

draft proposed American National Standard
for Information Systems —

Programming Languages —
Forth
(X3J14 dpANS-6 – June 30, 1993)

Secretariat

Computer and Business Equipment Manufacturers Association

American National Standards Institute, Inc.

WARNING

This is a draft of the proposed American National Standard. It is distributed solely for the purpose of review and comment and should not be used as a design document. It is inappropriate to claim compatibility with this draft standard.

Disclaimer

The X3J14 Technical Committee makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The Technical Committee shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Foreword

(This foreword is not a part of American National Standard X3.215-199x)

Forth is a language for direct communication between human beings and machines. Using natural-language diction and machine-oriented syntax, Forth provides an economical, productive environment for interactive compilation and execution of programs. Forth also provides low-level access to computer-controlled hardware, and the ability to extend the language itself. This extensibility allows the language to be quickly expanded and adapted to special needs and different hardware systems.

Forth was invented by Mr. Charles Moore to increase programmer productivity without sacrificing machine efficiency. Forth is a layered environment containing the elements of a computer language as well as those of an operating system and a machine monitor. This extensible, layered environment provides for highly interactive program development and testing.

In the interests of transportability of application software written in Forth, standardization efforts began in the mid-1970s by an international group of users and implementors who adopted the name "Forth Standards Team." This effort resulted in the Forth-77 Standard. As the language continued to evolve, an interim Forth-78 Standard was published by the Forth Standards Team. Following Forth Standards Team meetings in 1979, the Forth-79 Standard was published in 1980. Major changes were made by the Forth Standards Team in the Forth-83 Standard, which was published in 1983.

The first meeting of the Technical Committee on Forth Programming Systems was convened by the Organizing Committee of the X3J14 Forth Technical Committee on August 3, 1987, and has met subsequently on November 11-12, 1987, February 10-12, 1988, May 25-28, 1988, August 10-13, 1988, October 26-29, 1988, January 25-28, 1989, May 3-6, 1989, July 26-29, 1989, October 25-28, 1989, January 24-27, 1990, May 22-26, 1990, August 21-25, 1990, November 6-10, 1990, January 29-February 2, 1991, May 3-4, 1991, June 16-19, 1991, July 30-August 3, 1991, March 17-21, 1992, October 13-17, 1992, January 26-30, 1993, and June 28-30, 1993.

Membership

At the time it approved this draft of the proposed American National Standard, the Technical Committee X3J14 on the Forth Programming Language had the following members:

- Elizabeth Rather, Chair
- Mitch Bradley, acting Vice-Chair
- Don Colburn, Secretary
- John Rible, Technical Editor
- Len Zettel, Vocabulary Representative
- Greg Bailey, Technical Subcommittee Chair

<i>Organization Represented</i>	<i>Name of Representative</i>
ATHENA Programming, Inc.	Greg Bailey
	Howe Fong (Alt.)
Bradley Forthware, Inc.	Mitch Bradley
Creative Solutions, Inc.	Don Colburn
Ford Motor Company	Leonard F. Zettel, Jr.
FORTH, Inc.	Elizabeth Rather
	Dennis Ruffer (Alt.)
Institute for Applied Forth Research.....	Lawrence Forsley
	Horace Simmons (Alt.)
Johns Hopkins University, Applied Physics Lab.	John Hayes
Mephistopheles Systems.....	Dave Harralson
NASA/Goddard Space Flight Center.....	James Rash
Nomadic Software.....	John K. Stevenson
Unisyn, Inc.	Gary Betts
	Stephen Egbert (Alt.)
Up and Running	Martin Tracy
Vesta Technology.....	Jack Woehr
<i>Individual Members</i>	Loring Craymer
	John Rible
	J. E. (Jet) Thomas
<i>X3 Liasons</i>	Clyde R. Camp
	Kathleen McMillan

The following organizations and individuals have also participated in this project as Technical Committee members, alternates, or observers. The Technical Committee recognizes and respects their contributions.

Organizations

- | | |
|---------------------------------|--|
| British Columbia Inst. of Tech. | MCI Telecommunications Corp. |
| Computer Cowboys | Micromotion |
| Computer Sciences Corp. | MicroProcessor Engineering Ltd. |
| Computer Strategies, Inc. | National Inst. of Standards & Technology |
| Digalog Corp. | NCR Medical Systems Group |
| Embedded Sys. Programming Mag. | Performance Packages, Inc. |
| Forth Interest Group (FIG) | Purdue University |
| H.B. Pascal & Co., Inc. | Robert Berkey Services |
| Harris Semiconductor | Shaw Laboratories |
| IBM Corporation | Social Security Administration |
| IEEE | Software Engineering |
| Kelly Enterprises | Texas Instruments |
| Laboratory Microsystems, Inc. | The Dickens Company |
| Maxtor Corp. | |

Individuals

- | | | | |
|-----------------|------------------|----------------|----------------|
| David J. Angel | Ray Duncan | Charles Moore | James Ryland |
| Wil Baden | Douglas Fishman | Mike Nemeth | Dean Sanderson |
| Robert Berkey | Tom Hand | Harry Pascal | George Shaw |
| Ron Braithwaite | Gregory Ilg | Stephen Pelc | Gerald Shifrin |
| Jack Brown | Charles Keane | Dean Perrine | Robert Smith |
| Chris Colburn | Guy M. Kelly | David C. Petty | Tyler Sperry |
| Ted Dickens | Andrew Kobziar | Bill Ragsdale | Tom Zimmer |
| John Dorband | Martin Lascelles | Dennis Ruffer | |

Contents

1. Introduction	1
1.1 Purpose	1
1.2 Scope	1
1.2.1 Inclusions	1
1.2.2 Exclusions	1
1.3 Document organization	1
1.3.1 Word sets	1
1.3.2 Annexes	2
1.4 Future directions	2
1.4.1 New technology	2
1.4.2 Obsolescent features	2
2. Terms and notation	3
2.1 Definitions of terms	3
2.2 Notation	5
2.2.1 Numeric notation	5
2.2.2 Stack notation	5
2.2.3 Parsed-text notation	5
2.2.4 Glossary notation	6
3. Usage requirements	8
3.1 Data types	8
3.1.1 Data-type relationships	9
3.1.2 Character types	9
3.1.3 Single-cell types	10
3.1.4 Cell-pair types	11
3.1.5 System types	11
3.2 The implementation environment	12
3.2.1 Numbers	12
3.2.2 Arithmetic	12
3.2.3 Stacks	13
3.2.4 Operator terminal	14
3.2.5 Mass storage	14
3.2.6 Environmental queries	14
3.3 The Forth dictionary	15
3.3.1 Name space	15
3.3.2 Code space	15
3.3.3 Data space	15
3.4 The Forth text interpreter	17
3.4.1 Parsing	17
3.4.2 Finding definition names	18
3.4.3 Semantics	18
3.4.4 Possible actions on an ambiguous condition	19
3.4.5 Compilation	19
4. Documentation requirements	21
4.1 System documentation	21
4.1.1 Implementation-defined options	21
4.1.2 Ambiguous conditions	22
4.1.3 Other system documentation	23
4.2 Program documentation	23
4.2.1 Environmental dependencies	23
4.2.2 Other program documentation	23
5. Compliance and labeling	24
5.1 ANS Forth systems	24
5.1.1 System compliance	24

5.1.2	System labeling	24
5.2	ANS Forth programs	24
5.2.1	Program compliance	24
5.2.2	Program labeling	24
6.	Glossary	25
6.1	Core words	25
6.2	Core extension words	51
7.	The optional Block word set	61
8.	The optional Double-Number word set	67
9.	The optional Exception word set	72
10.	The optional Facility word set	76
11.	The optional File-Access word set	79
12.	The optional Floating-Point word set	89
13.	The optional Locals word set	105
14.	The optional Memory-Allocation word set	110
15.	The optional Programming-Tools word set	113
16.	The optional Search-Order word set	120
17.	The optional String word set	125
A.	Rationale (informative annex)	128
B.	Bibliography (informative annex)	181
C.	Perspective (informative annex)	183
C.1	Features of Forth	183
C.2	History of Forth	184
C.3	Hardware implementations of Forth	184
C.4	Standardization efforts	184
C.5	Programming in Forth	185
C.6	Multiprogrammed systems	191
C.7	Design and management considerations	191
C.8	Conclusion	192
D.	Compatibility analysis of ANS Forth (informative annex)	193
D.1	FIG Forth (circa 1978)	193
D.2	Forth 79	193
D.3	Forth 83	193
D.4	Recent developments	194
D.5	ANS Forth approach	194
D.6	Differences from Forth 83	195
D.6.1	Stack width	195
D.6.2	Number representation	195
D.6.3	Address units	196
D.6.4	Address increment for a cell is no longer two	196
D.6.5	Address alignment	197
D.6.6	Division/modulus rounding direction	198
D.6.7	Immediacy	198
D.6.8	Input character set	200
D.6.9	Shifting with UM/MOD	200
D.6.10	Vocabularies / wordlists	201
D.6.11	Multiprogramming impact	201
D.6.12	Words not provided in executable form	202
E.	ANS Forth portability guide (informative annex)	203
E.1	Introduction	203
E.2	Hardware peculiarities	203
E.2.1	Data/memory abstraction	203

E.2.2	Definitions	203
E.2.3	Addressing memory	204
E.2.4	Alignment problems	204
E.3	Number representation.....	205
E.3.1	Big endian vs. little endian.....	205
E.3.2	ALU organization	205
E.4	Forth system implementation	206
E.4.1	Definitions	206
E.4.2	Stacks	206
E.5	ROMed application disciplines and conventions	207
E.6	Summary	207
F.	Alphabetic list of words (informative annex).....	208

1. Introduction

1.1 Purpose

The purpose of this Standard is to promote the portability of Forth programs for use on a wide variety of computing systems, to facilitate the communication of programs, programming techniques, and ideas among Forth programmers, and to serve as a basis for the future evolution of the Forth language.

1.2 Scope

This Standard specifies an interface between a Forth System and a Forth Program by defining the words provided by a Standard System.

1.2.1 Inclusions

This Standard specifies:

- the forms that a program written in the Forth language may take;
- the rules for interpreting the meaning of a program and its data.

1.2.2 Exclusions

This Standard does not specify:

- the mechanism by which programs are transformed for use on computing systems;
- the operations required for setup and control of the use of programs on computing systems;
- the method of transcription of programs or their input or output data to or from a storage medium;
- the program and Forth system behavior when the rules of this Standard fail to establish an interpretation;
- the size or complexity of a program and its data that will exceed the capacity of any specific computing system or the capability of a particular Forth system;
- the physical properties of input/output records, files, and units;
- the physical properties and implementation of storage.

1.3 Document organization

1.3.1 Word sets

This Standard groups Forth words and capabilities into *word sets* under a name indicating some shared aspect, typically their common functional area. Each word set may have an extension, containing words that offer additional functionality. These words are not required in an implementation of the word set.

The “Core” word set, defined in sections 1 through 6, contains the required words and capabilities of a Standard System. The other word sets, defined in sections 7 through 17, are optional, making it possible to provide Standard Systems with tailored levels of functionality.

1.3.1.1 Text sections

Within each word set, section 1 contains introductory and explanatory material and section 2 introduces terms and notation used throughout the Standard. There are no requirements in these sections.

Sections 3 and 4 contain the usage and documentation requirements, respectively, for Standard Systems and Programs, while section 5 specifies their labeling.

1.3.1.2 Glossary sections

Section 6 of each word set specifies the required behavior of the definitions in the word set and the extensions word set.

1.3.2 Annexes

The annexes do not contain any required material.

Annex A provides some of the rationale behind the committee's decisions in creating this Standard, as well as implementation examples. It has the same section numbering as the body of the Standard to make it easy to relate each requirements section to its rationale section.

Annex B is a short bibliography on Forth.

Annex C provides an introduction to Forth.

Annex D discusses the compatibility of ANS Forth with earlier Forths, emphasizing the differences from Forth 83.

Annex E presents some techniques for writing portable programs in ANS Forth.

Annex F includes the words from all word sets in a single list, and serves as an index of ANS Forth words.

1.4 Future directions

1.4.1 New technology

This Standard adopts certain words and practices that are increasingly found in common practice. New words have also been adopted to ease creation of portable programs.

1.4.2 Obsolescent features

This Standard adopts certain words and practices that cause some previously used words to become obsolescent. Although retained here because of their widespread use, their use in new implementations or new programs is discouraged, because they may be withdrawn from future revisions of the Standard.

This Standard designates the following words as obsolescent:

6.2.0060 #TIB	15.6.2.1580 FORGET	6.2.2240 SPAN
6.2.0970 CONVERT	6.2.2040 QUERY	6.2.2290 TIB
6.2.1390 EXPECT		

2. Terms and notation

The phrase “See:” is used throughout this Standard to direct the reader to other sections of the Standard that have a direct bearing on the current section.

In this Standard, “shall” states a requirement on a system or program; conversely, “shall not” is a prohibition; “need not” means “is not required to”; “should” describes a recommendation of the Standard; and “may”, depending on context, means “is allowed to” or “might happen”.

Throughout the Standard, typefaces are used in the following manner:

- This proportional serif typeface is used for text, with *italic* used for symbols and the first appearance of new terms;
- A bold proportional sans-serif typeface is used for **headings**;
- A bold monospaced serif typeface is used for **Forth-language** text.

2.1 Definitions of terms

Terms defined in this section are used generally throughout this Standard. Additional terms specific to individual word sets are defined in those word sets. Other terms are defined at their first appearance, indicated by italic type. Terms not defined in this Standard are to be construed according to the *Dictionary for Information Systems*, ANSI X3.172-1990.

address unit: depending on context, either 1) the units into which a Forth address space is divided for the purposes of locating data objects such as characters and variables; 2) the physical memory storage elements corresponding to those units; 3) the contents of such a memory storage element; or 4) the units in which the length of a region of memory is expressed.

aligned address: the address of a memory location at which a character, cell, cell pair, or double-cell integer can be accessed.

ambiguous condition: a circumstance for which this Standard does not prescribe a specific behavior for Forth systems and programs.

Ambiguous conditions include such things as the absence of a needed delimiter while parsing, attempted access to a nonexistent file, or attempted use of a nonexistent word. An ambiguous condition also exists when a Standard word is passed values that are improper or out of range.

cell: the primary unit of information in the architecture of a Forth system.

cell pair: two cells that are treated as a single unit.

character: depending on context, either 1) a storage unit capable of holding a character; or 2) a member of a character set.

character-aligned address: the address of a memory location at which a character can be accessed.

character string: data space that is associated with a sequence of consecutive character-aligned addresses. Character strings usually contain text. Unless otherwise indicated, the term “string” means “character string.”

code space: the logical area of the dictionary in which word semantics are implemented.

compile: to transform source code into dictionary definitions.

compilation semantics: the behavior of a Forth definition when its name is encountered by the text interpreter in compilation state.

counted string: a data structure consisting of one character containing a length followed by zero or more contiguous data characters. Normally, counted strings contain text.

cross compiler: a system that compiles a program for later execution in an environment that may be physically and logically different from the compiling environment. In a cross compiler, the term “host” applies to the compiling environment, and the term “target” applies to the run-time environment.

current definition: the definition whose compilation has been started but not yet ended.

data field: the data space associated with a word defined via **CREATE**.

data space: the logical area of the dictionary that can be accessed.

data-space pointer: the address of the next available data space location, i.e., the value returned by **HERE**.

data stack: a stack that may be used for passing parameters between definitions. When there is no possibility of confusion, the data stack is referred to as “the stack”. Contrast with **return stack**.

data type: an identifier for the set of values that a data object may have.

defining word: a Forth word that creates a new definition when executed.

definition: a Forth execution procedure compiled into the dictionary.

dictionary: an extensible structure that contains definitions and associated data space.

display: to send one or more characters to the user output device.

environmental dependencies: a program’s implicit assumptions about a Forth system’s implementation options or underlying hardware. For example, a program that assumes a cell size greater than 16 bits is said to have an environmental dependency.

execution semantics: the behavior of a Forth definition when it is executed.

execution token: a value that identifies the execution semantics of a definition.

find: to search the dictionary for a definition name matching a given string.

immediate word: a Forth word whose compilation semantics are to perform its execution semantics.

implementation defined: denotes system behaviors or features that must be provided and documented by a system but whose further details are not prescribed by this Standard.

implementation dependent: denotes system behaviors or features that must be provided by a system but whose further details are not prescribed by this Standard.

input buffer: a region of memory containing the sequence of characters from the input source that is currently accessible to a program.

input source: the device, file, block, or other entity that supplies characters to refill the input buffer.

input source specification: a set of information describing a particular state of the input source, input buffer, and parse area. This information is sufficient, when saved and restored properly, to enable the nesting of parsing operations on the same or different input sources.

interpretation semantics: the behavior of a Forth definition when its name is encountered by the text interpreter in interpretation state.

keyboard event: a value received by the system denoting a user action at the user input device. The term “keyboard” in this document does not exclude other types of user input devices.

line: a sequence of characters followed by an actual or implied line terminator.

name space: the logical area of the dictionary in which definition names are stored.

number: in this Standard, “number” used without other qualification means “integer.” Similarly, “double number” means “double-cell integer.”

parse: to select and exclude a character string from the parse area using a specified set of delimiting characters, called delimiters.

parse area: the portion of the input buffer that has not yet been parsed, and is thus available to the system for subsequent processing by the text interpreter and other parsing operations.

pictured-numeric output: a number display format in which the number is converted using Forth words that resemble a symbolic “picture” of the desired output.

program: a complete specification of execution to achieve a specific function (application task) expressed in Forth source code form.

receive: to obtain characters from the user input device.

return stack: a stack that may be used for program execution nesting, do-loop execution, temporary storage, and other purposes.

standard word: a named Forth procedure, formally specified in this Standard.

user input device: the input device currently selected as the source of received data, typically a keyboard.

user output device: the output device currently selected as the destination of display data.

variable: a named region of data space located and accessed by its memory address.

word: depending on context, either 1) the name of a Forth definition; or 2) a parsed sequence of non-space characters, which could be the name of a Forth definition.

word list: a list of associated Forth definition names that may be examined during a dictionary search.

word set: a set of Forth definitions grouped together in this Standard under a name indicating some shared aspect, typically their common functional area.

2.2 Notation

2.2.1 Numeric notation

Unless otherwise stated, all references to numbers apply to signed single-cell integers. The inclusive range of values is shown as *{from...to}*. The allowable range for the contents of an address is shown in double braces, particularly for the contents of variables, e.g., **BASE** *{{2...36}}*.

2.2.2 Stack notation

Stack parameters input to and output from a definition are described using the notation:

(stack-id before -- after)

where *stack-id* specifies which stack is being described, *before* represents the stack-parameter data types before execution of the definition and *after* represents them after execution. The symbols used in *before* and *after* are shown in table 3.2.

The control-flow-stack *stack-id* is “C:”, the data-stack *stack-id* is “S:”, and the return-stack *stack-id* is “R:”. When there is no confusion, the data-stack *stack-id* may be omitted.

When there are alternate *after* representations, they are described by “*after₁ | after₂*”. The top of the stack is stack is to the right. Only those stack items required for or provided by execution of the definition are shown.

2.2.3 Parsed-text notation

If, in addition to using stack parameters, a definition parses text, that text is specified by an abbreviation from table 2.1 shown surrounded by double-quotes and placed between the *before* parameters and the “--” separator in the first stack described, e.g.,

(S: before “*parsed-text-abbreviation*” -- after).

Table 2.1 – Parsed text abbreviations

Abbreviation	Description
< <i>char</i> >	the delimiting character marking the end of the string being parsed
< <i>chars</i> >	zero or more consecutive occurrences of the character <i>char</i>
< <i>space</i> >	a delimiting space character
< <i>spaces</i> >	zero or more consecutive occurrences of the character <i>space</i>
< <i>quote</i> >	a delimiting double quote
< <i>paren</i> >	a delimiting right parenthesis
< <i>eol</i> >	an implied delimiter marking the end of a line
<i>cc</i>	a parsed sequence of arbitrary characters, excluding the delimiter character
<i>name</i>	a token delimited by space, equivalent to <i>cc</i> < <i>space</i> > or <i>cc</i> < <i>eol</i> >

2.2.4 Glossary notation

The glossary entries for each word set are listed in the standard ASCII collating sequence. Each glossary entry specifies an ANS Forth word and consists of two parts: an *index line* and the *semantic description* of the definition.

2.2.4.1 Glossary index line

The index line is a single-line entry containing, from left to right:

- Section number, the last four digits of which assign a unique sequential number to all words included in this Standard;
- **DEFINITION-NAME** in upper-case, mono-spaced, bold-face letters;
- Natural-language pronunciation in quotes if it differs from English;
- Word-set designator from table 2.2. The designation for extensions word sets includes “EXT”.

Table 2.2 – Word set designators

Word set	Designator
Core word set	CORE
Block word set	BLOCK
Double-Number word set	DOUBLE
Exception word set	EXCEPTION
Facility word set	FACILITY
File-Access word set	FILE
Floating-Point word set	FLOATING
Locals word set	LOCALS
Memory-Allocation word set	MEMORY
Programming-Tools word set	TOOLS
Search-Order word set	SEARCH
String-Handling word set	STRING

2.2.4.2 Glossary semantic description

The first paragraph of the semantic description contains a stack notation for each stack affected by execution of the word. The remaining paragraphs contain a text description of the semantics. See **3.4.3 Semantics**.

3. Usage requirements

A system shall provide all of the words defined in **6.1 Core Words**. It may also provide any words defined in the optional word sets and extensions word sets. No standard word provided by a system shall alter the system state in a way that changes the effect of execution of any other standard word except as provided in this Standard. A system may contain non-standard extensions, provided that they are consistent with the requirements of this Standard.

The implementation of a system may use words and techniques outside the scope of this Standard.

A system need not provide all words in executable form. The implementation may provide definitions, including definitions of words in the Core word set, in source form only. If so, the mechanism for adding the definitions to the dictionary is implementation defined.

A program that requires a system to provide words or techniques not defined in this Standard has an environmental dependency.

3.1 Data types

A data type identifies the set of permissible values for a data object. It is not a property of a particular storage location or position on a stack. Moving a data object shall not affect its type.

No data-type checking is required of a system. An ambiguous condition exists if an incorrectly typed data object is encountered.

Table 3.1 summarizes the data types used throughout this Standard. Multiple instances of the same type in the description of a definition are suffixed with a sequence digit subscript to distinguish them.

Table 3.1 – Data types

Symbol	Data type	Size on stack
<i>flag</i>	flag	1 cell
<i>true</i>	true flag	1 cell
<i>false</i>	false flag	1 cell
<i>char</i>	character	1 cell
<i>n</i>	signed number	1 cell
<i>+n</i>	non-negative number	1 cell
<i>u</i>	unsigned number	1 cell
<i>n u</i> ¹	number	1 cell
<i>x</i>	unspecified cell	1 cell
<i>xt</i>	execution token	1 cell
<i>addr</i>	address	1 cell
<i>a-addr</i>	aligned address	1 cell
<i>c-addr</i>	character-aligned address	1 cell
<i>d</i>	double-cell signed number	2 cells
<i>+d</i>	double-cell non-negative number	2 cells
<i>ud</i>	double-cell unsigned number	2 cells
<i>d ud</i> ²	double-cell number	2 cells
<i>xd</i>	unspecified cell pair	2 cells
<i>colon-sys</i>	definition compilation	implementation dependent
<i>do-sys</i>	do-loop structures	implementation dependent
<i>case-sys</i>	CASE structures	implementation dependent
<i>of-sys</i>	OF structures	implementation dependent
<i>orig</i>	control-flow origins	implementation dependent
<i>dest</i>	control-flow destinations	implementation dependent
<i>loop-sys</i>	loop-control parameters	implementation dependent

<i>nest-sys</i>	definition calls	implementation dependent
<i>i*x, j*x</i> ³	any data type	0 or more cells

- ¹ May be either a signed number or an unsigned number depending on context.
- ² May be either a double-cell signed number or a double-cell unsigned number depending on context.
- ³ May be an undetermined number of stack entries of unspecified type. For examples of use, see **6.1.1370 EXECUTE**, **6.1.2050 QUIT**.

3.1.1 Data-type relationships

Some of the data types are subtypes of other data types. A data type *i* is a subtype of type *j* if and only if the members of *i* are a subset of the members of *j*. The following list represents the subtype relationships using the phrase “*i* => *j*” to denote “*i* is a subtype of *j*.” The subtype relationship is transitive; if *i* => *j* and *j* => *k* then *i* => *k*:

```
+n => u => x;
+n => n => x;
char => +n;
a-addr => c-addr => addr => u;
flag => x;
xt => x;
+d => d => xd;
+d => ud => xd.
```

Any Forth definition that accepts an argument of type *i* shall also accept an argument that is a subtype of *i*.

3.1.2 Character types

Characters shall be at least one address unit wide, contain at least eight bits, and have a size less than or equal to cell size.

The characters provided by a system shall include the graphic characters {32..126}, which represent graphic forms as shown in table 3.2.

3.1.2.1 Graphic characters

A graphic character is one that is normally displayed (e.g., A, #, &, 6). These values and graphics, shown in table 3.2, are taken directly from ANS X3.4-1974 (ASCII) and ISO-646-1983, International Reference Version (IRV). The graphic forms of characters outside the hex range {20..7E} are implementation-defined. Programs that use the graphic hex 24 (the currency sign) have an environmental dependency.

The graphic representation of characters is not restricted to particular type fonts or styles. The graphics here are examples.

Table 3.2 – Standard graphic characters

Hex	IRV	ASCII	Hex	IRV	ASCII	Hex	IRV	ASCII	Hex	IRV	ASCII	Hex	IRV	ASCII	Hex	IRV	ASCII
20			30	0	0	40	@	@	50	P	P	60	`	`	70	p	p
21	!	!	31	1	1	41	A	A	51	Q	Q	61	a	a	71	q	q
22	"	"	32	2	2	42	B	B	52	R	R	62	b	b	72	r	r
23	#	#	33	3	3	43	C	C	53	S	S	63	c	c	73	s	s
24	□	\$	34	4	4	44	D	D	54	T	T	64	d	d	74	t	t
25	%	%	35	5	5	45	E	E	55	U	U	65	e	e	75	u	u
26	&	&	36	6	6	46	F	F	56	V	V	66	f	f	76	v	v
27	'	'	37	7	7	47	G	G	57	W	W	67	g	g	77	w	w
28	((38	8	8	48	H	H	58	X	X	68	h	h	78	x	x
29))	39	9	9	49	I	I	59	Y	Y	69	i	i	79	y	y
2A	*	*	3A	:	:	4A	J	J	5A	Z	Z	6A	j	j	7A	z	z
2B	+	+	3B	;	;	4B	K	K	5B	[[6B	k	k	7B	{	{
2C	,	,	3C	<	<	4C	L	L	5C	\	\	6C	l	l	7C		
2D	-	-	3D	=	=	4D	M	M	5D]]	6D	m	m	7D	}	}
2E	.	.	3E	>	>	4E	N	N	5E	^	^	6E	n	n	7E	~	~
2F	/	/	3F	?	?	4F	O	O	5F	_	_	6F	o	o			

3.1.2.2 Control characters

All non-graphic characters included in the implementation-defined character set are defined in this Standard as control characters. In particular, the characters {0..31}, which could be included in the implementation-defined character set, are control characters.

Programs that require the ability to send or receive control characters have an environmental dependency.

3.1.3 Single-cell types

The implementation-defined fixed size of a cell is specified in address units and the corresponding number of bits. See **E.2 Hardware peculiarities**.

Cells shall be at least one address unit wide and contain at least sixteen bits. The size of a cell shall be an integral multiple of the size of a character. Data-stack elements, return-stack elements, addresses, execution tokens, flags, and integers are one cell wide.

3.1.3.1 Flags

Flags may have one of two logical states, *true* or *false*. Programs that use flags as arithmetic operands have an environmental dependency.

A true flag returned by a standard word shall be a single-cell value with all bits set. A false flag returned by a standard word shall be a single-cell value with all bits clear.

3.1.3.2 Integers

The implementation-defined range of signed integers shall include {-32767..+32767}.

The implementation-defined range of non-negative integers shall include {0..32767}.

The implementation-defined range of unsigned integers shall include {0..65535}.

3.1.3.3 Addresses

An address identifies a location in data space with a size of one address unit, which a program may fetch from or store into except for the restrictions established in this Standard. The size of an address unit is specified in bits. Each distinct address value identifies exactly one such storage element. See **3.3.3 Data space**.

The set of character-aligned addresses, addresses at which a character can be accessed, is an implementation-defined subset of all addresses. Adding the size of a character to a character-aligned address shall produce another character-aligned address.

The set of aligned addresses is an implementation-defined subset of character-aligned addresses. Adding the size of a cell to an aligned address shall produce another aligned address.

3.1.3.4 Counted strings

A counted string in memory is identified by the address (*c-addr*) of its length character.

The length character of a counted string shall contain a binary representation of the number of data characters, between zero and the implementation-defined maximum length for a counted string. The maximum length of a counted string shall be at least 255.

3.1.3.5 Execution tokens

Different definitions may have the same execution token if the definitions are equivalent.

3.1.4 Cell-pair types

A cell pair in memory consists of a sequence of two contiguous cells. The cell at the lower address is the first cell, and its address is used to identify the cell pair. Unless otherwise specified, a cell pair on a stack consists of the first cell immediately above the second cell.

3.1.4.1 Double-cell integers

On the stack, the cell containing the most significant part of a double-cell integer shall be above the cell containing the least significant part.

The implementation-defined range of double-cell signed integers shall include $\{-2147483647..+2147483647\}$.

The implementation-defined range of double-cell non-negative integers shall include $\{0..2147483647\}$.

The implementation-defined range of double-cell unsigned integers shall include $\{0..4294967295\}$. Placing the single-cell integer zero on the stack above a single-cell unsigned integer produces a double-cell unsigned integer with the same value. See **3.2.1.1 Internal number representation**.

3.1.4.2 Character strings

A string is specified by a cell pair (*c-addr u*) representing its starting address and length in characters.

3.1.5 System types

The system data types specify permitted word combinations during compilation and execution.

3.1.5.1 System-compilation types

These data types denote zero or more items on the control-flow stack (see 3.2.3.2). The possible presence of such items on the data stack means that any items already there shall be unavailable to a program until the control-flow-stack items are consumed.

The implementation-dependent data generated upon beginning to compile a definition and consumed at its close is represented by the symbol *colon-sys* throughout this Standard.

The implementation-dependent data generated upon beginning to compile a do-loop structure such as **DO ... LOOP** and consumed at its close is represented by the symbol *do-sys* throughout this Standard.

The implementation-dependent data generated upon beginning to compile a **CASE ... ENDCASE** structure and consumed at its close is represented by the symbol *case-sys* throughout this Standard.

The implementation-dependent data generated upon beginning to compile an **OF ... ENDOF** structure and consumed at its close is represented by the symbol *of-sys* throughout this Standard.

The implementation-dependent data generated and consumed by executing the other standard control-flow words is represented by the symbols *orig* and *dest* throughout this Standard.

3.1.5.2 System-execution types

These data types denote zero or more items on the return stack. Their possible presence means that any items already on the return stack shall be unavailable to a program until the system-execution items are consumed.

The implementation-dependent data generated upon beginning to execute a definition and consumed upon exiting it is represented by the symbol *nest-sys* throughout this Standard.

The implementation-dependent loop-control parameters used to control the execution of do-loops are represented by the symbol *loop-sys* throughout this Standard. Loop-control parameters shall be available inside the do-loop for words that use or change these parameters, words such as **I**, **J**, **LEAVE** and **UNLOOP**.

3.2 The implementation environment

3.2.1 Numbers

3.2.1.1 Internal number representation

This Standard allows one's complement, two's complement, or sign-magnitude number representations and arithmetic. Arithmetic zero is represented as the value of a single cell with all bits clear.

The representation of a number as a compiled literal or in memory is implementation dependent.

3.2.1.2 Digit conversion

Numbers shall be represented externally by using characters from the standard character set.

Conversion between the internal and external forms of a digit shall behave as follows:

The value in **BASE** is the radix for number conversion. A digit has a value ranging from zero to one less than the contents of **BASE**. The digit with the value zero corresponds to the character "0". This representation of digits proceeds through the character set to the decimal value nine corresponding to the character "9". For digits beginning with the decimal value ten the graphic characters beginning with the character "A" are used. This correspondence continues up to and including the digit with the decimal value thirty-five which is represented by the character "Z". The conversion of digits outside this range is implementation defined.

3.2.1.3 Free-field number display

Free-field number display uses the characters described in digit conversion, without leading zeros, in a field the exact size of the converted string plus a trailing space. If a number is zero, the least significant digit is not considered a leading zero. If the number is negative, a leading minus sign is displayed.

Number display may use the pictured numeric output string buffer to hold partially converted strings (see **3.3.3.6 Other transient regions**).

3.2.2 Arithmetic

3.2.2.1 Integer division

Division produces a quotient q and a remainder r by dividing operand a by operand b . Division operations return q , r , or both. The identity $b*q + r = a$ shall hold for all a and b .

When unsigned integers are divided and the remainder is not zero, q is the largest integer less than the true quotient.

When signed integers are divided, the remainder is not zero, and a and b have the same sign, q is the largest integer less than the true quotient. If only one operand is negative, whether q is rounded toward

negative infinity (floored division) or rounded towards zero (symmetric division) is implementation defined.

Floored division is integer division in which the remainder carries the sign of the divisor or is zero, and the quotient is rounded to its arithmetic floor. Symmetric division is integer division in which the remainder carries the sign of the dividend or is zero and the quotient is the mathematical quotient “rounded towards zero” or “truncated.” Examples of each are shown in tables 3.3 and 3.4.

In cases where the operands differ in sign and the rounding direction matters, a program shall either include code generating the desired form of division, not relying on the implementation-defined default result, or have an environmental dependency on the desired rounding direction.

Table 3.3 – Floored Division Example

Dividend	Divisor	Remainder	Quotient
10	7	3	1
-10	7	4	-2
10	-7	-4	-2
-10	-7	-3	1

Table 3.4 – Symmetric Division Example

Dividend	Divisor	Remainder	Quotient
10	7	3	1
-10	7	-3	-1
10	-7	3	-1
-10	-7	-3	1

3.2.2.2 Other integer operations

In all integer arithmetic operations, both overflow and underflow shall be ignored. The value returned when either overflow or underflow occurs is implementation defined.

3.2.3 Stacks

3.2.3.1 Data stack

Objects on the data stack shall be one cell wide.

3.2.3.2 Control-flow stack

The control-flow stack is a last-in, first out list whose elements define the permissible matchings of control-flow words and the restrictions imposed on data-stack usage during the compilation of control structures.

The elements of the control-flow stack are system-compilation data types.

The control-flow stack may, but need not, physically exist in an implementation. If it does exist, it may be, but need not be, implemented using the data stack. The format of the control-flow stack is implementation defined. Since the control-flow stack may be implemented using the data stack, items placed on the data stack are unavailable to a program after items are placed on the control-flow stack and remain unavailable until the control-flow stack items are removed.

3.2.3.3 Return stack

Items on the return stack shall consist of one or more cells. A system may use the return stack in an implementation-dependent manner during the compilation of definitions, during the execution of do-loops, and for storing run-time nesting information.

A program may use the return stack for temporary storage during the execution of a definition subject to the following restrictions:

- A program shall not access values on the return stack (using **R@**, **R>**, **2R@** or **2R>**) that it did not place there using **>R** or **2>R**;
- A program shall not access from within a do-loop values placed on the return stack before the loop was entered;

- All values placed on the return stack within a do-loop shall be removed before **I**, **J**, **LOOP**, **+LOOP**, **UNLOOP**, or **LEAVE** is executed;
- All values placed on the return stack within a definition shall be removed before the definition is terminated or before **EXIT** is executed.

3.2.4 Operator terminal

See 1.2.2 Exclusions.

3.2.4.1 User input device

The method of selecting the user input device is implementation defined.

The method of indicating the end of an input line of text is implementation defined.

3.2.4.2 User output device

The method of selecting the user output device is implementation defined.

3.2.5 Mass storage

A system need not provide any standard words for accessing mass storage. If a system provides any standard word for accessing mass storage, it shall also implement the Block word set.

3.2.6 Environmental queries

The name spaces for **ENVIRONMENT?** and definitions are disjoint. Names of definitions that are the same as **ENVIRONMENT?** strings shall not impair the operation of **ENVIRONMENT?**. Table 3.5 contains the valid input strings and corresponding returned value for inquiring about the programming environment with **ENVIRONMENT?**.

Table 3.5 – Environmental Query Strings

String	Value data type	Constant?	Meaning
/COUNTED-STRING	<i>n</i>	yes	maximum size of a counted string, in characters
/HOLD	<i>n</i>	yes	size of the pictured numeric output string buffer, in characters
/PAD	<i>n</i>	yes	size of the scratch area pointed to by PAD , in characters
ADDRESS-UNIT-BITS	<i>n</i>	yes	size of one address unit, in bits
CORE	<i>flag</i>	no	true if complete core word set present (i.e., not a subset as defined in 5.1.1)
CORE-EXT	<i>flag</i>	no	true if core extensions word set present
FLOORED	<i>flag</i>	yes	true if floored division is the default
MAX-CHAR	<i>u</i>	yes	maximum value of any character in the implementation-defined character set
MAX-D	<i>d</i>	yes	largest usable signed double number
MAX-N	<i>n</i>	yes	largest usable signed integer
MAX-U	<i>u</i>	yes	largest usable unsigned integer
MAX-UD	<i>ud</i>	yes	largest usable unsigned double number
RETURN-STACK-CELLS	<i>n</i>	yes	maximum size of the return stack, in cells
STACK-CELLS	<i>n</i>	yes	maximum size of the data stack, in cells

If an environmental query (using **ENVIRONMENT?**) returns *false* (i.e., unknown) in response to a string, subsequent queries using the same string may return *true*. If a query returns *true* (i.e., known) in response

to a string, subsequent queries with the same string shall also return *true*. If a query designated as constant in the above table returns *true* and a value in response to a string, subsequent queries with the same string shall return *true* and the same value.

3.3 The Forth dictionary

Forth words are organized into a structure called the dictionary. While the form of this structure is not specified by the Standard, it can be described as consisting of three logical parts: a name space, a code space, and a data space. The logical separation of these parts does not require their physical separation.

A program shall not fetch from or store into locations outside data space. An ambiguous condition exists if a program addresses name space or code space.

3.3.1 Name space

The relationship between name space and data space is implementation dependent.

3.3.1.1 Word lists

The structure of a word list is implementation dependent. When duplicate names exist in a word list, the latest-defined duplicate shall be the one found during a search for the name.

3.3.1.2 Definition names

Definition names shall contain {1..31} characters. A system may allow or prohibit the creation of definition names containing non-standard characters.

Programs that use lower case for standard definition names or depend on the case-sensitivity properties of a system have an environmental dependency.

A program shall not create definition names containing non-graphic characters.

3.3.2 Code space

The relationship between code space and data space is implementation dependent.

3.3.3 Data space

Data space is the only logical area of the dictionary for which standard words are provided to allocate and access regions of memory. These regions are: contiguous regions, variables, text-literal regions, input buffers, and other transient regions, each of which is described in the following sections. A program may read from or write into these regions unless otherwise specified.

3.3.3.1 Address alignment

Most addresses used in ANS Forth are aligned addresses (indicated by *a-addr*) or character-aligned (indicated by *c-addr*). **ALIGNED**, **CHAR+**, and arithmetic operations can alter the alignment state of an address on the stack. **CHAR+** applied to an aligned address returns a character-aligned address that can only be used to access characters. Applying **CHAR+** to a character-aligned address produces the succeeding character-aligned address. Adding or subtracting an arbitrary number to an address can produce an unaligned address that shall not be used to fetch or store anything. The only way to find the next aligned address is with **ALIGNED**. An ambiguous condition exists when **@**, **!**, **,** (comma), **+!**, **2@**, or **2!** is used with an address that is not aligned, or when **C@**, **C!**, or **C,** is used with an address that is not character-aligned.

The definitions of **6.1.1000 CREATE** and **6.1.2410 VARIABLE** require that the definitions created by them return aligned addresses.

After definitions are compiled or the word **ALIGN** is executed the data-space pointer is guaranteed to be aligned.

3.3.3.2 Contiguous regions

A system guarantees that a region of data space allocated using **ALLOT**, **,** (comma), **C**, (c-comma), and **ALIGN** shall be contiguous with the last region allocated with one of the above words, unless the restrictions in the following paragraphs apply. The data-space pointer **HERE** always identifies the beginning of the next data-space region to be allocated. As successive allocations are made, the data-space pointer increases. A program may perform address arithmetic within contiguously allocated regions. The last region of data space allocated using the above operators may be released by allocating a corresponding negatively-sized region using **ALLOT**, subject to the restrictions of the following paragraphs.

CREATE establishes the beginning of a contiguous region of data space, whose starting address is returned by the **CREATED** definition. This region is terminated by compiling the next definition.

Since an implementation is free to allocate data space for use by code, the above operators need not produce contiguous regions of data space if definitions are added to or removed from the dictionary between allocations. An ambiguous condition exists if deallocated memory contains definitions.

3.3.3.3 Variables

The region allocated for a variable may be non-contiguous with regions subsequently allocated with **,** (comma) or **ALLOT**. For example, in:

```
VARIABLE X 1 CELLS ALLOT
```

the region **X** and the region **ALLOT**ted could be non-contiguous.

Some system-provided variables, such as **STATE**, are restricted to read-only access.

3.3.3.4 Text-literal regions

The text-literal regions, specified by strings compiled with **S**" and **C**", may be read-only.

A program shall not store into the text-literal regions created by **S**" and **C**" nor into any read-only system variable or read-only transient regions. An ambiguous condition exists when a program attempts to store into read-only regions.

3.3.3.5 Input buffers

The address, length, and content of the input buffer may be transient. A program shall not write into the input buffer. In the absence of any optional word sets providing alternative input sources, the input buffer is either the terminal-input buffer, used by **QUIT** to hold one line from the user input device, or a buffer specified by **EVALUATE**. In all cases, **SOURCE** returns the beginning address and length in characters of the current input buffer.

The minimum size of the terminal-input buffer shall be 80 characters.

The address and length returned by **SOURCE**, the string returned by **PARSE**, and directly computed input-buffer addresses are valid only until the text interpreter does I/O to refill the input buffer or the input source is changed.

A program may modify the size of the parse area by changing the contents of **>IN** within the limits imposed by this Standard. For example, if the contents of **>IN** are saved before a parsing operation and restored afterwards, the text that was parsed will be available again for subsequent parsing operations. The extent of permissible repositioning using this method depends on the input source (see **7.3.3 Block buffer regions** and **11.3.4 Input source**).

A program may directly examine the input buffer using its address and length as returned by **SOURCE**; the beginning of the parse area within the input buffer is indexed by the number in **>IN**. The values are valid for a limited time. An ambiguous condition exists if a program modifies the contents of the input buffer.

3.3.3.6 Other transient regions

The data space regions identified by **PAD**, **WORD**, and **#>** (the pictured numeric output string buffer) may be transient. Their addresses and contents may become invalid after:

- a definition is created via a defining word;
- definitions are compiled with **:** or **:NONAME**;
- data space is allocated using **ALLOT**, **,** (comma), **C**, (c-comma), or **ALIGN**.

The previous contents of the regions identified by **WORD** and **#>** may be invalid after each use of these words. Further, the regions returned by **WORD** and **#>** may overlap in memory. Consequently, use of one of these words can corrupt a region returned earlier by a different word. The other words that construct pictured numeric output strings (**<#**, **#**, **#S**, and **HOLD**) may also modify the contents of these regions. Words that display numbers may be implemented using pictured numeric output words. Consequently, **.** (dot), **.R**, **.S**, **?**, **D.**, **D.R**, **U.**, and **U.R** could also corrupt the regions.

The size of the scratch area whose address is returned by **PAD** shall be at least 84 characters. The contents of the region addressed by **PAD** are intended to be under the complete control of the user: no words defined in this Standard place anything in the region, although changing data-space allocations as described in **3.3.3.2 Contiguous regions** may change the address returned by **PAD**. Non-standard words provided by an implementation may use **PAD**, but such use shall be documented.

The size of the region identified by **WORD** shall be at least 33 characters.

The size of the pictured numeric output string buffer shall be at least $(2*n) + 2$ characters, where n is the number of bits in a cell. Programs that consider it a fixed area with unchanging access parameters have an environmental dependency.

3.4 The Forth text interpreter

Upon start-up, a system shall be able to interpret, as described by **6.1.2050 QUIT**, Forth source code received interactively from a user input device.

Such interactive systems usually furnish a "prompt" indicating that they have accepted a user request and acted on it. The implementation-defined Forth prompt should contain the word "OK" in some combination of upper or lower case.

Text interpretation (see **6.1.1360 EVALUATE** and **6.1.2050 QUIT**) shall repeat the following steps until either the parse area is empty or an ambiguous condition exists:

- a) Skip leading spaces and parse a *name* (see **3.4.1**);
- b) Search the dictionary name space (see **3.4.2**). If a definition name matching the string is found:
 - 1) if interpreting, perform the interpretation semantics of the definition (see **3.4.3.2**), and continue at a);
 - 2) if compiling, perform the compilation semantics of the definition (see **3.4.3.3**), and continue at a).
- c) If a definition name matching the string is not found, attempt to convert the string to a number (see **3.4.1.3**). If successful:
 - 1) if interpreting, place the number on the data stack, and continue at a);
 - 2) if compiling, compile code that when executed will place the number on the stack (see **6.1.1780 LITERAL**), and continue at a);
- d) If unsuccessful, an ambiguous condition exists (see **3.4.4**).

3.4.1 Parsing

Unless otherwise noted, the number of characters parsed may be from zero to the implementation-defined maximum length of a counted string.

If the parse area is empty, i.e., when the number in **>IN** is equal to the length of the input buffer, or contains no characters other than delimiters, the selected string is empty. Otherwise, the selected string begins with the next character in the parse area, which is the character indexed by the contents of **>IN**. An ambiguous condition exists if the number in **>IN** is greater than the size of the input buffer.

If delimiter characters are present in the parse area after the beginning of the selected string, the string continues up to and including the character just before the first such delimiter, and the number in **>IN** is changed to index immediately past that delimiter, thus removing the parsed characters and the delimiter from the parse area. Otherwise, the string continues up to and including the last character in the parse area, and the number in **>IN** is changed to the length of the input buffer, thus emptying the parse area.

Parsing may change the contents of **>IN**, but shall not affect the contents of the input buffer. Specifically, if the value in **>IN** is saved before starting the parse, resetting **>IN** to that value immediately after the parse shall restore the parse area without loss of data.

3.4.1.1 Delimiters

If the delimiter is the space character, hex 20 (**BL**), control characters may be treated as delimiters. The set of conditions, if any, under which a “space” delimiter matches control characters is implementation defined.

To skip leading delimiters is to pass by zero or more contiguous delimiters in the parse area before parsing.

3.4.1.2 Syntax

Forth has a simple, operator-ordered syntax. The phrase **A B C** returns values as if **A** were executed first, then **B** and finally **C**. Words that cause deviations from this linear flow of control are called control-flow words. Combinations of control-flow words whose stack effects are compatible form control-flow structures. Examples of typical use are given for each control-flow word in Annex A.

Forth syntax is extensible; for example, new control-flow words can be defined in terms of existing ones.

This Standard does not require a syntax or program-construct checker.

3.4.1.3 Text interpreter input number conversion

When converting input numbers, the text interpreter shall recognize both positive and negative numbers, with a negative number represented by a single minus sign, the character “-”, preceding the digits. The value in **BASE** is the radix for number conversion.

3.4.2 Finding definition names

A string matches a definition name if each character in the string matches the corresponding character in the string used as the definition name when the definition was created. The case sensitivity (whether or not the upper-case letters match the lower-case letters) is implementation defined. A system may be either case sensitive, treating upper- and lower-case letters as different and not matching, or case insensitive, ignoring differences in case while searching.

The matching of upper- and lower-case letters with alphabetic characters in character set extensions such as accented international characters is implementation defined.

A system shall be capable of finding the definition names defined by this Standard when they are spelled with upper-case letters.

3.4.3 Semantics

The semantics of a Forth definition are implemented by machine code or a sequence of execution tokens or other representations. They are largely specified by the stack notation in the glossary entries, which

shows what values shall be consumed and produced. The prose in each glossary entry further specifies the definition's behavior.

Each Forth definition may have several behaviors, described in the following sections. The terms "initiation semantics" and "run-time semantics" refer to definition fragments, and have meaning only within the individual glossary entries where they appear.

3.4.3.1 Execution semantics

The execution semantics of each Forth definition are specified in an "Execution:" section of its glossary entry. When a definition has only one specified behavior, the label is omitted.

Execution may occur implicitly, when the definition into which it has been compiled is executed, or explicitly, when its execution token is passed to **EXECUTE**. The execution semantics of a syntactically correct definition under conditions other than those specified in this Standard are implementation dependent.

Glossary entries for defining words include the execution semantics for the new definition in a "*name* Execution:" section.

3.4.3.2 Interpretation semantics

Unless otherwise specified in an "Interpretation:" section of the glossary entry, the interpretation semantics of a Forth definition are its execution semantics.

A system shall be capable of executing, in interpretation state, all of the definitions from the Core word set and any definitions included from the optional word sets or word set extensions whose interpretation semantics are defined by this Standard.

A system shall be capable of executing, in interpretation state, any new definitions created in accordance with **3. Usage requirements**.

3.4.3.3 Compilation semantics

Unless otherwise specified in a "Compilation:" section of the glossary entry, the compilation semantics of a Forth definition shall be to append its execution semantics to the execution semantics of the current definition.

3.4.4 Possible actions on an ambiguous condition

When an ambiguous condition exists, a system may take one or more of the following actions:

- ignore and continue;
- display a message;
- execute a particular word;
- set interpretation state and begin text interpretation;
- take other implementation-defined actions;
- take implementation-dependent actions.

The response to a particular ambiguous condition need not be the same under all circumstances.

3.4.5 Compilation

A program shall not attempt to nest compilation of definitions.

During the compilation of the current definition, a program shall not execute any defining word, **:NONAME**, or any definition that allocates dictionary data space. The compilation of the current definition may be suspended using [(left-bracket) and resumed using] (right-bracket). While the compilation of

the current definition is suspended, a program shall not execute any defining word, **:NONAME**, or any definition that allocates dictionary data space.

4. Documentation requirements

When it is impossible or infeasible for a system or program to define a particular behavior itself, it is permissible to state that the behavior is unspecifiable and to explain the circumstances and reasons why this is so.

4.1 System documentation

4.1.1 Implementation-defined options

The implementation-defined items in the following list represent characteristics and choices left to the discretion of the implementor, provided that the requirements of this Standard are met. A system shall document the values for, or behaviors of, each item.

- aligned address requirements (**3.1.3.3 Addresses**);
 - behavior of **6.1.1320 EMIT** for non-graphic characters;
 - character editing of **6.1.0695 ACCEPT** and **6.2.1390 EXPECT**;
 - character set (**3.1.2 Character types**, **6.1.1320 EMIT**, **6.1.1750 KEY**);
 - character-aligned address requirements (**3.1.3.3 Addresses**);
 - character-set-extensions matching characteristics (**3.4.2 Finding definition names**);
 - conditions under which control characters match a space delimiter (**3.4.1.1 Delimiters**);
 - format of the control-flow stack (**3.2.3.2 Control-flow stack**);
 - conversion of digits larger than thirty-five (**3.2.1.2 Digit conversion**);
 - display after input terminates in **6.1.0695 ACCEPT** and **6.2.1390 EXPECT**;
 - exception abort sequence (as in **6.1.0680 ABORT**");
 - input line terminator (**3.2.4.1 User input device**);
 - maximum size of a counted string, in characters (**3.1.3.4 Counted strings**, **6.1.2450 WORD**);
 - maximum size of a parsed string (**3.4.1 Parsing**);
 - maximum size of a definition name, in characters (**3.3.1.2 Definition names**);
 - maximum string length for **6.1.1345 ENVIRONMENT?**, in characters;
 - method of selecting **3.2.4.1 User input device**;
 - method of selecting **3.2.4.2 User output device**;
 - methods of dictionary compilation (**3.3 The Forth dictionary**);
 - number of bits in one address unit (**3.1.3.3 Addresses**);
 - number representation and arithmetic (**3.2.1.1 Internal number representation**);
 - ranges for n, +n, u, d, +d, ud (**3.1.3 Single-cell types**, **3.1.4 Cell-pair types**);
 - read-only data-space regions (**3.3.3 Data space**);
 - size of buffer at **6.1.2450 WORD** (**3.3.3.6 Other transient regions**);
 - size of one cell in address units (**3.1.3 Single-cell types**);
 - size of one character in address units (**3.1.2 Character types**);
 - size of the keyboard terminal input buffer (**3.3.3.5 Input buffers**);
 - size of the pictured numeric output string buffer (**3.3.3.6 Other transient regions**);
 - size of the scratch area whose address is returned by **6.2.2000 PAD** (**3.3.3.6 Other transient regions**);
 - system case-sensitivity characteristics (**3.4.2 Finding definition names**);
 - system prompt (**3.4 The Forth text interpreter**, **6.1.2050 QUIT**);
 - type of division rounding (**3.2.2.1 Integer division**, **6.1.0100 */**, **6.1.0110 */MOD**, **6.1.0230 /**, **6.1.0240 /MOD**, **6.1.1890 MOD**);
 - values of **6.1.2250 STATE** when true;
 - values returned after arithmetic overflow (**3.2.2.2 Other integer operations**);
-

- whether the current definition can be found after **6.1.1250 DOES>** (**6.1.0450 :**).

4.1.2 Ambiguous conditions

A system shall document the system action taken upon each of the general or specific ambiguous conditions identified in this Standard. See **3.4.4 Possible actions on an ambiguous condition**.

The following specific ambiguous conditions, indicated by an asterisk, are noted in the glossary entries of the relevant words. The following general ambiguous conditions could occur because of a combination of factors:

- a *name* is neither a valid definition name nor a valid number during text interpretation (**3.4 The Forth text interpreter**);
- a definition name exceeded the maximum length allowed (**3.3.1.2 Definition names**);
- addressing a region not listed in **3.3.3 Data Space**;
- argument type incompatible with specified input parameter, e.g., passing a *flag* to a word expecting an *n* (**3.1 Data types**);
- attempting to obtain the execution token, (e.g., with **6.1.0070 ' , 6.1.1550 FIND**, etc.) of a definition with undefined interpretation semantics;
- dividing by zero (**6.1.0100 */ , 6.1.0110 */MOD , 6.1.0230 / , 6.1.0240 /MOD , 6.1.1561 FM/MOD , 6.1.1890 MOD , 6.1.2214 SM/REM , 6.1.2370 UM/MOD , 8.6.1.1820 M*/**);
- insufficient data-stack space or return-stack space (stack overflow);
- insufficient space for loop-control parameters;
- insufficient space in the dictionary;
- interpreting a word with undefined interpretation semantics;
- modifying the contents of the input buffer or a string literal (**3.3.3.4 Text-literal regions , 3.3.3.5 Input buffers**);
- overflow of a pictured numeric output string;
- parsed string overflow;
- producing a result out of range, e.g., multiplication (using *****) results in a value too big to be represented by a single-cell integer (**6.1.0090 * , 6.1.0100 */ , 6.1.0110 */MOD , 6.1.0570 >NUMBER , 6.1.1561 FM/MOD , 6.1.2214 SM/REM , 6.1.2370 UM/MOD , 6.2.0970 CONVERT , 8.6.1.1820 M*/**);
- reading from an empty data stack or return stack (stack underflow);
- unexpected end of input buffer, resulting in an attempt to use a zero-length string as a *name*;
- * **>IN** greater than size of input buffer (**3.4.1 Parsing**);
- * **6.1.2120 RECURSE** appears after **6.1.1250 DOES>**;
- * argument input source different than current input source for **6.2.2148 RESTORE-INPUT**;
- * data space containing definitions is de-allocated (**3.3.3.2 Contiguous regions**);
- * data space read/write with incorrect alignment (**3.3.3.1 Address alignment**);
- * data-space pointer not properly aligned (**6.1.0150 , , 6.1.0860 C ,**);
- * less than **u+2** stack items (**6.2.2030 PICK , 6.2.2150 ROLL**);
- * loop-control parameters not available (**6.1.0140 +LOOP , 6.1.1680 I , 6.1.1730 J , 6.1.1760 LEAVE , 6.1.1800 LOOP , 6.1.2380 UNLOOP**);
- * most recent definition does not have a *name* (**6.1.1710 IMMEDIATE**);
- * *name* not defined by **6.2.2405 VALUE** used by **6.2.2295 TO**;
- * *name* not found (**6.1.0070 ' , 6.1.2033 POSTPONE , 6.1.2510 ['] , 6.2.2530 [COMPILE]**);
- * parameters are not of the same type (**6.1.1240 DO , 6.2.0620 ?DO , 6.2.2440 WITHIN**);
- * **6.1.2033 POSTPONE** or **6.2.2530 [COMPILE]** applied to **6.2.2295 TO**;
- * string longer than a counted string returned by **6.1.2450 WORD**;
- * *u* greater than or equal to the number of bits in a cell (**6.1.1805 LSHIFT , 6.1.2162 RSHIFT**);

- * word not defined via **6.1.1000 CREATE (6.1.0550 >BODY, 6.1.1250 DOES>)**;
- * words improperly used outside **6.1.0490 <#** and **6.1.0040 #>** (**6.1.0030 #**, **6.1.0050 #S**, **6.1.1670 HOLD**, **6.1.2210 SIGN**).

4.1.3 Other system documentation

A system shall provide the following information:

- list of non-standard words using **6.2.2000 PAD (3.3.3.6 Other transient regions)**;
- operator's terminal facilities available;
- program data space available, in address units;
- return stack space available, in cells;
- stack space available, in cells;
- system dictionary space required, in address units.

4.2 Program documentation

4.2.1 Environmental dependencies

A program shall document the following environmental dependencies, where they apply, and should document other known environmental dependencies:

- considering the pictured numeric output string buffer a fixed area with unchanging access parameters (**3.3.3.6 Other transient regions**);
- depending on the presence or absence of non-graphic characters in a received string (**6.1.0695 ACCEPT**, **6.2.1390 EXPECT**);
- relying on a particular rounding direction (**3.2.2.1 Integer division**);
- requiring a particular number representation and arithmetic (**3.2.1.1 Internal number representation**);
- requiring non-standard words or techniques (**3. Usage requirements**);
- requiring the ability to send or receive control characters (**3.1.2.2 Control characters**, **6.1.1750 KEY**);
- using control characters to perform specific functions (**6.1.1320 EMIT**, **6.1.2310 TYPE**);
- using flags as arithmetic operands (**3.1.3.1 Flags**);
- using lower case for standard definition names or depending on the case sensitivity of a system (**3.3.1.2 Definition names**);
- using the graphic character with a value of hex 24 (**3.1.2.1 Graphic characters**).

4.2.2 Other program documentation

A program shall also document:

- minimum operator's terminal facilities required;
 - whether a Standard System exists after the program is loaded.
-

5. Compliance and labeling

5.1 ANS Forth systems

5.1.1 System compliance

A system that complies with all the system requirements given in sections **3. Usage requirements** and **4.1 System documentation** and their sub-sections is a Standard System. An otherwise Standard System that provides only a portion of the Core words is a Standard System Subset. An otherwise Standard System (Subset) that fails to comply with one or more of the minimum values or ranges specified in **3. Usage requirements** and its sub-sections has environmental restrictions.

5.1.2 System labeling

A Standard System (Subset) shall be labeled an “ANS Forth System (Subset).” That label, by itself, shall not be applied to Standard Systems or Standard System Subsets that have environmental restrictions.

The phrase “with Environmental Restrictions” shall be appended to the label of a Standard System (Subset) that has environmental restrictions.

The phrase “Providing *name(s)* from the Core Extensions word set” shall be appended to the label of any Standard System that provides portions of the Core Extensions word set.

The phrase “Providing the Core Extensions word set” shall be appended to the label of any Standard System that provides all of the Core Extensions word set.

5.2 ANS Forth programs

5.2.1 Program compliance

A program that complies with all the program requirements given in sections **3. Usage requirements** and **4.2 Program documentation** and their sub-sections is a Standard Program.

5.2.2 Program labeling

A Standard Program shall be labeled an “ANS Forth Program.” That label, by itself, shall not be applied to Standard Programs that require the system to provide standard words outside the Core word set or that have environmental dependencies.

The phrase “with Environmental Dependencies” shall be appended to the label of Standard Programs that have environmental dependencies.

The phrase “Requiring *name(s)* from the Core Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Core Extensions word set.

The phrase “Requiring the Core Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Core Extensions word set.

6. Glossary

6.1 Core words

- 6.1.0010 !** "store" CORE
- (*x a-addr --*)
- Store *x* at *a-addr*.
- See: **3.3.3.1 Address alignment**.
- 6.1.0030 #** "number-sign" CORE
- (*ud₁ -- ud₂*)
- Divide *ud₁* by the number in **BASE** giving the quotient *ud₂* and the remainder *n*. (*n* is the least-significant digit of *ud₁*.) Convert *n* to external form and add the resulting character to the beginning of the pictured numeric output string. An ambiguous condition exists if # executes outside of a <# #> delimited number conversion.
- See: **6.1.0040 #>**, **6.1.0050 #s**, **6.1.0490 <#**.
- 6.1.0040 #>** "number-sign-greater" CORE
- (*xd -- c-addr u*)
- Drop *xd*. Make the pictured numeric output string available as a character string. *c-addr* and *u* specify the resulting character string. A program may replace characters within the string.
- See: **6.1.0030 #**, **6.1.0050 #s**, **6.1.0490 <#**.
- 6.1.0050 #s** "number-sign-s" CORE
- (*ud₁ -- ud₂*)
- Convert one digit of *ud₁* according to the rule for #. Continue conversion until the quotient is zero. *ud₂* is zero. An ambiguous condition exists if #s executes outside of a <# #> delimited number conversion.
- See: **6.1.0030 #**, **6.1.0040 #>**, **6.1.0490 <#**.
- 6.1.0070 '** "tick" CORE
- ("*<spaces>name*" -- *xt*)
- Skip leading space delimiters. Parse *name* delimited by a space. Find *name* and return *xt*, the execution token for *name*. An ambiguous condition exists if *name* is not found.
- When interpreting, ' **name EXECUTE** is equivalent to **name**.
- See: **3.4 The Forth text interpreter**, **3.4.1 Parsing**, **A.6.1.2033 POSTPONE**, **A.6.1.2510 [']**, **D.6.7 Immediacy**.
-

6.1.0080 ("paren" CORE

Compilation: Perform the execution semantics given below.

Execution: ("ccc<paren>" --)

Parse *ccc* delimited by a right parenthesis ")". (is an immediate word.

The number of characters in *ccc* may be zero to the number of characters in the parse area.

See: **3.4.1 Parsing, 11.6.1.0080** (.

6.1.0090 * "star" CORE

($n_1|u_1$ $n_2|u_2$ -- $n_3|u_3$)

Multiply $n_1|u_1$ by $n_2|u_2$ giving the product $n_3|u_3$.

6.1.0100 */ "star-slash" CORE

(n_1 n_2 n_3 -- n_4)

Multiply n_1 by n_2 producing the double-cell intermediate result d . Divide d by n_3 giving the single-cell quotient n_4 . An ambiguous condition exists if n_3 is zero or if the quotient n_4 lies outside the range of a signed number. If d and n_3 differ in sign, the implementation-defined result returned will be the same as that returned by either the phrase **>R M* R> FM/MOD SWAP DROP** or the phrase **>R M* R> SM/REM SWAP DROP**.

See: **3.2.2.1 Integer division**.

6.1.0110 */MOD "star-slash-mod" CORE

(n_1 n_2 n_3 -- n_4 n_5)

Multiply n_1 by n_2 producing the intermediate double-cell result d . Divide d by n_3 producing the single-cell remainder n_4 and the single-cell quotient n_5 . An ambiguous condition exists if n_3 is zero, or if the quotient n_5 lies outside the range of a single-cell signed integer. If d and n_3 differ in sign, the implementation-defined result returned will be the same as that returned by either the phrase **>R M* R> FM/MOD** or the phrase **>R M* R> SM/REM**.

See: **3.2.2.1 Integer division**.

6.1.0120 + "plus" CORE

($n_1|u_1$ $n_2|u_2$ -- $n_3|u_3$)

Add $n_2|u_2$ to $n_1|u_1$, giving the sum $n_3|u_3$.

See: **3.3.3.1 Address alignment**.

- 6.1.0130** **+**! "plus-store" CORE
 ($n|u\ a\text{-}addr\ \text{--}$)
 Add $n|u$ to the single-cell number at $a\text{-}addr$.
 See: **3.3.3.1 Address alignment**.
- 6.1.0140** **+LOOP** "plus-loop" CORE
 Interpretation: Interpretation semantics for this word are undefined.
 Compilation: ($C: do\text{-}sys\ \text{--}$)
 Append the run-time semantics given below to the current definition. Resolve the destination of all unresolved occurrences of **LEAVE** between the location given by $do\text{-}sys$ and the next location for a transfer of control, to execute the words following **+LOOP**.
 Run-time: ($n\ \text{--}$) ($R: loop\text{-}sys_1\ \text{--} | loop\text{-}sys_2$)
 An ambiguous condition exists if the loop control parameters are unavailable. Add n to the loop index. If the loop index did not cross the boundary between the loop limit minus one and the loop limit, continue execution at the beginning of the loop. Otherwise, discard the current loop control parameters and continue execution immediately following the loop.
 See: **6.1.1240 DO**, **6.1.1680 I**, **6.1.1760 LEAVE**.
- 6.1.0150** **,** "comma" CORE
 ($x\ \text{--}$)
 Reserve one cell of data space and store x in the cell. If the data-space pointer is aligned when **,** begins execution, it will remain aligned when **,** finishes execution. An ambiguous condition exists if the data-space pointer is not aligned prior to execution of **,**.
 See: **3.3.3 Data space**, **3.3.3.1 Address alignment**.
- 6.1.0160** **-** "minus" CORE
 ($n_1|u_1\ n_2|u_2\ \text{--}\ n_3|u_3$)
 Subtract $n_2|u_2$ from $n_1|u_1$, giving the difference $n_3|u_3$.
 See: **3.3.3.1 Address alignment**.
- 6.1.0180** **.** "dot" CORE
 ($n\ \text{--}$)
 Display n in free field format.
 See: **3.2.1.2 Digit conversion**, **3.2.1.3 Free-field number display**.
-

6.1.0190 `."` "dot-quote" CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (`"ccc<quote>" --`)

Parse `ccc` delimited by `"` (double-quote). Append the run-time semantics given below to the current definition.

Run-time: (`--`)

Display `ccc`.

See: **3.4.1 Parsing**, **6.2.0200** `.(`

6.1.0230 `/` "slash" CORE

(`n1 n2 -- n3`)

Divide n_1 by n_2 , giving the single-cell quotient n_3 . An ambiguous condition exists if n_2 is zero. If n_1 and n_2 differ in sign, the implementation-defined result returned will be the same as that returned by either the phrase `>R S>D R> FM/MOD SWAP DROP` or the phrase `>R S>D R> SM/REM SWAP DROP`.

See: **3.2.2.1 Integer division**.

6.1.0240 `/MOD` "slash-mod" CORE

(`n1 n2 -- n3 n4`)

Divide n_1 by n_2 , giving the single-cell remainder n_3 and the single-cell quotient n_4 . An ambiguous condition exists if n_2 is zero. If n_1 and n_2 differ in sign, the implementation-defined result returned will be the same as that returned by either the phrase `>R S>D R> FM/MOD` or the phrase `>R S>D R> SM/REM`.

See: **3.2.2.1 Integer division**.

6.1.0250 `0<` "zero-less" CORE

(`n -- flag`)

`flag` is true if and only if n is less than zero.

6.1.0270 `0=` "zero-equals" CORE

(`x -- flag`)

`flag` is true if and only if x is equal to zero.

6.1.0290 `1+` "one-plus" CORE

(`n1|u1 -- n2|u2`)

Add 1 to $n_1|u_1$ giving the sum $n_2|u_2$.

6.1.0300	1-	“one-minus”	CORE
		$(n_1 u_1 -- n_2 u_2)$	
		Subtract 1 from $n_1 u_1$ giving the difference $n_2 u_2$.	
6.1.0310	2!	“two-store”	CORE
		$(x_1 x_2 a-addr --)$	
		Store the cell pair $x_1 x_2$ at $a-addr$, with x_2 at $a-addr$ and x_1 at the next consecutive cell. It is equivalent to the sequence SWAP OVER ! CELL+ ! .	
		See: 3.3.3.1 Address alignment .	
6.1.0320	2*	“two-star”	CORE
		$(x_1 -- x_2)$	
		x_2 is the result of shifting x_1 one bit toward the most-significant bit, filling the vacated least-significant bit with zero.	
6.1.0330	2/	“two-slash”	CORE
		$(x_1 -- x_2)$	
		x_2 is the result of shifting x_1 one bit toward the least-significant bit, leaving the most-significant bit unchanged.	
6.1.0350	2@	“two-fetch”	CORE
		$(a-addr -- x_1 x_2)$	
		Fetch the cell pair $x_1 x_2$ stored at $a-addr$. x_2 is stored at $a-addr$ and x_1 at the next consecutive cell. It is equivalent to the sequence DUP CELL+ @ SWAP @ .	
		See: 3.3.3.1 Address alignment, 6.1.0310 2! .	
6.1.0370	2DROP	“two-drop”	CORE
		$(x_1 x_2 --)$	
		Drop cell pair $x_1 x_2$ from the stack.	
6.1.0380	2DUP	“two-dupe”	CORE
		$(x_1 x_2 -- x_1 x_2 x_1 x_2)$	
		Duplicate cell pair $x_1 x_2$.	
6.1.0400	2OVER	“two-over”	CORE
		$(x_1 x_2 x_3 x_4 -- x_1 x_2 x_3 x_4 x_1 x_2)$	
		Copy cell pair $x_1 x_2$ to the top of the stack.	

6.1.0430 **2SWAP** "two-swap" CORE

($x_1 x_2 x_3 x_4$ -- $x_3 x_4 x_1 x_2$)

Exchange the top two cell pairs.

6.1.0450 **:** "colon" CORE

(C: "<spaces>name" -- colon-sys)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name*, called a "colon definition." Enter compilation state and start the current definition, producing *colon-sys*. Append the initiation semantics given below to the current definition.

The execution semantics of *name* will be determined by the words compiled into the body of the definition. The current definition shall not be findable in the dictionary until it is ended (or until the execution of **DOES>** in some systems).

Initiation: ($i*x$ -- $i*x$) (R: -- nest-sys)

Save implementation-dependent information *nest-sys* about the calling definition. The stack effects $i*x$ represent arguments to *name*.

name Execution: ($i*x$ -- $j*x$)

Execute the definition *name*. The stack effects $i*x$ and $j*x$ represent arguments to and results from *name*, respectively.

See: **3.4 The Forth text interpreter, 3.4.1 Parsing, 3.4.5 Compilation, 6.1.1250 DOES>, 6.1.2500 [, 6.1.2540] , 15.6.2.0470 ;CODE.**

6.1.0460 **;** "semicolon" CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: colon-sys --)

Append the run-time semantics below to the current definition. End the current definition, allow it to be found in the dictionary and enter interpretation state, consuming *colon-sys*. If the data-space pointer is not aligned, reserve enough data space to align it.

Run-time: (--) (R: nest-sys --)

Return to the calling definition specified by *nest-sys*.

See: **3.4 The Forth text interpreter, 3.4.5 Compilation.**

6.1.0480 **<** "less-than" CORE

($n_1 n_2$ -- flag)

flag is true if and only if n_1 is less than n_2 .

See: **6.1.2340 U<.**

-
- 6.1.0490** <# "less-number-sign" CORE
 (--)
 Initialize the pictured numeric output conversion process.
 See: **6.1.0030 #**, **6.1.0040 #>**, **6.1.0050 #s**.
- 6.1.0530** = "equals" CORE
 (x_1 x_2 -- *flag*)
flag is true if and only if x_1 is bit-for-bit the same as x_2 .
- 6.1.0540** > "greater-than" CORE
 (n_1 n_2 -- *flag*)
flag is true if and only if n_1 is greater than n_2 .
 See: **6.2.2350 U>**.
- 6.1.0550** >BODY "to-body" CORE
 (*xt* -- *a-addr*)
a-addr is the data-field address corresponding to *xt*. An ambiguous condition exists if *xt* is not for a word defined via **CREATE**.
 See: **3.3.3 Data space**.
- 6.1.0560** >IN "to-in" CORE
 (-- *a-addr*)
a-addr is the address of a cell containing the offset in characters from the start of the input buffer to the start of the parse area.
- 6.1.0570** >NUMBER "to-number" CORE
 (ud_1 *c-addr*₁ u_1 -- ud_2 *c-addr*₂ u_2)
 ud_2 is the unsigned result of converting the characters within the string specified by *c-addr*₁ u_1 into digits, using the number in **BASE**, and adding each into ud_1 after multiplying ud_1 by the number in **BASE**. Conversion continues left-to-right until a character that is not convertible, including any "+" or "-", is encountered or the string is entirely converted. *c-addr*₂ is the location of the first unconverted character or the first character past the end of the string if the string was entirely converted. u_2 is the number of unconverted characters in the string. An ambiguous condition exists if ud_2 overflows during the conversion.
 See: **3.2.1.2 Digit conversion**.
-

6.1.0580 >R “to-r” CORE

Interpretation: Interpretation semantics for this word are undefined.

Execution: $(x \ --) \ (R: \ -- \ x)$

Move x to the return stack.

See: **3.2.3.3 Return stack**, **6.1.2060 R>**, **6.1.2070 R@**, **6.2.0340 2>R**, **6.2.0410 2R>**, **6.2.0415 2R@**.

6.1.0630 ?DUP “question-dupe” CORE

$(x \ -- \ 0 \ | \ x \ x)$

Duplicate x if it is non-zero.

6.1.0650 @ “fetch” CORE

$(a\text{-}addr \ -- \ x)$

x is the value stored at $a\text{-}addr$.

See: **3.3.3.1 Address alignment**.

6.1.0670 ABORT CORE

$(i*x \ --) \ (R: \ j*x \ --)$

Empty the data stack and perform the function of **QUIT**, which includes emptying the return stack, without displaying a message.

See: **9.6.2.0670 ABORT**.

6.1.0680 ABORT" “abort-quote” CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: $(\text{“}ccc\text{<quote>”} \ --)$

Parse ccc delimited by a " (double-quote). Append the run-time semantics given below to the current definition.

Run-time: $(i*x \ x_j \ -- \ | \ i*x) \ (R: \ j*x \ -- \ | \ j*x)$

Remove x_j from the stack. If any bit of x_j is not zero, display ccc and perform an implementation-defined abort sequence that includes the function of **ABORT**.

See: **3.4.1 Parsing**, **9.6.2.0680 ABORT"**.

6.1.0690 ABS “abs” CORE

$(n \ -- \ u)$

u is the absolute value of n .

6.1.0695 ACCEPT

CORE

 $(c\text{-}addr +n_1 \text{--} +n_2)$

Receive a string of at most $+n_1$ characters. An ambiguous condition exists if $+n_1$ is zero or greater than 32,767. Display graphic characters as they are received. A program that depends on the presence or absence of non-graphic characters in the string has an environmental dependency. The editing functions, if any, that the system performs in order to construct the string are implementation-defined.

Input terminates when an implementation-defined line terminator is received. When input terminates, nothing is appended to the string, and the display is maintained in an implementation-defined way.

$+n_2$ is the length of the string stored at $c\text{-}addr$.

6.1.0705 ALIGN

CORE

 (--)

If the data-space pointer is not aligned, reserve enough space to align it.

See: **3.3.3 Data space, 3.3.3.1 Address alignment.**

6.1.0706 ALIGNED

CORE

 $(addr \text{--} a\text{-}addr)$

$a\text{-}addr$ is the first aligned address greater than or equal to $addr$.

See: **3.3.3.1 Address alignment.**

6.1.0710 ALLOT

CORE

 $(n \text{--})$

If n is greater than zero, reserve n address units of data space. If n is less than zero, release $|n|$ address units of data space. If n is zero, leave the data-space pointer unchanged.

If the data-space pointer is aligned and n is a multiple of the size of a cell when **ALLOT** begins execution, it will remain aligned when **ALLOT** finishes execution.

If the data-space pointer is character aligned and n is a multiple of the size of a character when **ALLOT** begins execution, it will remain character aligned when **ALLOT** finishes execution.

See: **3.3.3 Data space.**

6.1.0720 AND

CORE

 $(x_1 x_2 \text{--} x_3)$

x_3 is the bit-by-bit logical “and” of x_1 with x_2 .

6.1.0750 BASE **CORE**

(-- *a-addr*)

a-addr is the address of a cell containing the current number-conversion radix {{2...36}}.

6.1.0760 BEGIN **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *dest*)

Put the next location for a transfer of control, *dest*, onto the control flow stack. Append the run-time semantics given below to the current definition.

Run-time: (--)

Continue execution.

See: **3.2.3.2 Control-flow stack**, **6.1.2140 REPEAT**, **6.1.2390 UNTIL**, **6.1.2430 WHILE**.

6.1.0770 BL **CORE**

“b-l”

(-- *char*)

char is the character value for a space.

6.1.0850 C! **CORE**

“c-store”

(*char c-addr* --)

Store *char* at *c-addr*. When character size is smaller than cell size, only the number of low-order bits corresponding to character size are transferred.

See: **3.3.3.1 Address alignment**

6.1.0860 C, **CORE**

“c-comma”

(*char* --)

Reserve space for one character in the data space and store *char* in the space. If the data-space pointer is character aligned when **C,** begins execution, it will remain character aligned when **C,** finishes execution. An ambiguous condition exists if the data-space pointer is not character-aligned prior to execution of **C,**.

See: **3.3.3 Data space**, **3.3.3.1 Address alignment**.

6.1.0870 C@ **CORE**

“c-fetch”

(*c-addr* -- *char*)

Fetch the character stored at *c-addr*. When the cell size is greater than character size, the unused high-order bits are all zeroes.

See: **3.3.3.1 Address alignment**.

-
- 6.1.0880 CELL+** “cell-plus” CORE
 ($a\text{-addr}_1$ -- $a\text{-addr}_2$)
 Add the size in address units of a cell to $a\text{-addr}_1$, giving $a\text{-addr}_2$.
 See: **3.3.3.1 Address alignment.**
- 6.1.0890 CELLS** CORE
 (n_1 -- n_2)
 n_2 is the size in address units of n_1 cells.
- 6.1.0895 CHAR** “char” CORE
 (“<spaces>name” -- char)
 Skip leading space delimiters. Parse *name* delimited by a space. Put the value of its first character onto the stack.
 See: **3.4.1 Parsing, 6.1.2520 [CHAR].**
- 6.1.0897 CHAR+** “char-plus” CORE
 ($c\text{-addr}_1$ -- $c\text{-addr}_2$)
 Add the size in address units of a character to $c\text{-addr}_1$, giving $c\text{-addr}_2$.
 See: **3.3.3.1 Address alignment.**
- 6.1.0898 CHARS** “chars” CORE
 (n_1 -- n_2)
 n_2 is the size in address units of n_1 characters.
- 6.1.0950 CONSTANT** CORE
 (x “<spaces>name” --)
 Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below.
name is referred to as a “constant.”
name Execution: (-- x)
 Place x on the stack.
 See: **3.4.1 Parsing.**
-

6.1.0980 COUNT **CORE**

(*c-addr*₁ -- *c-addr*₂ *u*)

Return the character string specification for the counted string stored at *c-addr*₁. *c-addr*₂ is the address of the first character after *c-addr*₁. *u* is the contents of the character at *c-addr*₁, which is the length in characters of the string at *c-addr*₂.

6.1.0990 CR **CORE**

“C-r”

(--)

Cause subsequent output to appear at the beginning of the next line.

6.1.1000 CREATE **CORE**

(“<*spaces*>*name*” --)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below. If the data-space pointer is not aligned, reserve enough data space to align it. The new data-space pointer defines *name*'s data field. **CREATE** does not allocate data space in *name*'s data field.

name Execution: (-- *a-addr*)

a-addr is the address of *name*'s data field. The execution semantics of *name* may be extended by using **DOES**>.

See: **3.3.3 Data space, 6.1.1250 DOES**>.

6.1.1170 DECIMAL **CORE**

(--)

Set the numeric conversion radix to ten (decimal).

6.1.1200 DEPTH **CORE**

(-- +*n*)

+*n* is the number of single-cell values contained in the data stack before +*n* was placed on the stack.

6.1.1240 DO

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *do-sys*)

Place *do-sys* onto the control-flow stack. Append the run-time semantics given below to the current definition. The semantics are incomplete until resolved by a consumer of *do-sys* such as **LOOP**.

Run-time: ($n_1|u_1$ $n_2|u_2$ --) (R: -- *loop-sys*)

Set up loop control parameters with index $n_2|u_2$ and limit $n_1|u_1$. An ambiguous condition exists if $n_1|u_1$ and $n_2|u_2$ are not both the same type. Anything already on the return stack becomes unavailable until the loop-control parameters are discarded.

See: **3.2.3.2 Control-flow stack**, **6.1.0140 +LOOP**, **6.1.1800 LOOP**.

6.1.1250 DOES>

"does"

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: *colon-sys*₁ -- *colon-sys*₂)

Append the run-time semantics below to the current definition. Whether or not the current definition is rendered findable in the dictionary by the compilation of **DOES>** is implementation defined. Consume *colon-sys*₁ and produce *colon-sys*₂. Append the initiation semantics given below to the current definition.

Run-time: (--) (R: *nest-sys*₁ --)

Replace the execution semantics of the most recent definition, referred to as *name*, with the *name* execution semantics given below. Return control to the calling definition specified by *nest-sys*₁. An ambiguous condition exists if *name* was not defined with **CREATE** or a user-defined word that calls **CREATE**.

Initiation: ($i*x$ -- $i*x$ *a-addr*) (R: -- *nest-sys*₂)

Save implementation-dependent information *nest-sys*₂ about the calling definition. Place *name*'s data field address on the stack. The stack effects $i*x$ represent arguments to *name*.

name Execution: ($i*x$ -- $j*x$)

Execute the portion of the definition that begins with the initiation semantics appended by the **DOES>** which modified *name*. The stack effects $i*x$ and $j*x$ represent arguments to and results from *name*, respectively.

See: **6.1.1000 CREATE**.

6.1.1260 DROP

CORE

(x --)

Remove x from the stack.

6.1.1290 DUP “dupe” CORE $(x \text{ -- } x x)$ Duplicate x .**6.1.1310 ELSE** CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: $(C: orig_1 \text{ -- } orig_2)$

Put the location of a new unresolved forward reference $orig_2$ onto the control flow stack. Append the run-time semantics given below to the current definition. The semantics will be incomplete until $orig_2$ is resolved (e.g., by **THEN**). Resolve the forward reference $orig_1$ using the location following the appended execution semantics.

Run-time: $(--)$ Continue execution at the location given by the resolution of $orig_2$.See: **6.1.1700 IF**, **6.1.2270 THEN**.**6.1.1320 EMIT** CORE $(x \text{ --})$

If x is a graphic character in the implementation-defined character set, display x . The effect of **EMIT** for all other values of x is implementation-defined.

When passed a character whose character-defining bits have a value between hex 20 and 7E inclusive, the corresponding standard character, specified by **3.1.2.1 Graphic characters**, is displayed. Because different output devices can respond differently to control characters, programs that use control characters to perform specific functions have an environmental dependency. Each **EMIT** deals with only one character.

See: **6.1.2310 TYPE**.**6.1.1345 ENVIRONMENT?** “environment-query” CORE $(c\text{-}addr\ u \text{ -- } false \mid i^*x\ true)$

$c\text{-}addr$ is the address of a character string and u is the string's character count. u may have a value in the range from zero to an implementation-defined maximum which shall not be less than 31. The character string should contain a keyword from **3.2.6 Environmental queries** or the optional word sets to be checked for correspondence with an attribute of the present environment. If the system treats the attribute as unknown, the returned flag is *false*; otherwise, the flag is *true* and the i^*x returned is of the type specified in the table for the attribute queried.

6.1.1360 EVALUATE

CORE

 $(i*x \ c\text{-addr} \ u \ -- \ j*x)$

Save the current input source specification. Store minus-one (-1) in **SOURCE-ID** if it is present. Make the string described by *c-addr* and *u* both the input source and input buffer, set **>IN** to zero, and interpret. When the parse area is empty, restore the prior input source specification. Other stack effects are due to the words **EVALUATED**.

6.1.1370 EXECUTE

CORE

 $(i*x \ xt \ -- \ j*x)$

Remove *xt* from the stack and perform the semantics identified by it. Other stack effects are due to the word **EXECUTED**.

See: **6.1.0070** ' , **6.1.2510** ['] .

6.1.1380 EXIT

CORE

Interpretation: Interpretation semantics for this word are undefined.

Execution: $(--) (R: \text{nest-sys} \ --)$

Return control to the calling definition specified by *nest-sys*. Before executing **EXIT** within a do-loop, a program shall discard the loop-control parameters by executing **UNLOOP**.

See: **3.2.3.3 Return stack**, **6.1.2380 UNLOOP**.

6.1.1540 FILL

CORE

 $(c\text{-addr} \ u \ char \ --)$

If *u* is greater than zero, store *char* in each of *u* consecutive characters of memory beginning at *c-addr*.

6.1.1550 FIND

CORE

 $(c\text{-addr} \ -- \ c\text{-addr} \ 0 \ | \ xt \ 1 \ | \ xt \ -1)$

Find the definition named in the counted string at *c-addr*. If the definition is not found, return *c-addr* and zero. If the definition is found, return its execution token *xt*. If the definition is immediate, also return one (1), otherwise also return minus-one (-1). For a given string, the values returned by **FIND** while compiling may differ from those returned while not compiling.

See: **3.4.2 Finding definition names**, **A.6.1.0070** ' , **A.6.1.2510** ['] , **A.6.1.2033 POSTPONE**, **D.6.7 Immediacy**.

6.1.1561 FM/MOD "f-m-slash-mod" **CORE**

($d_1 n_1$ -- $n_2 n_3$)

Divide d_1 by n_1 , giving the floored quotient n_3 and the remainder n_2 . Input and output stack arguments are signed. An ambiguous condition exists if n_1 is zero or if the quotient lies outside the range of a single-cell signed integer.

See: **3.2.2.1 Integer division**, **6.1.2214 SM/REM**, **6.1.2370 UM/MOD**.

6.1.1650 HERE **CORE**

(-- *addr*)

addr is the data-space pointer.

See: **3.3.3.2 Contiguous regions**.

6.1.1670 HOLD **CORE**

(*char* --)

Add *char* to the beginning of the pictured numeric output string. An ambiguous condition exists if **HOLD** executes outside of a <# #> delimited number conversion.

6.1.1680 I **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Execution: (-- $n|u$) (R: *loop-sys* -- *loop-sys*)

$n|u$ is a copy of the current (innermost) loop index. An ambiguous condition exists if the loop control parameters are unavailable.

6.1.1700 IF **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *orig*)

Put the location of a new unresolved forward reference *orig* onto the control flow stack. Append the run-time semantics given below to the current definition. The semantics are incomplete until *orig* is resolved, e.g., by **THEN** or **ELSE**.

Run-time: (*x* --)

If all bits of *x* are zero, continue execution at the location specified by the resolution of *orig*.

See: **3.2.3.2 Control flow stack**, **6.1.1310 ELSE**, **6.1.2270 THEN**.

6.1.1710 IMMEDIATE

CORE

(--)

Make the most recent definition an immediate word. An ambiguous condition exists if the most recent definition does not have a name.

See: **D.6.7 Immediacy**.

6.1.1720 INVERT

CORE

(x_1 -- x_2)

Invert all bits of x_1 , giving its logical inverse x_2 .

See: **6.1.1910 NEGATE**, **6.1.0270 0=**.

6.1.1730 J

CORE

Interpretation: Interpretation semantics for this word are undefined.

Execution: (-- $n|u$) (R: *loop-sys* -- *loop-sys*)

$n|u$ is a copy of the index of the next outer loop. An ambiguous condition exists if the loop control parameters of the next outer loop are unavailable.

6.1.1750 KEY

CORE

(-- *char*)

Receive one character *char*, a member of the implementation-defined character set. Keyboard events that do not correspond to such characters are discarded until a valid character is received, and those events are subsequently unavailable.

All standard characters can be received. Characters received by **KEY** are not displayed.

Any standard character returned by **KEY** has the numeric value specified in **3.1.2.1 Graphic characters**. Programs that require the ability to receive control characters have an environmental dependency.

See: **10.6.2.1307 EKEY**, **10.6.1.1755 KEY?**.

6.1.1760 LEAVE

CORE

Interpretation: Interpretation semantics for this word are undefined.

Execution: (--) (R: *loop-sys* --)

Discard the current loop control parameters. An ambiguous condition exists if they are unavailable. Continue execution immediately following the innermost syntactically enclosing **DO ... LOOP** or **DO ... +LOOP**.

See: **3.2.3.3 Return stack**, **6.1.0140 +LOOP**, **6.1.1800 LOOP**.

6.1.1780 LITERAL CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (x --)

Append the run-time semantics given below to the current definition.

Run-time: (-- x)

Place x on the stack.

6.1.1800 LOOP CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: *do-sys* --)

Append the run-time semantics given below to the current definition. Resolve the destination of all unresolved occurrences of **LEAVE** between the location given by *do-sys* and the next location for a transfer of control, to execute the words following the **LOOP**.

Run-time: (--) (R: *loop-sys*₁ -- | *loop-sys*₂)

An ambiguous condition exists if the loop control parameters are unavailable. Add one to the loop index. If the loop index is then equal to the loop limit, discard the loop parameters and continue execution immediately following the loop. Otherwise continue execution at the beginning of the loop.

See: **6.1.1240 DO**, **6.1.1680 I**, **6.1.1760 LEAVE**.

6.1.1805 LSHIFT CORE "l-shift"

(x_1 u -- x_2)

Perform a logical left shift of u bit-places on x_1 , giving x_2 . Put zero into the least significant bits vacated by the shift. An ambiguous condition exists if u is greater than or equal to the number of bits in a cell.

6.1.1810 M* CORE "m-star"

(n_1 n_2 -- d)

d is the signed product of n_1 times n_2 .

6.1.1870 MAX CORE

(n_1 n_2 -- n_3)

n_3 is the greater of n_1 and n_2 .

6.1.1880 MIN CORE

(n_1 n_2 -- n_3)

n_3 is the lesser of n_1 and n_2 .

6.1.1890 MOD **CORE**

$$(n_1 n_2 \text{ -- } n_3)$$

Divide n_1 by n_2 , giving the single-cell remainder n_3 . An ambiguous condition exists if n_2 is zero. If n_1 and n_2 differ in sign, the implementation-defined result returned will be the same as that returned by either the phrase **>R S>D R> FM/MOD DROP** or the phrase **>R S>D R> SM/REM DROP**.

See: **3.2.2.1 Integer division**.

6.1.1900 MOVE **CORE**

$$(addr_1 addr_2 u \text{ -- })$$

If u is greater than zero, copy the contents of u consecutive address units at $addr_1$ to the u consecutive address units at $addr_2$. After **MOVE** completes, the u consecutive address units at $addr_2$ contain exactly what the u consecutive address units at $addr_1$ contained before the move.

See: **17.6.1.0910 CMOVE**, **17.6.1.0920 CMOVE>**.

6.1.1910 NEGATE **CORE**

$$(n_1 \text{ -- } n_2)$$

Negate n_1 , giving its arithmetic inverse n_2 .

See: **6.1.1720 INVERT**, **6.1.0270 0=**.

6.1.1980 OR **CORE**

$$(x_1 x_2 \text{ -- } x_3)$$

x_3 is the bit-by-bit inclusive-or of x_1 with x_2 .

6.1.1990 OVER **CORE**

$$(x_1 x_2 \text{ -- } x_1 x_2 x_1)$$

Place a copy of x_1 on top of the stack.

6.1.2033 POSTPONE **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: ("*<spaces>name*" --)

Skip leading space delimiters. Parse *name* delimited by a space. Find *name*. Append the compilation semantics of *name* to the current definition. An ambiguous condition exists if *name* is not found.

See: **3.4.1 Parsing**.

6.1.2050 QUIT **CORE**

(--) (R: i*x --)

Empty the return stack, store zero in **SOURCE-ID** if it is present, make the user input device the input source, and enter interpretation state. Do not display a message. Repeat the following:

- Accept a line from the input source into the input buffer, set **>IN** to zero, and interpret.
- Display the implementation-defined system prompt if in interpretation state, all processing has been completed, and no ambiguous condition exists.

See: **3.4 The Forth text interpreter.**

6.1.2060 R> **CORE** “r-from”

Interpretation: Interpretation semantics for this word are undefined.

Execution: (-- x) (R: x --)

Move *x* from the return stack to the data stack.

See: **3.2.3.3 Return stack, 6.1.0580 >R, 6.1.2070 R@, 6.2.0340 2>R, 6.2.0410 2R>, 6.2.0415 2R@.**

6.1.2070 R@ **CORE** “r-fetch”

Interpretation: Interpretation semantics for this word are undefined.

Execution: (-- x) (R: x -- x)

Copy *x* from the return stack to the data stack.

See: **3.2.3.3 Return stack, 6.1.0580 >R, 6.1.2060 R>, 6.2.0340 2>R, 6.2.0410 2R>, 6.2.0415 2R@.**

6.1.2120 RECURSE **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (--)

Append the execution semantics of the current definition to the current definition. An ambiguous condition exists if **RECURSE** appears in a definition after **DOES>**.

See: **6.1.1250 DOES>, 6.1.2120 RECURSE.**

6.1.2140 REPEAT

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (*C: orig dest --*)

Append the run-time semantics given below to the current definition, resolving the backward reference *dest*. Resolve the forward reference *orig* using the location following the appended execution semantics.

Run-time: (--)

Continue execution at the location given by *dest*.

See: **6.1.0760 BEGIN**, **6.1.2430 WHILE**.

6.1.2160 ROT

"rote"

CORE

($x_1 x_2 x_3 -- x_2 x_3 x_1$)

Rotate the top three stack entries.

6.1.2162 RSHIFT

"r-shift"

CORE

($x_1 u -- x_2$)

Perform a logical right shift of *u* bit-places on x_1 , giving x_2 . Put zero into the most significant bits vacated by the shift. An ambiguous condition exists if *u* is greater than or equal to the number of bits in a cell.

6.1.2165 s"

"s-quote"

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: ("*ccc*<quote>" --)

Parse *ccc* delimited by " (double-quote). Append the run-time semantics given below to the current definition.

Run-time: (-- *c-addr u*)

Return *c-addr* and *u* describing a string consisting of the characters *ccc*. A program shall not alter the returned string.

See: **3.4.1 Parsing**, **6.2.0855 C"**, **11.6.1.2165 s"**.

6.1.2170 S>D

"s-to-d"

CORE

($n -- d$)

Convert the number *n* to the double-cell number *d* with the same numerical value.

6.1.2210 SIGN **CORE** $(n --)$

If n is negative, add a minus sign to the beginning of the pictured numeric output string. An ambiguous condition exists if **SIGN** executes outside of a **<# #>** delimited number conversion.

6.1.2214 SM/REM **CORE**

"s-m-slash-rem"

 $(d_1 n_1 -- n_2 n_3)$

Divide d_1 by n_1 , giving the symmetric quotient n_3 and the remainder n_2 . Input and output stack arguments are signed. An ambiguous condition exists if n_1 is zero or if the quotient lies outside the range of a single-cell signed integer.

See: **3.2.2.1 Integer division**, **6.1.1561**, **FM/MOD**, **6.1.2370**, **UM/MOD**.

6.1.2216 SOURCE **CORE** $(-- c-addr u)$

$c-addr$ is the address of, and u is the number of characters in, the input buffer.

6.1.2220 SPACE **CORE** $(--)$

Display one space.

6.1.2230 SPACES **CORE** $(n --)$

If n is greater than zero, display n spaces.

6.1.2250 STATE **CORE** $(-- a-addr)$

$a-addr$ is the address of a cell containing the compilation-state flag. **STATE** is true when in compilation state, false otherwise. The true value in **STATE** is non-zero, but is otherwise implementation-defined. Only the following standard words alter the value in **STATE**: **:** (colon), **;** (semicolon), **ABORT**, **QUIT**, **:NONAME**, **[** (left-bracket), and **]** (right-bracket).

Note: A program shall not directly alter the contents of **STATE**.

See: **3.4 The Forth text interpreter**, **6.1.0450** **:**, **6.1.0460** **;**, **6.1.0670** **ABORT**, **6.1.2050** **QUIT**, **6.1.2500** **[**, **6.1.2540** **]**, **6.2.0455** **:NONAME**, **15.6.2.2250** **STATE**.

6.1.2260 SWAP **CORE** $(x_1 x_2 -- x_2 x_1)$

Exchange the top two stack items.

6.1.2270 THEN

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (*C: orig --*)

Resolve the forward reference *orig* using the location of the execution semantics. Append the run-time semantics given below to the current definition.

Run-time: (--)

Continue execution.

See: **6.1.1310 ELSE**, **6.1.1700 IF**.

6.1.2310 TYPE

CORE

(*c-addr u --*)

If *u* is greater than zero, display the character string specified by *c-addr* and *u*.

When passed a character in a character string whose character-defining bits have a value between hex 20 and 7E inclusive, the corresponding standard character, specified by **3.1.2.1 graphic characters**, is displayed. Because different output devices can respond differently to control characters, programs that use control characters to perform specific functions have an environmental dependency.

See: **6.1.1320 EMIT**.

6.1.2320 U.

"u-dot"

CORE

(*u --*)

Display *u* in free field format.

6.1.2340 U<

"u-less-than"

CORE

(*u₁ u₂ -- flag*)

flag is true if and only if *u₁* is less than *u₂*.

See: **6.1.0480 <**.

6.1.2360 UM*

"u-m-star"

CORE

(*u₁ u₂ -- ud*)

Multiply *u₁* by *u₂*, giving the unsigned double-cell product *ud*. All values and arithmetic are unsigned.

6.1.2370 UM/MOD “u-m-slash-mod” **CORE**

(*ud* u_1 -- u_2 u_3)

Divide *ud* by u_1 , giving the quotient u_3 and the remainder u_2 . All values and arithmetic are unsigned. An ambiguous condition exists if u_1 is zero or if the quotient lies outside the range of a single-cell unsigned integer.

See: **3.2.2.1 Integer division**, **6.1.1561 FM/MOD**, **6.1.2214 SM/REM**.

6.1.2380 UNLOOP **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Execution: (--) (R: *loop-sys* --)

Discard the loop-control parameters for the current nesting level. An **UNLOOP** is required for each nesting level before the definition may be **EXITed**. An ambiguous condition exists if the loop-control parameters are unavailable.

See: **3.2.3.3 Return stack**.

6.1.2390 UNTIL **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: *dest* --)

Append the run-time semantics given below to the current definition, resolving the backward reference *dest*.

Run-time: (*x* --)

If all bits of *x* are zero, continue execution at the location specified by *dest*.

See: **6.1.0760 BEGIN**.

6.1.2410 VARIABLE **CORE**

(“<*spaces*>*name*” --)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below. Reserve one cell of data space at an aligned address.

name is referred to as a “variable”.

name Execution: (-- *a-addr*)

a-addr is the address of the reserved cell. A program is responsible for initializing the contents of the reserved cell.

See: **3.4.1 Parsing**.

6.1.2430 WHILE

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (*C: dest -- orig dest*)

Put the location of a new unresolved forward reference *orig* onto the control flow stack, under the existing *dest*. Append the run-time semantics given below to the current definition. The semantics are incomplete until *orig* and *dest* are resolved (e.g., by **REPEAT**).

Run-time: (*x --*)

If all bits of *x* are zero, continue execution at the location specified by the resolution of *orig*.

6.1.2450 WORD

CORE

(*char “<chars>ccc<char>” -- c-addr*)

Skip leading delimiters. Parse characters *ccc* delimited by *char*. An ambiguous condition exists if the length of the parsed string is greater than the implementation-defined length of a counted string.

c-addr is the address of a transient region containing the parsed word as a counted string. If the parse area was empty or contained no characters other than the delimiter, the resulting string has a zero length. A space, not included in the length, follows the string. A program may replace characters within the string.

Note: The requirement to follow the string with a space is obsolescent and is included as a concession to existing programs that use **CONVERT**. A program shall not depend on the existence of the space.

See: **3.3.3.6 Other transient regions, 3.4.1 Parsing.**

6.1.2490 XOR

“x-or”

CORE

(*x₁ x₂ -- x₃*)

x₃ is the bit-by-bit exclusive-or of *x₁* with *x₂*.

6.1.2500 [

“left-bracket”

CORE

Interpretation: Interpretation semantics for this word are undefined.

Compilation: Perform the execution semantics given below.

Execution: (--)

Enter interpretation state. [is an immediate word.

See: **3.4 The Forth text interpreter, 3.4.5 Compilation, 6.1.2540]**.

6.1.2510 ['] "bracket-tick" **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: ("<spaces>name" --)

Skip leading space delimiters. Parse *name* delimited by a space. Find *name*. Append the run-time semantics given below to the current definition.

An ambiguous condition exists if *name* is not found.

Run-time: (-- *xt*)

Place *name*'s execution token *xt* on the stack. The execution token returned by the compiled phrase "['] X" is the same value returned by "' X" outside of compilation state.

See: **3.4.1 Parsing, A.6.1.0070 ' , A.6.1.2033 POSTPONE, D.6.7 Immediacy.**

6.1.2520 [CHAR] "bracket-char" **CORE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: ("<spaces>name" --)

Skip leading space delimiters. Parse *name* delimited by a space. Append the run-time semantics given below to the current definition.

Run-time: (-- *char*)

Place *char*, the value of the first character of *name*, on the stack.

See: **3.4.1 Parsing, 6.1.0895 CHAR.**

6.1.2540] "right-bracket" **CORE**

(--)

Enter compilation state.

See: **3.4 The Forth text interpreter, 3.4.5 Compilation, 6.1.2500 [.**

6.2 Core extension words

6.2.0060 #TIB "number-t-i-b" CORE EXT

(-- *a-addr*)

a-addr is the address of a cell containing the number of characters in the terminal input buffer.

Note: This word is obsolescent and is included as a concession to existing implementations.

6.2.0200 .("dot-paren" CORE EXT

Compilation: Perform the execution semantics given below.

Execution: ("*ccc*<*paren*>" --)

Parse and display *ccc* delimited by a right parenthesis ")". .(is an immediate word.

See: **3.4.1 Parsing, 6.1.0190 ."**

6.2.0210 .R "dot-r" CORE EXT

(*n*₁ *n*₂ --)

Display *n*₁ right aligned in a field *n*₂ characters wide. If the number of characters required to display *n*₁ is greater than *n*₂, all digits are displayed with no leading spaces in a field as wide as necessary.

6.2.0260 0<> "zero-not-equals" CORE EXT

(*x* -- *flag*)

flag is true if and only if *x* is not equal to zero.

6.2.0280 0> "zero-greater" CORE EXT

(*n* -- *flag*)

flag is true if and only if *n* is greater than zero.

6.2.0340 2>R "two-to-r" CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Execution: (*x*₁ *x*₂ --) (R: -- *x*₁ *x*₂)

Transfer cell pair *x*₁ *x*₂ to the return stack. Semantically equivalent to **SWAP >R >R**.

See: **3.2.3.3 Return stack, 6.1.0580 >R, 6.1.2060 R>, 6.1.2070 R@, 6.2.0410 2R>, 6.2.0415 2R@.**

6.2.0410 2R> “two-r-from” **CORE EXT**

Interpretation: Interpretation semantics for this word are undefined.

Execution: $(-- x_1 x_2) (R: x_1 x_2 --)$

Transfer cell pair $x_1 x_2$ from the return stack. Semantically equivalent to **R> R> SWAP**.

See: **3.2.3.3 Return stack**, **6.1.0580 >R**, **6.1.2060 R>**, **6.1.2070 R@**, **6.2.0340 2>R**, **6.2.0415 2R@**.

6.2.0415 2R@ “two-r-fetch” **CORE EXT**

Interpretation: Interpretation semantics for this word are undefined.

Execution: $(-- x_1 x_2) (R: x_1 x_2 -- x_1 x_2)$

Copy cell pair $x_1 x_2$ from the return stack. Semantically equivalent to **R> R> 2DUP >R >R SWAP**.

See: **3.2.3.3 Return stack**, **6.1.0580 >R**, **6.1.2060 R>**, **6.1.2070 R@**, **6.2.0340 2>R**, **6.2.0410 2R>**.

6.2.0455 :NONAME “colon-no-name” **CORE EXT**

$(C: -- colon-sys) (S: -- xt)$

Create an execution token xt , enter compilation state and start the current definition, producing $colon-sys$. Append the initiation semantics given below to the current definition.

The execution semantics of xt will be determined by the words compiled into the body of the definition. This definition can be executed later by using xt **EXECUTE**.

If the control-flow stack is implemented using the data stack, $colon-sys$ shall be the topmost item on the data stack. See **3.2.3.2 Control-flow stack**.

Initiation: $(i*x -- i*x) (R: -- nest-sys)$

Save implementation-dependent information $nest-sys$ about the calling definition. The stack effects $i*x$ represent arguments to xt .

xt Execution: $(i*x -- j*x)$

Execute the definition specified by xt . The stack effects $i*x$ and $j*x$ represent arguments to and results from xt , respectively.

6.2.0500 <> “not-equals” **CORE EXT**

$(x_1 x_2 -- flag)$

$flag$ is true if and only if x_1 is not bit-for-bit the same as x_2 .

6.2.0620 ?DO "question-do" CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *do-sys*)

Put *do-sys* onto the control-flow stack. Append the run-time semantics given below to the current definition. The semantics are incomplete until resolved by a consumer of *do-sys* such as **LOOP**.

Run-time: ($n_1|u_1$ $n_2|u_2$ --) (R: -- | *loop-sys*)

If $n_1|u_1$ is equal to $n_2|u_2$, continue execution at the location given by the consumer of *do-sys*. Otherwise set up loop control parameters with index $n_2|u_2$ and limit $n_1|u_1$ and continue executing immediately following **?DO**. Anything already on the return stack becomes unavailable until the loop control parameters are discarded. An ambiguous condition exists if $n_1|u_1$ and $n_2|u_2$ are not both of the same type.

See: **3.2.3.2 Control-flow stack**, **6.1.0140 +LOOP**, **6.1.1240 DO**, **6.1.1680 I**, **6.1.1760 LEAVE**, **6.1.1800 LOOP**, **6.1.2380 UNLOOP**.

6.2.0700 AGAIN CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: *dest* --)

Append the run-time semantics given below to the current definition, resolving the backward reference *dest*.

Run-time: (--)

Continue execution at the location specified by *dest*. If no other control flow words are used, any program code after **AGAIN** will not be executed.

See: **6.1.0760 BEGIN**.

6.2.0855 c" "c-quote" CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: ("*ccc*<quote>" --)

Parse *ccc* delimited by " (double-quote) and append the run-time semantics given below to the current definition.

Run-time: (-- *c-addr*)

Return *c-addr*, a counted string consisting of the characters *ccc*. A program shall not alter the returned string.

See: **3.4.1 Parsing**, **6.1.2165 s"**, **11.6.1.2165 s"**.

6.2.0873 CASE

CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *case-sys*)

Mark the start of the **CASE ... OF ... ENDOF ... ENDCASE** structure. Append the run-time semantics given below to the current definition.

Run-time: (--)

Continue execution.

See: **6.2.1342 ENDCASE**, **6.2.1343 ENDOF**, **6.2.1950 OF**.

6.2.0945 COMPILE,

"compile-comma"

CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Execution: (*xt* --)

Append the execution semantics of the definition represented by *xt* to the execution semantics of the current definition.

6.2.0970 CONVERT

CORE EXT

(*ud*₁ *c-addr*₁ -- *ud*₂ *c-addr*₂)

*ud*₂ is the result of converting the characters within the text beginning at the first character after *c-addr*₁ into digits, using the number in **BASE**, and adding each digit to *ud*₁ after multiplying *ud*₁ by the number in **BASE**. Conversion continues until a character that is not convertible is encountered. *c-addr*₂ is the location of the first unconverted character. An ambiguous condition exists if *ud*₂ overflows.

Note: This word is obsolescent and is included as a concession to existing implementations. Its function is superseded by **6.1.0570 >NUMBER**.

See: **3.2.1.2 Digit conversion**.

6.2.1342 ENDCASE

"end-case"

CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: *case-sys* --)

Mark the end of the **CASE ... OF ... ENDOF ... ENDCASE** structure. Use *case-sys* to resolve the entire structure. Append the run-time semantics given below to the current definition.

Run-time: (*x* --)

Discard the case selector *x* and continue execution.

See: **6.2.0873 CASE**, **6.2.1343 ENDOF**, **6.2.1950 OF**.

6.2.1343 ENDOF "end-of" CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (*C*: *case-sys*₁ *of-sys* -- *case-sys*₂)

Mark the end of the ... **OF** ... **ENDOF** ... part of the **CASE** structure. The next location for a transfer of control resolves the reference given by *of-sys*. Append the run-time semantics given below to the current definition. Replace *case-sys*₁ with *case-sys*₂ on the control-flow stack, to be resolved by **ENDCASE**.

Run-time: (--)

Continue execution at the location specified by the consumer of *case-sys*₂.

See: **6.2.0873 CASE**, **6.2.1342 ENDCASE**, **6.2.1950 OF**.

6.2.1350 ERASE CORE EXT

(*addr u* --)

If *u* is greater than zero, clear all bits in each of *u* consecutive address units of memory beginning at *addr* .

6.2.1390 EXPECT CORE EXT

(*c-addr +n* --)

Receive a string of at most *+n* characters. Display graphic characters as they are received. A program that depends on the presence or absence of non-graphic characters in the string has an environmental dependency. The editing functions, if any, that the system performs in order to construct the string of characters are implementation-defined.

Input terminates when an implementation-defined line terminator is received or when the string is *+n* characters long. When input terminates, nothing is appended to the string and the display is maintained in an implementation-defined way.

Store the string at *c-addr* and its length in **SPAN**.

Note: This word is obsolescent and is included as a concession to existing implementations. Its function is superseded by **6.1.0695 ACCEPT**.

6.2.1485 FALSE CORE EXT

(-- *false*)

Return a *false* flag.

See: **3.1.3.1 Flags**

6.2.1660 HEX CORE EXT

(--)

Set contents of **BASE** to sixteen.

6.2.1850 MARKER

CORE EXT

("*<spaces>name*" --)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below.

name Execution: (--)

Restore all dictionary allocation and search order pointers to the state they had just prior to the definition of *name*. Remove the definition of *name* and all subsequent definitions. Restoration of any structures still existing that could refer to deleted definitions or deallocated data space is not necessarily provided. No other contextual information such as numeric base is affected.

See: **3.4.1 Parsing, 15.6.2.1580 FORGET.**

6.2.1930 NIP

CORE EXT

($x_1 x_2$ -- x_2)

Drop the first item below the top of stack.

6.2.1950 OF

CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *of-sys*)

Put *of-sys* onto the control flow stack. Append the run-time semantics given below to the current definition. The semantics are incomplete until resolved by a consumer of *of-sys* such as **ENDOF**.

Run-time: ($x_1 x_2$ -- | x_1)

If the two values on the stack are not equal, discard the top value and continue execution at the location specified by the consumer of *of-sys*, e.g., following the next **ENDOF**. Otherwise, discard both values and continue execution in line.

See: **6.2.0873 CASE, 6.2.1342 ENDCASE, 6.2.1343 ENDOF.**

6.2.2000 PAD

CORE EXT

(-- *c-addr*)

c-addr is the address of a transient region that can be used to hold data for intermediate processing.

See: **3.3.3.6 Other transient regions.**

6.2.2008 PARSE

CORE EXT

(*char* “*ccc<char>*” -- *c-addr u*)

Parse *ccc* delimited by the delimiter *char*.

c-addr is the address (within the input buffer) and *u* is the length of the parsed string. If the parse area was empty, the resulting string has a zero length.

See: **3.4.1 Parsing**.

6.2.2030 PICK

CORE EXT

($x_u \dots x_1 x_0$ *u* -- $x_u \dots x_1 x_0 x_u$)

Remove *u*. Copy the x_u to the top of the stack. An ambiguous condition exists if there are less than *u*+2 items on the stack before **PICK** is executed.

6.2.2040 QUERY

CORE EXT

(--)

Make the user input device the input source. Receive input into the terminal input buffer, replacing any previous contents. Make the result, whose address is returned by **TIB**, the input buffer. Set **>IN** to zero.

Note: This word is obsolescent and is included as a concession to existing implementations.

6.2.2125 REFILL

CORE EXT

(-- *flag*)

Attempt to fill the input buffer from the input source, returning a true flag if successful.

When the input source is the user input device, attempt to receive input into the terminal input buffer. If successful, make the result the input buffer, set **>IN** to zero, and return *true*. Receipt of a line containing no characters is considered successful. If there is no input available from the current input source, return *false*.

When the input source is a string from **EVALUATE**, return *false* and perform no other action.

See: **7.6.2.2125 REFILL**, **11.6.2.2125 REFILL**.

6.2.2148 RESTORE-INPUT

CORE EXT

($x_n \dots x_1$ *n* -- *flag*)

Attempt to restore the input source specification to the state described by x_1 through x_n . *flag* is true if the input source specification cannot be so restored.

An ambiguous condition exists if the input source represented by the arguments is not the same as the current input source.

See: **A.6.2.2182 SAVE-INPUT**.

6.2.2150 ROLL CORE EXT

$$(x_u x_{u-1} \dots x_0 \text{ -- } x_{u-1} \dots x_0 x_u)$$

Remove u . Rotate $u+1$ items on the top of the stack. An ambiguous condition exists if there are less than $u+2$ items on the stack before **ROLL** is executed.

6.2.2182 SAVE-INPUT CORE EXT

$$(\text{-- } x_n \dots x_1 n)$$

x_1 through x_n describe the current state of the input source specification for later use by **RESTORE-INPUT**.

6.2.2218 SOURCE-ID CORE EXT

"source-i-d"

$$(\text{-- } 0 \mid -1)$$

Identifies the input source as follows:

SOURCE-ID	Input source
-1	String (via EVALUATE)
0	User input device

See: **11.6.1.2218 SOURCE-ID**.

6.2.2240 SPAN CORE EXT

$$(\text{-- } a\text{-addr})$$

$a\text{-addr}$ is the address of a cell containing the count of characters stored by the last execution of **EXPECT**.

Note: This word is obsolescent and is included as a concession to existing implementations.

6.2.2290 TIB CORE EXT

"t-i-b"

$$(\text{-- } c\text{-addr})$$

$c\text{-addr}$ is the address of the terminal input buffer.

Note: This word is obsolescent and is included as a concession to existing implementations.

6.2.2295 TO CORE EXT

Interpretation: (*x* “<spaces>*name*” --)

Skip leading spaces and parse *name* delimited by a space. Store *x* in *name*. An ambiguous condition exists if *name* was not defined by **VALUE**.

Compilation: (“<spaces>*name*” --)

Skip leading spaces and parse *name* delimited by a space. Append the run-time semantics given below to the current definition. An ambiguous condition exists if *name* was not defined by **VALUE**.

Run-time: (*x* --)

Store *x* in *name*.

Note: An ambiguous condition exists if either **POSTPONE** or [**COMPILE**] is applied to **TO**.

See: **6.2.2405 VALUE**, **13.6.1.2295 TO**.

6.2.2298 TRUE CORE EXT

(-- *true*)

Return a *true* flag, a single-cell value with all bits set.

See: **3.1.3.1 Flags**.

6.2.2300 TUCK CORE EXT

(*x*₁ *x*₂ -- *x*₂ *x*₁ *x*₂)

Copy the first (top) stack item below the second stack item.

6.2.2330 U.R CORE EXT

“u-dot-r”

(*u n* --)

Display *u* right aligned in a field *n* characters wide. If the number of characters required to display *u* is greater than *n*, all digits are displayed with no leading spaces in a field as wide as necessary.

6.2.2350 U> CORE EXT

“u-greater-than”

(*u*₁ *u*₂ -- *flag*)

flag is true if and only if *u*₁ is greater than *u*₂.

See: **6.1.0540 >**.

6.2.2395 UNUSED CORE EXT

(-- *u*)

u is the amount of space remaining in the region addressed by **HERE**, in address units.

6.2.2405 VALUE

CORE EXT

 $(x \text{ “<spaces>name” } --)$

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below, with an initial value equal to *x*.

name is referred to as a “value”.

name Execution: $(-- x)$

Place *x* on the stack. The value of *x* is that given when *name* was created, until the phrase **x TO name** is executed, causing a new value of *x* to be associated with *name*.

See: **3.4.1 Parsing**.

6.2.2440 WITHIN

CORE EXT

 $(n_1|u_1 \ n_2|u_2 \ n_3|u_3 \ -- \ flag)$

Perform a comparison of a test value $n_1|u_1$ with a lower limit $n_2|u_2$ and an upper limit $n_3|u_3$, returning *true* if either $(n_2|u_2 < n_3|u_3 \text{ and } (n_2|u_2 \leq n_1|u_1 \text{ and } n_1|u_1 < n_3|u_3))$ or $(n_2|u_2 > n_3|u_3 \text{ and } (n_2|u_2 \leq n_1|u_1 \text{ or } n_1|u_1 < n_3|u_3))$ is true, returning *false* otherwise. An ambiguous condition exists if $n_1|u_1$, $n_2|u_2$, and $n_3|u_3$ are not all the same type.

6.2.2530 [COMPILE]

“bracket-compile”

CORE EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: $(\text{“<spaces>name” } --)$

Skip leading space delimiters. Parse *name* delimited by a space. Find *name*. If *name* has other than default compilation semantics, append them to the current definition; otherwise append the execution semantics of *name*. An ambiguous condition exists if *name* is not found.

See: **3.4.1 Parsing**.

**6.2.2535 **

“backslash”

CORE EXT

Compilation: Perform the execution semantics given below.

Execution: $(\text{“ccc<eol>”}--)$

Parse and discard the remainder of the parse area. \ is an immediate word.

See: **7.6.2.2535 **.

7. The optional Block word set

7.1 Introduction

7.2 Additional terms

block: 1024 characters of data on mass storage, designated by a block number.

block buffer: a block-sized region of data space where a block is made temporarily available for use.

The current block buffer is the block buffer most recently accessed by **BLOCK**, **BUFFER**, **LOAD**, **LIST**, or **THRU**.

7.3 Additional usage requirements

7.3.1 Environmental queries

Append table 7.1 to table 3.5.

See: **3.2.6 Environmental queries**.

Table 7.1 – Environmental Query Strings

String	Value data type	Constant?	Meaning
BLOCK	<i>flag</i>	no	block word set present
BLOCK-EXT	<i>flag</i>	no	block extensions word set present

7.3.2 Data space

A program may access memory within a valid block buffer.

See: **3.3.3 Data Space**.

7.3.3 Block buffer regions

The address of a block buffer returned by **BLOCK** or **BUFFER** is transient. A call to **BLOCK** or **BUFFER** may render a previously-obtained block-buffer address invalid, as may a call to any word that:

- parses;
- displays characters on the user output device, such as **TYPE** or **EMIT**;
- controls the user output device, such as **CR** or **AT-XY**;
- receives or tests for the presence of characters from the user input device such as **ACCEPT** or **KEY**;
- waits for a condition or event, such as **MS** or **EKEY**;
- manages the block buffers, such as **FLUSH**, **SAVE-BUFFERS**, or **EMPTY-BUFFERS**;
- performs any operation on a file or file-name directory that implies I/O, such as **REFILL** or any word that returns an *ior*;
- implicitly performs I/O, such as text interpreter nesting and un-nesting when files are being used (including un-nesting implied by **THROW**).

If the input source is a block, these restrictions also apply to the address returned by **SOURCE**.

Block buffers are uniquely assigned to blocks.

7.3.4 Parsing

The Block word set implements an alternative input source for the text interpreter. When the input source is a block, **BLK** shall contain the non-zero block number and the input buffer is the 1024-character buffer containing that block.

A block is conventionally displayed as 16 lines of 64 characters.

A program may switch the input source to a block by using **LOAD** or **THRU**. Input sources may be nested using **LOAD** and **EVALUATE** in any order.

A program may reposition the parse area within a block by manipulating **>IN**. More extensive repositioning can be accomplished using **SAVE-INPUT** and **RESTORE-INPUT**.

See: **3.4.1 Parsing**.

7.3.5 Possible action on an ambiguous condition

See: **3.4.4 Possible action on an ambiguous condition**.

- A system with the Block word set may set interpretation state and interpret a block.

7.4 Additional documentation requirements

7.4.1 System documentation

7.4.1.1 Implementation-defined options

- the format used for display by **7.6.2.1770 LIST** (if implemented);
- the length of a line affected by **7.6.2.2535 ** (if implemented).

7.4.1.2 Ambiguous conditions

- Correct block read was not possible;
- I/O exception in block transfer;
- Invalid block number (**7.6.1.0800 BLOCK**, **7.6.1.0820 BUFFER**, **7.6.1.1790 LOAD**);
- * A program directly alters the contents of **7.6.1.0790 BLK**;
- * No current block buffer for **7.6.1.2400 UPDATE**.

7.4.1.3 Other system documentation

- any restrictions a multiprogramming system places on the use of buffer addresses;
- the number of blocks available for source text and data.

7.4.2 Program documentation

- the number of blocks required by the program.

7.5 Compliance and labeling

7.5.1 ANS Forth systems

The phrase “Providing the Block word set” shall be appended to the label of any Standard System that provides all of the Block word set.

The phrase “Providing *name(s)* from the Block Extensions word set” shall be appended to the label of any Standard System that provides portions of the Block Extensions word set.

The phrase “Providing the Block Extensions word set” shall be appended to the label of any Standard System that provides all of the Block and Block Extensions word sets.

7.5.2 ANS Forth programs

The phrase “Requiring the Block word set” shall be appended to the label of Standard Programs that require the system to provide the Block word set.

The phrase “Requiring *name(s)* from the Block Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Block Extensions word set.

The phrase “Requiring the Block Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Block and Block Extensions word sets.

7.6 Glossary

7.6.1 Block words

7.6.1.0790 **BLK** “b-l-k” **BLOCK**

(-- *a-addr*)

a-addr is the address of a cell containing zero or the number of the mass-storage block being interpreted. If **BLK** contains zero, the input source is not a block and can be identified by **SOURCE-ID**, if **SOURCE-ID** is available. An ambiguous condition exists if a program directly alters the contents of **BLK**.

See: **7.3.3 Block buffer regions**.

7.6.1.0800 **BLOCK** **BLOCK**

(*u* -- *a-addr*)

a-addr is the address of the first character of the block buffer assigned to mass-storage block *u*. An ambiguous condition exists if *u* is not an available block number.

If block *u* is already in a block buffer, *a-addr* is the address of that block buffer.

If block *u* is not already in memory and there is an unassigned block buffer, transfer block *u* from mass storage to an unassigned block buffer. *a-addr* is the address of that block buffer.

If block *u* is not already in memory and there are no unassigned block buffers, unassign a block buffer. If the block in that buffer has been **UPDATED**, transfer the block to mass storage and transfer block *u* from mass storage into that buffer. *a-addr* is the address of that block buffer.

At the conclusion of the operation, the block buffer pointed to by *a-addr* is the current block buffer and is assigned to *u*.

7.6.1.0820 BUFFER **BLOCK**

(*u* -- *a-addr*)

a-addr is the address of the first character of the block buffer assigned to block *u*. The contents of the block are unspecified. An ambiguous condition exists if *u* is not an available block number.

If block *u* is already in a block buffer, *a-addr* is the address of that block buffer.

If block *u* is not already in memory and there is an unassigned buffer, *a-addr* is the address of that block buffer.

If block *u* is not already in memory and there are no unassigned block buffers, unassign a block buffer. If the block in that buffer has been **UPDATED**, transfer the block to mass storage. *a-addr* is the address of that block buffer.

At the conclusion of the operation, the block buffer pointed to by *a-addr* is the current block buffer and is assigned to *u*.

See: **7.6.1.0800 BLOCK**.

7.6.1.1360 EVALUATE **BLOCK**

Extend the semantics of **6.1.1360 EVALUATE** to include:

Store zero in **BLK**.

7.6.1.1559 FLUSH **BLOCK**

(--)

Perform the function of **SAVE-BUFFERS** then unassign all block buffers.

7.6.1.1790 LOAD **BLOCK**

(*i*x* *u* -- *j*x*)

Save the current input-source specification. Store *u* in **BLK** (thus making block *u* the input source and setting the input buffer to encompass its contents), set **>IN** to zero, and interpret. When the parse area is exhausted, restore the prior input source specification. Other stack effects are due to the words **LOAD**ed.

An ambiguous condition exists if *u* is zero or is not a valid block number.

See: **3.4 The Forth text interpreter**.

7.6.1.2180 SAVE-BUFFERS **BLOCK**

(--)

Transfer the contents of each **UPDATED** block buffer to mass storage. Mark all buffers as unmodified.

7.6.1.2400 UPDATE **BLOCK**

(--)

Mark the current block buffer as modified. An ambiguous condition exists if there is no current block buffer.

UPDATE does not immediately cause I/O.

See: **7.6.1.0800 BLOCK**, **7.6.1.0820 BUFFER**, **7.6.1.1559 FLUSH**, **7.6.1.2180 SAVE-BUFFERS**.

7.6.2 Block extension words**7.6.2.1330 EMPTY-BUFFERS** **BLOCK EXT**

(--)

Unassign all block buffers. Do not transfer the contents of any **UPDATED** block buffer to mass storage.

See: **7.6.1.0800 BLOCK**.

7.6.2.1770 LIST **BLOCK EXT**

(*u* --)

Display block *u* in an implementation-defined format. Store *u* in **SCR**.

See: **7.6.1.0800 BLOCK**.

7.6.2.2125 REFILL **BLOCK EXT**

(-- *flag*)

Extend the execution semantics of **6.2.2125 REFILL** with the following:

When the input source is a block, make the next block the input source and current input buffer by adding one to the value of **BLK** and setting **>IN** to zero. Return *true* if the new value of **BLK** is a valid block number, otherwise *false*.

See: **6.2.2125 REFILL**, **11.6.2.2125 REFILL**.

7.6.2.2190 SCR **BLOCK EXT**

“s-c-r”

(-- *a-addr*)

a-addr is the address of a cell containing the block number of the block most recently **LISTED**.

7.6.2.2280 THRU **BLOCK EXT**

(*i*x* *u*₁ *u*₂ -- *j*x*)

LOAD the mass storage blocks numbered *u*₁ through *u*₂ in sequence. Other stack effects are due to the words **LOADED**.

**7.6.2.2535 **

"backslash"

BLOCK EXT

Extend the semantics of **6.1.2535 ** to be:

Compilation: Perform the execution semantics given below.

Execution: ("ccc<eol>"--)

If **BLK** contains zero, parse and discard the remainder of the parse area; otherwise parse and discard the portion of the parse area corresponding to the remainder of the current line. \ is an immediate word.

See: **6.2.2535 **.

8. The optional Double-Number word set

8.1 Introduction

Sixteen-bit Forth systems often use double-length numbers. However, many Forths on small embedded systems do not, and many users of Forth on systems with a cell size of 32 bits or more find that the use of double-length numbers is much diminished. Therefore, the words that manipulate double-length entities have been placed in this optional word set.

8.2 Additional terms and notation

None.

8.3 Additional usage requirements

8.3.1 Environmental queries

Append table 8.1 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 8.1 – Environmental Query Strings

String	Value data type	Constant?	Meaning
DOUBLE	<i>flag</i>	no	double-number word set present
DOUBLE-EXT	<i>flag</i>	no	double-number extensions word set present

8.3.2 Text interpreter input number conversion

When the text interpreter processes a number that is immediately followed by a decimal point and is not found as a definition name, the text interpreter shall convert it to a double-cell number.

For example, entering **DECIMAL 1234** leaves the single-cell number **1234** on the stack, and entering **DECIMAL 1234.** leaves the double-cell number **1234 0** on the stack.

See: **3.4.1.3 Text interpreter input number conversion.**

8.4 Additional documentation requirements

8.4.1 System documentation

8.4.1.1 Implementation-defined options

- no additional requirements.

8.4.1.2 Ambiguous conditions

- * *d* outside range of *n* in **8.6.1.1140 D>S**.

8.4.1.3 Other system documentation

- no additional requirements.

8.4.2 Program documentation

- no additional requirements.

8.5 Compliance and labeling

8.5.1 ANS Forth systems

The phrase “Providing the Double-Number word set” shall be appended to the label of any Standard System that provides all of the Double-Number word set.

The phrase “Providing *name(s)* from the Double-Number Extensions word set” shall be appended to the label of any Standard System that provides portions of the Double-Number Extensions word set.

The phrase “Providing the Double-Number Extensions word set” shall be appended to the label of any Standard System that provides all of the Double-Number and Double-Number Extensions word sets.

8.5.2 ANS Forth programs

The phrase “Requiring the Double-Number word set” shall be appended to the label of Standard Programs that require the system to provide the Double-Number word set.

The phrase “Requiring *name(s)* from the Double-Number Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Double-Number Extensions word set.

The phrase “Requiring the Double-Number Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Double-Number and Double-Number Extensions word sets.

8.6 Glossary

8.6.1 Double-Number words

8.6.1.0360 **2CONSTANT** “two-constant” **DOUBLE**

(x_1 x_2 “<spaces>name” --)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below.

name is referred to as a “two-constant.”

name Execution: (-- x_1 x_2)

Place cell pair x_1 x_2 on the stack.

See: **3.4.1 Parsing**.

8.6.1.0390 **2LITERAL** “two-literal” **DOUBLE**

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (x_1 x_2 --)

Append the run-time semantics below to the current definition.

Run-time: (-- x_1 x_2)

Place cell pair x_1 x_2 on the stack.

8.6.1.0440	2VARIABLE	“two-variable”	DOUBLE
	(“<spaces>name” --)		
	Skip leading space delimiters. Parse <i>name</i> delimited by a space. Create a definition for <i>name</i> with the execution semantics defined below. Reserve two consecutive cells of data space.		
	<i>name</i> is referred to as a “two-variable.”		
	<i>name</i> Execution: (-- <i>a-addr</i>)		
	<i>a-addr</i> is the address of the first (lowest address) cell of two consecutive cells in data space reserved by 2VARIABLE when it defined <i>name</i> . A program is responsible for initializing the contents.		
	See: 3.4.1 Parsing, 6.1.2410 VARIABLE.		
8.6.1.1040	D+	“d-plus”	DOUBLE
	(<i>d</i> ₁ <i>ud</i> ₁ <i>d</i> ₂ <i>ud</i> ₂ -- <i>d</i> ₃ <i>ud</i> ₃)		
	Add <i>d</i> ₂ <i>ud</i> ₂ to <i>d</i> ₁ <i>ud</i> ₁ , giving the sum <i>d</i> ₃ <i>ud</i> ₃ .		
8.6.1.1050	D-	“d-minus”	DOUBLE
	(<i>d</i> ₁ <i>ud</i> ₁ <i>d</i> ₂ <i>ud</i> ₂ -- <i>d</i> ₃ <i>ud</i> ₃)		
	Subtract <i>d</i> ₂ <i>ud</i> ₂ from <i>d</i> ₁ <i>ud</i> ₁ , giving the difference <i>d</i> ₃ <i>ud</i> ₃ .		
8.6.1.1060	D.	“d-dot”	DOUBLE
	(<i>d</i> --)		
	Display <i>d</i> in free field format.		
8.6.1.1070	D.R	“d-dot-r”	DOUBLE
	(<i>d</i> <i>n</i> --)		
	Display <i>d</i> right aligned in a field <i>n</i> characters wide. If the number of characters required to display <i>d</i> is greater than <i>n</i> , all digits are displayed with no leading spaces in a field as wide as necessary.		
8.6.1.1075	D0<	“d-zero-less”	DOUBLE
	(<i>d</i> -- <i>flag</i>)		
	<i>flag</i> is true if and only if <i>d</i> is less than zero.		
8.6.1.1080	D0=	“d-zero-equals”	DOUBLE
	(<i>xd</i> -- <i>flag</i>)		
	<i>flag</i> is true if and only if <i>xd</i> is equal to zero.		

8.6.1.1090	D2* (xd_1 -- xd_2) xd_2 is the result of shifting xd_1 one bit toward the most-significant bit, filling the vacated least-significant bit with zero.	“d-two-star”	DOUBLE
8.6.1.1100	D2/ (xd_1 -- xd_2) xd_2 is the result of shifting xd_1 one bit toward the least-significant bit, leaving the most-significant bit unchanged.	“d-two-slash”	DOUBLE
8.6.1.1110	D< ($d_1 d_2$ -- <i>flag</i>) <i>flag</i> is true if and only if d_1 is less than d_2 .	“d-less-than”	DOUBLE
8.6.1.1120	D= ($xd_1 xd_2$ -- <i>flag</i>) <i>flag</i> is true if and only if xd_1 is bit-for-bit the same as xd_2 .	“d-equals”	DOUBLE
8.6.1.1140	D>S (d -- n) n is the equivalent of d . An ambiguous condition exists if d lies outside the range of a signed single-cell number.	“d-to-s”	DOUBLE
8.6.1.1160	DABS (d -- ud) ud is the absolute value of d .	“d-abs”	DOUBLE
8.6.1.1210	DMAX ($d_1 d_2$ -- d_3) d_3 is the greater of d_1 and d_2 .	“d-max”	DOUBLE
8.6.1.1220	DMIN ($d_1 d_2$ -- d_3) d_3 is the lesser of d_1 and d_2 .	“d-min”	DOUBLE

8.6.1.1230	DNEGATE	“d-negate”	DOUBLE
	$(d_1 \text{ -- } d_2)$		
	d_2 is the negation of d_1 .		
8.6.1.1820	M* /	“m-star-slash”	DOUBLE
	$(d_1 n_1 + n_2 \text{ -- } d_2)$		
	Multiply d_1 by n_1 producing the triple-cell intermediate result t . Divide t by $+n_2$ giving the double-cell quotient d_2 . An ambiguous condition exists if $+n_2$ is zero or negative, or the quotient lies outside of the range of a double-precision signed integer.		
8.6.1.1830	M+	“m-plus”	DOUBLE
	$(d_1 ud_1 n \text{ -- } d_2 ud_2)$		
	Add n to $d_1 ud_1$, giving the sum $d_2 ud_2$.		
8.6.2 Double-Number extension words			
8.6.2.0420	2ROT	“two-rote”	DOUBLE EXT
	$(x_1 x_2 x_3 x_4 x_5 x_6 \text{ -- } x_3 x_4 x_5 x_6 x_1 x_2)$		
	Rotate the top three cell pairs on the stack bringing cell pair $x_1 x_2$ to the top of the stack.		
8.6.2.1270	DU<	“d-u-less”	DOUBLE EXT
	$(ud_1 ud_2 \text{ -- } flag)$		
	$flag$ is true if and only if ud_1 is less than ud_2 .		

9. The optional Exception word set

9.1 Introduction

9.2 Additional terms and notation

None.

9.3 Additional usage requirements

9.3.1 THROW values

The **THROW** values $\{-255\dots-1\}$ shall be used only as assigned by this Standard. The values $\{-4095\dots-256\}$ shall be used only as assigned by a system.

If the File-Access or Memory-Allocation word sets are implemented, it is recommended that the non-zero values of *ior* lie within the range of system **THROW** values, as defined above. In an operating-system environment, this can sometimes be accomplished by “biasing” the range of operating-system exception-codes to fall within the **THROW** range.

Programs shall not define values for use with **THROW** in the range $\{-4095\dots-1\}$.

9.3.2 Exception frame

An exception frame is the implementation-dependent set of information recording the current execution state necessary for the proper functioning of **CATCH** and **THROW**. It often includes the depths of the data stack and return stack.

9.3.3 Exception stack

A stack used for the nesting of exception frames by **CATCH** and **THROW**. It may be, but need not be, implemented using the return stack.

9.3.4 Environmental queries

Append table 9.1 to table 3.5.

See: **3.2.6 Environmental queries**.

Table 9.1 – Environmental query strings

String	Value data type	Constant?	Meaning
EXCEPTION	<i>flag</i>	no	Exception word set present
EXCEPTION-EXT	<i>flag</i>	no	Exception extensions word set present

9.3.5 Possible actions on an ambiguous condition

A system choosing to execute **THROW** when detecting one of the ambiguous conditions listed in table 9.3.6 shall use the throw code listed there.

See: **3.4.4 Possible actions on an ambiguous condition**.

Table 9.2 – **THROW** code assignments

Code	Reserved for	Code	Reserved for
-1	ABORT	-3	stack overflow
-2	ABORT"	-4	stack underflow

-5	return stack overflow	-32	invalid name argument (e.g., TO xxx)
-6	return stack underflow	-33	block read exception
-7	do-loops nested too deeply during execution	-34	block write exception
-8	dictionary overflow	-35	invalid block number
-9	invalid memory address	-36	invalid file position
-10	division by zero	-37	file I/O exception
-11	result out of range	-38	non-existent file
-12	argument type mismatch	-39	unexpected end of file
-13	undefined word	-40	invalid BASE for floating point conversion
-14	interpreting a compile-only word	-41	loss of precision
-15	invalid FORGET	-42	floating-point divide by zero
-16	attempt to use zero-length string as a name	-43	floating-point result out of range
-17	pictured numeric output string overflow	-44	floating-point stack overflow
-18	parsed string overflow	-45	floating-point stack underflow
-19	definition name too long	-46	floating-point invalid argument
-20	write to a read-only location	-47	compilation word list deleted
-21	unsupported operation (e.g., AT-XY on a too-dumb terminal)	-48	invalid POSTPONE
-22	control structure mismatch	-49	search-order overflow
-23	address alignment exception	-50	search-order underflow
-24	invalid numeric argument	-51	compilation word list changed
-25	return stack imbalance	-52	control-flow stack overflow
-26	loop parameters unavailable	-53	exception stack overflow
-27	invalid recursion	-54	floating-point underflow
-28	user interrupt	-55	floating-point unidentified fault
-29	compiler nesting	-56	QUIT
-30	obsolescent feature	-57	exception in sending or receiving a character
-31	>BODY used on non- CREATED definition	-58	[IF] , [ELSE] , or [THEN] exception

9.3.6 Exception handling

There are several methods of coupling **CATCH** and **THROW** to other procedural nestings. The usual nestings are the execution of definitions, use of the return stack, use of loops, instantiation of locals and nesting of input sources (i.e., with **LOAD**, **EVALUATE**, or **INCLUDE-FILE**).

When a **THROW** returns control to a **CATCH**, the system shall un-nest not only definitions, but also, if present, locals and input source specifications, to return the system to its proper state for continued execution past the **CATCH**.

9.4 Additional documentation requirements

9.4.1 System documentation

9.4.1.1 Implementation-defined options

- Values used in the system by 9.6.1.0875 **CATCH** and 9.6.1.2275 **THROW** (9.3.1 **THROW** values, 9.3.5 Possible actions on an ambiguous condition).

9.4.1.2 Ambiguous conditions

- no additional requirements.

9.4.1.3 Other system documentation

- no additional requirements.

9.4.2 Program documentation

- no additional requirements.

9.5 Compliance and labeling

9.5.1 ANS Forth systems

The phrase “Providing the Exception word set” shall be appended to the label of any Standard System that provides all of the Exception word set.

The phrase “Providing *name(s)* from the Exception Extensions word set” shall be appended to the label of any Standard System that provides portions of the Exception Extensions word set.

The phrase “Providing the Exception Extensions word set” shall be appended to the label of any Standard System that provides all of the Exception and Exception Extensions word sets.

9.5.2 ANS Forth programs

The phrase “Requiring the Exception word set” shall be appended to the label of Standard Programs that require the system to provide the Exception word set.

The phrase “Requiring *name(s)* from the Exception Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Exception Extensions word set.

The phrase “Requiring the Exception Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Exception and Exception Extensions word sets.

9.6 Glossary

9.6.1 Exception words

9.6.1.0875 CATCH

EXCEPTION

(*i*x xt -- j*x 0 | i*x n*)

Push an exception frame on the exception stack and then execute the execution token *xt* (as with **EXECUTE**) in such a way that control can be transferred to a point just after **CATCH** if **THROW** is executed during the execution of *xt*.

If the execution of *xt* completes normally (i.e., the exception frame pushed by this **CATCH** is not popped by an execution of **THROW**) pop the exception frame and return zero on top of the data stack, above whatever stack items would have been returned by *xt* **EXECUTE**. Otherwise, the remainder of the execution semantics are given by **THROW**.

9.6.1.2275 THROW

EXCEPTION

$$(k*x\ n\ --\ k*x\ | \ i*x\ n)$$

If any bits of n are non-zero, pop the topmost exception frame from the exception stack, along with everything on the return stack above that frame. Then restore the input source specification in use before the corresponding **CATCH** and adjust the depths of all stacks defined by this Standard so that they are the same as the depths saved in the exception frame (i is the same number as the i in the input arguments to the corresponding **CATCH**), put n on top of the data stack, and transfer control to a point just after the **CATCH** that pushed that exception frame.

If the top of the stack is non zero and there is no exception frame on the exception stack, the behavior is as follows:

If n is minus-one (-1), perform the function of **6.1.0670 ABORT** (the version of **ABORT** in the Core word set), displaying no message.

If n is minus-two, perform the function of **6.1.0680 ABORT"** (the version of **ABORT"** in the Core word set), displaying the characters *ccc* associated with the **ABORT"** that generated the **THROW**.

Otherwise, the system may display an implementation-dependent message giving information about the condition associated with the **THROW** code n . Subsequently, the system shall perform the function of **6.1.0670 ABORT** (the version of **ABORT** in the Core word set).

9.6.2 Exception extension words**9.6.2.0670 ABORT**

EXCEPTION EXT

Extend the semantics of **6.1.0670 ABORT** to be:

$$(i*x\ --) (R: j*x\ --)$$

Perform the function of **-1 THROW**.

See: **6.1.0670 ABORT**.

9.6.2.0680 ABORT"

"abort-quote"

EXCEPTION EXT

Extend the semantics of **6.1.0680 ABORT"** to be:

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (*ccc*<quote> --)

Parse *ccc* delimited by a " (double-quote). Append the run-time semantics given below to the current definition.

Run-time: ($i*x\ x_j\ --\ | \ i*x$) (R: $j*x\ --\ | \ j*x$)

Remove x_j from the stack. If any bit of x_j is not zero, perform the function of **-2 THROW**, displaying *ccc* if there is no exception frame on the exception stack.

See: **3.4.1 Parsing, 6.1.0680 ABORT"**.

10. The optional Facility word set

10.1 Introduction

10.2 Additional terms and notation

None.

10.3 Additional usage requirements

10.3.1 Character types

Programs that use more than seven bits of a character by **10.6.2.1305 EKEY** have an environmental dependency.

See: **3.1.2 Character types**.

10.3.2 Environmental queries

Append table 10.1 to table 3.5.

See: **3.2.6 Environmental queries**.

Table 10.1 – Environmental query strings

String	Value data type	Constant?	Meaning
FACILITY	<i>flag</i>	no	facility word set present
FACILITY-EXT	<i>flag</i>	no	facility extensions word set present

10.4 Additional documentation requirements

10.4.1 System documentation

10.4.1.1 Implementation-defined options

- encoding of keyboard events (**10.6.2.1305 EKEY**);
- duration of a system clock tick;
- repeatability to be expected from execution of **10.6.2.1905 MS**.

10.4.1.2 Ambiguous conditions

- * **10.6.1.0742 AT-XY** operation can't be performed on user output device.

10.4.1.3 Other system documentation

- no additional requirements.

10.4.2 Program documentation

10.4.2.1 Environmental dependencies

- using more than seven bits of a character in **10.6.2.1305 EKEY**.

10.4.2.2 Other program documentation

- no additional requirements.

10.5 Compliance and labeling

10.5.1 ANS Forth systems

The phrase “Providing the Facility word set” shall be appended to the label of any Standard System that provides all of the Facility word set.

The phrase “Providing *name(s)* from the Facility Extensions word set” shall be appended to the label of any Standard System that provides portions of the Facility Extensions word set.

The phrase “Providing the Facility Extensions word set” shall be appended to the label of any Standard System that provides all of the Facility and Facility Extensions word sets.

10.5.2 ANS Forth programs

The phrase “Requiring the Facility word set” shall be appended to the label of Standard Programs that require the system to provide the Facility word set.

The phrase “Requiring *name(s)* from the Facility Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Facility Extensions word set.

The phrase “Requiring the Facility Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Facility and Facility Extensions word sets.

10.6 Glossary

10.6.1 Facility words

10.6.1.0742	AT-XY	“at-x-y”	FACILITY
	(<i>u₁ u₂ --</i>)		
	Perform implementation-dependent steps so that the next character displayed will appear in column <i>u₁</i> , row <i>u₂</i> of the user output device, the upper left corner of which is column zero, row zero. An ambiguous condition exists if the operation cannot be performed on the user output device with the specified parameters.		
10.6.1.1755	KEY?	“key-question”	FACILITY
	(<i>-- flag</i>)		
	If a character is available, return <i>true</i> . Otherwise, return <i>false</i> . If non-character keyboard events are available before the first valid character, they are discarded and are subsequently unavailable. The character shall be returned by the next execution of KEY .		
	After KEY? returns with a value of <i>true</i> , subsequent executions of KEY? prior to the execution of KEY or EKEY also return <i>true</i> , without discarding keyboard events.		
10.6.1.2005	PAGE		FACILITY
	(<i>--</i>)		
	Move to another page for output. Actual function depends on the output device. On a terminal, PAGE clears the screen and resets the cursor position to the upper left corner. On a printer, PAGE performs a form feed.		

10.6.2 Facility extension words

10.6.2.1305 EKEY “e-key” FACILITY EXT

(-- *u*)

Receive one keyboard event *u*. The encoding of keyboard events is implementation defined.

See: **10.6.1.1755 KEY?**, **6.1.1750 KEY**.

10.6.2.1306 EKEY>CHAR “e-key-to-char” FACILITY EXT

(*u* -- *u false* | *char true*)

If the keyboard event *u* corresponds to a character in the implementation-defined character set, return that character and *true*. Otherwise return *u* and *false*.

10.6.2.1307 EKEY? “e-key-question” FACILITY EXT

(-- *flag*)

If a keyboard event is available, returns *true*. Otherwise returns *false*. The event shall be returned by the next execution of **EKEY**.

After **EKEY?** returns with a value of *true*, subsequent executions of **EKEY?** prior to the execution of **KEY**, **KEY?** or **EKEY** also return *true*, referring to the same event.

10.6.2.1325 EMIT? “emit-question” FACILITY EXT

(-- *flag*)

flag is true if the user output device is ready to accept data and the execution of **EMIT** in place of **EMIT?** would not have suffered an indefinite delay. If the device status is indeterminate, *flag* is true.

10.6.2.1905 MS FACILITY EXT

(*u* --)

Wait at least *u* milliseconds.

Note: The actual length and variability of the time period depends upon the implementation-defined resolution of the system clock and upon other system and computer characteristics beyond the scope of this Standard.

10.6.2.2292 TIME&DATE “time-and-date” FACILITY EXT

(-- +*n*₁ +*n*₂ +*n*₃ +*n*₄ +*n*₅ +*n*₆)

Return the current time and date. +*n*₁ is the second {0...59}, +*n*₂ is the minute {0...59}, +*n*₃ is the hour {0...23}, +*n*₄ is the day {1...31} +*n*₅ is the month {1...12}, and +*n*₆ is the year (e.g., 1991).

11. The optional File-Access word set

11.1 Introduction

These words provide access to mass storage in the form of “files” under the following assumptions:

- files are provided by a host operating system;
- file names are represented as character strings;
- the format of file names is determined by the host operating system;
- an open file is identified by a single-cell file identifier (*fileid*);
- file-state information (e.g., position, size) is managed by the host operating system;
- file contents are accessed as a sequence of characters;
- file read operations return an actual transfer count, which can differ from the requested transfer count.

11.2 Additional terms

file-access method: a permissible means of accessing a file, such as “read/write” or “read only”.

file position: the character offset from the start of the file.

input file: the file, containing a sequence of lines, that is the input source.

11.3 Additional usage requirements

11.3.1 Data types

Append table 11.1 to table 3.1.

Table 11.1 – Data types

Symbol	Data type	Size on stack
<i>ior</i>	I/O results	1 cell
<i>fam</i>	file access method	1 cell
<i>fileid</i>	file identifiers	1 cell

11.3.1.1 File identifiers

File identifiers are implementation-dependent single-cell values that are passed to file operators to designate specific files. Opening a file assigns a file identifier, which remains valid until closed.

11.3.1.2 I/O results

I/O results are single-cell numbers indicating the result of I/O operations. A value of zero indicates that the I/O operation completed successfully; other values and their meanings are implementation-defined. Reaching the end of a file shall be reported as zero.

An I/O exception in the execution of a File-Access word that can return an I/O result shall not cause a **THROW**; exception indications are returned in the *ior*.

11.3.1.3 File access methods

File access methods are implementation-defined single-cell values.

11.3.1.4 File names

A character string containing the name of the file. The file name may include an implementation-dependent path name. The format of file names is implementation defined.

11.3.2 Blocks in files

If the File-Access word set is implemented, the Block word set shall be implemented.

Blocks may, but need not, reside in files. When they do:

- Block numbers may be mapped to one or more files by implementation-defined means. An ambiguous condition exists if a requested block number is not currently mapped;
- An **UPDATED** block that came from a file shall be transferred back to the same file.

11.3.3 Environmental queries

Append table 11.2 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 11.2 – Environmental query strings

String	Value data type	Constant?	Meaning
FILE	<i>flag</i>	no	file word set present
FILE-EXT	<i>flag</i>	no	file extensions word set present

11.3.4 Input source

The File-Access word set creates another input source for the text interpreter. When the input source is a text file, **BLK** shall contain zero, **SOURCE-ID** shall contain the *fileid* of that text file, and the input buffer shall contain one line of the text file.

Input with **INCLUDED**, **INCLUDE-FILE**, **LOAD** and **EVALUATE** shall be nestable in any order to at least eight levels.

A program that uses more than eight levels of input-file nesting has an environmental dependency.

See: **3.3.3.5 Input buffers, 9. Optional Exception word set.**

11.3.5 Other transient regions

The list of words using memory in transient regions is extended to include **11.6.1.2165 s**.

See: **3.3.3.6 Other transient regions.**

11.3.6 Parsing

When parsing from a text file using a space delimiter, control characters shall be treated the same as the space character.

Lines of at least 128 characters shall be supported. A program that requires lines of more than 128 characters has an environmental dependency.

A program may reposition the parse area within the input buffer by manipulating the contents of **>IN**. More extensive repositioning can be accomplished using **SAVE-INPUT** and **RESTORE-INPUT**.

See: **3.4.1 Parsing.**

11.4 Additional documentation requirements

11.4.1 System documentation

11.4.1.1 Implementation-defined options

- file access methods used by **11.6.1.0765 BIN**, **11.6.1.1010 CREATE-FILE**, **11.6.1.1970 OPEN-FILE**, **11.6.1.2054 R/O**, **11.6.1.2056 R/W**, and **11.6.1.2425 W/O**;
- file exceptions;
- file line terminator (**11.6.1.2090 READ-LINE**);
- file name format (**11.3.1.4 File names**);
- information returned by **11.6.2.1524 FILE-STATUS**;
- input file state after an exception (**11.6.1.1717 INCLUDE-FILE**, **11.6.1.1718 INCLUDED**);
- *ior* values and meaning (**11.3.1.2 I/O results**);
- maximum depth of file input nesting (**11.3.4 Input source**);
- maximum size of input line (**11.3.6 Parsing**);
- methods for mapping block ranges to files (**11.3.2 Blocks in files**);
- number of string buffers provided (**11.6.1.2165 S"**);
- size of string buffer used by **11.6.1.2165 S"**.

11.4.1.2 Ambiguous conditions

- * attempting to position a file outside its boundaries (**11.6.1.2142 REPOSITION-FILE**);
- * attempting to read from file positions not yet written (**11.6.1.2080 READ-FILE**, **11.6.1.2090 READ-LINE**);
- * *fileid* is invalid (**11.6.1.1717 INCLUDE-FILE**);
- * I/O exception reading or closing *fileid* (**11.6.1.1717 INCLUDE-FILE**, **11.6.1.1718 INCLUDED**);
- * named file cannot be opened (**11.6.1.1718 INCLUDED**);
- * requesting an unmapped block number (**11.3.2 Blocks in files**);
- * using **11.6.1.2218 SOURCE-ID** when **7.6.1.0790 BLK** is not zero.

11.4.1.3 Other system documentation

- no additional requirements.

11.4.2 Program documentation

11.4.2.1 Environmental dependencies

- requiring lines longer than 128 characters (**11.3.6 Parsing**);
- using more than eight levels of input-file nesting (**11.3.4 Input source**).

11.4.2.2 Other program documentation

- no additional requirements.

11.5 Compliance and labeling

11.5.1 ANS Forth systems

The phrase “Providing the File Access word set” shall be appended to the label of any Standard System that provides all of the File Access word set.

The phrase “Providing *name(s)* from the File Access Extensions word set” shall be appended to the label of any Standard System that provides portions of the File Access Extensions word set.

The phrase “Providing the File Access Extensions word set” shall be appended to the label of any Standard System that provides all of the File Access and File Access Extensions word sets.

11.5.2 ANS Forth programs

The phrase “Requiring the File Access word set” shall be appended to the label of Standard Programs that require the system to provide the File Access word set.

The phrase “Requiring *name(s)* from the File Access Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the File Access Extensions word set.

The phrase “Requiring the File Access Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the File Access and File Access Extensions word sets.

11.6 Glossary

11.6.1 File Access words

11.6.1.0080 (“paren” FILE

(“*ccc*<*paren*>” --)

Extend the semantics of **6.1.0080** (to include:

When parsing from a text file, if the end of the parse area is reached before a right parenthesis is found, refill the input buffer from the next line of the file, set **>IN** to zero, and resume parsing, repeating this process until either a right parenthesis is found or the end of the file is reached.

11.6.1.0765 **BIN** FILE

(*fam*₁ -- *fam*₂)

Modify the implementation-defined file access method *fam*₁ to additionally select a “binary”, i.e., not line oriented, file access method, giving access method *fam*₂.

See: **11.6.1.2054** R/O, **11.6.1.2056** R/W, **11.6.1.2425** W/O.

11.6.1.0900 **CLOSE-FILE** FILE

(*fileid* -- *ior*)

Close the file identified by *fileid*. *ior* is the implementation-defined I/O result code.

11.6.1.1010 CREATE-FILE FILE

(*c-addr u fam -- fileid ior*)

Create the file named in the character string specified by *c-addr* and *u*, and open it with file access method *fam*. The meaning of values of *fam* is implementation defined. If a file with the same name already exists, recreate it as an empty file.

If the file was successfully created and opened, *ior* is zero, *fileid* is its identifier, and the file has been positioned to the start of the file.

Otherwise, *ior* is the implementation-defined I/O result code and *fileid* is undefined.

11.6.1.1190 DELETE-FILE FILE

(*c-addr u -- ior*)

Delete the file named in the character string specified by *c-addr u*. *ior* is the implementation-defined I/O result code.

11.6.1.1520 FILE-POSITION FILE

(*fileid -- ud ior*)

ud is the current file position for the file identified by *fileid*. *ior* is the implementation-defined I/O result code. *ud* is undefined if *ior* is non-zero.

11.6.1.1522 FILE-SIZE FILE

(*fileid -- ud ior*)

ud is the size, in characters, of the file identified by *fileid*. *ior* is the implementation-defined I/O result code. This operation does not affect the value returned by **FILE-POSITION**. *ud* is undefined if *ior* is non-zero.

11.6.1.1717 INCLUDE-FILE FILE

(*i*x fileid -- j*x*)

Remove *fileid* from the stack. Save the current input source specification, including the current value of **SOURCE-ID**. Store *fileid* in **SOURCE-ID**. Make the file specified by *fileid* the input source. Store zero in **BLK**. Other stack effects are due to the words **INCLUDED**.

Repeat until end of file: read a line from the file, fill the input buffer from the contents of that line, set **>IN** to zero, and interpret.

Text interpretation begins at the file position where the next file read would occur.

When the end of the file is reached, close the file and restore the input source specification to its saved value.

An ambiguous condition exists if *fileid* is invalid, if there is an I/O exception reading *fileid*, or if an I/O exception occurs while closing *fileid*. When an ambiguous condition exists, the status (open or closed) of any files that were being interpreted is implementation-defined.

See: **11.3.4 Input source.**

11.6.1.1718 INCLUDED FILE

(*i*x c-addr u -- j*x*)

Remove *c-addr u* from the stack. Save the current input source specification, including the current value of **SOURCE-ID**. Open the file specified by *c-addr u*, store the resulting *fileid* in **SOURCE-ID**, and make it the input source. Store zero in **BLK**. Other stack effects are due to the words included.

Repeat until end of file: read a line from the file, fill the input buffer from the contents of that line, set **>IN** to zero, and interpret.

Text interpretation begins at the file position where the next file read would occur.

When the end of the file is reached, close the file and restore the input source specification to its saved value.

An ambiguous condition exists if the named file can not be opened, if an I/O exception occurs reading the file, or if an I/O exception occurs while closing the file. When an ambiguous condition exists, the status (open or closed) of any files that were being interpreted is implementation-defined.

See: **11.6.1.1717 INCLUDE-FILE.**

11.6.1.1970 OPEN-FILE FILE

(*c-addr u fam -- fileid ior*)

Open the file named in the character string specified by *c-addr u*, with file access method indicated by *fam*. The meaning of values of *fam* is implementation defined.

If the file is successfully opened, *ior* is zero, *fileid* is its identifier, and the file has been positioned to the start of the file.

Otherwise, *ior* is the implementation-defined I/O result code and *fileid* is undefined.

11.6.1.2054 R/O "r-O" FILE

(-- *fam*)

fam is the implementation-defined value for selecting the "read only" file access method.

See: **11.6.1.1010 CREATE-FILE, 11.6.1.1970 OPEN-FILE.**

11.6.1.2056 R/W "r-w" FILE

(--*fam*)

fam is the implementation-defined value for selecting the "read/write" file access method.

See: **11.6.1.1010 CREATE-FILE**, **11.6.1.1970 OPEN-FILE**.

11.6.1.2080 READ-FILE FILE

(*c-addr* u_1 *fileid* -- u_2 *ior*)

Read u_1 consecutive characters to *c-addr* from the current position of the file identified by *fileid*.

If u_1 characters are read without an exception, *ior* is zero and u_2 is equal to u_1 .

If the end of the file is reached before u_1 characters are read, *ior* is zero and u_2 is the number of characters actually read.

If the operation is initiated when the value returned by **FILE-POSITION** is equal to the value returned by **FILE-SIZE** for the file identified by *fileid*, *ior* is zero and u_2 is zero.

If an exception occurs, *ior* is the implementation-defined I/O result code, and u_2 is the number of characters transferred to *c-addr* without an exception.

An ambiguous condition exists if the operation is initiated when the value returned by **FILE-POSITION** is greater than the value returned by **FILE-SIZE** for the file identified by *fileid*, or if the requested operation attempts to read portions of the file not written.

At the conclusion of the operation, **FILE-POSITION** returns the next file position after the last character read.

11.6.1.2090 READ-LINE

FILE

(*c-addr* *u₁* *fileid* -- *u₂* *flag* *ior*)

Read the next line from the file specified by *fileid* into memory at the address *c-addr*. At most *u₁* characters are read. Up to two implementation-defined line terminating characters may be read into memory at the end of the line, but are not included in the count *u₂*. The line buffer provided by *c-addr* should be at least *u₁+2* characters long.

If the operation succeeded, *flag* is true and *ior* is zero. If a line terminator was received before *u₁* characters were read, then *u₂* is the number of characters, not including the line terminator, actually read ($0 \leq u_2 \leq u_1$). When $u_1 = u_2$ the line terminator has yet to be reached.

If the operation is initiated when the value returned by **FILE-POSITION** is equal to the value returned by **FILE-SIZE** for the file identified by *fileid*, *flag* is false, *ior* is zero, and *u₂* is zero. If *ior* is non-zero, an exception occurred during the operation and *ior* is the implementation-defined I/O result code.

An ambiguous condition exists if the operation is initiated when the value returned by **FILE-POSITION** is greater than the value returned by **FILE-SIZE** for the file identified by *fileid*, or if the requested operation attempts to read portions of the file not written.

At the conclusion of the operation, **FILE-POSITION** returns the next file position after the last character read.

11.6.1.2142 REPOSITION-FILE

FILE

(*ud* *fileid* -- *ior*)

Reposition the file identified by *fileid* to *ud*. *ior* is the implementation-defined I/O result code. An ambiguous condition exists if the file is positioned outside the file boundaries.

At the conclusion of the operation, **FILE-POSITION** returns the value *ud*.

11.6.1.2147 RESIZE-FILE

FILE

(*ud* *fileid* -- *ior*)

Set the size of the file identified by *fileid* to *ud*. *ior* is the implementation-defined I/O result code.

If the resultant file is larger than the file before the operation, the portion of the file added as a result of the operation might not have been written.

At the conclusion of the operation, **FILE-SIZE** returns the value *ud* and **FILE-POSITION** returns an unspecified value.

See: **11.6.1.2080 READ-FILE**, **11.6.1.2090 READ-LINE**.

11.6.1.2165 s" "s-quote" FILE

Extend the semantics of **6.1.2165 s"** to be:

Interpretation: ("*ccc*<quote>" -- *c-addr u*)

Parse *ccc* delimited by " (double quote). Store the resulting string *c-addr u* at a temporary location. The maximum length of the temporary buffer is implementation-dependent but shall be no less than 80 characters. Subsequent uses of **S"** may overwrite the temporary buffer. At least one such buffer shall be provided.

Compilation: ("*ccc*<quote>" --)

Parse *ccc* delimited by " (double quote). Append the run-time semantics given below to the current definition.

Run-time: (-- *c-addr u*)

Return *c-addr* and *u* that describe a string consisting of the characters *ccc*.

See: **3.4.1 Parsing**, **6.2.0855 C"**, **6.1.2165 S"**, **11.3.5 Other transient regions**.

11.6.1.2218 SOURCE-ID "source-i-d" FILE

(-- 0 | -1 | *fileid*)

Extend **6.2.2218 SOURCE-ID** to include text-file input as follows:

SOURCE-ID	Input source
<i>fileid</i>	Text file " <i>fileid</i> "
-1	String (via EVALUATE)
0	User input device

An ambiguous condition exists if **SOURCE-ID** is used when **BLK** contains a non-zero value.

11.6.1.2425 W/O "W-O" FILE

(-- *fam*)

fam is the implementation-defined value for selecting the "write only" file access method.

See: **11.6.1.1010 CREATE-FILE**, **11.6.1.1970 OPEN-FILE**.

11.6.1.2480 WRITE-FILE FILE

(*c-addr u fileid* -- *ior*)

Write *u* characters from *c-addr* to the file identified by *fileid* starting at its current position. *ior* is the implementation-defined I/O result code.

At the conclusion of the operation, **FILE-POSITION** returns the next file position after the last character written to the file, and **FILE-SIZE** returns a value greater than or equal to the value returned by **FILE-POSITION**.

See: **11.6.1.2080 READ-FILE**, **11.6.1.2090 READ-LINE**.

11.6.1.2485 WRITE-LINE FILE

(*c-addr* *u* *fileid* -- *ior*)

Write *u* characters from *c-addr* followed by the implementation-dependent line terminator to the file identified by *fileid* starting at its current position. *ior* is the implementation-defined I/O result code.

At the conclusion of the operation, **FILE-POSITION** returns the next file position after the last character written to the file, and **FILE-SIZE** returns a value greater than or equal to the value returned by **FILE-POSITION**.

See: **11.6.1.2080 READ-FILE**, **11.6.1.2090 READ-LINE**.

11.6.2 File-Access extension words**11.6.2.1524 FILE-STATUS** FILE EXT

(*c-addr* *u* -- *x* *ior*)

Return the status of the file identified by the character string *c-addr* *u*. If the file exists, *ior* is zero; otherwise *ior* is the implementation-defined I/O result code. *x* contains implementation-defined information about the file.

11.6.2.1560 FLUSH-FILE FILE EXT

(*fileid* -- *ior*)

Attempt to force any buffered information written to the file referred to by *fileid* to be written to mass storage, and the size information for the file to be recorded in the storage directory if changed. If the operation is successful, *ior* is zero. Otherwise, it is an implementation-defined I/O result code.

11.6.2.2125 REFILL FILE EXT

(-- *flag*)

Extend the execution semantics of **6.2.2125 REFILL** with the following:

When the input source is a text file, attempt to read the next line from the text-input file. If successful, make the result the current input buffer, set **>IN** to zero, and return *true*. Otherwise return *false*.

See: **6.2.2125 REFILL**, **7.6.2.2125 REFILL**.

11.6.2.2130 RENAME-FILE FILE EXT

(*c-addr*₁ *u*₁ *c-addr*₂ *u*₂ -- *ior*)

Rename the file named by the character string *c-addr*₁ *u*₁ to the name in the character string *c-addr*₂ *u*₂. *ior* is the implementation-defined I/O result code.

12. The optional Floating-Point word set

12.1 Introduction

12.2 Additional terms and notation

12.2.1 Definition of terms

float-aligned address: the address of a memory location at which a floating-point number can be accessed.

double-float-aligned address: the address of a memory location at which a 64-bit IEEE double-precision number can be accessed.

single-float-aligned address: the address of a memory location at which a 32-bit IEEE single-precision number can be accessed.

IEEE floating-point number: a floating-point number as defined in **ANSI/IEEE Standard 754-1985**.

12.2.2 Notation

12.2.2.1 Numeric notation

The following notation is used to define the syntax of the external representation of floating-point numbers:

- Each component of a floating-point number is defined with a rule consisting of the name of the component (italicized in angle-brackets, e.g., *<sign>*), the characters := and a concatenation of tokens and metacharacters;
- Tokens may be literal characters (in bold face, e.g., **E**) or rule names in angle brackets (e.g., *<digit>*);
- The metacharacter * is used to specify zero or more occurrences of the preceding token (e.g., *<digit>**);
- Tokens enclosed with [and] are optional (e.g., [*<sign>*]);
- Vertical bars separate choices from a list of tokens enclosed with braces (e.g., { + | - }).

12.2.2.2 Stack notation

Floating-point stack notation when the floating-point stack is separate from the data stack is:

(F: *before* -- *after*)

12.3 Additional usage requirements

12.3.1 Data types

Append table 12.1 to table 3.1.

Table 12.1 – Data Types

Symbol	Data type	Size on stack
<i>r</i>	floating-point number	implementation-defined
<i>f-addr</i>	float-aligned address	1 cell
<i>sf-addr</i>	single-float-aligned address	1 cell
<i>df-addr</i>	double-float-aligned address	1 cell

12.3.1.1 Addresses

The set of float-aligned addresses is an implementation-defined subset of the set of aligned addresses. Adding the size of a floating-point number to a float-aligned address shall produce a float-aligned address.

The set of double-float-aligned addresses is an implementation-defined subset of the set of aligned addresses. Adding the size of a 64-bit IEEE double-precision floating-point number to a double-float-aligned address shall produce a double-float-aligned address.

The set of single-float-aligned addresses is an implementation-defined subset of the set of aligned addresses. Adding the size of a 32-bit IEEE single-precision floating-point number to a single-float-aligned address shall produce a single-float-aligned address.

12.3.1.2 Floating-point numbers

The internal representation of a floating-point number, including the format and precision of the significand and the format and range of the exponent, is implementation defined.

Any rounding or truncation of floating-point numbers is implementation defined.

12.3.2 Floating-point operations

“Round to nearest” means round the result of a floating-point operation to the representable value nearest the result. If the two nearest representable values are equally near the result, the one having zero as its least significant bit shall be delivered.

“Round toward negative infinity” means round the result of a floating-point operation to the representable value nearest to and no greater than the result.

12.3.3 Floating-point stack

A last in, first out list that shall be used by all floating-point operators.

The width of the floating-point stack is implementation-defined. By default the floating-point stack shall be separate from the data and return stacks. A program may determine whether floating-point numbers are kept on the data stack by passing the string “**FLOATING-STACK**” to **ENVIRONMENT?**.

The size of a floating-point stack shall be at least 6 items.

A program that depends on the floating-point stack being larger than six items has an environmental dependency.

12.3.4 Environmental queries

Append table 12.2 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 12.2 – Environmental query strings

String	Value	Data type	Constant?	Meaning
FLOATING	<i>flag</i>		no	floating-point word set present
FLOATING-EXT	<i>flag</i>		no	floating-point extensions word set present
FLOATING-STACK	<i>n</i>		yes	If <i>n</i> = zero, floating-point numbers are kept on the data stack; otherwise <i>n</i> is the maximum depth of the separate floating-point stack.
MAX-FLOAT	<i>r</i>		yes	largest usable floating-point number

12.3.5 Address alignment

Since the address returned by a **CREATEd** word is not necessarily aligned for any particular class of floating-point data, a program shall align the address (to be float aligned, single-float aligned, or double-float aligned) before accessing floating-point data at the address.

See: **3.3.3.1 Address Alignment, 12.3.1.1 Addresses.**

12.3.6 Variables

A program may address memory in data space regions made available by **FVARIABLE**. These regions may be non-contiguous with regions subsequently allocated with **,** (comma) or **ALLOT**.

See: **3.3.3.3 Variables.**

12.3.7 Text interpreter input number conversion

If the Floating-Point word set is present in the dictionary and the current base is **DECIMAL**, the input number-conversion algorithm shall be extended to recognize floating-point numbers in this form:

Convertible string := <significand><exponent>

<significand> := [<sign>]<digits>[.<digits0>]

<exponent> := **E**[<sign>]<digits0>

<sign> := { + | - }

<digits> := <digit><digits0>

<digits0> := <digit>*

<digit> := { 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 }

These are examples of valid representations of floating-point numbers in program source:

1E 1.E 1.E0 +1.23E-1 -1.23E+1

See: **3.4.1.3 Text interpreter input number conversion, 12.6.1.0558 >FLOAT.**

12.4 Additional documentation requirements

12.4.1 System documentation

12.4.1.1 Implementation-defined options

- format and range of floating-point numbers (**12.3.1 Data types, 12.6.1.2143 REPRESENT**);
- results of **12.6.1.2143 REPRESENT** when *float* is out of range;
- rounding or truncation of floating-point numbers (**12.3.1.2 Floating-point numbers**);
- size of floating-point stack (**12.3.3 Floating-point stack**);
- width of floating-point stack (**12.3.3 Floating-point stack**).

12.4.1.2 Ambiguous conditions

- **DF@** or **DF!** is used with an address that is not double-float aligned;
- **F@** or **F!** is used with an address that is not float aligned;
- floating point result out of range (**12.6.1.1430 F/**);
- **SF@** or **SF!** is used with an address that is not single-float aligned;
- * **BASE** is not decimal (**12.6.1.2143 REPRESENT, 12.6.2.1427 F., 12.6.2.1513 FE., 12.6.2.1613 FS.**);
- * both arguments equal zero (**12.6.2.1489 FATAN2**);
- * $\cos(r_j)$ is zero as an argument for **12.6.2.1625 FTAN**;
- * *d* can't be precisely represented as *float* in **12.6.1.1130 D>F**;
- * dividing by zero (**12.6.1.14340**);

- * exponent too big for conversion (**12.6.2.1203 DF!**, **12.6.2.1204 DF@**, **12.6.2.2202 SF!**, **12.6.2.2203 SF@**);
- * *float* less than one (**12.6.2.1477 FACOSH**);
- * *float* less than or equal to minus-one (**12.6.2.1554 FLNP1**);
- * *float* less than or equal to zero (**12.6.2.1553 FLN**, **12.6.2.1557 FLOG**);
- * *float* less than zero (**12.6.2.1487 FASINE**, **12.6.2.1618 FSQRT**);
- * *float* magnitude greater than one (**12.6.2.1476 FACOS**, **12.6.2.1486 FASIN**, **12.6.2.1491 FATANH**);
- * integer part of *float* can't be represented by *d* in **12.6.1.1470 F>D**;
- * string larger than pictured-numeric output area (**12.6.2.1427 F.**, **12.6.2.1513 FE.**, **12.6.2.1613 FS.**).

12.4.1.3 Other system documentation

- no additional requirements.

12.4.2 Program documentation

12.4.2.1 Environmental dependencies

- requiring the floating-point stack to be larger than six items (**12.3.3 Floating-point stack**).

12.4.2.2 Other program documentation

- no additional requirements.

12.5 Compliance and labeling

12.5.1 ANS Forth systems

The phrase “Providing the Floating-Point word set” shall be appended to the label of any Standard System that provides all of the Floating-Point word set.

The phrase “Providing *name(s)* from the Floating-Point Extensions word set” shall be appended to the label of any Standard System that provides portions of the Floating-Point Extensions word set.

The phrase “Providing the Floating-Point Extensions word set” shall be appended to the label of any Standard System that provides all of the Floating-Point and Floating-Point Extensions word sets.

12.5.2 ANS Forth programs

The phrase “Requiring the Floating-Point word set” shall be appended to the label of Standard Programs that require the system to provide the Floating-Point word set.

The phrase “Requiring *name(s)* from the Floating-Point Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Floating-Point Extensions word set.

The phrase “Requiring the Floating-Point Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Floating-Point and Floating-Point Extensions word sets.

12.6 Glossary

12.6.1 Floating-Point words

12.6.1.0558 >FLOAT

“to-float”

FLOATING

$(c\text{-}addr\ u\ \text{--}\ true\ |\ false)$ (F: $--\ r\ |$) or $(c\text{-}addr\ u\ \text{--}\ r\ true\ |\ false)$

An attempt is made to convert the string specified by *c-addr* and *u* to internal floating-point representation. If the string represents a valid floating-point number in the syntax below, its value *r* and *true* are returned. If the string does not represent a valid floating-point number only *false* is returned.

A string of blanks should be treated as a special case representing zero.

The syntax of a convertible string := $\langle\text{significand}\rangle[\langle\text{exponent}\rangle]$

$\langle\text{significand}\rangle := [\langle\text{sign}\rangle]\{\langle\text{digits}\rangle[.\langle\text{digits0}\rangle]|\langle\text{digits}\rangle\}$

$\langle\text{exponent}\rangle := \langle\text{marker}\rangle\langle\text{digits0}\rangle$

$\langle\text{marker}\rangle := \{\langle e\text{-form}\rangle|\langle\text{sign-form}\rangle\}$

$\langle e\text{-form}\rangle := \langle e\text{-char}\rangle[\langle\text{sign-form}\rangle]$

$\langle\text{sign-form}\rangle := \{+|- \}$

$\langle e\text{-char}\rangle := \{\mathbf{D}|\mathbf{d}|\mathbf{E}|\mathbf{e}\}$

12.6.1.1130 D>F

“d-to-f”

FLOATING

$(d\ \text{--})$ (F: $--\ r$) or $(d\ \text{--}\ r)$

r is the floating-point equivalent of *d*. An ambiguous condition exists if *d* cannot be precisely represented as a floating-point value.

12.6.1.1400 F!

“f-store”

FLOATING

$(f\text{-}addr\ \text{--})$ (F: $r\ \text{--}$) or $(r\ f\text{-}addr\ \text{--})$

Store *r* at *f-addr*.

12.6.1.1410 F*

“f-star”

FLOATING

(F: $r_1\ r_2\ \text{--}\ r_3$) or $(r_1\ r_2\ \text{--}\ r_3)$

Multiply *r*₁ by *r*₂ giving *r*₃.

12.6.1.1420 F+

“f-plus”

FLOATING

(F: $r_1\ r_2\ \text{--}\ r_3$) or $(r_1\ r_2\ \text{--}\ r_3)$

Add *r*₁ to *r*₂ giving the sum *r*₃.

12.6.1.1425 F-

“f-minus”

FLOATING

(F: $r_1\ r_2\ \text{--}\ r_3$) or $(r_1\ r_2\ \text{--}\ r_3)$

Subtract *r*₂ from *r*₁, giving *r*₃.

12.6.1.1430	F/	“f-slash”	FLOATING
		$(F: r_1 r_2 -- r_3)$ or $(r_1 r_2 -- r_3)$	
		Divide r_1 by r_2 , giving the quotient r_3 . An ambiguous condition exists if r_2 is zero, or the quotient lies outside of the range of a floating-point number.	
12.6.1.1440	F0<	“f-zero-less-than”	FLOATING
		$(--flag)$ (F: $r --$) or $(r --flag)$	
		$flag$ is true if and only if r is less than zero.	
12.6.1.1450	F0=	“f-zero-equals”	FLOATING
		$(--flag)$ (F: $r --$) or $(r --flag)$	
		$flag$ is true if and only if r is equal to zero.	
12.6.1.1460	F<	“f-less-than”	FLOATING
		$(--flag)$ (F: $r_1 r_2 --$) or $(r_1 r_2 --flag)$	
		$flag$ is true if and only if r_1 is less than r_2 .	
12.6.1.1470	F>D	“f-to-d”	FLOATING
		$(--d)$ (F: $r --$) or $(r --d)$	
		d is the double-cell signed-integer equivalent of the integer portion of r . The fractional portion of r is discarded. An ambiguous condition exists if the integer portion of r cannot be precisely represented as a double-cell signed integer.	
12.6.1.1472	F@	“f-fetch”	FLOATING
		$(f-addr --)$ (F: $--r$) or $(f-addr --r)$	
		r is the value stored at $f-addr$.	
12.6.1.1479	FALIGN	“f-align”	FLOATING
		$(--)$	
		If the data-space pointer is not float aligned, reserve enough data space to make it so.	
12.6.1.1483	FALIGNED	“f-aligned”	FLOATING
		$(addr --f-addr)$	
		$f-addr$ is the first float-aligned address greater than or equal to $addr$.	

12.6.1.1492	FCONSTANT	“f-constant”	FLOATING
		$(\langle \text{spaces} \rangle \text{name} \text{ -- }) (\text{F: } r \text{ -- }) \text{ or } (r \langle \text{spaces} \rangle \text{name} \text{ -- })$	
		<p>Skip leading space delimiters. Parse <i>name</i> delimited by a space. Create a definition for <i>name</i> with the execution semantics defined below.</p> <p><i>name</i> is referred to as an “f-constant.”</p> <p><i>name</i> Execution: $(\text{-- }) (\text{F: } \text{-- } r) \text{ or } (\text{-- } r)$</p> <p>Place <i>r</i> on the floating-point stack.</p> <p>See: 3.4.1 Parsing.</p>	
12.6.1.1497	FDEPTH	“f-depth”	FLOATING
		$(\text{-- } +n)$	
		<p>+<i>n</i> is the number of values contained on the default separate floating-point stack. If floating-point numbers are kept on the data stack, +<i>n</i> is the current number of possible floating-point values contained on the data stack.</p>	
12.6.1.1500	FDROP	“f-drop”	FLOATING
		$(\text{F: } r \text{ -- }) \text{ or } (r \text{ -- })$	
		<p>Remove <i>r</i> from the floating-point stack.</p>	
12.6.1.1510	FDUP	“f-dupe”	FLOATING
		$(\text{F: } r \text{ -- } r r) \text{ or } (r \text{ -- } r r)$	
		<p>Duplicate <i>r</i>.</p>	
12.6.1.1552	FLITERAL	“f-literal”	FLOATING
		<p>Interpretation: Interpretation semantics for this word are undefined.</p> <p>Compilation: $(\text{F: } r \text{ -- }) \text{ or } (r \text{ -- })$</p> <p>Append the run-time semantics given below to the current definition.</p> <p>Run-time: $(\text{F: } \text{-- } r) \text{ or } (\text{-- } r)$</p> <p>Place <i>r</i> on the floating-point stack.</p>	
12.6.1.1555	FLOAT+	“float-plus”	FLOATING
		$(f\text{-addr}_1 \text{ -- } f\text{-addr}_2)$	
		<p>Add the size in address units of a floating-point number to <i>f-addr</i>₁, giving <i>f-addr</i>₂.</p>	

12.6.1.1556	FLOATS		FLOATING
	$(n_1 \text{ -- } n_2)$		
	n_2 is the size in address units of n_1 floating-point numbers.		
12.6.1.1558	FLOOR		FLOATING
	$(F: r_1 \text{ -- } r_2)$ or $(r_1 \text{ -- } r_2)$		
	Round r_1 to an integral value using the “round toward negative infinity” rule, giving r_2 .		
12.6.1.1562	FMAX	“f-max”	FLOATING
	$(F: r_1 r_2 \text{ -- } r_3)$ or $(r_1 r_2 \text{ -- } r_3)$		
	r_3 is the greater of r_1 and r_2 .		
12.6.1.1565	FMIN	“f-min”	FLOATING
	$(F: r_1 r_2 \text{ -- } r_3)$ or $(r_1 r_2 \text{ -- } r_3)$		
	r_3 is the lesser of r_1 and r_2 .		
12.6.1.1567	FNEGATE	“f-negate”	FLOATING
	$(F: r_1 \text{ -- } r_2)$ or $(r_1 \text{ -- } r_2)$		
	r_2 is the negation of r_1 .		
12.6.1.1600	FOVER	“f-over”	FLOATING
	$(F: r_1 r_2 \text{ -- } r_1 r_2 r_1)$ or $(r_1 r_2 \text{ -- } r_1 r_2 r_1)$		
	Place a copy of r_1 on top of the floating-point stack.		
12.6.1.1610	FROT	“f-rote”	FLOATING
	$(F: r_1 r_2 r_3 \text{ -- } r_2 r_3 r_1)$ or $(r_1 r_2 r_3 \text{ -- } r_2 r_3 r_1)$		
	Rotate the top three floating-point stack entries.		
12.6.1.1612	FROUND	“f-round”	FLOATING
	$(F: r_1 \text{ -- } r_2)$ or $(r_1 \text{ -- } r_2)$		
	Round r_1 to an integral value using the “round to nearest” rule, giving r_2 .		
	See: 12.3.2 Floating-point operations.		
12.6.1.1620	FSWAP	“f-swap”	FLOATING
	$(F: r_1 r_2 \text{ -- } r_2 r_1)$ or $(r_1 r_2 \text{ -- } r_2 r_1)$		
	Exchange the top two floating-point stack items.		

12.6.1.1630 FVARIABLE “f-variable” FLOATING

(“<spaces>name” --)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name* with the execution semantics defined below. Reserve 1 **FLOATS** address units of data space at a float-aligned address.

name is referred to as an “f-variable.”

name Execution: (--*f-addr*)

f-addr is the address of the data space reserved by **FVARIABLE** when it created *name*. A program is responsible for initializing the contents of the reserved space.

See: **3.4.1 Parsing**.

12.6.1.2143 REPRESENT FLOATING

(*c-addr u -- n flag₁ flag₂*) (F: *r --*) or (*r c-addr u -- n flag₁ flag₂*)

At *c-addr*, place the character-string external representation of the significand of the floating-point number *r*. Return the decimal-base exponent as *n*, the sign as *flag₁* and “valid result” as *flag₂*. The character string shall consist of the *u* most significant digits of the significand represented as a decimal fraction with the implied decimal point to the left of the first digit, and the first digit zero only if all digits are zero. The significand is rounded to *u* digits following the round to nearest rule; *n* is adjusted, if necessary, to correspond to the rounded magnitude of the significand. If *flag₂* is true then *r* was in the implementation-defined range of floating-point numbers. If *flag₁* is true then *r* is negative.

An ambiguous condition exists if the value of **BASE** is not decimal ten.

When *flag₂* is false, *n* and *flag₁* are implementation defined, as are the contents of *c-addr*. Under these circumstances, the string at *c-addr* shall consist of graphic characters.

See: **3.2.1.2 Digit conversion**, **6.1.0750 BASE**, **12.3.2 Floating-point operations**.

12.6.2 Floating-Point extension words**12.6.2.1203 DF!** “d-f-store” FLOATING EXT

(*df-addr --*) (F: *r --*) or (*r df-addr --*)

Store the floating-point number *r* as a 64-bit IEEE double-precision number at *df-addr*. If the significand of the internal representation of *r* has more precision than the IEEE double-precision format, it will be rounded using the “round to nearest” rule. An ambiguous condition exists if the exponent of *r* is too large to be accommodated in IEEE double-precision format.

See: **12.3.1.1 Addresses**, **12.3.2 Floating-point operations**.

12.6.2.1204 DF@ “d-f-fetch” FLOATING EXT

(*df-addr* --) (F: -- *r*) or (*df-addr* -- *r*)

Fetch the 64-bit IEEE double-precision number stored at *df-addr* to the floating-point stack as *r* in the internal representation. If the IEEE double-precision significand has more precision than the internal representation it will be rounded to the internal representation using the “round to nearest” rule. An ambiguous condition exists if the exponent of the IEEE double-precision representation is too large to be accommodated by the internal representation.

See: **2.3.1.1 Addresses, 12.3.2 Floating-point operations.**

12.6.2.1205 DFALIGN “d-f-align” FLOATING EXT

(--)

If the data-space pointer is not double-float aligned, reserve enough data space to make it so.

See: **12.3.1.1 Addresses.**

12.6.2.1207 DFALIGNED “d-f-aligned” FLOATING EXT

(*addr* -- *df-addr*)

df-addr is the first double-float-aligned address greater than or equal to *addr*.

See: **12.3.1.1 Addresses.**

12.6.2.1208 DFLOAT+ “d-float-plus” FLOATING EXT

(*df-addr*₁ -- *df-addr*₂)

Add the size in address units of a 64-bit IEEE double-precision number to *df-addr*₁, giving *df-addr*₂.

See: **12.3.1.1 Addresses.**

12.6.2.1209 DFLOATS “d-floats” FLOATING EXT

(*n*₁ -- *n*₂)

*n*₂ is the size in address units of *n*₁ 64-bit IEEE double-precision numbers.

12.6.2.1415 F** “f-star-star” FLOATING EXT

(F: *r*₁ *r*₂ -- *r*₃) or (*r*₁ *r*₂ -- *r*₃)

Raise *r*₁ to the power *r*₂, giving the product *r*₃.

12.6.2.1427 F.	"f-dot"	FLOATING EXT
$(--) (F: r --) \text{ or } (r --)$		
Display, with a trailing space, the top number on the floating-point stack using fixed-point notation:		
$[-] <digits> . <digits0>$		
An ambiguous condition exists if the value of BASE is not (decimal) ten or if the character string representation exceeds the size of the pictured numeric output string buffer.		
See: 12.6.1.0558 >FLOAT.		
12.6.2.1474 FABS	"f-abs"	FLOATING EXT
$(F: r_1 -- r_2) \text{ or } (r_1 -- r_2)$		
r_2 is the absolute value of r_1 .		
12.6.2.1476 FACOS	"f-a-cos"	FLOATING EXT
$(F: r_1 -- r_2) \text{ or } (r_1 -- r_2)$		
r_2 is the principal radian angle whose cosine is r_1 . An ambiguous condition exists if $ r_1 $ is greater than 1.		
12.6.2.1477 FACOSH	"f-a-cosh"	FLOATING EXT
$(F: r_1 -- r_2) \text{ or } (r_1 -- r_2)$		
r_2 is the floating-point value whose hyperbolic cosine is r_1 . An ambiguous condition exists if r_1 is less than one.		
12.6.2.1484 FALOG	"f-a-log"	FLOATING EXT
$(F: r_1 -- r_2) \text{ or } (r_1 -- r_2)$		
Raise ten to the power r_1 , giving r_2 .		
12.6.2.1486 FASIN	"f-a-sine"	FLOATING EXT
$(F: r_1 -- r_2) \text{ or } (r_1 -- r_2)$		
r_2 is the principal radian angle whose sine is r_1 . An ambiguous condition exists if $ r_1 $ is greater than one.		
12.6.2.1487 FASINH	"f-a-cinch"	FLOATING EXT
$(F: r_1 -- r_2) \text{ or } (r_1 -- r_2)$		
r_2 is the floating-point value whose hyperbolic sine is r_1 . An ambiguous condition exists if r_1 is less than zero.		

12.6.2.1488	FATAN	“f-a-tan”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the principal radian angle whose tangent is r_1 .		
12.6.2.1489	FATAN2	“f-a-tan-two”	FLOATING EXT
	(F: $r_1 r_2 \text{ -- } r_3$) or ($r_1 r_2 \text{ -- } r_3$)		
	r_3 is the radian angle whose tangent is r_1/r_2 . An ambiguous condition exists if r_1 and r_2 are zero.		
12.6.2.1491	FATANH	“f-a-tan-h”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the floating-point value whose hyperbolic tangent is r_1 . An ambiguous condition exists if r_1 is outside the range of -1E0 to 1E0.		
12.6.2.1493	FCOS	“f-cos”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the cosine of the radian angle r_1 .		
12.6.2.1494	FCOSH	“f-cosh”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the hyperbolic cosine of r_1 .		
12.6.2.1513	FE.	“f-e-dot”	FLOATING EXT
	(--) (F: $r \text{ --}$) or ($r \text{ --}$)		
	Display, with a trailing space, the top number on the floating-point stack using engineering notation, where the significand is greater than or equal to 1.0 and less than 1000.0 and the exponent is a multiple of three.		
	An ambiguous condition exists if the value of BASE is not (decimal) ten or if the character string representation exceeds the size of the pictured numeric output string buffer.		
	See: 6.1.0750 BASE , 12.3.2 Floating-point operations , 12.6.1.2143 REPRESENT .		
12.6.2.1515	FEXP	“f-e-x-p”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	Raise e to the power r_1 , giving r_2 .		

12.6.2.1516	FEXPM1	“f-e-x-p-m-one”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	Raise e to the power r_1 and subtract one, giving r_2 .		
12.6.2.1553	FLN	“f-l-n”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the natural logarithm of r_1 . An ambiguous condition exists if r_1 is less than or equal to zero.		
12.6.2.1554	FLNP1	“f-l-n-p-one”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the natural logarithm of the quantity r_1 plus one. An ambiguous condition exists if r_1 is less than or equal to negative one.		
12.6.2.1557	FLOG	“f-log”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the base-ten logarithm of r_1 . An ambiguous condition exists if r_1 is less than or equal to zero.		
12.6.2.1613	FS.	“f-s-dot”	FLOATING EXT
	(--) (F: $r \text{ --}$) or ($r \text{ --}$)		
	Display, with a trailing space, the top number on the floating-point stack in scientific notation: <significand><exponent>		
	where:		
	<significand> := [-]<digit>.<digits0>		
	<exponent> := E[-]<digits>		
	An ambiguous condition exists if the value of BASE is not (decimal) ten or if the character string representation exceeds the size of the pictured numeric output string buffer.		
	See: 6.1.0750 BASE , 12.3.2 Floating-point operations , 12.6.1.2143 REPRESENT .		
12.6.2.1614	FSIN	“f-sine”	FLOATING EXT
	(F: $r_1 \text{ -- } r_2$) or ($r_1 \text{ -- } r_2$)		
	r_2 is the sine of the radian angle r_1 .		

12.6.2.1616	FSINCOS	“f-sine-cos”	FLOATING EXT
	(F: r_1 -- r_2 r_3) or (r_1 -- r_2 r_3)		
	r_2 is the sine of the radian angle r_1 . r_3 is the cosine of the radian angle r_1 .		
12.6.2.1617	FSINH	“f-cinch”	FLOATING EXT
	(F: r_1 -- r_2) or (r_1 -- r_2)		
	r_2 is the hyperbolic sine of r_1 .		
12.6.2.1618	FSQRT	“f-square-root”	FLOATING EXT
	(F: r_1 -- r_2) or (r_1 -- r_2)		
	r_2 is the square root of r_1 . An ambiguous condition exists if r_1 is less than zero.		
12.6.2.1625	FTAN	“f-tan”	FLOATING EXT
	(F: r_1 -- r_2) or (r_1 -- r_2)		
	r_2 is the tangent of the radian angle r_1 . An ambiguous condition exists if $\cos(r_1)$ is zero.		
12.6.2.1626	FTANH	“f-tan-h”	FLOATING EXT
	(F: r_1 -- r_2) or (r_1 -- r_2)		
	r_2 is the hyperbolic tangent of r_1 .		
12.6.2.1640	F~	“f-proximate”	FLOATING EXT
	(-- <i>flag</i>) (F: r_1 r_2 r_3 --) or (r_1 r_2 r_3 -- <i>flag</i>)		
	If r_3 is positive, <i>flag</i> is true if the absolute value of (r_1 minus r_2) is less than r_3 .		
	If r_3 is zero, <i>flag</i> is true if the implementation-dependent encoding of r_1 and r_2 are exactly identical (positive and negative zero are unequal if they have distinct encodings).		
	If r_3 is negative, <i>flag</i> is true if the absolute value of (r_1 minus r_2) is less than the absolute value of r_3 times the sum of the absolute values of r_1 and r_2 .		
12.6.2.2035	PRECISION		FLOATING EXT
	(-- <i>u</i>)		
	Return the number of significant digits currently used by F. , FE. , or FS. as <i>u</i> .		
12.6.2.2200	SET-PRECISION		FLOATING EXT
	(<i>u</i> --)		
	Set the number of significant digits currently used by F. , FE. , or FS. to <i>u</i> .		

12.6.2.2202 **SF!** “s-f-store” FLOATING EXT

(*sf-addr* --) (F: *r* --) or (*r sf-addr* --)

Store the floating-point number *r* as a 32-bit IEEE single-precision number at *sf-addr*. If the significand of the internal representation of *r* has more precision than the IEEE single-precision format, it will be rounded using the “round to nearest” rule. An ambiguous condition exists if the exponent of *r* is too large to be accommodated by the IEEE single-precision format.

See: **12.3.1.1 Addresses**, **12.3.2 Floating-point operations**.

12.6.2.2203 **SF@** “s-f-fetch” FLOATING EXT

(*sf-addr* --) (F: -- *r*) or (*sf-addr* -- *r*)

Fetch the 32-bit IEEE single-precision number stored at *sf-addr* to the floating-point stack as *r* in the internal representation. If the IEEE single-precision significand has more precision than the internal representation, it will be rounded to the internal representation using the “round to nearest” rule. An ambiguous condition exists if the exponent of the IEEE single-precision representation is too large to be accommodated by the internal representation.

See: **12.3.1.1 Addresses**, **12.3.2 Floating-point operations**.

12.6.2.2204 **SFALIGN** “s-f-align” FLOATING EXT

(--)

If the data-space pointer is not single-float aligned, reserve enough data space to make it so.

See: **12.3.1.1 Addresses**.

12.6.2.2206 **SFALIGNED** “s-f-aligned” FLOATING EXT

(*addr* -- *sf-addr*)

sf-addr is the first single-float-aligned address greater than or equal to *addr*.

See: **12.3.1.1 Addresses**.

12.6.2.2207 **SFLOAT+** “s-float-plus” FLOATING EXT

(*sf-addr*₁ -- *sf-addr*₂)

Add the size in address units of a 32-bit IEEE single-precision number to *sf-addr*₁, giving *sf-addr*₂.

See: **12.3.1.1 Addresses**.

12.6.2.2208 SFLOATS

"s-floats"

FLOATING EXT

 $(n_1 \text{ -- } n_2)$ n_2 is the size in address units of n_1 32-bit IEEE single-precision numbers.See: **12.3.1.1 Addresses.**

13. The optional Locals word set

13.1 Introduction

See: **Annex A.13 The Locals Word Set.**

13.2 Additional terms and notation

None.

13.3 Additional usage requirements

13.3.1 Locals

A local is a data object whose execution semantics shall return its value, whose scope shall be limited to the definition in which it is declared, and whose use in a definition shall not preclude reentrancy or recursion.

13.3.2 Environmental queries

Append table 13.1 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 13.1 – Environmental query strings

String	Value data type	Constant?	Meaning
#LOCALS	<i>n</i>	yes	maximum number of local variables in a definition
LOCALS	<i>flag</i>	no	locals word set present
LOCALS-EXT	<i>flag</i>	no	locals extensions word set present

13.3.3 Processing locals

To support the locals word set, a system shall provide a mechanism to receive the messages defined by (**LOCAL**) and respond as described here.

During the compilation of a definition after **:** (colon), **:NONAME**, or **DOES>**, a program may begin sending local identifier messages to the system. The process shall begin when the first message is sent. The process shall end when the “last local” message is sent. The system shall keep track of the names, order, and number of identifiers contained in the complete sequence.

The system, upon receipt of a sequence of local-identifier messages, shall take the following actions at compile time:

- 1) Create temporary dictionary entries for each of the identifiers passed to (**LOCAL**), such that each identifier will behave as a *local*. These temporary dictionary entries shall vanish at the end of the definition, denoted by **;** (semicolon), **;CODE**, or **DOES>**. The system need not maintain these identifiers in the same way it does other dictionary entries as long as they can be found by normal dictionary searching processes. Furthermore, if the Search-Order word set is present, local identifiers shall always be searched before any of the word lists in any definable search order, and none of the Search-Order words shall change the locals’ privileged position in the search order. Local identifiers may reside in mass storage.
- 2) For each identifier passed to (**LOCAL**), the system shall generate an appropriate code sequence that does the following at execution time:

- a) Allocate a storage resource adequate to contain the value of a local. The storage shall be allocated in a way that does not preclude re-entrancy or recursion in the definition using the local.
- b) Initialize the value using the top item on the data stack. If more than one local is declared, the top item on the stack shall be moved into the first local identified, the next item shall be moved into the second, and so on.

The storage resource may be the return stack or may be implemented in other ways, such as in registers. The storage resource shall not be the data stack. Use of locals shall not restrict use of the data stack before or after the point of declaration.

- 3) Arrange that any of the legitimate methods of terminating execution of a definition, specifically `;` (semicolon), `;CODE, DOES>` or `EXIT`, will release the storage resource allocated for the locals, if any, declared in that definition. `ABORT` shall release all local storage resources, and `CATCH / THROW` (if implemented) shall release such resources for all definitions whose execution is being terminated.
- 4) Separate sets of locals may be declared in defining words before `DOES>` for use by the defining word, and after `DOES>` for use by the word defined.

A system implementing the Locals word set shall support the declaration of at least eight locals in a definition.

Immediate words in a program may use `(LOCAL)` to implement syntaxes for local declarations with the following restrictions:

- A program shall not compile any executable code into the current definition between the time `(LOCAL)` is executed to identify the first local for that definition and the time of sending the single required “last local” message;
- The position in program source at which the sequence of `(LOCAL)` messages is sent, referred to here as the point at which locals are declared, shall not lie within the scope of any control structure;
- Locals shall not be declared until values previously placed on the return stack within the definition have been removed;
- After a definition’s locals have been declared, a program may place data on the return stack. However, if this is done, locals shall not be accessed until those values have been removed from the return stack;
- Words that return execution tokens, such as `'` (tick), `[']`, or `FIND`, shall not be used with local names;
- A program that declares more than eight locals in a single definition has an environmental dependency;
- Locals may be accessed or updated within control structures, including do-loops;
- Local names shall not be referenced by `POSTPONE` and `[COMPILE]`.

See: [3.4 The Forth text interpreter](#).

13.4 Additional documentation requirements

13.4.1 System documentation

13.4.1.1 Implementation-defined options

- maximum number of locals in a definition ([13.3.3 Processing locals](#), [13.6.2.1795 LOCALS](#) |).

13.4.1.2 Ambiguous conditions

- * executing a named *local* while in interpretation state ([13.6.1.0086 \(LOCAL\)](#));
- * *name* not defined by `VALUE` or `LOCAL` ([13.6.1.2295 TO](#)).

13.4.1.3 Other system documentation

- no additional requirements.

13.4.2 Program documentation

13.4.2.1 Environmental dependencies

- declaring more than eight locals in a single definition (**13.3.3 Processing locals**).

13.4.2.2 Other program documentation

- no additional requirements.

13.5 Compliance and labeling

13.5.1 ANS Forth systems

The phrase “Providing the Locals word set” shall be appended to the label of any Standard System that provides all of the Locals word set.

The phrase “Providing *name(s)* from the Locals Extensions word set” shall be appended to the label of any Standard System that provides portions of the Locals Extensions word set.

The phrase “Providing the Locals Extensions word set” shall be appended to the label of any Standard System that provides all of the Locals and Locals Extensions word sets.

13.5.2 ANS Forth programs

The phrase “Requiring the Locals word set” shall be appended to the label of Standard Programs that require the system to provide the Locals word set.

The phrase “Requiring *name(s)* from the Locals Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Locals Extensions word set.

The phrase “Requiring the Locals Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Locals and Locals Extensions word sets.

13.6 Glossary

13.6.1 Locals words

13.6.1.0086 (LOCAL)

“Paren-local-paren”

LOCAL

Interpretation: Interpretation semantics for this word are undefined.

Execution: (*c-addr* *u* --)

When executed during compilation, (**LOCAL**) passes a message to the system that has one of two meanings. If *u* is non-zero, the message identifies a new *local* whose definition name is given by the string of characters identified by *c-addr u*. If *u* is zero, the message is “last local” and *c-addr* has no significance.

The result of executing (**LOCAL**) during compilation of a definition is to create a set of named *local* identifiers, each of which is a definition name, that only have execution semantics within the scope of that definition’s source.

local Execution: (-- *x*)

Push the *local*’s value, *x*, onto the stack. The *local*’s value is initialized as described in **13.3.3 Processing locals** and may be changed by preceding the *local*’s name with **TO**. An ambiguous condition exists when *local* is executed while in interpretation state.

Note: This word does not have special compilation semantics in the usual sense because it provides access to a system capability for use by other user-defined words that do have them. However, the locals facility as a whole and the sequence of messages passed defines specific usage rules with semantic implications that are described in detail in section **13.3.3 Processing locals**.

Note: This word is not intended for direct use in a definition to declare that definition’s locals. It is instead used by system or user compiling words. These compiling words in turn define their own syntax, and may be used directly in definitions to declare locals. In this context, the syntax for (**LOCAL**) is defined in terms of a sequence of compile-time messages and is described in detail in section **13.3.3 Processing locals**.

Note: The Locals word set modifies the syntax and semantics of **6.2.2295 TO** as defined in the Core Extensions word set.

See: **3.4 The Forth text interpreter**.

13.6.1.2295 TO

LOCAL

Extend the semantics of **6.2.2295 TO** to be:

Interpretation: (x “<spaces>name” --)

Skip leading spaces and parse *name* delimited by a space. Store x in *name*. An ambiguous condition exists if *name* was not defined by **VALUE**.

Compilation: (“<spaces>name” --)

Skip leading spaces and parse *name* delimited by a space. Append the run-time semantics given below to the current definition. An ambiguous condition exists if *name* was not defined by either **VALUE** or (**LOCAL**).

Run-time: (x --)

Store x in *name*.

Note: An ambiguous condition exists if either **POSTPONE** or [**COMPILE**] is applied to **TO**.

See: **3.4.1 Parsing**, **6.2.2295 TO**, **6.2.2405 VALUE**, **13.6.1.0086 (LOCAL)**.

13.6.2 Locals extension words**13.6.2.1795 LOCALS |**

“locals-bar”

LOCAL EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (“*name*₁” “*name*₂” ... “*name*_{*n*}” “|” --)

Create up to eight local identifiers as described in (**LOCAL**). The list of locals to be defined is terminated by |. Append the run-time semantics given below to the current definition.

Run-time: (x_n ... x_2 x_1 --)

Initialize up to eight local identifiers as described in **13.6.1.0086 (LOCAL)**, each of which takes as its initial value the top stack item, removing it from the stack. Identifier *name*_{*j*} is initialized with x_j , identifier *name*₂ with x_2 , etc. When invoked, each local will return its value. The value of a local may be changed using **13.6.1.2295 TO**.

14. The optional Memory-Allocation word set

14.1 Introduction

14.2 Additional terms and notation

None.

14.3 Additional usage requirements

14.3.1 I/O Results data type

I/O results are single-cell numbers indicating the result of I/O operations. A value of zero indicates that the I/O operation completed successfully; other values and their meanings are implementation-defined.

Append table 14.1 to table 3.1.

Table 14.1 – Data types

Symbol	Data type	Size on stack
<i>ior</i>	I/O results	1 cell

14.3.2 Environmental queries

Append table 14.2 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 14.2 – Environmental query strings

String	Value data type	Constant?	Meaning
MEMORY-ALLOC	<i>flag</i>	no	memory-allocation word set present
MEMORY-ALLOC-EXT	<i>flag</i>	no	memory-allocation extensions word set present

14.3.3 Allocated regions

A program may address memory in data space regions made available by **ALLOCATE** or **RESIZE** and not yet released by **FREE**.

See: **3.3.3 Data space.**

14.4 Additional documentation requirements

14.4.1 System documentation

14.4.1.1 Implementation-defined options

- values and meaning of *ior* (14.3.1 I/O Results data type, 14.6.1.0707 **ALLOCATE**, 14.6.1.1605 **FREE**, 14.6.1.2145 **RESIZE**).

14.4.1.2 Ambiguous conditions

- no additional requirements.

14.4.1.3 Other system documentation

- no additional requirements.

14.4.2 Program documentation

- no additional requirements.

14.5 Compliance and labeling

14.5.1 ANS Forth systems

The phrase “Providing the Memory-Allocation word set” shall be appended to the label of any Standard System that provides all of the Memory-Allocation word set.

The phrase “Providing *name(s)* from the Memory-Allocation Extensions word set” shall be appended to the label of any Standard System that provides portions of the Memory-Allocation Extensions word set.

The phrase “Providing the Memory-Allocation Extensions word set” shall be appended to the label of any Standard System that provides all of the Memory-Allocation and Memory-Allocation Extensions word sets.

14.5.2 ANS Forth programs

The phrase “Requiring the Memory-Allocation word set” shall be appended to the label of Standard Programs that require the system to provide the Memory-Allocation word set.

The phrase “Requiring *name(s)* from the Memory-Allocation Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Memory-Allocation Extensions word set.

The phrase “Requiring the Memory-Allocation Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Memory-Allocation and Memory-Allocation Extensions word sets.

14.6 Glossary

14.6.1 Memory-Allocation words

14.6.1.0707 ALLOCATE

MEMORY

(*u* -- *a-addr* *ior*)

Allocate *u* address units of contiguous data space. The data-space pointer is unaffected by this operation. The initial content of the allocated space is undefined.

If the allocation succeeds, *a-addr* is the aligned starting address of the allocated space and *ior* is zero.

If the operation fails, *a-addr* does not represent a valid address and *ior* is the implementation-defined I/O result code.

See: 6.1.1650 **HERE**, 14.6.1.1605 **FREE**, 14.6.1.2145 **RESIZE**.

14.6.1.1605 FREE

MEMORY

(a-addr -- ior)

Return the contiguous region of data space indicated by *a-addr* to the system for later allocation. *a-addr* shall indicate a region of data space that was previously obtained by **ALLOCATE** or **RESIZE**. The data-space pointer is unaffected by this operation.

If the operation succeeds, *ior* is zero. If the operation fails, *ior* is the implementation-defined I/O result code.

See: **6.1.1650 HERE**, **14.6.1.0707 ALLOCATE**, **14.6.1.2145 RESIZE**.

14.6.1.2145 RESIZE

MEMORY

(a-addr₁ u -- a-addr₂ ior)

Change the allocation of the contiguous data space starting at the address *a-addr₁*, previously allocated by **ALLOCATE** or **RESIZE**, to *u* address units. *u* may be either larger or smaller than the current size of the region. The data-space pointer is unaffected by this operation.

If the operation succeeds, *a-addr₂* is the aligned starting address of *u* address units of allocated memory and *ior* is zero. *a-addr₂* may be, but need not be, the same as *a-addr₁*. If they are not the same, the values contained in the region at *a-addr₁* are copied to *a-addr₂*, up to the minimum size of either of the two regions. If they are the same, the values contained in the region are preserved to the minimum of *u* or the original size. If *a-addr₂* is not the same as *a-addr₁*, the region of memory at *a-addr₁* is returned to the system according to the operation of **FREE**.

If the operation fails, *a-addr₂* equals *a-addr₁*, the region of memory at *a-addr₁* is unaffected, and *ior* is the implementation-defined I/O result code.

See: **6.1.1650 HERE**, **14.6.1.0707 ALLOCATE**, **14.6.1.1605 FREE**.

14.6.2 Memory-Allocation extension words

None

15. The optional Programming-Tools word set

15.1 Introduction

This word set extension contains words most often used during the development of applications.

15.2 Additional terms and notation

None.

15.3 Additional usage requirements

15.3.1 Environmental queries

Append table 15.1 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 15.1 – Environmental query strings

String	Value data type	Constant?	Meaning
TOOLS	<i>flag</i>	no	programming-tools word set present
TOOLS-EXT	<i>flag</i>	no	programming-tools extensions word set present

15.3.2 The Forth dictionary

A program using the words **CODE** or **;CODE** associated with assembler code has an environmental dependency on that particular instruction set and assembler notation.

Programs using the words **EDITOR** or **ASSEMBLER** require the Search Order word set or an equivalent implementation-defined capability.

See: **3.3 The Forth dictionary.**

15.4 Additional documentation requirements

15.4.1 System documentation

15.4.1.1 Implementation-defined options

- ending sequence for input following **15.6.2.0470 ;CODE** and **15.6.2.0930 CODE;**
- manner of processing input following **15.6.2.0470 ;CODE** and **15.6.2.0930 CODE;**
- search-order capability for **15.6.2.1300 EDITOR** and **15.6.2.0740 ASSEMBLER** (**15.3.3 The Forth dictionary**);
- source and format of display by **15.6.1.2194 SEE.**

15.4.1.2 Ambiguous conditions

- * deleting the compilation word-list (**15.6.2.1580 FORGET**);
- * fewer than $u+1$ items on control-flow stack (**15.6.2.1015**, **15.6.2.1020**);
- * *name* can't be found (**15.6.2.1580 FORGET**);
- * *name* not defined via **6.1.1000 CREATE** (**15.6.2.0470 ;CODE**);
- * **6.1.2033 POSTPONE** applied to **15.6.2.2532 [IF]**;

- * reaching the end of the input source before matching **15.6.2.2531** [**ELSE**] or **15.6.2.2533** [**THEN**] (**15.6.2.2532** [**IF**]);
- * removing a needed definition (**15.6.2.1580** **FORGET**).

15.4.1.3 Other system documentation

- no additional requirements.

15.4.2 Program documentation

15.4.2.1 Environmental dependencies

- using the words **15.6.2.0470** ;**CODE** or **15.6.2.0930** **CODE**.

15.4.2.2 Other program documentation

- no additional requirements.

15.5 Compliance and labeling

15.5.1 ANS Forth systems

The phrase “Providing the Programming-Tools word set” shall be appended to the label of any Standard System that provides all of the Programming-Tools word set.

The phrase “Providing *name(s)* from the Programming-Tools Extensions word set” shall be appended to the label of any Standard System that provides portions of the Programming-Tools Extensions word set.

The phrase “Providing the Programming-Tools Extensions word set” shall be appended to the label of any Standard System that provides all of the Programming-Tools and Programming-Tools Extensions word sets.

15.5.2 ANS Forth programs

The phrase “Requiring the Programming-Tools word set” shall be appended to the label of Standard Programs that require the system to provide the Programming-Tools word set.

The phrase “Requiring *name(s)* from the Programming-Tools Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Programming-Tools Extensions word set.

The phrase “Requiring the Programming-Tools Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Programming-Tools and Programming-Tools Extensions word sets.

15.6 Glossary

15.6.1 Programming-Tools words

15.6.1.0220 .S "dot-s" **TOOLS**

(--)

Copy and display the values currently on the data stack. The format of the display is implementation-dependent.

.S may be implemented using pictured numeric output words. Consequently, its use may corrupt the transient region identified by **#>**.

See: **3.3.3.6 Other transient regions.**

15.6.1.0600 ? "question" **TOOLS**

(*a-addr* --)

Display the value stored at *a-addr*.

? may be implemented using pictured numeric output words. Consequently, its use may corrupt the transient region identified by **#>**.

See: **3.3.3.6 Other transient regions.**

15.6.1.1280 DUMP **TOOLS**

(*addr u* --)

Display the contents of *u* consecutive addresses starting at *addr*. The format of the display is implementation dependent.

DUMP may be implemented using pictured numeric output words. Consequently, its use may corrupt the transient region identified by **#>**.

See: **3.3.3.6 Other Transient Regions.**

15.6.1.2194 SEE **TOOLS**

("*<spaces>name*" --)

Display a human-readable representation of the named word's definition. The source of the representation (object-code decompilation, source block, etc.) and the particular form of the display is implementation defined.

SEE may be implemented using pictured numeric output words. Consequently, its use may corrupt the transient region identified by **#>**.

See: **3.3.3.6 Other transient regions.**

15.6.1.2465 WORDS

TOOLS

(--)

List the definition names in the first word list of the search order. The format of the display is implementation-dependent.

WORDS may be implemented using pictured numeric output words. Consequently, its use may corrupt the transient region identified by #>.

See: **3.3.3.6 Other Transient Regions.**

15.6.2 Programming-Tools extension words**15.6.2.0470 ;CODE**

"semicolon-code"

TOOLS EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: *colon-sys* --)

Append the run-time semantics below to the current definition. End the current definition, allow it to be found in the dictionary, and enter interpretation state, consuming *colon-sys*.

Process subsequent characters in the parse area in an implementation-defined manner, thus generating corresponding machine code. Those characters typically represent source code in a programming language, usually some form of assembly language. The process continues, refilling the input buffer as needed, until an implementation-defined ending sequence is processed.

Run-time: (--) (R: *nest-sys* --)

Replace the execution semantics of the most recent definition with the *name* execution semantics given below. Return control to the calling definition specified by *nest-sys*. An ambiguous condition exists if the most recent definition was not defined with **CREATE** or a user-defined word that calls **CREATE**.

name Execution: (*i*x* -- *j*x*)

Perform the machine code sequence that was generated following **;CODE**.

See: **6.1.1250 DOES>**.

15.6.2.0702 AHEAD

TOOLS EXT

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (C: -- *orig*)

Put the location of a new unresolved forward reference *orig* onto the control flow stack. Append the run-time semantics given below to the current definition. The semantics are incomplete until *orig* is resolved (e.g., by **THEN**).

Run-time: (--)

Continue execution at the location specified by the resolution of *orig*.

15.6.2.0740 ASSEMBLER

TOOLS EXT

(--)

Replace the first word list in the search order with the **ASSEMBLER** word list.

See: **16. The optional Search-Order word set.**

15.6.2.0830 BYE

TOOLS EXT

(--)

Return control to the host operating system, if any.

15.6.2.0930 CODE

TOOLS EXT

("*<spaces>name*" --)

Skip leading space delimiters. Parse *name* delimited by a space. Create a definition for *name*, called a "code definition," with the execution semantics defined below.

Process subsequent characters in the parse area in an implementation-defined manner, thus generating corresponding machine code. Those characters typically represent source code in a programming language, usually some form of assembly language. The process continues, refilling the input buffer as needed, until an implementation-defined ending sequence is processed.

name Execution: (*i*x* -- *j*x*)

Execute the machine code sequence that was generated following **CODE**.

See: **3.4.1 Parsing.**

15.6.2.1015 CS-PICK

"c-s-pick"

TOOLS EXT

Interpretation: Interpretation semantics for this word are undefined.

Execution: (C: *dest_u* ... *orig₀*|*dest₀* -- *dest_u* ... *orig₀*|*dest₀* *dest_u*)
(S: *u* --)

Remove *u*. Copy *dest_u* to the top of the control-flow stack. An ambiguous condition exists if there are less than *u+1* items, each of which shall be an *orig* or *dest*, on the control-flow stack before **CS-PICK** is executed.

If the control-flow stack is implemented using the data stack, *u* shall be the topmost item on the data stack.

15.6.2.1020 CS-ROLL "c-s-roll" **TOOLS EXT**

Interpretation: Interpretation semantics for this word are undefined.

Execution: (C: $orig_u|dest_u\ orig_{u-1}|dest_{u-1} \dots orig_0|dest_0$ -- $orig_{u-1}|dest_{u-1} \dots orig_0|dest_0\ orig_u|dest_u$)
(S: u --)

Remove u . Rotate $u+1$ elements on top of the control-flow stack so that $orig_u|dest_u$ is on top of the control-flow stack. An ambiguous condition exists if there are less than $u+1$ items, each of which shall be an *orig* or *dest*, on the control-flow stack before **CS-ROLL** is executed.

If the control-flow stack is implemented using the data stack, u shall be the topmost item on the data stack.

15.6.2.1300 EDITOR **TOOLS EXT**

(--)

Replace the first word list in the search order with the **EDITOR** word list.

See: **16. The Optional Search-Order Word Set.**

15.6.2.1580 FORGET **TOOLS EXT**

("<spaces>name" --)

Skip leading space delimiters. Parse *name* delimited by a space. Find *name*, then delete *name* from the dictionary along with all words added to the dictionary after *name*. An ambiguous condition exists if *name* cannot be found.

If the Search-Order word set is present, **FORGET** searches the compilation word list. An ambiguous condition exists if the compilation word list is deleted.

An ambiguous condition exists if **FORGET** removes a word required for correct execution.

Note: This word is obsolescent and is included as a concession to existing implementations.

See: **3.4.1 Parsing.**

15.6.2.2250 STATE **TOOLS EXT**

(-- *a-addr*)

Extend the semantics of **6.1.2250 STATE** to allow **;CODE** to change the value in **STATE**. A program shall not directly alter the contents of **STATE**.

See: **3.4 The Forth text interpreter**, **6.1.0450 :**, **6.1.0460 ;**, **6.1.0670 ABORT**, **6.1.2050 QUIT**, **6.1.2250 STATE**, **6.1.2500 [**, **6.1.2540]**, **6.2.0455 :NONAME**, **15.6.2.0470 ;CODE**.

15.6.2.2531 [ELSE] "bracket-else" TOOLS EXT

Compilation: Perform the execution semantics given below.

Execution: ("<spaces>name ..." --)

Skipping leading spaces, parse and discard space-delimited words from the parse area, including nested occurrences of **[IF]** ... **[THEN]** and **[IF]** ... **[ELSE]** ... **[THEN]**, until the word **[THEN]** has been parsed and discarded. If the parse area becomes exhausted, it is refilled as with **REFILL**. **[ELSE]** is an immediate word.

See: **3.4.1 Parsing**.

15.6.2.2532 [IF] "bracket-if" TOOLS EXT

Compilation: Perform the execution semantics given below.

Execution: (*flag* | *flag* "<spaces>name ..." --)

If *flag* is *true*, do nothing. Otherwise, skipping leading spaces, parse and discard space-delimited words from the parse area, including nested occurrences of **[IF]** ... **[THEN]** and **[IF]** ... **[ELSE]** ... **[THEN]**, until either the word **[ELSE]** or the word **[THEN]** has been parsed and discarded. If the parse area becomes exhausted, it is refilled as with **REFILL**. **[IF]** is an immediate word.

An ambiguous condition exists if **[IF]** is **POSTPONED**, or if the end of the input buffer is reached and cannot be refilled before the terminating **[ELSE]** or **[THEN]** is parsed.

See: **3.4.1 Parsing**.

15.6.2.2533 [THEN] "bracket-then" TOOLS EXT

Compilation: Perform the execution semantics given below.

Execution: (--)

Does nothing. **[THEN]** is an immediate word.

16. The optional Search-Order word set

16.1 Introduction

16.2 Additional terms and notation

compilation word list: the word list into which new definition names are placed.

search order: a list of word lists specifying the order in which the dictionary will be searched.

16.3 Additional usage requirements

16.3.1 Data types

Word list identifiers are implementation-dependent single-cell values that identify word lists.

Append table 16.1 to table 3.1.

Table 16.1 – Data types

Symbol	Data type	Size on stack
<i>wid</i>	word list identifiers	1 cell

See: **3.1 Data types**, **3.4.2 Finding definition names**, **3.4 The Forth text interpreter**.

16.3.2 Environmental queries

Append table 16.2 to table 3.5.

See: **3.2.6 Environmental queries**.

Table 16.2 – Environmental query strings

String	Value data type	Constant?	Meaning
SEARCH-ORDER	<i>flag</i>	no	search-order word set present
SEARCH-ORDER-EXT	<i>flag</i>	no	search-order extensions word set present
WORDLISTS	<i>n</i>	yes	maximum number of word lists usable in the search order

16.3.3 Finding definition names

When searching a word list for a definition name, the system shall search each word list from its last definition to its first. The search may encompass only a single word list, as with **SEARCH-WORDLIST**, or all the word lists in the search order, as with the text interpreter and **FIND**.

Changing the search order shall only affect the subsequent finding of definition names in the dictionary.

A system with the Search-Order word set shall allow at least eight word lists in the search order.

An ambiguous condition exists if a program changes the compilation word list during the compilation of a definition or before modification of the behavior of the most recently compiled definition with **;CODE**, **DOES>**, or **IMMEDIATE**.

A program that requires more than eight word lists in the search order has an environmental dependency.

See: **3.4.2 Finding definition names**

16.3.4 Contiguous regions

The regions of data space produced by the operations described in 3.3.3.2 **Contiguous regions** may be non-contiguous if **WORDLIST** is executed between allocations.

16.4 Additional documentation requirements

16.4.1 System documentation

16.4.1.1 Implementation-defined options

- maximum number of word lists in the search order (**16.3.3 Finding definition names, 16.6.1.2197 SET-ORDER**);
- minimum search order (**16.6.1.2197 SET-ORDER, 16.6.2.1965 ONLY**).

16.4.1.2 Ambiguous conditions

- * changing the compilation word list (**16.3.3 Finding definition names**);
- * search order empty (**16.6.2.2037 PREVIOUS**);
- * too many word lists in search order (**16.6.2.0715 ALSO**).

16.4.1.3 Other system documentation

- no additional requirements.

16.4.2 Program documentation

16.4.2.1 Environmental dependencies

- requiring more than eight word-lists in the search order (**16.3.3 Finding definition names**).

16.4.2.2 Other program documentation

- no additional requirements.

16.5 Compliance and labeling

16.5.1 ANS Forth systems

The phrase “Providing the Search-Order word set” shall be appended to the label of any Standard System that provides all of the Search-Order word set.

The phrase “Providing *name(s)* from the Search-Order Extensions word set” shall be appended to the label of any Standard System that provides portions of the Search-Order Extensions word set.

The phrase “Providing the Search-Order Extensions word set” shall be appended to the label of any Standard System that provides all of the Search-Order and Search-Order Extensions word sets.

16.5.2 ANS Forth programs

The phrase “Requiring the Search-Order word set” shall be appended to the label of Standard Programs that require the system to provide the Search-Order word set.

The phrase “Requiring *name(s)* from the Search-Order Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the Search-Order Extensions word set.

The phrase “Requiring the Search-Order Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the Search-Order and Search-Order Extensions word sets.

16.6 Glossary

16.6.1 Search-Order words

16.6.1.1180 DEFINITIONS

SEARCH

(--)

Make the compilation word list the same as the first word list in the search order. Specifies that the names of subsequent definitions will be placed in the compilation word list. Subsequent changes in the search order will not affect the compilation word list.

See: **16.3.3 Finding Definition Names.**

16.6.1.1550 FIND

SEARCH

Extend the semantics of **6.1.1550 FIND** to be:

(*c-addr* -- *c-addr* 0 | *xt* 1 | *xt* -1)

Find the definition named in the counted string at *c-addr*. If the definition is not found after searching all the word lists in the search order, return *c-addr* and zero. If the definition is found, return *xt*. If the definition is immediate, also return one (1); otherwise also return minus-one (-1). For a given string, the values returned by **FIND** while compiling may differ from those returned while not compiling.

See: **3.4.2 Finding definition names, 6.1.0070 ' , 6.1.1550 FIND, 6.1.2033 POSTPONE, 6.1.2510 ['], D.6.7 Immediacy.**

16.6.1.1595 FORTH-WORDLIST

SEARCH

(-- *wid*)

Return *wid*, the identifier of the word list that includes all standard words provided by the implementation. This word list is initially the compilation word list and is part of the initial search order.

16.6.1.1643 GET-CURRENT

SEARCH

(-- *wid*)

Return *wid*, the identifier of the compilation word list.

16.6.1.1647 GET-ORDER

SEARCH

(-- *wid*_{*n*} ... *wid*₁ *n*)

Returns the number of word lists *n* in the search order and the word list identifiers *wid*_{*n*} ... *wid*₁ identifying these word lists. *wid*₁ identifies the word list that is searched first, and *wid*_{*n*} the word list that is searched last. The search order is unaffected.

16.6.1.2192 SEARCH-WORDLIST SEARCH

(*c-addr u wid -- 0 | xt 1 | xt -1*)

Find the definition identified by the string *c-addr u* in the word list identified by *wid*. If the definition is not found, return zero. If the definition is found, return its execution token *xt* and one (1) if the definition is immediate, minus-one (-1) otherwise.

16.6.1.2195 SET-CURRENT SEARCH

(*wid --*)

Set the compilation word list to the word list identified by *wid*.

16.6.1.2197 SET-ORDER SEARCH

(*wid_n ... wid₁ n --*)

Set the search order to the word lists identified by *wid_n ... wid₁*. Subsequently, word list *wid₁* will be searched first, with word list *wid_n* searched last. If *n* is zero, empty the search order. If *n* is minus one, set the search order to the implementation-defined minimum search order. The minimum search order shall include the words **FORTH-WORDLIST** and **SET-ORDER**. A system shall allow *n* to be at least eight.

16.6.1.2460 WORDLIST SEARCH

(*-- wid*)

Creates a new empty word list, returning its word list identifier *wid*. The new word list may be returned from a pool of preallocated word lists or may be dynamically allocated in data space. A system shall allow the creation of at least 8 new word lists in addition to any provided as part of the system.

16.6.2 Search-Order extension words**16.6.2.0715 ALSO** SEARCH EXT

(*--*)

Transform the search order consisting of *wid_n, ... wid₂, wid₁* (where *wid₁* is searched first) into *wid_n, ... wid₂, wid₁, wid₁*. An ambiguous condition exists if there are too many word lists in the search order.

16.6.2.1590 FORTH SEARCH EXT

(*--*)

Transform the search order consisting of *wid_n, ... wid₂, wid₁* (where *wid₁* is searched first) into *wid_n, ... wid₂, wid₁^{FORTH-WORDLIST}*.

16.6.2.1965 ONLY SEARCH EXT

(--)

Set the search order to the implementation-defined minimum search order. The minimum search order shall include the words **FORTH-WORDLIST** and **SET-ORDER**.

16.6.2.1985 ORDER SEARCH EXT

(--)

Display the word lists in the search order in their search order sequence, from first searched to last searched. Also display the word list into which new definitions will be placed. The display format is implementation dependent.

ORDER may be implemented using pictured numeric output words. Consequently, its use may corrupt the transient region identified by **#>**.

See **3.3.3.6 Other Transient Regions**.

16.6.2.2037 PREVIOUS SEARCH EXT

(--)

Transform the search order consisting of $wid_n, \dots, wid_2, wid_1$ (where wid_1 is searched first) into wid_n, \dots, wid_2 . An ambiguous condition exists if the search order was empty before **PREVIOUS** was executed.

17. The optional String word set

17.1 Introduction

17.2 Additional terms and notation

None.

17.3 Additional usage requirements

Append table 17.1 to table 3.5.

See: **3.2.6 Environmental queries.**

Table 17.1 – Environmental query strings

String	Value data type	Constant?	Meaning
STRING	<i>flag</i>	no	string word set present
STRING-EXT	<i>flag</i>	no	string extensions word set present

17.4 Additional documentation requirements

None.

17.5 Compliance and labeling

17.5.1 ANS Forth systems

The phrase “Providing the String word set” shall be appended to the label of any Standard System that provides all of the String word set.

The phrase “Providing *name(s)* from the String Extensions word set” shall be appended to the label of any Standard System that provides portions of the String Extensions word set.

The phrase “Providing the String Extensions word set” shall be appended to the label of any Standard System that provides all of the String and String Extensions word sets.

17.5.2 ANS Forth programs

The phrase “Requiring the String word set” shall be appended to the label of Standard Programs that require the system to provide the String word set.

The phrase “Requiring *name(s)* from the String Extensions word set” shall be appended to the label of Standard Programs that require the system to provide portions of the String Extensions word set.

The phrase “Requiring the String Extensions word set” shall be appended to the label of Standard Programs that require the system to provide all of the String and String Extensions word sets.

17.6 Glossary

17.6.1 String words

17.6.1.0170 **-TRAILING** “dash-trailing” STRING

(*c-addr* u_1 -- *c-addr* u_2)

If u_1 is greater than zero, u_2 is equal to u_1 less the number of spaces at the end of the character string specified by *c-addr* u_1 . If u_1 is zero or the entire string consists of spaces, u_2 is zero.

17.6.1.0245 **/STRING** “slash-string” STRING

(*c-addr*₁ u_1 *n* -- *c-addr*₂ u_2)

Adjust the character string at *c-addr*₁ by *n* characters. The resulting character string, specified by *c-addr*₂ u_2 , begins at *c-addr*₁ plus *n* characters and is u_1 minus *n* characters long.

17.6.1.0780 **BLANK** STRING

(*c-addr* u --)

If u is greater than zero, store the character value for space in u consecutive character positions beginning at *c-addr*.

17.6.1.0910 **CMOVE** “c-move” STRING

(*c-addr*₁ *c-addr*₂ u --)

If u is greater than zero, copy u consecutive characters from the data space starting at *c-addr*₁ to that starting at *c-addr*₂, proceeding character-by-character from lower addresses to higher addresses.

Contrast with: **17.6.1.0920** **CMOVE>**.

17.6.1.0920 **CMOVE>** “c-move-up” STRING

(*c-addr*₁ *c-addr*₂ u --)

If u is greater than zero, copy u consecutive characters from the data space starting at *c-addr*₁ to that starting at *c-addr*₂, proceeding character-by-character from higher addresses to lower addresses.

Contrast with: **17.6.1.0910** **CMOVE**.

17.6.1.0935 COMPARE

STRING

(*c-addr₁ u₁ c-addr₂ u₂ -- n*)

Compare the string specified by *c-addr₁ u₁* to the string specified by *c-addr₂ u₂*. The strings are compared, beginning at the given addresses, character by character, up to the length of the shorter string or until a difference is found. If the two strings are identical, *n* is zero. If the two strings are identical up to the length of the shorter string, *n* is minus-one (-1) if *u₁* is less than *u₂* and one (1) otherwise. If the two strings are not identical up to the length of the shorter string, *n* is minus-one (-1) if the first non-matching character in the string specified by *c-addr₁ u₁* has a lesser numeric value than the corresponding character in the string specified by *c-addr₂ u₂* and one (1) otherwise.

17.6.1.2191 SEARCH

STRING

(*c-addr₁ u₁ c-addr₂ u₂ -- c-addr₃ u₃ flag*)

Search the string specified by *c-addr₁ u₁* for the string specified by *c-addr₂ u₂*. If *flag* is true, a match was found at *c-addr₃* with *u₃* characters remaining. If *flag* is false there was no match and *c-addr₃* is *c-addr₁* and *u₃* is *u₁*.

17.6.1.2212 SLITERAL

STRING

Interpretation: Interpretation semantics for this word are undefined.

Compilation: (*c-addr₁ u --*)

Append the run-time semantics given below to the current definition.

Run-time: (*-- c-addr₂ u*)

Return *c-addr₂ u* describing a string consisting of the characters specified by *c-addr₁ u* during compilation. A program shall not alter the returned string.

17.6.2 String extension words

None

A. Rationale (informative annex)

A.1 Introduction

A.1.1 Purpose

A.1.2 Scope

This Standard is more extensive than previous industry standards for the Forth language. Several things made this necessary:

- the desire to resolve conflicts between previous standards;
- the need to eliminate semantic ambiguities and other inadequacies;
- the requirement to standardize common practice, where possible resolving divergences in a way that minimizes the cost of compliance;
- the desire to standardize common system techniques, including those germane to hardware.

The result of the effort to satisfy all of these objectives is a Standard arranged so that the required word set remains small. Thus ANS Forth can be provided for resource-constrained embedded systems. Words beyond those in the required word set are organized into a number of optional word sets and their extensions, enabling implementation of tailored systems that are Standard.

When judging relative merits, the members of the X3J14 Technical Committee were guided by the following goals (listed in alphabetic order):

Consistency	The Standard provides a functionally complete set of words with minimal functional overlap.
Cost of compliance	This goal includes such issues as common practice, how much existing code would be broken by the proposed change, and the amount of effort required to bring existing applications and systems into conformity with the Standard.
Efficiency	Execution speed, memory compactness.
Portability	Words chosen for inclusion should be free of system-dependent features.
Readability	Forth definition names should clearly delineate their behavior. That behavior should have an apparent simplicity which supports rapid understanding. Forth should be easily taught and support readily maintained code.
Utility	Be judged to have sufficiently essential functionality and frequency of use to be deemed suitable for inclusion.

A.1.3 Document organization

A.1.3.1 Word sets

From the beginning, the X3J14 Technical Committee faced not only conflicting ideas as to what “real” Forth is, but also conflicting needs of the various groups within the Forth community. At one extreme were those who pressed for a “bare” Forth. At the other extreme were those who wanted a “fat” Forth. Many were somewhere in between. All were convinced of the rightness of their own position and of the wrongness of at least one of the two extremes. The committee’s composition reflected this full range of interests.

The approach we have taken is to define a Core word set establishing a greatest lower bound for required system functionality and to provide a portfolio of optional word sets for special purposes. This simple approach parallels the fundamental nature of Forth as an extensible language, and thereby achieves a kind of meta-extensibility.

With this key, high-level compromise, regardless of the actual makeup of the individual word sets, a firm and workable framework is established for the long term. One may or may not agree that there should be a Locals word set, or that the word **COMPILE**, belongs in the Core Extensions word set. But at least there is a mechanism whereby such things can be included in a logical and orderly manner.

Several implications of this scheme of optional word sets are significant.

First, ANS Forth systems can continue to be implemented on a greater range of hardware than could be claimed by almost any other single language. Since only the Core word set is required, very limited hardware will be able to accommodate an ANS Forth implementation.

Second, a greater degree of portability of applications, and of programmers, is anticipated. The optional word sets standardize various functions (e.g., floating point) that were widely implemented before, but not with uniform definition names and methodologies, nor the same levels of completeness. With such words now standardized in the optional word sets, communications between programmers -- verbally, via magazine or journal articles, etc. -- will leap to a new level of facility, and the shareability of code and applications should rise dramatically.

Third, ANS Forth systems may be designed to offer the user the power to selectively, even dynamically, include or exclude one or more of the optional word sets or portions thereof. Also, lower-priced products may be offered for the user who needs the Core word set and not much more. Thus, virtually unlimited flexibility will be available to the user.

But these advantages have a price. The burden is on the user to decide what capabilities are desired, and to select product offerings accordingly, especially when portability of applications is important. We do not expect most implementors to attempt to provide all word sets, but rather to select those most valuable to their intended markets.

The basic requirement is that if the implementor claims to have a particular optional word set the entire required portion of that word set must be available. If the implementor wishes to offer only part of an optional word set, it is acceptable to say, for example, "This system offers portions of the [named] word set," particularly if the selected or excluded words are itemized clearly.

Each optional word set will probably appeal to a particular constituency. For example, scientists performing complex mathematical analysis may place a higher value on the Floating-Point word set than programmers developing simple embedded controllers. As in the case of the core extensions, we expect implementors to offer those word sets they expect will be valued by their users.

Optional word sets may be offered in source form or otherwise factored so that the user may selectively load them.

The extensions to the optional word sets include words which are deemed less essential to performing the primary activity supported by the word set, though clearly relevant to it. As in the case of the Core Extensions, implementors may selectively add itemized subsets of a word set extension providing the labeling doesn't mislead the user into thinking incorrectly that all words are present.

A.2 Terms and notation

A.2.1 Definitions of terms

ambiguous condition

The response of a Standard System to an ambiguous condition is left to the discretion of the implementor. A Standard System need not explicitly detect or report the occurrence of ambiguous conditions.

cross compiler

Cross-compilers may be used to prepare a program for execution in an embedded system, or may be used to generate Forth kernels either for the same or a different run-time environment.

data field

In earlier standards, data fields were known as “parameter fields”.

On subroutine threaded Forth systems, everything is object code. There are no traditional code or data fields. Only a word defined by **CREATE** or by a word that calls **CREATE** has a data field. Only a data field defined via **CREATE** can be manipulated portably.

word set

This Standard recognizes that some functions, while useful in certain application areas, are not sufficiently general to justify requiring them in all Forth systems. Further, it is helpful to group Forth words according to related functions. These issues are dealt with using the concept of word sets.

The “Core” word set contains the essential body of words in a Forth system. It is the only “required” word set. Other word sets defined in this Standard are optional additions to make it possible to provide Standard Systems with tailored levels of functionality.

A.2.2 Notation

A.2.2.2 Stack notation

The use of *-sys*, *orig*, and *dest* data types in stack effect diagrams conveys two pieces of information. First, it warns the reader that many implementations use the data stack in unspecified ways for those purposes, so that items underneath on either the control-flow or data stacks are unavailable. Second, in cases where *orig* and *dest* are used, explicit pairing rules are documented on the assumption that all systems will implement that model so that its results are equivalent to employment of some stack, and that in fact many implementations do use the data stack for this purpose. However, nothing in this Standard requires that implementations actually employ the data stack (or any other) for this purpose so long as the implied behavior of the model is maintained.

A.3 Usage requirements

Forth systems are unusually simple to develop, in comparison with compilers for more conventional languages such as C. In addition to Forth systems supported by vendors, public-domain implementations and implementation guides have been widely available for nearly twenty years, and a large number of individuals have developed their own Forth systems. As a result, a variety of implementation approaches have developed, each optimized for a particular platform or target market.

The X3J14 Technical Committee has endeavored to accommodate this diversity by constraining implementors as little as possible, consistent with a goal of defining a standard interface between an underlying Forth System and an application program being developed on it.

Similarly, we will not undertake in this section to tell you how to implement a Forth System, but rather will provide some guidance as to what the minimum requirements are for systems that can properly claim compliance with this Standard.

A.3.1 Data-types

Most computers deal with arbitrary bit patterns. There is no way to determine by inspection whether a cell contains an address or an unsigned integer. The only meaning a datum possesses is the meaning assigned by an application.

When data are operated upon, the meaning of the result depends on the meaning assigned to the input values. Some combinations of input values produce meaningless results: for instance, what meaning can be assigned to the arithmetic sum of the ASCII representation of the character “A” and a TRUE flag? The answer may be “no meaning”; or alternatively, that operation might be the first step in producing a checksum. Context is the determiner.

The discipline of circumscribing meaning which a program may assign to various combinations of bit patterns is sometimes called *data typing*. Many computer languages impose explicit data typing and have compilers that prevent ill-defined operations.

Forth rarely explicitly imposes data-type restrictions. Still, data types implicitly do exist, and discipline is required, particularly if portability of programs is a goal. In Forth, it is incumbent upon the programmer (rather than the compiler) to determine that data are accurately typed.

This section attempts to offer guidance regarding *de facto* data typing in Forth.

A.3.1.2 Character types

The correct identification and proper manipulation of the character data type is beyond the purview of Forth's enforcement of data type by means of stack depth. Characters do not necessarily occupy the entire width of their single stack entry with meaningful data. While the distinction between signed and unsigned character is entirely absent from the formal specification of Forth, the tendency in practice is to treat characters as short positive integers when mathematical operations come into play.

Standard Character Set

- 1) The storage unit for the character data type (**C@**, **C!**, **FILL**, etc.) must be able to contain unsigned numbers from 0 through 255.
- 2) An implementation is not required to restrict character storage to that range, but a Standard Program without environmental dependencies cannot assume the ability to store numbers outside that range in a "char" location.
- 3) The allowed number representations are two's-complement, one's-complement, and signed-magnitude. Note that all of these number systems agree on the representation of positive numbers.
- 4) Since a "char" can store small positive numbers and since the character data type is a sub-range of the unsigned integer data type, **C!** must store the n least-significant bits of a cell ($8 \leq n \leq \text{bits/cell}$). Given the enumeration of allowed number representations and their known encodings, "**TRUE xx C! xx C@**" must leave a stack item with some number of bits set, which will thus will be accepted as non-zero by **IF**.
- 5) For the purposes of input (**KEY**, **ACCEPT**, etc.) and output (**EMIT**, **TYPE**, etc.), the encoding between numbers and human-readable symbols is ISO646/IRV (ASCII) within the range from 32 to 126 (space to ~). EBCDIC is out (most "EBCDIC" computer systems support ASCII too). Outside that range, it is up to the implementation. The obvious implementation choice is to use ASCII control characters for the range from 0 to 31, at least for the "displayable" characters in that range (TAB, RETURN, LINEFEED, FORMFEED). However, this is not as clear-cut as it may seem, because of the variation between operating systems on the treatment of those characters. For example, some systems TAB to 4 character boundaries, others to 8 character boundaries, and others to preset tab stops. Some systems perform an automatic linefeed after a carriage return, others perform an automatic carriage return after a linefeed, and others do neither.

The codes from 128 to 255 may eventually be standardized, either formally or informally, for use as international characters, such as the letters with diacritical marks found in many European languages. One such encoding is the 8-bit ISO Latin-1 character set. The computer marketplace at large will eventually decide which encoding set of those characters prevails. For Forth implementations running under an operating system (the majority of those running on standard platforms these days), most Forth implementors will probably choose to do whatever the system does, without performing any remapping within the domain of the Forth system itself.

- 6) A Standard Program can depend on the ability to receive any character in the range 32 ... 126 through **KEY**, and similarly to display the same set of characters with **EMIT**. If a program must be able to receive or display any particular character outside that range, it can declare an environmental dependency on the ability to receive or display that character.

7) A Standard Program cannot use control characters in definition names. However, a Standard System is not required to enforce this prohibition. Thus, existing systems that currently allow control characters in words names from **BLOCK** source may continue to allow them, and programs running on those systems will continue to work. In text file source, the parsing action with space as a delimiter (e.g., **BL WORD**) treats control characters the same as spaces. This effectively implies that you cannot use control characters in definition names from text-file source, since the text interpreter will treat the control characters as delimiters. Note that this “control-character folding” applies only when space is the delimiter, thus the phrase “**CHAR) WORD**” may collect a string containing control characters.

Storage and retrieval

Characters are transferred from the data stack to memory by **C!** and from memory to the data stack by **C@**. A number of lower-significance bits equivalent to the implementation-dependent width of a *character* are transferred from a popped data stack entry to an address by the action of **C!** without affecting any bits which may comprise the higher-significance portion of the cell at the destination address; however, the action of **C@** clears all higher-significance bits of the data stack entry which it pushes that are beyond the implementation-dependent width of a character (which may include implementation-defined display information in the higher-significance bits). The programmer should keep in mind that operating upon arbitrary stack entries with words intended for the character data type may result in truncation of such data.

Manipulation on the stack

In addition to **C@** and **C!**, characters are moved to, from and upon the data stack by the following words:

>R ?DUP DROP DUP OVER PICK R> R@ ROLL ROT SWAP

Additional operations

The following mathematical operators are valid for character data:

+ - * / /MOD MOD

The following comparison and bitwise operators may be valid for characters, keeping in mind that display information cached in the most significant bits of characters in an implementation-defined fashion may have to be masked or otherwise dealt with:

**AND OR > < U> U< = <> 0= 0<> MAX MIN
LSHIFT RSHIFT**

A.3.1.3 Single-cell types

A single-cell stack entry viewed without regard to typing is the fundamental data type of Forth. All other data types are actually represented by one or more single-cell stack entries.

Storage and retrieval

Single-cell data are transferred from the stack to memory by **!**; from memory to the stack by **@**. All bits are transferred in both directions and no type checking of any sort is performed, nor does the Standard System check that a memory address used by **!** or **@** is properly aligned or properly sized to hold the datum thus transferred.

Manipulation on the stack

Here is a selection of the most important words which move single-cell data to, from and upon the data stack:

! @ >R ?DUP DROP DUP OVER PICK R> R@ ROLL ROT SWAP

Comparison operators

The following comparison operators are universally valid for one or more single cells:

= <> 0= 0<>

A.3.1.3.1 Flags

A **FALSE** flag is a single-cell datum with all bits unset, and a **TRUE** flag is a single-cell datum with all bits set. While Forth words which test flags accept any non-null bit pattern as true, there exists the concept of the *well-formed flag*. If an operation whose result is to be used as a flag may produce any bit-mask other than **TRUE** or **FALSE**, the recommended discipline is to convert the result to a well-formed flag by means of the Forth word **0<>** so that the result of any subsequent logical operations on the flag will be predictable.

In addition to the words which move, fetch and store single-cell items, the following words are valid for operations on one or more flag data residing on the data stack:

AND OR XOR INVERT

A.3.1.3.2 Integers

A single-cell datum may be treated by a Standard Program as a signed integer. Moving and storing such data is performed as for any single-cell data. In addition to the universally-applicable operators for single-cell data specified above, the following mathematical and comparison operators are valid for single-cell signed integers:

*** */ */MOD /MOD MOD + +! - / 1+ 1- ABS MAX MIN NEGATE
0< 0> < >**

Given the same number of bits, unsigned integers usually represent twice the number of absolute values representable by signed integers.

A single-cell datum may be treated by a Standard Program as an unsigned integer. Moving and storing such data is performed as for any single-cell data. In addition, the following mathematical and comparison operators are valid for single-cell unsigned integers:

UM* UM/MOD + +! - 1+ 1- * U< U>

A.3.1.3.3 Addresses

An address is uniquely represented as a single cell unsigned number and can be treated as such when being moved to, from, or upon the stack. Conversely, each unsigned number represents a unique address (which is not necessarily an address of accessible memory). This one-to-one relationship between addresses and unsigned numbers forces an equivalence between address arithmetic and the corresponding operations on unsigned numbers.

Several operators are provided specifically for address arithmetic:

CHAR+ CHARS CELL+ CELLS

and, if the floating-point word set is present:

FLOAT+ FLOATS SFLOAT+ SFLOATS DFLOAT+ DFLOATS

A Standard Program may never assume a particular correspondence between a Forth address and the physical address to which it is mapped.

A.3.1.3.4 Counted strings

The trend in ANS Forth is to move toward the consistent use of the “*c-addr u*” representation of strings on the stack. The use of the alternate “address of counted string” stack representation is discouraged. The traditional Forth words **WORD** and **FIND** continue to use the “address of counted string” representation for

historical reasons. The new word **C"**, added as a porting aid for existing programs, also uses the counted string representation.

Counted strings remain useful as a way to store strings in memory. This use is not discouraged, but when references to such strings appear on the stack, it is preferable to use the "*c-addr u*" representation.

A.3.1.3.5 Execution tokens

The association between an execution token and a definition is static. Once made, it does not change with changes in the search order or anything else. However it may not be unique, e.g., the phrases

```
' 1+ and
' CHAR+
```

might return the same value.

A.3.1.4 Cell-pair types

Storage and retrieval

Two operators are provided to fetch and store cell pairs:

```
2@ 2!
```

Manipulation on the stack

Additionally, these operators may be used to move cell pairs from, to and upon the stack:

```
2>R 2DROP 2DUP 2OVER 2R> 2SWAP 2ROT
```

Comparison

The following comparison operations are universally valid for cell pairs:

```
D= D0=
```

A.3.1.4.1 Double-Cell Integers

If a double-cell integer is to be treated as signed, the following comparison and mathematical operations are valid:

```
D+ D- D< D0< DABS DMAX DMIN DNEGATE M*/ M+
```

If a double-cell integer is to be treated as unsigned, the following comparison and mathematical operations are valid:

```
D+ D- UM/MOD DU<
```

A.3.1.4.2 Character strings

See: A.3.1.3.4 Counted Strings.

A.3.2 The Implementation environment

A.3.2.1 Numbers

Traditionally, Forth has been implemented on two's-complement machines where there is a one-to-one mapping of signed numbers to unsigned numbers—any single cell item can be viewed either as a signed or unsigned number. Indeed, the signed representation of any positive number is identical to the equivalent unsigned representation. Further, addresses are treated as unsigned numbers: there is no distinct pointer type. Arithmetic ordering on two's complement machines allows + and - to work on both signed and unsigned numbers. This arithmetic behavior is deeply embedded in common Forth practice. As a consequence of these behaviors, the likely ranges of signed and unsigned numbers for implementations hosted on each of the permissible arithmetic architectures is:

Arithmetic architecture	signed numbers	unsigned numbers
Two's complement	$-n-1$ to n	0 to $2n+1$
One's complement	$-n$ to n	0 to n
Signed magnitude	$-n$ to n	0 to n

where n is the largest positive signed number. For all three architectures, signed numbers in the 0 to N range are bitwise identical to the corresponding unsigned number. Note that unsigned numbers on a signed magnitude machine are equivalent to signed non-negative numbers as a consequence of the forced correspondence between addresses and unsigned numbers and of the required behavior of + and -.

For reference, these number representations may be defined by the way that **NEGATE** is implemented:

```
two's complement:  : NEGATE INVERT 1+ ;
one's complement:  : NEGATE INVERT ;
signed-magnitude: : NEGATE HIGH-BIT XOR ;
```

where **HIGH-BIT** is a bit mask with only the most-significant bit set. Note that all of these number systems agree on the representation of non-negative numbers.

Per **3.2.1.1 Internal number representation** and **6.1.0270 0=**, the implementor must ensure that no standard or supported word return negative zero for any numeric (non-Boolean or flag) result. Many existing programmer assumptions will be violated otherwise.

There is no requirement to implement circular unsigned arithmetic, nor to set the range of unsigned numbers to the full size of a cell. There is historical precedent for limiting the range of u to that of $+n$, which is permissible when the cell size is greater than 16 bits.

A.3.2.1.2 Digit conversion

For example, an implementation might convert the characters "a" through "z" identically to the characters "A" through "Z," or it might treat the characters "[" through "~" as additional digits with decimal values 35 through 71, respectively.

A.3.2.2 Arithmetic

A.3.2.2.1 division

The Forth-79 Standard specifies that the signed division operators ($/$, $/\mathbf{MOD}$, \mathbf{MOD} , $*/\mathbf{MOD}$, and $*/$) round non-integer quotients towards zero (symmetric division). Forth 83 changed the semantics of these operators to round towards negative infinity (floored division). Some in the Forth community have declined to convert systems and applications from the Forth-79 to the Forth-83 divide. To resolve this issue, an ANS Forth system is permitted to supply either floored or symmetric operators. In addition, ANS Forth systems must provide a floored division primitive (**FM/MOD**), a symmetric division primitive (**SM/REM**), and a mixed precision multiplication operator (**M***).

This compromise protects the investment made in current Forth applications; Forth-79 and Forth-83 programs are automatically compliant with ANS Forth with respect to division. In practice, the rounding direction rarely matters to applications. However, if a program requires a specific rounding direction, it can use the floored division primitive **FM/MOD** or the symmetric division primitive **SM/REM** to construct a division operator of the desired flavor. This simple technique can be used to convert Forth-79 and Forth-83 programs to ANS Forth without any analysis of the original programs.

A.3.2.2.2 Other integer operations

Whether underflow occurs depends on the data-type of the result. For example, the phrase `1 2 -` underflows if the result is unsigned and produces the valid signed result -1.

The requirement that control-flow words are properly balanced by other control-flow words makes reasonable the description of a compile-time implementation-defined *control-flow stack*. There is no prescription as to how the control-flow stack is implemented, e.g., data stack, linked list, special array. Each element of the control-flow stack mentioned above is the same size.

With these tools, the remaining basic control-structure elements, shown in **figure A.2**, can be defined. The stack notation used here for immediate words is (*compilation / execution*).

```

: WHILE ( dest -- orig dest / flag -- )
  \ conditional exit from loops
  POSTPONE IF      \ conditional forward branch
  1 CS-ROLL ; \ keep dest on top
IMMEDIATE

: REPEAT ( orig dest -- / -- )
  \ resolve a single WHILE and return to BEGIN
  POSTPONE AGAIN \ uncond. backward branch to dest
  POSTPONE THEN ; \ resolve forward branch from orig
IMMEDIATE

: ELSE ( orig1 -- orig2 / -- )
  \ resolve IF supplying alternate execution
  POSTPONE AHEAD \ unconditional forward branch orig2
  1 CS-ROLL \ put orig1 back on top
  POSTPONE THEN ; \ resolve forward branch from orig1
IMMEDIATE

```




Figure A.2 – Additional basic control-flow patterns.

Forth control flow provides a solution for well-known problems with strictly structured programming.

The basic control structures can be supplemented, as shown in the examples in **figure A.3**, with additional **WHILE**s in **BEGIN ... UNTIL** and **BEGIN ... WHILE ... REPEAT** structures. However, for each additional **WHILE** there must be a **THEN** at the end of the structure. **THEN** completes the syntax with **WHILE** and indicates where to continue execution when the **WHILE** transfers control. The use of more than one additional **WHILE** is possible but not common. Note that if the user finds this use of **THEN** undesirable, an alias with a more likable name could be defined.

Additional actions may be performed between the control flow word (the **REPEAT** or **UNTIL**) and the **THEN** that matches the additional **WHILE**. Further, if additional actions are desired for normal termination and early termination, the alternative actions may be separated by the ordinary Forth **ELSE**. The termination actions are all specified after the body of the loop.

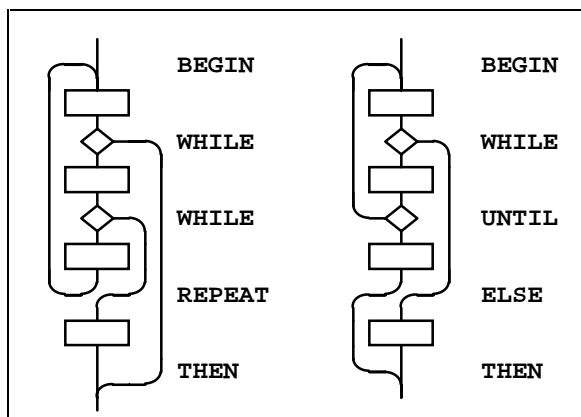


Figure A.3 – Extended control-flow pattern examples.

Note that **REPEAT** creates an anomaly when matching the **WHILE** with **ELSE** or **THEN**, most notable when compared with the **BEGIN...UNTIL** case. That is, there will be one less **ELSE** or **THEN** than there are **WHILE**s because **REPEAT** resolves one **THEN**. As above, if the user finds this count mismatch undesirable, **REPEAT** could be replaced in-line by its own definition.

Other loop-exit control-flow words, and even other loops, can be defined. The only requirements are that the control-flow stack is properly maintained and manipulated.

The simple implementation of the ANS Forth **CASE** structure below is an example of control structure extension. Note the maintenance of the data stack to prevent interference with the possible control-flow stack usage.

```

0 CONSTANT CASE IMMEDIATE ( init count of OFs )

: OF ( #of -- orig #of+1 / x -- )
  1+ ( count OFs )
  >R ( move off the stack in case the control-flow )
    ( stack is the data stack. )
  POSTPONE OVER POSTPONE = ( copy and test case value )
  POSTPONE IF ( add orig to control flow stack )
  POSTPONE DROP ( discards case value if = )
  R> ; ( we can bring count back now )
IMMEDIATE

: ENDOF ( orig1 #of -- orig2 #of )
  >R ( move off the stack in case the control-flow )
    ( stack is the data stack. )
  POSTPONE ELSE
  R> ; ( we can bring count back now )
IMMEDIATE

: ENDCASE ( orig1..orign #of -- )
  POSTPONE DROP ( discard case value )
  0 ?DO
    POSTPONE THEN
  LOOP ;
IMMEDIATE

```


A.3.2.3.3 Return stack

The restrictions in section **3.2.3.3 Return stack** are necessary if implementations are to be allowed to place loop parameters on the return stack.

A.3.2.6 Environmental queries

The size in address units of various data types may be determined by phrases such as **1 CHARS**. Similarly, alignment may be determined by phrases such as **1 ALIGN**.

The environmental queries are divided into two groups: those that always produce the same value and those that might not. The former groups include entries such as **MAX-N**. This information is fixed by the hardware or by the design of the Forth system; a user is guaranteed that asking the question once is sufficient.

The other group of queries are for things that may legitimately change over time. For example an application might test for the presence of the Double Number word set using an environment query. If it is missing, the system could invoke a system-dependent process to load the word set. The system is permitted to change **ENVIRONMENT?**'s database so that subsequent queries about it indicate that it is present.

Note that a query that returns an "unknown" response could produce a "known" result on a subsequent query.

A.3.3 The Forth dictionary

A Standard Program may redefine a standard word with a non-standard definition. The program is still Standard (since it can be built on any Standard System), but the effect is to make the combined entity (Standard System plus Standard Program) a non-standard system.

A.3.3.1 Name space

A.3.3.1.2 Definition names

The language in this section is there to ensure the portability of Standard Programs. If a program uses something outside the Standard that it does not provide itself, there is no guarantee that another implementation will have what the program needs to run. There is no intent whatsoever to imply that all Forth programs will be somehow lacking or inferior because they are not standard; some of the finest jewels of the programmer's art will be non-standard. At the same time, the committee is trying to ensure that a program labeled "Standard" will meet certain expectations, particularly with regard to portability.

In many system environments the input source is unable to supply certain non-graphic characters due to external factors, such as the use of those characters for flow control or editing. In addition, when interpreting from a text file, the parsing function specifically treats non-graphic characters like spaces; thus words received by the text interpreter will not contain embedded non-graphic characters. To allow implementations in such environments to call themselves Standard, this minor restriction on Standard Programs is necessary.

A Standard System is allowed to permit the creation of definition names containing non-graphic characters. Historically, such names were used for keyboard editing functions and "invisible" words.

A.3.3.2 Code space

A.3.3.3 Data space

The words **#TIB**, **>IN**, **BASE**, **BLK**, **SCR**, **SOURCE**, **SOURCE-ID**, **STATE**, and **TIB** contain information used by the Forth system in its operation and may be of use to the application. Any assumption made by the application about data available in the Forth system it did not store other than the data just listed is an environmental dependency.

There is no point in specifying (in the Standard) both what is and what is not addressable.

A Standard Program may NOT address:

- Directly into the data or return stacks;
- Into a definition's data field if not stored by the application.

The read-only restrictions arise because some Forth systems run from ROM and some share I/O buffers with other users or systems. Portable programs cannot know which areas are affected, hence the general restrictions.

A.3.3.3.1 Address alignment

Many processors have restrictions on the addresses that can be used by memory access instructions. For example, on a Motorola 68000, 16-bit or 32-bit data can be accessed only at even addresses. Other examples include RISC architectures where 16-bit data can be loaded or stored only at even addresses and 32-bit data only at addresses that are multiples of four.

An implementor of ANS Forth can handle these alignment restrictions in one of two ways. Forth's memory access words (@, !, +!, etc.) could be implemented in terms of smaller-width access instructions which have no alignment restrictions. For example, on a 68000 Forth with 16-bit cells, @ could be implemented with two 68000 byte-fetch instructions and a reassembly of the bytes into a 16-bit cell. Although this conceals hardware restrictions from the programmer, it is inefficient, and may have unintended side effects in some hardware environments. An alternate implementation of ANS Forth could define each memory-access word using the native instructions that most closely match the word's function. On a 68000 Forth with 16-bit cells, @ would use the 68000's 16-bit move instruction. In this case, responsibility for giving @ a correctly-aligned address falls on the programmer. A portable ANS Forth program must assume that alignment may be required and follow the requirements of this section.

A.3.3.3.2 Contiguous regions

The data space of a Forth system comes in discontinuous regions! The location of some regions is provided by the system, some by the program. Data space is contiguous within regions, allowing address arithmetic to generate valid addresses only within a single region. A Standard Program cannot make any assumptions about the relative placement of multiple regions in memory.

Section 3.3.3.2 does prescribe conditions under which contiguous regions of data space may be obtained. For example:

```
CREATE TABLE 1 C, 2 C, ALIGN 1000 , 2000 ,
```

makes a table whose address is returned by **TABLE**. In accessing this table,

```
TABLE C@           will return 1
TABLE CHAR+ C@     will return 2
TABLE 2 CHARS + ALIGNED @       will return 1000
TABLE 2 CHARS + ALIGNED CELL+ @ will return 2000.
```

Similarly,

```
CREATE DATA 1000 ALLOT
```

makes an array 1000 address units in size. A more portable strategy would define the array in application units, such as:

```
500 CONSTANT NCELLS
CREATE CELL-DATA NCELLS CELLS ALLOT
```

This array can be indexed like this:

```
: LOOK NCELLS 0 DO CELL-DATA I CELLS + ? LOOP ;
```

A.3.3.3.6 Other transient regions

In many existing Forth systems, these areas are at **HERE** or just beyond it, hence the many restrictions.

$(2*n)+2$ is the size of a character string containing the unpunctuated binary representation of the maximum double number with a leading minus sign and a trailing space.

Implementation note: Since the minimum value of n is 16, the absolute minimum size of the pictured numeric output string is 34 characters. But if your implementation has a larger n , you must also increase the size of the pictured numeric output string.

A.3.4 The Forth text interpreter

A.3.4.3 Semantics

The “initiation semantics” correspond to the code that is executed upon entering a definition, analogous to the code executed by **EXIT** upon leaving a definition. The “run-time semantics” correspond to code fragments, such as literals or branches, that are compiled inside colon definitions by words with explicit compilation semantics.

In a Forth cross-compiler, the execution semantics may be specified to occur in the host system only, the target system only, or in both systems. For example, it may be appropriate for words such as **CELLS** to execute on the host system returning a value describing the target, for colon definitions to execute only on the target, and for **CONSTANT** and **VARIABLE** to have execution behaviors on both systems. Details of cross-compiler behavior are beyond the scope of this Standard.

A.3.4.3.2 Interpretation semantics

For a variety of reasons, this Standard does not define interpretation semantics for every word. Examples of these words are **>R**, **.**, **DO**, and **IF**. Nothing in this Standard precludes an implementation from providing interpretation semantics for these words, such as interactive control-flow words. However, a Standard Program may not use them in interpretation state.

A.3.4.5 Compilation

Compiler recursion at the definition level consumes excessive resources, especially to support locals. The Technical Committee does not believe that the benefits justify the costs. Nesting definitions is also not common practice and won't work on many systems.

A.4 Documentation requirements

A.4.1 System documentation

A.4.2 Program documentation

A.5 Compliance and labeling

A.5.1 ANS Forth system

Section 5.1 defines the criteria that a system must meet in order to justify the label “ANS Forth System.” Briefly, the minimum requirement is that the system must “implement” the Core word set. There are several ways in which this requirement may be met. The most obvious is that all Core words may be in a pre-compiled kernel. This is not the only way of satisfying the requirement, however. For example, some words may be provided in source blocks or files with instructions explaining how to add them to the system if they are needed. So long as the words are provided in such a way that the user can obtain access to them with a clear and straightforward procedure, they may be considered to be present.

A Forth cross-compiler has many characteristics in common with an ANS Forth System, in that both use similar compiling tools to process a program. However, in order to fully specify an ANS Forth cross-compiler it would be necessary to address complex issues dealing with compilation and execution semantics in both host and target environments as well as ROMability issues. The level of effort to do this properly has proved to be impractical at this time. As a result, although it may be possible for a Forth cross-compiler to correctly prepare an ANS Forth program for execution in a target environment, it is inappropriate for a cross-compiler to be labeled an ANS Forth System.

A.5.2 ANS Forth programs

A.5.2.2 Program labeling

Declaring an environmental dependency should not be considered undesirable, merely an acknowledgment that the author has taken advantage of some assumed architecture. For example, most computers in common use are based on two's complement binary arithmetic. By acknowledging an environmental dependency on this architecture, a programmer becomes entitled to use the number -1 to represent all bits set without significantly restricting the portability of the program.

Because all programs require space for data and instructions, and time to execute those instructions, they depend on the presence of an environment providing those resources. It is impossible to predict how little of some of these resources (e.g. stack space) might be necessary to perform some task, so this Standard does not do so.

On the other hand, as a program requires increasing levels of resources, there will probably be successively fewer systems on which it will execute successfully. An algorithm requiring an array of 10^9 octets might run on fewer computers than one requiring only 10^3 .

Since there is also no way of knowing what minimum level of resources will be implemented in a system useful for at least some tasks, any program performing real work labeled simply an "ANS Forth Program" is unlikely to be labeled correctly.

A.6 Glossary

In this and following sections we present rationales for the handling of specific words: why we included them, why we placed them in certain word sets, or why we specified their names or meaning as we did.

Words in this section are organized by word set, retaining their index numbers for easy cross-referencing to the glossary.

Historically, many Forth systems have been written in Forth. Many of the words in Forth originally had as their primary purpose support of the Forth system itself. For example, **WORD** and **FIND** are often used as the principle instruments of the Forth text interpreter, and **CREATE** in many systems is the primitive for building dictionary entries. In defining words such as these in a standard way, we have endeavored not to do so in such a way as to preclude their use by implementors. One of the features of Forth that has endeared it to its users is that the same tools that are used to implement the system are available to the application programmer — a result of this approach is the compactness and efficiency that characterizes most Forth implementations.

A.6.1 Core words

A.6.1.0070 '

Typical use: ... ' name .

Many Forth systems use a state-smart tick. Many do not. ANS Forth follows the usage of Forth 83.

See A.3.4.2.2 Interpretation semantics, A.6.1.1550 **FIND**.

A.6.1.0080 (

Typical use: ... (ccc) ...

A.6.1.0140 +LOOP

Typical use:

```
: X ... limit first DO ... step +LOOP ;
```

A.6.1.0150 ,

The use of , (comma) for compiling execution tokens is not portable.

See 6.2.0945 **COMPILE,,**

A.6.1.0190 ."

Typical use: : X" ccc" ... ;

An implementation may define interpretation semantics for ." if desired. In one plausible implementation, interpreting ." would display the delimited message. In another plausible implementation, interpreting ." would compile code to display the message later. In still another plausible implementation, interpreting ." would be treated as an exception. Given this variation a Standard Program may not use ." while interpreting. Similarly, a Standard Program may not compile **POSTPONE** ." inside a new word, and then use that word while interpreting.

A.6.1.0320 2*

Historically, 2* has been implemented on two's-complement machines as a logical left-shift instruction. Multiplication by two is an efficient side-effect on these machines. However, shifting implies a knowledge of the significance and position of bits in a cell. While the name implies multiplication, most implementors have used a hardware left shift to implement 2*.

A.6.1.0330 2/

This word has the same common usage and misnaming implications as **2***. It is often implemented on two's-complement machines with a hardware right shift that propagates the sign bit.

A.6.1.0350 2@

With **2@** the storage order is specified by the Standard.

A.6.1.0450 :

Typical use: `: name ... ;`

In Forth 83, this word was specified to alter the search order. This specification is explicitly removed in this Standard. We believe that in most cases this has no effect; however, systems that allow many search orders found the Forth-83 behavior of colon very undesirable.

Note that colon does not itself invoke the compiler. Colon sets compilation state so that later words in the parse area are compiled.

A.6.1.0460 ;

Typical use: `: name ... ;`

One function performed by both `; and ;CODE` is to allow the current definition to be found in the dictionary. If the current definition was created by `:NONAME` the current definition has no definition name and thus cannot be found in the dictionary. If `:NONAME` is implemented the Forth compiler must maintain enough information about the current definition to allow `; and ;CODE` to determine whether or not any action must be taken to allow it to be found.

A.6.1.0550 >BODY

a-addr is the address that **HERE** would have returned had it been executed immediately after the execution of the **CREATE** that defined *xt*.

A.6.1.0680 ABORT"

Typical use: `: X ... test ABORT" ccc" ... ;`

A.6.1.0695 ACCEPT

Previous standards specified that collection of the input string terminates when either a "return" is received or when $+n_l$ characters have been received. Terminating when $+n_l$ characters have been received is difficult, expensive, or impossible to implement in some system environments. Consequently, a number of existing implementations do not comply with this requirement. Since line-editing and collection functions are often implemented by system components beyond the control of the Forth implementation, this Standard imposes no such requirement. A Standard Program may only assume that it can receive an input string with **ACCEPT** or **EXPECT**. The detailed sequence of user actions necessary to prepare and transmit that line are beyond the scope of this Standard.

Specification of a non-zero, positive integer count ($+n_l$) for **ACCEPT** allows some implementors to continue their practice of using a zero or negative value as a flag to trigger special behavior. Insofar as such behavior is outside the Standard, Standard Programs cannot depend upon it, but the Technical Committee doesn't wish to preclude it unnecessarily. Since actual values are almost always small integers, no functionality is impaired by this restriction.

ACCEPT and **EXPECT** perform similar functions. **ACCEPT** is recommended for new programs, and future use of **EXPECT** is discouraged.

It is recommended that all non-graphic characters be reserved for editing or control functions and not be stored in the input string.

Commonly, when the user is preparing an input string to be transmitted to a program, the system allows the user to edit that string and correct mistakes before transmitting the final version of the string. The editing function is supplied sometimes by the Forth system itself, and sometimes by external system software or hardware. Thus, control characters and functions may not be available on all systems. In the usual case, the end of the editing process and final transmission of the string is signified by the user pressing a “Return” or “Enter” key.

As in previous standards, **EXPECT** returns the input string immediately after the requested number of characters are entered, as well as when a line terminator is received. The “automatic termination after specified count of characters have been entered” behavior is widely considered undesirable because the user “loses control” of the input editing process at a potentially unknown time (the user does not necessarily know how many characters were requested from **EXPECT**). Thus **EXPECT** and **SPAN** have been made obsolescent and exist in the Standard only as a concession to existing implementations. If **EXPECT** exists in a Standard System it must have the “automatic termination” behavior.

ACCEPT does not have the “automatic termination” behavior of **EXPECT**. However, because external system hardware and software may perform the **ACCEPT** function, when a line terminator is received the action of the cursor, and therefore the display, is implementation-defined. It is recommended that the cursor remain immediately following the entered text after a line terminator is received.

A.6.1.0705 ALIGN

In this Standard we have attempted to provide transportability across various CPU architectures. One of the frequent causes of transportability problems is the requirement of cell-aligned addresses on some CPUs. On these systems, **ALIGN** and **ALIGNED** may be required to build and traverse data structures built with **C,**. Implementors may define these words as no-ops on systems for which they aren’t functional.

A.6.1.0706 ALIGNED

See **A.6.1.0705 ALIGN**.

A.6.1.0760 BEGIN

Typical use:

```
: X ... BEGIN ... test UNTIL ;
```

or

```
: X ... BEGIN ... test WHILE ... REPEAT ;
```

A.6.1.0770 BL

Because space is used throughout Forth as the standard delimiter, this word is the only way a program has to find and use the system value of “space”. The value of a space character can not be obtained with **CHAR**, for instance.

A.6.1.0880 CELL+

As with **ALIGN** and **ALIGNED**, the words **CELL** and **CELL+** were added to aid in transportability across systems with different cell sizes. They are intended to be used in manipulating indexes and addresses in integral numbers of cell-widths.

Example:

```
2VARIABLE DATA
```

```

0 100 DATA 2!
DATA @ . 100
DATA CELL+ @ . 0

```

A.6.1.0890 CELLS

See: **A.6.1.0880 CELL+**.

Example: **CREATE NUMBERS 100 CELLS ALLOT**

(Allots space in the array **NUMBERS** for 100 cells of data.)

A.6.1.0895 CHAR

Typical use: **... CHAR A CONSTANT "A" ...**

A.6.1.0950 CONSTANT

Typical use: **... DECIMAL 10 CONSTANT TEN ...**

A.6.1.1000 CREATE

The data-field address of a word defined by **CREATE** is given by the data-space pointer immediately following the execution of **CREATE**

Reservation of data field space is typically done with **ALLOT**.

Typical use: **... CREATE SOMETHING ...**

A.6.1.1240 DO

Typical use:

```
: X ... limit first DO ... LOOP ;
```

or

```
: X ... limit first DO ... step +LOOP ;
```

A.6.1.1250 DOES>

Typical use: **: X ... DOES> ... ;**

Following **DOES>**, a Standard Program may not make any assumptions regarding the ability to find either the name of the definition containing the **DOES>** or any previous definition whose name may be concealed by it. **DOES>** effectively ends one definition and begins another as far as local variables and control-flow structures are concerned. The compilation behavior makes it clear that the user is not entitled to place **DOES>** inside any control-flow structures.

A.6.1.1310 ELSE

Typical use: **: X ... test IF ... ELSE ... THEN ;**

A.6.1.1345 ENVIRONMENT?

In a Standard System that contains only the CORE word set, effective use of **ENVIRONMENT?** requires either its use within a definition, or the use of user-supplied auxiliary definitions. The CORE word set lacks both a direct method for collecting a string in interpretation state (**11.6.1.2165 S"** is in an optional word set) and also a means to test the returned flag in interpretation state (e.g. the optional **15.6.2.2532 [IF]**).

The combination of **6.1.1345 ENVIRONMENT?**, **11.6.1.2165 S**", **15.6.2.2532 [IF]**, **15.6.2.2531 [ELSE]**, and **15.6.2.2533 [THEN]** constitutes an effective suite of words for conditional compilation that works in interpretation state.

A.6.1.1360 EVALUATE

The Technical Committee is aware that this function is commonly spelled **EVAL**. However, there exist implementations that could suffer by defining the word as is done here. We also find **EVALUATE** to be more readable and explicit. There was some sentiment for calling this **INTERPRET**, but that too would have undesirable effects on existing code. The longer spelling was not deemed significant since this is not a word that should be used frequently in source code.

A.6.1.1380 EXIT

Typical use: `: X ... test IF ... EXIT THEN ... ;`

A.6.1.1550 FIND

One of the more difficult issues which the Committee took on was the problem of divorcing the specification of implementation mechanisms from the specification of the Forth language. Three basic implementation approaches can be quickly enumerated:

- 1) Threaded code mechanisms. These are the traditional approaches to implementing Forth, but other techniques may be used.
- 2) Subroutine threading with "macro-expansion" (code copying). Short routines, like the code for **DUP**, are copied into a definition rather than compiling a JSR reference.
- 3) Native coding with optimization. This may include stack optimization (replacing such phrases as **SWAP ROT +** with one or two machine instructions, for example), parallelization (the trend in the newer RISC chips is to have several functional subunits which can execute in parallel), and so on.

The initial requirement (inherited from Forth 83) that compilation addresses be compiled into the dictionary disallowed type 2 and type 3 implementations.

Type 3 mechanisms and optimizations of type 2 implementations were hampered by the explicit specification of immediacy or non-immediacy of all standard words. **POSTPONE** allowed de-specification of immediacy or non-immediacy for all but a few Forth words whose behavior must be **STATE**-independent.

One type 3 implementation, Charles Moore's cmForth, has both compiling and interpreting versions of many Forth words. At the present, this appears to be a common approach for type 3 implementations. The Committee felt that this implementation approach must be allowed. Consequently, it is possible that words without interpretation semantics can be found only during compilation, and other words may exist in two versions: a compiling version and an interpreting version. Hence the values returned by **FIND** may depend on **STATE**, and **'** and **[']** may be unable to find words without interpretation semantics.

A.6.1.1561 FM/MOD

By introducing the requirement for "floored" division, Forth 83 produced much controversy and concern on the part of those who preferred the more common practice followed in other languages of implementing division according to the behavior of the host CPU, which is most often symmetric (rounded toward zero). In attempting to find a compromise position, this Standard provides primitives for both common varieties, floored and symmetric (see **SM/REM**). **FM/MOD** is the floored version.

The Technical Committee considered providing two complete sets of explicitly named division operators, and declined to do so on the grounds that this would unduly enlarge and complicate the Standard. Instead, implementors may define the normal division words in terms of either **FM/MOD** or **SM/REM** providing they document their choice. People wishing to have explicitly named sets of operators are encouraged to do so. **FM/MOD** may be used, for example, to define:

```

: /_MOD ( n1 n2 -- n3 n4) >R S>D R> FM/MOD ;
: /_ ( n1 n2 -- n3) /_MOD SWAP DROP ;
: _MOD ( n1 n2 -- n3) /_MOD DROP ;
: */_MOD ( n1 n2 n3 -- n4 n5) >R M* R> FM/MOD ;
: */_ ( n1 n2 n3 -- n4 ) */_MOD SWAP DROP ;

```

A.6.1.1700 IF

Typical use:

```
: X ... test IF ... THEN ... ;
```

or

```
: X ... test IF ... ELSE ... THEN ... ;
```

A.6.1.1710 IMMEDIATE

Typical use: `: X ... ; IMMEDIATE`

A.6.1.1720 INVERT

The word **NOT** was originally provided in Forth as a flag operator to make control structures readable. Under its intended usage the following two definitions would produce identical results:

```

: ONE ( flag -- )
  IF ." true" ELSE ." false" THEN ;
: TWO ( flag -- )
  NOT IF ." false" ELSE ." true" THEN ;

```

This was common usage prior to the Forth-83 Standard which redefined **NOT** as a cell-wide one's-complement operation, functionally equivalent to the phrase `-1 XOR`. At the same time, the data type manipulated by this word was changed from a flag to a cell-wide collection of bits and the standard value for true was changed from "1" (rightmost bit only set) to "-1" (all bits set). As these definitions of **TRUE** and **NOT** were incompatible with their previous definitions, many Forth users continue to rely on the old definitions. Hence both versions are in common use.

Therefore, usage of **NOT** cannot be standardized at this time. The two traditional meanings of **NOT** – that of negating the sense of a flag and that of doing a one's complement operation – are made available by `O=` and **INVERT**, respectively.

A.6.1.1730 J

J may only be used with a nested `DO...LOOP`, `DO...+LOOP`, `?DO...LOOP`, or `?DO...+LOOP`, for example, in the form:

```
: X ... DO ... DO ... J ... LOOP ... +LOOP ... ;
```

A.6.1.1760 LEAVE

Note that **LEAVE** immediately exits the loop. No words following **LEAVE** within the loop will be executed. Typical use:

```
: X ... DO ... IF ... LEAVE THEN ... LOOP ... ;
```

A.6.1.1780 LITERAL

Typical use: `: X ... [x] LITERAL ... ;`

A.6.1.1800 LOOP

Typical use:

```
: X ... limit first DO ... LOOP ... ;
```

or

```
: X ... limit first ?DO ... LOOP ... ;
```

A.6.1.1810 M*

This word is a useful early step in calculation, going to extra precision conveniently. It has been in use since the Forth systems of the early 1970's.

A.6.1.1900 MOVE

CMOVE and **CMOVE>** are the primary move operators in Forth 83. They specify a behavior for moving that implies propagation if the move is suitably invoked. In some hardware, this specific behavior cannot be achieved using the best move instruction. Further, **CMOVE** and **CMOVE>** move characters; ANS Forth needs a move instruction capable of dealing with address units. Thus **MOVE** has been defined and added to the Core word set, and **CMOVE** and **CMOVE>** have been moved to the String word set.

A.6.1.2033 POSTPONE

Typical use:

```
: ENDF POSTPONE THEN ; IMMEDIATE
```

```
: X ... IF ... ENDF ... ;
```

POSTPONE replaces most of the functionality of **COMPILE** and **[COMPILE]**. **COMPILE** and **[COMPILE]** are used for the same purpose: postpone the compilation behavior of the next word in the parse area. **COMPILE** was designed to be applied to non-immediate words and **[COMPILE]** to immediate words. This burdens the programmer with needing to know which words in a system are immediate. Consequently, Forth standards have had to specify the immediacy or non-immediacy of all words covered by the Standard. This unnecessarily constrains implementors.

A second problem with **COMPILE** is that some programmers have come to expect and exploit a particular implementation, namely:

```
: COMPILE R> DUP @ , CELL+ >R ;
```

This implementation will not work on native code Forth systems. In a native code Forth using inline code expansion and peephole optimization, the size of the object code produced varies; this information is difficult to communicate to a "dumb" **COMPILE**. A "smart" (i.e., immediate) **COMPILE** would not have this problem, but this was forbidden in previous standards.

For these reasons, **COMPILE** has not been included in the Standard and **[COMPILE]** has been moved in favor of **POSTPONE**. Additional discussion can be found in Hayes, J.R., "Postpone", *Proceedings of the 1989 Rochester Forth Conference*.

A.6.1.2120 RECURSE

Typical use: `: X ... RECURSE ... ;`

This is Forth's recursion operator; in some implementations it is called **MYSELF**. The usual example is the coding of the factorial function.

```
: FACTORIAL ( +n1 -- +n2) DUP 2 > IF DUP 1-
  RECURSE * THEN ;
```

$n2 = n1(n1-1)(n1-2)...(2)(1)$, the product of $n1$ with all positive integers less than itself. While beloved of computer scientists, recursion makes unusually heavy use of both stacks and should therefore be used with caution. See alternate definition in **A.6.1.2140 REPEAT**.

A.6.1.2140 REPEAT

Typical use:

```

: FACTORIAL ( +n1 -- +n2 ) DUP
  BEGIN DUP 2 > WHILE
    1- SWAP OVER * SWAP
  REPEAT DROP ;

```

A.6.1.2165 S"

Typical use: `: X ... S" ccc" ... ;`

This word is found in many systems under the name " (quote). However, current practice is almost evenly divided on the use of ", with many systems using the execution semantics given here, while others return the address of a counted string. We attempt here to satisfy both camps by providing two words, **S"** and the Core Extension word **C"** so that users may have whichever behavior they expect with a simple renaming operation.

A.6.1.2214 SM/REM

See the previous discussion of division under **FM/MOD**. **SM/REM** is the symmetric-division primitive, which allows programs to define the following symmetric-division operators:

```

: /-REM ( n1 n2 -- n3 n4 ) >R S>D R> SM/REM ;
: /- ( n1 n2 -- n3 ) /-REM SWAP DROP ;
: -REM ( n1 n2 -- n3 ) /-REM DROP ;
: */-REM ( n1 n2 n3 -- n4 n5 ) >R M* R> SM/REM ;
: */- ( n1 n2 n3 -- n4 ) */-REM SWAP DROP ;

```

A.6.1.2216 SOURCE

SOURCE simplifies the process of directly accessing the input buffer by hiding the differences between its location for different input sources. This also gives implementors more flexibility in their implementation of buffering mechanisms for different input sources. The committee moved away from an input buffer specification consisting of a collection of individual variables, declaring **TIB** and **#TIB** obsolescent.

SOURCE in this form exists in F83, PolyFORTH, LMI's Forths and others. In conventional systems it is equivalent to the phrase

```
BLK @ IF BLK @ BLOCK 1024 ELSE TIB #TIB @ THEN.
```

A.6.1.2250 STATE

Although **EVALUATE**, **LOAD**, **INCLUDE-FILE**, and **INCLUDED** are not listed as words which alter **STATE**, the text interpreted by any one of these words could include one or more words which explicitly alter **STATE**. **EVALUATE**, **LOAD**, **INCLUDE-FILE**, and **INCLUDED** do not in themselves alter **STATE**.

STATE does not nest with text interpreter nesting. For example, the code sequence:

```
: FOO S" ]" EVALUATE ; FOO
```

will leave the system in compilation state. Similarly, after **LOADING** a block containing `]`, the system will be in compilation state.

Note that `]` does not affect the parse area and that the only effect that `:` has on the parse area is to parse a word. This entitles a program to use these words to set the state with known side-effects on the parse area. For example:

```
 : NOP : POSTPONE ; IMMEDIATE ;
  NOP ALIGN  NOP ALIGNED
```

Some non-ANS Forth compliant systems have `]` invoke a compiler loop in addition to setting **STATE**. Such a system would inappropriately attempt to compile the second use of **NOP**.

Also note that nothing in the Standard prevents a program from finding the execution tokens of `]` or `[` and using these to affect **STATE**. These facts suggest that implementations of `]` will do nothing but set **STATE** and a single interpreter/compiler loop will monitor **STATE**.

A.6.1.2270 THEN

Typical use:

```
 : X ... test IF ... THEN ... ;
```

or

```
 : X ... test IF ... ELSE ... THEN ... ;
```

A.6.1.2380 UNLOOP

Typical use:

```
 : X ...
   limit first DO
     ... test IF ... UNLOOP EXIT THEN ...
   LOOP ... ;
```

UNLOOP allows the use of **EXIT** within the context of `DO ... LOOP` and related do-loop constructs. **UNLOOP** as a function has been called **UNDO**. **UNLOOP** is more indicative of the action: nothing gets undone -- we simply stop doing it.

A.6.1.2390 UNTIL

Typical use: `: X ... BEGIN ... test UNTIL ... ;`

A.6.1.2410 VARIABLE

Typical use: `... VARIABLE XYZ ...`

A.6.1.2430 WHILE

Typical use: `: X ... BEGIN ... test WHILE ... REPEAT ... ;`

A.6.1.2450 WORD

Typical use: `char WORD ccc<char>`

A.6.1.2500 [

Typical use: `: X ... [4321] LITERAL ... ;`

A.6.1.2510 [']

Typical use: `: X ... ['] name ... ;`

See: A.6.1.1550 **FIND**.

A.6.1.2520 [CHAR]

Typical use: `: X ... [CHAR] ccc ... ;`

A.6.1.2540]

Typical use: `: X ... [1234] LITERAL ... ;`

A.6.2 Core extension words

The words in this collection fall into several categories:

- Words that are in common use but are deemed less essential than Core words (e.g., `0<>`);
- Words that are in common use but can be trivially defined from Core words (e.g., `FALSE`);
- Words that are primarily useful in narrowly defined types of applications or are in less frequent use (e.g., `PARSE`);
- Words that are being deprecated in favor of new words introduced to solve specific problems (e.g., `CONVERT`).

Because of the varied justifications for inclusion of these words, the Technical Committee does not encourage implementors to offer the complete collection, but to select those words deemed most valuable to their clientele.

A.6.2.0060 #TIB

The function of `#TIB` has been superseded by `SOURCE`.

A.6.2.0200 . (

Typical use: `.(ccc)`

A.6.2.0210 .R

In `.R`, “R” is short for RIGHT.

A.6.2.0340 2>R

Historically, `2>R` has been used to implement `DO`. Hence the order of parameters on the return stack.

The primary advantage of `2>R` is that it puts the top stack entry on the top of the return stack. For instance, a double-cell number may be transferred to the return stack and still have the most significant cell accessible on the top of the return stack.

A.6.2.0410 2R>

Note that `2R>` is not equivalent to `R> R>`. Instead, it mirrors the action of `2>R` (see **A.6.2.0340**).

A.6.2.0455 :NONAME

`:NONAME` allows a user to create an execution token with the semantics of a colon definition without an associated name. Previously, only `:` (colon) could create an execution token with these semantics. Thus, Forth code could only be compiled using the syntax of `:`, that is:

```
: NAME ... ;
```

`:NONAME` removes this constraint and places the Forth compiler in the hands of the programmer.

`:NONAME` can be used to create application-specific programming languages. One technique is to mix Forth code fragments with application-specific constructs. The application-specific constructs use

`:NONAME` to compile the Forth code and store the corresponding execution tokens in data structures.

The functionality of **:NONAME** can be built on any Forth system. For years, expert Forth programmers have exploited intimate knowledge of their systems to generate unnamed code fragments. Now, this function has been named and can be used in a portable program.

For example, **:NONAME** can be used to build a table of code fragments where indexing into the table allows executing a particular fragment. The declaration syntax of the table is:

```
:NONAME .. code for command 0 .. ; 0 CMD !
:NONAME .. code for command 1 .. ; 1 CMD !
...
:NONAME .. code for command 99 .. ; 99 CMD !

... 5 CMD @ EXECUTE ...
```

The definitions of the table building words are:

```
CREATE CMD-TABLE \ table for command execution tokens
100 CELLS ALLOT

: CMD ( n -- a-addr ) \ nth element address in table
CELLS CMD-TABLE + ;
```

As a further example, a defining word can be created to allow performance monitoring. In the example below, the number of times a word is executed is counted. **:** must first be renamed to allow the definition of the new **;**.

```
: DOCOLON ( -- )
\ Modify CREATED word to execute like a colon def
DOES> ( i*x a-addr -- j*x )
1 OVER +! \ count executions
CELL+ @ EXECUTE ; \ execute :NONAME definition
: OLD: : ; \ just an alias
OLD: : ( "name" -- a-addr xt colon-sys )
\ begins an execution-counting colon definition
CREATE HERE 0 , \ storage for execution counter
0 , \ storage for execution token
DOCOLON \ set run time for CREATED word
:NONAME ; \ begin unnamed colon definition
```

(Note the placement of **DOES>**: **DOES>** must modify the **CREATED** word and not the **:NONAME** definition, so **DOES>** must execute before **:NONAME**.)

```
OLD: ; ( a-addr xt colon-sys -- )
\ ends an execution-counting colon definition )
POSTPONE ; \ complete compilation of colon def
SWAP CELL+ ! ; \ save execution token
IMMEDIATE
```

The new **:** and **;** are used just like the standard ones to define words:

```
... : xxx ... ; ... xxx ...
```

Now however, these words may be “ticked” to retrieve the count (and execution token):

```
... ' xxx >BODY ? ...
```

A.6.2.0620 ?DO

Typical use: **: X ... limit first ?DO ... LOOP ... ;**

This word was added in response to many requests for a resolution of the difficulty introduced by Forth-83's **DO**, which on a 16-bit system will loop 65,535 times if given equal arguments. As this Standard also encourages 32-bit systems, this behavior can be intolerable. The Technical Committee considered applying these semantics to **DO**, but declined on the grounds that it might break existing code.

A.6.2.0700 **AGAIN**

Typical use: `: X ... BEGIN ... AGAIN ... ;`

Unless word-sequence has a way to terminate, this is an endless loop.

A.6.2.0855 **C"**

Typical use: `: X ... C" ccc" ... ;`

It is easy to convert counted strings to pointer/length but hard to do the opposite. **C"** is the only new word that uses the "address of counted string" stack representation. It is provided as an aid to porting existing programs to ANS Forth systems. It is relatively difficult to implement **C"** in terms of other standard words, considering its "compile string into the current definition" semantics.

Users of **C"** are encouraged to migrate their application code toward the consistent use of the preferred "*c-addr u*" stack representation with the alternate word **S"**. This may be accomplished by converting application words with counted string input arguments to use the preferred "*c-addr u*" representation, thus eliminating the need for **C"**.

See: A.3.1.3.4 Counted strings.

A.6.2.0873 **CASE**

Typical use:

```
: X ...
  CASE
    test1 OF ... ENDOF
    testn OF ... ENDOF
    ... ( default )
  ENDCASE ... ;
```

A.6.2.0945 **COMPILE,**

COMPILE, is the compilation equivalent of **EXECUTE**. In many cases, it is possible to compile a word by using **POSTPONE** without resorting to the use of **COMPILE,**. However, the use of **POSTPONE** requires that the name of the word must be known at compile time, whereas **COMPILE,** allows the word to be located at any time. It is sometime possible to use **EVALUATE** to compile a word whose name is not known until run time. This has two possible problems:

- **EVALUATE** is slower than **COMPILE,** because a dictionary search is required.
- The current search order affects the outcome of **EVALUATE**.

In traditional threaded-code implementations, compilation is performed by **,** (comma). This usage is not portable; it doesn't work for subroutine-threaded, native code, or relocatable implementations. Use of **COMPILE,** is portable.

In most systems it is possible to implement **COMPILE,** so it will generate code that is optimized to the same extent as code that is generated by the normal compilation process. However, in some implementations there are two different "tokens" corresponding to a particular definition name: the normal "execution token" that is used while interpreting or with **EXECUTE**, and another "compilation token" that is used while compiling. It is not always possible to obtain the compilation token from the execution token. In these implementations, **COMPILE,** might not generate code that is as efficient as normally compiled code.

A.6.2.0970 CONVERT

CONVERT may be defined as follows:

```
: CONVERT CHAR+ 65535 >NUMBER DROP ;
```

A.6.2.1342 ENDCASE

Typical use:

```
: X ...
CASE
  test1 OF ... ENDOF
  testn OF ... ENDOF
  ... ( default )
ENDCASE ... ;
```

A.6.2.1343 ENDOF

Typical use:

```
: X ...
CASE
  test1 OF ... ENDOF
  testn OF ... ENDOF
  ... ( default )
ENDCASE ... ;
```

A.6.2.1390 EXPECT

Specification of positive integer counts (+*n*) for **EXPECT** allows some implementors to continue their practice of using a zero or negative value as a flag to trigger special behavior. Insofar as such behavior is outside the Standard, Standard Programs cannot depend upon it, but the Technical Committee doesn't wish to preclude it unnecessarily. Since actual values are almost always small integers, no functionality is impaired by this restriction.

A.6.2.1850 MARKER

As dictionary implementations have gotten more elaborate and in some cases have used multiple address spaces, **FORGET** has become prohibitively difficult or impossible to implement on many Forth systems. **MARKER** greatly eases the problem by making it possible for the system to remember "landmark information" in advance that specifically marks the spots where the dictionary may at some future time have to be rearranged.

A.6.2.1950 OF

Typical use:

```
: X ...
CASE
  test1 OF ... ENDOF
  testn OF ... ENDOF
  ... ( default )
ENDCASE ... ;
```

A.6.2.2000 PAD

PAD has been available as scratch storage for strings since the earliest Forth implementations. It was brought to our attention that many programmers are reluctant to use **PAD**, fearing incompatibilities with

system uses. **PAD** is specifically intended as a programmer convenience, however, which is why we documented the fact that no standard words use it.

A.6.2.2008 **PARSE**

Typical use: *char* **PARSE** *ccc*<*char*>

The traditional Forth word for parsing is **WORD**. **PARSE** solves the following problems with **WORD**:

1. **WORD** always skips leading delimiters. This behavior is appropriate for use by the text interpreter, which looks for sequences of non-blank characters, but is inappropriate for use by words like **(**, **.(**, and **."**. Consider the following (flawed) definition of **.(** :

```
: .( [CHAR] ) WORD COUNT TYPE ; IMMEDIATE
```

This works fine when used in a line like:

```
.( HELLO) 5 .
```

but consider what happens if the user enters an empty string:

```
.( ) 5 .
```

The definition of **.(** shown above would treat the **)** as a leading delimiter, skip it, and continue consuming characters until it located another **)** that followed a non-**)** character, or until the parse area was empty. In the example shown, the **5 .** would be treated as part of the string to be printed.

With **PARSE**, we could write a correct definition of **.(** :

```
: .( [CHAR] ) PARSE TYPE ; IMMEDIATE
```

This definition avoids the “empty string” anomaly.

2. **WORD** returns its result as a counted string. This has four bad effects:
 - a) The characters accepted by **WORD** must be copied from the input buffer into a temporary buffer, in order to make room for the count character that must be at the beginning of the counted string. The copy step is inefficient, compared to **PARSE**, which leaves the string in the input buffer and doesn't need to copy it anywhere.
 - b) **WORD** must be careful not to store too many characters into the temporary buffer, thus overwriting something beyond the end of the buffer. This adds to the overhead of the copy step. (**WORD** may have to scan a lot of characters before finding the trailing delimiter.)
 - c) The count character limits the length of the string returned by **WORD** to 255 characters (longer strings can easily be stored in blocks!). This limitation does not exist for **PARSE**.
 - d) The temporary buffer is typically overwritten by the next use of **WORD**. This introduces a temporal dependency; the value returned by **WORD** is only valid for a limited duration. **PARSE** has a temporal dependency, too, related to the lifetime of the input buffer, but that is less severe in most cases than **WORD**'s temporal dependency.

The behavior of **WORD** with respect to skipping leading delimiters is useful for parsing blank-delimited names. Many system implementations include an additional word for this purpose, similar to **PARSE** with respect to the “*c-addr u*” return value, but without an explicit delimiter argument (the delimiter set is implicitly “white space”), and which does skip leading delimiters. A common description for this word is:

```
PARSE-WORD (“<spaces>name” -- c-addr u)
```

Skip leading spaces and parse *name* delimited by a space. *c-addr* is the address within the input buffer and *u* is the length of the selected string. If the parse area is empty, the resulting string has a zero length.

If both **PARSE** and **PARSE-WORD** are present, the need for **WORD** is largely eliminated.

A.6.2.2030 **PICK**

0 PICK is equivalent to **DUP** and **1 PICK** is equivalent to **OVER**.

A.6.2.2040 QUERY

The function of **QUERY** may be performed with **ACCEPT** and **EVALUATE**.

A.6.2.2125 REFILL

This word is a useful generalization of **QUERY**. Re-defining **QUERY** to meet this specification would have broken existing code. **REFILL** is designed to behave reasonably for all possible input sources. If the input source is coming from the user, as with **QUERY**, **REFILL** could still return a false value if, for instance, a communication channel closes so that the system knows that no more input will be available.

A.6.2.2150 ROLL

2 ROLL is equivalent to **ROT**, **1 ROLL** is equivalent to **SWAP** and **0 ROLL** is a null operation.

A.6.2.2182 SAVE-INPUT

SAVE-INPUT and **RESTORE-INPUT** allow the same degree of input source repositioning within a text file as is available with **BLOCK** input. **SAVE-INPUT** and **RESTORE-INPUT** "hide the details" of the operations necessary to accomplish this repositioning, and are used the same way with all input sources. This makes it easier for programs to reposition the input source, because they do not have to inspect several variables and take different action depending on the values of those variables.

SAVE-INPUT and **RESTORE-INPUT** are intended for repositioning within a single input source; for example, the following scenario is NOT allowed for a Standard Program:

```
: XX
  SAVE-INPUT CREATE
  S" RESTORE-INPUT" EVALUATE
  ABORT" couldn't restore input" ;
```

This is incorrect because, at the time **RESTORE-INPUT** is executed, the input source is the string via **EVALUATE**, which is not the same input source that was in effect when **SAVE-INPUT** was executed.

The following code is allowed:

```
: XX
  SAVE-INPUT CREATE
  S" .( Hello)" EVALUATE
  RESTORE-INPUT ABORT" couldn't restore input" ;
```

After **EVALUATE** returns, the input source specification is restored to its previous state, thus **SAVE-INPUT** and **RESTORE-INPUT** are called with the same input source in effect.

In the above examples, the **EVALUATE** phrase could have been replaced by a phrase involving **INCLUDE-FILE** and the same rules would apply.

The Standard does not specify what happens if a program violates the above rules. A Standard System might check for the violation and return an exception indication from **RESTORE-INPUT**, or it might fail in an unpredictable way.

The return value from **RESTORE-INPUT** is primarily intended to report the case where the program attempts to restore the position of an input source whose position cannot be restored. The keyboard might be such an input source.

Nesting of **SAVE-INPUT** and **RESTORE-INPUT** is allowed. For example, the following situation works as expected:

```
: XX
  SAVE-INPUT
  S" f1" INCLUDED
```

```

\ The file "f1" includes:
\   ... SAVE-INPUT ... RESTORE-INPUT ...
\ End of file "f1"
RESTORE-INPUT ABORT" couldn't restore input" ;

```

In principle, **RESTORE-INPUT** could be implemented to “always fail”, e.g.:

```

: RESTORE-INPUT ( x1 ... xn n -- flag )
0 ?DO DROP LOOP TRUE ;

```

Such an implementation would not be useful in most cases. It would be preferable for a system to leave **SAVE-INPUT** and **RESTORE-INPUT** undefined, rather than to create a useless implementation. In the absence of the words, the application programmer could choose whether or not to create “dummy” implementations or to work-around the problem in some other way.

Examples of how an implementation might use the return values from **SAVE-INPUT** to accomplish the save/restore function:

Input Source	possible stack values
block	>IN @ BLK @ 2
EVALUATE	>IN @ 1
keyboard	>IN @ 1
text file	>IN @ lo-pos hi-pos 3

These are examples only; a Standard Program may not assume any particular meaning for the individual stack items returned by **SAVE-INPUT**.

A.6.2.2290 TIB

The function of **TIB** has been superseded by **SOURCE**.

A.6.2.2295 TO

Historically, some implementations of **TO** have not explicitly parsed. Instead, they set a mode flag that is tested by the subsequent execution of name. ANS Forth explicitly requires that **TO** must parse, so that **TO**'s effect on the parse area will be predictable when it is used at the end of a line.

Typical use: **x TO name**

A.6.2.2298 TRUE

TRUE is equivalent to the phrase **0 0=**.

A.6.2.2405 VALUE

Typical use: **0 VALUE DATA**

```

: EXCHANGE ( n1 -- n2 ) DATA SWAP TO DATA ;

```

EXCHANGE leaves n_1 in **DATA** and returns the prior value n_2 .

A.6.2.2440 WITHIN

We describe **WITHIN** without mentioning circular number spaces (an undefined term) or providing the code. Here is a number line with the overflow point (o) at the far right and the underflow point (u) at the far left:

u-----o

There are two cases to consider: either the $n_2/u_2..n_3/u_3$ range straddles the overflow/underflow points or it does not. Lets examine the non-straddle case first:

u-----[.....)-----o

The [denotes n_2/u_2 , the) denotes n_3/u_3 , and the dots and [are numbers **WITHIN** the range. n_3/u_3 is greater than n_2/u_2 , so the following tests will determine if n_1/u_1 is **WITHIN** n_2/u_2 and n_3/u_3 :

$$n_2/u_2 \leq n_1/u_1 \text{ and } n_1/u_1 < n_3/u_3.$$

In the case where the comparison range straddles the overflow/underflow points:

u.....)-----[.....o

n_3/u_3 is less than n_2/u_2 and the following tests will determine if n_1/u_1 is **WITHIN** n_2/u_2 and n_3/u_3 :

$$n_2/u_2 \leq n_1/u_1 \text{ or } n_1/u_1 < n_3/u_3.$$

WITHIN must work for both signed and unsigned arguments. One obvious implementation does not work:

```

: WITHIN ( test low high -- flag )
  >R OVER < 0= ( test flag )
  SWAP R> < AND ;

```

Assume two's-complement arithmetic on a 16-bit machine, and consider the following test:

```

33000 32000 34000 WITHIN

```

The above implementation returns *false* for that test, even though the unsigned number 33000 is clearly within the range $\{32000..34000\}$.

The problem is that, in the incorrect implementation, the signed comparison $<$ gives the wrong answer when 32000 is compared to 33000, because when those numbers are treated as signed numbers, 33000 is treated as negative 32536, while 32000 remains positive.

Replacing $<$ with $U<$ in the above implementation makes it work with unsigned numbers, but causes problems with certain signed number ranges; in particular, the test:

```

1 -5 5 WITHIN

```

would give an incorrect answer.

For two's-complement machines that ignore arithmetic overflow (most machines), the following implementation works in all cases:

```

: WITHIN ( test low high -- ) OVER - >R - R> U< ;

```

A.6.2.2530 [COMPILE]

Typical use: `: name2 ... [COMPILE] name1 ... ; IMMEDIATE`

A.6.2.2535 \

Typical use: `5 CONSTANT THAT \ THIS IS A COMMENT ABOUT THAT`

A.7 The Block word set

Early Forth systems ran stand-alone, with no host OS. Blocks of 1024 bytes were designed as a convenient unit of disk, and most native Forth systems still use them. It is relatively easy to write a native disk driver that maps head/track/sector addresses to block numbers. Such disk drivers are extremely fast in comparison with conventional file-oriented operating systems, and security is high because there is no reliance on a disk map.

Today many Forth implementations run under host operating systems, because the compatibility they offer the user outweighs the performance overhead. Many people who use such systems prefer using host OS files only; however, people who use both native and non-native Forths need a compatible way of accessing disk. The Block Word set includes the most common words for accessing program source and data on disk.

In order to guarantee that Standard Programs that need access to mass storage have a mechanism appropriate for both native and non-native implementations, ANS Forth requires that the Block word set be available if any mass storage facilities are provided. On non-native implementations, blocks normally reside in host OS files.

A.7.2 Additional terms

block

Many Forth systems use blocks to contain program source. Conventionally such blocks are formatted for editing as 16 lines of 64 characters. Source blocks are often referred to as “screens”.

A.7.6 Glossary

A.7.6.2.2190 SCR

SCR is short for screen.

A.8 The Double-Number word set

Forth systems on 8-bit and 16-bit processors often find it necessary to deal with double-length numbers. But many Forths on small embedded systems do not, and many users of Forth on systems with a cell size of 32-bits or more find that the necessity for double-length numbers is much diminished. Therefore, we have factored the words that manipulate double-length entities into this optional word set.

Please note that the naming convention used in this word set conveys some important information:

1. Words whose names are of the form 2xxx deal with cell pairs, where the relationship between the cells is unspecified. They may be two-vectors, double-length numbers, or any pair of cells that it is convenient to manipulate together.
2. Words with names of the form Dxxx deal specifically with double-length integers.
3. Words with names of the form Mxxx deal with some combination of single and double integers. The order in which these appear on the stack is determined by long-standing common practice.

Refer to A.3.1 for a discussion of data types in Forth.

A.8.6 Glossary

A.8.6.1.0360 2CONSTANT

Typical use: **x1 x2 2CONSTANT name**

A.8.6.1.0390 2LITERAL

Typical use: **: X ... [x1 x2] 2LITERAL ... ;**

A.8.6.1.0440 2VARIABLE

Typical use: **2VARIABLE** *name*

A.8.6.1.1070 D.R

In **D.R**, the “R” is short for RIGHT.

A.8.6.1.1090 D2*

See **A.6.1.0320 2*** for applicable discussion.

A.8.6.1.1100 D2/

See **A.6.1.0330 2/** for applicable discussion.

A.8.6.1.1140 D>S

There exist number representations, e.g., the sign-magnitude representation, where reduction from double- to single-precision cannot simply be done with **DROP**. This word, equivalent to **DROP** on two’s complement systems, desensitizes application code to number representation and facilitates portability.

A.8.6.1.1820 M*/

M*/ was once described by Chuck Moore as the most useful arithmetic operator in Forth. It is the main workhorse in most computations involving double-cell numbers. Note that some systems allow signed divisors. This can cost a lot in performance on some CPUs. The requirement for a positive divisor has not proven to be a problem.

A.8.6.1.1830 M+

M+ is the classical method for integrating.

A.9 The Exception word set

CATCH and **THROW** provide a reliable mechanism for handling exceptions, without having to propagate exception flags through multiple levels of word nesting. It is similar in spirit to the “non-local return” mechanisms of many other languages, such as C’s **setjmp()** and **longjmp()**, and LISP’s **CATCH** and **THROW**. In the Forth context, **THROW** may be described as a “multi-level **EXIT**”, with **CATCH** marking a location to which a **THROW** may return.

Several similar Forth “multi-level **EXIT**” exception-handling schemes have been described and used in past years. It is not possible to implement such a scheme using only standard words (other than **CATCH** and **THROW**), because there is no portable way to “unwind” the return stack to a predetermined place.

THROW also provides a convenient implementation technique for the standard words **ABORT** and **ABORT"**, allowing an application to define, through the use of **CATCH**, the behavior in the event of a system **ABORT**.

This sample implementation of **CATCH** and **THROW** uses the non-standard words described below. They or their equivalents are available in many systems. Other implementation strategies, including directly saving the value of **DEPTH**, are possible if such words are not available.

SP@ (-- *addr*) returns the address corresponding to the top of data stack.

SP! (*addr* --) sets the stack pointer to *addr*, thus restoring the stack depth to the same depth that existed just before *addr* was acquired by executing **SP@**.

RP@ (-- *addr*) returns the address corresponding to the top of return stack.

RP! (*addr* --) sets the return stack pointer to *addr*, thus restoring the return stack depth to the same depth that existed just before *addr* was acquired by executing **RP@**.

```
VARIABLE HANDLER 0 HANDLER ! \ last exception handler
: CATCH ( xt -- exception# | 0 ) \ return addr on stack
  SP@ >R      ( xt ) \ save data stack pointer
  HANDLER @ >R ( xt ) \ and previous handler
  RP@ HANDLER ! ( xt ) \ set current handler
  EXECUTE      ( ) \ execute returns if no THROW
  R> HANDLER ! ( ) \ restore previous handler
  R> DROP      ( ) \ discard saved stack ptr
  0 ;          ( 0 ) \ normal completion

: THROW ( ??? exception# -- ??? exception# )
  ?DUP IF      ( exc# ) \ 0 THROW is no-op
    HANDLER @ RP ! ( exc# ) \ restore prev return stack
  R> HANDLER ! ( exc# ) \ restore prev handler
  R> SWAP >R    ( saved-sp ) \ exc# on return stack
  SP! DROP R>  ( exc# ) \ restore stack
  \ Return to the caller of CATCH because return
  \ stack is restored to the state that existed
  \ when CATCH began execution
  THEN ;
```

In a multi-tasking system, the **HANDLER** variable should be in the per-task variable area (i.e., a user variable).

This sample implementation does not explicitly handle the case in which **CATCH** has never been called (i.e., the **ABORT** behavior). One solution is to add the following code after the **IF** in **THROW**:

```
HANDLER @ 0= IF ( empty the stack ) QUIT THEN
```

Another solution is to execute **CATCH** within **QUIT**, so that there is always an “exception handler of last resort” present. For example:

```
: QUIT
  ( empty the return stack and )
  ( set the input source to the user input device )
  POSTPONE [
  BEGIN
    REFILL
  WHILE
    ['] INTERPRET CATCH
  CASE
    0 OF STATE @ 0= IF ." OK" THEN CR ENDOF
    -1 OF ( Aborted) ENDOF
    -2 OF ( display message from ABORT" ) ENDOF
    ( default ) DUP ." Exception # " .
  ENDCASE
  REPEAT BYE ;
```

This example assumes the existence of a system-implementation word **INTERPRET** that embodies the text interpreter semantics described in **3.4 The Forth text interpreter**. Note that this implementation of **QUIT** automatically handles the emptying of the stack and return stack, due to **THROW**'s inherent restoration of the data and return stacks. Given this definition of **QUIT**, it's easy to define:

```
: ABORT -1 THROW ;
```

In systems with other stacks in addition to the data and return stacks, the implementation of **CATCH** and **THROW** must save and restore those stack pointers as well. Such an “extended version” can be built on top

of this basic implementation. For example, with another stack pointer accessed with **FP@** and **FP!** only **CATCH** needs to be redefined:

```
: CATCH ( xt -- exception# | 0 )
  FP@ >R CATCH R> OVER IF FP! ELSE DROP THEN ;
```

No change to **THROW** is necessary in this case. Note that, as with all redefinitions, the redefined version of **CATCH** will only be available to definitions compiled after the redefinition of **CATCH**.

CATCH and **THROW** provide a convenient way for an implementation to “clean up” the state of open files if an exception occurs during the text interpretation of a file with **INCLUDE-FILE**. The implementation of **INCLUDE-FILE** may guard (with **CATCH**) the word that performs the text interpretation, and if **CATCH** returns an exception code, the file may be closed and the exception re**THROW**n so that the files being included at an outer nesting level may be closed also. Note that the Standard allows, but does not require, **INCLUDE-FILE** to close its open files if an exception occurs. However, it does require **INCLUDE-FILE** to unnest the input source specification if an exception is **THROW**n.

A.9.3 Additional usage requirements

One important use of an exception handler is to maintain program control under many conditions which **ABORT**. This is practicable only if a range of codes is reserved. Note that an application may overload many standard words in such a way as to **THROW** ambiguous conditions not normally **THROW**n by a particular system.

A.9.3.6 Exception handling

The method of accomplishing this coupling is implementation dependent. For example, **LOAD** could “know” about **CATCH** and **THROW** (by using **CATCH** itself, for example), or **CATCH** and **THROW** could “know” about **LOAD** (by maintaining input source nesting information in a data structure known to **THROW**, for example). Under these circumstances it is not possible for a Standard Program to define words such as **LOAD** in a completely portable way.

A.9.6 Glossary

A.9.6.1.2275 THROW

If **THROW** is executed with a non zero argument, the effect is as if the corresponding **CATCH** had returned it. In that case, the stack depth is the same as it was just before **CATCH** began execution. The values of the i^*x stack arguments could have been modified arbitrarily during the execution of xt . In general, nothing useful may be done with those stack items, but since their number is known (because the stack depth is deterministic), the application may **DROP** them to return to a predictable stack state.

Typical use:

```
: could-fail ( -- char )
  KEY DUP [CHAR] Q = IF 1 THROW THEN ;
: do-it ( a b -- c) 2DROP could-fail ;
: try-it ( --)
  1 2 ['] do-it CATCH IF
    ( x1 x2 ) 2DROP ." There was an exception" CR
  ELSE ." The character was " EMIT CR
  THEN ;
: retry-it ( -- )
  BEGIN 1 2 ['] do-it CATCH WHILE
    ( x1 x2) 2DROP ." Exception, keep trying" CR
  REPEAT ( char )
    ." The character was " EMIT CR ;
```

A.10 The Facility word set

A.10.6 Glossary

A.10.6.1.0742 AT-XY

Most implementors supply a method of positioning a cursor on a CRT screen, but there is great variance in names and stack arguments. This version is supported by at least one major vendor.

A.10.6.1.1755 KEY?

The Technical Committee has gone around several times on the stack effects. Whatever is decided will violate somebody's practice and penalize some machine. This way doesn't interfere with type-ahead on some systems, while requiring the implementation of a single-character buffer on machines where polling the keyboard inevitably results in the destruction of the character.

Use of **KEY** or **KEY?** indicates that the application does not wish to bother with non-character events, so they are discarded, in anticipation of eventually receiving a valid character. Applications wishing to handle non-character events must use **EKEY** and **EKEY?**. It is possible to mix uses of **KEY?** / **KEY** and **EKEY?** / **EKEY** within a single application, but the application must use **KEY?** and **KEY** only when it wishes to discard non-character events until a valid character is received.

A.10.6.2.1305 EKEY

EKEY provides a standard word to access a system-dependent set of "raw" keyboard events, including events corresponding to members of the standard character set, events corresponding to other members of the implementation-defined character set, and keystrokes that do not correspond to members of the character set.

EKEY assumes no particular numerical correspondence between particular event code values and the values representing standard characters. On some systems, this may allow two separate keys that correspond to the same standard character to be distinguished from one another.

In systems that combine both keyboard and mouse events into a single "event stream", the single number returned by **EKEY** may be inadequate to represent the full range of input possibilities. In such systems, a single "event record" may include a time stamp, the x,y coordinates of the mouse position, the keyboard state, and the state of the mouse buttons. In such systems, it might be appropriate for **EKEY** to return the address of an "event record" from which the other information could be extracted.

Also, consider a hypothetical Forth system running under MS-DOS on a PC-compatible computer. Assume that the implementation-defined character set is the "normal" 8-bit PC character set. In that character set, the codes from 0 to 127 correspond to ASCII characters. The codes from 128 to 255 represent characters from various non-English languages, mathematical symbols, and some graphical symbols used for line drawing. In addition to those characters, the keyboard can generate various other "scan codes", representing such non-character events as arrow keys and function keys.

There may be multiple keys, with different scan codes, corresponding to the same standard character. For example, the character representing the number "1" often appears both in the row of number keys above the alphabetic keys, and also in the separate numeric keypad.

When a program asks the MS-DOS operating system for a keyboard event, it receives either a single non-zero byte, representing a character, or a zero byte followed by a "scan code" byte, representing a non-character keyboard event (e.g., a function key).

EKEY represents each keyboard event as a single number, rather than as a sequence of numbers. For the system described above, the following would be a reasonable implementation of **EKEY** and related words:

The **MAX-CHAR** environmental query would return 255.

Assume the existence of a word **DOS-KEY** (*-- char*) which executes the MS-DOS “Direct STDIN Input” system call (Interrupt 21h, Function 07h) and a word **DOS-KEY?** (*-- flag*) which executes the MS-DOS “Check STDIN Status” system call (Interrupt 21h, Function 0Bh).

```

: EKEY? ( -- flag ) DOS-KEY? 0<> ;
: EKEY ( -- u )
  DOS-KEY ?DUP 0= IF
    DOS-KEY 256 +
  THEN ;
: EKEY>CHAR ( u -- u false | char true )
  DUP 256 >= IF      ( u )
    DUP 259 = IF      \ 259 is Ctrl-@ (ASCII NUL)
      DROP 0 TRUE EXIT \ so replace with character
    THEN FALSE EXIT   \ otherwise extended character
  THEN TRUE ;         \ normal extended ASCII char.
VARIABLE PENDING-CHAR -1 PENDING-CHAR !
: KEY? ( -- flag )
  PENDING-CHAR @ 0>= IF TRUE EXIT THEN
  BEGIN EKEY? WHILE
    EKEY EKEY>CHAR IF
      PENDING-CHAR ! TRUE EXIT
    THEN DROP
  REPEAT FALSE ;
: KEY ( -- char )
  PENDING-CHAR @ 0>= IF
    PENDING-CHAR @ -1 PENDING-CHAR ! EXIT
  THEN
  BEGIN EKEY EKEY>CHAR 0= WHILE
    DROP
  REPEAT ;

```

This is a full-featured implementation, providing the application program with an easy way to either handle non-character events (with **EKEY**), or to ignore them and to only consider “real” characters (with **KEY**).

Note that **EKEY** maps scan codes from 0 to 255 into numbers from 256 to 511. **EKEY** maps the number 259, representing the keyboard combination Ctrl-Shift-@, to the character whose numerical value is 0 (ASCII NUL). Many ASCII keyboards generate ASCII NUL for Ctrl-Shift-@, so we use that key combination for ASCII NUL (which is otherwise unavailable from MS-DOS, because the zero byte signifies that another scan-code byte follows).

One consequence of using the “Direct STDIN Input” system call (function 7) instead of the “STDIN Input” system call (function 8) is that the normal DOS “Ctrl-C interrupt” behavior is disabled when the system is waiting for input (Ctrl-C would still cause an interrupt while characters are being output). On the other hand, if the “STDIN Input” system call (function 8) were used to implement **EKEY**, Ctrl-C interrupts would be enabled, but Ctrl-Shift-@ would also cause an interrupt, because the operating system would treat the second byte of the 0,3 sequence as a Ctrl-C, even though the 3 is really a scan code and not a character. One “best of both worlds” solution is to use function 8 for the first byte received by **EKEY**, and function 7 for the scan code byte. For example:

```

: EKEY ( -- u )
  DOS-KEY-FUNCTION-8 ?DUP 0= IF
    DOS-KEY-FUNCTION-7 DUP 3 = IF
      DROP 0 ELSE 256 +

```

THEN
THEN ;

Of course, if the Forth implementor chooses to pass Ctrl-C through to the program, without using it for its usual interrupt function, then DOS function 7 is appropriate in both cases (and some additional care must be taken to prevent a typed-ahead Ctrl-C from interrupting the Forth system during output operations).

A Forth system might also choose a simpler implementation of **KEY**, without implementing **EKEY**, as follows:

```
: KEY ( -- char ) DOS-KEY ;
: KEY? ( -- flag ) DOS-KEY? 0<> ;
```

The disadvantages of the simpler version are:

- a) An application program that uses **KEY**, expecting to receive only valid characters, might receive a sequence of bytes (e.g., a zero byte followed by a byte with the same numerical value as the letter "A") that appears to contain a valid character, even though the user pressed a key (e.g., function key 4) that does not correspond to any valid character.
- b) An application program that wishes to handle non-character events will have to execute **KEY** twice if it returns zero the first time. This might appear to be a reasonable and easy thing to do. However, such code is not portable to other systems that do not use a zero byte as an "escape" code. Using the **EKEY** approach, the algorithm for handling keyboard events can be the same for all systems; the system dependencies can be reduced to a table or set of constants listing the system-dependent key codes used to access particular application functions. Without **EKEY**, the algorithm, not just the table, is likely to be system dependent.

Another approach to **EKEY** on MS-DOS is to use the BIOS "Read Keyboard Status" function (Interrupt 16h, Function 01h) or the related "Check Keyboard" function (Interrupt 16h, Function 11h). The advantage of this function is that it allows the program to distinguish between different keys that correspond to the same character (e.g. the two "1" keys). The disadvantage is that the BIOS keyboard functions read only the keyboard. They cannot be "redirected" to another "standard input" source, as can the DOS STDIN Input functions.

A.10.6.2.1306 EKEY>CHAR

EKEY>CHAR translates a keyboard event into the corresponding member of the character set, if such a correspondence exists for that event.

It is possible that several different keyboard events may correspond to the same character, and other keyboard events may correspond to no character.

A.10.6.2.1325 EMIT?

An indefinite delay is a device related condition, such as printer off-line, that requires operator intervention before the device will accept new data.

A.10.6.2.1905 MS

Although their frequencies vary, every system has a clock. Since many programs need to time intervals, this word is offered. Use of milliseconds as an internal unit of time is a practical "least common denominator" external unit. It is assumed implementors will use "clock ticks" (whatever size they are) as an internal unit and convert as appropriate.

A.10.6.2.2292 TIME&DATE

Most systems have a real-time clock/calendar. This word gives portable access to it.

A.11 The File-Access word set

Many Forth systems support access to a host file system, and many of these support interpretation of Forth from source text files. The Forth-83 Standard did not address host OS files. Nevertheless, a degree of similarity exists among modern implementations.

For example, files must be opened and closed, created and deleted. Forth file-system implementations differ mostly in the treatment and disposition of the exception codes, and in the format of the file-identification strings. The underlying mechanism for creating file-control blocks might or might not be visible. We have chosen to keep it invisible.

Files must also be read and written. Text files, if supported, must be read and written one line at a time. Interpretation of text files implies that they are somehow integrated into the text interpreter input mechanism. These and other requirements have shaped the file-access extensions word set.

Most of the existing implementations studied use simple English words for common host file functions: **OPEN**, **CLOSE**, **READ**, etc. Although we would have preferred to do likewise, there were so many minor variations in implementation of these words that adopting any particular meaning would have broken much existing code. We have used names with a suffix **-FILE** for most of these words. We encourage implementors to conform their single-word primitives to the ANS behaviors, and hope that if this is done on a widespread basis we can adopt better definition names in a future standard.

Specific rationales for members of this word set follow.

A.11.3 Additional usage requirements

A.11.3.2 Blocks in files

Many systems reuse file identifiers; when a file is closed, a subsequently opened file may be given the same identifier. If the original file has blocks still in block buffers, these will be incorrectly associated with the newly opened file with disastrous results. The block buffer system must be flushed to avoid this.

A.11.6 Glossary

A.11.6.1.0765 **BIN**

Some operating systems require that files be opened in a different mode to access their contents as an unstructured stream of binary data rather than as a sequence of lines.

The arguments to **READ-FILE** and **WRITE-FILE** are arrays of character storage elements, each element consisting of at least 8 bits. The Technical Committee intends that, in **BIN** mode, the contents of these storage elements can be written to a file and later read back without alteration. The Technical Committee has declined to address issues regarding the impact of “wide” characters on the File and Block word sets.

A.11.6.1.1010 **CREATE-FILE**

Typical use:

```
: X ... S" TEST.FTH" R/W CREATE-FILE
  ABORT" CREATE-FILE FAILED" ... ;
```

A.11.6.1.1717 **INCLUDE-FILE**

Here are two implementation alternatives for saving the input source specification in the presence of text file input:

- 1) Save the file position (as returned by **FILE-POSITION**) of the beginning of the line being interpreted. To restore the input source specification, seek to that position and re-read the line into the input buffer.

2) Allocate a separate line buffer for each active text input file, using that buffer as the input buffer. This method avoids the “seek and reread” step, and allows the use of “pseudo-files” such as pipes and other sequential-access-only communication channels.

A.11.6.1.1718 INCLUDED

Typical use: ... **S** " filename" INCLUDED ...

A.11.6.1.1970 OPEN-FILE

Typical use:

```
: X ... S" TEST.FTH" R/W OPEN-FILE
  ABORT" OPEN-FILE FAILED" ... ;
```

A.11.6.1.2080 READ-FILE

A typical sequential file-processing algorithm might look like:

```
BEGIN          ( )
  . . . READ-FILE THROW ( length )
  ?DUP WHILE   ( length )
  . . .       ( )
  REPEAT      ( )
```

In this example, **THROW** is used to handle (unexpected) exception conditions, which are reported as non-zero values of the *ior* return value from **READ-FILE**. End-of-file is reported as a zero value of the “length” return value.

A.11.6.1.2090 READ-LINE

Implementations are allowed to store the line terminator in the memory buffer in order to allow the use of line reading functions provided by host operating systems, some of which store the terminator. Without this provision, a temporary buffer might be needed. The two-character limitation is sufficient for the vast majority of existing operating systems. Implementations on host operating systems whose line terminator sequence is longer than two characters may have to take special action to prevent the storage of more than two terminator characters.

Standard Programs may not depend on the presence of any such terminator sequence in the buffer.

A typical line-oriented sequential file-processing algorithm might look like:

```
BEGIN          ( )
  . . . READ-LINE THROW ( length not-eof-flag )
  WHILE        ( length )
  . . .       ( )
  REPEAT DROP  ( )
```

In this example, **THROW** is used to handle (unexpected) I/O exception condition, which are reported as non-zero values of the “*ior*” return value from **READ-LINE**.

READ-LINE needs a separate end-of-file flag because empty (zero-length) lines are a routine occurrence, so a zero-length line cannot be used to signify end-of-file.

A.11.6.1.2165 S"

Typical use: ... **S** " ccc" ...

The interpretation semantics for **S"** are intended to provide a simple mechanism for entering a string in the interpretation state. Since an implementation may choose to provide only one buffer for interpreted strings, an interpreted string is subject to being overwritten by the next execution of **S"** in interpretation

state. It is intended that no standard words other than **S"** should in themselves cause the interpreted string to be overwritten. However, since words such as **EVALUATE**, **LOAD**, **INCLUDE-FILE** and **INCLUDED** can result in the interpretation of arbitrary text, possibly including instances of **S"**, the interpreted string may be invalidated by some uses of these words.

When the possibility of overwriting a string can arise, it is prudent to copy the string to a "safe" buffer allocated by the application.

Programs wishing to parse in the fashion of **S"** are advised to use **PARSE** or **WORD COUNT** instead of **S"**, preventing the overwriting of the interpreted string buffer.

A.12 The Floating-Point word set

The Technical Committee has considered many proposals dealing with the inclusion and makeup of the Floating-Point Word Sets in ANS Forth. Although it has been argued that ANS Forth should not address floating-point arithmetic and numerous Forth applications do not need floating-point, there are a growing number of important Forth applications from spread sheets to scientific computations that require the use of floating-point arithmetic. Initially the Technical Committee adopted proposals that made the Forth Vendors Group Floating-Point Standard, first published in 1984, the framework for inclusion of Floating-Point in ANS Forth. There is substantial common practice and experience with the Forth Vendors Group Floating-Point Standard. Subsequently the Technical Committee adopted proposals that placed the basic floating-point arithmetic, stack and support words in the Floating-Point word set and the floating-point transcendental functions in the Floating-Point Extensions word set. The Technical Committee also adopted proposals that:

- Changed names for clarity and consistency; e.g., **REALS** to **FLOATS**, and **REAL+** to **FLOAT+** .
- Removed words; e.g., **FPICK** .
- Added words for completeness and increased functionality; e.g., **FSINCOS**, **F~**, **DF@**, **DF!**, **SF@** and **SF!**

Several issues concerning the Floating-Point word set were resolved by consensus in the Technical Committee:

Floating-point stack: By default the floating-point stack is separate from the data and return stacks; however, an implementation may keep floating-point numbers on the data stack. A program can determine whether floating-point numbers are kept on the data stack by passing the string **FLOATING-STACK** to **ENVIRONMENT?** It is the experience of several members of the Technical Committee that with proper coding practices it is possible to write floating-point code that will run identically on systems with a separate floating-point stack and with floating-point numbers kept on the data stack.

Floating-point input: The current base must be **DECIMAL**. Floating-point input is not allowed in an arbitrary base. All floating-point numbers to be interpreted by an ANS Forth system must contain the exponent indicator "E" (see **12.3.7 Text interpreter input number conversion**). Consensus in the Technical Committee deemed this form of floating-point input to be in more common use than the alternative that would have a floating-point input mode that would allow numbers with embedded decimal points to be treated as floating-point numbers.

Floating-point representation: Although the format and precision of the significand and the format and range of the exponent of a floating-point number are implementation defined in ANS Forth, the Floating-Point Extensions word set contains the words **DF@**, **SF@**, **DF!**, and **SF!** for fetching and storing double- and single-precision IEEE floating-point-format numbers to memory. The IEEE floating-point format is commonly used by numeric math co-processors and for exchange of floating-point data between programs and systems.

A.12.3 Additional usage requirements

A.12.3.5 Address alignment

In defining custom floating-point data structures, be aware that **CREATE** doesn't necessarily leave the data space pointer aligned for various floating-point data types. Programs may comply with the requirement for the various kinds of floating-point alignment by specifying the appropriate alignment both at compile-time and execution time. For example:

```

: FCONSTANT ( F: r -- )
  CREATE FALIGN HERE 1 FLOATS ALLOT F!
DOES> ( F: -- r ) FALIGNED F@ ;

```

A.12.3.7 Text interpreter input number conversion

The Technical Committee has more than once received the suggestion that the text interpreter in Standard Forth systems should treat numbers that have an embedded decimal point, but no exponent, as floating-point numbers rather than double cell numbers. This suggestion, although it has merit, has always been voted down because it would break too much existing code; many existing implementations put the full digit string on the stack as a double number and use other means to inform the application of the location of the decimal point.

A.12.6 Glossary

A.12.6.1.0558 >FLOAT

>FLOAT enables programs to read floating-point data in legible ASCII format. It accepts a much broader syntax than does the text interpreter since the latter defines rules for composing source programs whereas >FLOAT defines rules for accepting data. >FLOAT is defined as broadly as is feasible to permit input of data from ANS Forth systems as well as other widely used standard programming environments.

This is a synthesis of common FORTRAN practice. Embedded spaces are explicitly forbidden in much scientific usage, as are other field separators such as comma or slash.

While >FLOAT is not required to treat a string of blanks as zero, this behavior is strongly encouraged, since a future version of ANS Forth may include such a requirement.

A.12.6.1.1427 F.

For example, `1E3 F.` displays `1000.`

A.12.6.1.1492 FCONSTANT

Typical use: `r FCONSTANT name`

A.12.6.1.1552 FLITERAL

Typical use: `: X ... [... (r)] FLITERAL ... ;`

A.12.6.1.1630 FVARIABLE

Typical use: `FVARIABLE name`

A.12.6.1.2143 REPRESENT

This word provides a primitive for floating-point display. Some floating-point formats, including those specified by IEEE-754, allow representations of numbers outside of an implementation-defined range. These include plus and minus infinities, denormalized numbers, and others. In these cases we expect that

REPRESENT will usually be implemented to return appropriate character strings, such as “+infinity” or “nan”, possibly truncated.

A.12.6.2.1489 **FATAN2**

FSINCOS and **FATAN2** are a complementary pair of operators which convert angles to 2-vectors and vice-versa. They are essential to most geometric and physical applications since they correctly and unambiguously handle this conversion in all cases except null vectors, even when the tangent of the angle would be infinite.

FSINCOS returns a Cartesian unit vector in the direction of the given angle, measured counter-clockwise from the positive X-axis. The order of results on the stack, namely y underneath x, permits the 2-vector data type to be additionally viewed and used as a ratio approximating the tangent of the angle. Thus the phrase **FSINCOS F/** is functionally equivalent to **FTAN**, but is useful over only a limited and discontinuous range of angles, whereas **FSINCOS** and **FATAN2** are useful for all angles. This ordering has been found convenient for nearly two decades, and has the added benefit of being easy to remember. A corollary to this observation is that vectors in general should appear on the stack in this order.

The argument order for **FATAN2** is the same, converting a vector in the conventional representation to a scalar angle. Thus, for all angles, **FSINCOS FATAN2** is an identity within the accuracy of the arithmetic and the argument range of **FSINCOS**. Note that while **FSINCOS** always returns a valid unit vector, **FATAN2** will accept any non-null vector. An ambiguous condition exists if the vector argument to **FATAN2** has zero magnitude.

A.12.6.2.1516 **FEXPM1**

This function allows accurate computation when its arguments are close to zero, and provides a useful base for the standard exponential functions. Hyperbolic functions such as $\cosh(x)$ can be efficiently and accurately implemented by using **FEXPM1**; accuracy is lost in this function for small values of x if the word **FEXP** is used.

An important application of this word is in finance; say a loan is repaid at 15% per year; what is the daily rate? On a computer with single precision (six decimal digit) accuracy:

1. Using **FLN** and **FEXP**:

FLN of 1.15 = 0.139762,
 divide by 365 = 3.82910E-4,
 form the exponent using **FEXP** = 1.00038, and
 subtract one (1) and convert to percentage = 0.038%.

Thus we only have two digit accuracy.

2. Using **FLNP1** and **FEXPM1**:

FLNP1 of 0.15 = 0.139762, (this is the same value as in the first example, although with the argument closer to zero it may not be so)
 divide by 365 = 3.82910E-4,
 form the exponent and subtract one (1) using **FEXPM1** = 3.82983E-4, and
 convert to percentage = 0.0382983%.

This is full six digit accuracy.

The presence of this word allows the hyperbolic functions to be computed with usable accuracy. For example, the hyperbolic sine can be defined as:

```
: FSINH ( r1 -- r2 )
  FEXPM1 FDUP FDUP 1.0E0 F+ F/ F+ 2.0E0 F/ ;
```

A.12.6.2.1554 FLNP1

This function allows accurate compilation when its arguments are close to zero, and provides a useful base for the standard logarithmic functions. For example, **FLN** can be implemented as:

```
: FLN 1.0E0 F- FLNP1 ;
```

See: **A.12.2.1516 FEXPM1**.

A.12.6.2.1616 FSINCOS

See: **A.12.2.1489 FATAN2**.

A.12.6.2.1640 F~

This provides the three types of “floating point equality” in common use -- “close” in absolute terms, exact equality as represented, and “relatively close”.

A.13 The Locals word set

The Technical Committee has had a problem with locals. It has been argued forcefully that ANS Forth should say nothing about locals since:

- there is no clear accepted practice in this area;
- not all Forth programmers use them or even know what they are; and
- few implementations use the same syntax, let alone the same broad usage rules and general approaches.

It has also been argued, it would seem equally forcefully, that the lack of any standard approach to locals is precisely the reason for this lack of accepted practice since locals are at best non-trivial to implement in a portable and useful way. It has been further argued that users who have elected to become dependent on locals tend to be locked into a single vendor and have little motivation to join the group that it is hoped will “broadly accept” ANS Forth unless the Standard addresses their problems.

Since the Technical Committee has been unable to reach a strong consensus on either leaving locals out or on adopting any particular vendor’s syntax, it has sought some way to deal with an issue that it has been unable to simply dismiss. Realizing that no single mechanism or syntax can simultaneously meet the desires expressed in all the locals proposals that have been received, it has simplified the problem statement to be to define a locals mechanism that:

- is independent of any particular syntax;
- is user extensible;
- enables use of arbitrary identifiers, local in scope to a single definition;
- supports the fundamental cell size data types of Forth; and
- works consistently, especially with respect to re-entrancy and recursion.

This appears to the Technical Committee to be what most of those who actively use locals are trying to achieve with them, and it is at present the consensus of the Technical Committee that if ANS Forth has anything to say on the subject this is an acceptable thing for it to say.

This approach is proposed as one that can be used with a small amount of user coding to implement some, but not all, of the locals schemes in use. The following coding examples illustrate how it can be used to implement two syntaxes.

1. A proposed syntax: (**LOCAL** *name*) with additional usage rules:

```
: LOCAL ( "name" ) BL WORD COUNT (LOCAL) ; IMMEDIATE
: END-LOCALS 0 0 (LOCAL) ; IMMEDIATE
```

```

: EXAMPLE ( n -- n n2 n3) LOCAL N END-LOCALS
N DUP N * DUP N * ;

```

2. The syntax used by systems of Creative Solutions, Inc.:

```

: LOCALS| ( "name...name |" -- )
BEGIN BL WORD COUNT OVER C@
[CHAR] | - OVER 1 - OR
WHILE (LOCAL)
REPEAT 2DROP 0 0 (LOCAL); IMMEDIATE
: EXAMPLE ( n -- n2 n3)
LOCALS| N | N DUP N * DUP N * ;

```

Other syntaxes can be implemented, although some will admittedly require considerably greater effort or in some cases program conversion. Yet other approaches to locals are completely incompatible due to gross differences in usage rules and in some cases even scope identifiers. For example, the complete local scheme in use at Johns Hopkins had elaborate semantics that cannot be duplicated in terms of this model.

To reinforce the intent of section 13, here are two examples of actual use of locals. The first illustrates correct usage:

```

(1.1) : { BEGIN BL WORD COUNT OVER C@
[CHAR] } - OVER 1 - OR
WHILE (LOCAL)
REPEAT 2DROP 0 0 (LOCAL) ; IMMEDIATE
(1.2) : JOE ( a b c -- n ) >R 2* R> 2DUP + 0
{ ANS 2B+C C 2B A }
2 0 DO 1 ANS + I + TO ANS ANS . LOOP
ANS . 2B+C . C . 2B . A . ANS ;
(1.3) 100 300 10 JOE .

```

(1.1) defines a local declaration syntax that surrounds the list of locals with braces. It doesn't do anything fancy, such as reordering locals or providing initial values for some of them, so locals are initialized from the stack in the default order. (1.2) illustrates a use of this syntax in the definition of **JOE**. Note that work is performed at execution time in that definition before locals are declared. It's OK to use the return stack as long as whatever is placed there is removed before the declarations begin.

Note that before declaring locals, **B** is doubled, a subexpression (**2B+C**) is computed, and an initial value (zero) for **ANS** is provided. After locals have been declared, **JOE** proceeds to use them. Note that locals may be accessed and updated within do-loops. The effect of interpreting (1.3) is to display the following values:

```

1 (ANS the first time through the loop),
3 (ANS the second time),
3 (ANS) 610 (2B+C). 10 (C), 600 (2B) 100 (A), and
3 (ANS left on the stack by JOE).

```

The *names* of the locals vanish after **JOE** has been compiled. The *storage and meaning* of locals appear when **JOE**'s locals are declared and vanish as **JOE** returns to its caller at **;** (semicolon).

A second set of examples illustrates various things that break the rules. We assume that the definitions of **LOCAL** and **END-LOCALS** above are present, along with **{** from the preceding example.

```

(2.1) : ZERO 0 POSTPONE LITERAL POSTPONE LOCAL ; IMMEDIATE
(2.2) : JOE ( a b ) ZERO TEMP LOCAL B 1+ LOCAL A+ ZERO ANSWER ;
(2.3) : BOB ( a b c d ) { D C } { B A } ;

```

Here are two definitions with various violations of rule 13.3.3 (1). In (2.2) the declaration of **TEMP** is legal and creates a local whose initial value is zero. It's OK because the executable code that **ZERO**

generates precedes the first use of (**LOCAL**) in the definition. However, the **1+** preceding the declaration of **A+** is illegal. Likewise the use of **ZERO** to define **ANSWER** is illegal because it generates executable code between uses of (**LOCAL**). Finally, **JOE** terminates illegally (no **END-LOCALS**). **BOB** in (2.3) violates the rule against declaring two sets of locals.

```
(2.4) : ANN ( a b -- b) DUP >R DUP IF { B A } THEN R> ;
```

```
(2.5) : JANE ( a b -- n) { B A } A B + >R A B - R> / ;
```

ANN in (2.4) violates two rules. The **IF ... THEN** around the declaration of its locals violates **13.3.3** (2), and the copy of **B** left on the return stack before declaring locals violates **13.3.3** (3). **JANE** in (2.5) violates **13.3.3** (4) by accessing locals after placing the sum of **A** and **B** on the return stack without first removing that sum.

```
(2.6) : CHRIS ( a b) { B A } [ ' ] A EXECUTE
      5 [ ' ] B >BODY ! [ ' A ] LITERAL LEE ;
```

CHRIS in (2.6) illustrates three violations of **13.3.3** (5). The attempt to **EXECUTE** the local called **A** is inconsistent with some implementations. The store into **B** via **>BODY** is likely to cause tragic results with many implementations; moreover, if locals are in registers they can't be addressed as memory no matter what is written.

The third violation, in which an execution token for a definition's local is passed as an argument to the word **LEE** would, if allowed, have the unpleasant implication that **LEE** could **EXECUTE** the token and obtain a value for **A** from the particular execution of **CHRIS** that called **LEE** this time.

A.13.3 Additional usage requirements

Rule **13.3.3** (4) could be relaxed without affecting the integrity of the rest of this structure. **13.3.3** (3) could not be.

13.3.3 (2) forbids the use of the data stack for local storage because no usage rules have been articulated for programmer users in such a case. Of course, if the data stack is somehow employed in such a way that there are no usage rules, then the locals are invisible to the programmer, are logically not on the stack, and the implementation conforms.

The minimum required number of locals can (and should) be adjusted to minimize the cost of compliance for existing users of locals.

Access to previously declared local variables is prohibited by Section **13.3.3**(4) until any data placed onto the return stack by the application has been removed, due to the possible use of the return stack for storage of locals.

Authorization for a Standard Program to manipulate the return stack (e.g., via **>R R>**) while local variables are active overly constrains implementation possibilities. The consensus of users of locals was that Local facilities represent an effective functional replacement for return stack manipulation, and restriction of standard usage to only one method was reasonable.

Access to Locals within **DO. . LOOPS** is expressly permitted as an additional requirement of conforming systems by Section **13.3.3**(7). Although words, such as (**LOCALS**), written by a System Implementor, may require inside knowledge of the internal structure of the return stack, such knowledge is not required of a user of compliant Forth systems.

A.13.6 Glossary

A.13.6.1.2295 TO

Typical use: **x TO name**

See: **A.6.2.2295 TO**.

A.13.6.2.1795 LOCALS |

A possible implementation of this word and an example of usage is given in **A.13**. It is intended as an example only; any implementation yielding the described semantics is acceptable.

A.14 The Memory-Allocation word set

The Memory-Allocation word set provides a means for acquiring memory other than the contiguous data space that is allocated by **ALLOT**. In many operating system environments it is inappropriate for a process to pre-allocate large amounts of contiguous memory (as would be necessary for the use of **ALLOT**). The Memory-Allocation word set can acquire memory from the system at any time, without knowing in advance the address of the memory that will be acquired.

A.15 The Programming-Tools word set

These words have been in widespread common use since the earliest Forth systems.

Although there are environmental dependencies intrinsic to programs using an assembler, virtually all Forth systems provide such a capability. Insofar as many Forth programs are intended for real-time applications and are intrinsically non-portable for this reason, the Technical Committee believes that providing a standard window into assemblers is a useful contribution to Forth programmers.

Similarly, the programming aids **DUMP**, etc., are valuable tools even though their specific formats will differ between CPUs and Forth implementations. These words are primarily intended for use by the programmer, and are rarely invoked in programs.

One of the original aims of Forth was to erase the boundary between “user” and “programmer” -- to give all possible power to anyone who had occasion to use a computer. Nothing in the above labeling or remarks should be construed to mean that this goal has been abandoned.

A.15.6 Glossary**A.15.6.1.0220 .S**

.S is a debugging convenience found on almost all Forth systems. It is universally mentioned in Forth texts.

A.15.6.1.2194 SEE

SEE acts as an on-line form of documentation of words, allowing modification of words by decompiling and regenerating with appropriate changes.

A.15.6.1.2465 WORDS

WORDS is a debugging convenience found on almost all Forth systems. It is universally referred to in Forth texts.

A.15.6.2.0470 ;CODE

Typical use: `: namex ... <create> ... ;CODE ...`

where **namex** is a defining word, and **<create>** is **CREATE** or any user defined word that calls **CREATE**.

A.15.6.2.0930 CODE

Some Forth systems implement the assembly function by adding an **ASSEMBLER** word list to the search order, using the text interpreter to parse a postfix assembly language with lexical characteristics similar to Forth source code. Typically, in such systems, assembly ends when a word **END-CODE** is interpreted.

A.15.6.2.1015 CS-PICK

The intent is to reiterate a *dest* on the control-flow stack so that it can be resolved more than once. For example:

```
\ Conditionally transfer control to beginning of loop
\ This is similar in spirit to C's "continue" statement.
: ?REPEAT ( dest -- dest ) \ Compilation
      ( flag -- ) \ Execution
      0 CS-PICK POSTPONE UNTIL ; IMMEDIATE
: XX ( -- ) \ Example use of ?REPEAT
  BEGIN
  ...
  flag ?REPEAT ( Go back to BEGIN if flag is false )
  ...
  flag ?REPEAT ( Go back to BEGIN if flag is false )
  ...
  flag UNTIL ( Go back to BEGIN if flag is false )
  ... ;
```

A.15.6.2.1020 CS-ROLL

The intent is to modify the order in which the *origs* and *dests* on the control-flow stack are to be resolved by subsequent control-flow words. For example, **WHILE** could be implemented in terms of **IF** and **CS-ROLL**, as follows:

```
: WHILE ( dest -- orig dest )
  POSTPONE IF 1 CS-ROLL ; IMMEDIATE
```

A.15.6.2.1580 FORGET

Typical use: ... **FORGET** name ...

FORGET assumes that all the information needed to restore the dictionary to its previous state is inferable somehow from the forgotten word. While this may be true in simple linear dictionary models, it is difficult to implement in other Forth systems; e.g., those with multiple address spaces. For example, if Forth is embedded in ROM, how does **FORGET** know how much RAM to recover when an array is forgotten? A general and preferred solution is provided by **MARKER**.

A.15.6.2.2531 [ELSE]

Typical use: ... **flag** [**IF**] ... [**ELSE**] ... [**THEN**] ...

A.15.6.2.2532 [IF]

Typical use: ... **flag** [**IF**] ... [**ELSE**] ... [**THEN**] ...

A.15.6.2.2533 [THEN]

Typical use: ... **flag** [**IF**] ... [**ELSE**] ... [**THEN**] ...

Software that runs in several system environments often contains some source code that is environmentally dependent. Conditional compilation -- the selective inclusion or exclusion of portions of the source code at compile time -- is one technique that is often used to assist in the maintenance of such source code.

Conditional compilation is sometimes done with “smart comments” -- definitions that either skip or do not skip the remainder of the line based on some test. For example:

```

\ If 16-Bit? contains TRUE, lines preceded by 16BIT\
\ will be skipped. Otherwise, they will not be skipped.
VARIABLE 16-BIT?
: 16BIT\ 16-BIT? @ IF POSTPONE \ THEN ; IMMEDIATE

```

This technique works on a line by line basis, and is good for short, isolated variant code sequences.

More complicated conditional compilation problems suggest a nestable method that can encompass more than one source line at a time. The words included in the ANS Forth Programmers-Tool-kit extensions word set are useful for this purpose. The implementation given below works with any input source (keyboard, **EVALUATE**, **BLOCK**, or text file).

```

: [ELSE] ( -- )
  1 BEGIN
    BEGIN BL WORD COUNT DUP WHILE \ level
      2DUP S" [IF]" COMPARE 0= IF \ level adr len
      2DROP 1+ \ level'
    ELSE \ level adr len
      2DUP S" [ELSE]" COMPARE 0= IF \ level adr len
      2DROP 1- DUP IF 1+ THEN \ level'
    ELSE \ level adr len
      S" [THEN]" COMPARE 0= IF \ level
      1- \ level'
    THEN
    THEN
      THEN ?DUP 0= IF EXIT THEN \ level'
    REPEAT 2DROP \ level
    REFILL 0= UNTIL \ level
    DROP ; IMMEDIATE
: [IF] ( flag -- )
  0= IF POSTPONE [ELSE] THEN ; IMMEDIATE
: [THEN] ( -- ) ; IMMEDIATE

```

A.16 The Search-Order word set

Search-order specification and control mechanisms vary widely. The FIG-Forth, Forth-79, polyFORTH, and Forth-83 vocabulary and search order mechanisms are all mutually incompatible. The complete list of incompatible mechanisms, in use or proposed, is much longer. The **ALSO/ONLY** scheme described in a Forth-83 Experimental Proposal has substantial community support. However, many consider it to be fundamentally flawed, and oppose it vigorously.

Recognizing this variation, this Standard specifies a new “primitive” set of tools from which various schemes may be constructed. This primitive search-order word set is intended to be a portable “construction set” from which search-order words may be built, rather than a user interface. **ALSO/ONLY** or the various “vocabulary” schemes supported by the major Forth vendors can be defined in terms of the primitive search-order word set.

The encoding for word list identifiers *wid* might be a small-integer index into an array of word-list definition records, the data-space address of such a record, a user-area offset, the execution token of a Forth-83 style sealed vocabulary, the link-field address of the first definition in a word list, or anything else. It is entirely up to the system implementor.

In some systems the interpretation of numeric literals is controlled by including “pseudo word lists” that recognize numbers at the end of the search order. This technique is accommodated by the “default search order” behavior of **SET-ORDER** when given an argument of -1. In a system using the traditional implementation of **ALSO/ONLY**, the minimum search order would be equivalent to the word **ONLY**.

There has never been a portable way to restore a saved search order. F83 (not Forth 83) introduced the word **PREVIOUS**, which almost made it possible to “unload” the search order by repeatedly executing the phrase **CONTEXT @ PREVIOUS**. The search order could be “reloaded” by repeating **ALSO CONTEXT !**. Unfortunately there was no portable way to determine how many word lists were in the search order.

ANS Forth has removed the word **CONTEXT** because in many systems its contents refer to more than one word list, compounding portability problems.

Note that **:** (colon) no longer affects the search order. The previous behavior, where the compilation word list replaces the first word list of the search order, can be emulated with the following redefinition of **:** (colon).

```
 : : GET-ORDER SWAP DROP GET-CURRENT SWAP SET-ORDER : ;
```

A.16.2 Additional terms

search order

Note that the use of the term “list” does not necessarily imply implementation as a linked list.

A.16.3.3 Finding definition names

In other words, the following is not guaranteed to work:

```
 : FOO ... [ ... SET-CURRENT ] ... RECURSE ... ;  
 IMMEDIATE
```

RECURSE, **;** (semicolon), and **IMMEDIATE** may or may not need information stored in the compilation word list.

A.16.6 Glossary

A.16.6.1.2192 SEARCH-WORDLIST

The string argument to **SEARCH-WORDLIST** is represented by *c-addr u*, rather than by just *c-addr* as with **FIND**. The committee wishes to establish *c-addr u* as the preferred representation of a string on the stack, and has adopted that representation for all new functions that accept string arguments. While this decision may cause the implementation of **SEARCH-WORDLIST** to be somewhat more difficult in existing systems, the committee feels that the additional difficulty is minor.

When **SEARCH-WORDLIST** fails to find the word, it does not return the string, as does **FIND**. This is in accordance with the general principle that Forth words consume their arguments.

A.16.6.2.0715 ALSO

Here is an implementation of **ALSO/ONLY** in terms of the primitive search-order word set.

```
WORDLIST CONSTANT ROOT    ROOT SET-CURRENT  
: DO-VOCABULARY ( -- ) \ Implementation factor  
DOES> @ >R          ( ) ( R: widnew )  
  GET-ORDER SWAP DROP ( wid1 ... widn-1 n )  
  R> SWAP SET-ORDER ;  
: DISCARD ( x1 .. xu u - ) \ Implementation factor  
  0 ?DO DROP LOOP ;          \ DROP u+1 stack items  
CREATE FORTH  FORTH-WORDLIST , DO-VOCABULARY  
: VOCABULARY ( name -- )  
  WORDLIST CREATE , DO-VOCABULARY ;  
: ALSO ( -- ) GET-ORDER OVER SWAP 1+ SET-ORDER ;
```



```

: PREVIOUS ( -- ) GET-ORDER
  SWAP DROP 1- SET-ORDER ;
: DEFINITIONS ( -- ) GET-ORDER
  OVER SET-CURRENT DISCARD ;
: ONLY ( -- ) ROOT ROOT 2 SET-ORDER ;
\ Forth-83 version; just removes ONLY
: SEAL ( -- ) GET-ORDER 1- SET-ORDER DROP ;
\ F83 and F-PC version; leaves only CONTEXT
: SEAL ( -- ) GET-ORDER OVER 1 SET-ORDER DISCARD ;

```

The preceding definition of **ONLY** in terms of a “**ROOT**” word list follows F83 usage, and assumes that the default search order just includes **ROOT** and **FORTH**. A more portable definition of **FORTH** and **ONLY**, without the assumptions, is:

```

<omit the ... WORDLIST CONSTANT ROOT ... line>
CREATE FORTH GET-ORDER OVER , DISCARD DO-VOCABULARY
: ONLY ( -- ) -1 SET-ORDER ;

```

Here is a simple implementation of **GET-ORDER** and **SET-ORDER**, including a corresponding definition of **FIND**. The implementations of **WORDLIST**, **SEARCH-WORDLIST**, **GET-CURRENT** and **SET-CURRENT** depend on system details and are not given here.

```

16 CONSTANT #VOCS
VARIABLE #ORDER
CREATE CONTEXT #VOCS CELLS ALLOT
: GET-ORDER ( -- wid1 .. widn n )
  #ORDER @ 0 ?DO
    #ORDER @ I - 1- CELLS CONTEXT + @
  LOOP
  #ORDER @ ;
: SET-ORDER ( wid1 .. widn n -- )
  DUP -1 = IF
    DROP <push system default word lists and n>
  THEN
  DUP #ORDER !
  0 ?DO I CELLS CONTEXT + ! LOOP ;
: FIND ( c-addr -- c-addr 0 | w 1 | w -1 )
  0 ( c-addr 0 )
  #ORDER @ 0 ?DO
    OVER COUNT ( c-addr 0 c-addr' u )
    I CELLS CONTEXT + @ ( c-addr 0 c-addr' u wid)
    SEARCH-WORDLIST ( c-addr 0; 0 | w 1 | w -1 )
    ?DUP IF ( c-addr 0; w 1 | w -1 )
      2SWAP 2DROP LEAVE ( w 1 | w -1 )
    THEN ( c-addr 0 )
  LOOP ; ( c-addr 0 | w 1 | w -1 )

```

In an implementation where the dictionary search mechanism uses a hash table or lookup cache to reduce the search time, **SET-ORDER** might need to reconstruct the hash table or flush the cache.

A.17 The String word set

A.17.6 Glossary

A.17.6.1.0245 /STRING

/STRING is used to remove or add characters relative to the “left” end of the character string. Positive values of n will exclude characters from the string while negative values of n will include characters to the left of the string. **/STRING** is a natural factor of **WORD** and commonly available.

A.17.6.1.0910 CMOVE

If $c\text{-}addr_2$ lies within the source region, memory propagation occurs. ($c\text{-}addr_2$ lies within the source region if $c\text{-}addr_2$ is not less than $c\text{-}addr_1$ and $c\text{-}addr_2$ is less than the quantity $c\text{-}addr_1 u \text{ CHARS } +$).

Typical use: Assume a character string at address 100: “ABCD”. Then after
100 DUP CHAR+ 3 CMOVE the string at address 100 is “AAAA”.

Rationale for **CMOVE** and **CMOVE>** follows **MOVE**.

A.17.6.1.0920 CMOVE>

If $c\text{-}addr_1$ lies within the destination region, memory propagation occurs. ($c\text{-}addr_1$ lies within the destination region if $c\text{-}addr_1$ is greater than or equal to $c\text{-}addr_2$ and if $c\text{-}addr_2$ is less than the quantity $c\text{-}addr_1 u \text{ CHARS } +$.)

Typical use: Assume a character string at address 100: “ABCD”. Then after
100 DUP CHAR+ SWAP 3 CMOVE> the string at address 100 is “DDDD”.

A.17.6.1.0935 COMPARE

Existing Forth systems perform string comparison operations using words that differ in spelling, input and output arguments, and case sensitivity. One in widespread use was chosen.

A.17.6.1.2191 SEARCH

Existing Forth systems perform string searching operations using words that differ in spelling, input and output arguments, and case sensitivity. One in widespread use was chosen.

A.17.6.1.2212 SLITERAL

The current functionality of **6.1.2165 S"** may be provided by the following definition:

```
: S" ( "ccc<quote>" -- )
  [CHAR] " PARSE POSTPONE SLITERAL ; IMMEDIATE
```

B. Bibliography (informative annex)

Industry standards

Forth-77 Standard, Forth Users Group, FST-780314.

Forth-78 Standard, Forth International Standards Team.

Forth-79 Standard, Forth Standards Team.

Forth-83 Standard and Appendices, Forth Standards Team.

The standards referenced in this section were developed by the Forth Standards Team, a volunteer group which included both implementors and users. This was a volunteer organization operating under its own charter and without any formal ties to ANSI, IEEE or any similar standards body. Several members of the Forth Standards Team have also been members of the X3J14 Technical Committee.

Books

Brodie, L. *Starting FORTH*. Englewood Cliffs, NJ: Prentice Hall, 1981.

Brodie, L. *Thinking FORTH*. Englewood Cliffs, NJ: Prentice Hall, 1984.

Feierbach, G. and Thomas, P. *Forth Tools & Applications*. Reston, VA: Reston Computer Books, 1985.

Haydon, Dr. Glen B. *All About FORTH, Third Edition*. La Honda, CA: 1990.

Kelly, Mahlon G. and Spies, N. *FORTH: A Text and Reference*. Englewood Cliffs, NJ: Prentice Hall, 1986.

Knecht, K. *Introduction to Forth*. Indiana: Howard Sams & Co., 1982.

Koopman, P. *Stack Computers, The New Wave*. Chichester, West Sussex, England: Ellis Horwood Ltd. 1989

Martin, Thea, editor. *A Bibliography of Forth References, Third Edition*. Rochester, New York: Institute of Applied Forth Research, 1987.

McCabe, C. K. *Forth Fundamentals* (2 volumes). Oregon: Dilithium Press, 1983.

Pountain, R. *Object Oriented Forth*. New York: Academic Press, 1987.

Ouverson, Marlin, editor. *Dr. Dobbs Toolbook of Forth*. Redwood City, CA: M&T Press, Vol. 1, 1986; Vol. 2, 1987.

Terry, J. D. *Library of Forth Routines and Utilities*. New York: Shadow Lawn Press, 1986

Tracy, M. and Anderson, A. *Mastering Forth* (2nd ed). New York: Brady Books, 1989.

Winfield, A. *The Complete Forth*. New York: Wiley Books, 1983.

Journals, magazines and newsletters

Forsley, Lawrence P., Conference Chairman. *Rochester Forth Conference Proceedings*. Rochester, New York: Institute of Applied Forth Research, 1981 to present.

Forsley, Lawrence P., Editor-in-Chief. *The Journal of Forth Application and Research*. Rochester, New York: Institute of Applied Forth Research, 1983 to present.

Frenger, Paul, editor. *SIGForth Newsletter*. New York, NY: Association for Computing Machinery, 1989 to present.

Ouverson, Marlin, editor. *Forth Dimensions*. San Jose, CA: The Forth Interest Group, 1978 to present.

Reiling, Robert, editor. *FORML Conference Proceedings*. San Jose, CA: The Forth Interest Group, 1980 to present.

Ting, Dr. C. H., editor. *More on Forth Engines*. San Mateo, CA: Offete Enterprises, 1986 to present.

Selected articles

Kelly, Guy M. "Forth." *McGraw-Hill Personal Computer Programming Encyclopedia – Languages and Operation Systems*. New York: McGraw-Hill, 1985.

Kogge, P. M. "An Architectural Trail to Threaded Code Systems." *IEEE Computer* (March, 1982).

Moore, C. H. "The Evolution of FORTH – An Unusual Language." *Byte* (August 1980).

Rather, E. D. "Forth Programming Language." *Encyclopedia of Physical Science & Technology* (Vol. 5). New York: Academic Press, 1987.

Rather, E. D. "FORTH." *Computer Programming Management*. Auerbach Publishers, Inc., 1985.

C. Perspective (informative annex)

The purpose of this section is to provide an informal overview of Forth as a language, illustrating its history, most prominent features, usage, and common implementation techniques. Nothing in this section should be considered as binding upon either implementors or users. A list of books and articles is given in Annex **B** for those interested in learning more about Forth.

C.1 Features of Forth

Forth provides an interactive programming environment. Its primary uses have been in scientific and industrial applications such as instrumentation, robotics, process control, graphics and image processing, artificial intelligence and business applications. The principal advantages of Forth include rapid, interactive software development and efficient use of computer hardware.

Forth is often spoken of as a language because that is its most visible aspect. But in fact, Forth is both more and less than a conventional programming language: more in that all the capabilities normally associated with a large portfolio of separate programs (compilers, editors, etc.) are included within its range and less in that it lacks (deliberately) the complex syntax characteristic of most high-level languages.

The original implementations of Forth were stand-alone systems that included functions normally performed by separate operating systems, editors, compilers, assemblers, debuggers and other utilities. A single simple, consistent set of rules governed this entire range of capabilities. Today, although very fast stand-alone versions are still marketed for many processors, there are also many versions that run co-resident with conventional operating systems such as MS-DOS and UNIX.

Forth is not derived from any other language. As a result, its appearance and internal characteristics may seem unfamiliar to new users. But Forth's simplicity, extreme modularity, and interactive nature offset the initial strangeness, making it easy to learn and use. A new Forth programmer must invest some time mastering its large command repertoire. After a month or so of full-time use of Forth, that programmer could understand more of its internal working than is possible with conventional operating systems and compilers.

The most unconventional feature of Forth is its *extensibility*. The programming process in Forth consists of defining new "words" – actually new commands in the language. These may be defined in terms of previously defined words, much as one teaches a child concepts by explaining them in terms of previously understood concepts. Such words are called "high-level definitions." Alternatively, new words may also be defined in assembly code, since most Forth implementations include an assembler for the host processor.

This extensibility facilitates the development of special application languages for particular problem areas or disciplines.

Forth's extensibility goes beyond just adding new commands to the language. With equivalent ease, one can also add new *kinds* of words. That is, one may create a word which itself will define words. In creating such a defining word the programmer may specify a specialized behavior for the words it will create which will be effective at compile time, at run-time, or both. This capability allows one to define specialized data types, with complete control over both structure and behavior. Since the run-time behavior of such words may be defined either in high-level or in code, the words created by this new defining word are equivalent to all other kinds of Forth words in performance. Moreover, it is even easy to add new *compiler directives* to implement special kinds of loops or other control structures.

Most professional implementations of Forth are written in Forth. Many Forth systems include a "meta-compiler" which allows the user to modify the internal structure of the Forth system itself.

C.2 History of Forth

Forth was invented by Charles H. Moore. A direct outgrowth of Moore's work in the 1960's, the first program to be called Forth was written in about 1970. The first complete implementation was used in 1971 at the National Radio Astronomy Observatory's 11-meter radio telescope in Arizona. This system was responsible for pointing and tracking the telescope, collecting data and recording it on magnetic tape, and supporting an interactive graphics terminal on which an astronomer could analyze previously recorded data. The multi-tasking nature of the system allowed all these functions to be performed concurrently, without timing conflicts or other interference – a very advanced concept for that time.

The system was so useful that astronomers from all over the world began asking for copies. Its use spread rapidly, and in 1976 Forth was adopted as a standard language by the International Astronomical Union.

In 1973, Moore and colleagues formed FORTH, Inc. to explore commercial uses of the language. FORTH, Inc. developed multi-user versions of Forth on minicomputers for diverse projects ranging from data bases to scientific applications such as image processing. In 1977, FORTH, Inc. developed a version for the newly introduced 8-bit microprocessors called "microFORTH," which was successfully used in embedded microprocessor applications in the United States, Britain and Japan.

Stimulated by the volume marketing of microFORTH, a group of computer hobbyists in Northern California became interested in Forth, and in 1978 formed the Forth Interest Group (FIG). They developed a simplified model which they implemented on several microprocessors and published listings and disks at very low cost. Interest in Forth spread rapidly, and today there are chapters of the Forth Interest Group throughout the U.S. and in over fifteen countries.

By 1980, a number of new Forth vendors had entered the market with versions of Forth based upon the FIG model. Primarily designed for personal computers, these relatively inexpensive Forth systems have been distributed very widely.

C.3 Hardware implementations of Forth

The internal architecture of Forth simulates a computer with two stacks, a set of registers, and other standardized features. As a result, it was almost inevitable that someone would attempt to build a hardware representation of an actual Forth computer.

In the early 1980's, Rockwell produced a 6502-variant with Forth primitives in on-board ROM, the Rockwell 65F11. This chip has been used successfully in many embedded microprocessor applications. In the mid-1980's Zilog developed the z8800 (Super8) which offered ENTER (nest), EXIT (unnest) and NEXT in microcode.

In 1981, Moore undertook to design a chip-level implementation of the Forth virtual machine. Working first at FORTH, Inc. and subsequently with the start-up company NOVIX, formed to develop the chip, Moore completed the design in 1984, and the first prototypes were produced in early 1985. More recently, Forth processors have been developed by Harris Semiconductor Corp., Johns Hopkins University, and others.

C.4 Standardization efforts

The first major effort to standardize Forth was a meeting in Utrecht in 1977. The attendees produced a preliminary standard, and agreed to meet the following year. The 1978 meeting was also attended by members of the newly formed Forth Interest Group. In 1979 and 1980 a series of meetings attended by both users and vendors produced a more comprehensive standard called Forth 79.

Although Forth 79 was very influential, many Forth users and vendors found serious flaws in it, and in 1983 a new standard called Forth 83 was released.

Encouraged by the widespread acceptance of Forth 83, a group of users and vendors met in 1986 to investigate the feasibility of an American National Standard. The X3J14 Technical Committee for ANS Forth held its first meeting in 1987. This Standard is the result.

C.5 Programming in Forth

Forth is an English-like language whose elements (called “words”) are named data items, procedures, and defining words capable of creating data items with customized characteristics. Procedures and defining words may be defined in terms of previously defined words or in machine code, using an embedded assembler.

Forth “words” are functionally analogous to subroutines in other languages. They are also equivalent to commands in other languages – Forth blurs the distinction between linguistic elements and functional elements.

Words are referred to either from the keyboard or in program source by name. As a result, the term “word” is applied both to program (and linguistic) units and to their text names. In parsing text, Forth considers a word to be any string of characters bounded by spaces. There are a few special characters that cannot be included in a word or start a word: space (the universal delimiter), CR (which ends terminal input), and backspace or DEL (for backspacing during keyboard input). Many groups adopt naming conventions to improve readability. Words encountered in text fall into three categories: defined words (i.e., Forth routines), numbers, and undefined words. For example, here are four words:

```
HERE      DOES>      !      8493
```

The first three are standard-defined words. This means that they have entries in Forth’s dictionary, described below, explaining what Forth is to do when these words are encountered. The number “8493” will presumably not be found in the dictionary, and Forth will convert it to binary and place it on its push-down stack for parameters. When Forth encounters an undefined word and cannot convert it to a number, the word is returned to the user with an exception message.

Architecturally, Forth words adhere strictly to the principles of “structured programming”:

- Words must be defined before they are used.
- Logical flow is restricted to sequential, conditional, and iterative patterns. Words are included to implement the most useful program control structures.
- The programmer works with many small, independent modules (words) for maximum testability and reliability.

Forth is characterized by five major elements: a dictionary, two push-down stacks, interpreters, an assembler, and virtual storage. Although each of these may be found in other systems, the combination produces a synergy that yields a powerful and flexible system.

C.5.1 The Forth dictionary

A Forth program is organized into a dictionary that occupies most of the memory used by the system. This dictionary is a threaded list of variable-length items, each of which defines a word. The content of each definition depends upon the type of word (data item, constant, sequence of operations, etc.). The dictionary is extensible, usually growing toward high memory. On some multi-user systems individual users have private dictionaries, each of which is connected to a shared system dictionary.

Words are added to the dictionary by “defining words”, of which the most commonly used is **:** (colon). When **:** is executed, it constructs a definition for the word that follows it. In classical implementations¹, the content of this definition is a string of addresses of previously defined words which will be executed in turn whenever the word being defined is invoked. The definition is terminated by **;** (semicolon). For example, here is a definition:

```
: RECEIVE ( -- addr n ) PAD DUP 32 ACCEPT ;
```

The name of the new word is **RECEIVE**. The comment (in parentheses) indicates that it requires no parameters and will return an address and count on the data stack. When **RECEIVE** is executed, it will

¹ Other common implementation techniques include direct translation to code and other types of tokens.

perform the words in the remainder of the definition in sequence. The word **PAD** places on the stack the address of a scratch pad used to handle strings. **DUP** duplicates the top stack item, so we now have two copies of the address. The number 32 is also placed on the stack. The word **ACCEPT** takes an address (provided by **PAD**) and length (32) on the stack, accepts from the keyboard a string of up to 32 characters which will be placed at the specified address, and returns the number of characters received. The copy of the scratch-pad address remains on the stack below the count so that the routine that called **RECEIVE** can use it to pick up the received string.

C.5.2 Push-down stacks

The example above illustrates the use of push-down stacks for passing parameters between Forth words. Forth maintains two push-down stacks, or LIFO lists. These provide communication between Forth words plus an efficient mechanism for controlling logical flow. A stack contains 16-bit items on 8-bit and 16-bit computers, and 32-bit items on 32-bit processors. Double-cell numbers occupy two stack positions, with the most-significant part on top. Items on either stack may be addresses or data items of various kinds. Stacks are of indefinite size, and usually grow towards low memory.

Although the structure of both stacks is the same, they have very different uses. The user interacts most directly with the Data Stack, which contains arguments passed between words. This function replaces the calling sequences used by conventional languages. It is efficient internally, and makes routines intrinsically re-entrant. The second stack is called the Return Stack, as its main function is to hold return addresses for nested definitions, although other kinds of data are sometimes kept there temporarily.

The use of the Data Stack (often called just “the stack”) leads to a notation in which operands precede operators. The word **ACCEPT** in the example above took an address and count from the stack and left another address there. Similarly, a word called **BLANK** expects an address and count, and will place the specified number of space characters (20H) in the region starting at that address. Thus,

```
PAD 25 BLANK
```

will fill the scratch region whose address is pushed on the stack by **PAD** with 25 spaces. Application words are usually defined to work similarly. For example,

```
100 SAMPLES
```

might be defined to record 100 measurements in a data array.

Arithmetic operators also expect values and leave results on the stack. For example, **+** adds the top two numbers on the stack, replacing them both by their sum. Since results of operations are left on the stack, operations may be strung together without a need to define variables to use for temporary storage.

C.5.3 Interpreters

Forth is traditionally an interpretive system, in that program execution is controlled by data items rather than machine code. Interpreters can be slow, but Forth maintains the high speed required of real-time applications by having two levels of interpretation.

The first is the text interpreter, which parses strings from the terminal or mass storage and looks each word up in the dictionary. When a word is found it is executed by invoking the second level, the address interpreter.

The second is an “address interpreter.” Although not all Forth systems are implemented in this way, it was the first and is still the primary implementation technology. For a small cost in performance, an address interpreter can yield a very compact object program, which has been a major factor in Forth’s wide acceptance in embedded systems and other applications where small object size is desirable.

The address interpreter processes strings of addresses or tokens compiled in definitions created by **:** (colon), by executing the definition pointed to by each. The content of most definitions is a sequence of addresses of previously defined words, which will be executed by the address interpreter in turn. Thus, when the word **RECEIVE** (defined above) is executed, the word **PAD**, the word **DUP**, the literal 32, and the

word **ACCEPT** will be executed in sequence. The process is terminated by the semicolon. This execution requires no dictionary searches, parsing, or other logic, because when **RECEIVE** was *compiled* the dictionary was searched for each word, and its address (or other token) was placed in the next successive cell of the entry. The text was not stored in memory, even in condensed form.

The address interpreter has two important properties. First, it is fast. Although the actual speed depends upon the specific implementation, professional implementations are highly optimized, often requiring only one or two machine instructions per address. On most benchmarks, a good Forth implementation substantially out-performs interpretive languages such as BASIC or LISP, and will compare favorably with other compiled high-level languages.

Second, the address interpreter makes Forth definitions extremely compact, as each reference requires only one cell. In comparison, a subroutine call constructed by most compilers involves instructions for handling the calling sequence (unnecessary in Forth because of the stack) before and after a **CALL** or **JSR** instruction and address.

Most of the words in a Forth dictionary will be defined by **:** (colon) and interpreted by the address interpreter. Most of Forth itself is defined this way.

C.5.4 Assembler

Most implementations of Forth include a macro assembler for the CPU on which they run. By using the defining word **CODE** the programmer can create a definition whose behavior will consist of executing actual machine instructions. **CODE** definitions may be used to do I/O, implement arithmetic primitives, and do other machine-dependent or time-critical processing. When using **CODE** the programmer has full control over the CPU, as with any other assembler, and **CODE** definitions run at full machine speed.

This is an important feature of Forth. It permits explicit computer-dependent code in manageable pieces with specific interfacing conventions that are machine-independent. To move an application to a different processor requires re-coding only the **CODE** words, which will interact with other Forth words in exactly the same manner.

Forth assemblers are so compact (typically a few Kbytes) that they can be resident in the system (as are the compiler, editor, and other programming tools). This means that the programmer can type in short **CODE** definitions and execute them immediately. This capability is especially valuable in testing custom hardware.

C.5.5 Virtual memory

The final unique element of Forth is its way of using disk or other mass storage as a form of “virtual memory” for data and program source. As in the case of the address interpreter, this approach is historically characteristic of Forth, but is by no means universal. Disk is divided into 1024-byte blocks. Two or more buffers are provided in memory, into which blocks are read automatically when referred to. Each block has a fixed block number, which in native systems is a direct function of its physical location. If a block is changed in memory, it will be automatically written out when its buffer must be reused. Explicit reads and writes are not needed; the program will find the data in memory whenever it accesses it.

Block-oriented disk handling is efficient and easy for native Forth systems to implement. As a result, blocks provide a completely transportable mechanism for handling program source and data across both native and co-resident versions of Forth on different host operating systems.

Definitions in program source blocks are compiled into memory by the word **LOAD**. Most implementations include an editor, which formats a block for display into 16 lines of 64 characters each, and provides commands modifying the source. An example of a Forth source block is given in Fig. C.1 below.

Source blocks have historically been an important element in Forth style. Just as Forth definitions may be considered the linguistic equivalent of sentences in natural languages, a block is analogous to a paragraph.

A block normally contains definitions related to a common theme, such as “vector arithmetic.” A comment on the top line of the block identifies this theme. An application may selectively load the blocks it needs.

Blocks are also used to store data. Small records can be combined into a block, or large records spread over several blocks. The programmer may allocate blocks in whatever way suits the application, and on native systems can increase performance by organizing data to minimize disk head motion. Several Forth vendors have developed sophisticated file and data base systems based on Forth blocks.

Versions of Forth that run co-resident with a host OS often implement blocks in host OS files. Others use the host files exclusively. The Standard requires that blocks be available on systems providing any disk access method, as they are the only means of referencing disk that can be transportable across both native and co-resident implementations.

C.5.6 Programming environment

Although this Standard does not require it, most Forth systems include a resident editor. This enables a programmer to edit source and recompile it into executable form without leaving the Forth environment. As it is easy to organize an application into layers, it is often possible to recompile only the topmost layer (which is usually the one currently under development), a process which rarely takes more than a few seconds.

Most Forth systems also provide resident interactive debugging aids, not only including words such as those in **15. The optional Programming-Tools word set**, but also having the ability to examine and change the contents of **VARIABLES** and other data items and to execute from the keyboard most of the component words in both the underlying Forth system and the application under development.

The combination of resident editor, integrated debugging tools, and direct executability of most defined words leads to a very interactive programming style, which has been shown to shorten development time.

C.5.7 Advanced programming features

One of the unusual characteristics of Forth is that the words the programmer defines in building an application become integral elements of the language itself, adding more and more powerful application-oriented features.

For example, Forth includes the words **VARIABLE** and **2VARIABLE** to name locations in which data may be stored, as well as **CONSTANT** and **2CONSTANT** to name single and double-cell values. Suppose a programmer finds that an application needs arrays that would be automatically indexed through a number of two-cell items. Such an array might be called **2ARRAY**. The prefix “2” in the name indicates that each element in this array will occupy two cells (as would the contents of a **2VARIABLE** or **2CONSTANT**). The prefix “2”, however, has significance only to a human and is no more significant to the text interpreter than any other character that may be used in a definition name.

Such a definition has two parts, as there are two “behaviors” associated with this new word **2ARRAY**, one at compile time, and one at run or execute time. These are best understood if we look at how **2ARRAY** is used to define its arrays, and then how the array might be used in an application. In fact, this is how one would design and implement this word.

Beginning the top-down design process, here’s how we would like to use **2ARRAY**:

```
100 2ARRAY RAW    50 2ARRAY REFINED
```

In the first case, we are defining an array 100 elements long, whose name is **RAW**. In the second, the array is 50 elements long, and is named **REFINED**. In each case, a size parameter is supplied to **2ARRAY** on the data stack (Forth’s text interpreter automatically puts numbers there when it encounters them), and the name of the word immediately follows. This order is typical of Forth defining words.

When we use **RAW** or **REFINED**, we would like to supply on the stack the index of the element we want, and get back the address of that element on the stack. Such a reference would characteristically take place

in a loop. Here's a representative loop that accepts a two-cell value from a hypothetical application word **DATA** and stores it in the next element of **RAW**:

```
: ACQUIRE 100 0 DO DATA I RAW 2! LOOP ;
```

The name of this definition is **ACQUIRE**. The loop begins with **DO**, ends with **LOOP**, and will execute with index values running from 0 through 99. Within the loop, **DATA** gets a value. The word **I** returns the current value of the loop index, which is the argument to **RAW**. The address of the selected element, returned by **RAW**, and the value, which has remained on the stack since **DATA**, are passed to the word **2!** (pronounced "two-store"), which stores two stack items in the address.

Now that we have specified exactly what **2ARRAY** does and how the words it defines are to behave, we are ready to write the two parts of its definition:

```
: 2ARRAY ( n -- ) CREATE 2* CELLS ALLOT
DOES> ( i a -- a' ) SWAP 2* CELLS + ;
```

The part of the definition before the word **DOES>** specifies the "compile-time" behavior, that is, what the **2ARRAY** will do when it is used to define a word such as **RAW**. The comment indicates that this part expects a number on the stack, which is the size parameter. The word **CREATE** constructs the definition for the new word. The phrase **2* CELLS** converts the size parameter from two-cell units to the internal addressing units of the system (normally characters). **ALLOT** then allocates the specified amount of memory to contain the data to be associated with the newly defined array.

The second line defines the "run-time" behavior that will be shared by all words defined by **2ARRAY**, such as **RAW** and **REFINED**. The word **DOES>** terminates the first part of the definition and begins the second part. A second comment here indicates that this code expects an index and an address on the stack, and will return a different address. The index is supplied on the stack by the caller (of **RAW** in the example), while the address of the content of a word defined in this way (the **ALLOT**ted region) is automatically pushed on top of the stack before this section of the code is to be executed. This code works as follows: **SWAP** reverses the order of the two stack items, to get the index on top. **2* CELLS** converts the index to the internal addressing units as in the compile-time section, to yield an offset from the beginning of the array. The word **+** then adds the offset to the address of the start of the array to give the effective address which is the desired result.

Given this basic definition, one could easily modify it to do more sophisticated things. For example, the compile-time code could be changed to initialize the array to zeros, spaces, or any other desired initial value. The size of the array could be compiled at its beginning, so that the run-time code could compare the index against it to ensure it is within range. Or the entire array could be made to reside on disk instead of main memory. *None of these changes would affect the run-time usage we have specified in any way.* This illustrates a little of the flexibility available with these defining words.

C.5.8 A programming example

Figure C.1 contains a typical block of Forth source. It represents a portion of an application that controls a bank of eight LEDs used as indicator lamps on an instrument, and indicates some of the ways in which Forth definitions of various kinds combine in an application environment. This example was coded for a STD-bus system with an 8088 processor and a millisecond clock, which is also used in the example.

The LEDs are interfaced through a single 8-bit port whose address is 40H. This location is defined as a **CONSTANT** on Line 1, so that it may be referred to by name; should the address change, one need only adjust the value of this constant. The word **LIGHTS** returns this address on the stack. The definition **LIGHT** takes a value on the stack and sends it to the device. The nature of this value is a bit mask, whose bits correspond directly to the individual lights. Thus, the command:

```
255 LIGHT
```

will turn on all lights, while:

```
0 LIGHT
```


The time to compile this block of code on that system was about half a second, including the time to fetch it from disk. So it is quite practical (and normal practice) for a programmer to simply type in a definition and try it immediately.

In addition, one always has the capability of communicating with external devices directly. The first thing one would do when told about the lamps would be to type:

HEX FF 40 OUTPUT

and see if all the lamps come on. If not, the presumption is that something is amiss with the hardware, since this phrase directly transmits the “all ones” mask to the device. This type of direct interaction is useful in applications involving custom hardware, as it reduces hardware debugging time.

C.6 Multiprogrammed systems

Multiprogrammed Forth systems have existed since about 1970. The earliest public Forth systems propagated the “hooks” for this capability despite the fact that many did not use them. Nevertheless the underlying assumptions have been common knowledge in the community, and there exists considerable common ground among these multiprogrammed systems. These systems are not just language processors, but contain operating system characteristics as well. Many of these integrated systems run entirely stand-alone, performing all necessary operating system functions.

Some Forth systems are very fast, and can support both multi-tasking and multi-user operation even on computers whose hardware is usually thought incapable of such advanced operation. For example, one producer of telephone switchboards is running over 50 tasks on a Z80. There are several multiprogrammed products for PC's, some of which even support multiple users. Even on computers that are commonly used in multi-user operations, the number of users that can be supported may be much larger than expected. One large data-base application running on a single 68000 has over 100 terminals updating and querying its data-base, with no significant degradation.

Multi-user systems may also support multiple programmers, each of which has a private dictionary, stacks, and a set of variables controlling that task. The private dictionary is linked to a shared, re-entrant dictionary containing all the standard Forth functions. The private dictionary can be used to develop application code which may later be integrated into the shared dictionary. It may also be used to perform functions requiring text interpretation, including compilation and execution of source code.

C.7 Design and management considerations

Just as the choice of building materials has a strong effect on the design and construction of a building, the choice of language and operating system will affect both application design and project management decisions.

Conventionally, software projects progress through four stages: analysis, design, coding, and testing. A Forth project necessarily incorporates these activities as well. Forth is optimized for a project-management methodology featuring small teams of skilled professionals. Forth encourages an iterative process of “successive prototyping” wherein high-level Forth is used as an executable design tool, with “stubs” replacing lower-level routines as necessary (e.g., for hardware that isn't built yet).

In many cases successive prototyping can produce a sounder, more useful product. As the project progresses, implementors learn things that could lead to a better design. Wiser decisions can be made if true relative costs are known, and often this isn't possible until prototype code can be written and tried.

Using Forth can shorten the time required for software development, and reduce the level of effort required for maintenance and modifications during the life of the product as well.

C.8 Conclusion

Forth has produced some remarkable achievements in a variety of application areas. In the last few years its acceptance has grown rapidly, particularly among programmers looking for ways to improve their productivity and managers looking for ways to simplify new software-development projects.

D. Compatibility analysis of ANS Forth (informative annex)

Prior to ANS Forth, there were several industry standards for Forth. The most influential are listed here in chronological order, along with the major differences between ANS Forth and the most recent, Forth 83.

D.1 FIG Forth (circa 1978)

FIG Forth was a “model” implementation of the Forth language developed by the Forth Interest Group (FIG). In FIG Forth, a relatively small number of words were implemented in processor-dependent machine language and the rest of the words were implemented in Forth. The FIG model was placed in the public domain, and was ported to a wide variety of computer systems. Because the bulk of the FIG Forth implementation was the same across all machines, programs written in FIG Forth enjoyed a substantial degree of portability, even for “system-level” programs that directly manipulate the internals of the Forth system implementation.

FIG Forth implementations were influential in increasing the number of people interested in using Forth. Many people associate the implementation techniques embodied in the FIG Forth model with “the nature of Forth”.

However, FIG Forth was not necessarily representative of commercial Forth implementations of the same era. Some of the most successful commercial Forth systems used implementation techniques different from the FIG Forth “model”.

D.2 Forth 79

The Forth-79 Standard resulted from a series of meetings from 1978 to 1980, by the Forth Standards Team, an international group of Forth users and vendors. (Interim versions known as Forth 77 and Forth 78 were released by the group).

Forth 79 described a set of words defined on a 16-bit, twos-complement, unaligned, linear byte-addressing virtual machine. It prescribed an implementation technique known as “indirect threaded code”, and used the ASCII character set.

The Forth-79 Standard served as the basis for several public domain and commercial implementations, some of which are still available and supported today.

D.3 Forth 83

The Forth-83 Standard, also by the Forth Standards Team, was released in 1983. Forth 83 attempted to fix some of the deficiencies of Forth 79.

Forth 83 was similar to Forth 79 in most respects. However, Forth 83 changed the definition of several well-defined features of Forth 79. For example, the rounding behavior of integer division, the base value of the operands of **PICK** and **ROLL**, the meaning of the address returned by **'**, the compilation behavior of **'**, the value of a “true” flag, the meaning of **NOT**, and the “chaining” behavior of words defined by **VOCABULARY** were all changed. Forth 83 relaxed the implementation restrictions of Forth 79 to allow any kind of threaded code, but it did not fully allow compilation to native machine code (this was not specifically prohibited, but rather was an indirect consequence of another provision).

Many new Forth implementations were based on the Forth-83 Standard, but few “strictly compliant” Forth-83 implementations exist.

Although the incompatibilities resulting from the changes between Forth 79 and Forth 83 were usually relatively easy to fix, a number of successful Forth vendors did not convert their implementations to be Forth 83 compliant. For example, the most successful commercial Forth for Apple Macintosh computers is based on Forth 79.

D.4 Recent developments

Since the Forth-83 Standard was published, the computer industry has undergone rapid and profound changes. The speed, memory capacity, and disk capacity of affordable personal computers have increased by factors of more than 100. 8-bit processors have given way to 16-bit processors, and now 32-bit processors are commonplace.

The operating systems and programming-language environments of small systems are much more powerful than they were in the early 80's.

The personal-computer marketplace has changed from a predominantly "hobbyist" market to a mature business and commercial market.

Improved technology for designing custom microprocessors has resulted in the design of numerous "Forth chips", computers optimized for the execution of the Forth language.

The market for ROM-based embedded control computers has grown substantially.

In order to take full advantage of this evolving technology, and to better compete with other programming languages, many recent Forth implementations have ignored some of the "rules" of previous Forth standards. In particular:

- 32-bit Forth implementations are now common.
- Some Forth systems adopt the address-alignment restrictions of the hardware on which they run.
- Some Forth systems use native-code generation, microcode generation, and optimization techniques, rather than the traditional "threaded code".
- Some Forth systems exploit segmented addressing architectures, placing portions of the Forth "dictionary" in different segments.
- More and more Forth systems now run in the environment of another "standard" operating system, using OS text files for source code, rather than the traditional Forth "blocks".
- Some Forth systems allow external operating system software, windowing software, terminal concentrators, or communications channels to handle or preprocess user input, resulting in deviations from the input editing, character set availability, and screen management behavior prescribed by Forth 83.

Competitive pressure from other programming languages (predominantly "C") and from other Forth vendors have led Forth vendors to optimizations that do not fit in well with the "virtual machine model" implied by existing Forth standards.

D.5 ANS Forth approach

The ANS Forth committee addressed the serious fragmentation of the Forth community caused by the differences between Forth 79 and Forth 83, and the divergence from either of these two industry standards caused by marketplace pressures.

Consequently, the committee has chosen to base its compatibility decisions not upon a strict comparison with the Forth-83 Standard, but instead upon consideration of the variety of existing implementations, especially those with substantial user bases and/or considerable success in the marketplace.

The committee feels that, if ANS Forth prescribes stringent requirements upon the virtual machine model, as did the previous standards, then many implementors will choose not to comply with ANS Forth. The committee hopes that ANS Forth will serve to unify rather than to further divide the Forth community, and thus has chosen to encompass rather than invalidate popular implementation techniques.

Many of the changes from Forth 83 are justified by this rationale. Most fall into the category that "an ANS Forth Standard Program may not assume x", where "x" is an entitlement resulting from the virtual machine model prescribed by the Forth-83 Standard. The committee feels that these restrictions are reasonable, especially considering that a substantial number of existing Forth implementations do not

correctly implement the Forth-83 virtual model, thus the Forth-83 entitlements exist “in theory” but not “in practice”.

Another way of looking at this is that while ANS Forth acknowledges the diversity of current Forth practice, it attempts to document the similarity therein. In some sense, ANS Forth is thus a “description of reality” rather than a “prescription for a particular virtual machine”.

Since there is no previous American National Standard for Forth, the action requirements prescribed by section 3.4 of X3/SD-9, “Policy and Guidelines”, regarding previous standards do not apply.

The following discussion describes differences between ANS Forth and Forth 83. In most cases, Forth 83 is representative of Forth 79 and FIG Forth for the purposes of this discussion. In many of these cases, however, ANS Forth is more representative of the existing state of the Forth industry than the previously-published standards.

D.6 Differences from Forth 83

D.6.1 Stack width

Forth 83 specifies that stack items occupy 16 bits. This includes addresses, flags, and numbers. ANS Forth specifies that stack items are at least 16 bits; the actual size must be documented by the implementation.

Words affected: all arithmetic, logical and addressing operators

Reason: 32-bit machines are becoming commonplace. A 16-bit Forth system on a 32-bit machine is not competitive.

Impact: Programs that assume 16-bit stack width will continue to run on 16-bit machines; ANS Forth does not require a different stack width, but simply allows it. Many programs will be unaffected (but see “address unit”).

Transition/Conversion: Programs which use bit masks with the high bits set may have to be changed, substituting either an implementation-defined bit-mask constant, or a procedure to calculate a bit mask in a stack-width-independent way. Here are some procedures for constructing width-independent bit masks:

```

1  CONSTANT LO-BIT
TRUE 1 RSHIFT  INVERT  CONSTANT HI-BIT
: LO-BITS  ( n -- mask )
  0 SWAP  0 ?DO  1 LSHIFT  LO-BIT OR  LOOP ;
: HI-BITS  ( n -- mask )
  0 SWAP  0 ?DO  1 RSHIFT  HI-BIT OR  LOOP ;

```

Programs that depend upon the “modulo 65536” behavior implicit in 16-bit arithmetic operations will need to be rewritten to explicitly perform the modulus operation in the appropriate places. The committee believes that such assumptions occur infrequently. Examples: some checksum or CRC calculations, some random number generators and most fixed-point fractional math.

D.6.2 Number representation

Forth 83 specifies two’s-complement number representation and arithmetic. ANS Forth also allows one’s-complement and signed-magnitude.

Words affected: all arithmetic and logical operators, **LOOP**, **+LOOP**

Reason: Some computers use one’s-complement or signed-magnitude. The committee did not wish to force Forth implementations for those machines to emulate two’s-complement arithmetic, and thus incur severe performance penalties. The experience of some committee members with such machines

indicates that the usage restrictions necessary to support their number representations are not overly burdensome.

Impact: An ANS Forth Standard Program may declare an “environmental dependency on two’s-complement arithmetic”. This means that the otherwise-Standard Program is only guaranteed to work on two’s-complement machines. Effectively, this is not a severe restriction, because the overwhelming majority of current computers use two’s-complement. The committee knows of no Forth-83 compliant implementations for non-two’s-complement machines at present, so existing Forth-83 programs will still work on the same class of machines on which they currently work.

Transition/Conversion: Existing programs wishing to take advantage of the possibility of ANS Forth Standard Systems on non-two’s-complement machines may do so by eliminating the use of arithmetic operators to perform logical functions, by deriving bit-mask constants from bit operations as described in the section about stack width, by restricting the usage range of unsigned numbers to the range of positive numbers, and by using the provided operators for conversion from single numbers to double numbers.

D.6.3 Address units

Forth 83 specifies that each unique address refers to an 8-bit byte in memory. ANS Forth specifies that the size of the item referred to by each unique address is implementation-defined, but, by default, is the size of one character. Forth 83 describes many memory operations in terms of a number of bytes. ANS Forth describes those operations in terms of a number of either characters or address units.

Words affected: those with “address unit” arguments

Reason: Some machines, including the most popular Forth chip, address 16-bit memory locations instead of 8-bit bytes.

Impact: Programs may choose to declare an environmental dependency on byte addressing, and will continue to work on the class of machines for which they now work. In order for a Forth implementation on a word-addressed machine to be Forth 83 compliant, it would have to simulate byte addressing at considerable cost in speed and memory efficiency. The committee knows of no such Forth-83 implementations for such machines, thus an environmental dependency on byte addressing does not restrict a Standard Program beyond its current de facto restrictions.

Transition/Conversion: The new **CHARS** and **CHAR+** address arithmetic operators should be used for programs that require portability to non-byte-addressed machines. The places where such conversion is necessary may be identified by searching for occurrences of words that accept a number of address units as an argument (e.g., **MOVE** , **ALLOT**).

D.6.4 Address increment for a cell is no longer two

As a consequence of Forth-83’s simultaneous specification of 16-bit stack width and byte addressing, the number two could reliably be used in address calculations involving memory arrays containing items from the stack. Since ANS Forth requires neither 16-bit stack width nor byte addressing, the number two is no longer necessarily appropriate for such calculations.

Words affected: @ ! +! 2+ 2* 2- +LOOP

Reason: See reasons for “Address Units” and “Stack Width”

Impact: In this respect, existing programs will continue to work on machines where a stack cell occupies two address units when stored in memory. This includes most machines for which Forth 83 compliant implementations currently exist. In principle, it would also include 16-bit-word-addressed machines with 32-bit stack width, but the committee knows of no examples of such machines.

Transition/Conversion: The new **CELLS** and **CELL+** address arithmetic operators should be used for portable programs. The places where such conversion is necessary may be identified by searching for the character “2” and determining whether or not it is used as part of an address calculation. The following substitutions are appropriate within address calculations:

Old	New
2+ or 2 +	CELL+
2* or 2 *	CELLS
2- or 2 -	1 CELLS -
2/ or 2 /	1 CELLS /
2	1 CELLS

The number “2” by itself is sometimes used for address calculations as an argument to **+LOOP**, when the loop index is an address. When converting the word **2/** which operates on negative dividends, one should be cognizant of the rounding method used.

D.6.5 Address alignment

Forth 83 imposes no restriction upon the alignment of addresses to any boundary. ANS Forth specifies that a Standard System may require alignment of addresses for use with various “@” and “!” operators.

Words Affected: ! +! 2! 2@ @ ? ,

Reason: Many computers have hardware restrictions that favor the use of aligned addresses. On some machines, the native memory-access instructions will cause an exception trap if used with an unaligned address. Even on machines where unaligned accesses do not cause exception traps, aligned accesses are usually faster.

Impact: All of the ANS Forth words that return addresses suitable for use with aligned “@” and “!” words must return aligned addresses. In most cases, there will be no problem. Problems can arise from the use of user-defined data structures containing a mixture of character data and cell-sized data.

Many existing Forth systems, especially those currently in use on computers with strong alignment requirements, already require alignment. Much existing Forth code that is currently in use on such machines has already been converted for use in an aligned environment.

Transition/Conversion: There are two possible approaches to conversion of programs for use on a system requiring address alignment.

The easiest approach is to redefine the system’s aligned “@” and “!” operators so that they do not require alignment. For example, on a 16-bit little-endian byte-addressed machine, unaligned “@” and “!” could be defined:

```
: @ ( addr -- x ) DUP C@ SWAP CHAR+ C@ 8 LSHIFT OR ;
: ! ( x addr -- ) OVER 8 RSHIFT OVER CHAR+ C! C! ;
```

These definitions, and similar ones for “+!”, “2@”, “2!”, “,”, and “?” as needed, can be compiled before an unaligned application, which will then work as expected.

This approach may conserve memory if the application uses substantial numbers of data structures containing unaligned fields.

Another approach is to modify the application’s source code to eliminate unaligned data fields. The ANS Forth words **ALIGN** and **ALIGNED** may be used to force alignment of data fields. The places where such alignment is needed may be determined by inspecting the parts of the application where data structures (other than simple variables) are defined, or by “smart compiler” techniques (see the “Smart Compiler” discussion below).

This approach will probably result in faster application execution speed, at the possible expense of increased memory utilization for data structures.

implementations, the choice was generally quite clear (with the single exception of the word **LEAVE**), and the standard could specify which words should be immediate. However, some of the currently popular implementation techniques, such as native-code generation with optimization, require the immediacy attribute on a different set of words than the set of immediate words of a threaded code implementation. ANS Forth, acknowledging the validity of these other implementation techniques, specifies the immediacy attribute in as few cases as possible.

When the membership of the set of immediate words is unclear, the decision about whether to use **COMPILE** or **[COMPILE]** becomes unclear. Consequently, ANS Forth provides a “general purpose” replacement word **POSTPONE** that serves the purpose of the vast majority of uses of both **COMPILE** and **[COMPILE]**, without requiring that the user know whether or not the “postponed” word is immediate.

Similarly, the use of **'** and **[']** with compiling words is unclear if the precise compilation behavior of those words is not specified, so ANS Forth does not permit a Standard Program to use **'** or **[']** with compiling words.

The traditional (non-immediate) definition of the word **COMPILE** has an additional problem. Its traditional definition assumes a threaded code implementation technique, and its behavior can only be properly described in that context. In the context of ANS Forth, which permits other implementation techniques in addition to threaded code, it is very difficult, if not impossible, to describe the behavior of the traditional **COMPILE**. Rather than changing its behavior, and thus breaking existing code, ANS Forth does not include the word **COMPILE**. This allows existing implementations to continue to supply the word **COMPILE** with its traditional behavior, if that is appropriate for the implementation.

Impact: **[COMPILE]** remains in ANS Forth, since its proper use does not depend on knowledge of whether or not a word is immediate. (Use of **[COMPILE]** with a non-immediate word is and has always been a no-op). Whether or not you need to use **[COMPILE]** requires knowledge of whether or not its target word is immediate, but it is always safe to use **[COMPILE]**. **[COMPILE]** is no longer in the (required) core word set, having been moved to the Core-Extensions word set, but the committee anticipates that most vendors will supply it anyway.

In nearly all cases, it is correct to replace both **[COMPILE]** and **COMPILE** with **POSTPONE**. Uses of **[COMPILE]** and **COMPILE** that are not suitable for “mindless” replacement by **POSTPONE** are quite infrequent, and fall into the following two categories:

- 1) Use of **[COMPILE]** with non-immediate words. This is sometimes done with the words **'** (tick, which was immediate in Forth 79 but not in Forth 83) and **LEAVE** (which was immediate in Forth 83 but not in Forth 79), in order to force the compilation of those words without regard to whether you are using a Forth 79 or Forth 83 system.
- 2) Use of the phrase “**COMPILE [COMPILE] <immediate word>**” to “doubly postpone” an immediate word.

Transition/Conversion: Many ANS Forth implementations will continue to implement both **[COMPILE]** and **COMPILE** in forms compatible with existing usage. In those environments, no conversion is necessary.

For complete portability, uses of **COMPILE** and **[COMPILE]** should be changed to **POSTPONE**, except in the rare cases indicated above. Uses of **[COMPILE]** with non-immediate words may be left as-is, and the program may declare a requirement for the word **[COMPILE]** from the Core Extensions word set, or the **[COMPILE]** before the non-immediate word may be simply deleted if the target word is known to be non-immediate.

Uses of the phrase “**COMPILE [COMPILE] <immediate-word>**” may be handled by introducing an “intermediate word” (**XX** in the example below) and then postponing that word. For example:

```
: ABC COMPILE [COMPILE] IF ;
```

changes to:

```
: XX POSTPONE IF ;
```

```
: ABC POSTPONE XX ;
```

A non-standard case can occur with programs that “switch out of compilation state” to explicitly compile a thread in the dictionary following a **COMPILE** . For example:

```
: XYZ COMPILE [ ' ABC , ] ;
```

This depends heavily on knowledge of exactly how **COMPILE** and the threaded-code implementation works. Cases like this cannot be handled mechanically; they must be translated by understanding exactly what the code is doing, and rewriting that section according to ANS Forth restrictions.

Use the phrase **POSTPONE [COMPILE]** to replace **[COMPILE] [COMPILE]**.

D.6.8 Input character set

Forth 83 specifies that the full 7-bit ASCII character set is available through **KEY** . ANS Forth restricts it to the graphic characters of the ASCII set, with codes from hex 20 to hex 7E inclusive.

Words Affected: **KEY**

Reason: Many system environments “consume” certain control characters for such purposes as input editing, job control, or flow control. A Forth implementation cannot always control this system behavior.

Impact: Standard Programs which require the ability to receive particular control characters through **KEY** must declare an environmental dependency on the input character set.

Transition/Conversion: For maximum portability, programs should restrict their required input character set to only the graphic characters. Control characters may be handled if available, but complete program functionality should be accessible using only graphic characters.

As stated above, an environmental dependency on the input character set may be declared. Even so, it is recommended that the program should avoid the requirement for particularly-troublesome control characters, such as control-S and control-Q (often used for flow control, sometimes by communication hardware whose presence may be difficult to detect), ASCII NUL (difficult to type on many keyboards), and the distinction between carriage return and line feed (some systems translate carriage returns into line feeds, or vice versa).

D.6.9 Shifting with **UM/MOD**

Given Forth-83’s two’s-complement nature, and its requirement for floored (round toward minus infinity) division, shifting is equivalent to division. Also, two’s-complement representation implies that unsigned division by a power of two is equivalent to logical right-shifting, so **UM/MOD** could be used to perform a logical right-shift.

Words Affected: **UM/MOD**

Reason: The problem with **UM/MOD** is a result of allowing non-two’s-complement number representations, as already described.

ANS Forth provides the words **LSHIFT** and **RSHIFT** to perform logical shifts. This is usually more efficient, and certainly more descriptive, than the use of **UM/MOD** for logical shifting.

Impact: Programs running on ANS Forth systems with two’s-complement arithmetic (the majority of machines), will not experience any incompatibility with **UM/MOD** . Existing Forth-83 Standard programs intended to run on non-two’s-complement machines will not be able to use **UM/MOD** for shifting on a non-two’s-complement ANS Forth system. This should not affect a significant number of existing programs (perhaps none at all), since the committee knows of no existing Forth-83 implementations on non-two’s-complement machines.

Transition/Conversion: A program that requires **UM/MOD** to behave as a shift operation may declare an environmental dependency on two’s-complement arithmetic.

A program that cannot declare an environmental dependency on two's-complement arithmetic may require editing to replace incompatible uses of **UM/MOD** with other operators defined within the application.

D.6.10 Vocabularies / wordlists

ANS Forth does not define the words **VOCABULARY**, **CONTEXT**, and **CURRENT**, which were present in Forth 83. Instead, ANS Forth defines a primitive word set for search order specification and control, including words which have not existed in any previous standard.

Forth-83's "**ALSO/ONLY**" experimental search order word set is specified for the most part as the extension portion of the ANS Forth Search Order word set.

Words Affected: **VOCABULARY CONTEXT CURRENT**

Reason: Vocabularies are an area of much divergence among existing systems. Considering major vendors' systems and previous standards, there are at least 5 different and mutually incompatible behaviors of words defined by **VOCABULARY**. Forth 83 took a step in the direction of "run-time search-order specification" by declining to specify a specific relationship between the hierarchy of compiled vocabularies and the run-time search order. Forth 83 also specified an experimental mechanism for run-time search-order specification, the **ALSO/ONLY** scheme. **ALSO/ONLY** was implemented in numerous systems, and has achieved some measure of popularity in the Forth community.

However, several vendors refuse to implement it, citing technical limitations. In an effort to address those limitations and thus hopefully make **ALSO/ONLY** more palatable to its critics, the committee specified a simple "primitive word set" that not only fixes some of the objections to **ALSO/ONLY**, but also provides sufficient power to implement **ALSO/ONLY** and all of the other search-order word sets that are currently popular.

The Forth 83 **ALSO/ONLY** word set is provided as an optional extension to the search-order word set. This allows implementors that are so inclined to provide this word set, with well-defined standard behavior, but does not compel implementors to do so. Some vendors have publicly stated that they will not implement **ALSO/ONLY**, no matter what, and one major vendor stated an unwillingness to implement ANS Forth at all if **ALSO/ONLY** is mandated. The committee feels that its actions are prudent, specifying **ALSO/ONLY** to the extent possible without mandating its inclusion in all systems, and also providing a primitive search-order word set that vendors may be more likely to implement, and which can be used to synthesize **ALSO/ONLY**.

Transition/Conversion: Since Forth 83 did not mandate precise semantics for **VOCABULARY**, existing Forth-83 Standard programs cannot use it except in a trivial way. Programs can declare a dependency on the existence of the Search Order word set, and can implement whatever semantics are required using that word set's primitives. Forth 83 programs that need **ALSO/ONLY** can declare a dependency on the Search Order Extensions word set, or can implement the extensions in terms of the Search Order word set itself.

D.6.11 Multiprogramming impact

Forth 83 marked words with "multiprogramming impact" by the letter "M" in the first lines of their descriptions. ANS Forth has removed the "M" designation from the word descriptions, moving the discussion of multiprogramming impact to this non-normative annex.

Words affected: none

Reason: The meaning of "multiprogramming impact" is precise only in the context of a specific model for multiprogramming. Although many Forth systems do provide multiprogramming capabilities using a particular round-robin, cooperative, block-buffer sharing model, that model is not universal. Even assuming the classical model, the "M" designations did not contain enough information to enable writing of applications that interacted in a multiprogrammed system.

Practically speaking, the “M” designations in Forth 83 served to document usage rules for block buffer addresses in multiprogrammed systems. These addresses often become meaningless after a task has relinquished the CPU for any reason, most often for the purposes of performing I/O, awaiting an event, or voluntarily sharing CPU resources using the word **PAUSE**. It was essential that portable applications respect those usage rules to make it practical to run them on multiprogrammed systems; failure to adhere to the rules could easily compromise the integrity of other applications running on those systems as well as the applications actually in error. Thus, “M” appeared on all words that by design gave up the CPU, with the understanding that other words NEVER gave it up.

These usage rules have been explicitly documented in the **BLOCK** word set where they are relevant. The “M” designations have been removed entirely.

Impact: In practice, none.

In the sense that any application that depends on multiprogramming must consist of at least two tasks that share some resource(s) and communicate between themselves, Forth 83 did not contain enough information to enable writing of a standard program that **DEPENDED** on multiprogramming. This is also true of ANS Forth.

Non-multiprogrammed applications in Forth 83 were required to respect usage rules for **BLOCK** so that they could be run properly on multiprogrammed systems. The same is true of ANS Forth.

The only difference is the documentation method used to define the **BLOCK** usage rules. The Technical Committee believes that the current method is clearer than the concept of “multiprogramming impact.”

Transition/Conversion: none needed.

D.6.12 Words not provided in executable form

ANS Forth allows an implementation to supply some words in source code or “load as needed” form, rather than requiring all supplied words to be available with no additional programmer action.

Words affected: all

Reason: Forth systems are often used in environments where memory space is at a premium. Every word included in the system in executable form consumes memory space. The committee believes that allowing standard words to be provided in source form will increase the probability that implementors will provide complete ANS Forth implementations even in systems designed for use in constrained environments.

Impact: In order to use a Standard Program with a given ANS Forth implementation, it may be necessary to precede the program with an implementation-dependent “preface” to make “source form” words executable. This is similar to the methods that other computer languages require for selecting the library routines needed by a particular application.

In languages like C, the goal of eliminating unnecessary routines from the memory image of an application is usually accomplished by providing libraries of routines, using a “linker” program to incorporate only the necessary routines into an executable application. The method of invoking and controlling the linker is outside the scope of the language definition.

Transition/Conversion: Before compiling a program, the programmer may need to perform some action to make the words required by that program available for execution.

E. ANS Forth portability guide (informative annex)

E.1 Introduction

The most popular architectures used to implement Forth have had byte-addressed memory, 16-bit operations, and two's-complement number representation. The Forth-83 Standard dictates that these particular features must be present in a Forth-83 Standard system and that Forth-83 programs may exploit these features freely.

However, there are many beasts in the architectural jungle that are bit addressed or cell addressed, or prefer 32-bit operations, or represent numbers in one's complement. Since one of Forth's strengths is its usefulness in "strange" environments on "unusual" hardware with "peculiar" features, it is important that a Standard Forth run on these machines too.

A primary goal of the ANS Forth Standard is to increase the types of machines that can support a Standard Forth. This is accomplished by allowing some key Forth terms to be implementation-defined (e.g., how big is a cell?) and by providing Forth operators (words) that conceal the implementation. This frees the implementor to produce the Forth system that most effectively utilizes the native hardware. The machine independent operators, together with some programmer discipline, enable a programmer to write Forth programs that work on a wide variety of machines.

The remainder of this Annex provides guidelines for writing portable ANS Forth programs. The first section describes ways to make a program hardware independent. It is difficult for someone familiar with only one machine architecture to imagine the problems caused by transporting programs between dissimilar machines. Consequently, examples of specific architectures with their respective problems are given. The second section describes assumptions about Forth implementations that many programmers make, but can't be relied upon in a portable program.

E.2 Hardware peculiarities

E.2.1 Data/memory abstraction

Data and memory are the stones and mortar of program construction. Unfortunately, each computer treats data and memory differently. The ANS Forth Systems Standard gives definitions of data and memory that apply to a wide variety of computers. These definitions give us a way to talk about the common elements of data and memory while ignoring the details of specific hardware. Similarly, ANS Forth programs that use data and memory in ways that conform to these definitions can also ignore hardware details. The following sections discuss the definitions and describe how to write programs that are independent of the data/memory peculiarities of different computers.

E.2.2 Definitions

Three terms defined by ANS Forth are address unit, cell, and character. The address space of an ANS Forth system is divided into an array of address units; an address unit is the smallest collection of bits that can be addressed. In other words, an address unit is the number of bits spanned by the addresses *addr* and *addr*+1. The most prevalent machines use 8-bit address units. Such "byte addressed" machines include the Intel 8086 and Motorola 68000 families. However, other address unit sizes exist. There are machines that are bit addressed and machines that are 4-bit nibble addressed. There are also machines with address units larger than 8-bits. For example, several Forth-in-hardware computers are cell addressed.

The cell is the fundamental data type of a Forth system. A cell can be a single-cell integer or a memory address. Forth's parameter and return stacks are stacks of cells. Forth 83 specifies that a cell is 16-bits. In ANS Forth the size of a cell is an implementation-defined number of address units. Thus, an ANS Forth implemented on a 16-bit microprocessor could use a 16-bit cell and an implementation on a 32-bit machine could use a 32-bit cell. Also 18-bit machines, 36-bit machines, etc., could support ANS Forth systems with 18 or 36-bit cells respectively. In all of these systems, **DUP** does the same thing: it

duplicates the top of the data stack. **!** (store) behaves consistently too: given two cells on the data stack it stores the second cell in the memory location designated by the top cell.

Similarly, the definition of a character has been generalized to be an implementation-defined number of address units (but at least eight bits). This removes the need for a Forth implementor to provide 8-bit characters on processors where it is inappropriate. For example, on an 18-bit machine with a 9-bit address unit, a 9-bit character would be most convenient. Since, by definition, you can't address anything smaller than an address unit, a character must be at least as big as an address unit. This will result in big characters on machines with large address units. An example is a 16-bit cell addressed machine where a 16-bit character makes the most sense.

E.2.3 Addressing memory

ANS Forth eliminates many portability problems by using the above definitions. One of the most common portability problems is addressing successive cells in memory. Given the memory address of a cell, how do you find the address of the next cell? In Forth 83 this is easy: **2 + .** This code assumes that memory is addressed in 8-bit units (bytes) and a cell is 16-bits wide. On a byte-addressed machine with 32-bit cells the code to find the next cell would be **4 + .** The code would be **1+** on a cell-addressed processor and **16 +** on a bit-addressed processor with 16-bit cells. ANS Forth provides a next-cell operator named **CELL+** that can be used in all of these cases. Given an address, **CELL+** adjusts the address by the size of a cell (measured in address units). A related problem is that of addressing an array of cells in an arbitrary order. A defining word to create an array of cells using Forth 83 would be:

```
: ARRAY CREATE 2* ALLOT DOES> SWAP 2* + ;
```

Use of **2*** to scale the array index assumes byte addressing and 16-bit cells again. As in the example above, different versions of the code would be needed for different machines. ANS Forth provides a portable scaling operator named **CELLS**. Given a number *n*, **CELLS** returns the number of address units needed to hold *n* cells. A portable definition of array is:

```
: ARRAY CREATE CELLS ALLOT  
DOES> SWAP CELLS + ;
```

There are also portability problems with addressing arrays of characters. In Forth 83 (and in the most common ANS Forth implementations), the size of a character will equal the size of an address unit. Consequently addresses of successive characters in memory can be found using **1+** and scaling indices into a character array is a no-op (i.e., **1 ***). However, there are cases where a character is larger than an address unit. Examples include (1) systems with small address units (e.g., bit- and nibble-addressed systems), and (2) systems with large character sets (e.g., 16-bit characters on a byte-addressed machine). **CHAR+** and **CHARS** operators, analogous to **CELL+** and **CELLS** are available to allow maximum portability.

ANS Forth generalizes the definition of some Forth words that operate on chunks of memory to use address units. One example is **ALLOT**. By prefixing **ALLOT** with the appropriate scaling operator (**CELLS**, **CHARS**, etc.), space for any desired data structure can be allocated (see definition of array above). For example:

```
CREATE ABUFFER 5 CHARS ALLOT ( allot 5 character buffer)
```

The memory-block-move word also uses address units:

```
source destination 8 CELLS MOVE ( move 8 cells)
```

E.2.4 Alignment problems

Not all addresses are created equal. Many processors have restrictions on the addresses that can be used by memory access instructions. This Standard does not require an implementor of an ANS Forth to make alignment transparent; on the contrary, it requires (in Section 3.3.3.1 **Address Alignment**) that an ANS Forth program assume that character and cell alignment may be required.

One of the most common problems caused by alignment restrictions is in creating tables containing both characters and cells. When `,` (comma) or `C,` is used to initialize a table, data is stored at the data-space pointer. Consequently, it must be suitably aligned. For example, a non-portable table definition would be:

```
CREATE ATABLE 1 C, X , 2 C, Y ,
```

On a machine that restricts 16-bit fetches to even addresses, `CREATE` would leave the data space pointer at an even address, the `1 C,` would make the data space pointer odd, and `,` (comma) would violate the address restriction by storing `X` at an odd address. A portable way to create the table is:

```
CREATE ATABLE 1 C, ALIGN X , 2 C, ALIGN Y ,
```

`ALIGN` adjusts the data space pointer to the first aligned address greater than or equal to its current address. An aligned address is suitable for storing or fetching characters, cells, cell pairs, or double-cell numbers.

After initializing the table, we would also like to read values from the table. For example, assume we want to fetch the first cell, `X`, from the table. `ATABLE CHAR+` gives the address of the first thing after the character. However this may not be the address of `X` since we aligned the dictionary pointer between the `C,` and the `,`. The portable way to get the address of `X` is:

```
ATABLE CHAR+ ALIGNED
```

`ALIGNED` adjusts the address on top of the stack to the first aligned address greater than or equal to its current value.

E.3 Number representation

Different computers represent numbers in different ways. An awareness of these differences can help a programmer avoid writing a program that depends on a particular representation.

E.3.1 Big endian vs. little endian

The constituent bits of a number in memory are kept in different orders on different machines. Some machines place the most-significant part of a number at an address in memory with less-significant parts following it at higher addresses. Other machines do the opposite — the least-significant part is stored at the lowest address. For example, the following code for a 16-bit 8086 “little endian” Forth would produce the answer 34 (hex):

```
VARIABLE FOO HEX 1234 FOO ! FOO C@
```

The same code on a 16-bit 68000 “big endian” Forth would produce the answer 12 (hex). A portable program cannot exploit the representation of a number in memory.

A related issue is the representation of cell pairs and double-cell numbers in memory. When a cell pair is moved from the stack to memory with `2!`, the cell that was on top of the stack is placed at the lower memory address. It is useful and reasonable to manipulate the individual cells when they are in memory.

E.3.2 ALU organization

Different computers use different bit patterns to represent integers. Possibilities include binary representations (two’s complement, one’s complement, sign magnitude, etc.) and decimal representations (BCD, etc.). Each of these formats creates advantages and disadvantages in the design of a computer’s arithmetic logic unit (ALU). The most commonly used representation, two’s complement, is popular because of the simplicity of its addition and subtraction algorithms.

Programmers who have grown up on two’s complement machines tend to become intimate with their representation of numbers and take some properties of that representation for granted. For example, a trick to find the remainder of a number divided by a power of two is to mask off some bits with `AND`. A common application of this trick is to test a number for oddness using `1 AND`. However, this will not work on a one’s complement machine if the number is negative (a portable technique is `2 MOD`).

The remainder of this section is a (non-exhaustive) list of things to watch for when portability between machines with binary representations other than two's complement is desired.

To convert a single-cell number to a double-cell number, ANS Forth provides the operator **S>D**. To convert a double-cell number to single-cell, Forth programmers have traditionally used **DROP**. However, this trick doesn't work on sign-magnitude machines. For portability a **D>S** operator is available. Converting an unsigned single-cell number to a double-cell number can be done portably by pushing a zero on the stack.

E.4 Forth system implementation

During Forth's history, an amazing variety of implementation techniques have been developed. The ANS Forth Standard encourages this diversity and consequently restricts the assumptions a user can make about the underlying implementation of an ANS Forth system. Users of a particular Forth implementation frequently become accustomed to aspects of the implementation and assume they are common to all Forths. This section points out many of these incorrect assumptions.

E.4.1 Definitions

Traditionally, Forth definitions have consisted of the name of the Forth word, a dictionary search link, data describing how to execute the definition, and parameters describing the definition itself. These components are called the name, link, code, and parameter fields¹. No method for accessing these fields has been found that works across all of the Forth implementations currently in use. Therefore, ANS Forth severely restricts how the fields may be used. Specifically, a portable ANS Forth program may not use the name, link, or code field in any way. Use of the parameter field (renamed to data field for clarity) is limited to the operations described below.

Only words defined with **CREATE** or with other defining words that call **CREATE** have data fields. The other defining words in the Standard (**VARIABLE**, **CONSTANT**, **:**, etc.) might not be implemented with **CREATE**. Consequently, a Standard Program must assume that words defined by **VARIABLE**, **CONSTANT**, **:**, etc., may have no data fields. There is no way for a Standard Program to modify the value of a constant or to change the meaning of a colon definition. The **DOES>** part of a defining word operates on a data field. Since only **CREATED** words have data fields, **DOES>** can only be paired with **CREATE** or words that call **CREATE**.

In ANS Forth, **FIND**, **[']** and **'** (tick) return an unspecified entity called an "execution token". There are only a few things that may be done with an execution token. The token may be passed to **EXECUTE** to execute the word ticked or compiled into the current definition with **COMPILE, .** The token can also be stored in a variable and used later. Finally, if the word ticked was defined via **CREATE**, **>BODY** converts the execution token into the word's data-field address.

One thing that definitely cannot be done with an execution token is use **!** or **,** to store it into the object code of a Forth definition. This technique is sometimes used in implementations where the object code is a list of addresses (threaded code) and an execution token is also an address. However, ANS Forth permits native code implementations where this will not work.

E.4.2 Stacks

In some Forth implementations, it is possible to find the address of a stack in memory and manipulate the stack as an array of cells. This technique is not portable, however. On some systems, especially Forth-in-hardware systems, the stacks might be in a part of memory that can't be addressed by the program or might not be in memory at all. Forth's parameter and return stacks must be treated as stacks.

A Standard Program may use the return stack directly only for temporarily storing values. Every value examined or removed from the return stack using **R@**, **R>**, or **2R>** must have been put on the stack

¹ These terms are not defined in the standard. They are mentioned here for historical continuity.

explicitly using **>R** or **2>R**. Even this must be done carefully since the system may use the return stack to hold return addresses and loop-control parameters. Section 3.2.3.3 of the Standard has a list of restrictions.

E.5 ROMed application disciplines and conventions

When a Standard System provides a data space which is uniformly readable and writeable we may term this environment “RAM-only”.

Programs designed for ROMed application must divide data space into at least two parts: a writeable and readable uninitialized part, called “RAM”, and a read-only initialized part, called “ROM”. A third possibility, a writeable and readable initialized part, normally called “initialized RAM”, is not addressed by this discipline. A Standard Program must explicitly initialize the RAM data space as needed.

The separation of data space into RAM and ROM is meaningful only during the generation of the ROMed program. If the ROMed program is itself a standard development system, it has the same taxonomy as an ordinary RAM-only system.

The words affected by conversion from a RAM-only to a mixed RAM and ROM environment are:

, (comma) **ALIGN** **ALIGNED** **ALLOT** **C**, **CREATE** **HERE** **UNUSED**
 (**VARIABLE** always accesses the RAM data space.)

With the exception of **,** (comma) and **C**, these words are meaningful in both RAM and ROM data space.

To select the data space, these words could be preceded by selectors **RAM** and **ROM**. For example:

ROM **CREATE** **ONES** **32** **ALLOT** **ONES** **32** **1** **FILL** **RAM**

would create a table of ones in the ROM data space. The storage of data into RAM data space when generating a program for ROM would be an ambiguous condition.

A straightforward implementation of these selectors would maintain separate address counters for each space. A counter value would be returned by **HERE** and altered by **,** (comma), **C**, **ALIGN**, and **ALLOT**, with **RAM** and **ROM** simply selecting the appropriate address counter. This technique could be extended to additional partitions of the data space.

E.6 Summary

The ANS Forth Standard cannot and should not force anyone to write a portable program. In situations where performance is paramount, the programmer is encouraged to use every trick in the book. On the other hand, if portability to a wide variety of systems is needed, ANS Forth provides the tools to accomplish this. There is probably no such thing as a completely portable program. A programmer, using this guide, should intelligently weigh the tradeoffs of providing portability to specific machines. For example, machines that use sign-magnitude numbers are rare and probably don't deserve much thought. But, systems with different cell sizes will certainly be encountered and should be provided for. In general, making a program portable clarifies both the programmer's thinking process and the final program.

F. Alphabetic list of words (informative annex)

In the following list, the last, four-digit, part of the reference number establishes a sequence corresponding to the alphabetic ordering of all standard words. The first two or three parts indicate the word set and glossary section in which the word is defined.

6.1.0010	!	“store”	CORE	25
6.1.0030	#	“number-sign”	CORE	25
6.1.0040	#>	“number-sign-greater”	CORE	25
6.1.0050	#S	“number-sign-s”	CORE	25
6.2.0060	#TIB	“number-t-i-b”	CORE EXT	51
6.1.0070	'	“tick”	CORE	25
6.1.0080	(.....	“paren”	CORE	26
11.6.1.0080	(.....	“paren”	FILE	82
13.6.1.0086	(LOCAL)	“Paren-local-paren”	LOCAL	108
6.1.0090	*	“star”	CORE	26
6.1.0100	*/	“star-slash”	CORE	26
6.1.0110	*/MOD	“star-slash-mod”	CORE	26
6.1.0120	+	“plus”	CORE	26
6.1.0130	+	“plus-store”	CORE	27
6.1.0140	+LOOP	“plus-loop”	CORE	27
6.1.0150	,	“comma”	CORE	27
6.1.0160	-	“minus”	CORE	27
17.6.1.0170	-TRAILING	“dash-trailing”	STRING	126
6.1.0180	“dot”	CORE	27
6.1.0190	."	“dot-quote”	CORE	28
6.2.0200	.(.....	“dot-paren”	CORE EXT	51
6.2.0210	.R	“dot-r”	CORE EXT	51
15.6.1.0220	.S	“dot-s”	TOOLS	115
6.1.0230	/	“slash”	CORE	28
6.1.0240	/MOD	“slash-mod”	CORE	28
17.6.1.0245	/STRING	“slash-string”	STRING	126
6.1.0250	0<	“zero-less”	CORE	28
6.2.0260	0<>	“zero-not-equals”	CORE EXT	51
6.1.0270	0=	“zero-equals”	CORE	28
6.2.0280	0>	“zero-greater”	CORE EXT	51
6.1.0290	1+	“one-plus”	CORE	28
6.1.0300	1-	“one-minus”	CORE	29
6.1.0310	2!	“two-store”	CORE	29
6.1.0320	2*	“two-star”	CORE	29
6.1.0330	2/	“two-slash”	CORE	29
6.2.0340	2>R	“two-to-r”	CORE EXT	51
6.1.0350	2@	“two-fetch”	CORE	29
8.6.1.0360	2CONSTANT	“two-constant”	DOUBLE	68
6.1.0370	2DROP	“two-drop”	CORE	29
6.1.0380	2DUP	“two-dupe”	CORE	29
8.6.1.0390	2LITERAL	“two-literal”	DOUBLE	68
6.1.0400	2OVER	“two-over”	CORE	29
6.2.0410	2R>	“two-r-from”	CORE EXT	52
6.2.0415	2R@	“two-r-fetch”	CORE EXT	52
8.6.2.0420	2ROT	“two-rote”	DOUBLE EXT	71
6.1.0430	2SWAP	“two-swap”	CORE	30
8.6.1.0440	2VARIABLE	“two-variable”	DOUBLE	69

6.1.0450	:	“colon”	CORE	30
6.2.0455	:NONAME	“colon-no-name”	CORE EXT	52
6.1.0460	;	“semicolon”	CORE	30
15.6.2.0470	;CODE	“semicolon-code”	TOOLS EXT	116
6.1.0480	<	“less-than”	CORE	30
6.1.0490	<#	“less-number-sign”	CORE	31
6.2.0500	<>	“not-equals”	CORE EXT	52
6.1.0530	=	“equals”	CORE	31
6.1.0540	>	“greater-than”	CORE	31
6.1.0550	>BODY	“to-body”	CORE	31
12.6.1.0558	>FLOAT	“to-float”	FLOATING	93
6.1.0560	>IN	“to-in”	CORE	31
6.1.0570	>NUMBER	“to-number”	CORE	31
6.1.0580	>R	“to-r”	CORE	32
15.6.1.0600	?	“question”	TOOLS	115
6.2.0620	?DO	“question-do”	CORE EXT	53
6.1.0630	?DUP	“question-dupe”	CORE	32
6.1.0650	@	“fetch”	CORE	32
6.1.0670	ABORT		CORE	32
9.6.2.0670	ABORT		EXCEPTION EXT	75
6.1.0680	ABORT"	“abort-quote”	CORE	32
9.6.2.0680	ABORT"	“abort-quote”	EXCEPTION EXT	75
6.1.0690	ABS	“abs”	CORE	32
6.1.0695	ACCEPT		CORE	33
6.2.0700	AGAIN		CORE EXT	53
15.6.2.0702	AHEAD		TOOLS EXT	116
6.1.0705	ALIGN		CORE	33
6.1.0706	ALIGNED		CORE	33
14.6.1.0707	ALLOCATE		MEMORY	111
6.1.0710	ALLOT		CORE	33
16.6.2.0715	ALSO		SEARCH EXT	123
6.1.0720	AND		CORE	33
15.6.2.0740	ASSEMBLER		TOOLS EXT	117
10.6.1.0742	AT-XY	“at-x-y”	FACILITY	77
6.1.0750	BASE		CORE	34
6.1.0760	BEGIN		CORE	34
11.6.1.0765	BIN		FILE	82
6.1.0770	BL	“b-l”	CORE	34
17.6.1.0780	BLANK		STRING	126
7.6.1.0790	BLK	“b-l-k”	BLOCK	63
7.6.1.0800	BLOCK		BLOCK	63
7.6.1.0820	BUFFER		BLOCK	64
15.6.2.0830	BYE		TOOLS EXT	117
6.1.0850	C!	“c-store”	CORE	34
6.2.0855	C"	“c-quote”	CORE EXT	53
6.1.0860	C,	“c-comma”	CORE	34
6.1.0870	C@	“c-fetch”	CORE	34
6.2.0873	CASE		CORE EXT	54
9.6.1.0875	CATCH		EXCEPTION	74
6.1.0880	CELL+	“cell-plus”	CORE	35
6.1.0890	CELLS		CORE	35
6.1.0895	CHAR	“char”	CORE	35
6.1.0897	CHAR+	“char-plus”	CORE	35
6.1.0898	CHARS	“chars”	CORE	35

11.6.1.0900	CLOSE-FILE	FILE	82
17.6.1.0910	CMOVE	"c-move"	STRING126
17.6.1.0920	CMOVE>	"c-move-up"	STRING126
15.6.2.0930	CODE	TOOLS EXT	117
17.6.1.0935	COMPARE	STRING	127
6.2.0945	COMPILE,	"compile-comma"	CORE EXT54
6.1.0950	CONSTANT	CORE	35
6.2.0970	CONVERT	CORE EXT	54
6.1.0980	COUNT	CORE	36
6.1.0990	CR	"c-r"	CORE36
6.1.1000	CREATE	CORE	36
11.6.1.1010	CREATE-FILE	FILE	83
15.6.2.1015	CS-PICK	"c-s-pick"	TOOLS EXT117
15.6.2.1020	CS-ROLL	"c-s-roll"	TOOLS EXT118
8.6.1.1040	D+	"d-plus"	DOUBLE69
8.6.1.1050	D-	"d-minus"	DOUBLE69
8.6.1.1060	D.	"d-dot"	DOUBLE69
8.6.1.1070	D.R	"d-dot-r"	DOUBLE69
8.6.1.1075	D0<	"d-zero-less"	DOUBLE69
8.6.1.1080	D0=	"d-zero-equals"	DOUBLE69
8.6.1.1090	D2*	"d-two-star"	DOUBLE70
8.6.1.1100	D2/	"d-two-slash"	DOUBLE70
8.6.1.1110	D<	"d-less-than"	DOUBLE70
8.6.1.1120	D=	"d-equals"	DOUBLE70
12.6.1.1130	D>F	"d-to-f"	FLOATING93
8.6.1.1140	D>S	"d-to-s"	DOUBLE70
8.6.1.1160	DABS	"d-abs"	DOUBLE70
6.1.1170	DECIMAL	CORE	36
16.6.1.1180	DEFINITIONS	SEARCH	122
11.6.1.1190	DELETE-FILE	FILE	83
6.1.1200	DEPTH	CORE	36
12.6.2.1203	DF!	"d-f-store"	FLOATING EXT97
12.6.2.1204	DF@	"d-f-fetch"	FLOATING EXT98
12.6.2.1205	DFALIGN	"d-f-align"	FLOATING EXT98
12.6.2.1207	DFALIGNED	"d-f-aligned"	FLOATING EXT98
12.6.2.1208	DFLOAT+	"d-float-plus"	FLOATING EXT98
12.6.2.1209	DFLOATS	"d-floats"	FLOATING EXT98
8.6.1.1210	DMAX	"d-max"	DOUBLE70
8.6.1.1220	DMIN	"d-min"	DOUBLE70
8.6.1.1230	DNEGATE	"d-negate"	DOUBLE71
6.1.1240	DO	CORE	37
6.1.1250	DOES>	"does"	CORE37
6.1.1260	DROP	CORE	37
8.6.2.1270	DU<	"d-u-less"	DOUBLE EXT71
15.6.1.1280	DUMP	TOOLS	115
6.1.1290	DUP	"dupe"	CORE38
15.6.2.1300	EDITOR	TOOLS EXT	118
10.6.2.1305	EKEY	"e-key"	FACILITY EXT78
10.6.2.1306	EKEY>CHAR	"e-key-to-char"	FACILITY EXT78
10.6.2.1307	EKEY?	"e-key-question"	FACILITY EXT78
6.1.1310	ELSE	CORE	38
6.1.1320	EMIT	CORE	38
10.6.2.1325	EMIT?	"emit-question"	FACILITY EXT78
7.6.2.1330	EMPTY-BUFFERS	BLOCK EXT	65

Collating Sequence: !"#\$%&'()*+,-./digits ; < = > ? @ ALPHA [\] ^ _ ` alpha { | } ~

6.2.1342	ENDCASE	“end-case”	CORE EXT	54
6.2.1343	ENDOF	“end-of”	CORE EXT	55
6.1.1345	ENVIRONMENT?	“environment-query”	CORE	38
6.2.1350	ERASE		CORE EXT	55
6.1.1360	EVALUATE		CORE	39
7.6.1.1360	EVALUATE		BLOCK	64
6.1.1370	EXECUTE		CORE	39
6.1.1380	EXIT		CORE	39
6.2.1390	EXPECT		CORE EXT	55
12.6.1.1400	F!	“f-store”	FLOATING	93
12.6.1.1410	F*	“f-star”	FLOATING	93
12.6.2.1415	F**	“f-star-star”	FLOATING EXT	98
12.6.1.1420	F+	“f-plus”	FLOATING	93
12.6.1.1425	F-	“f-minus”	FLOATING	93
12.6.2.1427	F.	“f-dot”	FLOATING EXT	99
12.6.1.1430	F/	“f-slash”	FLOATING	94
12.6.1.1440	F0<	“f-zero-less-than”	FLOATING	94
12.6.1.1450	F0=	“f-zero-equals”	FLOATING	94
12.6.1.1460	F<	“f-less-than”	FLOATING	94
12.6.1.1470	F>D	“f-to-d”	FLOATING	94
12.6.1.1472	F@	“f-fetch”	FLOATING	94
12.6.2.1474	FABS	“f-abs”	FLOATING EXT	99
12.6.2.1476	FACOS	“f-a-cos”	FLOATING EXT	99
12.6.2.1477	FACOSH	“f-a-cosh”	FLOATING EXT	99
12.6.1.1479	FALIGN	“f-align”	FLOATING	94
12.6.1.1483	FALIGNED	“f-aligned”	FLOATING	94
12.6.2.1484	FALOG	“f-a-log”	FLOATING EXT	99
6.2.1485	FALSE		CORE EXT	55
12.6.2.1486	FASIN	“f-a-sine”	FLOATING EXT	99
12.6.2.1487	FASINH	“f-a-cinch”	FLOATING EXT	99
12.6.2.1488	FATAN	“f-a-tan”	FLOATING EXT	100
12.6.2.1489	FATAN2	“f-a-tan-two”	FLOATING EXT	100
12.6.2.1491	FATANH	“f-a-tan-h”	FLOATING EXT	100
12.6.1.1492	FCONSTANT	“f-constant”	FLOATING	95
12.6.2.1493	FCOS	“f-cos”	FLOATING EXT	100
12.6.2.1494	FCOSH	“f-cosh”	FLOATING EXT	100
12.6.1.1497	FDEPTH	“f-depth”	FLOATING	95
12.6.1.1500	FDROP	“f-drop”	FLOATING	95
12.6.1.1510	FDUP	“f-dupe”	FLOATING	95
12.6.2.1513	FE.	“f-e-dot”	FLOATING EXT	100
12.6.2.1515	FEXP	“f-e-x-p”	FLOATING EXT	100
12.6.2.1516	FEXPM1	“f-e-x-p-m-one”	FLOATING EXT	101
11.6.1.1520	FILE-POSITION		FILE	83
11.6.1.1522	FILE-SIZE		FILE	83
11.6.2.1524	FILE-STATUS		FILE EXT	88
6.1.1540	FILL		CORE	39
6.1.1550	FIND		CORE	39
16.6.1.1550	FIND		SEARCH	122
12.6.1.1552	FLITERAL	“f-literal”	FLOATING	95
12.6.2.1553	FLN	“f-l-n”	FLOATING EXT	101
12.6.2.1554	FLNP1	“f-l-n-p-one”	FLOATING EXT	101
12.6.1.1555	FLOAT+	“float-plus”	FLOATING	95
12.6.1.1556	FLOATS		FLOATING	96
12.6.2.1557	FLOG	“f-log”	FLOATING EXT	101

12.6.1.1558	FLOOR	FLOATING96
7.6.1.1559	FLUSH	BLOCK64
11.6.2.1560	FLUSH-FILE	FILE EXT88
6.1.1561	FM/MOD“f-m-slash-mod”	CORE40
12.6.1.1562	FMAX“f-max”	FLOATING96
12.6.1.1565	FMIN“f-min”	FLOATING96
12.6.1.1567	FNEGATE“f-negate”	FLOATING96
15.6.2.1580	FORGET	TOOLS EXT118
16.6.2.1590	FORTH	SEARCH EXT123
16.6.1.1595	FORTH-WORDLIST	SEARCH122
12.6.1.1600	FOVER“f-over”	FLOATING96
14.6.1.1605	FREE	MEMORY112
12.6.1.1610	FROT“f-rote”	FLOATING96
12.6.1.1612	FROUND“f-round”	FLOATING96
12.6.2.1613	FS.“f-s-dot”	FLOATING EXT101
12.6.2.1614	FSIN“f-sine”	FLOATING EXT101
12.6.2.1616	FSINCOS“f-sine-cos”	FLOATING EXT102
12.6.2.1617	FSINH“f-cinch”	FLOATING EXT102
12.6.2.1618	FSQRT“f-square-root”	FLOATING EXT102
12.6.1.1620	FSWAP“f-swap”	FLOATING96
12.6.2.1625	FTAN“f-tan”	FLOATING EXT102
12.6.2.1626	FTANH“f-tan-h”	FLOATING EXT102
12.6.1.1630	FVARIABLE“f-variable”	FLOATING97
12.6.2.1640	F~“f-proximate”	FLOATING EXT102
16.6.1.1643	GET-CURRENT	SEARCH122
16.6.1.1647	GET-ORDER	SEARCH122
6.1.1650	HERE	CORE40
6.2.1660	HEX	CORE EXT55
6.1.1670	HOLD	CORE40
6.1.1680	I	CORE40
6.1.1700	IF	CORE40
6.1.1710	IMMEDIATE	CORE41
11.6.1.1717	INCLUDE-FILE	FILE83
11.6.1.1718	INCLUDED	FILE84
6.1.1720	INVERT	CORE41
6.1.1730	J	CORE41
6.1.1750	KEY	CORE41
10.6.1.1755	KEY?“key-question”	FACILITY77
6.1.1760	LEAVE	CORE41
7.6.2.1770	LIST	BLOCK EXT65
6.1.1780	LITERAL	CORE42
7.6.1.1790	LOAD	BLOCK64
13.6.2.1795	LOCALS“locals-bar”	LOCAL EXT109
6.1.1800	LOOP	CORE42
6.1.1805	LSHIFT“l-shift”	CORE42
6.1.1810	M*“m-star”	CORE42
8.6.1.1820	M*/“m-star-slash”	DOUBLE71
8.6.1.1830	M+“m-plus”	DOUBLE71
6.2.1850	MARKER	CORE EXT56
6.1.1870	MAX	CORE42
6.1.1880	MIN	CORE42
6.1.1890	MOD	CORE43
6.1.1900	MOVE	CORE43
10.6.2.1905	MS	FACILITY EXT78

6.1.1910	NEGATE	CORE	43
6.2.1930	NIP	CORE EXT	56
6.2.1950	OF	CORE EXT	56
16.6.2.1965	ONLY	SEARCH EXT	124
11.6.1.1970	OPEN-FILE	FILE	84
6.1.1980	OR	CORE	43
16.6.2.1985	ORDER	SEARCH EXT	124
6.1.1990	OVER	CORE	43
6.2.2000	PAD	CORE EXT	56
10.6.1.2005	PAGE	FACILITY	77
6.2.2008	PARSE	CORE EXT	57
6.2.2030	PICK	CORE EXT	57
6.1.2033	POSTPONE	CORE	43
12.6.2.2035	PRECISION	FLOATING EXT	102
16.6.2.2037	PREVIOUS	SEARCH EXT	124
6.2.2040	QUERY	CORE EXT	57
6.1.2050	QUIT	CORE	44
11.6.1.2054	R/O "r-o"	FILE	84
11.6.1.2056	R/W "r-w"	FILE	85
6.1.2060	R> "r-from"	CORE	44
6.1.2070	R@ "r-fetch"	CORE	44
11.6.1.2080	READ-FILE	FILE	85
11.6.1.2090	READ-LINE	FILE	86
6.1.2120	RECURSE	CORE	44
6.2.2125	REFILL	CORE EXT	57
7.6.2.2125	REFILL	BLOCK EXT	65
11.6.2.2125	REFILL	FILE EXT	88
11.6.2.2130	RENAME-FILE	FILE EXT	88
6.1.2140	REPEAT	CORE	45
11.6.1.2142	REPOSITION-FILE	FILE	86
12.6.1.2143	REPRESENT	FLOATING	97
14.6.1.2145	RESIZE	MEMORY	112
11.6.1.2147	RESIZE-FILE	FILE	86
6.2.2148	RESTORE-INPUT	CORE EXT	57
6.2.2150	ROLL	CORE EXT	58
6.1.2160	ROT "rote"	CORE	45
6.1.2162	RSHIFT "r-shift"	CORE	45
6.1.2165	S" "s-quote"	CORE	45
11.6.1.2165	S" "s-quote"	FILE	87
6.1.2170	S>D "s-to-d"	CORE	45
7.6.1.2180	SAVE-BUFFERS	BLOCK	64
6.2.2182	SAVE-INPUT	CORE EXT	58
7.6.2.2190	SCR "s-c-r"	BLOCK EXT	65
17.6.1.2191	SEARCH	STRING	127
16.6.1.2192	SEARCH-WORDLIST	SEARCH	123
15.6.1.2194	SEE	TOOLS	115
16.6.1.2195	SET-CURRENT	SEARCH	123
16.6.1.2197	SET-ORDER	SEARCH	123
12.6.2.2200	SET-PRECISION	FLOATING EXT	102
12.6.2.2202	SF! "s-f-store"	FLOATING EXT	103
12.6.2.2203	SF@ "s-f-fetch"	FLOATING EXT	103
12.6.2.2204	SFALIGN "s-f-align"	FLOATING EXT	103
12.6.2.2206	SFALIGNED "s-f-aligned"	FLOATING EXT	103
12.6.2.2207	SFLOAT+ "s-float-plus"	FLOATING EXT	103

12.6.2.2208	SFLOATS "s-floats"	FLOATING EXT	104
6.1.2210	SIGN	CORE	46
17.6.1.2212	SLITERAL	STRING	127
6.1.2214	SM/REM "s-m-slash-rem"	CORE	46
6.1.2216	SOURCE	CORE	46
6.2.2218	SOURCE-ID "source-i-d"	CORE EXT	58
11.6.1.2218	SOURCE-ID "source-i-d"	FILE	87
6.1.2220	SPACE	CORE	46
6.1.2230	SPACES	CORE	46
6.2.2240	SPAN	CORE EXT	58
6.1.2250	STATE	CORE	46
15.6.2.2250	STATE	TOOLS EXT	118
6.1.2260	SWAP	CORE	46
6.1.2270	THEN	CORE	47
9.6.1.2275	THROW	EXCEPTION	75
7.6.2.2280	THRU	BLOCK EXT	65
6.2.2290	TIB "t-i-b"	CORE EXT	58
10.6.2.2292	TIME&DATE "time-and-date"	FACILITY EXT	78
6.2.2295	TO	CORE EXT	59
13.6.1.2295	TO	LOCAL	109
6.2.2298	TRUE	CORE EXT	59
6.2.2300	TUCK	CORE EXT	59
6.1.2310	TYPE	CORE	47
6.1.2320	U. "u-dot"	CORE	47
6.2.2330	U.R "u-dot-r"	CORE EXT	59
6.1.2340	U< "u-less-than"	CORE	47
6.2.2350	U> "u-greater-than"	CORE EXT	59
6.1.2360	UM* "u-m-star"	CORE	47
6.1.2370	UM/MOD "u-m-slash-mod"	CORE	48
6.1.2380	UNLOOP	CORE	48
6.1.2390	UNTIL	CORE	48
6.2.2395	UNUSED	CORE EXT	59
7.6.1.2400	UPDATE	BLOCK	65
6.2.2405	VALUE	CORE EXT	60
6.1.2410	VARIABLE	CORE	48
11.6.1.2425	W/O "w-o"	FILE	87
6.1.2430	WHILE	CORE	49
6.2.2440	WITHIN	CORE EXT	60
6.1.2450	WORD	CORE	49
16.6.1.2460	WORDLIST	SEARCH	123
15.6.1.2465	WORDS	TOOLS	116
11.6.1.2480	WRITE-FILE	FILE	87
11.6.1.2485	WRITE-LINE	FILE	88
6.1.2490	XOR "x-or"	CORE	49
6.1.2500	[..... "left-bracket"	CORE	49
6.1.2510	['] "bracket-tick"	CORE	50
6.1.2520	[CHAR] "bracket-char"	CORE	50
6.2.2530	[COMPILE] "bracket-compile"	CORE EXT	60
15.6.2.2531	[ELSE] "bracket-else"	TOOLS EXT	119
15.6.2.2532	[IF] "bracket-if"	TOOLS EXT	119
15.6.2.2533	[THEN] "bracket-then"	TOOLS EXT	119
6.2.2535	\ "backslash"	CORE EXT	60
7.6.2.2535	\ "backslash"	BLOCK EXT	66
6.1.2540] "right-bracket"	CORE	50

Collating Sequence: ! " # \$ % & ' () * + , - . / digits ; : < = > ? @ ALPHA [\] ^ _ ` alpha { | } ~