name(n)*  $\Rightarrow$  stack & table

The four-byte *index* specifies a package stored in the table. The one-byte operand *n* specifies the length of the print name of a symbol. The name follows, one character per byte, with the first byte of the print name being the first read. The name is interned in the specified package, and the resulting symbol is both pushed onto the stack and added to the table.

10 FOP-SYMBOL-IN-BYTE-PACKAGE-SAVE *index(1)*  $n(4)$  *name(n)*  $\Rightarrow$  stack & table

The one-byte *index* specifies a package stored in the table. The four-byte operand *n* specifies the length of the print name of a symbol. The name follows, one character per byte, with the first byte of the print name being the first read. The name is interned in the specified package, and the resulting symbol is both pushed onto the stack and added to the table.

11 FOP-SMALL-SYMBOL-IN-BYTE-PACKAGE-SAVE *index(1)*  $n(1)$  *name(n)*  $\Rightarrow$  stack & table

The one-byte *index* specifies a package stored in the table. The one-byte operand *n* specifies the length of the print name of a symbol. The name follows, one character per byte, with the first byte of the print name being the first read. The name is interned in the specified package, and the resulting symbol is both pushed onto the stack and added to the table.

12 Unused.

13 FOP-DEFAULT-PACKAGE *index*(4)

A package stored in the table entry specified by *index* is made the default package for future FOP-SYMBOL and FOP-SMALL-SYMBOL interning operations. (These package FOPs may change or disappear as the package system is determined.)

14 FOP-PACKAGE  $\Rightarrow$  table

An item is popped from the stack; it must be a symbol. The package of that name is located and pushed onto the table.

15 FOP-LIST *length*(1)  $\Rightarrow$  stack

The unsigned operand *length* specifies a number of operands to be popped from the stack. These are made into a list of that length, and the list is pushed onto the stack. The first item popped from the stack becomes the last element of the list, and so on. Hence an iterative loop can start with the empty list and perform "pop an item and cons it onto the list" *length* times. (Lists of length greater than 255 can be made by using FOP-LIST\* repeatedly.)

16 FOP-LIST\* *length*(1)  $\Rightarrow$  stack

This is like FOP-LIST except that the constructed list is terminated not by ( ) (the empty list), but by an item popped from the stack before any others are. Therefore *length*+1 items are popped in all. Hence an iterative loop can start with a popped item and perform "pop an item and cons it onto the list" *length*+1 times.

17-24 FOP-LIST-1, FOP-LIST-2, ..., FOP-LIST-8

FOP-LIST-*k* is like FOP-LIST with a byte containing *k* following it. These exist purely to reduce the size of Fasload files. Measurements need to be made to determine the useful values of *k*.

25-32 FOP-LIST\*-1, FOP-LIST\*-2, ..., FOP-LIST\*-8

FOP-LIST\*-*k* is like FOP-LIST\* with a byte containing *k* following it. These exist purely to reduce the size of Fasload files. Measurements need to be made to determine the useful values of *k*.

33 FOP-INTEGER *n*(4) *value*(*n*)  $\Rightarrow$  stack

A four-byte unsigned operand specifies the number of following bytes. These bytes define the value of a signed integer in two's-complement form. The first byte of the value is the least significant byte.

34 FOP-SMALL-INTEGER *n*(1) *value*(*n*)  $\Rightarrow$  stack

A one-byte unsigned operand specifies the number of following bytes. These bytes define the value of a signed integer in two's-complement form. The first byte of the value is the least significant byte.

35 FOP-WORD-INTEGER *value*(4)  $\Rightarrow$  stack

A four-byte integer (in the range  $-2^{31}$  to  $2^{31}-1$ ) follows the operation code. A LISP integer (fixnum or bignum) with that value is constructed and pushed onto the stack.

- 36 FOP-BYTE-INTEGER  $value(1) \Rightarrow stack$   
 A one-byte signed integer (in the range -128 to 127) follows the operation code. A LISP integer (fixnum or bignum) with that value is constructed and pushed onto the stack.
- 37 FOP-STRING  $n(4) name(n) \Rightarrow stack$   
 The four-byte operand  $n$  specifies the length of a string to construct. The characters of the string follow, one per byte. The constructed string is pushed onto the stack.
- 38 FOP-SMALL-STRING  $n(1) name(n) \Rightarrow stack$   
 The one-byte operand  $n$  specifies the length of a string to construct. The characters of the string follow, one per byte. The constructed string is pushed onto the stack.
- 39 FOP-VECTOR  $n(4) \Rightarrow stack$   
 The four-byte operand  $n$  specifies the length of a vector of LISP objects to construct. The elements of the vector are popped off the stack; the first one popped becomes the last element of the vector. The constructed vector is pushed onto the stack.
- 40 FOP-SMALL-VECTOR  $n(1) \Rightarrow stack$   
 The one-byte operand  $n$  specifies the length of a vector of LISP objects to construct. The elements of the vector are popped off the stack; the first one popped becomes the last element of the vector. The constructed vector is pushed onto the stack.
- 41 FOP-UNIFORM-VECTOR  $n(4) \Rightarrow stack$   
 The four-byte operand  $n$  specifies the length of a vector of LISP objects to construct. A single item is popped from the stack and used to initialize all elements of the vector. The constructed vector is pushed onto the stack.
- 42 FOP-SMALL-UNIFORM-VECTOR  $n(1) \Rightarrow stack$   
 The one-byte operand  $n$  specifies the length of a vector of LISP objects to construct. A single item is popped from the stack and used to initialize all elements of the vector. The constructed vector is pushed onto the stack.
- 43 FOP-INT-VECTOR  $n(4) size(1) count(1) data(\lceil n/count \rceil \lceil size*count/8 \rceil) \Rightarrow stack$   
 The four-byte operand  $n$  specifies the length of a vector of unsigned integers to be constructed. Each integer is  $size$  bits big, and are packed in the data stream in sections of  $count$  apiece. Each section occupies an integral number of bytes. If the bytes of a section are lined up in a row, with the first byte read at the right, and successive bytes placed to the left, with the bits within a byte being arranged so that the low-order bit is to the right, then the integers of the section are successive groups of  $size$  bits, starting from the right and running across byte boundaries. (In other words, this is a consistent right-to-left convention.) Any bits wasted at the left end of a section are ignored, and any wasted groups in the last section are ignored. It is permitted for the loading implementation to use a vector format providing more precision than is required by  $size$ . For example, if  $size$  were 3, it would be permitted to use a vector of 4-bit integers, or even vector of general LISP objects filled with integer LISP objects. However, an implementation is expected to use the most restrictive format that will suffice, and is expected to reconstruct objects identical to those output if the Fasload file was produced by the same implementation. (For the PERQ U-vector formats, one would have  $size$  an element of {1, 2, 4, 8, 16}, and  $count = 32/size$ ; words could be read directly into the U-vector. This operation provides a

very general format whereby almost any conceivable implementation can output in its preferred packed format, and another can read it meaningfully; by checking at the beginning for good cases, loading can still proceed quickly.) The constructed vector is pushed onto the stack.

44 FOP-UNIFORM-INT-VECTOR  $n(4)$   $size(1)$   $value(\lceil size/8 \rceil) \Rightarrow$  stack

The four-byte operand  $n$  specifies the length of a vector of unsigned integers to construct. Each integer is  $size$  bits big, and is initialized to the value of the operand  $value$ . The constructed vector is pushed onto the stack.

45 FOP-FLOAT  $n(1)$   $exponent(\lceil n/8 \rceil)$   $m(1)$   $mantissa(\lceil m/8 \rceil) \Rightarrow$  stack

The first operand  $n$  is one unsigned byte, and describes the number of *bits* in the second operand  $exponent$ , which is a signed integer in two's-complement format. The high-order bits of the last (most significant) byte of  $exponent$  shall equal the sign bit. Similar remarks apply to  $m$  and  $mantissa$ . The value denoted by these four operands is  $mantissa \times 2^{exponent - \text{length}(mantissa)}$ . A floating-point number shall be constructed which has this value, and then pushed onto the stack. That floating-point format should be used which is the smallest (most compact) provided by the implementation which nevertheless provides enough accuracy to represent both the exponent and the mantissa correctly.

46-51 Unused

52 FOP-ALTER  $index(1)$

Two items are popped from the stack: call the first *newval* and the second *object*. The component of *object* specified by *index* is altered to contain *newval*. The precise operation depends on the type of *object*:

List                    A zero *index* means alter the car (perform RPLACA), and *index*=1 means alter the cdr (RPLACD).

Symbol                By definition these indices have the following meaning, and have nothing to do with the actual representation of symbols in a given implementation:

0                      Alter value cell.

1                      Alter function cell.

2                      Alter property list (!).

Vector (of any kind)

Alter component number *index* of the vector.

String

Alter character number *index* of the string.

53 FOP-EVAL  $\Rightarrow$  stack

Pop an item from the stack and evaluate it (give it to EVAL). Push the result back onto the stack.

## 54 FOP-EVAL-FOR-EFFECT

Pop an item from the stack and evaluate it (give it to EVAL). The result is ignored.

55 FOP-FUNCALL *nargs*(1)  $\Rightarrow$  stack

Pop *nargs*+1 items from the stack and apply the last one popped as a function to all the rest as arguments (the first one popped being the last argument). Push the result back onto the stack.

56 FOP-FUNCALL-FOR-EFFECT *nargs*(1)

Pop *nargs*+1 items from the stack and apply the last one popped as a function to all the rest as arguments (the first one popped being the last argument). The result is ignored.

57 FOP-CODE-FORMAT *id*(1)

The operand *id* is a unique identifier specifying the format for following code objects. The operations FOP-CODE and its relatives may not occur in a group until after FOP-CODE-FORMAT has appeared; there is no default format. This is provided so that several compiled code formats may co-exist in a file, and so that a loader can determine whether or not code was compiled by the correct compiler for the implementation being loaded into. So far the following code format identifiers are defined:

0 PERQ

1 VAX

58 FOP-CODE *nitems*(4) *size*(4) *code*(*size*)  $\Rightarrow$  stack

A compiled function is constructed and pushed onto the stack. This object is in the format specified by the most recent occurrence of FOP-CODE-FORMAT. The operand *nitems* specifies a number of items to pop off the stack to use in the "boxed storage" section. The operand *code* is a string of bytes constituting the compiled executable code.

59 FOP-SMALL-CODE *nitems*(1) *size*(2) *code*(*size*)  $\Rightarrow$  stack

A compiled function is constructed and pushed onto the stack. This object is in the format specified by the most recent occurrence of FOP-CODE-FORMAT. The operand *nitems* specifies a number of items to pop off the stack to use in the "boxed storage" section. The operand *code* is a string of bytes constituting the compiled executable code.

## 60 FOP-STATIC-HEAP

Until further notice operations which allocate data structures may allocate them in the static area rather than the dynamic area. (The default area for allocation is the dynamic area; this default is reset whenever a new group is begun. This command is of an advisory nature; implementations with no static heap can ignore it.)

## 61 FOP-DYNAMIC-HEAP

Following storage allocation should be in the dynamic area.

62 FOP-VERIFY-TABLE-SIZE *size*(4)

If the current size of the table is not equal to *size*, then an inconsistency has been detected. This operation is inserted into a Fasload file purely for error-checking purposes. It is good practice for a compiler to output this at least at the end of every group, if not more often.

## 63 FOP-VERIFY-EMPTY-STACK

If the stack is not currently empty, then an inconsistency has been detected. This operation is inserted into a Fasload file purely for error-checking purposes. It is good practice for a compiler to output this at least at the end of every group, if not more often.

## 64 FOP-END-GROUP

This is the last operation of a group. If this is not the last byte of the file, then a new group follows; the next nine bytes must be "FASL FILE".

65 FOP-POP-FOR-EFFECT *stack* ⇒

One item is popped from the stack.

66 FOP-MISC-TRAP ⇒ *stack*

A trap object is pushed onto the stack.

## 67 FOP-READ-ONLY-HEAP

Following storage allocation may be in a read-only heap. (For symbols, the symbol itself may not be in a read-only area, but its print name (a string) may be. This command is of an advisory nature; implementations with no read-only heap can ignore it, or use a static heap.)

68 FOP-CHARACTER *character(3)* ⇒ *stack*

The three bytes specify the 24 bits of a Spice Lisp character object. The bytes, lowest first, represent the code, control, and font bits. A character is constructed and pushed onto the stack.

69 FOP-SHORT-CHARACTER *character(1)* ⇒ *stack*

The one byte specifies the lower eight bits of a spice lisp character object (the code). A character is constructed with zero control and zero font attributes and pushed onto the stack.

70 FOP-RATIO ⇒ *stack*

Creates a ratio from two integers popped from the stack. The denominator is popped first, the numerator second.

71 FOP-COMPLEX ⇒ *stack*

Creates a complex number from two numbers popped from the stack. The imaginary part is popped first, the real part second.

72 FOP-LINK-ADDRESS-FIXUP *nargs(1) restp(1) offset(4)* ⇒ *stack*

Valid only for when FOP-CODE-FORMAT corresponds to the Vax. This operation pops a symbol and a code object from the stack and pushes a modified code object back onto the stack according to the needs of the runtime Vax code linker.

73 FOP-LINK-FUNCTION-FIXUP *offset(4)* ⇒ *stack*

Valid only for when FOP-CODE-FORMAT corresponds to the Vax. This operation pops a symbol and a code object from the stack and pushes a modified code object back onto the stack according to the needs of the runtime Vax code linker.

74 FOP-FSET

Pops the top two things off of the stack and uses them as arguments to FSET (i.e. SETF of SYMBOL-FUNCTION).

255 FOP-END-HEADER

Indicates the end of a group header, as described above.

## Appendix II

# The Opcode Definition File

```

;;; -*- Lisp -*-
;;;
;;; Instruction definitions for Spice Lisp.
;;; Written by Skef Wholey.
;;;
;;; This file contains information about the instruction set and is
;;; used by the microassembler, the compiler, the error system, and the
;;; disassembler.
;;;

(defvar *1byte-instruction-table*
  (make-array 256)
  "Table used to find the name of a 1 byte long instruction given its
  opcode.")

(defvar *2byte-instruction-table*
  (make-array 256)
  "Table used to find the name of a 2 byte long instruction given the second
  byte of its opcode.")

(defvar *instruction-list* () "List of the instruction names.")

;;; We do this random setq so that the right thing happens when a new
;;; definition file is loaded.

(setq *instruction-list* ())

(defun definstruction (name opcode
  &optional (type 'read) (operand 'stack) offset)
  "Defines an instruction with the given Name (a symbol) and Opcode.
  Opcode may be either a single integer or a list of integers. Type
  should be one of Read, Write, Read-Modify-Write, Long-Branch, or
  Short-Branch. Operand defaults to Stack. Instructions which don't
  really have operands are considered to be Read Stack operations.
  Operand should be one of Stack, PSIC, NSIC, AL, Long-AL, Constant,
  Long-Constant, Symbol, Long-Symbol, or Ignore. If the instruction has an
  implied offset, that should be specified with the Offset."
  (if (fixnum opcode)
      (setq opcode (list opcode)))
  (if (not (listp opcode))
      (error "The opcode for ~S must be either an integer or a list." name))
  (setf (get name '%instruction-opcode) opcode)
  (setf (get name '%instruction-length)
        (+ (length opcode)
           (cond ((memq type '(read write read-modify-write))
                  (cond (offset 0)
                        ((memq operand '(stack ignore))
                         0)
                        ((memq operand '(psic nsic al constant symbol))
                         1)
                        ((memq operand '(long-al long-constant long-symbol))
                         2)))))))

```



```

                2)
                (t
                 (error "~S is a losing operand." operand))))
                ((eq type 'long-branch) 2)
                ((eq type 'short-branch) 1)
                ((eq type 'long-dispatch) 3)
                ((eq type 'short-dispatch) 4)
                (t (error "~S is a losing type." type))))
(setf (get name '%instruction-type) type)
(setf (get name '%instruction-operand) operand)
(setf (get name '%instruction-offset) offset)
(push name *instruction-list*)
(if (= (car opcode) 254)
    (setf (aref *2byte-instruction-table* (cadr opcode)) name)
    (setf (aref *1byte-instruction-table* (car opcode)) name)))

```

```

;;; Definstructionclass is used to define a class of instructions, i.e. a
;;; set of instructions that perform the same operation on operands in
;;; different places. Each instruction in the class has its %Instruction-Gr
;;; property set to the Stack-Form.

```

```

(defun definstructionclass (stack-form &rest other-forms)
  (setf (get stack-form '%instruction-class) other-forms)
  (do ((forms other-forms (cdr forms))
        ((null forms))
        (let ((glob (cdar forms))
              (if (listp glob)
                  (do ((subforms glob (cdr subforms))
                        ((null subforms))
                            (setf (get (cdar subforms) '%instruction-group) stack-form))
                            (setf (get glob '%instruction-group) stack-form))))))

```

```

;;; 1Byte generates a definstruction for a one-byte instruction.

```

```

(defvar *1byte-instruction-counter* nil
  "Counter used to generate unique 1 byte long instructions.")

(defmacro 1byte (name . other-stuff)
  "Generates a Definstruction for the Name and Other-Stuff with a unique
  one-byte opcode."
  '(definstruction ,name ,(prog1 *1byte-instruction-counter*
                                (incf *1byte-instruction-counter*))
    . ,other-stuff))

```

```

;;; 2Byte generates a definstruction for a two-byte instruction.

```

```

(defvar *2byte-instruction-counter* nil
  "Counter used to generate unique 2 byte long instructions.")

(defmacro 2byte (name . other-stuff)
  "Generates a Definstruction for the Name and Other-Stuff with a unique
  one-byte opcode."
  '(definstruction ,name '(254 ,(prog1 *2byte-instruction-counter*
                                       (incf *2byte-instruction-counter*)))
    . ,other-stuff))

```

```

    ,other-stuff))

;;; Set the counts:

(eval-when (compile)
  (setq *1byte-instruction-counter* 1)
  (setq *2byte-instruction-counter* 0))

;;; InstrSynonym defines a synonym for an instruction.

(defmacro instrsynonym (for is)
  '(progn (setf (get ,is '%instruction-offset)
                (get ,for '%instruction-offset))
          (setf (get ,is '%instruction-destination)
                (get ,for '%instruction-destination))
          (setf (get ,is '%instruction-type)
                (get ,for '%instruction-type))
          (setf (get ,is '%instruction-length)
                (get ,for '%instruction-length))
          (setf (get ,is '%instruction-opcode)
                (get ,for '%instruction-opcode))))

;;; Allocation:

(2byte 'get-allocation-space)
(2byte 'set-allocation-space)
(2byte 'alloc-bit-vector)
(2byte 'alloc-i-vector)
(2byte 'alloc-string)
(2byte 'alloc-bignum)
(2byte 'float-long)
(2byte 'make-complex)
(2byte 'make-ratio)
(2byte 'alloc-g-vector)
(definstructionclass 'vector
  '(psic . vector-psic))
(2byte 'vector)
(2byte 'vector-psic 'read 'psic)
(2byte 'alloc-function)
(2byte 'alloc-array)
(2byte 'alloc-symbol)
(1byte 'cons)
(definstructionclass 'list
  '(psic . list-psic))
(2byte 'list)
(2byte 'list-psic)
(definstructionclass 'list*
  '(psic . list*-psic))
(2byte 'list*)
(2byte 'list*-psic 'read 'psic)

;;; Stack manipulation:

(definstructionclass 'push

```

```
'(psic . ((psic . push-psic)
          (0 . push-psic0)
          (1 . push-psic1)
          (2 . push-psic2)
          (3 . push-psic3)
          (4 . push-psic4)
          (5 . push-psic5)
          (6 . push-psic6)
          (7 . push-psic7)
          (8 . push-psic8)))
'(nsic . push-nsic)
'(al . ((al . push-al)
        (0 . push-a10)
        (1 . push-a11)
        (2 . push-a12)
        (3 . push-a13)
        (4 . push-a14)
        (5 . push-a15)
        (6 . push-a16)
        (7 . push-a17)
        (8 . push-a18)
        (9 . push-a19)
        (10 . push-a110)
        (11 . push-a111)
        (12 . push-a112)
        (13 . push-a113)))
'(longal . push-longal)
'(c . ((c . push-c)
        (1 . push-c1)
        (2 . push-c2)
        (3 . push-c3)
        (4 . push-c4)
        (5 . push-c5)
        (6 . push-c6)))
'(longc . push-longc)
'(s . push-s)
'(longs . push-longs))
(1byte 'push-psic 'read 'psic)
(instrsynonym 'set-0 'push-psic0)
(1byte 'push-psic1 'read 'psic 1)
(1byte 'push-psic2 'read 'psic 2)
(1byte 'push-psic3 'read 'psic 3)
(1byte 'push-psic4 'read 'psic 4)
(1byte 'push-psic5 'read 'psic 5)
(1byte 'push-psic6 'read 'psic 6)
(1byte 'push-psic7 'read 'psic 7)
(1byte 'push-psic8 'read 'psic 8)
(1byte 'push-nsic 'read 'nsic)
(1byte 'push-al 'read 'al)
(1byte 'push-a10 'read 'al 0)
(1byte 'push-a11 'read 'al 1)
(1byte 'push-a12 'read 'al 2)
(1byte 'push-a13 'read 'al 3)
(1byte 'push-a14 'read 'al 4)
(1byte 'push-a15 'read 'al 5)
```

```

(1byte 'push-a16 'read 'a1 6)
(1byte 'push-a17 'read 'a1 7)
(1byte 'push-a18 'read 'a1 8)
(1byte 'push-a19 'read 'a1 9)
(1byte 'push-a110 'read 'a1 10)
(1byte 'push-a111 'read 'a1 11)
(1byte 'push-a112 'read 'a1 12)
(1byte 'push-a113 'read 'a1 13)
(2byte 'push-longa1 'read 'long-a1)
(1byte 'push-c 'read 'constant)
(1byte 'push-c1 'read 'constant)
(1byte 'push-c2 'read 'constant)
(1byte 'push-c3 'read 'constant)
(1byte 'push-c4 'read 'constant)
(1byte 'push-c5 'read 'constant)
(1byte 'push-c6 'read 'constant)
(2byte 'push-longc 'read 'long-constant)
(1byte 'push-s 'read 'symbol)
(2byte 'push-longs 'read 'long-symbol)
(definstructionclass 'pop
  '(a1 . ((a1 . pop-a1)
          (0 . pop-a10)
          (1 . pop-a11)
          (2 . pop-a12)
          (3 . pop-a13)
          (4 . pop-a14)
          (5 . pop-a15)
          (6 . pop-a16)
          (7 . pop-a17)))
  '(longa1 . pop-longa1)
  '(s . pop-s)
  '(longs . pop-longs)
  '(ignore . pop-ignore))
(1byte 'pop-a1 'write 'a1)
(1byte 'pop-a10 'write 'a1 0)
(1byte 'pop-a11 'write 'a1 1)
(1byte 'pop-a12 'write 'a1 2)
(1byte 'pop-a13 'write 'a1 3)
(1byte 'pop-a14 'write 'a1 4)
(1byte 'pop-a15 'write 'a1 5)
(1byte 'pop-a16 'write 'a1 6)
(1byte 'pop-a17 'write 'a1 7)
(2byte 'pop-longa1 'write 'long-a1)
(1byte 'pop-s 'write 'symbol)
(2byte 'pop-longs 'write 'long-symbol)
(1byte 'pop-ignore 'write 'ignore)
(2byte 'exchange)
(definstructionclass 'copy
  '(a1 . ((a1 . copy-a1)
          (2 . copy-a12)
          (3 . copy-a13)
          (4 . copy-a14)
          (5 . copy-a15))))
(2byte 'copy)
(1byte 'copy-a1 'write 'a1)

```

```
(1byte 'copy-a12 'write 'a1 2)
(1byte 'copy-a13 'write 'a1 3)
(1byte 'copy-a14 'write 'a1 4)
(1byte 'copy-a15 'write 'a1 5)
(definstructionclass 'npop
  '(nsic . npop-nsic))
(1byte 'npop)
(1byte 'npop-nsic 'write 'nsic)
(definstructionclass 'bind-null
  '(c . bind-null-c))
(2byte 'bind-null)
(2byte 'bind-null-c 'read 'constant)
(definstructionclass 'bind
  '(c . bind-c))
(2byte 'bind)
(1byte 'bind-c 'read 'constant)
(definstructionclass 'unbind
  '(psic . unbind-psic))
(2byte 'unbind)
(2byte 'unbind-psic 'read 'psic)
```

::: List manipulation:

```
(definstructionclass 'car
  '(a1 . car-a1))
(1byte 'car)
(1byte 'car-a1 'read 'a1)
(definstructionclass 'cdr
  '(a1 . cdr-a1))
(1byte 'cdr)
(1byte 'cdr-a1 'read 'a1)
(definstructionclass 'cadr
  '(a1 . cadr-a1))
(1byte 'cadr)
(1byte 'cadr-a1 'read 'a1)
(definstructionclass 'caddr
  '(a1 . caddr-a1))
(1byte 'caddr)
(1byte 'caddr-a1 'read 'a1)
(definstructionclass 'cdar
  '(a1 . cdar-a1))
(1byte 'cdar)
(1byte 'cdar-a1 'read 'a1)
(definstructionclass 'caar
  '(a1 . caar-a1))
(1byte 'caar)
(1byte 'caar-a1 'read 'a1)
(definstructionclass 'set-cdr
  '(a1 . set-cdr-a1)
  '(s . set-cdr-s))
(1byte 'set-cdr-a1 'read-modify-write 'a1)
(1byte 'set-cdr-s 'read-modify-write 'symbol)
(definstructionclass 'set-caddr
  '(a1 . set-caddr-a1)
  '(s . set-caddr-s))
```

```
(1byte 'set-cddr-a1 'read-modify-write 'a1)
(2byte 'set-cddr-s 'read-modify-write 'symbol)
(definstructioclclass 'spread
  '(a1 . spread-a1))
(2byte 'spread)
(2byte 'spread-a1 'read 'a1)
(definstructioclclass 'replace-car
  '(a1 . replace-car-a1))
(1byte 'replace-car)
(1byte 'replace-car-a1 'read 'a1)
(definstructioclclass 'replace-cdr
  '(a1 . replace-cdr-a1))
(1byte 'replace-cdr)
(1byte 'replace-cdr-a1 'read 'a1)
(2byte 'assoc)
(2byte 'assq)
(2byte 'member)
(2byte 'memq)
(definstructioclclass 'endp
  '(a1 . endp-a1))
(2byte 'endp)
(1byte 'endp-a1 'read 'a1)
(2byte 'getf)
```

::: Symbol manipulation:

```
(1byte 'get-value)
(1byte 'set-value)
(1byte 'get-definition)
(2byte 'set-definition)
(1byte 'get-plist)
(2byte 'set-plist)
(1byte 'get-pname)
(2byte 'get-package)
(2byte 'set-package)
(2byte 'boundp)
(2byte 'fboundp)
```

::: Array manipulation:

```
(2byte 'vector-length)
(2byte 'g-vector-length)
(2byte 'simple-string-length)
(2byte 'simple-bit-vector-length)
(2byte 'get-vector-subtype)
(2byte 'set-vector-subtype)
(2byte 'get-vector-access-code)
(2byte 'shrink-vector)
(2byte 'typed-vref)
(2byte 'typed-vset)
(2byte 'header-length)
(2byte 'header-ref)
(2byte 'header-set)
(2byte 'aref1)
(1byte 'svref)
```

```
(1byte 'schar)
(1byte 'sbit)
(2byte 'aset1)
(definstructionclass 'svset
  '(ignore . (svset-ignore)))
(1byte 'svset)
(1byte 'svset-ignore)
(definstructionclass 'scharset
  '(ignore . (scharset-ignore)))
(1byte 'scharset)
(1byte 'scharset-ignore)
(1byte 'sbitset)
(2byte 'bit-bash)
(2byte 'byte-blt)
(2byte 'find-character)
(2byte 'find-character-with-attribute)
(2byte 'sxhash-simple-string)
```

::: Type predicates:

```
(1byte 'get-type)
(2byte 'get-space)
(2byte 'bit-vector-p)
(2byte 'simple-bit-vector-p)
(2byte 'simple-integer-vector-p)
(1byte 'stringp)
(1byte 'simple-string-p)
(2byte 'bignump)
(2byte 'long-float-p)
(2byte 'complexp)
(2byte 'ratiop)
(2byte 'integerp)
(2byte 'rationalp)
(2byte 'floatp)
(2byte 'numberp)
(2byte 'general-vector-p)
(1byte 'simple-vector-p)
(2byte 'compiled-function-p)
(1byte 'arrayp)
(1byte 'vectorp)
(2byte 'complex-array-p)
(1byte 'symbolp)
(1byte 'listp)
(1byte 'atom)
(1byte 'consp)
(1byte 'fixnump)
(2byte 'short-float-p)
(2byte 'characterp)
```

::: Arithmetic and Logic:

```
(2byte 'integer-length)
(2byte 'float-short)
(2byte 'realpart)
(2byte 'imagpart)
```

```

(2byte 'numerator)
(2byte 'denominator)
(2byte 'decode-float)
(2byte 'scale-float)
(definstructionclass '=
  '(al . ==-al)
  '(psic . ==-psic))
(1byte '=)
(1byte '==-al 'read 'al)
(1byte '==-psic 'read 'psic)
(definstructionclass '<
  '(al . <-al)
  '(psic . <-psic))
(1byte '<')
(1byte '<-al 'read 'al)
(1byte '<-psic 'read 'psic)
(definstructionclass '>
  '(al . >-al)
  '(psic . >-psic))
(1byte '>')
(1byte '>-al 'read 'al)
(1byte '>-psic 'read 'psic)
(1byte 'truncate)
(definstructionclass '+'
  '(al . +-al)
  '(psic . ((psic . +-psic)
            (1 . +-psic1)
            (2 . +-psic2))))
(1byte '+)
(1byte '+-psic 'read 'psic)
(1byte '+-psic1 'read-modify-write 'psic 1)
(1byte '+-psic2 'read 'psic 2)
(1byte '+-al 'read 'al)
(definstructionclass '-'
  '(al . --al)
  '(psic . ((psic . --psic)
            (1 . --psic1)
            (2 . --psic2))))
(1byte '-')
(1byte '--psic 'read 'psic)
(1byte '--psic1 'read-modify-write 'psic 1)
(1byte '--psic2 'read 'psic 2)
(1byte '--al 'read 'al)
(1byte '*')
#+Common
(1byte '/' )
#-Common
(1byte '//)
(definstructionclass '1+
  '(al . 1+-al))
(instrsynonym '+-psic1 '1+)
(1byte '1+-al 'read-modify-write 'al)
(definstructionclass '1-
  '(al . 1--al))
(instrsynonym '--psic1 '1-)

```



```

(1byte '1--al 'read-modify-write 'al)
(1byte 'negate)
(2byte 'abs)
(2byte 'logand)
(2byte 'logior)
(2byte 'logxor)
(2byte 'lognot)
(2byte 'boole)
(2byte 'ash)
(2byte 'ldb)
(2byte 'mask-field)
(2byte 'dpb)
(2byte 'deposit-field)
(2byte 'lsh)
(2byte 'logldb)
(2byte 'logdpb)

```

::: Branching and dispatching:

```

(1byte 'branch-forward 'short-branch)
(2byte 'long-branch-forward 'long-branch)
(1byte 'branch-backward 'short-branch)
(2byte 'long-branch-backward 'long-branch)
(1byte 'branch-null-forward 'short-branch)
(2byte 'long-branch-null-forward 'long-branch)
(1byte 'branch-not-null-forward 'short-branch)
(2byte 'long-branch-not-null-forward 'long-branch)
(1byte 'branch-null-backward 'short-branch)
(2byte 'long-branch-null-backward 'long-branch)
(1byte 'branch-not-null-backward 'short-branch)
(2byte 'long-branch-not-null-backward 'long-branch)
(1byte 'branch-save-not-null-forward 'short-branch)
(2byte 'long-branch-save-not-null-forward 'long-branch)
(1byte 'branch-save-not-null-backward 'short-branch)
(2byte 'long-branch-save-not-null-backward 'long-branch)
(2byte 'dispatch 'short-dispatch)
(2byte 'long-dispatch 'long-dispatch)

```

::: Function call and return:

```

(definstructiionclass 'call
  '(c . ((c . call-c)
          (1 . call-c1)
          (2 . call-c2)
          (3 . call-c3)
          (4 . call-c4))))
(2byte 'call)
(1byte 'call-c 'read 'constant)
(1byte 'call-c1 'read 'constant 1)
(1byte 'call-c2 'read 'constant 2)
(1byte 'call-c3 'read 'constant 3)
(1byte 'call-c4 'read 'constant 4)
(definstructiionclass 'call-0
  '(c . call-0-c))
(2byte 'call-0)

```

```

(1byte 'call-0-c 'read 'constant)
(definstructionclass 'call-multiple
  '(c . call-multiple-c))
(1byte 'call-multiple)
(1byte 'call-multiple-c 'read 'constant)
(1byte 'start-call)
(definstructionclass 'push-last
  '(a1 . ((a1 . push-last-a1)
          (0 . push-last-a10)
          (1 . push-last-a11)
          (2 . push-last-a12)
          (3 . push-last-a13))))
(instrsynonym 'start-call 'push-last)
(1byte 'push-last-a1 'read 'a1)
(1byte 'push-last-a10 'read 'a1 0)
(1byte 'push-last-a11 'read 'a1 1)
(1byte 'push-last-a12 'read 'a1 2)
(1byte 'push-last-a13 'read 'a1 3)
(definstructionclass 'return
  '(a1 . return-a1))
(1byte 'return)
(1byte 'return-a1 'read 'a1)
(2byte 'escape-return)
(2byte 'break-return)
(2byte 'catch)
(2byte 'catch-multiple)
(2byte 'catch-all)
(2byte 'throw)

```

::: Miscellaneous:

```

(1byte 'eq)
(1byte 'eq1)
(1byte 'set-null)
(1byte 'set-t)
(1byte 'set-0)
(1byte 'make-predicate)
(1byte 'not-predicate)
(2byte 'values-to-n)
(2byte 'n-to-values)
(2byte 'force-values)
(2byte 'flush-values)

```

::: System hacking:

```

(2byte 'make-immediate-type)
(2byte '8bit-system-ref)
(2byte '8bit-system-set)
(2byte '16bit-system-ref)
(2byte '16bit-system-set)
(2byte 'collect-garbage)
(2byte 'newspace-bit)
(2byte 'kernel-trap)
(2byte 'halt)
(2byte 'arg-in-frame)

```

```
(2byte 'active-call-frame)
(2byte 'set-call-frame)
(2byte 'current-open-frame)
(2byte 'set-open-frame)
(2byte 'current-stack-pointer)
(2byte 'current-binding-pointer)
(2byte 'read-control-stack)
(2byte 'write-control-stack)
(2byte 'read-binding-stack)
(2byte 'write-binding-stack)

(setq *1byte-instruction-counter* #.*1byte-instruction-counter*)
(setq *2byte-instruction-counter* #.*2byte-instruction-counter*)

(format t "[~3D 1-byte instructions have been defined.]"
  (1- #.*1byte-instruction-counter*))
(terpri)
(format t "[~3D 2-byte instructions have been defined.]"
  #.*2byte-instruction-counter*)
```

# Index

%SP-Internal-Apply 13, 47  
%SP-Internal-Error 13  
%SP-Internal-Throw-Tag 13  
%SP-Software-Interrupt-Handler 13

\* 36

+ 36

- 36

/ 36

1+ 37

1- 37

16bit-System-Ref 43

16bit-System-Set 43

8bit-System-Ref 43

8bit-System-Set 43

< 35

= 35

> 35

Abs 37

Accent message space 4

Access-type codes 11

Active frame 16

Active-Call-Frame 44

Active-Catch register 14

Active-Catch-Frame 44

Active-Code register 14

Active-Frame register 14

Active-Function register 14

Alloc-Array 24

Alloc-Bignum 23

Alloc-Bit-Vector 23

Alloc-Function 24

Alloc-G-Vector 24

Alloc-I-Vector 23

Alloc-String 23

Alloc-Symbol 24

Arefl 30

Arg-In-Frame 44

Array format 7, 10

Array header format 12

ArrayP 34

Arrays 12

Asetl 31

Ash 38

Assoc 27

Assq 27

Bignum format 7, 12

BignumP 33  
Bind 26  
Bind-Null 26  
Binding stack format 17  
Binding stack space 8  
Binding-Stack pointer 4  
Binding-Stack-Pointer register 14  
Bit numbering 2  
Bit-Bash 32  
Bit-Vector format 7  
Bit-Vector-P 33  
Boole 37  
Boundp 29  
Branch 39  
Branch-Not-Null 39  
Branch-Null 39  
Branch-Save-Not-Null 39  
Break-Return 41  
Byte code formats 21  
Byte codes 21  
Byte numbering 2  
Byte-BLT 32  
  
Caar 26  
Cadr 26  
Call 40, 46  
Call Header format 5  
Call-0 40, 46, 47  
Call-Header 5  
Call-Multiple 40, 46  
Car 26  
CAref2 31  
CAref3 31  
CAset2 32  
CAset3 32  
Catch 17, 41, 49  
Catch frames 17  
Catch header format 5  
Catch-All 41  
Catch-All object 5, 49  
Catch-Frame 5  
Catch-Multiple 41  
Cdar 26  
Cddr 26  
Cdr 26  
CE (contents of effective address) 22  
Character object 5  
CharacterP 34  
Clean-Space pointer 18  
Code vector 15  
Collect-Garbage 43  
Compiled-Function-P 34  
Complex number format 7  
Complex-Array-P 34  
ComplexP 33  
Cons 24  
ConsP 34  
Constants in code 15  
Control registers 14  
Control stack space 8

Control-stack format 16  
Control-Stack pointer 4  
Control-Stack-Pointer register 14  
Copy 26  
Current-Binding-Pointer 45  
Current-Open-Frame 44  
Current-Stack-Pointer 44  
  
Decode-Float 35  
Definition cell 6  
DESTRUCT 10  
Denominator 35  
Deposit-Field 38  
Dispatch 39  
Dpb 38  
  
E (effective address) 22  
Effective address 21  
Endp 27  
Eq 41  
Eq1 41  
Errors 53  
Escape to macrocode convention 51  
Escape-Return 40  
Exchange 25  
  
FBoundp 29  
Find-Character 32  
Find-Character-With-Attribute 32  
Fixnum format 4  
FixnumP 34  
Float-Long 35  
Float-Short 35  
Floating point formats 4, 7  
FloatP 33  
Flonum formats 4, 7  
Flush-Values 42  
Force-Values 42  
Forwarding pointers 8  
Free-Storage pointer 18  
Function object format 8, 10  
  
G-Vector format 7  
G-Vector-Length 29  
Garbage Collection 18  
GC-Forward pointer 8  
General-Vector format 7, 9  
General-Vector-P 34  
Get-Allocation-Space 23  
Get-Definition 28  
Get-Package 28  
Get-Plist 28  
Get-Pname 28  
Get-Space 43  
Get-Type 43  
Get-Value 28  
Get-Vector-Access-Code 30  
Get-Vector-Subtype 29  
GetF 28

Hairy stuff 46  
Halt 44  
Hash tables 10  
Header-Length 30  
Header-Ref 30  
Header-Set 30  
  
I-Vector format 7  
Imagpart 35  
Immediate object format 3  
Integer-Length 35  
Integer-Vector format 7, 10  
IntegerP 33  
Interrupts 57  
  
Kernel traps 57  
Kernel-Trap 44  
  
Ldb 38  
Lisp objects 3  
List 24  
List cell 6  
List\* 24  
ListP 34  
Logand 37  
Logdpb 38  
Logior 37  
Logldb 38  
Lognot 37  
Logxor 37  
Long-Float-P 33  
Long-Flonum format 7  
Lsh 38  
  
Macro instruction formats 21  
Macro instruction set 21  
Make-Complex 23  
Make-Immediate-Type 43  
Make-Predicate 42  
Make-Ratio 23  
Mask-Field 38  
Member 27  
Memq 27  
Misc type codes 4  
Misc-Binding-Stack-Pointer 4  
Misc-Control-Stack-Pointer 4  
Misc-System-Table-Pointer 4  
Misc-Trap 4  
Multiple values 48  
  
N-To-Values 42  
Negate 37  
Newspace-Bit 43  
NIL 13  
Non-Local Exits 49  
Not-Predicate 42  
NPop 26  
NumberP 33  
Numerator 35

Open frame 16  
Open-Frame register 14

Package cell 6  
PC register 14  
Perq quadword alignment 9  
Plist cell 6  
Pname cell 6  
Pointer object format 3, 5  
Pop 25  
Print name cell 6  
Program Counter register 14  
Property list cell 6  
Purification 19  
Push 25  
Push-Last 40, 47

Quadword alignment 9

Ratio format 7  
RationalP 33  
RatioP 33  
Read-Binding-Stack 45  
Read-Control-Stack 45  
Read-only space 6  
Realpart 35  
Replace-Car 27  
Replace-Cdr 27  
Return 40, 48  
Runtime Environment 14

SBit 30  
SBitset 31  
Scale-Float 35  
Scavenger 19  
SChar 30  
SCharset 31  
Set-0 42  
Set-Allocation-Space 23  
Set-Call-Frame 44  
Set-Cddr 27  
Set-Cdr 27  
Set-Definition 28  
Set-Lpop 27  
Set-LPush 24  
Set-Newspace-For-Type 43  
Set-Null 42  
Set-Open-Frame 44  
Set-Package 29  
Set-Plist 28  
Set-T 42  
Set-Value 28  
Set-Vector-Subtype 29  
Short-Float format 4  
Short-Float-P 34  
Shrink-Vector 30  
Simple-Bit-Vector-Length 29  
Simple-Bit-Vector-P 33  
Simple-Integer-Vector-P 33  
Simple-String-Length 29



Simple-String-P 33  
Simple-Vector-P 34  
Space codes 4, 6  
Special binding stack space 8  
Spread 27  
Stack spaces 8  
Start-Call 40, 47  
Static space 6  
Storage management 18  
String format 7, 12  
StringP 33  
SVref 30  
SVset 31  
SVset\* 31  
SXI hash-Simple-String 32  
Symbol 6  
SymbolP 34  
System table pointer 4  
System table space 4, 8

T 13  
Throw 41, 49  
TOS register 14  
Transporter 18  
Trap code 4  
Trapping to the kernel 57  
Truncate 36, 52  
Type codes 3  
Typed-Vref 30  
Typed-Vset 30

Unbind 26  
Unwind-Protect 49

Value cell 6  
Values-Marker 5  
Values-Γσ-N 42  
Vector 24  
Vector format 7  
Vector-Length 29  
VectorP 34  
Vectors 9  
Virtual memory 5

Write-Binding-Stack 45  
Write-Control-Stack 45