

NAME

snobol4 – SNOBOL4 interpreter

SYNOPSIS

snobol4 [*options* ...] [*file(s)* ...] [*params* ...]

DESCRIPTION

This manual page describes a port of the original Bell Telephone Labs (BTL) Macro Implementation of SNOBOL4 (MAINBOL) to machines with ILP32 (32-bit int/long/pointer) or LP64 (64-bit long/pointer) C compilers by Philip L. Budne. The language and its implementation are described in [1] and [2]. Extensions from Catspaw SNOBOL4+, SPITBOL and SITBOL have been added. This page discusses only the changes/extensions.

Limitations

All aspects of the language are implemented except;

- Trapping of arithmetic exceptions.
- **LOAD()** can be used to access external functions on most platforms, but not all. External functions can be statically linked (poor man's loading) into the snobol4 executable **ALL** platforms. See /usr/local/lib/snobol4/load.txt for more information.

Changes

The following behaviors have been changed from the original Macro SNOBOL4;

- Listings are disabled by default. Default listing side (when enabled by **-LIST** or the **-l** command option is **LEFT**. Listings are directed to standard output.
- Error messages, the startup banner and statistics are directed to standard error. Compilation error messages (including erroneous lines) appear on standard error as well as in the listings. Error messages now reference the source file name and line number.
- Character set (see below).
- The **PUNCH** output variable no longer exists (see **TERMINAL** variable below).
- I/O is not performed using FORTRAN I/O. The 3rd argument to the **OUTPUT()** and **INPUT()** functions are interpreted as a string of I/O options (see below).
- Control lines and comment characters are valid after the end of string (;) statement separator. Listing statement numbers show the statement number of the **LAST** statement on the line (rather than the first).
- Setting the **&ABEND** keyword causes a core dump upon termination!
- The value of the **&CODE** keyword determines the exit status of the **snobol4** application.
- The **DATE()** function returns strings of the form: *MM/DD/YYYY HH:MM:SS*. See extensions section for arguments to the **DATE()** function.
- Keyword **&STLIMIT** now defaults to -1. When **&STLIMIT** is less than zero there is no limit to the number of statements executed, and **&STCOUNT** is not incremented.
- **VALUE** tracing applies to variables modified by immediate value assignment (**\$** operator) and value assignment (**.** operator) during pattern matching.
- The **BACKSPACE()** function is not implemented. Use the **SET()** function instead.
- I/O unit numbers up to 256 can be used.
- Attempts to output on a closed unit generates a fatal "Output error"

Character set

snobol4(1) is 8-bit clean, and uses the native character set. Any 8-bit byte is accepted as a SNOBOL datum or in a string constant of a SNOBOL source program. The value of the SNOBOL protected keyword **&ALPHABET** is a 256-character string of all bytes from 0 to 255, in ascending order.

On ASCII based systems, any character with the 8th bit set is treated as "alphabetic", and can start, or be used in identifiers and labels. This includes characters from the "upper half" of national character sets and all bytes resulting from the UTF-8 encoding of Unicode characters.

Programs may be entered in mixed case; By default lower case identifiers are folded to upper case (see **&CASE** and **-CASE** extensions below). Case folding is performed by using the C library **islower**(3) test, and then using **toupper**(3) to convert the lower-case characters to upper case. When using UTF-8 encoded characters in code, case folding should be disabled, to prevent any bytes which appear to be lower case in the current locale from being modified.

The following operator character sequences are permitted and represent a cross between PDP-10 Macro SNOBOL (a.k.a. DECBOL), SITBOL and Catspaw SPITBOL usage;

Exponentiation:	^ **
Alternation:	!
Unary negation:	~ \
Assignment:	= _
Comment line:	* # ; !
Continuation line:	+ .

Both square brackets ([]) and angle brackets (<>) may be used to subscript arrays and tables. The **TAB** (ASCII 9) character is accepted as whitespace. Note that the use of the pound sign for comments allows use of the shell interpreter sequence at the top of a file (ie; "#!/usr/local/bin/snobol4 -b"). Underscore (_) and period (.) are legal within identifiers and labels.

Extensions

ARRAY/TABLE access

Multiple **ARRAY** and/or **TABLE** index operations may appear in a row, without having to resort to use of the **ELEMENT** function, so long as no intervening spaces (or line continuations) appear.

BREAKX()

The **BREAKX()** function is a pattern function used for fast scanning. **BREAKX(str)** is equivalent to **BREAK(str) ARBNO(LEN(1) BREAK(str))**. In other words **BREAKX** matches a sequence of ever larger strings terminated by a break set. **BREAKX** can be used as a faster matching replacement for **ARB; BREAKX('S') 'STRING'** always runs faster than **ARB 'STRING'** since it only attempts matching **'STRING'** at locations where an **'S'** has been detected.

Case folding

By default the compiler folds identifiers and directives to upper case, so programs can be entered in either case. To disable case folding use the directive **-CASE 0** or **-CASE**. To re-enable case folding use directive **-CASE n** where *n* is a non-zero integer. The status of case folding may be examined and controlled from a running program by the unprotected system keyword **&CASE**.

CHAR()

The **CHAR()** function takes an integer from 0 to 255 and returns the *n*'th character in **&ALPHABET**.

DATE()

For compatibility with new versions of Catspaw SPITBOL **DATE(0)** returns strings of the form *MM/DD/YY HH:MM:SS*, and **DATE(2)** returns strings of the form *YYYY-MM-YY HH:MM:SS*. With any other arguments **DATE()** returns strings of the form *MM/DD/YYYY HH:MM:SS*.

-ERROR/-NOERRORS

Directives **-ERROR** and **-NOERRORS** control execution of program with compiler errors. If the **-ERROR** directive is given, the program will be executed (but any attempt to execute a statement with a compiler error will cause a fatal execution error). By default programs with compiler errors will not be started, this can be restored using **-NOERRORS**.

&ERRTEXT

After an otherwise fatal error is curtailed due to a non-zero value in **&ERRLIMIT**, the protected keyword **&ERRTEXT** will contain the error message.

-EXECUTE/-NOEXECUTE

Directives **-EXECUTE** and **-NOEXECUTE** control execution of programs. If the **-NOEXECUTE** directive is given, the program will be not executed after compilation. **-EXECUTE** Cancels any previous **-NOEXECUTE**.

FREEZE()/THAW()

The **FREEZE()** function prohibits creation of new entries in the referenced table. This is useful once a table has been initialized to avoid creating empty entries on lookups that fail. This can greatly improve program speed, since frozen tables will not become clogged with empty entries. Lookups for uninitialized entries will return the null string. Attempts to assign to a non-existent entry will cause a "Variable not present where required" error. The **THAW()** function restores normal entry creation behavior.

FUNCTION()

The **FUNCTION()** predicate evaluates it's argument as a string (with case folding), and returns the null string if a function with that name exists and fails if it does not. The **FUNCTION()** predicate exists in SITBOL.

>RACE

If set to a non-zero value, each time a garbage collection is run, a trace message is output indicating the source file and line number of the current statement, how long the GC took, and how many units of storage are now free.

HOST()

A limited simulation of the SPITBOL **HOST()** function (with liberal implementation specific extensions) is included.

The **-INCLUDE** file **host.sno** contains symbolic defines for these (and many other) function codes.

HOST() with no parameters returns a string describing the system the program is running on. The string contains three parts, separated by colons. The first part describes the physical architecture, the second describes the operating system, and the third describes the language implementation name. Example: *Power Macintosh:Darwin 7.9.0:CSNOBOLA 1.1*

HOST(0) returns a string containing the command line parameter supplied to the **-u** option, if any. If no **-u** option was given, **HOST(0)** returns the concatenation of all user parameters following the input filename(s).

HOST(1,string) passes the string to the **system(3)** c library function, and returns the subprocess exit status.

HOST(2,n) for integer *n* returns the *n*'th command line argument (regardless of whether the argument was the command name, an option, a filename or a user parameter) as a string, or failure if *n* is out of range.

HOST(3) returns an integer for use with **HOST(2)** indicating the first command line argument available as a user parameter.

HOST(4,string) returns the value of the environment variable named *string*.

-INCLUDE

The **-INCLUDE** directive causes the compiler to interpolate the contents of the named file enclosed in single or double quotes. Any filename will be included only once, this can be overridden by appending a trailing space to the filename. Trailing spaces are removed from the filename before use. If the file is not found in the current working directory an attempt will be made to find it in the directory specified by the **SNOLIB** environment variable, or if that is not set, a predetermined library directory.

-COPY is a synonym for **-INCLUDE** for compatibility with SPITBOL/370.

IO_FINDUNIT()

The **IO_FINDUNIT()** function returns an unused I/O unit number for use with the **INPUT()** or **OUTPUT()** functions. **IO_FINDUNIT()** is meant for use in subroutines which can be reused. **IO_FINDUNIT()** will never return a unit number below 20.

LABEL()

The **LABEL()** predicate evaluates its argument as a string (with case folding), and returns the null string if a label with that name has been defined, and fails if it does not. The **LABEL()** predicate was copied from SITBOL and Steve Duff's version of Macro SPITBOL.

>RACE

keyword enables Garbage Collection tracing if non-zero. If positive, the value of **>RACE** will be decremented after it is tested.

&LINE/&FILE/&LASTLINE/&LASTFILE

The **&LINE** and **&FILE** keywords can be used to determine the source file and file line associated with the current statement. The **&LASTLINE** and **&LASTFILE** return the source file and file line associated with the previous statement.

-LINE

The **-LINE** directive can be used to alter SNOBOL's idea of the current source file and line (ie; for use by preprocessors). **-LINE** takes a line number and an optional quoted string filename.

Lexical comparison

A full set of lexical (string) comparison predicates have been added to complement the standard **LGT()** function; **LEQ()**, **LGE()**, **LLE()**, **LLT()**, **LNE()**.

LPAD()/RPAD()

The **LPAD()** and **RPAD()** functions take the first argument (subject) string, and pad it out to the length specified in the second argument, using the first character of the optional third argument. If the third argument is missing, or is the null string, spaces will be used for padding. The subject will be returned unmodified if already long enough.

Named files

Filenames can be supplied to the **INPUT()** and **OUTPUT()** functions via an optional fourth argument. If the filename begins with a vertical bar (`|`), the remainder is used as a shell command whose stdin (in the case of **OUTPUT()**) or stdout (in the case of **INPUT()**) will be connected to the file variable via a pipe. If a pipe is opened by **INPUT()** input in "update" mode, the connection will be bi-directional (on systems with socketpair and Unix-domain sockets). The filename `-` (hyphen) is interpreted as stdin on **INPUT()** and stdout on **OUTPUT()**. The magic filenames `/dev/stdin`, `/dev/stdout`, and `/dev/stderr` refer to the current process standard input, standard output and standard error I/O streams respectively regardless of whether those special filenames exist on your system. The magic pathname `/dev/fd/n`, opens a new I/O stream associated with file descriptor number *n*. The magic pathname `/tcp/hostname/service` can be used to open connection to a TCP server. If the path ends in the optional suffix `/priv` the local address will be bound to a port number below 1024, if process privileges allow. `/udp/hostname/service` behaves similarly for UDP. The magic pathname `/dev/tmpfile` opens an anonymous temporary file for reading and writing, see **tmpfile(3)**. On VMS, Win32, and MS-DOS (when compiled with DJGPP), the pathnames

/dev/null and /dev/tty are magical, and refer to the null device, and the user's terminal/console, respectively.

ORD()

The **ORD()** function returns the ordinal value (zero to 255) of the first character in its string argument (the inverse of the **CHAR()** function).

&PARM

The entire command line is available via the **&PARM** protected keyword for compatibility with Catspaw SNOBOL4+. Use of the SPITBOL compatible **HOST()** function is probably preferable.

REAL numbers in **INTEGER** contexts

REAL numbers (or strings convertible to **REAL**) are accepted in all contexts which previously required an **INTEGER** (or string convertible to **INTEGER**). Contexts include **TABLE()**, **ITEM()**, array indices, **INPUT()**, **OUTPUT()**, **SET()**, keyword values, **CHAR()**, **RPAD()**, **LPAD()**, **FIELD()**, **COLLECT()**, **DUMP()**, **DUPL()**, **OPSYN()**, **SUBSTR()**.

REVERSE()

REVERSE() returns its subject string in reverse order.

Scientific notation

REAL number syntax has been expanded to allow exponents of the form:

ANY('Ee') ('+' / '-' / ') SPAN('0123456789'). Exponential format reals need not contain a decimal point.

SERV_LISTEN()

The **SERV_LISTEN()** function makes SNOBOL4 into a network server process, and takes three **STRING** arguments: **FAMILY**, **TYPE**, **SERVICE**. **FAMILY** must be either "inet" for an Internet Protocol v4 socket, "inet6" for an Internet Protocol v6 socket, or "unix" for a local ("unix domain") socket. The second argument, **TYPE** must be "stream", and the third argument, **SERVICE** must be a port number or service name (for an internet socket), or a pathname (for a "unix" socket). **SERV_LISTEN()** listens for incoming requests, accepts them, then "forks" a child process and returns an integer file descriptor which can be opened for bidirectional I/O using a "/dev/fd/n" magic pathname. The original ("parent") process never returns from the **SERV_LISTEN()** call. This function is only available on systems with the "fork" system call, which makes a child process which is an identical copy of the parent process.

SET()

The **SET()** function can be used to seek the file pointer of an open file. The first argument is an I/O unit number, the second is an integer offset. The third argument, an integer determines from **whence** the file pointer will be adjusted. If whence is zero the starting point is the beginning of the file, if whence is one, the starting point is the current file pointer, and if whence is two, the starting point is the end of the file. **SET()** returns the new file pointer value. On systems with 64-bit file pointers and 32-bit integers (ie; 4.4BSD on i386) the return value will be truncated to 32-bits, and only the first and last 4 gigabytes of a file can be accessed directly.

SITBOL file functions

FILE(string) is a predicate which returns the null string if its argument is the name of a file that exists, and fails if it does not. **DELETE(string)** is a predicate which tries to remove the file named by its argument, and fails if it cannot. **RENAME(string1,string2)** is a predicate which attempts to rename the file named by *string2* to the file named by *string1*. Unlike the **SITBOL** version, if the target file exists, it will be removed.

SNOBOL4+ real functions

EXP(), **LOG()** and **CHOP()** functions are available for compatibility with SNOBOL4+. **EXP()** returns the value e^{**x} , **LOG()** returns the natural logarithm of its **REAL** argument, and **CHOP()** truncates the fractional part of its **REAL** argument (rounding towards zero), and returns a **REAL**.

SORT()/RSORT()

The **SORT()** and **RSORT()** functions take two arguments. The first can be either an array or a table. If the first argument is an array, it may be singly-dimensioned in which case the second argument, if non-null should indicate the name of a field of a programmer defined data type to use to access the sort key. Otherwise the first argument should be a table or a doubly-dimensioned array, in which case the second argument may an integer indicating the array column on which to sort. If the second argument is null, it is taken to be 1. The array (or table) is not modified; a new array is allocated and returned. **SORT()** sorts elements in ascending order, while **RSORT()** sorts in descending order.

Example: for a table *TAB* of integers, indexed by strings being used to tabulate word counts *FREQ* = *RSORT(TAB,2)* returns an array such that *FREQ<1,1>* contains the most frequent word while *FREQ<1,2>* contains the number of occurrences of that word. While *WORDS* = *SORT(TAB,1)* returns an array with the rows by the lexicographical ordering of the words; *WORDS<1,1>* contains the lexicographically first word and *WORDS<1,2>* contains the number of occurrences of that word.

SPITBOL operators

The SPITBOL scan (?) and assignment (=) operators have been added. A pattern match can appear within an expression, and returns the matched string as it's value. Similarly assignment can appear in an expression, and returns the assigned value. An assignment after a scan (ie; *STRING ? PATTERN = VALUE*) performs a scan and replace. Assignment is right associative, and has the lowest precedence, while scan is left associative and has a precedence just higher than assignment.

The SPITBOL selection/alternative construction can be used in any expression. It consists of a comma separated list of expressions inside parentheses. The expressions are evaluated until one succeeds, and it's value is returned. Abuse of this construction may result in incomprehensible code.

The type **NUMERIC** with **CONVERT()** and the removal of leading spaces from strings converted to numbers (implicitly or explicitly) are also legal when SPITBOL extensions are enabled. SPITBOL extensions can be enabled and disabled using the **-PLUSOPS** directive. **-PLUSOPS 0** disables SPITBOL operators, while **-PLUSOPS** or **-PLUSOPS *n*** where *n* is a non-zero integer enables them. SPITBOL extensions are enabled by default.

SQRT()

The **SQRT()** function is available for compatibility with SPARC SPITBOL. **SQRT()** fails if the argument is negative, but does not cause a fatal error.

SSET()

Experimental "scaled set" function. Takes arguments unit, offset, whence, and scale. The first three are analogous to the same arguments for the **SET()** function. The last parameter is used as a multiplicative scaling factor on the **offset** parameter, and as a divisor on the return value. When used in combination with relative **SET()** calls (whence of one), any file offset can be achieved, even when system file offsets are larger than can be represented in a SNOBOL4 **INTEGER**. Support for "Large Files" is enabled when available, but not all file systems support them.

SUBSTR()

SUBSTR() takes a subject string as it's first argument, and returns the substring starting at the position specified by the second argument (one-based) with a length specified by the third argument. If the third argument is missing or zero, the remainder of the string is returned.

TERMINAL I/O variable

The variable **TERMINAL** is associated with the standard error file descriptor for both input and output.

Trig functions

SIN(), **COS()** and **TAN()** functions are available for compatibility with SPARC SPITBOL and take arguments in radians.

&UCASE/&LCASE

Protected keywords **&UCASE** and **&LCASE** contain upper and lower case characters respectively.

VDIFFER()

The **VDIFFER()** function takes two arguments, if they **DIFFER()**, the first argument's value is returned. This is intended to be used in contexts where *DIFFER(X) X* would otherwise have been used. The **VDIFFER()** function was copied from Steve Duff's version of Macro SPITBOL.

I/O Associations

I/O is performed by associating a variable name with a numbered I/O unit using the **INPUT()** and **OUTPUT()** functions. The following associations are available by default;

Variable	Unit	Association
INPUT	5	standard input
OUTPUT	6	standard output
TERMINAL	7	standard error (output)
TERMINAL	8	/dev/tty (input)

I/O Options

The third argument of the **INPUT()** and **OUTPUT()** functions is interpreted as a string of single letter options, commas are ignored. Some options effect only the I/O variable named in the first argument, others effect any variable associated with the unit number in the second argument.

digits A span of digits will set the input record length for the named I/O variable. This controls the maximum string that will be returned for regular text I/O, and the number of bytes returned for binary I/O. Record length is per-variable; multiple variables may be associated with the same unit, but with different record lengths.

A For **OUTPUT()** the unit will be opened for append access (no-op for **INPUT()**).

B The unit will be opened for binary access. On input newline characters have no special meaning; the number of bytes transferred depends on record length (see above). On output no newline is appended. For terminal devices, all I/O to this unit will be done without special processing for line editing or EOF, while characters which deliver signals (interrupt, kill, suspend) are still processed. Units opened on the same terminal device entry operate independently; some can use binary mode, while others operate in text mode.

C Character at a time I/O. A synonym for **B,1**.

T Terminal mode. No newline characters are added on output, and any newline characters are returned on input. Terminal mode effects only the referenced unit. Terminal mode is useful for outputting prompts in interactive programs.

Q Quiet mode. Turns off input echo on terminals. Effects only input from this unit.

U Update mode. The unit is opened for both input and output.

W Unbuffered writes. Each output variable assignment causes an I/O transfer to occur, rather than collecting the data in a buffer for efficiency.

OPTIONS

-b Toggle startup banner output (by default on).

-d DDD Allocate "dynamic storage" region of *DDD* descriptors for program code and data. A suffix of **k** multiplies the number by 1024, a suffix of **m** multiplies the number by 1048576.

A larger dynamic region may result in fewer garbage collections (storage regenerations), however large values may cause execution to slow down when large amounts of garbage collect. Most programs do not need an increased dynamic region to run. If your program terminates with an “Insufficient storage to continue” message you need to increase the dynamic storage region size.

- f** Toggle folding of identifiers to upper case (see **-CASE** and **&CASE**).
- g** Enable garbage collection tracing (sets **>RACE** to -1).
- h** Give help. Shows usage message, includes default sizes for “dynamic region” and pattern match stack.
- k** Toggle running programs with compilation errors (see **-ERROR** and **-NOERRORS** extensions). By default programs with compilation errors will not be run.
- l** Re-enable listing to stdout. (default is **-UNLIST**). Default listing side is **LEFT**.
- n** Toggle running programs after compilation (see **-EXECUTE** and **-NOEXECUTE** extensions). By default programs are run after compilation.
- p** Toggle SPITBOL extensions (also controlled by **-PLUSOPS**).
- r** Toggle reading **INPUT** from input file(s) after **END** label. Otherwise **INPUT** defaults (back) to standard input after program compilation is complete.
- s** Toggle termination statistics (off by default).
- u *params*** specifies a parameter string available via **HOST(0)**.
- v** Show version and exit.
- Terminates processing items as options. Any remaining strings are treated as files or user parameters.
- M** Specifies that all items left on the command line after option processing is complete are to be treated as filenames. The files are read in turn until an **END** statement is found (Any remaining data is available via the **INPUT** variable if the **-r** option is also given). A **--** terminates processing of arguments as files, and makes the remaining arguments available as user parameters (see the **HOST()** function).
- P *DDD*** Allocate *DDD* descriptors for the pattern match stack. A suffix of **k** multiplies the number by 1024, a suffix of **m** multiplies the number by 1048576. The pattern match stack is used to save backtracking and conditional assignment information during pattern matching. If your program terminates with an “Overflow during pattern matching” message (Error 16) you need to increase the pattern match stack size.
- S *DDD*** Allocate *DDD* descriptors for the interpreter stack. A suffix of **k** multiplies the number by 1024, a suffix of **m** multiplies the number by 1048576. The interpreter stack is used for saving data, and passing parameters to internal procedures. If your program terminates with an “Stack overflow” message (Error 21) you need to increase the interpreter stack size. A common reason for needing additional stack space is for tracing deeply nested **DATA()** structures during garbage collection.

SEE ALSO

snobol4dbm(3),snobol4tcl(3),snolib(3).

<http://www.snobol4.org>

All things SNOBOL4 related.

<http://www.snobol4.com>

Catspaw: commercial SPITBOL implementations, Free SNOBOL4+ for DOS.

<http://www.snobol4.org/doc/burks/tutorial/contents.htm>

SNOBOL4 language tutorial (from Catspaw Vanilla SNOBOL4)

<http://www.snobol4.org/doc/burks/manual/contents.htm>

Catspaw Vanilla SNOBOL4 manual.

<ftp://ftp.snobol4.com/spitman.pdf>

Catspaw Macro SPITBOL manual

[1] R. E. Griswold, J. F. Poage, and I. P. Polonsky

The SNOBOL4 Programming Language, 2nd ed., Prentice-Hall Inc., 1971.

Reference manual for Macro SNOBOL4. (soon to be available as a PDF!!!)

[2] R. E. Griswold,

The Macro Implementation of SNOBOL4, W. H. Freeman and Co., 1972.

Book describing the implementation techniques used in Macro SNOBOL4.

<ftp://ftp.snobol4.org/snobol/doc/>

Directory with sources for University of Arizona SNOBOL4 memos, including formatted versions of memo s4d58, which describes each pseudo-instruction in the Snobol Implementation Language (SIL).

AUTHOR

Philip L. Budne

with a little help from:

R. E. Griswold, J. F. Poage, and I. P. Polonsky

Mark Emmer (code from SNOBOL4+)

Viktors Berstis (code from Minnesota SNOBOL4)

DIAGNOSTICS

“Insufficient storage to continue” (error 20)

“Dynamic” storage exhausted. See the **-d** option.

“Stack overflow” (error 21)

Interpreter stack exhausted. See the **-S** option.

BUGS

I/O retains some record oriented flavor.

I/O is still tied to unit numbers.

“Dynamic” storage cannot be expanded after startup.

Integer math can never “fail”, even on overflow.

Oversize integer constants may not be detected.