






*C Language
Reference Manual*



SiliconGraphics
Computer Systems



C Language Reference Manual

Version 1.0

Document Number 007-0701-010

Technical Publications:

Robert Reimann

Engineering:

Greg Boyd

Jim Terhorst

© Copyright 1987, Silicon Graphics, Inc.

All rights reserved.

This document contains proprietary information of Silicon Graphics, Inc., and is protected by Federal copyright law. The information may not be disclosed to third parties or copied or duplicated in any form, in whole or in part, without prior written consent of Silicon Graphics, Inc.

The information in this document is subject to change without notice.

C Language Reference Manual

Version 1.0

Document Number 007-0701-010

Silicon Graphics, Inc.

Mountain View, California

Contents

1. Introduction	1-1
2. Lexical Conventions	2-1
2.1 Comments	2-1
2.2 Identifiers (Names)	2-1
2.3 Keywords	2-2
2.4 Constants	2-2
2.4.1 Integer Constants	2-2
2.4.2 Explicit Long Constants	2-3
2.4.3 Character Constants	2-3
2.4.4 Floating Constants	2-4
2.4.5 Enumeration Constants	2-4
2.5 String Literals	2-4
2.6 Syntax Notation	2-4
3. Storage Class and Type	3-1
3.1 Storage Class	3-1
3.2 Type	3-2
3.3 Objects and lvalues	3-3
4. Operator Conversions	4-1
4.1 Characters and Integers	4-1
4.2 Float and Double	4-2
4.3 Floating and Integral	4-2
4.4 Pointers and Integers	4-2
4.5 Unsigned	4-3
4.6 Arithmetic Conversions	4-3
4.7 Void	4-4
5. Expressions and Operators	5-1
5.1 Primary Expressions	5-2
5.2 Unary Operators	5-4
5.3 Multiplicative Operators	5-6
5.4 Additive Operators	5-7
5.5 Shift Operators	5-8

5.6	Relational Operators	5-8
5.7	Equality Operators	5-9
5.8	Bitwise AND Operator	5-9
5.9	Bitwise Exclusive OR Operator	5-9
5.10	Bitwise Inclusive OR Operator	5-10
5.11	Logical AND Operator	5-10
5.12	Logical OR Operator	5-10
5.13	Conditional Operator	5-11
5.14	Assignment Operators	5-11
5.15	Comma Operator	5-12
6.	Declarations	6-1
6.1	Storage Class Specifiers	6-1
6.2	Type Specifiers	6-2
6.3	Declarators	6-3
6.4	Meaning of Declarators	6-4
6.5	Structure and Union Declarations	6-7
6.6	Enumeration Declarations	6-9
6.7	Initialization	6-11
6.8	Type Names	6-13
6.9	Implicit Declarations	6-14
6.10	<i>typedef</i>	6-14
7.	Statements	7-1
7.1	Expression Statement	7-1
7.2	Compound Statement or Block	7-1
7.3	Conditional Statement	7-2
7.4	while Statement	7-2
7.5	do Statement	7-3
7.6	for Statement	7-3
7.7	switch Statement	7-4
7.8	break Statement	7-5
7.9	continue Statement	7-5
7.10	return Statement	7-6
7.11	goto Statement	7-6
7.12	Labeled Statement	7-6
7.13	Null Statement	7-7
8.	External Definitions	8-1

8.1	External Function Definitions	8-1
8.2	External Data Definitions	8-2
9.	Scope Rules	9-1
9.1	Lexical Scope	9-1
9.2	Scope of Externals	9-2
10.	Compiler Control Lines	10-1
10.1	Token Replacement	10-1
10.2	File Inclusion	10-2
10.3	Conditional Compilation	10-3
10.4	Line Control	10-4
11.	Types Revisited	11-1
11.1	Structures and Unions	11-1
11.2	Functions	11-2
11.3	Arrays, Pointers, and Subscripting	11-3
11.4	Explicit Pointer Conversions	11-4
12.	Constant Expressions	12-1
13.	Portability Considerations	13-1
14.	Syntax Summary	14-1
14.1	Expressions	14-1
14.2	Declarations	14-3
14.3	Statements	14-6
14.4	External Definitions	14-7
14.5	Preprocessor	14-8
A:	C on the IRIS-4D	A-2
A.1	<i>vararg.h</i> Macros	A-1
A.2	Deviations	A-2
A.3	Extensions	A-2
A.4	Translation Limits	A-3
A.5	Storage Mapping	A-4
	A.5.1 Alignment, Size, and Value Ranges	A-4
	A.5.2 Arrays, Structures, and Unions	A-5
	A.5.3 Storage Classes	A-9
A.6	Compiler Options	A-10



1. Introduction

This document contains a summary of the grammar and syntax rules of the C Programming Language as implemented on the IRIS-4D Series workstations. Appendix A discusses details of MIPS C as implemented for the IRIS-4D Series.



2. Lexical Conventions

There are six classes of tokens: identifiers, keywords, constants, string literals, operators, and other separators. Blanks, tabs, new-lines, and comments (collectively, “white space”) as described below are ignored except as they serve to separate tokens. Some white space is required to separate otherwise adjacent identifiers, keywords, and constants.

If the input stream has been parsed into tokens up to a given character, the next token is taken to include the longest string of characters that could possibly constitute a token.

2.1 Comments

The characters `/*` introduce a comment that terminates with the characters `*/`. Comments do not nest.

2.2 Identifiers (Names)

An identifier is a sequence of letters and digits. The first character must be a letter. The underscore (`_`) counts as a letter. Uppercase and lowercase letters are different. There is no limit on the length of a name. Other implementations may collapse case distinctions for external names, and may reduce the number of significant characters for both external and non-external names.

2.3 Keywords

The following identifiers are reserved for use as keywords and may not be used otherwise:

asm	const	external	long	static	void
auto	default	float	register	struct	volatile
break	do	for	return	switch	while
case	double	goto	short	typedef	
char	else	if	signed	union	
continue	enum	int	sizeof	unsigned	

Some implementations also reserve the word **fortran**.

2.4 Constants

There are several kinds of constants. Each has a type; an introduction to types is given in Chapter 3, Storage Class and Type.

2.4.1 Integer Constants

An integer constant consisting of a sequence of digits is taken to be octal if it begins with **0** (digit zero). An octal constant consists of the digits **0** through **7** only. A sequence of digits preceded by **0x** or **0X** (digit zero) is taken to be a hexadecimal integer. The hexadecimal digits include **a** or **A** through **f** or **F** with values 10 through 15. Otherwise, the integer constant is taken to be decimal. A decimal constant whose value exceeds the largest signed machine integer is taken to be **long**; an octal or hex constant that exceeds the largest unsigned machine integer is likewise taken to be **long**. Otherwise, integer constants are **int**.

2.4.2 Explicit Long Constants

A decimal, octal, or hexadecimal integer constant immediately followed by `l` (letter ell) or `L` is a long constant. As discussed below, integer and long values may be considered identical.

2.4.3 Character Constants

A character constant is a character enclosed in single quotes, as in `'x'`. The value of a character constant is the numerical value of the character in the machine's character set. Certain nongraphic characters, the single quote (`'`) and the backslash (`\`), may be represented according to the escape sequences shown in Table 2-1.

Character Name	Symbol	Escape Sequence
new-line	NL (LF)	<code>\n</code>
horizontal tab	HT	<code>\t</code>
vertical tab	VT	<code>\v</code>
backspace	BS	<code>\b</code>
carriage return	CR	<code>\r</code>
form feed	FF	<code>\f</code>
backslash	<code>\</code>	<code>\\</code>
single quote	<code>'</code>	<code>\'</code>
corresponding character	<i>ddd</i>	<code>\ddd</code>

Table 2-1. Escape Sequences for Nongraphic Characters

The escape `\ddd` consists of the backslash followed by 1, 2, or 3 octal digits that are taken to specify the value of the desired character. A special case of this construction is `\0` (not followed by a digit), which indicates the ASCII character NUL. If the character following a backslash is not one of those specified, the behavior is undefined. An explicit new-line character is illegal in a character constant. The type of a character constant is `int`.

2.4.4 Floating Constants

A floating constant consists of an integer part, a decimal point, a fraction part, an *e* or *E*, and an optionally signed integer exponent. The integer and fraction parts both consist of a sequence of digits. Either the integer part or the fraction part (not both) may be missing. Either the decimal point or the *e* and the exponent (not both) may be missing. Every floating constant has type **double**.

2.4.5 Enumeration Constants

Names declared as enumerators (see Section 6.5, Structure and Union Declarations and Section 6.6, Enumeration Declarations) have type **int**.

2.5 String Literals

A string literal is a sequence of characters surrounded by double quotes, as in "...". A string literal has type "array of **char**" and storage class **static** (see Chapter 3) and is initialized with the given characters. The compiler places a null byte (**\0**) at the end of each string literal so that programs that scan the string literal can find its end. In a string literal, the double quote character (") must be preceded by a ****; in addition, the same escapes as described for character constants may be used.

A **** and the immediately following new-line are ignored. All string literals, even when written identically, are distinct.

2.6 Syntax Notation

Syntactic categories are indicated by *italic* type and literal words and characters by **bold** type. Alternative categories are listed on separate lines. An optional entry is indicated by the subscript "opt", so that

$$\{ \textit{expression}_{opt} \}$$

indicates an optional expression enclosed in braces.

3. Storage Class and Type

The C language bases the interpretation of an identifier upon two attributes of the identifier: its storage class and its type. The storage class determines the location and lifetime of the storage associated with an identifier; the type determines the meaning of the values found in the identifier's storage.

3.1 Storage Class

There are four declarable storage classes:

- automatic
- static
- external
- register

Automatic variables are local to each invocation of a block (see Section 7.2, Compound Statement or Block) and are discarded upon exit from the block. Static variables are local to a block but retain their values upon reentry to a block even after control has left the block. External variables exist and retain their values throughout the execution of the entire program and may be used for communication between functions, even separately compiled functions. Register variables are (if possible) stored in the fast registers of the machine; like automatic variables, they are local to each block and disappear on exit from the block.

3.2 Type

The C language supports several fundamental types of objects. Objects declared as characters (**char**) are large enough to store any member of the implementation's character set. If a genuine character from that character set is stored in a **char** variable, its value is equivalent to the integer code for that character. Other quantities may be stored into character variables, but the implementation is machine dependent. In particular, **char** may be signed or unsigned by default. In this implementation the default is unsigned.

Up to three sizes of integer, declared **short int**, **int**, and **long int**, are available. Longer integers provide no less storage than shorter ones, but the implementation may make either short integers or long integers, or both, equivalent to plain integers. Plain integers have the natural size suggested by the host machine architecture. The other sizes are provided to meet special needs. The sizes are shown in Table 3-1.

Type	Size
char	8 bits
int	32
short	16
long	32
float	32
double	64
float range	$\pm 10^{\pm 38}$
double range	$\pm 10^{\pm 308}$

Table 3-1. Hardware Characteristics

The properties of **enum** types (see Chapter 6, Declarations) are identical to those of some integer types. The implementation may use the range of values to determine how to allot storage.

Unsigned integers, declared **unsigned**, obey the laws of arithmetic modulo 2^n where n is the number of bits in the representation.

Single-precision floating point (**float**) and double precision floating point (**double**) may be synonymous in some implementations. This is not the case in this implementation.

Because objects of the foregoing types can usefully be interpreted as numbers, they will be referred to as arithmetic types. **Char**, **int** of all sizes whether **unsigned** or not, and **enum** will collectively be called integral types. The **float** and **double** types will collectively be called floating types. Arithmetic types and pointers will be collectively referred to as **scalar** types.

The **void** type specifies an empty set of values. It is used as the type returned by functions that generate no value.

Besides the fundamental arithmetic types, there is a conceptually infinite class of derived types constructed from the fundamental types in the following ways:

- arrays of objects of most types
- functions that return objects of a given type
- pointers to objects of a given type
- structures containing a sequence of objects of various types
- unions capable of containing any one of several objects of various types

In general these methods of constructing objects can be applied recursively.

3.3 Objects and lvalues

An object is a manipulatable region of storage. An lvalue is an expression referring to an object. An obvious example of an lvalue expression is an identifier. There are operators that yield lvalues: for example, if **E** is an expression of pointer type, then ***E** is an lvalue expression referring to the object to which **E** points. The name "lvalue" comes from the assignment expression **E1 = E2** in which the left operand **E1** must be an lvalue expression. The discussion of each operator below indicates whether it expects lvalue operands and whether it yields an lvalue.

(

(

(

4. Operator Conversions

A number of operators may, depending on their operands, cause conversion of the value of an operand from one type to another. This part explains the result to be expected from such conversions. The conversions demanded by most ordinary operators are summarized under Arithmetic Conversions in this chapter. The summary will be supplemented as required by the discussion of each operator.

4.1 Characters and Integers

A character or a short integer may be used wherever an integer may be used. In all cases the value is converted to an integer. Conversion of a shorter integer to a longer preserves sign. It is guaranteed that a member of the standard character set is non-negative.

On machines that treat characters as signed, the characters of the ASCII set are all non-negative. However, a character constant specified with an octal escape suffers sign extension and may appear negative; for example, `'\377'` has the value `-1`.

When a longer integer is converted to a shorter integer or to a `char`, it is truncated on the left. Excess bits are simply discarded.

4.2 Float and Double

All floating arithmetic in C is carried out in double precision. Whenever a **float** appears in an expression it is lengthened to **double** by zero padding its fraction. When a **double** must be converted to **float**, for example by an assignment, the **double** is rounded before truncation to **float** length. This result is undefined if it cannot be represented as a **float**. This implementation supports an extension (invoked by the `-float` switch) which will limit the precision in which floating-point expressions are computed to **float**, unless the expression contains doubles.

4.3 Floating and Integral

Conversions of floating values to integral type are rather machine dependent. In particular, the direction of truncation of negative numbers varies. The result is undefined if it will not fit in the space provided.

Conversions of integral values to floating type behave well. Some loss of accuracy occurs if the destination lacks sufficient bits.

4.4 Pointers and Integers

An expression of integral type may be added to or subtracted from a pointer; in such a case, the first is converted as specified in the discussion of the addition operator. Two pointers to objects of the same type may be subtracted; in this case, the result is converted to an integer as specified in the discussion of the subtraction operator.

4.5 Unsigned

Whenever an unsigned integer and a plain integer are combined, the plain integer is converted to unsigned and the result is unsigned. The value is the least unsigned integer congruent to the signed integer (modulo 2^{wordsize}). In a 2's complement representation, this conversion is conceptual; and there is no actual change in the bit pattern.

When an unsigned **short** integer is converted to **long**, the value of the result is the same numerically as that of the unsigned integer. Thus, the conversion amounts to padding with zeros on the left.

4.6 Arithmetic Conversions

A great many operators cause conversions and yield result types in a similar way. This pattern will be called the "usual arithmetic conversions."

1. First, any operands of type **char** or **short** are converted to **int**, and any operands of type **unsigned char** or **unsigned short** are converted to **unsigned int**.
2. Then, if either operand is **float** or **double**, both are converted to **double** and that is the type of the result. (Floats are not extended to doubles in some instances if the `-float` switch is used. See Section A.6, Compiler Options.)
3. Otherwise, if either operand is **unsigned long**, the other is converted to **unsigned long** and that is the type of the result.
4. Otherwise, if either operand is **long**, the other is converted to **long** and that is the type of the result.
5. Otherwise, if one operand is **long**, and the other is **unsigned int**, they are both converted to **unsigned long** and that is the type of the result.
6. Otherwise, if either operand is **unsigned**, the other is converted to **unsigned** and that is the type of the result.
7. Otherwise, both operands must be **int**, and that is the type of the result.

4.7 Void

The (nonexistent) value of a **void** object may not be used in any way, and neither explicit nor implicit conversion may be applied. Because a void expression denotes a nonexistent value, such an expression may be used only as an expression statement (see Section 7.1, Expression Statement) or as the left operand of a comma expression (see Section 5.15, Comma Operator).

An expression may be converted to type **void** by use of a cast. For example, this makes explicit the discarding of the value of a function call used as an expression statement.

5. Expressions and Operators

The precedence of expression operators is the same as the order of the major subsections of this section, highest precedence first. Thus, for example, the expressions referred to as the operands of + (see Additive Operators in this chapter) are those expressions defined under Primary Expressions, Unary Operators, and Multiplicative Operators in this chapter. Within each subpart, the operators have the same precedence. Left- or right-associativity is specified in each subsection for the operators discussed therein. The precedence and associativity of all the expression operators are summarized in the grammar of Chapter 14, Syntax Summary.

Otherwise, the order of evaluation of expressions is undefined. In particular, the compiler considers itself free to compute subexpressions in the order it believes most efficient even if the subexpressions involve side effects. Expressions involving a commutative and associative operator (*, +, &, |, ^) may be rearranged arbitrarily even in the presence of parentheses; to force a particular order of evaluation, an explicit temporary must be used.

The handling of overflow and divide check in expression evaluation is undefined. Most existing implementations of C ignore integer overflows; treatment of division by 0 and all floating-point exceptions varies between machines and is usually adjustable by a library function.

5.1 Primary Expressions

Primary expressions involving `..`, `->`, subscripting, and function calls group left to right.

primary-expression:
identifier
constant
string literal
(expression)
primary-expression [expression]
primary-expression (expression-list^{opt})
primary-expression . identifier
primary-expression -> identifier

expression-list:
expression
expression-list , expression

An identifier is a primary expression provided it has been suitably declared as discussed below. Its type is specified by its declaration. If the type of the identifier is "array of ...", then the value of the identifier expression is a pointer to the first object in the array; and the type of the expression is "pointer to ...". Moreover, an array identifier is not an lvalue expression. Likewise, an identifier that is declared "function returning ...", when used except in the function-name position of a call, is converted to "pointer to function returning ...".

A constant is a primary expression. Its type may be **int**, **long**, or **double** depending on its form. Character constants have type **int** and floating constants have type **double**.

A string literal is a primary expression. Its type is originally "array of **char**", but following the same rule given above for identifiers, this is modified to "pointer to **char**" and the result is a pointer to the first character in the string literal. (There is an exception in certain initializers; see Section 6.7, Initialization.)

A parenthesized expression is a primary expression whose type and value are identical to those of the unadorned expression. The presence of parentheses does not affect whether the expression is an lvalue.

A primary expression followed by an expression in square brackets is a primary expression. The intuitive meaning is that of a subscript. Usually,

the primary expression has type "pointer to ...", the subscript expression is `int`, and the type of the result is "...". The expression `E1[E2]` is identical (by definition) to `*((E1)+(E2))`. All the clues needed to understand this notation are contained in this subpart together with the discussions on identifiers in "Unary Operators" and "Additive Operators" in this chapter, `*` and `+`, respectively. The implications are summarized in Section 11.3, Arrays, Pointers, and Subscripting.

A function call is a primary expression followed by parentheses containing a possibly empty, comma-separated list of expressions that constitute the actual arguments to the function. The primary expression must be of type "function returning ...", and the result of the function call is of type "...". As indicated below, a hitherto unseen identifier followed immediately by a left parenthesis is contextually declared to represent a function returning an integer.

If a corresponding function prototype is in force which specifies a type for the argument being evaluated, it is converted to that type. Otherwise, the argument is converted according to the following **default argument promotions**:

- type `float` is converted to `double`
- types `unsigned short` and `unsigned char` are converted to `unsigned int`
- types `signed short` and `signed char` are converted to `signed int`
- array and function names are converted to `nd` functions

In preparing for the call to a function, a copy is made of each actual parameter. Thus, all argument passing in C is strictly by value. A function may change the values of its formal parameters, but these changes cannot affect the values of the actual parameters. It is possible to pass a pointer on the understanding that the function may change the value of the object to which the pointer points. An array name is a pointer expression. As the order of evaluation of arguments is undefined by the language, this implementation reserves the right to evaluate arguments in whatever order is considered optimal. In particular, side effects of the argument evaluation (e.g. postincrement) may be delayed until after all arguments are evaluated. Recursive calls to any function are permitted.

A primary expression followed by a dot followed by an identifier is an expression. The first expression must be a structure or a union, and the

identifier must name a member of the structure or union. The value is the named member of the structure or union, and it is an lvalue if the first expression is an lvalue.

A primary expression followed by an arrow (built from `-` and `>`) followed by an identifier is an expression. The first expression must be a pointer to a structure or a union and the identifier must name a member of that structure or union. The result is an lvalue referring to the named member of the structure or union to which the pointer expression points. Thus the expression `E1->MOS` is the same as `(*E1).MOS`. Structures and unions are discussed in Chapter 6.

5.2 Unary Operators

Expressions with unary operators group right to left.

unary-expression:
* *expression*
& *lvalue*
- *expression*
! *expression*
~ *expression*
++ *lvalue*
-- *lvalue*
lvalue ++
lvalue --
(*type-name*) *expression*
sizeof *expression*
sizeof (*type-name*)

The unary `*` operator means "indirection"; the expression must be a pointer, and the result is an lvalue referring to the object to which the expression points. If the type of the expression is "pointer to ...", the type of the result is "...".

The result of the unary `&` operator is a pointer to the object referred to by the lvalue. If the type of the lvalue is "...", the type of the result is "pointer to ...".

The result of the unary `-` operator is the negative of its operand. The usual arithmetic conversions are performed. The negative of an unsigned quantity is computed by subtracting its value from 2^n where n is the number of bits in the corresponding signed type.

There is no unary `+` operator.

The result of the logical negation operator `!` is one if the value of its operand is zero, zero if the value of its operand is nonzero. The type of the result is `int`. It is applicable to any arithmetic type or to pointers.

The `~` operator yields the 1's complement of its operand. The usual arithmetic conversions are performed. The type of the operand must be integral.

The object referred to by the lvalue operand of prefix `++` is incremented. The value is the new value of the operand but is not an lvalue. The expression `++x` is equivalent to `x += 1`. See the discussions "Additive Operators" and "Assignment Operators" for information on conversions.

The lvalue operand of prefix `--` is decremented analogously to the prefix `++` operator.

When postfix `++` is applied to an lvalue, the result is the value of the object referred to by the lvalue. After the result is noted, the object is incremented in the same manner as for the prefix `++` operator. The type of the result is the same as the type of the lvalue expression.

When postfix `--` is applied to an lvalue, the result is the value of the object referred to by the lvalue. After the result is noted, the object is decremented in the manner as for the prefix `--` operator. The type of the result is the same as the type of the lvalue expression.

An expression preceded by the parenthesized name of a data type causes conversion of the value of the expression to the named type. This construction is called a cast. Type names are described in Section 6.8, Type Names.

The `sizeof` operator yields the size in bytes of its operand. (A byte is undefined by the language except in terms of the value of `sizeof`. However, in all existing implementations, a byte is the space required to hold a `char`.) When applied to an array, the result is the total number of bytes in the array. The size is determined from the declarations of the objects in the expression. This expression is semantically an **unsigned** constant and may be used anywhere a constant is required. Its major use is in communication with routines like storage allocators and I/O systems.

The `sizeof` operator may also be applied to a parenthesized type name. In that case it yields the size in bytes of an object of the indicated type.

The construction `sizeof(type)` is taken to be a unit, so the expression `sizeof(type)-2` is the same as `(sizeof(type))-2`.

5.3 Multiplicative Operators

The multiplicative operators `*`, `/`, and `%` group left to right. The usual arithmetic conversions are performed.

multiplicative expression:
*expression * expression*
expression / expression
expression % expression

The binary `*` operator indicates multiplication. The `*` operator is associative, and expressions with several multiplications at the same level may be rearranged by the compiler. The binary `/` operator indicates division. The operands must be arithmetic.

The binary `%` operator yields the remainder from the division of the first expression by the second. The operands must be integral.

When positive integers are divided, truncation is toward 0; but the form of truncation is machine-dependent if either operand is negative. On all machines covered by this manual, the remainder has the same sign as the dividend. It is always true that $(a/b)*b + a \% b$ is equal to a (if b is not 0).

5.4 Additive Operators

The additive operators + and – group left to right. The usual arithmetic conversions are performed. In general, the operands must be of arithmetic type; however there are some additional type possibilities for each operator.

additive-expression:

expression + expression

expression – expression

The result of the + operator is the sum of the operands. A pointer to an object in an array and a value of any integral type may be added. The latter is in all cases converted to an address offset by multiplying it by the length of the object to which the pointer points. The result is a pointer of the same type as the original pointer that points to another object in the same array, appropriately offset from the original object. Thus if **P** is a pointer to an object in an array, the expression **P+1** is a pointer to the next object in the array. No further type combinations are allowed for pointers.

The + operator is associative, and expressions with several additions at the same level may be rearranged by the compiler.

The result of the – operator is the difference of the operands. The usual arithmetic conversions are performed. Additionally, a value of any integral type may be subtracted from a pointer, and then the same conversions for addition apply.

If two pointers to objects of the same type are subtracted, the result is converted (by division by the length of the object) to an **int** representing the number of objects separating the pointed-to objects. This conversion will in general give unexpected results unless the pointers point to objects in the same array, since pointers, even to objects of the same type, do not necessarily differ by a multiple of the object length.

5.5 Shift Operators

The shift operators `<<` and `>>` group left to right. Both perform the usual arithmetic conversions on their operands, each of which must be integral. Then the right operand is converted to `int`; the type of the result is that of the left operand. The result is undefined if the right operand is negative or greater than or equal to the length of the object in bits.

shift-expression:

expression << expression
expression >> expression

The value of `E1<<E2` is `E1` (interpreted as a bit pattern) left-shifted `E2` bits. Vacated bits are 0 filled. The value of `E1>>E2` is `E1` right-shifted `E2` bit positions. The right shift is guaranteed to be logical (0 fill) if `E1` is **unsigned**; otherwise, it may be arithmetic.

5.6 Relational Operators

The relational operators group left to right.

relational-expression:

expression < expression
expression > expression
expression <= expression
expression >= expression

The operators `<` (less than), `>` (greater than), `<=` (less than or equal to), and `>=` (greater than or equal to) all yield 0 if the specified relation is false and 1 if it is true. The operands must be arithmetic or compatible pointers. The type of the result is `int`. The usual arithmetic conversions are performed. Two pointers may be compared; the result depends on the relative locations in the address space of the pointed-to objects. Pointer comparison is portable only when the pointers point to objects in the same array.

5.7 Equality Operators

equality-expression:

expression == expression

expression != expression

The == (equal to) and the != (not equal to) operators are exactly analogous to the relational operators except for their lower precedence. (Thus $a < b == c < d$ is 1 whenever $a < b$ and $c < d$ have the same truth value.) The operands must be arithmetic or compatible pointers.

A pointer may be compared to an integer only if the integer is the constant 0. A pointer to which 0 has been assigned is guaranteed not to point to any object and will appear to be equal to 0. In conventional usage, such a pointer is considered to be null.

5.8 Bitwise AND Operator

and-expression:

expression & expression

The & operator is associative, and expressions involving & may be rearranged. The usual arithmetic conversions are performed. The result is the bitwise AND function of the operands. The operator applies only to integral operands.

5.9 Bitwise Exclusive OR Operator

exclusive-or-expression:

expression ^ expression

The ^ operator is associative, and expressions involving ^ may be rearranged. The usual arithmetic conversions are performed; the result is the bitwise exclusive OR function of the operands. The operator applies only to integral operands.

5.10 Bitwise Inclusive OR Operator

inclusive-or-expression:
expression | expression

The `|` operator is associative, and expressions involving `|` may be rearranged. The usual arithmetic conversions are performed; the result is the bitwise inclusive OR function of its operands. The operator applies only to integral operands.

5.11 Logical AND Operator

logical-and-expression:
expression && expression

The `&&` operator groups left to right. It returns 1 if both its operands evaluate to nonzero, 0 otherwise. Unlike `&`, `&&` guarantees left to right evaluation; moreover, the second operand is not evaluated if the first operand evaluates to 0.

The operands need not have the same type, but must be scalar. The result is always `int`.

5.12 Logical OR Operator

logical-or-expression:
expression || expression

The `||` operator groups left to right. It returns 1 if either of its operands evaluates to nonzero, 0 otherwise. Unlike `|`, `||` guarantees left to right evaluation; moreover, the second operand is not evaluated if the value of the first operand evaluates to nonzero.

The operands need not have the same type, but each must be scalar. The result is always `int`.

5.13 Conditional Operator

conditional-expression:

expression ? expression : expression

Conditional expressions group right to left. The first expression is evaluated; and if it is nonzero, the result is the value of the second expression, otherwise that of third expression. If possible, the usual arithmetic conversions are performed to bring the second and third expressions to a common type. If both are structures or unions of the same type, the result has the type of the structure or union. If both pointers are of the same type, the result has the common type. Otherwise, one must be a pointer and the other the constant 0, and the result has the type of the pointer. Only one of the second and third expressions is evaluated.

5.14 Assignment Operators

There are a number of assignment operators, all of which group right to left. All require an lvalue as their left operand, and the type of an assignment expression is that of its left operand. The value is the value stored in the left operand after the assignment has taken place.

assignment-expression:

lvalue = expression

lvalue += expression

lvalue -= expression

*lvalue *= expression*

lvalue /= expression

lvalue %= expression

lvalue >>= expression

lvalue <<= expression

lvalue &= expression

lvalue = expression

lvalue /= expression

In the simple assignment with =, the value of the expression replaces that of the object referred to by the lvalue. If both operands have arithmetic type, the right operand is converted to the type of the left preparatory to the assignment. Second, both operands may be structures or unions of the same

type. Finally, if the left operand is a pointer, the right operand must in general be a pointer of the same type. However, the constant 0 may be assigned to a pointer; it is guaranteed that this value will produce a null pointer distinguishable from a pointer to any object.

The behavior of an expression of the form $E1 \text{ op } = E2$ may be inferred by taking it as equivalent to $E1 = E1 \text{ op } (E2)$; however, $E1$ is evaluated only once. In $+=$ and $-=$, the left operand may be a pointer, in which case the (integral) right operand is converted as explained in "Additive Operators" in this chapter. All right operands and all nonpointer left operands must have arithmetic type.

5.15 Comma Operator

comma-expression:
expression , expression

A pair of expressions separated by a comma is evaluated left to right, and the value of the left expression is discarded. The type and value of the result are the type and value of the right operand. This operator groups left to right. In contexts where comma is given a special meaning, e.g., in lists of actual arguments to functions (see Section 5.1, Primary Expressions) and lists of initializers (see Section 6.7, Initialization), the comma operator as described in this section can only appear in parentheses. For example,

$f(a, (t=3, t+2), c)$

has three arguments, the second of which has the value 5.

6. Declarations

Declarations are used to specify the interpretation that C gives to each identifier; they do not necessarily reserve storage associated with the identifier. Declarations have the form

declaration:
*decl-specifiers declarator-list*_{opt} ;

The declarators in the declarator-list contain the identifiers being declared. The decl-specifiers consist of a sequence of type and storage class specifiers.

decl-specifiers:
type-specifier decl-specifiers
*sc-specifier decl-specifiers*_{opt}

The list must be self-consistent in a way described below.

6.1 Storage Class Specifiers

The sc-specifiers are:

sc-specifier:
auto
static
extern
register
typedef

The **typedef** specifier does not reserve storage and is called a "storage class specifier" only for syntactic convenience. See Section 6.10, **typedef** for more information. The meanings of the various storage classes were discussed in Chapter 2, Lexical Conventions.

The **auto**, **static**, and **register** declarations also serve as definitions in that they cause an appropriate amount of storage to be reserved. In the **extern** case, there must be an external definition (see Chapter 8, External Definitions) for the given identifiers somewhere outside the function in which they are declared.

A **register** declaration is best thought of as an **auto** declaration, together with a hint to the compiler that the variables declared will be heavily used. Only the first few such declarations in each function are effective. Moreover, only variables of certain types will be stored in registers. One other restriction applies to variables declared using register storage class: the address-of operator, **&**, cannot be applied to them. Smaller, faster programs can be expected if register declarations are used appropriately.

At most, one **sc-specifier** may be given in a declaration. If the **sc-specifier** is missing from a declaration, it is taken to be **auto** inside a function, **extern** outside. Exception: functions are never automatic.

6.2 Type Specifiers

The type-specifiers are

type-specifier:
 struct-or-union-specifier
 typedef-name
 enum-specifier
basic-type-specifier:
 basic-type
 basic-type basic-type-specifiers
basic-type:
 char
 short
 int
 long
 signed
 unsigned
 float
 double
 void

At most one of the words **long** or **short** may be specified in conjunction with **int**; the meaning is the same as if **int** were not mentioned. The word **long** may be specified in conjunction with **float**; the meaning is the same as **double**. Either **signed** or **unsigned** may be specified alone, or in conjunction with **int** or any of its short or long varieties, or with **char**.

Otherwise, at most one type-specifier may be given in a declaration. In particular, adjectival use of **long**, **short**, or **unsigned** is not permitted with **typedef** names. If the type-specifier is missing from a declaration, it is taken to be **int**.

Specifiers for structures, unions, and enumerations are discussed in later in this chapter. Declarations with **typedef** names are discussed in "typedef" in this chapter.

6.3 Declarators

The declarator-list appearing in a declaration is a comma-separated sequence of declarators, each of which may have an initializer:

declarator-list:
init-declarator
init-declarator , declarator-list

init-declarator:
declarator initializer_{opt}

Initializers are discussed in Section 6.7, Initialization. The specifiers in the declaration indicate the type and storage class of the objects to which the declarators refer. Declarators have the syntax:

declarator:
identifier
(declarator)
** declarator*
declarator (parameter-type-list_{opt})
declarator [constant-expression_{opt}]

The grouping is the same as in expressions.

6.4 Meaning of Declarators

Each declarator is taken to be an assertion that when a construction of the same form as the declarator appears in an expression, it yields an object of the indicated type and storage class.

Each declarator contains exactly one identifier; it is this identifier that is declared. If an unadorned identifier appears as a declarator, then it has the type indicated by the specifier heading the declaration.

A declarator in parentheses is identical to the unadorned declarator, but the binding of complex declarators may be altered by parentheses. See the examples below.

Now imagine a declaration

T D1

where **T** is a type-specifier (like `int`, etc.) and **D1** is a declarator. Suppose this declaration makes the identifier have type "... **T**", where the "..." is empty if **D1** is just a plain identifier (so that the type of `x` in "`int x`" is just `int`). Then if **D1** has the form

***D**

the type of the contained identifier is "... pointer to **T**."

If **D1** has the form

D(*parameter-type-list*_{opt})

then the contained identifier has the type "... function returning **T**."

If **D1** has the form

D[*constant-expression*]

or

D[]

then the contained identifier has type "... array of **T**." In the first case, the constant expression is an expression whose value is determinable at compile time, whose type is `int`, and whose value is positive. (Constant expressions are defined precisely in Chapter 12.) When several "array of" specifications are adjacent, a multi-dimensional array is created; the constant expressions that specify the bounds of the arrays may be missing only for the first member of the sequence. This elision is useful when the array is external

and the actual definition, which allocates storage, is given elsewhere. The first constant expression may also be omitted when the declarator is followed by initialization. In this case the size is calculated from the number of initial elements supplied.

An array may be constructed from one of the basic types, from a pointer, from a structure or union, or from another array (to generate a multi-dimensional array).

A *parameter-type-list* declares the types of, and may declare identifiers for, the formal parameters of a function. When a function is invoked for which a function prototype is in scope, each actual parameter is converted to the type of the corresponding formal parameter specified in the function prototype. After this conversion, any parameter of size less than that of type **int** is widened according to the default argument promotions. In particular, **floats** specified in the type list are not converted to **double** before the call. If the list terminates with an ellipsis (...), only the parameters specified in the prototype have their types checked; additional parameters are converted according to the default argument promotions (see Section 5.1). Otherwise, the number of parameters appearing in the parameter list at the point of call must agree in number with those in the function prototype. Neither the absence of a *parameter-type-list* nor the definition of the prototype function indicate any information about the number and/or types of the function parameters which is used during argument conversion, i.e., only the information supplied by the prototype is used to supercede the default argument promotions.

The following is an example of function prototypes:

```
double foo(int *first, float second, ... );
int *fip(int a, long l, int (*ff)(float));
```

The first prototype declares a function **foo**, returning a **double**, which has (at least) two parameters: a pointer to an **int**, and a **float**. Further parameters may appear in an instance of the function, and are unspecified. The second prototype declares a function **fip**, which returns a pointer to an **int**. **fip** has three parameters: an **int**, a **long**, and a pointer to a function returning an **int** which has a single (**float**) argument.

Not all the possibilities allowed by the syntax above are actually permitted. The restrictions are as follows: functions may not return arrays or functions although they may return pointers; there are no arrays of functions although there may be arrays of pointers to functions. Likewise, a structure or union may not contain a function; but it may contain a pointer to a function.

As an example, the declaration

```
int i, *ip, f(), *fip(), (*pfi());
```

declares an integer **i**, a pointer **ip** to an integer, a function **f** returning an integer, a function **fip** returning a pointer to an integer, and a pointer **pfi** to a function, which returns an integer. It is especially useful to compare the last two. The binding of ***fip()** is ***(fip())**. The declaration suggests, and the same construction in an expression requires, the calling of a function **fip**, and then using indirection through the (pointer) result to yield an integer. In the declarator **(*pfi)()**, the extra parentheses are necessary, as they are also in an expression, to indicate that indirection through a pointer to a function yields a function, which is then called; it returns an integer.

As another example,

```
float fa[17], *afp[17];
```

declares an array of **float** numbers and an array of pointers to **float** numbers. Finally,

```
static int x3d[3][5][7];
```

declares a static 3-dimensional array of integers, with rank $3 \times 5 \times 7$. In complete detail, **x3d** is an array of three items; each item is an array of five arrays; each of the latter arrays is an array of seven integers. Any of the expressions **x3d**, **x3d[i]**, **x3d[i][j]**, **x3d[i][j][k]** may reasonably appear in an expression. The first three have type "array" and the last has type **int**.

6.5 Structure and Union Declarations

A structure is an object consisting of a sequence of named members. Each member may have any type. A union is an object that may, at a given time, contain any one of several members. Structure and union specifiers have the same form.

struct-or-union-specifier:
struct-or-union { struct-decl-list }
struct-or-union identifier { struct-decl-list }
struct-or-union identifier

struct-or-union:
struct
union

The struct-decl-list is a sequence of declarations for the members of the structure or union:

struct-decl-list:
struct-declaration
struct-declaration struct-decl-list

struct-declaration:
type-specifier struct-declarator-list ;

struct-declarator-list:
struct-declarator
struct-declarator , struct-declarator-list

In the usual case, a struct-declarator is just a declarator for a member of a structure or union. A structure member may also consist of a specified number of bits. Such a member is also called a field; its length, a non-negative constant expression, is separated from the field name by a colon.

struct-declarator:
declarator
declarator : constant-expression
: constant-expression

Within a structure, the objects declared have addresses that increase as the declarations are read left to right. Each non-field member of a structure begins on an addressing boundary appropriate to its type; therefore, there may be unnamed holes in a structure. Field members are packed into machine integers; they do not straddle words. A field that does not fit into

the space remaining in a word is put into the next word. No field may be wider than a word. (See Section A.5 for the sizes of basic types.)

A struct-declarator with no declarator, only a colon and a width, indicates an unnamed field useful for padding to conform to externally-imposed layouts. As a special case, a field with a width of 0 specifies alignment of the next field at an implementation dependent boundary.

The language does not restrict the types of things that are declared as fields. Moreover, even `int` fields may be considered to be unsigned. For these reasons, it is strongly recommended that fields be declared as **unsigned** where that is the intent. There are no arrays of fields, and the address-of operator, `&`, may not be applied to them, so that there are no pointers to fields.

A union may be thought of as a structure all of whose members begin at offset 0 and whose size is sufficient to contain any of its members. At most, one of the members can be stored in a union at any time.

A structure or union specifier of the second form, that is, one of

```
struct identifier { struct-decl-list }  
union identifier { struct-decl-list }
```

declares the identifier to be the *structure tag* (or union tag) of the structure specified by the list. A subsequent declaration may then use the third form of specifier, one of

```
struct identifier  
union identifier
```

Structure tags allow definition of self-referential structures. Structure tags also permit the long part of the declaration to be given once and used several times. It is illegal to declare a structure or union that contains an instance of itself, but a structure or union may contain a pointer to an instance of itself.

The third form of a structure or union specifier may be used prior to a declaration that gives the complete specification of the structure or union in situations in which the size of the structure or union is unnecessary. The size is unnecessary in two situations: when a pointer to a structure or union is being declared and when a **typedef** name is declared to be a synonym for a structure or union. This, for example, allows the declaration of a pair of structures that contain pointers to each other.

The names of members do not conflict with each other or with ordinary variables. A particular member name may not be used twice in the same structure, but it may be used in several different structures in the same scope. Names which are used for tags do not conflict with other names or with names used for tags in an enclosing scope.

A simple but important example of a structure declaration is the following binary tree structure:

```
struct tnode
{
    char tword[20];
    int count;
    struct tnode *left;
    struct tnode *right;
};
```

which contains an array of 20 characters, an integer, and two pointers to similar structures. Once this declaration has been given, the declaration

struct tnode s, *sp;

declares **s** to be a structure of the given sort and **sp** to be a pointer to a structure of the given sort. With these declarations, the expression

sp->count

refers to the **count** field of the structure to which **sp** points;

s.left

refers to the left subtree pointer of the structure **s**; and

s.right->tword[0]

refers to the first character of the **tword** member of the right subtree of **s**.

6.6 Enumeration Declarations

Enumeration variables and constants have integral type.

enum-specifier:
enum { *enum-list* }
enum *identifier* { *enum-list* }
enum *identifier*

enum-list:
enumerator
enum-list , *enumerator*

enumerator:
identifier
identifier = *constant-expression*

The identifiers in an enum-list are declared as constants and may appear wherever constants are required. If no enumerators with = appear, then the values of the corresponding constants begin at 0 and increase by 1 as the declaration is read from left to right. An enumerator with = gives the associated identifier the value indicated; subsequent identifiers continue the progression from the assigned value.

The names of enumerators in the same scope must all be distinct from each other and from those of ordinary variables.

The role of the identifier in the enum-specifier is entirely analogous to that of the structure tag in a struct-specifier; it names a particular enumeration. For example,

```
enum color { chartreuse, burgundy, claret=20, winedark };  
...  
enum color *cp, col;  
...  
col = claret;  
cp = &col;  
...  
if (*cp == burgundy) ...
```

makes **color** the enumeration-tag of a type describing various colors, and then declares **cp** as a pointer to an object of that type and **col** as an object of that type. The possible values are drawn from the set {0,1,20,21}.

6.7 Initialization

A declarator may specify an initial value for the identifier being declared. The initializer is preceded by = and consists of an expression or a list of values nested in braces.

```
initializer:  
= expression  
= { initializer-list }  
= { initializer-list , }
```

```
initializer-list:  
expression  
initializer-list , initializer-list  
{ initializer-list }  
{ initializer-list , }
```

All the expressions in an initializer for a static or external variable must be constant expressions, which are described in Chapter 12, Constant Expressions, or expressions that reduce to the address of a previously declared variable, possibly offset by a constant expression. Automatic or register variables may be initialized by arbitrary expressions involving constants and previously declared variables and functions.

Static and external variables that are not initialized are guaranteed to start off as zero. The value of automatic and register variables that are not initialized is undefined.

When an initializer applies to a scalar (a pointer or an object of arithmetic type), it consists of a single expression, perhaps in braces. The initial value of the object is taken from the expression; the same conversions as for assignment are performed.

When the declared variable is an aggregate (a structure or array), the initializer consists of a brace-enclosed, comma-separated list of initializers for the members of the aggregate written in increasing subscript or member order. If the aggregate contains subaggregates, this rule applies recursively to the members of the aggregate. If there are fewer initializers in the list than there are members of the aggregate, then the aggregate is padded with zeros. It is not permitted to initialize unions or automatic aggregates.

Braces may in some cases be omitted. If the initializer begins with a left brace, then the succeeding comma-separated list of initializers initializes the members of the aggregate; it is erroneous for there to be more initializers

than members. If, however, the initializer does not begin with a left brace, then only enough elements from the list are taken to account for the members of the aggregate; any remaining members are left to initialize the next member of the aggregate of which the current aggregate is a part.

A final abbreviation allows a **char** array to be initialized by a string literal. In this case successive characters of the string literal initialize the members of the array.

For example,

```
int x[] = { 1, 3, 5 };
```

declares and initializes **x** as a one-dimensional array that has three members, since no size was specified and there are three initializers.

```
float y[4][3] =
{
    { 1, 3, 5 },
    { 2, 4, 6 },
    { 3, 5, 7 },
};
```

is a completely-bracketed initialization: 1, 3, and 5 initialize the first row of the array **y**[0], namely **y**[0][0], **y**[0][1], and **y**[0][2]. Likewise, the next two lines initialize **y**[1] and **y**[2]. The initializer ends early and therefore **y**[3] is initialized with 0. Precisely, the same effect could have been achieved by

```
float y[4][3] =
{
    1, 3, 5, 2, 4, 6, 3, 5, 7
};
```

The initializer for **y** begins with a left brace but that for **y**[0] does not; therefore, three elements from the list are used. Likewise, the next three are taken successively for **y**[1] and **y**[2]. Also,

```
float y[4][3] =
{
    { 1 }, { 2 }, { 3 }, { 4 }
};
```

initializes the first column of **y** (regarded as a two-dimensional array) and leaves the rest 0.

Finally,

```
char msg[] = "Syntax error on line %s\n";
```

shows a character array whose members are initialized with a string literal. The length of the string (or size of the array) includes the terminating NUL character, `\0`.

6.8 Type Names

In three contexts (to specify type conversions explicitly by means of a cast, in a function prototype, and as an argument of `sizeof`), it is desired to supply the name of a data type. This is accomplished using a "type name", which in essence is a declaration for an object of that type that omits the name of the object.

type-name:

type-specifier abstract-declarator

abstract-declarator:

empty

(abstract-declarator)

** abstract-declarator*

abstract-declarator (parameter-type-list_{opt})

abstract-declarator [constant-expression_{opt}]

To avoid ambiguity, in the construction

(abstract-declarator)

the abstract-declarator is required to be nonempty. Under this restriction, it is possible to identify uniquely the location in the abstract-declarator where the identifier would appear if the construction were a declarator in a declaration. The named type is then the same as the type of the hypothetical identifier. For example,

```
int
int *
int *[3]
int (*)[3]
int *()
int *()
int (*[3])()
```

name respectively the types "integer", "pointer to integer", "array of three pointers to integers", "pointer to an array of three integers", "function returning pointer to integer", "pointer to function returning an integer", and "array of three pointers to functions returning an integer."

6.9 Implicit Declarations

It is not always necessary to specify both the storage class and the type of identifiers in a declaration. The storage class is supplied by the context in external definitions and in declarations of formal parameters and structure members. In a declaration inside a function, if a storage class but no type is given, the identifier is assumed to be **int**; if a type but no storage class is indicated, the identifier is assumed to be **auto**. An exception to the latter rule is made for functions because **auto** functions do not exist. If the type of an identifier is "function returning ...", it is implicitly declared to be **extern**.

In an expression, an identifier followed by (and not already declared is contextually declared to be "function returning int".

6.10 typedef

Declarations whose "storage class" is **typedef** do not define storage but instead define identifiers that can be used later as if they were type keywords naming fundamental or derived types.

typedef-name:
identifier

Within the scope of a declaration involving **typedef**, each identifier appearing as part of any declarator therein becomes syntactically equivalent to the type keyword naming the type associated with the identifier in the way described in Section 6.4, Meaning of Declarators.

For example, after

```
typedef int MILES, *KCLICKSP;  
typedef struct { double re, im; } complex;
```

the constructions

```
MILES distance;  
extern KCLICKSP metricp;  
complex z, *zp;
```

are all legal declarations; the type of **distance** is **int**, that of **metricp** is "pointer to **int**", and that of **z** is the specified structure. The **zp** is a pointer to such a structure.

The **typedef** does not introduce brand-new types, only synonyms for types that could be specified in another way. Thus in the example above **distance** is considered to have exactly the same type as any other **int** object.

C

C

C

7. Statements

Except as indicated, statements are executed in sequence.

7.1 Expression Statement

Most statements are expression statements, which have the form

expression ;

Usually expression statements are assignments or function calls.

7.2 Compound Statement or Block

So that several statements can be used where one is expected, the compound statement (also, and equivalently, called "block") is provided:

compound-statement:
{ declaration-list_{opt} statement-list_{opt} }

declaration-list:
declaration
declaration declaration-list

statement-list:
statement
statement statement-list

If any of the identifiers in the declaration-list were previously declared, the outer declaration is hidden for the duration of the block, after which it resumes its force.

Any initializations of **auto** or **register** variables are performed each time the block is entered at the top. It is currently possible (but a bad practice) to transfer into a block; in that case the initializations are not performed. Initializations of **static** variables are performed only once when the program begins execution. Inside a block, **extern** declarations do not reserve storage so their initialization is not permitted.

7.3 Conditional Statement

The two forms of the conditional statement are

```
if ( expression ) statement  
if ( expression ) statement else statement
```

In both cases, the expression is evaluated; if it is nonzero, the first substatement is executed. If it is zero, the second substatement is executed. An **else** clause which follows multiple sequential **else-less if** statements is associated with the most recent one in the same block (i.e., not in an enclosed block).

7.4 while Statement

The **while** statement has the form

```
while ( expression ) statement
```

The substatement is executed repeatedly so long as the value of the expression remains nonzero. The test takes place before each execution of the statement.

7.5 do Statement

The **do** statement has the form

```
do statement while ( expression );
```

The substatement is executed repeatedly until the value of the expression becomes 0. The test takes place after each execution of the statement.

7.6 for Statement

The **for** statement has the form:

```
for ( exp-1opt ; exp-2opt ; exp-3opt ) statement
```

Except for the behavior of **continue**, this statement is equivalent to

```
exp-1 ;  
while ( exp-2 )  
{  
    statement  
    exp-3 ;  
}
```

Thus the first expression specifies initialization for the loop; the second specifies a test, made before each iteration, such that the loop is exited when the expression becomes 0. The third expression often specifies an incrementing that is performed after each iteration.

Any or all of the expressions may be dropped. A missing *exp-2* makes the implied **while** clause equivalent to **while(1)**; other missing expressions are simply dropped from the expansion above.

7.7 switch Statement

The **switch** statement causes control to be transferred to one of several statements depending on the value of an expression. It has the form

switch (*expression*) *statement*

The usual arithmetic conversion is performed on the expression, but the result must be **int**. The statement is typically compound. Any statement within the statement may be labeled with one or more case prefixes as follows:

case *constant-expression* :

where the constant expression must be **int**. No two of the case constants in the same switch may have the same value. Constant expressions are precisely defined in Chapter 12.

There may also be at most one statement prefix of the form

default :

which properly goes at the end of the case constants.

When the **switch** statement is executed, its expression is evaluated and compared in turn with each case constant. If one of the case constants is equal to the value of the expression, control is passed to the statement following the matched case prefix. If no case constant matches the expression and if there is a **default** prefix, control passes to the statement prefixed by **default**. If no case matches and if there is no **default**, then none of the statements in the switch is executed.

The prefixes **case** and **default** do not alter the flow of control, which continues unimpeded across such prefixes. That is, once a case constant is matched, all **case** statements (and the **default**) from there to the end of the **switch** are executed. To exit from a switch, see Section 7.8, **break**.

Usually, the statement that is the subject of a switch is compound. Declarations may appear at the head of this statement, but initializations of automatic or register variables are ineffective.

A simple example of a complete **switch** statement is:

```
switch (c) {
    case 'o':
        oflag = TRUE;
        break;
    case 'p':
        pflag = TRUE;
        break;
    case 'r':
        rflag = TRUE;
        break;
    default :
        (void) fprintf(stderr,
            "Unknown option\n");
        exit (2);
}
```

7.8 break Statement

The statement **break** ; causes termination of the smallest enclosing **while**, **do**, **for**, or **switch** statement; control passes to the statement following the terminated statement.

7.9 continue Statement

The statement **continue** ; causes control to pass to the loop-continuation portion of the smallest enclosing **while**, **do**, or **for** statement; that is to the end of the loop. More precisely, in each of the statements

<pre>while (...)</pre>	<pre>do</pre>	<pre>for (...)</pre>
<pre>{</pre>	<pre>{</pre>	<pre>{</pre>
<pre>...</pre>	<pre>...</pre>	<pre>...</pre>
<pre>contin: ;</pre>	<pre>contin: ;</pre>	<pre>contin: ;</pre>
<pre>}</pre>	<pre>} while (...);</pre>	<pre>}</pre>

a **continue** is equivalent to **goto contin**. (Following the **contin:** is a null statement; see Section 7.13, Null Statement.)

7.10 return Statement

A function returns to its caller by means of the **return** statement, which has one of the forms

```
return ;  
return expression ;
```

In the first case (which is implicit if the end of the function is reached without executing a **return**) the returned value is undefined. In the second case, the value of the expression is returned to the caller of the function. If required, the expression is converted, as if by assignment, to the type of function in which it appears.

7.11 goto Statement

Control may be transferred unconditionally by means of the statement

```
goto identifier ;
```

The identifier must be a label (see Section 7.12, Labeled Statement.) located in the current function.

7.12 Labeled Statement

Any statement may be preceded by label prefixes of the form

```
identifier :
```

which serve to declare the identifier as a label. The only use of a label is as a target of a **goto**. The scope of a label is the current function, excluding any subblocks in which the same identifier has been redeclared. See Chapter 9, Scope Rules.

7.13 Null Statement

The null statement has the form

;

A null statement is useful to carry a label just before the } of a compound statement or to supply a null body to a looping statement such as **while**.



8. External Definitions

A C program consists of a sequence of external definitions. An external definition declares an identifier to have storage class **extern** (by default) or perhaps **static**, and a specified type. The type-specifier (see Section 6.2, Type Specifiers) may also be empty, in which case the type is taken to be **int**. The scope of external definitions persists to the end of the file in which they are declared just as the effect of declarations persists to the end of a block. The syntax of external definitions is the same as that of all declarations except that only at this level may the code for functions be given.

8.1 External Function Definitions

Function definitions have the form

function-definition:
decl-specifiers _{*opt*} *function-declarator* *function-body*

The only sc-specifiers allowed among the decl-specifiers are **extern** or **static**; see Section 9.2, Scope of Externals for the distinction between them. A function declarator is similar to a declarator for a "function returning ..." except that it lists the formal parameters of the function being defined.

function-declarator:
declarator (*parameter-list* _{*opt*})

parameter-list:
identifier
identifier , *parameter-list*

The function-body has the form

function-body:
*declaration-list*_{opt} *compound-statement*

The identifiers in the parameter list, and only those identifiers, may be declared in the declaration list. Any identifiers whose type is not given are taken to be **int**. The only storage class that may be specified is **register**; if it is specified, the corresponding actual parameter will be copied, if possible, into a register at the outset of the function.

A simple example of a complete function definition is

```
int max(a, b, c)
    int a, b, c;
{
    int m;

    m = (a > b) ? a : b;
    return((m > c) ? m : c);
}
```

Here **int** is the type-specifier; **max(a, b, c)** is the function-declarator; **int a, b, c;** is the declaration-list for the formal parameters; { ... } is the block giving the code for the statement.

The C program converts all **float** actual parameters to **double**, so formal parameters declared **float** have their declaration adjusted to read **double**. All **char** and **short** formal parameter declarations are similarly adjusted to read **int**. Also, since a reference to an array in any context (in particular as an actual parameter) is taken to mean a pointer to the first element of the array, declarations of formal parameters declared "array of ..." are adjusted to read "pointer to"

8.2 External Data Definitions

An external data definition has the form

data-definition:
declaration

The storage class of such data may be **extern** (which is the default) or **static**, but not **auto** or **register**.

9. Scope Rules

A C program need not all be compiled at the same time. The source text of the program may be kept in several files, and precompiled routines may be loaded from libraries. Communication among the functions of a program may be carried out both through explicit calls and through manipulation of external data.

Therefore, there are two kinds of scopes to consider: first, what may be called the lexical scope of an identifier, which is essentially the region of a program during which it may be used without drawing "undefined identifier" diagnostics; and second, the scope associated with external identifiers, which is characterized by the rule that references to the same external identifier are references to the same object.

9.1 Lexical Scope

The lexical scope of identifiers declared in external definitions persists from the definition through the end of the source file in which they appear. The lexical scope of identifiers that are formal parameters persists through the function with which they are associated. The lexical scope of identifiers declared at the head of a block persists until the end of the block. The lexical scope of labels is the whole of the function in which they appear.

In all cases, however, if an identifier is explicitly declared at the head of a block, including the block constituting a function, any declaration of that identifier outside the block is suspended until the end of the block.

Remember also that tags, identifiers associated with ordinary variables, and identities associated with structure and union members form three disjoint classes which do not conflict. Members and tags follow the same scope rules as other identifiers. The **enum** constants are in the same class as

ordinary variables and follow the same scope rules. The `typedef` names are in the same class as ordinary identifiers. They may be redeclared in inner blocks, but an explicit type must be given in the inner declaration:

```
typedef float distance;
...
{
    int distance;
    ...
}
```

The `int` must be present in the second declaration, or it would be taken to be a declaration with no declarators and type `distance`.

9.2 Scope of Externals

If a function refers to an identifier declared to be **extern**, then somewhere among the files or libraries constituting the complete program there must be at least one external definition for the identifier. All functions in a given program that refer to the same external identifier refer to the same object, so care must be taken that the type and size specified in the definition are compatible with those specified by each function that references the data.

It is illegal to explicitly initialize any external identifier more than once in the set of files and libraries comprising a multi-file program. It is legal to have more than one data definition for any external non-function identifier; explicit use of **extern** does not change the meaning of an external declaration.

In restricted environments, the use of the **extern** storage class takes on an additional meaning. In these environments, the explicit appearance of the **extern** keyword in external data declarations of identities without initialization indicates that the storage for the identifiers is allocated elsewhere, either in this file or another file. It is required that there be exactly one definition of each external identifier (without **extern**) in the set of files and libraries comprising a multiple-file program.

Identifiers declared **static** at the top level in external definitions are not visible in other files. Functions may be declared **static**.

10. Compiler Control Lines

The C compilation system contains a preprocessor capable of macro substitution, conditional compilation, and inclusion of named files. Lines beginning with # communicate with this preprocessor. There may be any number of blanks and horizontal tabs between the # and the directive, but no additional material (such as comments) is permitted. These lines have syntax independent of the rest of the language; they may appear anywhere and have effect that lasts (independent of scope) until the end of the source program file.

10.1 Token Replacement

A control line of the form

```
#define identifier token-stringopt
```

causes the preprocessor to replace subsequent instances of the identifier with the given string of tokens. Semicolons in or at the end of the token-string are part of that string. A line of the form

```
#define identifier(identifier, ... ) token-stringopt
```

where there is no space between the first identifier and the (, is a macro definition with arguments. There may be zero or more formal parameters. Subsequent instances of the first identifier followed by a (, a sequence of tokens delimited by commas, and a) are replaced by the token string in the definition. Each occurrence of an identifier mentioned in the formal parameter list of the definition is replaced by the corresponding token string from the call. The actual arguments in the call are token strings separated by commas; however, commas in quoted strings or protected by parentheses do not separate arguments. The number of formal and actual parameters must be the same. Strings and character constants in the token-string are

scanned for formal parameters, but strings and character constants in the rest of the program are not scanned for defined identifiers to replace.

In both forms the replacement string is rescanned for more defined identifiers. In both forms a long definition may be continued on another line by writing `\` at the end of the line to be continued. This facility is most valuable for definition of “manifest constants”, as in

```
#define TABSIZE 100  
  
int table[TABSIZE];
```

A control line of the form

`#undef identifier`

causes the identifier’s preprocessor definition (if any) to be forgotten.

If a **`#defined`** identifier is the subject of a subsequent **`#define`** with no intervening **`#undef`**, then the two token-strings are compared textually. If the two token-strings are not identical (all white space is considered as equivalent), then the identifier is considered to be redefined.

10.2 File Inclusion

A control line of the form

`#include "filename"`

causes the replacement of that line by the entire contents of the file *filename*. The named file is searched for first in the directory of the file containing the **`#include`**, and then in a sequence of specified or standard places.

Alternatively, a control line of the form

`#include <filename>`

searches only the specified or standard places and not the directory of the **`#include`**. (How the places are specified is not part of the language. See **`cpp(1)`** for a description of how to specify additional libraries.)

`#includes` may be nested.

10.3 Conditional Compilation

A compiler control line of the form

#if restricted-constant-expression

checks whether the restricted-constant expression evaluates to nonzero. (Constant expressions are discussed in Chapter 12, Constant Expressions; the following additional restrictions apply here: the constant expression may not contain **sizeof**, casts, or an enumeration constant.)

A restricted-constant expression may also contain the additional unary expression

defined identifier

or

defined (identifier)

which evaluates to one if the identifier is currently defined in the preprocessor and zero if it is not.

All currently defined identifiers in restricted-constant-expressions are replaced by their token-strings (except those identifiers modified by **defined**) just as in normal text. The restricted-constant expression will be evaluated only after all expressions have finished. During this evaluation, all undefined (to the procedure) identifiers evaluate to zero.

A control line of the form

#ifdef identifier

checks whether the identifier is currently defined in the preprocessor; i.e., whether it has been the subject of a **#define** control line. It is equivalent to **#if defined (identifier)**.

A control line of the form

#ifndef identifier

checks whether the identifier is currently undefined in the preprocessor. It is equivalent to **#if !defined (identifier)**.

All three forms are followed by an arbitrary number of lines, possibly containing a control line

#else

and then by a control line

#endif

If the checked condition is true, then any lines between **#else** and **#endif** are ignored. If the checked condition is false, then any lines between the test and a **#else** or, lacking a **#else**, the **#endif** are ignored.

Another control directive is

#elif *restricted-constant-expression*

An arbitrary number of **#elif** directives can be included between **#if**, **#ifdef**, or **#ifndef** and **#else**, or **#endif** directives. These constructions may be nested.

10.4 Line Control

For the benefit of other preprocessors that generate C programs, a line of the form

#line *constant* "*filename*"

causes the compiler to believe, for purposes of error diagnostics, that the line number of the next source line is given by the constant and the current input file is named by "*filename*". If "*filename*" is absent, the remembered file name does not change.

11. Types Revisited

This part summarizes the operations that can be performed on objects of certain types.

11.1 Structures and Unions

Structures and unions may be assigned, passed as arguments to functions, and returned by functions. Other plausible operators, such as equality comparison and structure casts, are not implemented.

In a reference to a structure or union member, the name on the right of the \rightarrow or the \cdot must specify a member of the aggregate named or pointed to by the expression on the left. In general, a member of a union may not be inspected unless the value of the union has been assigned using that same member. However, one special guarantee is made by the language in order to simplify the use of unions: if a union contains several structures that share a common initial sequence and if the union currently contains one of these structures, it is permitted to inspect the common initial part of any of the contained structures.

For example, the following is a legal fragment:

```
union
{
    struct
    {
        int      type;
    } n;
    struct
    {
        int      type;
        int      intnode;
    } ni;
    struct
    {
        int      type;
        float    floatnode;
    } nf;
} u;
...
u.nf.type = FLOAT;
u.nf.floatnode = 3.14;
...
if (u.n.type == FLOAT)
    ... sin(u.nf.floatnode) ...
```

11.2 Functions

There are only two things that can be done with a function: call it or take its address. If the name of a function appears in an expression not in the function-name position of a call, a pointer to the function is generated. Thus, to pass one function to another, one might say

```
int f();
...
g(f);
```

Then the definition of **g** might read

```
g (funcp)
    int (*funcp) ();
{
    ...
    (*funcp) ();
    ...
}
```

Notice that **f** must be declared explicitly in the calling routine since its appearance in **g(f)** was not followed by (.

11.3 Arrays, Pointers, and Subscripting

Every time an identifier of array type appears in an expression, it is converted into a pointer to the first member of the array. Because of this conversion, arrays are not lvalues. By definition, the subscript operator **[]** is interpreted in such a way that **E1[E2]** is identical to ***((E1)+(E2))**. Because of the conversion rules that apply to **+**, if **E1** is an array and **E2** an integer, then **E1[E2]** refers to the **E2** -th member of **E1**. Therefore, despite its asymmetric appearance, subscripting is a commutative operation.

A consistent rule is followed in the case of multidimensional arrays. If **E** is an *n*-dimensional array of rank **i×j×...×k**, then **E** appearing in an expression is converted to a pointer to an (n-1)-dimensional array with rank **j×...×k**. If the ***** operator, either explicitly or implicitly as a result of subscripting, is applied to this pointer, the result is the pointed-to (n-1)-dimensional array, which itself is immediately converted into a pointer.

For example, consider **int x[3][5]**; Here **x** is a 3×5 array of integers. When **x** appears in an expression, it is converted to a pointer to (the first of three) 5-membered arrays of integers. In the expression **x[i]**, which is equivalent to ***(x+i)**, **x** is first converted to a pointer as described; then **i** is converted to the type of **x**, which involves multiplying **i** by the length the object to which the pointer points, namely 5-integer objects. The results are added and indirection applied to yield an array (of five integers) which in turn is converted to a pointer to the first of the integers. If there is another subscript, the same argument applies again; this time the result is an integer.

Arrays in C are stored row-wise (last subscript varies fastest) and the first subscript in the declaration helps determine the amount of storage consumed by an array. Arrays play no other part in subscript calculations.

11.4 Explicit Pointer Conversions

Certain conversions involving pointers are permitted but have implementation-dependent aspects. They are all specified by means of an explicit type-conversion operator, see Section 5.2, Unary Operators and Section 6.8, Type Names.

A pointer may be converted to any of the integral types large enough to hold it. Whether an **int** or **long** is required is machine dependent. The mapping function is also machine dependent but is intended to be unsurprising to those who know the addressing structure of the machine.

An object of integral type may be explicitly converted to a pointer. The mapping always carries an integer converted from a pointer back to the same pointer but is otherwise machine dependent.

A pointer to one type may be converted to a pointer to another type. The resulting pointer may cause addressing exceptions upon use if the subject pointer does not refer to an object suitably aligned in storage. It is guaranteed that a pointer to an object of a given size may be converted to a pointer to an object of a smaller size and back again without change.

For example, a storage-allocation routine might accept a size (in bytes) of an object to allocate, and return a **char** pointer; it might be used in this way.

```
extern char *alloc();
double *dp;

dp = (double *) alloc(sizeof(double));
*dp = 22.0 / 7.0;
```

The **alloc** must ensure (in a machine-dependent way) that its return value is suitable for conversion to a pointer to **double**; then the use of the function is portable.

12. Constant Expressions

In several places C requires expressions that evaluate to a constant: after case, as array bounds, and in initializers. In the first two cases, the expression can involve only integer constants, character constants, casts to integral types, enumeration constants, and **sizeof** expressions, possibly connected by the binary operators

+ - * / % & | << >> == != < > <= >= && ||

or by the unary operators

-

or by the ternary operator

?:

Parentheses can be used for grouping but not for function calls.

More latitude is permitted for initializers; besides constant expressions as discussed above, one can also use floating constants and arbitrary casts and can also apply the unary & operator to external or static objects and to external or static arrays subscripted with a constant expression. The unary & can also be applied implicitly by appearance of unsubscripted arrays and functions. The basic rule is that initializers must evaluate either to a constant or to the address of a previously declared external or static object plus or minus a constant.



13. Portability Considerations

Certain parts of C are inherently machine dependent. The following list of potential trouble spots is not meant to be all-inclusive but to point out the main ones.

Purely hardware issues like word size and the properties of floating point arithmetic and integer division have proven in practice to be not much of a problem. Other facets of the hardware are reflected in differing implementations. Some of these, particularly sign extension (converting a negative character into a negative integer) and the order in which bytes are placed in a word, are nuisances that must be carefully watched. Most of the others are only minor problems.

The number of **register** variables that can actually be placed in registers varies from machine to machine as does the set of valid types. Nonetheless, the compilers all do things properly for their own machine; excess or invalid **register** declarations are ignored.

The order of evaluation of function arguments is not specified by the language. The order in which side effects take place is also unspecified.

Since character constants are really objects of type **int**, multicharacter character constants may be permitted. The specific implementation is very machine dependent because the order in which characters are assigned to a word varies from one machine to another.

Fields are assigned to words and characters to integers right to left on some machines and left to right on other machines. These differences are invisible to isolated programs that do not indulge in type punning (e.g., by converting an **int** pointer to a **char** pointer and inspecting the pointed-to storage) but must be accounted for when conforming to externally-imposed storage layouts.



14. Syntax Summary

This summary of C syntax is intended more for aiding comprehension than as an exact statement of the language. In particular, the syntax of function prototypes is not included. (See Section 6.4.)

14.1 Expressions

The basic expressions are:

expression:
primary
** expression*
& lvalue
- expression
! expression
expression
++ lvalue
-- lvalue
lvalue ++
lvalue --
sizeof expression
sizeof (type-name)
(type-name) expression
expression binop expression
expression ? expression : expression
lvalue asgnop expression
expression , expression

primary:
identifier
constant
string literal
(expression)
primary (expression-list opt)
primary [expression]
primary . identifier
primary -> identifier

lvalue:
identifier
primary [expression]
lvalue . identifier
primary -> identifier
** expression*
(lvalue)

The primary-expression operators

() [] . ->

have highest priority and group left to right. The unary operators

* & - ! ~ ++ -- sizeof (*type-name*)

have priority below the primary operators but higher than any binary operator and group right to left. Binary operators group left to right; they have priority decreasing as indicated below.

binop:
 * / %
 + -
 >> <<
 < > <= >=
 == !=
 &
 ^
 |
 &&
 ||

The conditional operator groups right to left.

Assignment operators all have the same priority and all group right to left.

asgnop:

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

The comma operator has the lowest priority and groups left to right.

14.2 Declarations

declaration:

decl-specifiers init-declarator-list_{opt} ;

decl-specifiers:

*type-specifier decl-specifiers_{opt}
sc-specifier decl-specifiers_{opt}*

sc-specifier:

auto
static
extern
register
typedef

type-specifier:

*struct-or-union-specifier
typedef-name
enum-specifier*

basic-type-specifier:

*basic-type
basic-type basic-type-specifiers*

basic-type:

char
short
int
long
unsigned
float
double
void

enum-specifier:

enum { *enum-list* }
enum *identifier* { *enum-list* }
enum *identifier*

enum-list:

enumerator
enum-list , *enumerator*

enumerator:

identifier
identifier = *constant-expression*

init-declarator-list:

init-declarator
init-declarator , *init-declarator-list*

init-declarator:

declarator *initializer* _{*opt*}

declarator:

identifier
(*declarator*)
* *declarator*
declarator ()
declarator [*constant-expression* _{*opt*}]

struct-or-union-specifier:
struct { *struct-decl-list* }
struct *identifier* { *struct-decl-list* }
struct *identifier*
union { *struct-decl-list* }
union *identifier* { *struct-decl-list* }
union *identifier*

struct-decl-list:
struct-declaration
struct-declaration struct-decl-list

struct-declaration:
type-specifier struct-declarator-list ;

struct-declarator-list:
struct-declarator
struct-declarator , struct-declarator-list

struct-declarator:
declarator
declarator : constant-expression
: constant-expression

initializer:
= *expression*
= { *initializer-list* }
= { *initializer-list* , }

initializer-list:
expression
initializer-list , initializer-list
{ *initializer-list* }
{ *initializer-list* , }

type-name:
type-specifier abstract-declarator

abstract-declarator:
empty
(*abstract-declarator*)
* *abstract-declarator*
abstract-declarator ()
abstract-declarator [*constant-expression*_{opt}]

typedef-name:
identifier

14.3 Statements

compound-statement:
{ declaration-list_{opt} statement-list_{opt} }

declaration-list:
declaration
declaration declaration-list

statement-list:
statement
statement statement-list

statement:
compound-statement
expression ;
if (expression) statement
if (expression) statement else statement
while (expression) statement
do statement while (expression) ;
for (exp_{opt} ; exp_{opt} ; exp_{opt}) statement
switch (expression) statement
case constant-expression : statement
default : statement
break ;
continue ;
return ;
return expression ;
goto identifier ;
identifier : statement
;

14.4 External Definitions

program:

external-definition

external-definition program

external-definition:

function-definition

data-definition

function-definition:

decl-specifier *opt* *function-declarator* *function-body*

function-declarator:

declarator (*parameter-list* *opt*)

parameter-list:

identifier

identifier , *parameter-list*

function-body:

declaration-list *opt* *compound-statement*

data-definition:

extern *declaration* ;

static *declaration* ;

14.5 Preprocessor

```
#define identifier token-stringopt  
#define identifier(identifier,...)token-stringopt  
#undef identifier  
#include "filename"  
#include <filename>  
#if restricted-constant-expression  
#ifdef identifier  
#ifndef identifier  
#elif restricted-constant-expression  
#else  
#endif  
#line constant "filename"
```

Appendix A: C on the IRIS-4D

The C language supported by the IRIS-4D Series *cc* compiler is an implementation of the language defined in *The C Programming Language* by Kernighan and Ritchie (Prentice Hall, 1978). This appendix covers the following topics:

- specifying *vararg* macros, a requirement for all functions that take a variable number of arguments
- deviations and extensions to the C language, as defined in *The C Programming Language* by Kernighan and Ritchie (Prentice-Hall)
- translation limits
- storage mapping

A.1 *vararg.h* Macros

If a function takes a variable number of arguments (for example, the C library functions *printf* and *scanf*), you must use the macros (defined in the *varargs.h* header file) shown below.

```
/*  @(#)varargs.h  1.2  */
typedef char *va_list;
#define va_dcl int va_alist;
#define va_start(list) list = (char *) &va_alist
#define va_end(list)
#define va_arg(list, mode) ((mode *) (list = \
    (char *) (sizeof(mode) > 4 ? (int)list + 2*8 - 1 & -8 \
    : (int)list + 2*4 - 1 & -4)))[-1]
```

The *va_dcl* macro declares the formal parameter *va_alist*, which is either the format descriptor for the remaining parameters or a parameter itself.

va_start must be called within the body of the function whose argument list is to be traversed. The function can then transverse the list or pass its *va_list* pointer to other functions to transverse the list. The *type* of the *va_start* argument is *va_list*; it is defined by the *typedef* statement in *varargs.h*.

The *va_arg* macro accesses the value of an argument rather than obtaining its address. The macro handles those type names that can be transformed into the appropriate pointer type by appending an asterisk (*), which handles most simple cases. The argument type in a variable argument list must never be an integer type smaller than **int**, and the type of the first argument must not be floating point.

For more information on the *varargs.h* macros, see the *varargs* man page in the *IRIS-4D User's Reference Manual*.

A.2 Deviations

IRIS-4D Series C does not support the **entry** keyword, which has no defined use. Additionally, IRIS-4D Series C does not support the **asm** keyword, as implemented by some C compilers to allow for the inclusion of assembly language instructions.

A.3 Extensions

Extensions to IRIS-4D Series C include the following:

- the **enumeration** type, a set of values represented by identifiers called enumeration constants; enumeration constants are specified when the type is defined. For information on the alignment, size, and value ranges of the **enumeration** type, see the section on “Storage Mapping” in this appendix.
- the **void** type, which allows you to specify that no value be returned from a function.

- the **volatile** type modifier, which is used when programming I/O devices and the **signed** type. In addition, the **const** keyword has been reserved for future use. For more information on the **volatile** modifier, see Section A.5, Storage Mapping section of this appendix.
- function prototypes which specify the types of some or all of the function's formal parameters. These may be used to force implicit coercion of actual parameters to the types of the corresponding formal parameters declared in the prototype. For more information, see Chapter 6.

A.4 Translation Limits

Table A-1 shows the maximum limits imposed on certain items by the C compiler.

C Specification	Maximum
Nesting levels	30
Compound statements	
Iterations	
Selections	
Conditional compilations	
Type modifiers	9
Case labels	500
Function call parameters	150
Significant characters	32
External identifiers	
Internal identifiers	

Table A-1. IRIS-4D Series C Limitations

A.5 Storage Mapping

This section describes how the compiler maps C variables into storage and contains the following topics:

- Alignment, Size, and Value Ranges
- Arrays, Structures, and Unions
- Storage Classes

A.5.1 Alignment, Size, and Value Ranges

Table A-2 describes how the C compiler implements size, alignment, and value ranges for the data types.

Type	Size	Alignment	Value Range	
			Signed	Unsigned
int	32 bits	Word ¹	-2^{31} to $2^{31} - 1$	0 to $2^{32} - 1$
long	32 bits	Word ¹	-2^{31} to $2^{31} - 1$	0 to $2^{32} - 1$
enum	32 bits	Word ¹	-2^{31} to $2^{31} - 1$	0 to $2^{32} - 1$
short	16 bits	Halfword ²	-32,768..32,767	0..65,535
char ⁴	8 bits	Byte	-128..127	0..255
float ⁵	32 bits	Word ¹	See note.	0..255
double ⁶	64 bits	Doubleword ³	See note.	0..255
pointer	32 bit	Word ¹		0 to $2^{32} - 1$

¹Byte boundary divisible by four.
²Byte boundary divisible by two.
³Byte boundary divisible by eight.
⁴char is assumed to be unsigned, unless the signed attribute is used.
⁵IEEE single precision. See note following this table for valid ranges.
⁶IEEE double precision. See note following this table for valid ranges.

Table A-2. Size, Alignment, and Value Ranges for C Data Types

NOTE: Approximate valid ranges for **float** and **double** are:

Maximum Value	
float	$3.40282356 \times 10^{38}$
double	$1.7976931348623158 \times 10^{308}$

Minimum Values		
	Denormalized	Normalized
float	$1.40129846 \times 10^{-46}$	$1.17549429 \times 10^{-38}$
double	$4.9406564584124654 \times 10^{-324}$	$2.2250738585072012 \times 10^{-308}$

For characters to be treated as signed, either use the compiler option **-signed**, or use the keyword **signed** in conjunction with **char**, as shown in the following example:

```
signed char c;
```

The header files *limits.h* and *float.h* (usually found in */usr/include*) contain C macros that define minimum and maximum values for the various data types. Refer to these files for the macro names and values.

A.5.2 Arrays, Structures, and Unions

Arrays. Arrays have the same boundary requirements as the data type specified for the array. The size of an array is the size of the data type multiplied by the number of elements. For example, for the following declaration:

```
double x[2][3]
```

The size of the resulting array is 48 ($2 \times 3 \times 8$, where 8 is the size of the **double** floating point type).

Structures. Each member of a structure begins at an offset from the structure base. The offset corresponds to the order in which a member is declared; the first member is at offset 0.

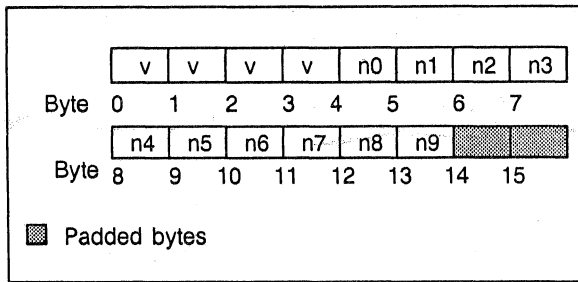
The size of a structure in the object file is the size of its combined members plus padding added, where necessary, by the compiler. The following rules apply to structures:

- Structures must align on the same boundary as that required by the member with the most restrictive boundary requirement. The boundary requirements by degree of restrictiveness are: byte, halfword, word, and doubleword, with doubleword being the most restrictive.
- The compiler terminates the structure on the same alignment boundary on which it begins. For example, if a structure begins on an even-byte boundary, it also ends on an even-byte boundary.

For example, the following structure:

```
struct S {  
    int v;  
    char n[10];  
}
```

is mapped out in storage as follows:



Note that the length of the structure is 16 bytes, even though the byte count as defined by the `int v` and the `char n` component is only 14. Because `int` has a stricter boundary requirement (word boundary) than `char` (byte boundary), the structure must end on a word boundary (a byte offset divisible by four). The compiler therefore adds two bytes of padding to meet this requirement.

An array of data structures illustrates the reason for this requirement. For example, if the above structure were the element-type of an array, some of the `int v` components wouldn't be aligned properly without the two-byte pad.

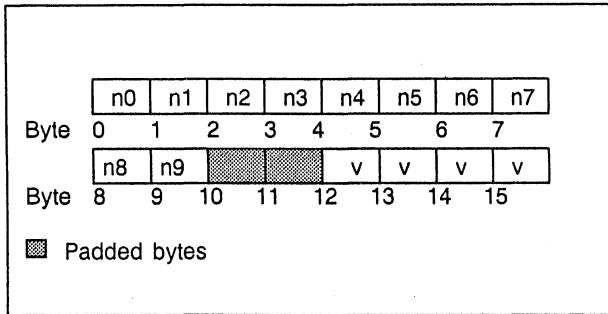
Alignment requirements may cause padding to appear in the middle of a structure. For example, by rearranging the structure in the last example to the following:

```

struct S {
    char n[10];
    int v;
}

```

The compiler maps the structures as follows:



Note that the size of the structure remains 16 bytes, but two bytes of padding follow the `n` component to align `v` on a word boundary.


Bit fields are packed from the most significant bit to least significant bit in a word and can be no longer than 32 bits; bit fields can be signed or unsigned. The following structure:

```

typedef struct {
    unsigned offset      :12;
    unsigned page       :10;
    unsigned segment    : 9;
    unsigned supervisor : 1;
} virtual_address;

```


is mapped out as follows:




The compiler moves fields that overlap a word boundary to the next word.

The compiler aligns a nonbit field that follows a bit-field declaration to the next boundary appropriate for its type. For example, the following structure:

```
struct {
    unsigned a :3;
    char      b;
    short     c;
}x;
```

is mapped out as follows:



Note that five bits of padding are added after **unsigned a** so that **char b** aligns on a byte boundary, as required.

Unions. A union must align on the same boundary as the member with the most restrictive boundary requirement. The boundary requirements by increasing degree of restrictiveness are: byte, halfword, word, and doubleword. For example, a union containing **char**, **int**, and **double** data types must align on a doubleword boundary, as required by the **double** data type.

A.5.3 Storage Classes

Auto. An **auto** declaration indicates that storage is allocated at execution and exists only for the duration of that block activation.

Static. The compiler allocates storage for a **static** declaration at compile time. This allocation remains fixed for the duration of the program. Static variables reside in the program bss section if they are not initialized, otherwise they are placed in the data section.

Register. The compiler allocates variables with the **register** storage class to registers. For programs compiled using the **-O** (optimize) option, the optimization phase of the compiler tries to assign all variables to registers, regardless of the storage class specified.

Extern. The **extern** storage class indicates that the variable refers to storage defined elsewhere in an external data definition. The compiler doesn't allocate storage to **extern** variable declarations; it uses the following logic in defining and referencing them:

Extern is omitted. If an initializer is present, a definition for the symbol is emitted. Having two or more such definitions among all the files comprising a program results in an error at link time or before. If no initializer is present, a common definition is emitted. Any number of common definitions of the same identifier may coexist.

Extern is present. The compiler assumes that declaration refers to a name defined elsewhere. A declaration having an initializer is illegal. If a declared identifier is never used, the compiler doesn't issue an external reference to the linker.

Volatile. The **volatile** storage class is specified for those variables that may be modified in ways unknown to the compiler. For example, **volatile** might be specified for an object corresponding to a memory mapped input/output port or an object accessed by an asynchronously interrupting function. Except for expression evaluation, no phase of the compiler optimizes any of the code dealing with **volatile** objects.

NOTE. If a pointer specified as **volatile** is assigned to another pointer without the volatile specification, the compiler treats the other pointer as non-volatile. In the following example:

```
volatile int *i;
int *j;
.
.
(volatile*)j = i;
3108282356*10
```

the compiler treats *j* as a non-volatile pointer and the object it points to as non-volatile, and may optimize it.

The compiler option **-volatile** causes all objects to be compiled as volatile.

A.6 Compiler Options

The following is a list of the compiler options specific to MIPS C. For a list of general options, see Chapter 1 of the *IRIS-4D Series Compiler Guide*.

- signed** Causes all **char** declarations to be **signed char** declarations; the default is to treat **char** as **unsigned char**.
- volatile** Causes all variables to be treated as **volatile**.
- varargs** Prints warnings for lines that may require the *varargs.h* macros.
- float** Normally, expressions involving floating-point are always computed in double-precision in C, even if the highest precision involved is float. This option allows the compiler to evaluate floating-point expressions which do not involve double-precision data in single-precision. This option has no effect on the promotion of floats to doubles when they are passed as parameters. **NOTE:** This switch is non-standard and may not be supported across product lines.



Silicon Graphics, Inc.

Date _____

Your name _____

Title _____

Department _____

Company _____

Address _____

Phone _____

COMMENTS

Manual title and version _____

Please list any errors, inaccuracies, or omissions you have found in this manual

Please list any suggestions you may have for improving this manual



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 45 MOUNTAIN VIEW, CA

POSTAGE WILL BE PAID BY ADDRESSEE

Silicon Graphics, Inc.

Attention: Technical Publications

2011 Stierlin Road

Mountain View, CA 94043-1321

