

B: Version 4.0

Reference Manual

User's Reference to B

By R.P. Gurd

University of Waterloo,

Waterloo, Ontario, Canada.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57

This edition of the manual applies to B version 4.

The section on machine-level linkage conventions in chapter 6 has been totally rewritten to reflect the change in the run time stack format.

Parts of Chapter 7 have been changed to reflect the new conventions for accessing files with OPEN.

Table of Contents

1.	Introduction	1
2.	Basic symbols	2
2.1.	Identifiers	2
2.2.	Comments	3
2.3.	Keywords	3
2.4.	Constants	3
2.4.1.	Decimal constants	3
2.4.2.	Octal constants	4
2.4.3.	Floating point decimal constants	4
2.4.4.	ASCII character constants	4
2.4.5.	BCD character constants	4
2.4.6.	String constants	5
2.4.7.	Escape sequences	5
2.5.	Source file inclusion	6
2.6.	Compiler directives	6
3.	The building blocks of a B program	7
3.1.	Manifest constants	8
3.2.	External definitions	9
3.3.	Function definition	12
4.	Statements	14
4.1.	Null statement	14
4.2.	Expression statement	15
4.3.	Storage declaration	15
4.3.1.	Storage types	15
4.3.2.	Extrn	16
4.3.3.	Auto	16
4.3.4.	Labels	17
4.4.	Transfer of control	17
4.4.1.	Goto	17
4.4.2.	Return	18
4.4.3.	Break	18
4.4.4.	Next	18
4.5.	Conditional statement	19
4.5.1.	If	19
4.6.	Iterative statements	19
4.6.1.	Repeat	20
4.6.2.	While	20
4.6.3.	Do-while	20
4.6.4.	For	20
4.7.	Switch statement	21
5.	Expressions	24
5.1.	Primary expressions	24
5.1.1.	Subscripting	25
5.1.2.	Function calls	26
5.2.	Rvalues and lvalues	26
5.3.	Unary operators	27
5.4.	Binary operators	30
5.5.	Shift operators	30
5.6.	Bitwise and	30

5.7.	Bitwise exclusive or	30
5.8.	Bitwise or	31
5.9.	Multiplicative operators	31
5.10.	Additive operators	32
5.11.	Relational operators	32
5.12.	Logical and	32
5.13.	Logical or	33
5.14.	Query operator	33
5.15.	Assignment operators	33
6.	Implementation-dependent information	35
6.1.	Linkage conventions	35
6.1.1.	Function call	35
6.1.2.	Entry	35
6.1.3.	Exit	37
6.2.	Internal representation of objects	37
7.	The B library	38
7.1.	.BSET - redirection of i/o	38
7.2.	Introduction to input/output	40
7.2.1.	Units	40
7.2.2.	Unit opening	41
7.2.3.	Unit closing	44
7.2.4.	Unit switching	44
7.3.	Sequential stream i/o	45
7.3.1.	Terminal vs file	46
7.3.2.	Stream i/o functions	47
7.4.	Random file i/o	50
7.5.	String operations	50
7.5.1.	Random string processing	51
7.5.2.	Sequential string access	51
7.5.3.	String utilities	52
7.6.	Storage allocation	53
7.7.	Character code conversion	54
7.8.	Call FORTRAN	54
7.9.	DRLs and MMEs	55
8.	Using B	56
8.1.	Compiling/running	56
8.2.	Debugging	58
8.3.	Compiler/loader interface	62
8.4.	Using tabs for readability	62
8.5.	Some pitfalls	63
Appendix A	- Escape sequences	65
Appendix B	- Binding strength of operators	66
Appendix C	- B compiler error message	67
Appendix D	- Index of B library routines	72

1. Introduction.

This manual describes the programming language B accepted by the compiler written at the University of Waterloo by R. Braga. It also introduces the runtime package written at the University of Waterloo, largely by T.J. Thompson.

A derivative of BCPL, B was designed and first implemented by D.M. Ritchie and K.L. Thompson, of Bell Telephone Laboratories, Inc., Murray Hill, N.J. The original implementation of the runtime package is due to S.C. Johnson, also of Bell Labs.

The present version of the compiler differs from the original by incorporating an expanded SWITCH statement, adding floating point operators, adding proper logical operators, and altering the order of evaluation of operators.

The runtime package works in both TSS and batch and will read almost any "media code" found in the GCOS environment. It uses EIS instructions whenever it is appropriate. Note that the language itself has no constructs for input/output; all i/o is done by function calls.

B is a typeless language, in which the compiler always assumes the type of a variable is suitable to the operator acting upon it. B has a large set of operators, providing integer, bitwise, logical and floating point operations.

The machine word is the basic unit of computation. The word size on the Honeywell Series 60 Level 66 or Series 6000 machines is 36 bits.

A B program consists of procedures called functions. Any function may call another function or, since local variables are allocated on a stack, call itself recursively. All functions may selectively access a global pool of externals.

A function is composed of a set of one or more statements. A statement is composed of permissible combinations of keywords and expressions.

Although this is a reference manual, and definitely not a tutorial, it is organized such that you will often find examples which involve material covered later on. This is intentional. Such examples should be ignored at first reading, but will hopefully be beneficial when you refer to the manual again.

You should try to keep your first efforts at B programming as simple as possible, in order to minimize the difficulties you may encounter. If you are ever unsure of some feature of the language or run-time package, try it out by writing and running a simple little program which exercises only that feature. The time you take to do so may save you a lot of trouble in the long run.

Before starting on your first major B program, you would be well advised to have a close look at the source code of a well-written B program or two in order to get some idea of exactly how things are done.

2. Basic symbols.

B is very much oriented towards use of the ASCII character set in which each character occupies nine bits (four characters per word). There is support for representing character constants in the BCD character set, in which one character occupies six bits (six characters per word).

Certain characters, such as '{' or '}', do not appear on some terminal keyboards. Escape sequences for these are defined in Appendix A.

The compiler treats its input as an unbroken stream of characters. Any form feed, tab, or newline character is converted to the space character, except when it occurs inside a string or character constant. Newlines are counted, so the compiler can tell you on what line it detected an error. Line length may be arbitrary. The compiler knows nothing about any "sequence field", such as is supported by certain card-oriented compilers.

Line-numbered source files are permitted. If the first character of the first file is a digit, the compiler assumes the program being compiled has line numbers. If so, on each line of the file, it attempts to form a line number by collecting numeric characters until a non-numeric character is found. The number is used in error messages pertaining to that line.

2.1. Identifiers.

An identifier or name is formed from the characters a-z, A-Z, 0-9, underscore ('_') or dot ('.'), and must begin with a non-digit. Since the run time package uses names containing at least one dot, you should avoid conflicts by not using the dot character in names defined in your program.

Names may be arbitrarily long, but only the first eight characters are significant. For external names or external references, only the first six characters are significant, due to the restriction imposed by the GCOS and TSS loaders.

Normally, the compiler ignores case distinctions, so that the identifiers "SUM", "sum" and "Sum" would be considered to be the same thing. You may specify an option to the compile command which forces the compiler to respect case distinctions, but then all keywords **must** appear in lower case.

2.2. Comments.

The beginning of a comment is signalled by the appearance of a `/*` in the input stream. It is ended with the first occurrence of a `*/` any number of characters or lines later. For example,

```
/*  
 * this is a comment  
*/
```

Comments may not be nested.

2.3. Keywords.

B uses 15 keywords, which may be categorized as follows:

- 1) identifier scope keywords:
 AUTO EXTRN
- 2) execution flow control keywords:
 IF ELSE FOR WHILE REPEAT SWITCH DO
- 3) transfer keywords:
 RETURN BREAK GOTO NEXT
- 4) switch statement keywords:
 CASE DEFAULT

The compiler does not allow you to use any keyword as an identifier. In particular, beware of inadvertently using NEXT as an identifier.

2.4. Constants.

You may define octal, decimal, floating point, ASCII character, BCD character, or string constants in your program. The form of a constant is well-defined, in that it is possible to unambiguously differentiate between the various types of constants.

2.4.1. Decimal constants.

A decimal constant consists of an integer number, which may not contain leading zeroes. For example:

```
25    4737  981    32
```


2.4.2. Octal constants.

An octal constant consists of an integer number, preceded by a zero and formed only from the digits zero through seven. For example:

01 077 026 0400000 077777777777

2.4.3. Floating point decimal constants.

A floating-point constant is any number containing a decimal point. It must not begin with a decimal point, but may have leading zeroes and may be followed by the letter 'e' and a possibly signed integer exponent. Examples:

3.2 1. 0.5 1.e5 3.e5 4.987e-2

2.4.4. ASCII character constants.

An ASCII character constant consists of from one to four characters inside single quotes. The result is a word which contains the internal form of the ASCII characters, right-adjusted and left-padded with zero bits. Some examples:

'a' 'abc' 'abcd'

The compiler counts characters inside character constants and issues an error message if there are more than four.

2.4.5. BCD character constants.

A BCD constant consists of from one to six characters enclosed by grave accent characters. The result is a word containing the characters transliterated to BCD, right justified, and left-padded with zero bits. Characters which do not have an exact equivalent in the BCD set are converted to BCD blanks. Here are three examples of BCD constants:

'a' 'ot' '123456'

If your terminal does not have a grave accent character you may alternately alternately write a BCD constant like an ASCII character constant, except preceded by a dollar sign, as in

\$'a' '\$ot' '\$123456'

Note that the runtime package provides functions to transliterate between ASCII and BCD and that the i/o function PRINTF will take a BCD format specification.

2.4.6. String constants.

A string constant is any string of characters enclosed in double quotes. For example:

```
"this is a string"
```

```
""
```

```
"the above is the null string"
```

When processing a string, the compiler packs the characters of the string four per word and always appends one extra character, an ASCII null (000), to mark the end of the string.

The value of a string is quite different from the other types of constants. The value of a floating-point, octal, decimal or character constant is a word containing the internal representation of the given constant. The value of a string constant is a word containing the address of the string.

In constructing the string, the compiler gobbles all characters it sees, translating escape sequences if necessary, until it finds a closing, unescaped double quote. The rule which says tabs and form feeds are ignored does not, of course, apply in this case, but real newlines (as opposed to escaped newlines) are treated specially. If a real newline is preceded by a '*', both '*' and newline are thrown away, so you can enter a very long line. If the newline is not preceded by a '*', it is kept, but a warning message is issued, on the grounds that you probably forgot the closing string quote. To get a newline in a string without drawing a warning, use the escape '*n'.

2.4.7. Escape sequences.

Escape sequences, beginning with the character '*', are defined to allow you to use, in a string or character constant, characters which would otherwise be inconvenient or impossible to enter. For example, if you wanted to place a double quote inside a string constant, you would use the escape '*\"'. The special end-of-string character (ASCII null) is escaped as '*0'. The newline character is escaped as '*n'. A newline is taken as a carriage return and a line feed when output to a terminal. Arbitrary nine-bit characters can be generated with '*#nnn', where nnn is one to three octal digits. The complete set of escape sequences is given in Appendix A.

2.5. Source file inclusion.

If the compiler encounters in the source program a line of the form

```
%filename
```

it suspends processing of the current file and begins collecting input from the specified file. When end-of-file is encountered in the included file, processing in the original file resumes at the next line following the file inclusion request. Such included files may themselves contain "%filename" requests pointing to other files.

The '%' character must be the first character on the line, not just the first non-blank character. If the line has a line number, the '%' must immediately follow the line number.

The file name given may be in any of the forms acceptable in the TSS environment, such as

```
%temp  
%/main.b  
%fbaggins/dif.b  
%fbaggins/s/dif.b  
010%fbaggins/s/dif.b
```

This allows you to keep parts of a large module broken up into easily manageable files, while retaining the ability to compile the files together. File inclusion is also often used to bring in a file of manifest definitions, such as TSS Derail equivalences, which a variety of possibly unrelated programs might find useful.

2.6. Compiler directives.

Any line whose first character is a '#' is assumed to be one of the compiler directives shown below. In each case, "<text>" denotes a string of characters which begins with any non-blank character.

```
#title <text>
```

will place "<text>" in the comments field of any \$ OBJECT card written from the time the directive is encountered.

```
#lbl <text>
```

will place "<text>", truncated to eight characters if necessary, in the "deck name" field of any \$ OBJECT or \$ DKEND card written from the time the directive is encountered. If this directive has not been encountered when it comes time to generate an object deck, then the current file name is used instead.

#ttldat <text>

After being truncated to six characters if necessary, "<text>" is used to fill in the "ttl" date field of any \$ OBJECT card written from the time the directive is encountered.

#copyright <text>

is taken as comments.

Here is an example of the use of all four directives:

```
#title tss login subsystem - .tslog
#lbl tlga
#ttldat 771209
#copyright (c) by the University of Waterloo, 1977.
```

3. The building blocks of a B program.

A complete B program consists of at least one, but usually many, modules. A module is any of the following:

- 1) a manifest constant identifier definition.
- 2) an external (global) variable definition.
- 3) a function body definition.

Modules may appear in any order at all, with the sole proviso that the definition of a manifest identifier must appear before the identifier is first used.

Manifest definitions are used to associate a name with a compile time constant.

External definitions are used to create a global pool of possibly initialized identifiers. This pool might be used to declare large blocks of memory, or to declare identifiers that must be accessible to more than one B function. Since an external is global in scope, any B function may use it, but only after declaring its intention to do so in an EXTRN statement.

A function definition is used to declare a component of the executable code of the program. The definition includes the name the function will be called by, the arguments it will be called with, and the statements which define what it will reference and what it will do.

3.1. Manifest constants.

A manifest constant has the general form

```
name = text;
```

where "name" is any valid identifier and "text" is the collection of characters between the equals sign and the semicolon.

When a manifest is defined, the compiler enters the identifier in the symbol table and associates it with the "text", which it keeps in an internal buffer. Absolutely no processing of the "text" is done at the time of definition.

When the compiler reads an identifier, it first checks to see if the identifier is a manifest. If so, the action taken is to substitute the text of the manifest for the identifier. For this reason, it is not possible to speak of redefining a manifest and any inadvertent attempt to do so will usually result in a syntax error. Substitution effectively takes place before the syntax analyzer scans the line. Manifests may be used anywhere, including inside later manifest definitions.

Because the compiler does not analyze the text of the manifest until substitution takes place, it is possible for the text to refer to a manifest which is defined after it, as long as the definitions of both appear before the first use.

The use of manifests has no effect on the order of expression evaluation. For example, look at the definitions

```
A = 1;  
B = A+A;  
C = B*B;
```

When "C" is used somewhere else in the program, the compiler will actually get "1+1*1+1", which will be evaluated as three and not four, as one might mistakenly assume.

The compiler permits nesting up to 10 levels deep of manifests inside other manifests.

Normally, however, you will find manifests quite natural to use. Here are some examples of manifests:

```
VECSIZE = 63;  
vec [ VECSIZE ];  
...  
for( i = 0; i <= VECSIZE; ++i ) sum += vec[i];
```

The manifest "vecsize" is used to establish the size of the vector "vec" at compile time, and later used to control iteration in a FOR statement.


```

3      /* binary list structure */
4      NULL = -1;    EMPTY = 0;
5      CONTENTS = 0;
6      LEFT_PTR = 1;
7      RIGHT_PTR = 2;
8      ...
9      printree( ptr )
10         if( ptr != NULL)
11            {
12                printree( LEFT_PTR[ptr] );
13                print_contents( CONTENTS[ptr] );
14                printree( RIGHT_PTR[ptr] );
15            }
16      /* end printree */

```

Here, manifests are being used to make the code involved in traversing a binary tree more meaningful. Manifests are often used in this way to give names to the elements of a structure, which may be an array of fixed size, or a dynamically allocated block of storage.

A common convention, used in this document, is to differentiate manifest identifiers from other identifiers by always showing the manifest identifier in upper case, and to show all others in lower case. It is also considered good practice to group all manifests for a program together at the beginning of the source code.

3.2. External definitions.

We will first look at the possible forms of external definitions, then look at several examples.

If an external is defined and possibly initialized in more than one place, the first one encountered during loading is the one which is used.

These are the possible forms of external definitions:

name;

A single word is allocated and initialized to zero.

name { ival };

"Name" is defined as a single word and initialized with the single value ival.

Ival may be any legal constant expression, in which case "name" has the value of its result. A constant expression may be either a string constant or an expression formulated with any legal combination of numeric or character constants, binary operators, unary operators and parentheses, following the rules in the chapter on expressions. Alternatively, ival may be a name, in which case the value of the ival is a word containing the address of the name given. A function name may be used.

name (ival, ival, ...);

Allocates space for as many words as there are ival's. This is in effect a vector which does not have a word set aside as a pointer to it; its address is "&name", rather than just "name", which in this case refers to the first ival. This is the way a vector is set up in FORTRAN, but is not the same as a B vector.

name [const-expr];

This is the first of several forms of B vector declarations. "Name" is defined as a pointer to a vector whose length is a number of words which is the value of the constant expression plus one (since all B vectors start subscripting at zero). The zeroth element of the vector is initialized to zero; the initial contents of the remaining cells are undefined.

The const-expr in brackets may be any expression which is a legal combination of numeric or character constants, unary operators, binary operators, and parentheses. It is up to you to make sure the value of the expression is reasonable, since the compiler's grammar lets it accept things like floating point constants and negative numbers, which give absurd results. For all practical purposes, "const-expr" must be such that it gives an integer result.

name [] { ival, ival, ... };

"Name" is defined as a pointer to a vector whose length is the number of initial values.

name [const-expr] { ival, ival, ... };

"Name" is defined as a pointer to a vector whose length is the maximum of the result of the constant expression plus one and the number of initial values. The contents of those vector elements, which do not have corresponding initial values, are undefined.

For compatibility with a previous version of the compiler, B also accepts an ival or ival list which is not surrounded by braces. In this case, the compiler does not permit a constant expression to appear. Only a numeric, character or string constant is acceptable, although a numeric constant may be prefixed by an integer unary minus sign.

Here are some examples of external definitions:

a { 10 };

One word of storage is allocated, initialized to the decimal constant 10 and associated with the name "a".

b [] { 'ab', 'cde', 'fghi' };

One word is associated with the name "b" and initialized with a pointer to a vector of three words. The first element of the vector, referred to as "b[0]", is initialized with the character constant 'ab'. The other two elements, "b[1]" and "b[2]", are initialized with 'cde' and 'fghi', respectively.

c { 'ab', 'abc' };

"c" is defined and associated with a word containing 'ab'. The word immediately following is initialized to the constant 'abc', but is not associated with any name. This is in effect a vector which does not have a word set aside as a pointer to it.

d[63];

Defines "d" and associates it with a word containing a pointer to a vector of 64 words, whose initial contents are undefined.

e[10] { a, b, c, d };

Declares "e" and associates it with a vector of 11 words. Words zero, one, two and three are initialized with the addresses of the externals "a", "b", "c" and "d", respectively. The contents of the remaining elements are undefined.

f { "a string" };

This is the usual way of defining an external string with an initial value. "f" is defined and initialized with a pointer to the storage occupied by the string constant "a string".

g[] { "pascal/library", "pascal/compiler", -1 };

This sets up a vector of strings with an end marker. It defines "g" and associates it with a word containing a pointer to a vector of three words. The cell "g[0]" is initialized with a pointer to the storage occupied by the string constant "pascal/library". The cell "g[1]" is initialized in a similar manner, while "g[2]", the last element of the three-word vector, is initialized to the decimal constant -1.

B does not allow you to explicitly declare arrays with more than one dimension. Usually if you need more than one dimension, you build it at run time by calling the library

function GETMATRIX, which will obtain storage, construct the necessary edge vectors and return a pointer to the array. However, in spite of the fact that you cannot declare such an array, it is possible to construct one as an initialized external! The secret is that any ival (inside braces) may be replaced by an ival or ival list surrounded by braces. The compiler then constructs the ival list and places a pointer to it in the original ival list. The maximum nesting depth is seven. For example,

```
x[ ]{
    { 00, 01, 02 },
    { 10, 11, 12 },
    { 20, 21, 22 }
};
```

In this case, "x" ends up being initialized as a pointer to a vector containing three pointers. Each pointer points to a vector of three words. In an expression, the value of "x[0]" is a pointer to the first vector of three words, while the value of "x[1][2]" is 12.

3.3. Function definition.

B functions serve a purpose similar to the subroutine in FORTRAN or the procedure in ALGOL. The mechanism of the function call involves very little cost in overhead and permits recursion.

A working B program always contains at least one function, called MAIN, and usually others, since the language is designed to encourage modular or structured programming.

The general form of a function definition is

```
name( arg1, arg2, ... ) statement
```

The name must be a valid identifier and is automatically defined as an external by the compiler.

The formal arguments consist of a possibly empty list of identifiers separated by commas. Each argument is implicitly declared as an automatic (local) variable and storage for it is allocated on the runtime stack frame. Note that, although you may not declare a vector as a formal argument, you can always use an argument in a subscripting operation, as if it were a vector pointer.

"Statement" defines what actions the function will take. Most often, it is a compound statement, consisting of a set of statements enclosed by braces. The rules for formulating statements are presented in the next chapter.

When writing the code for a function, there are a few things you should keep in mind.

The caller will always pass its arguments strictly by value. This means that altering an argument has no effect on the state of the caller. However, if an argument is a

pointer, you can change the state of the caller by indirecting through the pointer using either the unary indirection operator or subscripting.

The function may at any time return a one word value using the RETURN statement. The caller and the callee do not have to agree on whether or not a value is returned. If a value is returned, but not expected, the value is ignored. If a value is expected, but not returned, the value is undefined.

A function can determine the actual number of arguments it is called with by invoking the library function NARGS. For instance, the statement

```
x = nargs();
```

would assign to the variable "x" the number of arguments supplied to the current invocation of the function.

The availability of NARGS lets you write functions which may take a variable number of arguments. Most of the time, this means that such a function is called with fewer arguments than are defined for it, in which case one of the first things such a function does is to establish default values for the arguments it does not have.

The other case, in which a function is called with more arguments than are defined for it, is somewhat trickier, and should be avoided, unless you know precisely what you are doing.

The function called MAIN is the entry point to your program from the B runtime initialization routine. If not present, the TSS loader prints the message:

```
<w> main undefined
```

For details on how your main function is invoked, see the chapter on the run time library.

4. Statements.

Statements are used to define the actions taken by a B function. They may appear only in the body of a function definition. In certain cases, the definition of a statement is recursive, in that a statement may appear inside a statement. In this chapter, you will see that in many cases, one may use an expression in a statement. Since the rules for formulating expressions are discussed in the next chapter, we merely note here that an expression may be a statement, but a statement may not appear in an expression.

In every case where a statement is permitted, it may be replaced by a compound statement, consisting of one or more statements enclosed in curly braces, as in

```
{
  statement1
  statement2
}
```

The compiler does not permit a null compound statement like

```
{ }
```

All statements, except compound statements, must end with a semicolon.

In the formal definitions which follow, reserved words are underlined and parentheses, where shown, are required. Also, "statement" implies either a statement ended by a semicolon or else a compound statement surrounded by braces.

4.1. Null statement.

```
;
```

The null statement does absolutely nothing. It is typically used to supply a null body to a WHILE statement, as in

```
while( putchar( getchar() ) );
```

or to provide a convenient place on which to hang a label.

4.2. Expression statement.

```
expression ;
```

Any valid B expression followed by a semicolon is acceptable as a statement. To be meaningful, the expression will usually involve an assignment operation or function call, as in

```
x = min(a,b) + x;  
open( "/myfile", "r" );  
++i;
```

but the compiler will happily accept statements which do absolutely nothing, such as

```
a < b;  
open;  
i;
```

Remember that, in B, assignment is an operator in an expression, not a statement.

4.3. Storage declaration.

4.3.1. Storage types.

Before discussing the statements pertaining to storage declaration or reference, we will briefly look at how storage is allocated in a B program.

External storage consists of the global pool of externals declared in the manner described previously. For a function to use one of these externals, the name must be referenced in an EXTRN statement.

Automatic storage consists of local variables which are created anew on the runtime stack each time the function is called and which disappear when the function returns. Automatic storage is unique to each invocation of a function.

Internal (local static) storage is allocated within a function body and is common to all invocations of a function. The label, which is never explicitly declared, is the only permitted instance of internal storage.

Constants used inside functions are allocated as internal storage, but the compiler will not accept constructs that would result in directly changing a constant's value.

Finally, there is a pool of free storage which can be dynamically allocated by the library function GETVEC and dynamically released by the library function RLSEVEC. This free area is automatically grown as required up to the limits imposed by the operating system.

Any identifier used in a function body must be a formal argument, a label, or previously referenced in an EXTRN or AUTO statement. The only exception is a function name used in a function call, since the compiler automatically types as external any name immediately followed by a left parenthesis.

AUTO and EXTRN statements may appear anywhere in a function body, but you should group them at the beginning of the function.

4.3.2. Extrn.

```
extrn name1, name2, ... ;
```

This statement allows the function to begin using the names previously defined as externals (see chapter 2). Although an identifier may be externally declared as a vector, you should not indicate this in the EXTRN statement, since B lets you use any cell in a subscripting operation.

4.3.3. Auto.

```
auto name1, name2[const-expr], ... ;
```

The AUTO statement is used to declare local storage, which is unique to each invocation of the function. For example,

```
auto x;  
auto i, j, x[10];
```

A vector declaration is legal in an AUTO statement, but the size of the vector must be a constant expression, since it is established at compile time.

Const-expr is any legal combination of numeric or character constants, unary operators, binary operators and parentheses. Make sure what you use is sensible, because the compiler accepts constructs like "auto x[-1]" which lead to undefined results. Normally, one would expect a constant expression which is not a simple numeric constant to involve a manifest constant, as in

```
MAX = 10;  
...  
auto x[MAX*2], y[MAX + 7];
```

If you need dynamic vector allocation, you must use the library function GETVEC to obtain it from the free storage area.

An AUTO statement which declares a vector is executable in the sense that, when it is encountered, it initializes

the pointer to "n+1" words. The initial contents of an AUTO vector or other AUTO variables are always undefined.

Because AUTO variables are allocated on the stack, and because the stack size is finite, you should be cautious about declaring vectors, whose size is greater than 64 words, as auto variables. Although the compiler command will let you change the default stack size of 500 words, it may be preferable to either use an external, or else allocate it from free storage.

4.3.4. Labels.

Any unique identifier followed by a colon and preceding a statement is defined as a label. For example:

```
again: ;
nxt: x = getchar();
```

A statement may be preceded by as many labels as appear to be necessary, as in

```
lab1: lab2: lab3: printf("hi there");
```

4.4. Transfer of control.

The GOTO statement does the obvious thing. The RETURN statement is used to exit from a function. The NEXT and BREAK statements greatly simplify loop control.

4.4.1. Goto.

```
goto identifier ;
```

will cause a function to transfer control to the statement which has the label "identifier".

If "label" has not already appeared, it is defined as one. It is a fatal error if it is not used as a label in the function body.

It is legal to transfer to any location inside a function body, including into or out of a compound statement. It is almost never a good idea to transfer into a compound statement, because the action is difficult to follow and because it can lead to unpleasant surprises.

Because B is a typeless language, the compiler has no way of knowing whether the identifier you supply in a GOTO statement really turns out to be a valid label at runtime, so it is perfectly legal, but probably erroneous, to say

```
extrn b;
...
goto b;
```


Never try to pass a label as an argument to a function and use it to transfer to another function. The program will end up in one function, but with a different function's stack pointer, resulting in immediate or eventual disaster.

4.4.2. Return.

```
return ;
```

```
return ( expression ) ;
```

The RETURN statement ends the execution of a function and results in return to the caller. Upon return, all temporary storage in use by the particular invocation of the function disappears.

The first form of the RETURN statement merely returns control. The second form causes a one word value to be returned also.

Note that the construct

```
return();
```

is not permitted by the compiler (it gives a syntax error).

A simple RETURN statement is supplied implicitly at the end of a B function body.

The library function EXIT is also available, should your program wish to terminate execution at a point other than after the last statement of MAIN.

4.4.3. Break.

```
break ;
```

The effect of BREAK is to drop out of the most recent innermost enclosing WHILE, FOR, SWITCH, REPEAT, or DO-WHILE statement. The compiler generates a fatal error if a BREAK statement is not inside one of these.

4.4.4. Next.

```
next ;
```

NEXT is a directive to skip all further statements in the most recent enclosing WHILE, FOR, REPEAT, or DO-WHILE loop, and transfer to the test which determines whether looping should continue. Note that NEXT is only legal in a SWITCH statement if the SWITCH is itself inside one of these looping statements.

4.5. Conditional statement.

4.5.1. If.

`if (expression) statement`

If the result of the expression is non-zero, then the statement is executed. The parentheses around the expression are mandatory.

`IF (expression) statement ELSE statement`

If the result of the expression is non-zero, the first statement executes, otherwise the second statement executes.

In the case of nested IF statements where there are fewer ELSE s than IF s, the compiler associates the ELSE with the closest IF at the same level of nesting.

`if (...) if (...) s1 else s2`

resolves to

`if(...)`

Think of IF s and ELSE s being placed on a pushdown stack as they appear. An ELSE which you pull off the stack always goes with the next IF pulled off.

Here are some examples of IF statements:

`if(a) y = x;`

`if(a < 2) y = a; else y = 0;`

`if(a != b) z = g(y);`

`else`

`{`

`a += x;`

`b -= y;`

`}`

4.6. Iterative statements.

A REPEAT iterates a statement until a BREAK statement is encountered or a GOTO causes control to pass outside the loop. WHILE iterates a statement as long as an expression is non-zero, testing at the top of the loop. DO-WHILE iterates a statement until an expression is non-zero, testing at the bottom of the loop. A FOR uses three expressions to initialize, test and modify in controlling a loop.

4.6.1. Repeat.

`repeat statement`

The REPEAT merely executes the statement forever. The statement is almost invariably compound. NEXT and BREAK statements are legal inside a REPEAT.

4.6.2. While.

`while (expression) statement`

If the result of the evaluation of the expression is non-zero, the statement associated with the WHILE is executed. After execution of the statement, the expression is re-evaluated again and, if the result is again non-zero, the statement is executed again. In other words, while the result of the expression is non-zero, the statement is executed. When the result of the expression is zero, control passes to the next statement following the WHILE statement.

BREAK and NEXT statements are legal in a WHILE statement.

4.6.3. Do-while.

`do statement while (expression);`

The DO-WHILE provides a loop with a test at the bottom of the loop. It is equivalent to:

```
repeat
{
    statement
    if( !expression ) break;
}
```

BREAK and NEXT statements are legal in a DO-WHILE statement.

4.6.4. For.

`for (expr1; expr2; expr3) statement`

The FOR statement may be used to set, test and increment a variable in order to control a loop. The FOR statement is equivalent to


```

1
2
3     expr1;
4     while ( expr2 )
5     {
6         statement
7         expr3;
8     }

```

The first expression, which might initialize a controlling variable, is evaluated. Then, if and while the result of second expression (usually a test) is non-zero, the statement is executed. Before returning to re-evaluate the second expression, the third expression, which might increment a controlling variable, is evaluated.

Both BREAK and NEXT are legal in a FOR statement. The effect of NEXT is to pass control to the evaluation of the third expression.

Any or all of the expressions may be null, and they need not necessarily involve the same controlling variable, if any. Note that the second expression is always treated as a logical expression. Some examples:

```

23     for( i = 0; i < 10; ++i ) x[i] = j[i];
24
25     for( i = 10; i <= x; i += 2 )
26         for( j = 1; j < y; ++j )
27             g[i][j] = f( i + j );
28
29     for( ; i < n; ++i ) y[i] = z[n - i];
30
31     NULL = 0;
32     NEXT = 1;
33     DATA = 0;
34     ...
35     for( p = startlist; p != NULL; p = p[NEXT]; )
36         if( p[DATA] >= x ) break;
37
38
39

```

4.7. Switch statement.

The SWITCH provides a conditional branch depending on the one word result of an expression. The SWITCH has the following formal syntax:

```

44     switch ( expression ) statement
45
46
47
48
49

```


The statement is always compound and special labels are allowed inside the statement to point to where to start processing for a given case, as in

```
switch ( expression )
{
    case const-expr: statement
    case const-expr :: const-expr: statement
        break ;
    case <rel op> const-expr: statement
        /* rel op. is one of <, <=, >=, > */
    default : statement
}
```

The SWITCH evaluates the expression and compares the result with the constant or constant bounds in each CASE label. It selects a case, if there is one, and begins executing the compound statement at the statement immediately following the appropriate CASE label. If the expression result fits no case, execution continues at the label DEFAULT (if supplied) or at the next statement following the SWITCH, if DEFAULT is not supplied.

Once a case is selected, execution always falls through into the next case, unless a statement which alters the control flow is encountered.

Usually, a BREAK is used. It causes control to go to the statement following the SWITCH.

A statement may have more than one label or CASE label, just as a label or CASE label may be followed by more than one statement.

As shown above, a CASE may be satisfied by 1) a single value, 2) a range of values which include the endpoints, or 3) an upper or lower bound. Overlapping bounds draw a fatal diagnostic from the compiler.

By const-expr we mean as usual any legal combination of numeric or character constants, unary operators, binary operators and parentheses which can be evaluated at compile time as some constant value. String constants are not permitted in this context.

Any attempt to SWITCH on floating-point values will not work, since the generated code performs integer comparisons.

The compiler will construct a jump table for the SWITCH statement if the ratio of 1) the maximum case value minus the minimum to 2) the number of case labels is between one and two.

1
2
3 As an example, here is a function which uses a SWITCH to
4 determine if a character is legal for a B identifier.
5

```
6     alphnum( c )  
7         switch( c )  
8         {  
9             case 'A' :: 'Z' :  
10                /* converts upper case to lower */  
11                /* and falls through to return */  
12                c |= ' ';  
13             case 'a' :: 'z' :  
14             case '0' :: '9' :  
15             case '.':  
16             case '_':  
17                 return( c );  
18                 /* would use break if return not used */  
19             default:  
20                 return( 0 );  
21         }  
22     /* end of alphnum */  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57
```


5. Expressions.

Expressions in B are constructed according to rules which govern combinations of operators, identifiers, square brackets, and parentheses. B has a large set of operators, which are described in this section.

Because B is typeless, the compiler always assumes a given operation on a word is appropriate. Although this tends to force you to do more checking yourself, it also gives you the scope to do almost anything you want. This typeless characteristic often causes trouble for beginning users of B, because the compiler happily accepts possibly erroneous operations, such as adding one to a function name, or using a pointer as a function call. Such is the price of freedom.

The compiler takes no responsibility for the validity of expressions. There is no runtime monitoring of possible arithmetic overflows or faults. In the B run time environment, overflow faults are masked out, but, in the Pascal run time environment, they are not. A divide error (like dividing by zero) will result in a fault.

Expressions are evaluated according to an order of binding which includes both the hierarchy of evaluation, which determines the order of evaluating different types of operators, and grouping, which determines the order in which operators of the same type are evaluated. We will discuss the hierarchy from highest (evaluated first) to lowest and mention the grouping rule for each type. The results are summarized in Appendix A and the explain file "explain b binding".

5.1. Primary expressions.

The primary expression is the basic building block used to construct expressions. It is defined recursively as follows:

name

A legal identifier is a primary expression.

constant

Any legal constant constitutes a primary expression.

primary[expr]

A subscripting operation, which is a primary expression followed by an expression in square brackets, is a primary expression.

primary(arglist)

A function call operation, which consists of a primary expression followed by an open parenthesis, is a primary expression. The open parenthesis must be followed by a possibly empty set of arguments, consisting of comma-separated expressions, followed by a close parenthesis.

(expr)

Any expression enclosed by parentheses which is not a function argument list is a primary expression. This lets you use parentheses to alter the order of binding.

Here are some examples of simple primary expressions:

x getchar() (a + b) 6 x[i] 6[x]

In cases where a primary expression is composed itself of another primary expression, grouping occurs from left to right. For example, look at

x[i][j] x[i]()

In the first case, "x" is treated as a pointer to a vector of vectors. In the second case, "x" is treated as a pointer to a vector of functions, one of which is to be called. In both cases, "x[i]" is evaluated first, placed in a temporary, call it "y", and then the remainder of the expression is evaluated as "y[j]" or "y()", respectively.

5.1.1. Subscripting.

Subscripting is not restricted to use with variables originally declared as vectors. It is a completely general operation which may be applied using any two arbitrary expressions.

To help you understand how subscripting works in B, take a look at the primary expression

a[i]

One of the variables is supposed to be a pointer, while the other is supposed to be an offset, but it does not matter which! The reason for this is that B gets a pointer to the cell "a[i]" simply by adding "a" and "i" together. If the value of the cell is required, the compiler gets it by using the pointer. Therefore it is perfectly legal to alternatively say:

i[a]

anywhere you could have said "a[i]".

5.1.2. Function calls.

As you can see above, the general form of a function call is:

```
primary( expr1, expr2, ..., exprn )
```

Most commonly, "primary" is just the name of the function to call, but the generality of expression is there to permit you construct and use vectors or lists which contain functions to be called.

A function call primary may always be assumed to return a value. It is up to the programmer to make sure that a value is returned when one is wanted or that a value is only wanted when one is returned.

It is also up to the programmer to make sure that a function is called with as many arguments as it needs. It is safe to call a function with more or fewer arguments than are defined for it, assuming the called function is prepared for such contingencies.

Note that it is the parentheses surrounding the argument list which tell the compiler the operation is a function call, so they must always be present. For instance, say you have a function called PROC which requires no arguments. To call it, you say

```
proc()
```

But if you say only

```
proc
```

no function call will take place, because none is implied.

5.2. Rvalues and lvalues.

When we come to the assignment statement, or to operators which perform implicit assignment, it becomes necessary to distinguish between the address of a thing and its contents.

An rvalue is the contents of a word. Any expression in B may be evaluated for an rvalue. For example, the rvalue of a subscripting operation is the word addressed by the sum of pointer and offset.

Everywhere in this manual where we say "expression", we mean an expression whose result is some rvalue.

An lvalue is the address of a word. Only a name, a subscripting operation, or a primary expression prefixed by the unary indirection operator '*' may be evaluated for an lvalue. The lvalue of a subscripting operation is the address formed by the adding pointer and offset.

1
2
3 It is convenient to think of an lvalue as an expression
4 which is legal to the left of an assignment operator and of
5 an rvalue as an expression which is legal to the right, as
6 long as you remember that both may also appear in other cir-
7 cumstances.

8
9 Context determines whether an expression is evaluated
10 for its rvalue or its lvalue. For example, look at the as-
11 signment

12
13 $a[3] = 2 + x$

14
15 The expression on the right yields an rvalue which is the
16 sum of the contents of "x" and the constant "2". "a[3]" must
17 be able to, and does, yield an lvalue which is the address
18 of the place to put the sum.

19 Conversely, it is illegal to say either of

20
21 $6 = x$
22 $(a + b) = x$

23
24 because the expressions which are on the left of the assign-
25 ment operator are not permitted to have an lvalue. If they
26 could have an lvalue, you could then in the first case
27 change the value of the constant, or in the second case make
28 a meaningless assignment.

29 30 5.3. Unary operators.

31 A unary operator acts upon a unary expression to
32 transform it in some manner. A "unary expression" is either
33 a primary expression or a primary expression already modi-
34 fied by one or more unary operators. In the definitions
35 below, "rvalue" or "lvalue" must be a unary expression.
36 Unary operators are applied from left to right. Except for
37 the unary indirection operator, the result of applying a
38 unary operator is always an rvalue.

39 The following unary operators are defined:

40 #rvalue

41 Assumes the value of the expression to be integer and
42 converts it to single precision floating point.

43 ##rvalue

44 Converts single precision floating point to integer.

45 ~rvalue

46 One's complement. Converts all zero bits of its
47 operand to ones and all one bits to zeros.

48 -rvalue

49 Results in the arithmetic negation (two's complement)
50 of the operand.

1
2
3 #rvalue

4 Results in the floating point negation of the operand
5 word.

6 !rvalue

7 Logical not. The result is zero if the operand is
8 non-zero; otherwise, the result is one.

9 *rvalue

10 The indirection operator. Takes the rvalue but uses
11 it as an lvalue. This is the only case in which a
12 unary operation returns an lvalue. Thus any primary
13 expression prefixed by a "*" may appear on the left
14 hand side of an assignment. "*6 = x" stores the con-
15 tents of x in location six. "y = *x" stores the con-
16 tents of the word pointed at by x into y.

17 &lvalue

18 The address operator. Forces the program to generate
19 the lvalue of the expression, then use it as an
20 rvalue. For instance, "&x" is an rvalue which con-
21 tains the address of x in the lower 18 bits, while
22 "&6" is illegal, because "6" may not have an lvalue.

23 ++lvalue

24 Adds one to the rvalue, before using it. Each of the
25 auto increment/decrement operators require an lvalue
26 as its operand. An lvalue is required, because of the
27 implicit action of assignment, but the result is al-
28 ways an rvalue.

29 lvalue++

30 Adds one to the rvalue, after using it.

31 --lvalue

32 Subtracts one from the rvalue, then uses it.

33 lvalue--

34 Subtracts one from the rvalue, after using it.

35 @primary

36 The at-sign operator is used to force the use of
37 Honeywell hardware indirection. Its effect is to OR
38 the indirect bit into the last generated instruction
39 for an expression (rvalue or lvalue). The instruction
40 affected is usually a load or store. It was most com-
41 monly used to access characters using tallies, by in-
42 directing to a word with tally modification and the
43 address of a tally word. Normally, you should not use
44 it.
45
46
47
48
49
50
51
52
53
54
55
56
57

There is a fundamental relationship between the "*" operator and subscripting which should help you understand how addressing works in B. The following are exactly equivalent everywhere:

$$a[b] \quad \Leftrightarrow \quad *(a+b) \quad \Leftrightarrow \quad b[a]$$

To be able to write or understand B programs, it is vital that you understand the validity of this relationship.

Here are some examples involving unary expressions:

`++i`

Adds one to the value of "i". Frequently used shorthand for "i = i + 1".

`a[b][c]++`

Forms an address by adding together "a" and "b", picking up the word pointed at and then adding "c" to the contents. It is equivalent to "`*(*(a + b) + c)`". If this were part of a larger expression, the word pointed at would be loaded into a temporary. Then the contents of the addressed matrix element is incremented by one.

`&a[i]`

Forms the address of the cell "a[i]" by adding the values of "a" and "i". That is, it is evaluated as "a + i".

`y = *(&x)`

The verbose way of saying "y = x".

`x = *p++`

The word which word pointer "p" points at is copied into "x", then "p" is incremented by one to point at the next word. That is, "p" is used, then incremented.

`x = ++*p;`

The word pointed at by "p" is incremented by one and then copied into "x".

`***p = x`

The contents of "x" are copied into the word pointed at by "p", but only after "p" has been incremented by one to point to the next word. That is, "p" is first incremented, then used.

`*6 = 2`

Places the value two in location six. "`*6`" is the same as "`0[6]`" or "`6[0]`". This kind of construct is used to access locations in the slave program prefix.

`x[a[b] + 1]`

Gets the contents of "a[b]" into a temporary and adds one to the temporary to get the subscript. The contents of "x" and the subscript are added together, yielding the address of the element of "x" to be used.

5.4. Binary operators.

All other operators are binary operators, which means they require both a left and a right operand. Each operand must be an rvalued expression.

With one exception, the order in which the two operands are evaluated is undefined, so don't have the evaluation of one side depend on a side effect generated by the other side (in a function call, for instance). Logical operators are the only exception. Their operands are always taken strictly from left to right.

The code generated for floating point operations is correct, but not blindingly efficient, since the compiler generates a separate load and store for each use of a floating-point operand. However, it is there if you need it. For non-casual use of these operators, it is probably better idea to either call a Fortran or Pascal routine to do the job, or else program in some other language.

5.5. Shift operators.

`expr << expr`

The left operand is taken as the one word bit pattern to be logically left shifted. The right operand supplies the number of bits to shift. If negative, or greater than 127, the result is undefined.

`expr >> expr`

Logical right shift according to the same rules. No arithmetic right shift is defined in the language, but you may use the library function ARS.

Shift operators group from left to right.

5.6. Bitwise and.

`expr & expr`

The "&" operator takes the bitwise "and" of its two 36 bit operands. If bit *i* of both operands is one, then bit *i* of the result is one. Otherwise bit *i* of the result is zero.

5.7. Bitwise exclusive or.

`expr ^ expr`

Takes the bitwise "exclusive or" of its two 36 bit operands. If bit *i* is on in one, but not in the other, then bit *i* of the result is on.

5.8. Bitwise or.

`expr | expr`

This takes the bitwise "or" of its two operands, such that if bit *i* is on in either of the two operands or both, bit *i* in the result is on also.

The following is a summary chart of the results of bitwise operations. The table shows the effect of each operation on one bit.

operands		results		
a	b	and	or	exor
0	0	0	0	0
0	1	0	1	1
1	0	0	1	1
1	1	1	1	0

The three bitwise operators group from left to right.

5.9. Multiplicative operators.

`expr / expr`

Integer division of the first integer operand by the second. Will result in a divide check abort if the right operand is zero. The result is zero if the left operand is less than the right. The result is truncated towards zero if the right operand does not divide evenly into the left. The result is positive if the operands are both positive or both negative; otherwise, it is negative.

`expr % expr`

Results in the integer remainder of the integer division of the first operand by the second. If the remainder is non-zero, it has the same sign as the left operand.

`expr * expr`

Integer multiplication.

`expr #/ expr`

Single precision floating point divide.

`expr #* expr`

Single precision floating point multiply. All floating point operators assume floating-point operands.

Multiplicative operators group from left to right.

1
2
3 5.10. Additive operators.

4 These provide integer or floating point addition and
5 subtraction.

6
7 `expr + expr`

8 Integer add.

9 `expr - expr`

10 Integer subtract.

11 `expr #+ expr`

12 Single precision floating point add.

13 `expr #- expr`

14 Single precision floating point subtract.

15
16 Additive operators group left to right.

17
18 5.11. Relational operators.

19
20 `expr == expr`

21 Equal.

22 `expr != expr`

23 Not equal.

24 `expr < expr`

25 Less than.

26 `expr <= expr`

27 Less than or equal.

28 `expr > expr`

29 Greater than.

30 `expr >= expr`

31 Greater than or equal.

32
33 The result is one if the given relation between two integer
34 operands is true, and zero otherwise.

35
36 The following operators perform the same function for
37 floating point operands:

38
39 `#== #!= #< #<= #> #>=`

40
41
42 5.12. Logical and.

43
44 `expr && expr`

45 The result is an integer one if the result of both
46 expressions is non-zero, and zero otherwise.

47
48 The left-hand expression is always evaluated first. If its
49 result is zero, the result of the expression is zero and the
50 right-hand expression is not evaluated.

5.13. Logical or.

expr || expr

The result is an integer one if the result of either expression or both is non-zero, and zero otherwise.

The left-hand expression is always evaluated first. If the result is non-zero, then the result of the expression is non-zero and the right-hand expression is not evaluated.

5.14. "Query" operator.

expr1 ? expr2 : expr3

The first expression is evaluated. If the result is non-zero, the second expression is evaluated and returned, while the third expression is ignored. If the result of the first expression is zero, the third is evaluated and returned, while the second is ignored.

This is analagous to "if(expr1) expr2; else expr3", but has the advantage that it may be used in an expression. For example, a function to calculate the maximum of two numbers might be coded as:

```
max( a, b ) return( a > b ? a : b );
```

Grouping is left to right, so that

```
a ? b : c ? d : e
```

is equivalent to

```
a ? b : (c ? d : e)
```

5.15. Assignment operators.

lvalue = expr

Takes the one word result of the evaluation of "expr" and stores it in the word addressed by the lvalue.

lvalue <op>= expr

Is equivalent to the assignment

```
lvalue = rvalue <op> ( expr )
```

where <op> can be any one of

```
* / % + - << >> & ^ |
```

Note that neither floating point nor relational operators are included.

For example,

```
x *= a + b;
```

is the same as

```
x = x*(a + b);
```

In all cases, the expression is evaluated first, even though the operator in the assignment may have higher binding strength than an operator in the expression.

Assignments group right to left:

```
x = y = 0;
```

is taken as

```
x = (y = 0);
```

Remember that assignment is an operation, not a statement, and so is legal almost anywhere, including conditional expressions, such as

```
if( (x = y[i++]) == z )...
```

Note that parentheses are used in this case to alter the order of precedence. These are required in this case because the assignment operators have the lowest precedence, which means that they are evaluated last.

6. Implementation-dependent information.

The information in this chapter is subject to change.

6.1. Linkage conventions.

The B compiler's mechanism of performing a function call is rather different from the standard Honeywell calling sequence. In this chapter, we will describe the calling conventions in detail, so you can attempt to write or understand functions written in GMAP for the B environment.

The stack format described here is also used by PASCAL/66.

6.1.1. Function call.

The standard B function call looks like

```
    tsx1    sub  
    zero    s,n
```

where "s" is amount by which the stack pointer should be bumped and "n" is the number of words of arguments supplied. The compiler generates these numbers for B functions.

When a B function is running, its stack pointer points to a word of return information. Just below the stack pointer are two words, which are used by any function which is called from the current function. Above that point is a fixed size area for arguments and for AUTO variables, all addressed relative to the stack pointer. Stack space above the AUTO variables is used to hold temporaries created during expression evaluation. When one function calls another, the stack pointer is moved so the callee does not affect the state of the caller.

Before executing the TSX1 instruction to transfer to the function, the caller must first set up the argument values. The first and second arguments are loaded into the A and Q registers, respectively. Other arguments, if present, must be placed in the stack in such a way that they are available to the caller once the stack gets bumped. This is shown in the next section.

6.1.2. Entry.

The general convention for a subroutine entry looks like this:

	symdef	sub	
	symref	.entry	
sub	tra	*+1	value of function name
	tsx0	.entry	adjust stack, save linkage
	zero	frm,dbg	frame size, debug table
	staq	1,7	first two args passed in the AQ

"Frm" indicates the maximum number of words of stack that will be used by the function. A pointer to this word is saved in the stack, so that, if there occurs an asynchronous event (such as a fault), it will be possible to locate the top of the stack. "Frm" is also used in the check for stack overflow. "Dbg" is a pointer to the local debug symbol table, if present, or else zero, if not. It is used by the debugger and by the profiler, .PROFILE.

The code executed by .ENTRY (without the stack overflow check) looks like this:

stx1	-1,7	return for traceback
stx1	-2,7	actual return address
sxl0	-1,7	save pointer to frame size
ldx2	0,1	get "s" from function call
sar7	0,2,7	save old stack pointer
awd	0,2,7	bump stack
tra	1,0	return

There are two things to notice about this sequence of code. First, note that the return address is stored twice. This allows the actual return address to be modified if it is necessary to intercept a function return, while still allowing the debugger to print a correct traceback. SETEXIT and ALLOCATE are examples of library functions which set up interceptions of function returns. Second, note that return information is saved in the caller's stack frame. This means that certain library routines can avoid the overhead of a call to .ENTRY but still get access to two words of stack space "for free".

Register usage conventions are as follows. Address register seven is reserved for the stack pointer. Index register six is reserved for the coroutine package. Address registers four and five are reserved for use by the i/o package. Index register five may be used, as long as its value on entry is restored on return. The remaining registers (ar0-ar3, x0-x4) may be used freely, but note that a function which does not invoke .ENTRY is responsible for preserving the return address in index register one. You can use index register seven freely, except in the Pascal run time environment, in which you may use it freely only after copying over the display it points to.

By convention, the first two arguments are passed in the A and Q register. It is the responsibility of the called function to store them if necessary.

The hardware will only allow a STAQ instruction to work correctly if the address of the store is at an even word boundary. To ensure that the STAQ will work, the stack pointer is always initially set to an odd address, and must always be incremented by an even amount.

Once the STAQ is done, the stack is organized as follows:

-2,,7	-> used in next call
-1,,7	-> used in next call
0,,7	-> previous stack pointer
1,,7	-> first argument (initially in A register)
2,,7	-> second argument (initially in Q register)
n,,7	-> nth argument (placed in stack by caller)
n+1,,7	-> start of auto variables and temporaries

6.1.3. Exit.

When a B function has done its job, it returns using the sequence:

symref	.retrn	
tsx0	.retrn	(lar7 0,,7 - restore stack pointer)
rem		(ldx1 -2,,7 - get return address)
rem		(tra 1,1 - return)

Prior to this, the function may load a one word value into the Q register, which becomes the value of the function call.

A TSX0 instruction is used so that the interactive debugger, if in use, can determine the address at which the function returned, in order to find out the name of the function returning.

6.2. Internal representation of objects.

The material in this section is intended to help you understand how the code generated by B works internally.

The value of a function name is an external word containing a transfer instruction with the address of the first word of the function body in the upper 18 bits. Transfer of control to a function body always occurs by transferring to a word, which is expected to contain a transfer to the start of the actual function body.

The value of a label is a word containing, in the upper 18 bits, the address of the place to go to in the function body and zeros in the lower 18 bits. The transfer involved in a GOTO occurs indirectly through the label word.

The value of a pointer or address is a word, whose bottom 18 bits are taken as an address.

The value of string constant is a pointer to the text

of the string.

7. The B library.

One of the big advantages in using B is the availability of a large library of useful functions which simplify your programming problems and also supply a reasonable interface to the GCOS/TSS environment.

Every B library function you could reasonably expect to use has an explain file under "explain b lib". There is also an index of all documented routines.

Only the routines you need to get started using B will be discussed here and even then not all options may be treated.

Some functions may return a value; this is indicated in this section by showing an assignment to indicate a value is returned. Also, some functions are called with a variable number of arguments. If you want to use an optional argument, you must usually also specify any preceding optional arguments also. Optional arguments are shown enclosed in square brackets. For example, if a function is shown as

```
return = func( arg1 [ , arg2 , arg3 ] );
```

and you want to use "arg3", then you must also use "arg2". Sometimes, as you will see, the first argument, usually a "unit", may be optional. In this case, the called function craftily examines the first argument to see if it is a number valid for "unit"; if it is not, it adjusts its argument references accordingly.

7.1. .BSET - redirection of i/o.

Before your main program is entered, the run time initialization routine calls a function named .BSET which "predigests" the command line for the user program and also sets up any "redirection of i/o" requested on the command line.

In batch, .BSET looks for the command line on filecode CZ. In your jcl, you might have something like

```
$      data      cz  
command arg1 arg2 ...
```

BSET breaks the command line up into "arguments". An argument is either a string of non-blank characters, a quoted string, or a redirect request.

A redirect request has three forms:

<filename

.BSET will open the file for reading on B unit 0. Input for unit 0 will come from this file, rather than the terminal.

>filename

The filename is opened for writing. Output to B unit 1 will go to this file, rather than the terminal.

>>filename

Same as above, except that if the file already exists, output is appended to the file.

A quoted string is delimited by either single or double quotes. To get a quote inside a quoted string, either use the quote which is not the delimiter or else put in two of the delimiter characters.

.BSET collects the arguments which are not redirect requests and builds a vector of pointers to those strings. MAIN is later called by

```
main( argc, argv );
```

where "argc" is the number of arguments collected and "argv" is a pointer to the vector of strings. "argv[argc]" always contains the constant -1.

In addition, .BSET builds an external vector called .ARGTYPE, each cell of which contains a character giving some indication of the type of argument in the corresponding ARGV string:

type	character
string	' '
string in single quotes	'*''
string in double quotes	'"'
-option	'-'
+option	'+'
possibly signed number	'0'
string with = in it	'='

For example, look at the command line

```
go -r /myfile "a string" >b.out
```

MAIN will be called with ARGV set to four, since ">b.out" is not included. All writes on unit 1 will go to the file b.out, which is created if necessary. The ARGV and .ARGTYPE vectors are set up as follows:

argv[0] = "go"	.argtyp[0] = ' '
argv[1] = "-r"	.argtyp[1] = '-'
argv[2] = "/myfile"	.argtyp[2] = ' '
argv[3] = "a string"	.argtyp[3] = '"'

argv[4] = -1

and the contents of the remaining elements of the ARGV vector are undefined.

If you do not want to have .BSET, simply supply your own function definition of ".bset();" which will replace the library version.

Normally, however, you will want .BSET, because it greatly simplifies the task of handling command lines. In fact, there are even more powerful facilities built into .BSET for scanning command lines with arguments of specified types. As well, you may call .BSET to scan an arbitrary string. For full details, see the explain file "explain b lib .bset".

7.2. Introduction to input/output.

The largest class of functions in the B library are those concerned with input and output. Sequential input routines will read, and convert to ASCII if necessary, any sequential file in standard system format, including media 0, 2 or 3 BCD, media 5, 6 or 7 ASCII and media 1 compressed source decks (comdks). Output is ASCII (media 6) unless special precautions are taken.

The i/o package will create output files if necessary and "grow" them as required up to their maximum size or to the limit of the file space quota for a userid.

7.2.1. Units.

A B program may have several files open for reading or writing at the same time. Each file is associated with a number called a "unit", to which every i/o call implicitly or explicitly refers.

There are five units whose function is predefined and may not be altered by the user.

-5

This is an input unit whose origin is always the terminal in TSS or file code I* in batch. It may be used to force reading from the terminal or file code I*, even though the standard input may have been redirected to come from a file. Normally, you should use unit zero instead.

-4

This is an output unit whose destination is always the terminal in TSS or file code P* in batch. It is most often used to avoid possible redirection of i/o by forcing error messages to appear on a hard copy device. Normally you should use unit one instead.

-3

Used for console input in batch only.

-2

Used for console output in batch only.

-1

In TSS, all output directed at this unit behaves as if it were typed at command level. In batch, output to unit -1 goes to the execution report.

Unit zero is initialized as the standard input unit. In TSS, this is the terminal but input may come from a file if redirection of i/o is used. In batch, reads on unit zero come from file code I*, if it is defined and if input was not redirected. If I* is not present and there is no input redirection, unit zero is placed in the end-of-file condition.

Similarly, unit one is initialized as the standard output unit. In TSS, this is again the terminal and is subject to redirection of i/o. In batch, output to unit one goes to the printer, unless redirected.

Units two through 44 may be assigned by or to you, using the file opening call, usually to permanent or temporary disk files. It is permissible to open units zero or one, but the usual practice is to leave them alone, so they may be redirected.

7.2.2. Unit opening.

Before any i/o may be performed on a unit, it must be initialized by a call to OPEN, which is of the following form:

```
ret = open( [unit,] filename, action );
```

Normally, "unit" is not supplied, and OPEN finds a free unit, which it returns. If you do specify a unit, and that unit is already open, the state of the previous unit is saved on a stack; when the current use of the unit is closed, the previous state is restored and i/o may continue on that unit as if there had been no interruption.

"filename" is a pointer to a string containing the usual catalog/file string (e.g. "fbaggins/s/test.b"). An alternate name in quotes is used, if present. Permissions are effectively ignored.

"action" is also a pointer to a string containing characters which specify the access permissions required and the type of i/o to be done on the unit.

"ret" is the value returned by OPEN. If non-negative, the open succeeded and "ret" contains the number of the unit just opened. "ret" is negative and no unit is opened if there was a file access error or if there was an OPEN error.

Although not quite all of the various things accepted by OPEN are dealt with in this chapter, they are treated

fully in the explain file for OPEN.

Access actions: Most of the time, the only thing you need to specify to OPEN is how you want to access the unit. Here are the alternatives:

r

The 'r', for read, means you want to read on the unit.

w

The 'w', for write, means you want to write on the unit.

a

The 'a', for append, means you want to write on the unit, but by appending to what is already there. If there is nothing already there, or if appending is inappropriate, the result is the same as if you had used the 'w' action.

For ordinary sequential file processing, one of the above is the only action you need usually specify.

Mode actions: OPEN offers three ways to specify the mode of the unit being opened. They are as follows:

b

The 'b', for block i/o, requires that the unit being opened be a random access file.

s

This is used for so-called "string i/o". The "filename" argument is taken as a pointer to the start of a block or words in memory. The stream i/o functions will place characters into this block of memory, rather than transmitting them to a file.

i

Open the file according to its mode. It is then up to the program to determine, with the help of the function FILDES, whether the file is random or sequential and then make whatever i/o calls seem appropriate.

If none of these is supplied, the assumption is that a file is wanted and that it should be sequential, as is usually the case.

Error actions: Normally, OPEN never returns an error status, since the default action is to abort the program with a reasonably understandable error message.

The OPEN function lets you specify options which allow you to handle either file opening errors or file i/o errors or both. These options are in the form of characters which may appear in the action string:

1
2
3 e Arranges things so that a negative status is returned
4 on an i/o error.

5
6 f Sets things up so a negative status is returned on an
7 OPEN or file access error.

8
9 m Normally, when an error status is returned, no error
10 message is printed. The inclusion of the 'm' option
11 in the action string will cause the appropriate
12 routine to display an error message before returning
13 bad status to the caller. It has no effect if neither
14 the 'e' or 'f' action is used.

15
16
17 For instance, if you said

18
19 `open("/myfile", "rfm");`

20
21 and "/myfile" could not be opened with read permission, OPEN
22 would print an error message, then return bad status to the
23 caller.

24 In the case of OPEN errors, you will most likely get
25 back a number which is the negative of the file system error
26 status. For example, OPEN would return -5 for "permissions
27 denied". In addition, OPEN itself is liable to return any of
28 the following special error statuses:

29
30 -64 - too few arguments
31 -65 - no free unit
32 -66 - open append error

33
34
35
36 File access conventions: There are a number of file
37 access/create conventions for TSS which you should be aware
38 of:

39 1. Search rule: A quick access name is one which con-
40 tains no slashes or dollar signs and does not have an alt-
41 name. It must also be less than nine characters long; if
42 not, it is considered to be in error. If you are opening a
43 file and the name is a quick access name (e.g. "b.out"), the
44 file accessor first searches the AFT for a file of that
45 name. If not found, the file accessor searches for a file of
46 that name in the User-Master-Catalog of the current userid.
47 If the file name is not of the quick-access form, it is as-
48 sumed to be the name of a permanent file.

49 2. Create rule: If the search fails and the request is
50 to write or append, OPEN will attempt to create the file. If
51 the file name is of the quick access form, OPEN will try to
52 create it as temporary; otherwise it tries to create a per-
53 manent file.

54 3. AFT rule: If the file was already in the AFT when

accessed, or if the file is a temporary file created during the access, it is left there when the unit is closed. If the file was not in the AFT initially, and the file is permanent, it is removed from the AFT when the unit is closed. You may override this by including in the action string either the character 't' (for transient) to force deaccess, or the character 'k' (for keep) to force the file to be kept in the AFT.

7.2.3. Unit closing.

When you are through with a unit, you may want to close it explicitly by calling CLOSE:

```
close( unit );
```

For sequential stream output units, CLOSE flushes the output buffer if necessary, with an end-of-file mark written if appropriate. A unit associated with a disk file has the file deaccessed, if required. CLOSE releases the i/o vector after checking to see if there was a prior use of the unit which had been interrupted and saved. If there was, it is restored and i/o may then proceed on that unit as if there had been no interruption; otherwise, the unit is free for further allocation.

When your MAIN function terminates, or when you call EXIT directly, all open units are closed automatically. Note that if you hit break, things are set up to call EXIT, unless your program has established its own break handling procedure.

Any attempt to read data from a unit which is not open results in the library routine called returning a value which indicates nothing happened. Similarly, output to a unit which is not open tends to vanish.

7.2.4. Unit switching.

Since some input/output calls may not specify a unit directly, the i/o package maintains a default input unit and a default output unit to which these calls implicitly refer.

Initially, the i/o package is set to read from the standard input (unit zero) and write on the standard output (unit one).

Any successful call to OPEN changes the default input or output unit.

A library function which takes a unit as one of its arguments will change the reading/writing unit for the duration of the call and restore the previous value before returning to the caller.

1
2
3 User control over unit switching is supplied by

4
5 `old.unit = .read([new.unit]);`

6 If "new.unit" is given, it becomes the current read
7 unit. The number of the old read unit is returned.

8
9 `old.unit = .write([new.unit]);`

10 Works the same way as .READ, except that it applies
11 to the default write unit.

12 13 7.3. Sequential stream i/o.

14
15
16 This section describes the body of routines oriented towards
17 handling input/output on terminals or standard system format
18 sequential files.

19 The i/o package reads almost any media and arranges
20 things so that the using program sees only a stream of ASCII
21 characters. For instance, BCD printer slews and strange es-
22 capes in media 3 files are correctly detected and converted
23 on input, as are ASCII slews in media 7 print image files.
24 Compressed source decks are handled correctly, but the way
25 it handles object decks is probably not very useful. Media 0
26 is always taken as variable-length BCD.

27 On output, the i/o package writes media 6 ASCII unless
28 special action is taken as described in the explain file for
29 OPEN. If writing to SYSOUT in batch, output is media three
30 (BCD printer format).

31 You have the option of opening a unit to read, to write
32 or to append as follows:

33
34
35 `unit = open(filename, "r");`

36 open a unit for reading, requesting read/concurrent
37 permission.

38
39 `unit = open(filename, "w");`

40 will open a unit for writing, requesting
41 write/concurrent permission.

42
43 `unit = open(filename, "a");`

44 will open a unit for writing so that data written by
45 the program is appended to the end of the file. If
46 the file is null to start with, OPEN treats the si-
47 tuation just like a regular open for writing.

7.3.1. Terminal vs file.

There are a few special features of and differences between file and terminal i/o of which you should be aware.

A logical record consists of a string of characters followed by a "record terminator", which is one of '*n', '*r', or '*f'. Of these, '*n' is never present on an input file; instead, the '*n' is automatically supplied by the i/o package to indicate the end of a record. On input from a terminal, you separate logical records (lines) by using either the "return" or "line-feed" key.

End of file on a file is signalled by the presence of a special record at the end of the file. End of file on a terminal is signalled by a line whose first (and usually only) character is an ASCII file separator (FS - octal 034) character, the same as is used by TSS GFRC. On most ASCII terminals, including the Teleray and Volker-Craig, a FS character is transmitted by typing the 'cntl' and '\' (backslash) keys simultaneously, followed by a carriage return. On certain others, you get an FS by typing 'cntl', 'shift', and 'L'. You can't signal end of file on a 2741.

Sequential file output is written in GFRC standard system format with 320 word blocks. An end of line ('*n') character written to a sequential disk file signals the end of a record but is not itself placed in the record; all other record terminators do get written out. A new logical record is started following the receipt of any record terminator. If more than 1272 characters are written in a logical record, the i/o package will generate partitioned records to permit the logical record to span more than one physical block. A 320 word buffer is written out only when it is necessary to start a new buffer or when CLOSE is called.

Terminal output is buffered by TSS. However, if your program issues a read from a terminal, the i/o package arranges that all output sent appears on the terminal before it is "unlocked" for input.

When writing to the batch console, a '*v' is translated to a '*n', but does not cause the current line to be flushed, in order to let you either write a couple of lines or else write a line and then read a line, without having to worry about some other process affecting the console between writes or between read and write.

7.3.2. Stream i/o functions.

`char got = getchar();`

Returns the next character from the current read unit. Returns the character '*0' if the current reading unit is closed or at end-of-file. Most of the input routines described here behave as if they make repetitive calls to GETCHAR.

`char = ungetc(char);`

Sends a character back to the current read unit, so that the next call to GETCHAR will return the last character put back.

`char = getc(unit);`

Same as GETCHAR, except the reading unit is switched to "unit" for the duration of the call, then restored.

`char_out = putchar(char);`

Sends the character supplied as its argument to the current writing unit. PUTCHAR also returns its argument word as its value. The argument word may actually contain up to four non-zero characters. PUTCHAR will output as many characters as there are in the word.

`char = putc(unit, char);`

Same as PUTCHAR, except PUTC switches writing units for the duration of the call.

`string = getstring([unit,] string [,maxl]);`

GETSTRING gets the next line of input, or the remainder of the current line of input if GETCHAR has already been called. The newline at the end of the line is not returned; instead it is replaced by a '*0' to mark the end of the string. "string" is taken as a pointer to a vector long enough to hold the string. "unit" is used if supplied; otherwise the current read unit is used. If "maxl" is given, only the first "maxl" characters are returned, with a '*0' tacked on to the end. If the unit is closed or at end-of-file, GETSTRING returns zero; otherwise it returns "string". The string is the nullstring if a line contains only a newline. If GETCHAR was not called, GETSTRING has the effect of returning the next line of input from the terminal or the next logical record from a file.

2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57

```
string = getline( [unit,] string [, maxlen] );
```

Same as GETSTRING, except the line terminating '*n' is included in the string, just before the string ending '*0'.

```
printf( [unit,] format-string, arg1, arg2, ... );
```

PRINTF is the most frequently used means of doing output in the B library. If "unit" is not supplied, the default writing unit is used. If "unit" is given, PRINTF temporarily switches writing units, but restores the original state upon return. "format" is a string describing how the arguments are to be output. It may contain any combination of literal characters and formats. A format is of the form "%nnx", where "nn" is an optional count, and "x" is one of the following characters:

b - The corresponding argument is taken as a pointer to a string of BCD characters, which is to be translated to ASCII and printed. Since BCD strings do not have a string terminator, a count of six is assumed if not supplied. Trailing blanks are stripped.

c - The corresponding argument is printed as an ASCII character. The count option is not applicable. The argument word may actually contain up to four non-zero ASCII characters, which will be printed.

d - The corresponding argument is taken as a decimal integer which is converted to a string and output.

f - The argument is taken as a floating-point number to be converted and output. If your program does not use at least one floating-point operator, you must include an "extrn .float;" to force the loading of the floating-point output routine.

o - The contents of the argument word are output in octal.

s - The corresponding argument is taken as a pointer to an ASCII string, which is transmitted to the output unit stripped of its trailing '*0'.

If a character in the format string is not part of a format, it is printed as it appears. If a format does not have a corresponding argument, it is printed as a literal string. To print out '%', you must use '%%'. There is more to PRINTF than is given here; for full details, see the explain file "explain b lib printf".

string = putstring([unit,] string [,maxl]);

Works in much the same way as GETSTRING. The '*0' which marks the end of the string is not output. If you want to write out a logical record and the string to be output does not end with a '*n', you should usually follow a call to PUTSTRING by a "putchar('*n');".

number = getnum();

GETNUM returns the next, possibly signed integer number from the input stream. It calls GETCHAR until it has skipped over all blanks, tabs or newlines. If the character is not a digit or a sign, zero is returned, indicating no number was found. If a sign was found, and the next character is not a digit, zero is returned again. Otherwise it collects numeric characters until a non-digit is found, then converts the numeric string it has collected to binary and returns that number as its value. The external GETN.A has the value 1 if a valid number was found. The external GETN.L contains the last character read.

putnum(number);

PUTNUM converts the assumed binary integer which is its argument to a character string and directs the string to the current output unit.

status = eof([unit]);

Returns a non-zero value if "unit" is at end of file. If "unit" is not given, the current reading unit is used. Once a unit is open, end of file is set only after an attempted read results in the detection of that condition. If "unit" is given, and the unit is an output disk file, a logical end of file is written and a new block begun.

There are a few other routines which must be mentioned, but which will not be described in detail here. READF does formatted input, somewhat like PRINTF in reverse. It also supplies the only convenient means of reading in a floating point number. GETNUM and PUTUCT can read or write octal numbers.

GETREC and PUTREC allow you to obtain/transmit a logical record, including record control word, without any intervening processing by the i/o package. You must understand standard system format before you attempt to use GETREC or PUTREC.

7.4. Random file i/o.

When using random-access files, your program is responsible for all input or output done on the file. The basic unit of i/o is the sector of 64 words. Any number of sectors may be read or written at one time.

To open a random file for reading, use the call

```
unit = open( filename, "rb" );
```

The rules are the same as for the regular OPEN call, except that the character 'b' in the action string indicates that the file is to be accessed as random. The action 'b' stands for "block" - we would have used 'r' for "random", but it is already taken for "read".

To open a file for writing or reading and writing, use

```
unit = open( filename, "wb" );
```

Reading is accomplished by

```
status = read( unit, buffer, sector, nwords );
```

"buffer" is a pointer to a vector into which the data will be read, "sector" indicates at what sector in the file the read will start, and "nwords" indicates how many words will be transferred. The first sector number in the file is zero. If the status returned is non-negative it is a count of the number of words transmitted; otherwise it is the negative major (bad) status from the i/o.

Writing is accomplished by

```
status = write( unit, buffer, sector, nwords );
```

The arguments have the same meaning as those for READ. The GCOS disk i/o system always writes a multiple of 64 words. If the number of words you transmit is not a multiple of 64, the unused fraction will be filled with zeros on writing.

7.5. String operations.

The B compiler recognizes the existence of strings only in that it handles string constants. All operations on strings are handled by function calls. Recall that a string is a sequence of characters packed four to a word in a vector and terminated by the ASCII '*0'.

Functions are available to permit processing characters in a string in a "random" manner or character by character. We will also mention several useful string utilities.

7.5.1. "Random" string processing.

```
ch = char( string, i );
```

Returns the *i*th character in a string pointed to by "string". The count always starts at zero.

```
ch = lchar( string, i, char );
```

Replaces the *i*th character in the string pointed at by "string" with the character "char" and returns as its value the character supplied.

Where characters are in BCD format (6 characters per word), you should use the function CHARB instead of CHAR and LCHARB instead of LCHAR. The calling sequences are the same, but remember that they take and return BCD characters, not ASCII.

7.5.2. Sequential string access.

By using one of the following calls, it is possible to "open" a string in such a manner that calls to regular sequential i/o routines place characters in or return characters from a string. This method is faster than using CHAR/LCHAR, because the implementation uses hardware "tallies". The action 's' stands for "string".

```
unit = open( string, "rs" [,pos] );
```

Opens a string so that calls to GETCHAR will return characters in the string. A call to GETSTRING returns all characters up to but not including the next '*n' or else up to the terminating '*0'. When the string is exhausted, the unit is in EOF status. If you want to start getting characters at some point other than the first character position, use the optional starting character position "pos". Any library function which obtains characters from an i/o unit will also work even if the unit is a string.

```
unit = open( string, "ws" [,pos] );
```

Opens a string so that calls to PUTCHAR place characters in the string. PRINTF, PUTSTRING, PUTNUM and other functions will also send characters to the string. The function of "pos" is the same as described above.

```
unit = open( string, "as" );
```

Locates the terminating '*0' of "string" and sets things up so you start writing into the string at that point. It is up to you to make sure that the vector pointed at by "string" is large enough to hold whatever your program puts into it.

2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57

```
print( string, format, arg1, arg2, ... );
```

Same as PRINTF, except it directs its output to a string, instead of an open output unit.

When you call CLOSE on a unit open for output to a string, a terminating '*0' is placed in the string.

7.5.3. String utilities.

```
string = concat( string, s1, s2, ... );
```

Concatenates the strings "s1" through "sn" together and places them in string "string". The output string may be used as input, providing it appears first in the list of strings to be concatenated. When called with only two arguments, CONCAT efficiently copies one string into another. CONCAT returns its first argument as its value.

```
val = nullstring( string );
```

Returns a non-zero value if the string contains only the end-of-string character '*0' and zero otherwise.

```
val = equal( string1, string2 );
```

Returns a non-zero value if the two strings supplied are identical, and zero otherwise.

```
count = length( string );
```

Returns the number of characters in the string pointed at by "string", not including the terminating '*0'.

```
result = compare( string1, string2 )
```

This compares two strings. If they are equal, COMPARE returns zero. Otherwise, it returns a value one greater than the position in "string1" where they differed. The returned value is positive if "string1" is greater than "string2", or negative if "string1" is less than "string2".

```
pos = any( c, string [,i] );
```

If the character "c" appears anywhere in the string "string", with scanning starting at character position "i" (which defaults to zero if not supplied), then the position of the character in the string is returned. If the character is not found, then the value "-1" is returned.

newpos = getarg(arg, string, pos [,delim]);

Starting at position "pos" in "string", GETARG obtains the next group of characters ending with a blank and places it in the string pointed to by "arg". It returns the position in the string where the scan stopped, so that it can be called repeatedly to obtain successive "arguments" from the string. If "delim" is supplied, it is taken as a pointer to a string containing the delimiters which will cause the scan to stop; the string must include a blank if you want the scan to stop on a blank. Leading blanks are ignored. GETARG is useful for scanning a command line.

There are a number of other functions which perform string operations, including NUMARG, which scans off numbers instead of character strings; ADDCHA, which appends a character to the end of a string; and READF, which can do formatted input from a string. All of these have explain files.

7.6. Storage allocation.

Library functions are supplied which allow you to dynamically obtain or release memory in the free storage pool.

addr = getvec(n);

Obtains from free memory a vector of length "n" plus one words and returns a pointer to the vector. Once acquired in this manner, the block may be referenced using subscripting, as in

```
x = getvec( 63 );  
x[1] = a[3];
```

rlsevec(addr, n);

Undoes a GETVEC by releasing the "n" plus one words pointed to by "addr" back to the free memory area.

All memory allocation is done by manipulating a free list. The free list initially includes the so-called "core hole". You can return via RLSEVEC any space which is not on the free list, as long as the address of the space is greater than the load address of RLSEVEC. If you attempt to release memory which is already on the free list, in whole or in part, RLSEVEC will immediately abort.

GETVEC obtains more storage from the operating system as required. In TSS, a subsystem is aborted with the message "not enough core to run job" if a request for memory cannot be satisfied. In batch, GETVEC aborts with a "OK" abort code if a request for memory is denied.

Finally, we will mention three useful routines, each

described fully by an explain file, which use these calls: GETMATRIX will construct and return a pointer to a multidimensional matrix; ALLOCATE can be used to obtain a dynamic array which will automatically disappear when a function returns; and RELMEM will release any free memory back to the operating system, in order to reduce program size.

7.7. Character code conversion.

Two functions are supplied to let you transliterate BCD into ASCII and vice versa. A BCD string consists of a vector of words containing the characters packed six to a word, left-adjusted and padded with blanks. There is no equivalent to the '*0' in a BCD string.

```
ptr = ascbcd( output, count, input );
```

Takes "count" characters from the ASCII string "input", transliterates them to BCD and places them in "output". If a '*0' is encountered before "count" is exhausted, blanks are supplied and also used to pad the BCD string to a word boundary. "input" and "output" must be pointers. "output" is returned.

```
ptr = bcdasc( output, input, count );
```

Takes "count" BCD characters from "input", transliterates them to ASCII and places them in the ASCII string "output". Any trailing blanks are deleted and the end of string delimiter '*0' is placed at the end of the ASCII string. "input" and "output" must be pointers. "output" is returned.

7.8. Call FORTRAN.

The function CALLF provides the ability to call FORTRAN subroutines, or any routine which uses the GCOS CALL conventions. However, routines so called must not be called recursively.

```
intval = callf( &routine, &arg1, &arg2, ... );
```

Converts its arguments to the form of a standard GCOS CALL macro and calls the named "routine". "routine" must be referenced in an EXTRN statement and must be passed by address as shown. The arguments must be passed by address also. That is, if an argument is not a vector pointer, you must say "&arg". Constant values must be assigned to a temporary before being given to CALLF, since you can't say something like "&2". The value of a CALLF is the logical or integer value returned, if the routine called is a function subroutine.


```
floatval = callff( &routine, &arg1, &arg2, ... );
```

Works exactly the same way as CALLF, except that it must be used for function subroutines which return a floating point result.

As usual, you are responsible for ensuring the correct number and type of arguments are passed.

7.9. DRLs and MMEs.

For those (hopefully rare) occasions when you need to perform system calls directly, here are the functions provided to let your B program execute DRL or MME system calls in a reasonable manner:

```
drl.drl(number [,arg1, arg2, ...] )
```

Allows direct access to the DRL functions. "Number" is the DRL number to be executed, and any following arguments are the words to follow the DRL. The A and Q registers are set to the values of the externals DRL.A and DRL.Q respectively. After the DRL has been executed, these externals are set to the contents of the A and the Q. It is possible to use a DRL which requires an error exit or a place to go to, since the DRL is executed in the stack, using the stack pointer of the caller. For example:

```
%b/manif/drls
```

```
...
auto buf, vec[2];
...
buf = getvec(600);
open( 9, "gcos3/gcos-hi-use", "r" );
aft.name( 9, vec );
vec[2] = ".mbrt3";
drl.drl(restor_, buf<<18 | 1, buf<<18 | 1,
        (tra&0777777000000)| (buf+512));
tra:
printf( "%24b*n", buf+4+status*4 );
rlsevec(buf,600);
```

This code sequence, which obtains a batch error message by locating it in the batch error message module, uses the value of a label to supply a return address to DRL RESTOR. Note that, in this particular case, you could have called the library function .RESTR, replacing the call to AFT.NAME and DRL.DRL with ".restr(9, ".mbrt3", 1, buf);".

mme.mme(number [,arg1, arg2, ...]);

Functions in exactly the same manner as DRL.DRL, except that it uses externals called MME.A and MME.Q.

8. Using B.

8.1. Compiling/running.

The B command is the main tool for preparing B programs. If given a source file, it will call the compiler to read the source and generate a set of object modules. It may call the random library editor RANEDIT to place or replace modules in a library. Unless there are fatal compilation errors, it always calls the TSS loader to prepare a load module. Only the most common use of the B command is discussed here. For full details, see the TSS explain file "explain b command".

To compile and load a source file, just type

```
b srcfile
```

where "srcfile" is the name of a sequential file containing B source statements. If there were no fatal errors, the load module is left in a random file called ".h", which is created as temporary if necessary. If you have a quick-access permanent file called ".h", it is used instead. This file is "grown" automatically by the TSS loader as required.

You could have forced the B command to initiate execution by entering

```
b -go source-file
```

but, if your program plans to interpret a command line, it is preferable to use the command GO, like

```
go arg1 arg2 ....
```

which will run the load module in ".h" with the given command line.

There are a few other options in the B command which you may find useful. If you find you are allocating too many AUTO variables, so that your program violates the stack limit and overwrites your own code, you can specify a larger stack by using the "Stack=nnn" option, as in

```
b src.b stack=700
```


Fully debugged production programs need not carry the debug tables with them when running. Use the "-Nodebug" option to turn off the loading of debug tables.

```
b -nodebug src.b
```

If you change one routine in a program, you usually have to recompile the whole program. It is sometimes more convenient to store routines in a random library. That way, you need only recompile one routine or one group of routines to make a change. The B command provides an interface with the RANEDIT subroutine library editor. To start with, if you type

```
b src.b ranelib=/lib -clear
```

the routine or routines in src.b will be edited into the library "lib", which will be created if necessary according to the usual B file accessing conventions and initialized (cleared) by RANEDIT. To add new routines or replace old ones, you simply enter

```
b src1.b ranelib=/lib
```

Your program can load from a user library by specifying the "Library=" option, as in

```
b src.b library=/lib
```

When the loader is called, the library specified by the "r=" option is searched along with any other libraries given using the "l=" option. Libraries are always searched in the order given on the command line. To delete routines from a library, it is necessary to use the TSS command RANEDIT (see the TSS explain file).

The options to the B command are of the forms

```
keyword=string  
-keyword
```

In both cases, the keyword may be abbreviated using the following rule: In the explain file, a keyword is shown with upper and lower case letters. A valid abbreviation must include those letters which are in upper case, along with any other letters in the order in which they appear. For example, valid abbreviations of the "Ranelib=filename" option include

```
rane=filename  
rlib=filename  
r=filename
```


Various options may, of course, be combined onto one command line and abbreviated, as in the following example:

```
b cmdlib/s/roff h=cmdlib/roff l=b/xlib -n -d
```

This command line uses the option "Hstar=filename", which allows you to designate the file into which the generated load module will be placed.

8.2. Debugging.

Rarely will you find that a newly-written B program works the first time you try to run it. Typically, it may do nothing, it may go into an infinite loop, or it may abort with memory fault or some such hardware-detected error.

The most useful tool for debugging programs is BOFF, which stands for "B Obscure Feature Finder". You can use BOFF in three modes: to inspect and/or patch the core image file prepared by the loader, to monitor the progress of your program as it is running, or to inspect a TSS dump file, called "abrt", after your program has died.

In this section, we will look at some of the basic procedures to follow when debugging a program, using BOFF. For full details on BOFF, see the explain files.

Usually, the first thing that happens is that you try to run your program and it aborts, leaving an abort file containing the dump in the AFT. To look at the dump, type

```
boff -abort
```

BOFF will display the fault type, the value of the instruction counter at the time of the fault, and, if debug tables are available, the abort location expressed as an offset from the load table name whose address is closest to and less than the abort location. For a batch abort file, the fault type is always "incorrect primitive"; however, you can determine the actual fault by looking at the execution report.

In BOFF, displayed numbers are shown in octal, but when snapping memory locations you have the option of changing the form of the display. A number which BOFF reads from the terminal may only be an integer, which is taken as octal if it has a leading zero, and decimal otherwise. In other words, every time you type in an octal number, you have to stick a zero in front of it.

The first thing you will want to do with an abort dump is to try to get a traceback of calls in reverse order. Type

```
:t
```

For each call, the traceback shows the stack pointer used during the call, the function name, the arguments, and the symbolic address from which it was called.

There are a couple of things you should keep in mind when looking at a traceback. You may see function names in the upper part which are not part of your program; this means your program died in a library function. The arguments displayed for a library routine may not look like what you expect; this may be due to your program passing bad arguments, but it is also quite common for a library routine to modify an argument or use it for a different purpose.

If you are lucky, your problem will be a bad argument and a solution will suggest itself. More likely, you will want to inspect the state of the various variables in the program. You can look at any external value at any time, but you can only look at the arguments and AUTO variables of the current function, which is initially the topmost function in the traceback. BOFF provides two commands which let you move the current function context either up or down in the traceback list.

For example, suppose you have the following traceback:

```
00221  open (0150)      rets to 003306 (readf +0410)
00151  readf (01162,07777777777766) rets to 001160 (func+010)
00147  func (07777777777766) rets to 001140 (main +06)
00145  main (056150000132)   rets to 001321 (..... +0114)
```

The program has died in the library routine OPEN, which was called from another library function, READF. In fact, the cause of this particular abort was a bad argument to READF. The initial context is that of OPEN, which, being a library routine, probably has no debug tables. To switch the context to MAIN, tell BOFF to move down the stack three functions by giving the command

```
3:d
```

From there, to get to FUNC's context, you would tell BOFF to move up one function by typing

```
1:u
```

As you can see, context positioning on the stack is specified by giving the number to move, followed by a colon, followed by the direction, which is either 'u' or 'd'.

Suppose you tell BOFF to move up or down too far? BOFF will leave you in a reasonable place. To find out where, type

```
:va
```

which will show the name of the function whose context is current, and, if there is a debug table for the function, the names and values of the arguments, automatic variables, and labels of the function. Remember that "va" means "View

Autos".

To snap memory in BOFF, use a command formed as follows:

```
[addr][,nws][\mode]
```

where the square brackets indicate an option.

The "addr" must be an lvalue, which means that it must be in either the form "name" or the form "*nnn", where "nnn" is a number. If you do not supply an "addr", BOFF uses the last address displayed plus one.

The "nws" is usually a number. If not given, it defaults to one.

The "mode" is a possibly null string of characters, which can include:

- a - address format
- b - BCD constant
- c - ASCII character constant
- d - decimal
- f - floating point
- i - interpret as machine instruction
- o - octal
- s - display B string

If the string is null, the mode is octal. If '\ ' is not available, you can use ' ' instead. If "mode" is not present, the last mode is used. The initial mode is octal.

To snap a variable called "x", you might type

```
x
```

To snap it in octal, BCD, and ASCII, you would say

```
x\obc
```

To snap the contents of the ten word vector "y[10]", you would enter

```
*y,11
```

To snap the vector in octal and address format, you would type

```
*y,11\oa
```

To snap the first 20 words of the function MAIN, you would say

```
main,20\i
```

The lvalue of the expression "main" is the address of the Waterloo

routine "main", which is where BOFF starts snapping. To snap the next 20 words, enter one of the following two alternatives:

```
*(&main+20),20  
,20
```

Since no "mode" is given, the last mode is used.

To snap 20 words, starting from location 110 (octal), you would type

```
*0110,20
```

When you go on to read the full BOFF documentation, you will find that you can construct arbitrary expressions composed of names, integer numbers, and a subset of the B language operators.

At this point, we have covered the basic information you need to know to use BOFF on an abort file. It is more likely that you will want to use BOFF on your program as it is running. In this mode, not only can you use everything mentioned so far, but you can also set breakpoints and execute function calls from BOFF.

To run your program with BOFF, give the command

```
boff:.h arg arg ...
```

BOFF will set a breakpoint at MAIN, so you will get control in BOFF just before MAIN is executed.

When you get control, the current context for the ":va" command will be MAIN, the function at which the breakpoint is placed. The commands for moving up or down the stack, for snapping memory, and for printing a traceback will otherwise work as described above.

BOFF makes it very convenient to put a breakpoint on a function. If your program consists of a number of small functions, rather than one monolithic MAIN function, it will probably be easier to debug.

To put a breakpoint on a function, use the ":bf" command as follows:

```
name:bf
```

where "name" is the name of the function. To delete such a breakpoint, use the command:

```
name:bd
```

You may also want to know about what value a function returns. You can set a function return breakpoint, but only at the function entry breakpoint for that function, by typing the command

:br

Finally, you can call, from BOFF, any function defined in the program, with arguments. For instance, you might say

```
printf("%c %c*n", 'abcd', c )
```

This example shows that not only can BOFF read numbers and variable names, but it can also read B string and character constants. BOFF will always print out the value returned by the called function. If the callee does not deliberately return a value, the displayed value can be disregarded.

8.3. Compiler/loader interface.

When processing a source program, the compiler generates not one but a set of object decks and places them onto a temporary file called "b*" - the input file for the loader. This file is also used as input to the random library editor RANEDIT, if it is called.

The compiler always generates an object deck containing the B stack area, which is either 500 words or the size specified in the "Stack=nnn" option. This object deck also contains the externals defined before the first function definition, if any.

A separate object deck is generated for each function. The deck includes a SYMDEF for the function and a SYMREF for each name mentioned in an EXTRN statement or used in a function call.

An object deck is also generated for each group of externals between function bodies and one for the group of externals after the last function body, if any. The deck includes a SYMDEF for each defined external and a SYMREF for any external name referenced in an initializer list.

8.4. Using tabs for readability.

When you type in a B program, you will probably want to leave indentations in order to make clear the order of nesting of your source statements.

Spaces are the logical thing to use, but they are tedious to type and it is difficult to be consistent. At Waterloo, we usually suggest you use an ASCII tab character as one unit of indentation. This has the advantage that, when you use TLIST to get a listing of the source, the TLIST command automatically expands tabs into the right number of blanks, so your program comes out indented the way you want. Also, you can use the "OTD" directive inside the QED text editor so that it, too, expands tabs when displaying a line.

It is not possible to use a "tab character" which is not an ASCII tab.

8.5. Some pitfalls.

If you have a floating point value, it is not a good idea to say

```
if( floatval ) ...
```

because the code generated checks to see if "floatval" is logically non-zero, rather than to see if it is equal to a floating-point zero. Since a floating point zero may not in fact be a word containing all zero bits (since the exponent may be non-zero and is part of the word) it is better to try

```
if( floatval != 0.0 ) ...
```

In general, floating point is tricky to use in B, since there can be no type checking. You must constantly watch out for erroneous constructs such as using "-3.0" instead of "#-3.0".

Also, here is a common pitfall in the use of string constants. If you say

```
auto x[20];  
x = "a string";
```

The cell x is changed to point to the storage occupied by the string and you lose the ability to address the 21 words originally reserved for the vector. What you really want to say is

```
auto x;  
x = "a string";
```

Alternatively, if you had wanted to initialize the vector with the string, you should have used the library function CONCAT to copy in the string:

```
auto x[20];  
concat( x, "a string");
```

or else you could have defined "x" as an initialized external.

Finally, something should be said about the size of a vector and the length of a string.

When you declare a vector of size "n", you know that it will actually occupy "n + 1" words, because the vector is indexed starting at zero. Every library routine to which you must pass the size of a vector observes exactly the same convention.

In practice, if one needs a vector of size "n", then one declares it to be of size "n", and then ignores the zeroth or the nth word. Thus a FOR loop indexing through the vector might run in either of two ways:

```
for( i = 0; i < n; ++i ) ...
```

```
for( i = 1; i <= n; ++i ) ...
```

Strings also can be indexed into, using library functions, using a zero origin, but at first it might appear to you that a string with "n" characters in it has length "n", rather than "n + 1". For example, the string "abcdef" has six characters and its length is six. But recall that, by definition, a string is terminated by a '*0' character, which you do not see. If you include the trailing '*0' in the count, then a string of length "n" actually contains "n + 1" characters.

All library functions which require the length of a string need the number of characters, not including the '*0'. The library function LENGTH returns just that number.

APPENDIX A

Escape sequences

There are two sets of escape sequences, one for use inside string or character constants, and the other for use outside.

1. Escape sequences are used in character constants and strings to obtain characters which for one reason or another are hard to represent directly. Here are the escapes:

*0	end of string (ASCII NUL = 000)
*e	end of string (ASCII NUL = 000)
*{	{ - left curly brace
*}	} - right curly brace
*[[- left square bracket
*]] - right square bracket
*t	tab
**	*
*'	'
*"	"
*n	newline
*r	carriage return (no line feed)
*f	ASCII formfeed
*b	backspace
*v	vertical tab
*x	rubout (octal 177)
*#nnn	nnn is 1-3 character octal number

2. The following are escapes used outside character and string constants on terminals (such as the 2741) which do not have on their keyboards some of the characters used by B. If you use QED, it is nicer to use the QED escapes for these characters, so that when you shift to an ASCII terminal you can see the characters the way they ought to appear.

\$({ - left curly brace
\$)	} - right curly brace
\$([- left square bracket
\$>] - right square bracket
\$+	- or-bar
\$-	^ - up-arrow (or cent-sign)
\$a	@ - at-sign
\$'	` - grave accent

APPENDIX B

Binding strength of operators

Operators are listed from highest to lowest binding strength; there is no order within groups. Operators of equal strength bind left to right or right to left as indicated.

```
name const primary[expr] primary(arglist) (expr) [LR]
++ -- * & - ! ~ #- # ## (unary) [RL]
>> << [LR]
& [LR]
^ [LR]
| [LR]
* / % #* #/ (binary) [LR]
+ - #- #+ [LR]
== != > < <= >= #== #!= #> #< #<= #== #!= #>=
&&
||
?: [RL]
= += -= etc. (all assignment operators) [RL]
```


APPENDIX C

B compiler error message

This is a list of diagnostics known to be generated by the B compiler. There may be others.

In each description, "nn" means a line number, while "name" is some identifier name. The name of the source file is usually also given.

Any message not preceded by "warning: " is a fatal error. If there is a fatal error, neither the loader nor the random library editor will be called.

Diagnostics:

syntax error at line nn [in file <name>]

This is the most common diagnostic and it could mean almost any kind of error. Most often, it means a semicolon is missing or the number of open curly braces "{" does not match the number of close curly braces "}", in which case the line number will be the number of the last line in the last file being processed plus one. This may be due to neglecting to end a string constant, character constant or comment. You also get this message if you use a keyword in an inappropriate context, such as an AUTO statement, if you neglect to define a manifest, or if you attempt to redefine a manifest.

<identifier> undefined in function <name>

An identifier in the named function has not been referenced by an EXTRN or AUTO statement and has not been used as a label. The line number given is the last line of the function being compiled.

warning: /* inside comment ...

This is a warning only, but there will probably be a syntax error later on, since comments may not be nested. After reading a "/*", the compiler skips all text until a "*/" is encountered; if there is a comment inside a comment, then the compiler will attempt to compile the remainder of the outside comment.

end of file in comment

This usually indicates that you forgot to end a comment with the terminating "*/".

warning: newline in constant not preceded by *

The most probable cause is that you forgot to terminate a string or character constant with the appropriate delimiter. If this is the case, you will surely get a syntax error later. If you want a "real" newline inside the constant, but no warning, use the escape sequence '\n'. If the constant is a string constant which is too long to fit on one line, precede the newline with a '*'; the newline will be discarded. When the warning is issued, the newline is

3 kept.
4 invalid octal constant
5 An integer beginning with the digit zero, which is
6 thus assumed to be an octal constant, contains a
7 character other than the digits zero through seven.
8 character constant too long
9 A character constant may not contain more than four
10 characters, although each character may be a two
11 character escape sequence.
12 bcd constant too long
13 A BCD constant contains more than six characters.
14 exponent too large in constant
15 The exponent of a floating point constant is too
16 large or too small to represent in the hardware.
17 attempt zero division
18 In evaluating the constant part of an expression, the
19 right operand of a division or remainder operator was
20 found to be the constant zero.
21 invalid & prefix
22 The "&" operator has been used in an invalid context,
23 such as "&x = y".
24 warning: found ++r-value
25 warning: found --r-value
26 You get this if you say something like "++x++".
27 invalid \$ escape sequence
28 An escape sequence beginning with '\$' is not known to
29 the compiler.
30 invalid unary operator
31 The compiler discovered you trying to use a binary
32 operator in a unary manner.
33 invalid =
34 invalid *=
35 invalid >>=
36 The expression on the left hand side of an assignment
37 operator does not have an lvalue.
38 invalid ++
39 invalid --
40 The expression operated upon by the '++' or '--'
41 operator does not have an lvalue.
42 invalid label
43 A name used as a label has previously been declared
44 as EXTRN or AUTO in the current function.
45 invalid break
46 The compiler found a BREAK statement which was not
47 inside a FOR, WHILE, DO-WHILE, REPEAT or SWITCH
48 statement.
49 invalid next
50 The compiler found a NEXT statement which was not in-
51 side a FOR, WHILE, DO-WHILE or REPEAT statement.
52 invalid constant expression
53 Will happen if you try to use a string constant in a
54 constant expression.

invalid operator

This is one of those "cannot happen" messages. If it does happen, please submit an error report.

auto array too large

You attempted to declare an auto vector with a dimension greater than 1000 words. It is better to use an external vector or else GETVEC the space, since AUTO variables are allocated on the stack and stack space is limited.

extrn array too large

This will happen if you declare an external vector like "x[3.0];".

invalid case

A CASE label is not inside a SWITCH statement.

invalid default

A DEFAULT label is not inside a SWITCH statement.

default already supplied

More than one DEFAULT label in a SWITCH statement.

invalid case operator

The only bound operators permitted in a CASE are <, >, >=, and <=.

%filename ignored- too many open files

This usually happens when you include a file which includes itself.

bad input character: <ddd> (octal)

A character encountered in the input stream outside of a string or character constant has no meaning for the compiler. This might be a backspace or some control character typed in by mistake. Since it may be non-printing, the value of the offending character is displayed in octal.

rewrite this expression

A subscripting expression is too involved for the code generator to handle. Try breaking up the expression into more than one statement.

manifest nesting too deep

This will occur when you have manifest constants whose evaluation involves other manifest constants. This will occur if you have a series of manifest definitions, each of which is defined in terms of the previous manifest. This is ok in GMAP, but not in B.

warning: program size > 32k

One of the object decks generated will require more than 32K words to load. You may get this warning if you declare several very large external vectors. However, it might also mean the loader will be aborted by TSS due to "not enough core to run job".

expression too complex

no tree space

no stack space

An expression is too complex for the compiler to evaluate. Try simplifying it by breaking it up into

two or more expressions.

The constant <ddd> occurs in two case labels

The same constant appears in more than one CASE label in a SWITCH statement. The value of the offending constant is printed in decimal.

the upper range <ddd> overlaps the lower range <ddd>

The compiler has detected overlapping bounds inside a SWITCH statement. The values of the bounds are displayed in decimal.

The constant <ddd> is in the range <ddd>::<ddd>

The compiler has detected a CASE constant which is in the range of a range case or relational case, in a SWITCH statement. The numbers are given in decimal. If something conflicts with a relational case, then the bounds generated for the relation are shown. For example, the bounds for "case > 0:" would be "1::34359738367".

Initializers nested too deeply

An external declaration has initializers in braces nested to a depth greater than seven.

external redefined

auto variable redefined

label redefined

auto array name redefined

The compiler has detected an attempt to redefine a symbol which has already been defined to the current function body.

no space for symdef

There are too many external definitions; try dividing them into two groups by either compiling them separately or placing a function in between. This error is almost never encountered.

no space for symref

There are too many external references in a function definition; try simplification. This error is almost never encountered.

warning: #<text> ignored

A line beginning with a '#', which is taken to be a compiler directive, does not contain a recognizable directive. The line is ignored.

TSS loader warning messages:

<w> name undefined

This is a loader message, which indicates that an external variable referenced by one of your functions, or a library function, remains undefined after all libraries have been searched. If your program references the named external it will abort with a MME fault in TSS, or with a USER'S L1 MME GEBORT in batch.

<w> name loaded previously

The loader has discovered a function or external with

1
2
3 the same name as one already loaded. The most
4 probable reason is that you have two or more dif-
5 ferent names which, when truncated to six characters
6 end up being the same. The loader ignores all but the
7 first. Make sure all your externals and function
8 names are unique in their first six characters.
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57

APPENDIX D

Index of B library routines

7	.ABBRV	check for valid abbreviations
8	.BOFF	define a debugger breakpoint
9	.BSET	parse string into arguments
10	.FEDIT	edit a filename for printing
11	.GROW	grow a file
12	.GSNMB	generate a snumb
13	.IDENT	write an ident image for a backdoor file
14	.IOSTA	find IO status for a unit
15	.LEAVE	leave file accessed on close
16	.PROFILE	generate profile of a B program
17	.READ	reference or change the current read unit
18	.REMOV	force deaccess of file on close
19	.RESTR	Load an element from an hstar or system
20	.TABS	set tabs for the current output unit
21	.TRACE	set trace output unit
22	.WRAP	user specified termination actions
23	.WRITE	reference or change the current write unit
24	ABORT	abort job, producing dump and returning status
25	ABS	absolute value of an integer
26	ACC.FILE	access file
27	ACCLIB	access the system libraries
28	ADDCHAR	adds a character to the end of a B string
29	ADDVEC	replace a vector by a new, larger one
30	AFT.NAME	return aftname/filecode
31	ALLOCATE	A simple garbage collecting storage allocator
32	ANY	check if a character appears in a string
33	APPLY	call arbitrary function with arbitrary arguments
34	ARS	arithmetic right shift
35	ASCBCD	convert an ASCII string to a BCD vector
36	ATTACH	attach file to task file list
37	BACKSPACE	back up output unit by one character
38	BACK.DOOR	submit a file to the sysout backdoor queue
39	BCDASC	convert characters from BCD to ASCII
40	BCDADD	add two bcd numbers
41	BCDSUB	subtract two BCD numbers
42	BINBCD	convert a binary number to bcd
43	C.READ	force unit to be read unit
44	C.WRITE	force unit to be write unit
45	CALLF	call FORTRAN program from B routine
46	CALLFF	call FORTRAN function returning floating-point value
47	CATSCAF	turn pathname into Filsys format
48	CHAR	extract ASCII character from string
49	CHARB	extract BCD character from string
50	CHARP	set up a character pointer into a string
51	CHCKSM	compute a checksum
52	CLOSE	close currently open unit
53	CMPC	EIS string compare
54	CMPLOG	compare two values logically

1
 2
 3 CMPVEC compare one B vector to another
 4 COLUMN find current output column
 5 COMPARE compare two B strings
 6 CONCAT concatenate a series of strings
 7 COPY copy contents of one vector into another
 8 COPYCH character copy and substring function
 9 DATE return current date in ASCII
 10 DATEJUL conversion from normal date form to Julian form
 11 DATESI convert date in ASCII string to a standard form
 12 DATEV return current date & time as vector of integers
 13 DAYMON convert date to dd/mmm/yy format
 14 DIV integer divide with uniform direction of truncation
 15 DRL.DRL execute a given TSS Derail (system call)
 16 DRLJSTS get status of batch job
 17 DTOA decimal (integers) to ASCII
 18 DUMP dump vector
 19 DUMPA dump an array
 20 EBCASC convert string from EBCDIC to ASCII
 21 EOF test input end of file, or write output end of file
 22 EQUAL compare two strings for equality
 23 ERROR type an error message, then exit
 24 EXIT end job and return status
 25 EXTERNALS useful externals in the B library
 26 FILDES get file descriptor word
 27 FLUSH force i/o to end of line
 28 FPINPUT ASCII to floating point binary
 29 FPO floating point output
 30 FSFILE space input file forward one file
 31 GETARG extract (command) arguments from a string
 32 GETBIN read vector of binary data from sequential file
 33 GETCHAR/GETC read a character
 34 GETDATE turn ASCII date into the form mm/dd/yy
 35 GETLINE read a line from an input unit
 36 GETMATRIX dynamically allocate a matrix
 37 GETMEDIA find media code of file
 38 GETNUMB read a number from the current input unit
 39 GETREC get next logical record, with rcw, from a file
 40 GETRCP Get record pointer
 41 GETSTR read a string from an input unit
 42 GETTAPE ask for a tape from GCOS
 43 GETUMC get userid of current user
 44 GETVEC dynamically allocate a vector
 45 GNUMBER extract number from string
 46 GOTOSS execute a TSS command, never to return
 47 GTB convert gray code to binary (execute a gtb instruction)
 48 HIST The B histogram package
 49 HISTDESTROY Free space used by a histogram
 50 HISTINIT Allocate and initialize a histogram
 51 HISTOGRAM Add a point to a histogram
 52 HISTPRINT Print accumulated histogram
 53 ID.DRUN find ID of executing DRUN
 54 INCRUN determine if a user is in a crun

1 INTREQUEST handle interrupt (break) requests
 2 INTSS tell whether running batch or TSS
 3 IOERRORS change default handling of file access errors
 4 JULDATE conversion from Julian to normal date form
 5 LCHAR replace ASCII character in string
 6 LCHARB replace BCD character in string
 7 LENGTH return the length of a string
 8 LINUMB return the current line number of a file
 9 LOWERCASE turn alphabets in a string to lower case
 10 LSTAR stuff the byte addressed by a charp pointer
 11 MAIN entry to your program from b setup
 12 MAX maximum value of list of integers
 13 MIN minimum of a list of integers
 14 MME.MME execute a given batch MME (system call)
 15 MOVELR move characters left to right using an EIS MLR
 16 MOVERL move characters right to left using an EIS MRL
 17 NARGS return number of arguments to a function
 18 NOBRKS count number of times break key was hit
 19 NULLSTRING check for null string
 20 NUMARG extract numeric argument from character string
 21 OPEN open a file or string for i/o
 22 OVERFLOW test and reset overflow indicator
 23 PASUST execute a TSS DRL PASUST
 24 PEEK copy memory into your address space
 25 PNMATCH perform simple pattern match
 26 PRINT do a PRINTF into a string
 27 PRINTF formatted print
 28 PROMPT prompt for input at terminal
 29 PUTASC dump bytes in ASCII
 30 PUTBCD output the contents of a BCD string
 31 PUTBIN write vector of binary data to sequential file
 32 PUTCHAR/PUTC output a character
 33 PUTNUMB output a decimal number
 34 PUTOCT output octal numbers
 35 PUTREC output unprocessed record to a sequential file
 36 PUTSTR write a string to current output unit
 37 QSORT quick sort
 38 RAN.RD do disc i/o
 39 RAND generate pseudo-random numbers
 40 RD.LNK/WR.LNK Low level sequential disk I/O
 41 RD.RAN/WR.RAN Low level random disk I/O
 42 READ/WRITE unit oriented binary i/o
 43 READF formatted character stream input
 44 RELMEM release unused memory from the core hole
 45 REM please see explain
 46 REMOV remove file from AFT given pathname
 47 REREAD back up input to beginning of line
 48 RESET non-local goto
 49 RETFIL remove a file from the aft
 REWIND rewind an open file
 RLSEVEC please see expl
 ROTATE rotate word N bits to left or right

RSCR read system controller clock
RSTPSW turn off switch word bits
SBAR find size of allocated memory
SCAF convert ASCII pathname to BCD catfile stack
SCAN extract delimited substring of a string
SCM scan characters under mask
SETMEDIA change media code of output file
SETPSW set switch word by bitwise OR
SHELLSORT a Shell sort
SIDATE convert standard date to ASCII string
SLEEP wait for specified interval
SMC.HASH compute smc hash bucket for a given userid
STAR get the byte addressed by a charp pointer
STRING string i/o support
STRIP cause line numbers to be stripped on input
SWAPDESCR change program descriptors
SYSTEM execute a TSS command
T2741 check if terminal is a 2741
TABSET specify tab stops for the current output unit
TALLY create tally to BCD string
TALLYB create tally to ASCII string
TAPE tape I/O in B
TASK submit a task job via Derail TASK
TICK return cpu time for current user
TIME get time in pulses, or convert it to a string
TR9T09 translate to any 9 bit character code from any other
TRACE B function call-return tracer
TRIM trim trailing blanks off a string
TRTEST translate and test characters (fast character scan)
TTYN determine if I/O is to terminal
UNGETC place character in input stream
UPPERCASE turn alphabets in a string to upper case
VECTOR create a B vector and initialize contents
WDLENG return word length in bits
XLATE translate one character code to another
ZERO initialise a B vector to some value