

Cforall Reference Manual and Rationale

Glen Ditchfield

DRAFT
May 10, 2015

© 2015 Glen Ditchfield

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Contents

Introduction	1
3 Terms, definitions, and symbols	2
6 Language	3
6.1 Notation	3
6.2 Concepts	3
6.2.1 Scopes of identifiers	3
6.2.2 Linkage of identifiers	3
6.3 Conversions	4
6.3.1 Arithmetic operands	4
6.3.2 Other operands	4
6.3.3 Conversion Cost	6
6.4 Lexical elements	7
6.4.1 Keywords	7
6.4.2 Identifiers	7
6.5 Expressions	9
6.5.1 Primary expressions	10
6.5.2 Postfix operators	11
6.5.3 Unary operators	22
6.5.4 Cast operators	30
6.5.5 Multiplicative operators	30
6.5.6 Additive operators	33
6.5.7 Bitwise shift operators	36
6.5.8 Relational operators	36
6.5.9 Equality operators	38
6.5.10 Bitwise AND operator	40
6.5.11 Bitwise exclusive OR operator	41
6.5.12 Bitwise inclusive OR operator	41
6.5.13 Logical AND operator	42
6.5.14 Logical OR operator	42
6.5.15 Conditional operator	43
6.5.16 Assignment operators	45
6.5.17 Comma operator	57
6.6 Constant expressions	57
6.7 Declarations	57
6.7.2 Type specifiers	58
6.7.3 Type qualifiers	61

6.7.9	Initialization	63
6.7.11	Specification definitions	63
6.7.12	Type declarations	65
6.8	Statements and blocks	71
6.8.3	Expression and null statements	71
6.8.4	Selection statements	71
6.8.5	Iteration statements	71
6.8.6	Jump statements	72
6.10	Preprocessing directives	72
6.10.8	Predefined macro names	72
A	Examples	73
A.1	C types	73
A.1.1	Scalar, arithmetic, and integral types	73
A.1.2	Modifiable types	73
A.1.3	Pointer and array types	74
A.2	Relationships between operations	76
A.2.1	Relational and equality operators	76
A.2.2	Arithmetic and integer operations	77
B	TODO	78

Introduction

This document is a reference manual and rationale for Cforall, a polymorphic extension of the C programming language. It makes frequent reference to the ISO/IEC C standard [1], and occasionally compares Cforall to C++ [5].

The manual deliberately imitates the ordering of the ISO/IEC C standard (although the section numbering differs). Unfortunately, this means that the manual contains more “forward references” than usual, and that it will be hard to follow if the reader does not have a copy of the ISO/IEC C standard near-by. For a gentle introduction to Cforall, see the companion document “An Overview of Cforall” [4].

□ Commentary (like this) is quoted with quads. Commentary usually deals with subtle points, the rationale behind a rule, and design decisions. □

Chapter 3

Terms, definitions, and symbols

Terms from the ISO/IEC C standard used in this document have the same meaning as in the ISO/IEC C standard.

Chapter 6

Language

6.1 Notation

The syntax notation used in this document is the same as is used in the ISO/IEC C standard, with one exception: ellipsis in the definition of a nonterminal, as in “*declaration: ...*”, indicates that these rules extend a previous definition, which occurs in this document or in the ISO/IEC C standard.

6.2 Concepts

6.2.1 Scopes of identifiers

Cforall’s scope rules differ from C’s in one major respect: a declaration of an identifier may overload outer declarations of lexically identical identifiers in the same name space, instead of hiding them. The outer declaration is hidden if the two declarations have compatible type, or if one declares an array type and the other declares a pointer type and the element type and pointed-at type are compatible, or if one has function type and the other is a pointer to a compatible function type, or if one declaration is a **type** or **typedef** declaration and the other is not. The outer declaration becomes visible when the scope of the inner declaration terminates.

- Hence, a Cforall program can declare an **int v** and a **float v** in the same scope; a C++ program can not. □

6.2.2 Linkage of identifiers

Cforall’s linkage rules differ from C’s in only one respect: instances of a particular identifier with external or internal linkage do not necessarily denote the same object or function. Instead, in the set of translation units and libraries that constitutes an entire program, any two instances of a particular identifier with external linkage denote the same object or function if they have compatible types, or if one declares an array type and the other declares a pointer type and the element type and pointed-at type are compatible, or if one has function type and the other is a pointer to a compatible function type. Within one translation unit, each instance of an identifier with internal linkage denotes the same object or function in the same circumstances. Identifiers with no linkage always denote unique entities.

- A Cforall program can declare an **extern int v** and an **extern float v**; a C program cannot. □

6.3 Conversions

Cforall defines situations where values of one type are automatically converted to another type. These conversions are called *implicit conversions*. The programmer can request *explicit conversions* using cast expressions.

6.3.1 Arithmetic operands

6.3.1.8 Safe arithmetic conversions

In C, a pattern of conversions known as the *usual arithmetic conversions* is used with most binary arithmetic operators to convert the operands to a common type and determine the type of the operator's result. In Cforall, these conversions play a role in overload resolution, and collectively are called the *safe arithmetic conversions*.

Let int_r and $unsigned_r$ be the signed and unsigned integer types with integer conversion rank r . Let $unsigned_{mr}$ be the unsigned integer type with maximal rank.

The following conversions are *direct* safe arithmetic conversions.

- The integer promotions.
- For every rank r greater than or equal to the rank of `int`, conversion from int_r to $unsigned_r$.
- For every rank r greater than or equal to the rank of `int`, where int_{r+1} exists and can represent all values of $unsigned_r$, conversion from $unsigned_r$ to int_{r+1} .
- Conversion from $unsigned_{mr}$ to `float`.
- Conversion from an enumerated type to its compatible integer type.
- Conversion from `float` to `double`, and from `double` to `long double`.
- Conversion from `float _Complex` to `double _Complex`, and from `double _Complex` to `long double _Complex`.
- Conversion from `float _Imaginary` to `double _Imaginary`, and from `double _Imaginary` to `long double _Imaginary`, if the implementation supports imaginary types.

If type T can be converted to type U by a safe direct arithmetic conversion and type U can be converted to type V by a safe arithmetic conversion, then the conversion from T to type V is an *indirect* safe arithmetic conversion.

- Note that ISO/IEC C does not include conversion from real types to complex types in the usual arithmetic conversions, and Cforall does not include them as safe conversions. □

6.3.2 Other operands

6.3.2.4 Anonymous structures and unions

If an expression's type is a pointer to a structure or union type that has a member that is an anonymous structure or an anonymous union, it can be implicitly converted to a pointer to the anonymous structure's or anonymous union's type. The result of the conversion is a pointer to the member.

Examples

```

struct point {
    int x, y;
};
void move_by(struct point* p1, struct point* p2) {
    p1->x += p2.x;
    p1->y += p2.y;
}

struct color_point {
    enum { RED, BLUE, GREEN } color;
    struct point;
} cp1, cp2;
move_to(&cp1, &cp2);

```

Thanks to implicit conversion, the two arguments that `move_by()` receives are pointers to `cp1`'s second member and `cp2`'s second member.

6.3.2.5 Specialization

A function or value whose type is polymorphic may be converted to one whose type is less polymorphic by binding values to one or more of its inferred parameters. Any value that is legal for the inferred parameter may be used, including other inferred parameters.

If, after the inferred parameter binding, an assertion parameter has no inferred parameters in its type, then an object or function must be visible at the point of the specialization that has the same identifier as the assertion parameter and has a type that is compatible with or can be specialized to the type of the assertion parameter. The assertion parameter is bound to that object or function.

The type of the specialization is the type of the original with the bound inferred parameters and the bound assertion parameters replaced by their bound values.

Examples

The type

```
forall(type T, type U) void (*)(T, U);
```

can be specialized to (among other things)

```
forall(type T) void (*)(T, T);    /* U bound to T */
forall(type T) void (*)(T, real); /* U bound to real */
forall(type U) void (*)(real, U); /* T bound to real */
void f(real, real);               /* both bound to real */

```

The type

```
forall(type T | T ?+?(T,T)) T (*)(T);
```

can be specialized to (among other things)

```
int (*)(int);                    /* T bound to int, and */
                                  /* T ?+?(T,T) bound to int ?+?(int,int) */

```

6.3.2.6 Safe conversions

A *direct safe conversion* is one of the following conversions:

- a direct safe arithmetic conversion;
- from any object type or incomplete type to `void`;
- from a pointer to any non-void type to a pointer to `void`;
- from a pointer to any type to a pointer to a more qualified version of the type;
- from a pointer to a structure or union type to a pointer to the type of a member of the structure or union that is an anonymous structure or an anonymous union;
- within the scope of an initialized type declaration, conversions between a type and its implementation or between a pointer to a type and a pointer to its implementation.

Conversions that are not safe conversions are *unsafe conversions*.

- As in C, there is an implicit conversion from `void*` to any pointer type. This is clearly dangerous, and C++ does not have this implicit conversion. Cforall keeps it, in the interest of remaining as pure a superset of C as possible, but discourages it by making it unsafe. □

6.3.3 Conversion Cost

The *conversion cost* of a safe conversion¹ is a measure of how desirable or undesirable it is. It is defined as follows.

- The cost of a conversion from any type to itself is 0.
- The cost of a direct safe conversion is 1.
- The cost of an indirect safe arithmetic conversion is the smallest number of direct conversions needed to make up the conversion.

Examples

In the following, assume an implementation that does not provide any extended integer types.

- The cost of an implicit conversion from `int` to `long` is 1. The cost of an implicit conversion from `long` to `double` is 3, because it is defined in terms of conversions from `long` to `unsigned long`, then to `float`, and then to `double`.
- If `int` can represent all the values of `unsigned short`, then the cost of an implicit conversion from `unsigned short` to `unsigned` is 2: `unsigned short` to `int` to `unsigned`. Otherwise, `unsigned short` is converted directly to `unsigned`, and the cost is 1.
- If `long` can represent all the values of `unsigned`, then the conversion cost of `unsigned` to `long` is 1. Otherwise, the conversion is an unsafe conversion, and its conversion cost is undefined.

¹Unsafe conversions do not have defined conversion costs.

6.4 Lexical elements

6.4.1 Keywords

Syntax

```
keyword: ...
        forall
        lvalue
        spec
        dtype
        ftype
        type
```

6.4.2 Identifiers

Cforall allows operator overloading by associating operators with special function identifiers. Furthermore, the constants “0” and “1” have special status for many of C’s data types (and for many programmer-defined data types as well), so Cforall treats them as overloadable identifiers. Programmers can use these identifiers to declare functions and objects that implement operators and constants for their own types.

6.4.2.3 Constant identifiers

Syntax

```
identifier: ...
          0
          1
```

The tokens “0” and “1” are identifiers. No other tokens defined by the rules for integer constants are considered to be identifiers.

- Why “0” and “1”? Those integers have special status in C. All scalar types can be incremented and decremented, which is defined in terms of adding or subtracting 1. The operations “&&”, “||”, and “!” can be applied to any scalar arguments, and are defined in terms of comparison against 0. A *constant-expression* that evaluates to 0 is effectively compatible with every pointer type.

In C, the integer constants 0 and 1 suffice because the integer promotion rules can convert them to any arithmetic type, and the rules for pointer expressions treat constant expressions evaluating to 0 as a special case. However, user-defined arithmetic types often need the equivalent of a 1 or 0 for their functions or operators, polymorphic functions often need 0 and 1 constants of a type matching their polymorphic parameters, and user-defined pointer-like types may need a null value. Defining special constants for a user-defined type is more efficient than defining a conversion to the type from `_Bool`.

Why *just* “0” and “1”? Why not other integers? No other integers have special status in C. A facility that let programmers declare specific constants—“`const Rational 12`”, for instance—would not be much of an improvement. Some facility for defining the creation of values of programmer-defined types from arbitrary integer tokens would be needed. The complexity of such a feature doesn’t seem worth the gain. □

6.4.2.4 Operator identifiers

Table 6.1 lists the programmer-definable operator identifiers and the operations they are associated with. Functions that are declared with (or pointed at by function pointers that are declared with) these identifiers can be called by expressions that use the operator tokens and syntax, or the operator identifiers and “function call” syntax. The relationships between operators and function calls are discussed in descriptions of the operators.

?[?]	subscripting	?%?	remainder	?^?	exclusive OR
?()	function call	?+?	addition	? ?	inclusive OR
?++	postfix increment	?-?	subtraction	?=?	simple assignment
?--	postfix decrement	?<<?	left shift	?*=?	multiplication assignment
++?	prefix increment	?>>?	right shift	?/=?	division assignment
--?	prefix decrement	?<?	less than	?%=?	remainder assignment
*?	dereference	?<=?	less than or equal	?+=?	addition assignment
+?	unary plus	?>=?	greater than or equal	?-=?	subtraction assignment
-?	arithmetic negation	?>?	greater than	?<<=?	left-shift assignment
~?	bitwise negation	?==?	equality	?>>=?	right-shift assignment
!?	logical complement	?!=?	inequality	?&=?	bitwise AND assignment
?*?	multiplication	?&?	bitwise AND	?^=?	exclusive OR assignment
?/?	division			? =?	inclusive OR assignment

Table 6.1: Operator Identifiers

□ Operator identifiers are made up of the characters of the operator token, with question marks added to mark the positions of the arguments of operators. The question marks serve as mnemonic devices; programmers can not create new operators by arbitrarily mixing question marks and other non-alphabetic characters. Note that prefix and postfix versions of the increment and decrement operators are distinguished by the position of the question mark. □

□ The use of “?” in identifiers means that some C programs are not Cforall programs. For instance, the sequence of characters “(i < 0)?--i:i” is legal in a C program, but a Cforall compiler will detect a syntax error because it will treat “?--” as an identifier, not as the two tokens “?” and “--”. □

□ Certain operators *cannot* be defined by the programmer:

- The logical operators “&&” and “||”, and the conditional operator “?:”. These operators do not always evaluate their operands, and hence can not be properly defined by functions unless some mechanism like call-by-name is added to the language. Note that the definitions of “&&” and “||” say that they work by checking that their arguments are unequal to 0, so defining “!=” and “0” for user-defined types is enough to allow them to be used in logical expressions.
- The comma operator. It is a control-flow operator like those above. Changing its meaning seems pointless and confusing.
- The “address of” operator. It would seem useful to define a unary “&” operator that returns values of some programmer-defined pointer-like type. The problem lies with the type of the operator. Consider the expression “p = &x”, where x is of type T and p has the programmer-defined type T_ptr. The expression might be treated as a call to the unary function “&?”. Now what is the type of the function’s parameter? It can not be T, because then x would

be passed by value, and there is no way to create a useful pointer-like result from a value. Hence the parameter must have type `T*`. But then the expression must be rewritten as “`p = &?(&x)`”—which doesn’t seem like progress!

The rule for address-of expressions would have to be something like “keep applying address-of functions until you get one that takes a pointer argument, then use the built-in operator and stop”. It seems simpler to define a conversion function from `T*` to `T_ptr`.

- The `sizeof` operator. It is already defined for every object type, and intimately tied into the language’s storage allocation model. Redefining it seems pointless.
- The “member of” operators “.” and “->”. These are not really infix operators, since their right “operand” is not a value or object.
- Cast operators. Anything that can be done with an explicit cast can be done with a function call. The difference in syntax is small.

□

6.5 Expressions

`Cforall` allows operators and identifiers to be overloaded. Hence, each expression can have a number of *interpretations*, each of which has a different type. The interpretations that are potentially executable are called *valid interpretations*. The set of interpretations depends on the kind of expression and on the interpretations of the subexpressions that it contains. The rules for determining the valid interpretations of an expression are discussed below for each kind of expression. Eventually the context of the outermost expression chooses one interpretation of that expression to be executed.

An *ambiguous interpretation* is an interpretation which does not specify the exact object or function denoted by every identifier in the expression. An expression can have some interpretations that are ambiguous and others that are unambiguous. An expression that is chosen to be executed shall not be ambiguous.

The *best valid interpretations* are the valid interpretations that use the fewest unsafe conversions. Of these, the best are those where the functions and objects involved are the least polymorphic. Of these, the best have the lowest total conversion cost, including all implicit conversions in the argument expressions. Of these, the best have the highest total conversion cost for the implicit conversions (if any) applied to the argument expressions. If there is no single best valid interpretation, or if the best valid interpretation is ambiguous, then the resulting interpretation is ambiguous.

□ `Cforall`’s rules for selecting the best interpretation are designed to allow overload resolution to mimic C’s operator semantics. In C, the “usual arithmetic conversions” are applied to the operands of binary operators if necessary to convert the operands to a common type (roughly speaking, the “smallest” type that can hold both arguments). In `Cforall`, those conversions are “safe”. The “fewest unsafe conversions” rule ensures that the usual conversions are done, if possible. The “lowest total expression cost” rule chooses the proper common type. The odd-looking “highest argument conversion cost” rule ensures that, when unary expressions must be converted, conversions of function results are preferred to conversion of function arguments: `(double)-i` will be preferred to `-(double)i`.

The “least polymorphic” rule reduces the number of polymorphic function calls, since such functions are presumably more expensive than monomorphic functions and since the more specific function is presumably more appropriate. It also gives preference to monomorphic values (such as the `int 0`) over polymorphic values (such as the null pointer `0`). However, interpretations that call polymorphic functions are preferred to interpretations that perform unsafe conversions, because those conversions potentially lose accuracy or violate strong typing.

There are two notable differences between Cforall’s overload resolution rules and the rules for C++ defined in [5]. First, the result type of a function plays a role. In C++, a function call must be completely resolved based on the arguments to the call in most circumstances. In Cforall, a function call may have several interpretations, each with a different result type, and the interpretations of the containing context choose among them. Second, safe conversions are used to choose among interpretations of all sorts of functions; in C++, the “usual arithmetic conversions” are a separate set of rules that apply only to the built-in operators. □

Expressions involving certain operators are considered to be equivalent to function calls. A transformation from “operator” syntax to “function call” syntax is defined by *rewrite rules*. Each operator has a set of predefined functions that overload its identifier. Overload resolution determines which member of the set is executed in a given expression. The functions have internal linkage and are implicitly declared with file scope. The predefined functions and rewrite rules are discussed below for each of these operators.

□ Predefined functions and constants have internal linkage because that simplifies optimization in traditional compile-and-link environments. For instance, “`an_int + an_int`” is equivalent to “`?+?(an_int, an_int)`”. If integer addition has not been redefined in the current scope, a compiler can generate code to perform the addition directly. If predefined functions had external linkage, this optimization would be difficult. □

□ Since each subsection describes the interpretations of an expression in terms of the interpretations of its subexpressions, this chapter can be taken as describing an overload resolution algorithm that uses one bottom-up pass over an expression tree. Such an algorithm was first described (for Ada) by Baker [2]. It is extended here to handle polymorphic functions and arithmetic conversions. The overload resolution rules and the predefined functions have been chosen so that, in programs that do not introduce overloaded declarations, expressions will have the same meaning in C and in Cforall. □

□ Expression syntax is quoted from the ISO/IEC C standard. The syntax itself defines the precedence and associativity of operators. The sections are arranged in decreasing order of precedence, with all operators in a section having the same precedence. □

6.5.1 Primary expressions

Syntax

primary-expression:
identifier
constant
string-literal
(expression)
generic-selection

Predefined Identifiers

```
const int          1;
const int          0;
forall(dtype DT) DT *const 0;
forall(ftype FT) FT *const 0;
```

Semantics

The valid interpretations of an *identifier* are given by the visible declarations of the identifier.

A *constant* or *string-literal* has one valid interpretation, which has the type and value defined by ISO/IEC C. The predefined integer identifiers “1” and “0” have the integer values 1 and 0, respectively. The other two predefined “0” identifiers are bound to polymorphic pointer values that, when specialized with a data type or function type respectively, produce a null pointer of that type.

A parenthesised expression has the same interpretations as the contained *expression*.

Examples

The expression `(void*)0` specializes the (polymorphic) null pointer to a null pointer to `void`. `(const void*)0` does the same, and also uses a safe conversion from `void*` to `const void*`. In each case, the null pointer conversion is better than the unsafe conversion of the integer 0 to a pointer.

- Note that the predefined identifiers have addresses.

`Cforall` does not have C’s concept of “null pointer constants”, which are not typed values but special strings of tokens. The C token “0” is an expression of type `int` with the value “zero”, and it *also* is a null pointer constant. Similarly, “`(void*)0`” is an expression of type `(void*)` whose value is a null pointer, and it also is a null pointer constant. However, in C, “`(void*)(void*)0`” is *not* a null pointer constant, even though it is null-valued, a pointer, and constant! The semantics of C expressions contain many special cases to deal with subexpressions that are null pointer constants.

`Cforall` handles these cases through overload resolution. The declaration

```
forall(dtype DT) DT *const 0;
```

means that `0` is a polymorphic object, and contains a value that can have *any* pointer-to-object type or pointer-to-incomplete type. The only such value is the null pointer. Therefore the type *alone* is enough to identify a null pointer. Where C defines an operator with a special case for the null pointer constant, `Cforall` defines predefined functions with a polymorphic object parameter.

□

6.5.1.1 Generic selection**Constraints**

The best interpretation of the controlling expression shall be unambiguous, and shall have type compatible with at most one of the types named in its generic association list. If a generic selection has no default generic association, the best interpretation of its controlling expression shall have type compatible with exactly one of the types named in its generic association list.

Semantics

A generic selection has the same interpretations as its result expression.

6.5.2 Postfix operators**Syntax**

postfix-expression:

primary-expression

postfix-expression [*expression*]

postfix-expression (*argument-expression-list*_{opt})

postfix-expression . *identifier*

```

postfix-expression -> identifier
postfix-expression ++
postfix-expression --
( type-name ) \{ initializer-list \}
( type-name ) \{ initializer-list , \}
argument-expression-list:
assignment-expression
argument-expression-list , assignment-expression

```

Rewrite Rules

```

a[b] ⇒ ?[?](b, a)  if a has integer type
a[b] ⇒ ?[?](a, b)  otherwise
a( arguments ) ⇒ ?()(a, arguments )
a++ ⇒ ?++(&a)
a-- ⇒ ?--(&a)

```

6.5.2.1 Array subscripting

Predefined Identifiers

```

forall(type T) lvalue T           ?[?](T*, ptrdiff_t);
forall(type T) lvalue _Atomic T   ?[?]( _Atomic T*, ptrdiff_t);
forall(type T) lvalue const T     ?[?](const T*, ptrdiff_t);
forall(type T) lvalue restrict T  ?[?](restrict T*, ptrdiff_t);
forall(type T) lvalue volatile T  ?[?](volatile T*, ptrdiff_t);
forall(type T) lvalue _Atomic const T ?[?]( _Atomic const T*, ptrdiff_t);
forall(type T) lvalue _Atomic restrict T ?[?]( _Atomic restrict T*, ptrdiff_t);
forall(type T) lvalue _Atomic volatile T ?[?]( _Atomic volatile T*, ptrdiff_t);
forall(type T) lvalue const restrict T ?[?](const restrict T*, ptrdiff_t);
forall(type T) lvalue const volatile T ?[?](const volatile T*, ptrdiff_t);
forall(type T) lvalue restrict volatile T
    ?[?](restrict volatile T*, ptrdiff_t);
forall(type T) lvalue _Atomic const restrict T
    ?[?]( _Atomic const restrict T*, ptrdiff_t);
forall(type T) lvalue _Atomic const volatile T
    ?[?]( _Atomic const volatile T*, ptrdiff_t);
forall(type T) lvalue _Atomic restrict volatile T
    ?[?]( _Atomic restrict volatile T*, ptrdiff_t);
forall(type T) lvalue const restrict volatile T
    ?[?](const restrict volatile T*, ptrdiff_t);
forall(type T) lvalue _Atomic const restrict volatile T
    ?[?]( _Atomic const restrict volatile T*, ptrdiff_t);

```

Semantics

The interpretations of subscript expressions are the interpretations of the corresponding function call expressions.

□ C defines subscripting as pointer arithmetic in a way that makes $a[i]$ and $i[a]$ equivalent. Cforall provides the equivalence through a rewrite rule to reduce the number of overloads of $?[?]$.

Subscript expressions are rewritten as function calls that pass the first parameter by value. This is somewhat unfortunate, since array-like types tend to be large. The alternative is to use the rewrite rule “ $a[b] \Rightarrow ?[?](\&a, b)$ ”. However, C semantics forbid this approach: the a in “ $a[b]$ ” can be an arbitrary pointer value, which does not have an address.

The repetitive form of the predefined identifiers shows up a deficiency of Cforall’s type system. Type qualifiers are not included in type values, so polymorphic functions that take pointers to arbitrary types often come in one flavor for each possible qualification of the pointed-at type. □

6.5.2.2 Function calls

Semantics

A *function designator* is an interpretation of an expression that has function type. The *postfix-expression* in a function call may have some interpretations that are function designators and some that are not.

For those interpretations of the *postfix-expression* that are not function designators, the expression is rewritten and becomes a call of a function named “ $?()$ ”. The valid interpretations of the rewritten expression are determined in the manner described below.

Each combination of function designators and argument interpretations is considered. For those interpretations of the *postfix-expression* that are monomorphic function designators, the combination has a valid interpretation if the function designator accepts the number of arguments given, and each argument interpretation matches the corresponding explicit parameter:

- if the argument corresponds to a parameter in the function designator’s prototype, the argument interpretation must have the same type as the corresponding parameter, or be implicitly convertible to the parameter’s type
- if the function designator’s type does not include a prototype or if the argument corresponds to “ \dots ” in a prototype, a default argument promotion is applied to it.

The type of the valid interpretation is the return type of the function designator.

For those combinations where the interpretation of the *postfix-expression* is a polymorphic function designator and the function designator accepts the number of arguments given, there shall be at least one set of *implicit arguments* for the implicit parameters such that

- If the declaration of the implicit parameter uses type-class **type**, the implicit argument must be an object type; if it uses **dtype**, the implicit argument must be an object type or an incomplete type; and if it uses **ftype**, the implicit argument must be a function type.
- if an explicit parameter’s type uses any implicit parameters, then the corresponding explicit argument must have a type that is (or can be safely converted to) the type produced by substituting the implicit arguments for the implicit parameters in the explicit parameter type.
- the remaining explicit arguments must match the remaining explicit parameters, as described for monomorphic function designators.
- for each assertion parameter in the function designator’s type, there must be an object or function with the same identifier that is visible at the call site and whose type is compatible with or can be specialized to the type of the assertion declaration.

There is a valid interpretation for each such set of implicit parameters. The type of each valid interpretation is the return type of the function designator with implicit parameter values substituted for the implicit arguments.

A valid interpretation is ambiguous if the function designator or any of the argument interpretations is ambiguous.

Every valid interpretation whose return type is not compatible with any other valid interpretation's return type is an interpretation of the function call expression.

Every set of valid interpretations that have mutually compatible result types also produces an interpretation of the function call expression. The type of the interpretation is the composite type of the types of the valid interpretations, and the value of the interpretation is that of the best valid interpretation.

□ One desirable property of a polymorphic programming language is *generalizability*: the ability to replace an abstraction with a more general but equivalent abstraction without requiring changes in any of the uses of the original[3]. For instance, it should be possible to replace a function “`int f(int);`” with “`forall(type T) T f(T);`” without affecting any calls of `f`.

`Cforall` does not fully possess this property, because unsafe conversions are not done when arguments are passed to polymorphic parameters. Consider

```
float  g(float, float);
int    i;
float  f;
double d;
f = g(f, f); /* (1) */
f = g(i, f); /* (2) (safe conversion to float) */
f = g(d, f); /* (3) (unsafe conversion to float) */
```

If `g` was replaced by “`forall(type T) T g(T,T);`”, the first and second calls would be unaffected, but the third would change: `f` would be converted to `double`, and the result would be a `double`.

Another example is the function “`void h(int*);`”. This function can be passed a `void*` argument, but the generalization “`forall(type T) void h(T*);`” can not. In this case, `void` is not a valid value for `T` because it is not an object type. If unsafe conversions were allowed, `T` could be inferred to be *any* object type, which is undesirable. □

Examples

A function called “`?()`” might be part of a numerical differentiation package.

```
extern type      Derivative;
extern double    ?()(Derivative, double);
extern Derivative derivative_of( double (*f)(double) );
extern double    sin(double);

Derivative sin_dx = derivative_of(sin);
double        d;
d = sin_dx(12.9);
```

Here, the only interpretation of `sin_dx` is as an object of type `Derivative`. For that interpretation, the function call is treated as “`?()(sin_dx, 12.9)`”.

```
int f(long);      /* (1) */
int f(int, int); /* (2) */
int f(int*);     /* (3) */
```

```
int i = f(5);      /* calls (1) */
```

Function (1) provides a valid interpretation of “f(5)”, using an implicit `int` to `long` conversion. The other functions do not, since the second requires two arguments, and since there is no implicit conversion from `int` to `int*` that could be used with the third function.

```
forall(type T) T h(T);
double d = h(1.5);
```

“1.5” is a `double` constant, so `T` is inferred to be `double`, and the result of the function call is a `double`.

```
forall(type T, type U) void g(T,U);    /* (4) */
forall(type T)         void g(T,T);    /* (5) */
forall(type T)         void g(T,long);  /* (6) */
void g(long, long);    /* (7) */
double d;
int i;
int* p;

g(d,d);    /* calls (5) */
g(d,i);    /* calls (6) */
g(i,i);    /* calls (7) */
g(i,p);    /* calls (4) */
```

The first call has valid interpretations for all four versions of `g`. (6) and (7) are discarded because they involve unsafe `double`-to-`long` conversions. (5) is chosen because it is less polymorphic than (4).

For the second call, (7) is again discarded. Of the remaining interpretations for (4), (5), and (6) (with `i` converted to `long`), (6) is chosen because it is the least polymorphic.

The third call has valid interpretations for all of the functions; (7) is chosen since it is not polymorphic at all.

The fourth call has no interpretation for (5), because its arguments must have compatible type. (4) is chosen because it does not involve unsafe conversions.

```
forall(type T) T min(T,T);
double         max(double, double);
spec min_max(T) {
  T min(T,T);
  T max(T,T);
}
forall(type U | min_max(U)) void shuffle(U,U);
shuffle(9, 10);
```

The only possibility for `U` is `double`, because that is the type used in the only visible `max` function. 9 and 10 must be converted to `double`, and `min` must be specialized with `T` bound to `double`.

```
extern void q(int);    /* (8) */
extern void q(void*);  /* (9) */
extern void r();
q(0);
r(0);
```

The `int 0` could be passed to (8), or the `(void*)` specialization of the null pointer `0` could be passed to (9). The former is chosen because the `int 0` is less polymorphic. For the same reason, `int 0` is passed to `r()`, even though it has *no* declared parameter types.

6.5.2.3 Structure and union members

Semantics

In the member selection expression “ $s.m$ ”, there shall be at least one interpretation of s whose type is a structure type or union type containing a member named m . If two or more interpretations of s have members named m with mutually compatible types, then the expression has an ambiguous interpretation whose type is the composite type of the types of the members. If an interpretation of s has a member m whose type is not compatible with any other s ’s m , then the expression has an interpretation with the member’s type. The expression has no other interpretations.

The expression “ $p \rightarrow m$ ” has the same interpretations as the expression “ $(*p).m$ ”.

6.5.2.4 Postfix increment and decrement operators

Predefined Identifiers

```

_Bool          ?++(volatile _Bool*),
               ?++(_Atomic volatile _Bool*);
char           ?++(volatile char*),
               ?++(_Atomic volatile char*);
signed char    ?++(volatile signed char*),
               ?++(_Atomic volatile signed char*);
unsigned char  ?++(volatile signed char*),
               ?++(_Atomic volatile signed char*);
_Bool          ?++(volatile _Bool*),
               ?++(_Atomic volatile _Bool*);
short int      ?++(volatile short int*),
               ?++(_Atomic volatile short int*);
unsigned short int ?++(volatile unsigned short int*),
               ?++(_Atomic volatile unsigned short int*);
int            ?++(volatile int*),
               ?++(_Atomic volatile int*);
unsigned int    ?++(volatile unsigned int*),
               ?++(_Atomic volatile unsigned int*);
long int       ?++(volatile long int*),
               ?++(_Atomic volatile long int*);
long unsigned int ?++(volatile long unsigned int*),
               ?++(_Atomic volatile long unsigned int*);
long long int  ?++(volatile long long int*),
               ?++(_Atomic volatile long long int*);
long long unsigned int ?++(volatile long long unsigned int*),
               ?++(_Atomic volatile long long unsigned int*);
float          ?++(volatile float*),
               ?++(_Atomic volatile float*);
double         ?++(volatile double*),
               ?++(_Atomic volatile double*);
long double    ?++(volatile long double*),
               ?++(_Atomic volatile long double*);

forall(type T) T
  * ?++(T* restrict volatile*),
  * ?++(T* _Atomic restrict volatile*);

```

```
forall(type T) _Atomic T
  * ?+(_Atomic T* restrict volatile*),
  * ?+(_Atomic T* _Atomic restrict volatile*);

forall(type T) const T
  * ?+(const T* restrict volatile*),
  * ?+(const T* _Atomic restrict volatile*);

forall(type T) volatile T
  * ?+(volatile T* restrict volatile*),
  * ?+(volatile T* _Atomic restrict volatile*);

forall(type T) restrict T
  * ?+(restrict T* restrict volatile*),
  * ?+(restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic const T
  * ?+(_Atomic const T* restrict volatile*),
  * ?+(_Atomic const T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict T
  * ?+(_Atomic restrict T* restrict volatile*),
  * ?+(_Atomic restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic volatile T
  * ?+(_Atomic volatile T* restrict volatile*),
  * ?+(_Atomic volatile T* _Atomic restrict volatile*);

forall(type T) const restrict T
  * ?+(const restrict T* restrict volatile*),
  * ?+(const restrict T* _Atomic restrict volatile*);

forall(type T) const volatile T
  * ?+(const volatile T* restrict volatile*),
  * ?+(const volatile T* _Atomic restrict volatile*);

forall(type T) restrict volatile T
  * ?+(restrict volatile T* restrict volatile*),
  * ?+(restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict T
  * ?+(_Atomic const restrict T* restrict volatile*),
  * ?+(_Atomic const restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic const volatile T
  * ?+(_Atomic const volatile T* restrict volatile*),
  * ?+(_Atomic const volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict volatile T
```

```

    * ?+((_Atomic restrict volatile T* restrict volatile*),
    * ?+((_Atomic restrict volatile T* _Atomic restrict volatile*));

forall(type T) const restrict volatile T
    * ?+(const restrict volatile T* restrict volatile*),
    * ?+(const restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict volatile T
    * ?+((_Atomic const restrict volatile T* restrict volatile*),
    * ?+((_Atomic const restrict volatile T* _Atomic restrict volatile*);

_Bool          ?--(volatile _Bool*),
               ?--(_Atomic volatile _Bool*);
char           ?--(volatile char*),
               ?--(_Atomic volatile char*);
signed char    ?--(volatile signed char*),
               ?--(_Atomic volatile signed char*);
unsigned char  ?--(volatile signed char*),
               ?--(_Atomic volatile signed char*);
_Bool         ?--(volatile _Bool*),
               ?--(_Atomic volatile _Bool*);
short int     ?--(volatile short int*),
               ?--(_Atomic volatile short int*);
unsigned short int ?--(volatile unsigned short int*),
               ?--(_Atomic volatile unsigned short int*);
int           ?--(volatile int*),
               ?--(_Atomic volatile int*);
unsigned int  ?--(volatile unsigned int*),
               ?--(_Atomic volatile unsigned int*);
long int     ?--(volatile long int*),
               ?--(_Atomic volatile long int*);
long unsigned int ?--(volatile long unsigned int*),
               ?--(_Atomic volatile long unsigned int*);
long long int ?--(volatile long long int*),
               ?--(_Atomic volatile long long int*);
long long unsigned ?--(volatile long long unsigned int*),
               ?--(_Atomic volatile long long unsigned int*);
float        ?--(volatile float*),
               ?--(_Atomic volatile float*);
double       ?--(volatile double*),
               ?--(_Atomic volatile double*);
long double  ?--(volatile long double*),
               ?--(_Atomic volatile long double*);

forall(type T) T
    * ?--(T* restrict volatile*),
    * ?--(T* _Atomic restrict volatile*);

forall(type T) _Atomic T

```

```

    * ?--(_Atomic T* restrict volatile*),
    * ?--(_Atomic T* _Atomic restrict volatile*);

forall(type T) const T
    * ?--(const T* restrict volatile*),
    * ?--(const T* _Atomic restrict volatile*);

forall(type T) volatile T
    * ?--(volatile T* restrict volatile*),
    * ?--(volatile T* _Atomic restrict volatile*);

forall(type T) restrict T
    * ?--(restrict T* restrict volatile*),
    * ?--(restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic const T
    * ?--(_Atomic const T* restrict volatile*),
    * ?--(_Atomic const T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict T
    * ?--(_Atomic restrict T* restrict volatile*),
    * ?--(_Atomic restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic volatile T
    * ?--(_Atomic volatile T* restrict volatile*),
    * ?--(_Atomic volatile T* _Atomic restrict volatile*);

forall(type T) const restrict T
    * ?--(const restrict T* restrict volatile*),
    * ?--(const restrict T* _Atomic restrict volatile*);

forall(type T) const volatile T
    * ?--(const volatile T* restrict volatile*),
    * ?--(const volatile T* _Atomic restrict volatile*);

forall(type T) restrict volatile T
    * ?--(restrict volatile T* restrict volatile*),
    * ?--(restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict T
    * ?--(_Atomic const restrict T* restrict volatile*),
    * ?--(_Atomic const restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic const volatile T
    * ?--(_Atomic const volatile T* restrict volatile*),
    * ?--(_Atomic const volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict volatile T
    * ?--(_Atomic restrict volatile T* restrict volatile*),

```

```

    * ?--(_Atomic restrict volatile T* _Atomic restrict volatile*);

forall(type T) const restrict volatile T
    * ?--(const restrict volatile T* restrict volatile*),
    * ?--(const restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict volatile T
    * ?--(_Atomic const restrict volatile T* restrict volatile*),
    * ?--(_Atomic const restrict volatile T* _Atomic restrict volatile*);

```

For every extended integer type X there exist

```

X ?++(volatile X*), ?++(_Atomic volatile X*),
  ?--(volatile X*), ?--(_Atomic volatile X*);

```

For every complete enumerated type E there exist

```

E ?++(volatile E*), ?++(_Atomic volatile E*),
  ?--(volatile E*), ?--(_Atomic volatile E*);

```

□ Note that “++” and “--” are rewritten as function calls that are given a pointer to that operand. (This is true of all operators that modify an operand.) As Hamish Macdonald has pointed out, this forces the modified operand of such expressions to be an lvalue. This partially enforces the C semantic rule that such operands must be *modifiable* lvalues. □

□ In C, a semantic rule requires that pointer operands of increment and decrement be pointers to object types. Hence, `void*` objects cannot be incremented. In Cforall, the restriction follows from the use of a `type` parameter in the predefined function definitions, as opposed to `dtype`, since only object types can be inferred arguments corresponding to the type parameter `T`. □

Semantics

First, each interpretation of the operand of an increment or decrement expression is considered separately. For each interpretation that is a bit-field or is declared with the `register` storage-class specifier, the expression has one valid interpretation, with the type of the operand, and the expression is ambiguous if the operand is.

For the remaining interpretations, the expression is rewritten, and the interpretations of the expression are the interpretations of the corresponding function call. Finally, all interpretations of the expression produced for the different interpretations of the operand are combined to produce the interpretations of the expression as a whole; where interpretations have compatible result types, the best interpretations are selected in the manner described for function call expressions.

Examples

```

volatile short int vs; vs++; // rewritten as ?++( &(vs) )
short int s;          s++;
const short int cs;   cs++;
_Atomic short int as; as++;

```

Since `&(vs)` has type `volatile short int*`, the best valid interpretation of `vs++` calls the `?++` function with the `volatile short *` parameter. `s++` does the same, applying the safe conversion from `short int*`

to `volatile short int*`. Note that there is no conversion that adds an `_Atomic` qualifier, so the `_Atomic volatile short int` overloading does not provide a valid interpretation.

There is no safe conversion from `const short int*` to `volatile short int*`, and no `?++` function that accepts a `const *` parameter, so `cs++` has no valid interpretations.

The best valid interpretation of `as++` calls the short `?++` function with the `_Atomic volatile short int*` parameter, applying a safe conversion to add the `volatile` qualifier.

```
char * const restrict volatile * restrict volatile pqpc; pqpc++
char * * restrict volatile ppc; ppc++;
```

Since `&(pqpc)` has type `char* const restrict volatile* restrict volatile*`, the best valid interpretation of `pqpc++` calls the polymorphic `?++` function with the `const restrict volatile T* restrict volatile*` parameter, inferring `T` to be `char *`.

`ppc++` calls the same function, again inferring `T` to be `char *`, and using the safe conversions from `T` to `T const restrict volatile`.

□ Increment and decrement expressions show up a deficiency of Cforall's type system. There is no such thing as a pointer to a register object or bit-field. Therefore, there is no way to define a function that alters them, and hence no way to define increment and decrement functions for them. As a result, the semantics of increment and decrement expressions must treat them specially. This holds true for all of the operators that may modify such objects. □

□ The polymorphic overloads for pointer increment and decrement can be understood by considering increasingly complex types.

1. “`char * p; p++;`”. The argument to `?++` has type `char * *`, and the result has type `char *`. The expression would be valid if `?++` were declared by

```
forall(type T) T* ?++(T * *);
```

with `T` inferred to be `char`.

2. “`char *restrict volatile qp; qp++`”. The result again has type `char *`, but the argument now has type `char *restrict volatile *`, so it cannot be passed to the hypothetical function declared in point 1. Hence the actual predefined function is

```
forall(type T) T* ?++(T* restrict volatile*);
```

which also accepts a `char * *` argument, because of the safe conversions that add `volatile` and `restrict` qualifiers. (The parameter is not `const`-qualified, so constant pointers cannot be incremented.)

3. “`char *_Atomic ap; ap++`”. The result again has type `char *`, but no safe conversion adds an `_Atomic` qualifier, so the function in point 2 is not applicable. A separate overloading of `?++` is required.
4. “`char const volatile * pq; pq++`”. Here the result has type `char const volatile *`, so a new overloading is needed:

```
forall(type T)
  T const volatile * ?++(T const volatile *restrict volatile *);
```

One overloading is needed for each combination of qualifiers in the pointed-at type.

5. “`float *restrict * prp; prp++`”. The `restrict` qualifier is handled just like `const` and `volatile` in the previous case:

```
forall(type T)
  T restrict * ?++(T restrict *restrict volatile *);
```

with T inferred to be `float *`. This looks odd, because ISO/IEC C contains a constraint that requires restrict-qualified types to be pointer-to-object types, and T is not syntactically a pointer type. Cforall loosens the constraint.

□

6.5.2.5 Compound literals

Semantics

A compound literal has one interpretation, with the type given by the *type-name* of the compound literal.

6.5.3 Unary operators

Syntax

unary-expression:

- postfix-expression*
- `++ unary-expression`
- `-- unary-expression`
- unary-operator cast-expression*
- `sizeof unary-expression`
- `sizeof (type-name)`

unary-operator: one of

`& * + - ~ !`

Rewrite Rules

```
*a ⇒ *(a)
+a ⇒ +(a)
-a ⇒ -(a)
~a ⇒ ~(a)
!a ⇒ !(a)
++a ⇒ ++(&(a))
--a ⇒ --(&(a))
```

6.5.3.1 Prefix increment and decrement operators

Predefined Identifiers

```
_Bool          ++?(volatile _Bool*),
               ++?(_Atomic volatile _Bool*);
char           ++?(volatile char*),
               ++?(_Atomic volatile char*);
signed char    ++?(volatile signed char*),
               ++?(_Atomic volatile signed char*);
unsigned char  ++?(volatile signed char*),
               ++?(_Atomic volatile signed char*);
_Bool         ++?(volatile _Bool*),
               ++?(_Atomic volatile _Bool*);
short int     ++?(volatile short int*),
               ++?(_Atomic volatile short int*);
```

```

unsigned short int ++?(volatile unsigned short int*),
                    ++?(_Atomic volatile unsigned short int*);
int                ++?(volatile int*),
                    ++?(_Atomic volatile int*);
unsigned int       ++?(volatile unsigned int*),
                    ++?(_Atomic volatile unsigned int*);
long int          ++?(volatile long int*),
                    ++?(_Atomic volatile long int*);
long unsigned int ++?(volatile long unsigned int*),
                    ++?(_Atomic volatile long unsigned int*);
long long int     ++?(volatile long long int*),
                    ++?(_Atomic volatile long long int*);
long long unsigned ++?(volatile long long unsigned int*),
                    ++?(_Atomic volatile long long unsigned int*);
float             ++?(volatile float*),
                    ++?(_Atomic volatile float*);
double            ++?(volatile double*),
                    ++?(_Atomic volatile double*);
long double       ++?(volatile long double*),
                    ++?(_Atomic volatile long double*);

forall(type T) T
    * ++?(T* restrict volatile*),
    * ++?(T* _Atomic restrict volatile*);

forall(type T) _Atomic T
    * ++?(_Atomic T* restrict volatile*),
    * ++?(_Atomic T* _Atomic restrict volatile*);

forall(type T) const T
    * ++?(const T* restrict volatile*),
    * ++?(const T* _Atomic restrict volatile*);

forall(type T) volatile T
    * ++?(volatile T* restrict volatile*),
    * ++?(volatile T* _Atomic restrict volatile*);

forall(type T) restrict T
    * ++?(restrict T* restrict volatile*),
    * ++?(restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic const T
    * ++?(_Atomic const T* restrict volatile*),
    * ++?(_Atomic const T* _Atomic restrict volatile*);

forall(type T) _Atomic volatile T
    * ++?(_Atomic volatile T* restrict volatile*),
    * ++?(_Atomic volatile T* _Atomic restrict volatile*);

```

```

forall(type T) _Atomic restrict T
    * ++?(_Atomic restrict T* restrict volatile*),
    * ++?(_Atomic restrict T* _Atomic restrict volatile*);

forall(type T) const volatile T
    * ++?(const volatile T* restrict volatile*),
    * ++?(const volatile T* _Atomic restrict volatile*);

forall(type T) const restrict T
    * ++?(const restrict T* restrict volatile*),
    * ++?(const restrict T* _Atomic restrict volatile*);

forall(type T) restrict volatile T
    * ++?(restrict volatile T* restrict volatile*),
    * ++?(restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const volatile T
    * ++?(_Atomic const volatile T* restrict volatile*),
    * ++?(_Atomic const volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict T
    * ++?(_Atomic const restrict T* restrict volatile*),
    * ++?(_Atomic const restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict volatile T
    * ++?(_Atomic restrict volatile T* restrict volatile*),
    * ++?(_Atomic restrict volatile T* _Atomic restrict volatile*);

forall(type T) const restrict volatile T
    * ++?(const restrict volatile T* restrict volatile*),
    * ++?(const restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict volatile T
    * ++?(_Atomic const restrict volatile T* restrict volatile*),
    * ++?(_Atomic const restrict volatile T* _Atomic restrict volatile*);

_Bool            --?(volatile _Bool*),
                 --?(_Atomic volatile _Bool*);
char             --?(volatile char*),
                 --?(_Atomic volatile char*);
signed char      --?(volatile signed char*),
                 --?(_Atomic volatile signed char*);
unsigned char    --?(volatile signed char*),
                 --?(_Atomic volatile signed char*);
_Bool            --?(volatile _Bool*),
                 --?(_Atomic volatile _Bool*);
short int        --?(volatile short int*),
                 --?(_Atomic volatile short int*);
unsigned short int --?(volatile unsigned short int*),

```

```

int                --?(_Atomic volatile unsigned short int*);
                  --?(volatile int*),
                  --?(_Atomic volatile int*);
unsigned int       --?(volatile unsigned int*),
                  --?(_Atomic volatile unsigned int*);
long int           --?(volatile long int*),
                  --?(_Atomic volatile long int*);
long unsigned int  --?(volatile long unsigned int*),
                  --?(_Atomic volatile long unsigned int*);
long long int      --?(volatile long long int*),
                  --?(_Atomic volatile long long int*);
long long unsigned --?(volatile long long unsigned int*),
                  --?(_Atomic volatile long long unsigned int*);
float              --?(volatile float*),
                  --?(_Atomic volatile float*);
double             --?(volatile double*),
                  --?(_Atomic volatile double*);
long double        --?(volatile long double*),
                  --?(_Atomic volatile long double*);

forall(type T) T
  * --?(T* restrict volatile*),
  * --?(T* _Atomic restrict volatile*);

forall(type T) _Atomic T
  * --?(_Atomic T* restrict volatile*),
  * --?(_Atomic T* _Atomic restrict volatile*);

forall(type T) const T
  * --?(const T* restrict volatile*),
  * --?(const T* _Atomic restrict volatile*);

forall(type T) volatile T
  * --?(volatile T* restrict volatile*),
  * --?(volatile T* _Atomic restrict volatile*);

forall(type T) restrict T
  * --?(restrict T* restrict volatile*),
  * --?(restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic const T
  * --?(_Atomic const T* restrict volatile*),
  * --?(_Atomic const T* _Atomic restrict volatile*);

forall(type T) _Atomic volatile T
  * --?(_Atomic volatile T* restrict volatile*),
  * --?(_Atomic volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict T

```

```

* --?(_Atomic restrict T* restrict volatile*),
* --?(_Atomic restrict T* _Atomic restrict volatile*);

forall(type T) const volatile T
* --?(const volatile T* restrict volatile*),
* --?(const volatile T* _Atomic restrict volatile*);

forall(type T) const restrict T
* --?(const restrict T* restrict volatile*),
* --?(const restrict T* _Atomic restrict volatile*);

forall(type T) restrict volatile T
* --?(restrict volatile T* restrict volatile*),
* --?(restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const volatile T
* --?(_Atomic const volatile T* restrict volatile*),
* --?(_Atomic const volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict T
* --?(_Atomic const restrict T* restrict volatile*),
* --?(_Atomic const restrict T* _Atomic restrict volatile*);

forall(type T) _Atomic restrict volatile T
* --?(_Atomic restrict volatile T* restrict volatile*),
* --?(_Atomic restrict volatile T* _Atomic restrict volatile*);

forall(type T) const restrict volatile T
* --?(const restrict volatile T* restrict volatile*),
* --?(const restrict volatile T* _Atomic restrict volatile*);

forall(type T) _Atomic const restrict volatile T
* --?(_Atomic const restrict volatile T* restrict volatile*),
* --?(_Atomic const restrict volatile T* _Atomic restrict volatile*);

```

For every extended integer type X there exist

```

X  ++?(volatile X*),
   ++?(_Atomic volatile X*),
   --?(volatile X*),
   --?(_Atomic volatile X*);

```

For every complete enumerated type E there exist

```

E  ++?(volatile E*),
   ++?(_Atomic volatile E*),
   ?--(volatile E*),
   ?--(_Atomic volatile E*);

```

Semantics

The interpretations of prefix increment and decrement expressions are determined in the same way as the interpretations of postfix increment and decrement expressions.

6.5.3.2 Address and indirection operators

Predefined Identifiers

```

forall(type T) lvalue T          *?(T*);
forall(type T) _Atomic lvalue T  *?(_Atomic T*);
forall(type T) const lvalue T    *?(const T*);
forall(type T) volatile lvalue T *?(volatile T*);
forall(type T) restrict lvalue T *?(restrict T*);
forall(type T) _Atomic const lvalue T *?(_Atomic const T*);
forall(type T) _Atomic volatile lvalue T *?(_Atomic volatile T*);
forall(type T) _Atomic restrict lvalue T *?(_Atomic restrict T*);
forall(type T) const volatile lvalue T *?(const volatile T*);
forall(type T) const restrict lvalue T *?(const restrict T*);
forall(type T) restrict volatile lvalue T *?(restrict volatile T*);
forall(type T) _Atomic const volatile lvalue T
    *?(_Atomic const volatile T*);
forall(type T) _Atomic const restrict lvalue T
    *?(_Atomic const restrict T*);
forall(type T) _Atomic restrict volatile lvalue T
    *?(_Atomic restrict volatile T*);
forall(type T) const restrict volatile lvalue T
    *?(const restrict volatile T*);
forall(type T) _Atomic const restrict volatile lvalue T
    *?(_Atomic const restrict volatile T*);

forall(ftype FT) FT *?(FT*);

```

Semantics

When the “&” operator is applied to an expression, the operand shall have exactly one unambiguous interpretation. The “&” expression has one interpretation which is of type T^* , where T is the type of the operand.

The interpretations of an indirection expression are the interpretations of the corresponding function call.

6.5.3.3 Unary arithmetic operators

Predefined Identifiers

```

int
    +?(int),
    -?(int),
    ~?(int);
unsigned int
    +?(unsigned int),
    -?(unsigned int),
    ~?(unsigned int);
long int
    +?(long int),
    -?(long int),
    ~?(long int);
long unsigned int

```

```

    +?(long unsigned int),
    -?(long unsigned int),
    ~?(long unsigned int);
long long int
    +?(long long int),
    -?(long long int),
    ~?(long long int);
long long unsigned int
    +?(long long unsigned int),
    -?(long long unsigned int),
    ~?(long long unsigned int);
float
    +?(float),
    -?(float);
double
    +?(double),
    -?(double);
long double
    +?(long double),
    -?(long double);
_Complex float
    +?(_Complex float),
    -?(_Complex float);
_Complex double
    +?(_Complex double),
    -?(_Complex double);
_Complex long double
    +?(_Complex long double),
    -?(_Complex long double);

int !(int),
    !(unsigned int),
    !(long),
    !(long unsigned int),
    !(long long int),
    !(long long unsigned int),
    !(float),
    !(double),
    !(long double),
    !?(_Complex float),
    !?(_Complex double),
    !?(_Complex long double);

forall(dtype DT) int !(const restrict volatile DT*);
forall(dtype DT) int !?(_Atomic const restrict volatile DT*);
forall(ftype FT) int !?(FT*);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```

X +?(X), -?(X), ~?(X);
int !?(X);

```

Semantics

The interpretations of a unary arithmetic expression are the interpretations of the corresponding function call.

Examples

```
long int li;
void eat_double(double);

eat_double(-li); /* ⇒ eat_double( -(li) ); */
```

The valid interpretations of “-li” (assuming no extended integer types exist) are

interpretation	result type	expression conversion cost
??((int)li)	int	(unsafe)
??((unsigned)li)	unsigned int	(unsafe)
??((long)li)	long	0
??((long unsigned int)li)	long unsigned int	1
??((long long int)li)	long long int	2
??((long long unsigned int)li)	long long unsigned int	3
??((float)li)	float	4
??((double)li)	double	5
??((long double)li)	long double	6
??((_Complex float)li)	float	(unsafe)
??((_Complex double)li)	double	(unsafe)
??((_Complex long double)li)	long double	(unsafe)

The valid interpretations of the eat_double call, with the cost of the argument conversion and the cost of the entire expression, are

interpretation	argument cost	expression cost
eat_double((double)-??((int)li))	7	(unsafe)
eat_double((double)-??((unsigned)li))	6	(unsafe)
eat_double((double)-?(li))	5	0 + 5 = 5
eat_double((double)-??((long unsigned int)li))	4	1 + 4 = 5
eat_double((double)-??((long long int)li))	3	2 + 3 = 5
eat_double((double)-??((long long unsigned int)li))	2	3 + 2 = 5
eat_double((double)-??((float)li))	1	4 + 1 = 5
eat_double((double)-??((double)li))	0	5 + 0 = 5
eat_double((double)-??((long double)li))	(unsafe)	(unsafe)
eat_double((double)-??((_Complex float)li))	(unsafe)	(unsafe)
eat_double((double)-??((_Complex double)li))	(unsafe)	(unsafe)
eat_double((double)-??((_Complex long double)li))	(unsafe)	(unsafe)

Each has result type void, so the best must be selected. The interpretations involving unsafe conversions are discarded. The remainder have equal expression conversion costs, so the “highest argument conversion cost” rule is invoked, and the chosen interpretation is eat_double((double)-?(li)).

6.5.3.4 The sizeof and _Alignof operators**Constraints**

The operand of sizeof shall not be type, dtype, or ftype.

Semantics

When the `sizeof` and `_Alignof` operators are applied to an expression, the expression shall have exactly one unambiguous interpretation. The expression has one interpretation, which has an integer type.

When `sizeof` is applied to an identifier declared by a *type-declaration* or a *type-parameter*, it yields the size in bytes of the type that implements the operand. When the operand is an opaque type or an inferred type parameter, the expression is not a constant expression.

When `_Alignof` is applied to an identifier declared by a *type-declaration* or a *type-parameter*, it yields the alignment requirement of the type that implements the operand. When the operand is an opaque type or an inferred type parameter, the expression is not a constant expression.

□

```

type   Pair = struct { int first, second; };
size_t p_size = sizeof(Pair);    /* constant expression */

extern type Rational;
size_t c_size = sizeof(Rational); /* non-constant expression */

forall(type T) T f(T p1, T p2) {
    size_t t_size = sizeof(T);    /* non-constant expression */
    /* ... */
}

```

“`sizeof Rational`”, although not statically known, is fixed. Within `f()`, “`sizeof(T)`” is fixed for each call of `f()`, but may vary from call to call. □

6.5.4 Cast operators**Syntax**

cast-expression:

```

    unary-expression
    ( type-name ) cast-expression

```

Constraints

The *type-name* in a *cast-expression* shall not be `type`, `dtype`, or `ftype`.

Semantics

In a cast expression “*(type-name) e*”, if *type-name* is the type of an interpretation of *e*, then that interpretation is the only interpretation of the cast expression; otherwise, *e* shall have some interpretation that can be converted to *type-name*, and the interpretation of the cast expression is the cast of the interpretation that can be converted at the lowest cost. The cast expression’s interpretation is ambiguous if more than one interpretation can be converted at the lowest cost or if the selected interpretation is ambiguous.

□ Casts can be used to eliminate ambiguity in expressions by selecting interpretations of subexpressions, and to specialize polymorphic functions and values. □

6.5.5 Multiplicative operators**Syntax**

multiplicative-expression:
cast-expression
*multiplicative-expression * cast-expression*
multiplicative-expression / cast-expression
multiplicative-expression % cast-expression

Rewrite Rules

$a * b \Rightarrow ??(a,b)$
 $a / b \Rightarrow ??(a,b)$
 $a \% b \Rightarrow ??(a,b)$

Predefined Identifiers

int	??(int, int), ??(int, int), ??(int, int);
unsigned int	??(unsigned int, unsigned int), ??(unsigned int, unsigned int), ??(unsigned int, unsigned int);
long int	??(long int, long int), ??(long, long), ??(long, long);
long unsigned int	??(long unsigned int, long unsigned int), ??(long unsigned int, long unsigned int), ??(long unsigned int, long unsigned int);
long long int	??(long long int, long long int), ??(long long int, long long int), ??(long long int, long long int);
long long unsigned int	??(long long unsigned int, long long unsigned int), ??(long long unsigned int, long long unsigned int), ??(long long unsigned int, long long unsigned int);
float	??(float, float), ??(float, float);
double	??(double, double), ??(double, double);
long double	??(long double, long double), ??(long double, long double);
_Complex float	??(float, _Complex float), ??(float, _Complex float), ??(_Complex float, float), ??(_Complex float, float), ??(_Complex float, _Complex float), ??(_Complex float, _Complex float);
_Complex double	??(double, _Complex double), ??(double, _Complex double), ??(_Complex double, double), ??(_Complex double, double), ??(_Complex double, _Complex double),

```

        ?/?(_Complex double, _Complex double);
_Complex long double  ?*(long double, _Complex long double),
        ?/(long double, _Complex long double),
        ?*( _Complex long double, long double),
        ?/( _Complex long double, long double),
        ?*( _Complex long double, _Complex long double),
        ?/( _Complex long double, _Complex long double);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```
X  ?*(X),  ?/(X),  ?/(X);
```

□ ISO/IEC C does not include conversions from the real types to complex types in the usual arithmetic conversions. Instead it specifies conversion of the result of binary operations on arguments from mixed type domains. `Cforall`'s predefined operators match that pattern. □

Semantics

The interpretations of multiplicative expressions are the interpretations of the corresponding function call.

Examples

```

int  i;
long li;
void eat_double(double);
eat_double( li % i );

```

“`li % i`” is rewritten as “`?(li,i)`”. The valid interpretations of `?(li, i)`, the cost of converting their arguments, and the cost of converting the result to `double` (assuming no extended integer types are present) are

interpretation	argument cost	result cost
<code>?((int)li, i)</code>	(unsafe)	6
<code>?((unsigned)li, (unsigned)i)</code>	(unsafe)	5
<code>?(li, (long)i)</code>	1	4
<code>?((long unsigned)li, (long unsigned)i)</code>	3	3
<code>?((long long)li, (long long)i)</code>	5	2
<code>?((long long unsigned)li, (long long unsigned)i)</code>	7	1

The best interpretation of `eat_double(li, i)` is `eat_double((double)?(li, (long)i))`, which has no unsafe conversions and the lowest total cost.

□ ISO/IEC C defines most arithmetic operations to apply an integer promotion to any argument that belongs to a type that has an integer conversion rank less than that of `int`. If `s` is a `short int`, “`s*s`” does not have type `short int`; it is treated as “`((int)s)*((int)s)`”, and has type `int`. `Cforall` matches that pattern; it does not predefine “`short ?*(short,short)`”.

These “missing” operators limit polymorphism. Consider

```

forall(type T | T ?*(T,T) ) T square(T);
short s;
square(s);

```

Since `Cforall` does not define a multiplication operator for `short int`, `square(s)` is treated as `square((int)s)`, and the result has type `int`. This is mildly surprising, but it follows the ISO/IEC C operator pattern.

A more troubling example is

```
forall(type T | ??(T,T) ) T product(T[], int n);
short sa[5];
product(sa, 5);
```

This has no valid interpretations, because `Cforall` has no conversion from “array of `short int`” to “array of `int`”. The alternatives in such situations include

- Defining monomorphic overloads of `product` for `short` and the other “small” types.
- Defining “`short ??(short,short)`” within the scope containing the call to `product`.
- Defining `product` to take as an argument a conversion function from the “small” type to the operator’s argument type.

□

6.5.6 Additive operators

Syntax

additive-expression:
multiplicative-expression
additive-expression + multiplicative-expression
additive-expression - multiplicative-expression

Rewrite Rules

$a + b \Rightarrow ??(a,b)$
 $a - b \Rightarrow ?-(a,b)$

Predefined Identifiers

<code>int</code>	<code>??(int, int),</code> <code>?-(int, int);</code>
<code>unsigned int</code>	<code>??(unsigned int, unsigned int),</code> <code>?-(unsigned int, unsigned int);</code>
<code>long int</code>	<code>??(long int, long int),</code> <code>?-(long int, long int);</code>
<code>long unsigned int</code>	<code>??(long unsigned int, long unsigned int),</code> <code>?-(long unsigned int, long unsigned int);</code>
<code>long long int</code>	<code>??(long long int, long long int),</code> <code>?-(long long int, long long int);</code>
<code>long long unsigned int</code>	<code>??(long long unsigned int, long long unsigned int),</code> <code>?-(long long unsigned int, long long unsigned int);</code>
<code>float</code>	<code>??(float, float),</code> <code>?-(float, float);</code>
<code>double</code>	<code>??(double, double),</code> <code>?-(double, double);</code>

```

long double          ?+?(long double, long double),
                    ?-?(long double, long double);
_Complex float      ?+?(_Complex float, float),
                    ?-?(_Complex float, float),
                    ?+?(float, _Complex float),
                    ?-?(float, _Complex float),
                    ?+?(_Complex float, _Complex float),
                    ?-?(_Complex float, _Complex float);
_Complex double     ?+?(_Complex double, double),
                    ?-?(_Complex double, double),
                    ?+?(double, _Complex double),
                    ?-?(double, _Complex double),
                    ?+?(_Complex double, _Complex double),
                    ?-?(_Complex double, _Complex double);
_Complex long double ?+?(_Complex long double, long double),
                    ?-?(_Complex long double, long double),
                    ?+?(long double, _Complex long double),
                    ?-?(long double, _Complex long double),
                    ?+?(_Complex long double, _Complex long double),
                    ?-?(_Complex long double, _Complex long double);

forall(type T) T
    * ?+?(T*,      ptrdiff_t),
    * ?+?(ptrdiff_t, T*),
    * ?-?(T*,      ptrdiff_t);

forall(type T) _Atomic T
    * ?+?(_Atomic T*, ptrdiff_t),
    * ?+?(ptrdiff_t, _Atomic T*),
    * ?-?(_Atomic T*, ptrdiff_t);

forall(type T) const T
    * ?+?(const T*, ptrdiff_t),
    * ?+?(ptrdiff_t, const T*),
    * ?-?(const T*, ptrdiff_t);

forall(type T) restrict T
    * ?+?(restrict T*, ptrdiff_t),
    * ?+?(ptrdiff_t, restrict T*),
    * ?-?(restrict T*, ptrdiff_t);

forall(type T) volatile T
    * ?+?(volatile T*, ptrdiff_t),
    * ?+?(ptrdiff_t, volatile T*),
    * ?-?(volatile T*, ptrdiff_t);

forall(type T) _Atomic const T
    * ?+?(_Atomic const T*, ptrdiff_t),
    * ?+?(ptrdiff_t, _Atomic const T*),

```

```

    * ?-?(_Atomic const T*, ptrdiff_t);

forall(type T) _Atomic restrict T
    * ?+?(_Atomic restrict T*, ptrdiff_t),
    * ?+?(ptrdiff_t, _Atomic restrict T*),
    * ?-?(_Atomic restrict T*, ptrdiff_t);

forall(type T) _Atomic volatile T
    * ?+?(_Atomic volatile T*, ptrdiff_t),
    * ?+?(ptrdiff_t, _Atomic volatile T*),
    * ?-?(_Atomic volatile T*, ptrdiff_t);

forall(type T) const restrict T
    * ?+?(const restrict T*, ptrdiff_t),
    * ?+?(ptrdiff_t, const restrict T*),
    * ?-?(const restrict T*, ptrdiff_t);

forall(type T) const volatile T
    * ?+?(const volatile T*, ptrdiff_t),
    * ?+?(ptrdiff_t, const volatile T*),
    * ?-?(const volatile T*, ptrdiff_t);

forall(type T) restrict volatile T
    * ?+?(restrict volatile T*, ptrdiff_t),
    * ?+?(ptrdiff_t, restrict volatile T*),
    * ?-?(restrict volatile T*, ptrdiff_t);

forall(type T) _Atomic const restrict T
    * ?+?(_Atomic const restrict T*, ptrdiff_t),
    * ?+?(ptrdiff_t, _Atomic const restrict T*),
    * ?-?(_Atomic const restrict T*, ptrdiff_t);

forall(type T) ptrdiff_t
    * ?-?(const restrict volatile T*, const restrict volatile T*),
    * ?-?(_Atomic const restrict volatile T*, _Atomic const restrict volatile T*);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

X `?+?(X)`, `?-?(X)`;

Semantics

The interpretations of additive expressions are the interpretations of the corresponding function calls.

□ `ptrdiff_t` is an implementation-defined identifier defined in `<stddef.h>` that is synonymous with a signed integral type that is large enough to hold the difference between two pointers. It seems reasonable to use it for pointer addition as well. (This is technically a difference between Cforall and C, which only specifies that pointer addition uses an *integral* argument.) Hence it is also used for subscripting, which is defined in terms of pointer addition. The ISO/IEC C standard uses `size_t` in several cases where a library function takes an argument that is used as a subscript, but `size_t` is unsuitable here because it is an unsigned type. □

6.5.7 Bitwise shift operators

Syntax

shift-expression:
additive-expression
shift-expression << *additive-expression*
shift-expression >> *additive-expression*

Rewrite Rules

$a \ll b \Rightarrow ?\ll?(a,b)$
 $a \gg b \Rightarrow ?\gg?(a,b)$

Predefined Identifiers

<code>int</code>	<code>?<<?(int, int),</code> <code>?>>?(int, int);</code>
<code>unsigned int</code>	<code>?<<?(unsigned int, int),</code> <code>?>>?(unsigned int, int);</code>
<code>long int</code>	<code>?<<?(long int, int),</code> <code>?>>?(long int, int);</code>
<code>long unsigned int</code>	<code>?<<?(long unsigned int, int),</code> <code>?>>?(long unsigned int, int);</code>
<code>long long int</code>	<code>?<<?(long long int, int),</code> <code>?>>?(long long int, int);</code>
<code>long long unsigned int</code>	<code>?<<?(long long unsigned int, int),</code> <code>?>>?(long long unsigned int, int);</code>

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

$X \text{ ?<<?}(X, \text{int}), \text{ ?>>?}(X, \text{int});$

- The bitwise shift operators break the usual pattern: they do not convert both operands to a common type. The right operand only undergoes integer promotion. □

Semantics

The interpretations of a bitwise shift expression are the interpretations of the corresponding function calls.

6.5.8 Relational operators

Syntax

relational-expression:
shift-expression
relational-expression < *shift-expression*
relational-expression > *shift-expression*
relational-expression <= *shift-expression*
relational-expression >= *shift-expression*

Rewrite Rules

```

a < b  ⇒ ?<?(a,b)
a > b  ⇒ ?>?(a,b)
a <= b ⇒ ?<=?(a,b)
a >= b ⇒ ?>=?(a,b)

```

Predefined Identifiers

```

int ?<?(int, int),
    ?<=?(int, int),
    ?>?(int, int),
    ?>=?(int, int);
int ?<?(unsigned int, unsigned int),
    ?<=?(unsigned int, unsigned int),
    ?>?(unsigned int, unsigned int),
    ?>=?(unsigned int, unsigned int);
int ?<?(long int, long int),
    ?<=?(long int, long int),
    ?>?(long int, long int),
    ?>=?(long int, long int);
int ?<?(long unsigned int, long unsigned),
    ?<=?(long unsigned int, long unsigned),
    ?>?(long unsigned int, long unsigned),
    ?>=?(long unsigned int, long unsigned);
int ?<?(long long int, long long int),
    ?<=?(long long int, long long int),
    ?>?(long long int, long long int),
    ?>=?(long long int, long long int);
int ?<?(long long unsigned int, long long unsigned),
    ?<=?(long long unsigned int, long long unsigned),
    ?>?(long long unsigned int, long long unsigned),
    ?>=?(long long unsigned int, long long unsigned);
int ?<?(float, float),
    ?<=?(float, float),
    ?>?(float, float),
    ?>=?(float, float);
int ?<?(double, double),
    ?<=?(double, double),
    ?>?(double, double),
    ?>=?(double, double);
int ?<?(long double, long double),
    ?<=?(long double, long double),
    ?>?(long double, long double),
    ?>=?(long double, long double);

forall(dtype DT) int
    ?<?(const restrict volatile DT*, const restrict volatile DT*),
    ?<?(_Atomic const restrict volatile DT*, _Atomic const restrict volatile DT*),

```

```

?<=? (const restrict volatile DT*, const restrict volatile DT*),
?<=? (_Atomic const restrict volatile DT*, _Atomic const restrict volatile DT*),
?>?(const restrict volatile DT*, const restrict volatile DT*),
?>?(_Atomic const restrict volatile DT*, _Atomic const restrict volatile DT*),
?>=? (const restrict volatile DT*, const restrict volatile DT*),
?>=? (_Atomic const restrict volatile DT*, _Atomic const restrict volatile DT*);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```

int ?<?(X, X),
?<=?(X, X),
?<?(X, X),
?>=?(X, X);

```

Semantics

The interpretations of a relational expression are the interpretations of the corresponding function call.

6.5.9 Equality operators

Syntax

equality-expression:
relational-expression
equality-expression == *relational-expression*
equality-expression != *relational-expression*

Rewrite Rules

```

a == b ⇒ ?==?(a,b)
a != b ⇒ ?!=?(a,b)

```

Predefined Identifiers

```

int ?==?(int, int),
?!=?(int, int),
?==?(unsigned int, unsigned int),
?!=?(unsigned int, unsigned int),
?==?(long int, long int),
?!=?(long int, long int),
?==?(long unsigned int, long unsigned int),
?!=?(long unsigned int, long unsigned int),
?==?(long long int, long long int),
?!=?(long long int, long long int),
?==?(long long unsigned int, long long unsigned int),
?!=?(long long unsigned int, long long unsigned int),
?==?(float, float),
?!=?(float, float),
?==?(_Complex float, float),
?!=?(_Complex float, float),
?==?(float, _Complex float),

```

```

?!=?(float, _Complex float),
?==?(_Complex float, _Complex float),
?!=?(_Complex float, _Complex float),
?==?(double, double),
?!=?(double, double),
?==?(_Complex double, double),
?!=?(_Complex double, double),
?==?(double, _Complex double),
?!=?(double, _Complex double),
?==?(_Complex double, _Complex double),
?!=?(_Complex double, _Complex double),
?==?(long double, long double),
?!=?(long double, long double),
?==?(_Complex long double, long double),
?!=?(_Complex long double, long double),
?==?(long double, _Complex long double),
?!=?(long double, _Complex long double),
?==?(_Complex long double, _Complex long double),
?!=?(_Complex long double, _Complex long double);

forall(dtype DT) int
    ?==?(const restrict volatile DT*,    const restrict volatile DT*),
    ?!=?(const restrict volatile DT*,    const restrict volatile DT*),
    ?==?(const restrict volatile DT*,    const restrict volatile void*),
    ?!=?(const restrict volatile DT*,    const restrict volatile void*),
    ?==?(const restrict volatile void*,  const restrict volatile DT*),
    ?!=?(const restrict volatile void*,  const restrict volatile DT*),
    ?==?(const restrict volatile DT*,    forall(dtype DT2) const DT2*),
    ?!=?(const restrict volatile DT*,    forall(dtype DT2) const DT2*),
    ?==?(forall(dtype DT2) const DT2*,   const restrict volatile DT*),
    ?!=?(forall(dtype DT2) const DT2*,   const restrict volatile DT*),
    ?==?(forall(dtype DT2) const DT2*,   forall(dtype DT3) const DT3*),
    ?!=?(forall(dtype DT2) const DT2*,   forall(dtype DT3) const DT3*),

    ?==?(_Atomic const restrict volatile DT*, _Atomic const restrict volatile DT*),
    ?!=?(_Atomic const restrict volatile DT*, _Atomic const restrict volatile DT*),
    ?==?(_Atomic const restrict volatile DT*, const restrict volatile void*),
    ?!=?(_Atomic const restrict volatile DT*, const restrict volatile void*),
    ?==?(const restrict volatile void*,      _Atomic const restrict volatile DT*),
    ?!=?(const restrict volatile void*,      _Atomic const restrict volatile DT*),
    ?==?(_Atomic const restrict volatile DT*, forall(dtype DT2) const DT2*),
    ?!=?(_Atomic const restrict volatile DT*, forall(dtype DT2) const DT2*),
    ?==?(forall(dtype DT2) const DT2*,      _Atomic const restrict volatile DT*),
    ?!=?(forall(dtype DT2) const DT2*,      _Atomic const restrict volatile DT*);

forall(ftype FT) int
    ?==?(FT*, FT*),
    ?!=?(FT*, FT*),
    ?==?(FT*, forall(ftype FT2) FT2*),

```

```

?!=?(FT*, forall(ftype FT2) FT2*),
?==?(forall(ftype FT2) FT2*, FT*),
?!=?(forall(ftype FT2) FT2*, FT*),
?==?(forall(ftype FT2) FT2*, forall(ftype FT3) FT3*),
?!=?(forall(ftype FT2) FT2*, forall(ftype FT3) FT3*);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```

int ?==?(X, X),
    ?!=?(X, X);

```

□ The polymorphic equality operations come in three styles: comparisons between pointers of compatible types, between pointers to `void` and pointers to object types or incomplete types, and between the null pointer constant and pointers to any type. In the last case, a special constraint rule for null pointer constant operands has been replaced by a consequence of the `Cforall` type system. □

Semantics

The interpretations of an equality expression are the interpretations of the corresponding function call.

The result of an equality comparison between two pointers to predefined functions or predefined values is implementation-defined.

□ The implementation-defined status of equality comparisons allows implementations to use one library routine to implement many predefined functions. These optimization are particularly important when the predefined functions are polymorphic, as is the case for most pointer operations. □

6.5.10 Bitwise AND operator

Syntax

AND-expression:
equality-expression
AND-expression & equality-expression

Rewrite Rules

$a \& b \Rightarrow \text{?&?}(a,b)$

Predefined Identifiers

```

int                ?&?(int, int);
unsigned int       ?&?(unsigned int, unsigned int);
long int           ?&?(long int, long int);
long unsigned int  ?&?(long unsigned int, long unsigned int);
long long int      ?&?(long long int, long long int);
long long unsigned int ?&?(long long unsigned int, long long unsigned int);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```

int ?&?(X, X);

```

Semantics

The interpretations of a bitwise AND expression are the interpretations of the corresponding function call.

6.5.11 Bitwise exclusive OR operator**Syntax**

exclusive-OR-expression:
AND-expression
exclusive-OR-expression \wedge *AND-expression*

Rewrite Rules

$$a \wedge b \Rightarrow \text{??}(a, b)$$
Predefined Identifiers

```
int                ??(int, int);
unsigned int       ??(unsigned int, unsigned int);
long int           ??(long int, long int);
long unsigned int  ??(long unsigned int, long unsigned int);
long long int      ??(long long int, long long int);
long long unsigned int ??(long long unsigned int, long long unsigned int);
```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```
int ??(X, X);
```

Semantics

The interpretations of a bitwise exclusive OR expression are the interpretations of the corresponding function call.

6.5.12 Bitwise inclusive OR operator**Syntax**

inclusive-OR-expression:
exclusive-OR-expression
inclusive-OR-expression $|$ *exclusive-OR-expression*

Rewrite Rules

$$a | b \Rightarrow \text{?}|?(a, b)$$

Predefined Identifiers

```

int                ?|?(int, int);
unsigned int       ?|?(unsigned int, unsigned int);
long int           ?|?(long int, long int);
long unsigned int  ?|?(long unsigned int, long unsigned int);
long long int      ?|?(long long int, long long int);
long long unsigned int ?|?(long long unsigned int, long long unsigned int);

```

For every extended integer type X with integer conversion rank greater than the rank of `int` there exist

```
int ?|?(X, X);
```

Semantics

The interpretations of a bitwise inclusive OR expression are the interpretations of the corresponding function call.

6.5.13 Logical AND operator**Syntax**

logical-AND-expression:
inclusive-OR-expression
logical-AND-expression **&&** *inclusive-OR-expression*

Semantics

The operands of the expression “ $a \ \&\& \ b$ ” are treated as “`(int)((a)!=0)`” and “`(int)((b)!=0)`”, which shall both be unambiguous. The expression has only one interpretation, which is of type `int`.

□ When the operands of a logical expression are values of built-in types, and “`!=`” has not been redefined for those types, the compiler can optimize away the function calls.

A common C idiom omits comparisons to 0 in the controlling expressions of loops and `if` statements. For instance, the loop below iterates as long as `rp` points at a `Rational` value that is non-zero.

```

extern type          Rational;
extern const Rational 0;
extern int           ?!=?(Rational, Rational);
Rational             *rp;

while (rp && *rp) { /* ... */ }

```

The logical expression calls the `Rational` inequality operator, passing it `*rp` and the `Rational 0`, and getting a 1 or 0 as a result. In contrast, C++ would apply a programmer-defined `Rational-to-int` conversion to `*rp` in the equivalent situation. The conversion to `int` would produce a general integer value, which is unfortunate, and possibly dangerous if the conversion was not written with this situation in mind. □

6.5.14 Logical OR operator**Syntax**

logical-OR-expression:
logical-AND-expression
logical-OR-expression **||** *logical-AND-expression*

Semantics

The operands of the expression “ $a \ || \ b$ ” are treated as “ $(\text{int})((a) \neq 0)$ ” and “ $(\text{int})((b) \neq 0)$ ”, which shall both be unambiguous. The expression has only one interpretation, which is of type `int`.

6.5.15 Conditional operator**Syntax**

conditional-expression:
logical-OR-expression
logical-OR-expression ? expression : conditional-expression

Semantics

In the conditional expression “ $a?b:c$ ”, if the second and third operands have void type, then the expression has one interpretation, with type `void`, and is interpreted as

```
(void)((int)((a) != 0) ? (void)(b) : (void)(c))
```

Otherwise, it is treated as if it were the call “`cond((a) != 0, b, c)`”, with `cond` declared as

```
forall(type T) T cond(int, T, T);

forall(dtype D) void
  * cond(int, D*, void*),
  * cond(int, void*, D*);

forall(dtype D) _atomic void
  * cond(int, _Atomic D*, _Atomic void*),
  * cond(int, _Atomic void*, _Atomic D*);

forall(dtype D) const void
  * cond(int, const D*, const void*),
  * cond(int, const void*, const D*);

forall(dtype D) restrict void
  * cond(int, restrict D*, restrict void*),
  * cond(int, restrict void*, restrict D*);

forall(dtype D) volatile void
  * cond(int, volatile D*, volatile void*),
  * cond(int, volatile void*, volatile D*);

forall(dtype D) _Atomic const void
  * cond(int, _Atomic const D*, _Atomic const void*),
  * cond(int, _Atomic const void*, _Atomic const D*);

forall(dtype D) _Atomic restrict void
  * cond(int, _Atomic restrict D*, _Atomic restrict void*),
  * cond(int, _Atomic restrict void*, _Atomic restrict D*);
```

```

forall(dtype D) _Atomic volatile void
    * cond(int, _Atomic volatile D*, _Atomic volatile void*),
    * cond(int, _Atomic volatile void*, _Atomic volatile D*);

forall(dtype D) const restrict void
    * cond(int, const restrict D*, const restrict void*),
    * cond(int, const restrict void*, const restrict D*);

forall(dtype D) const volatile void
    * cond(int, const volatile D*, const volatile void*),
    * cond(int, const volatile void*, const volatile D*);

forall(dtype D) restrict volatile void
    * cond(int, restrict volatile D*, restrict volatile void*),
    * cond(int, restrict volatile void*, restrict volatile D*);

forall(dtype D) _Atomic const restrict void
    * cond(int, _Atomic const restrict D*, _Atomic const restrict void*),
    * cond(int, _Atomic const restrict void*, _Atomic const restrict D*);

forall(dtype D) _Atomic const volatile void
    * cond(int, _Atomic const volatile D*, _Atomic const volatile void*),
    * cond(int, _Atomic const volatile void*, _Atomic const volatile D*);

forall(dtype D) _Atomic restrict volatile void
    * cond(int, _Atomic restrict volatile D*,
            _Atomic restrict volatile void*),
    * cond(int, _Atomic restrict volatile void*,
            _Atomic restrict volatile D*);

forall(dtype D) const restrict volatile void
    * cond(int, const restrict volatile D*,
            const restrict volatile void*),
    * cond(int, const restrict volatile void*,
            const restrict volatile D*);

forall(dtype D) _Atomic const restrict volatile void
    * cond(int, _Atomic const restrict volatile D*,
            _Atomic const restrict volatile void*),
    * cond(int, _Atomic const restrict volatile void*,
            _Atomic const restrict volatile D*);

```

□ The object of the above is to apply the usual arithmetic conversions when the second and third operands have arithmetic type, and to combine the qualifiers of the second and third operands if they are pointers. □

Examples

```
#include <stdlib.h>
```

```
int i;
long l;
rand() ? i : l;
```

The best interpretation infers the expression's type to be `long` and applies the safe `int-to-long` conversion to `i`.

```
const int *cip;
volatile int *vip;
rand() ? cip : vip;
```

The expression has type `const volatile int*`, with safe conversions applied to the second and third operands to add `volatile` and `const` qualifiers, respectively.

```
rand() ? cip : 0;
```

The expression has type `const int*`, with a specialization conversion applied to `0`.

6.5.16 Assignment operators

Syntax

assignment-expression:
conditional-expression
unary-expression assignment-operator assignment-expression
assignment-operator: one of
`= *= /= %= += -= <<= >>= &= ^= |=`

Rewrite Rules

Let “ \leftarrow ” be any of the assignment operators. Then

$$a \leftarrow b \Rightarrow ?\leftarrow?(\&(a), b)$$

Semantics

Each interpretation of the left operand of an assignment expression is considered separately. For each interpretation that is a bit-field or is declared with the `register` storage class specifier, the expression has one valid interpretation, with the type of the left operand. The right operand is cast to that type, and the assignment expression is ambiguous if either operand is. For the remaining interpretations, the expression is rewritten, and the interpretations of the assignment expression are the interpretations of the corresponding function call. Finally, all interpretations of the expression produced for the different interpretations of the left operand are combined to produce the interpretations of the expression as a whole; where interpretations have compatible result types, the best interpretations are selected in the manner described for function call expressions.

6.5.16.1 Simple assignment

Predefined Identifiers

```
_Bool
  ??(volatile _Bool*, _Bool),
  ??(volatile _Bool*, forall(dtype D) D*),
  ??(volatile _Bool*, forall(ftype F) F*),
```

```

    ??(_Atomic volatile _Bool*, _Bool),
    ??(_Atomic volatile _Bool*, forall(dtype D) D*),
    ??(_Atomic volatile _Bool*, forall(ftype F) F*);
char
    ??(volatile char*, char),
    ??(_Atomic volatile char*, char);
unsigned char
    ??(volatile unsigned char*, unsigned char),
    ??(_Atomic volatile unsigned char*, unsigned char);
signed char
    ??(volatile signed char*, signed char),
    ??(_Atomic volatile signed char*, signed char);
short int
    ??(volatile short int*, short int),
    ??(_Atomic volatile short int*, short int);
unsigned short
    ??(volatile unsigned int*, unsigned int),
    ??(_Atomic volatile unsigned int*, unsigned int);
int
    ??(volatile int*, int),
    ??(_Atomic volatile int*, int);
unsigned int
    ??(volatile unsigned int*, unsigned int),
    ??(_Atomic volatile unsigned int*, unsigned int);
long int
    ??(volatile long int*, long int),
    ??(_Atomic volatile long int*, long int);
unsigned long int
    ??(volatile unsigned long int*, unsigned long int),
    ??(_Atomic volatile unsigned long int*, unsigned long int);
long long int
    ??(volatile long long int*, long long int),
    ??(_Atomic volatile long long int*, long long int);
unsigned long long int
    ??(volatile unsigned long long int*, unsigned long long int),
    ??(_Atomic volatile unsigned long long int*, unsigned long long int);
float
    ??(volatile float*, float),
    ??(_Atomic volatile float*, float);
double
    ??(volatile double*, double),
    ??(_Atomic volatile double*, double);
long double
    ??(volatile long double*, long double),
    ??(_Atomic volatile long double*, long double);
_Complex float
    ??(volatile float*, float),
    ??(_Atomic volatile float*, float);
_Complex double

```

```

    ?=(volatile double*, double),
    ?=?( _Atomic volatile double*, double);
_Complex long double
    ?=(volatile _Complex long double*, _Complex long double),
    ?=?( _Atomic volatile _Complex long double*, _Atomic _Complex long double);

forall(ftype FT) FT
    * ?=(FT* volatile*, FT*),
    * ?=(FT* volatile*, forall(ftype F) F*);

forall(ftype FT) FT const
    * ?=(FT const* volatile*, FT const*),
    * ?=(FT const* volatile*, forall(ftype F) F*);

forall(ftype FT) FT volatile
    * ?=(FT volatile* volatile*, FT*),
    * ?=(FT volatile* volatile*, forall(ftype F) F*);

forall(ftype FT) FT const
    * ?=(FT const volatile* volatile*, FT const*),
    * ?=(FT const volatile* volatile*, forall(ftype F) F*);

forall(dtype DT) DT
    * ?=(DT* restrict volatile*, DT*),
    * ?=(DT* restrict volatile*, void*),
    * ?=(DT* restrict volatile*, forall(dtype D) D*),
    * ?=(DT* _Atomic restrict volatile*, DT*),
    * ?=(DT* _Atomic restrict volatile*, void*),
    * ?=(DT* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT _Atomic
    * ?=?( _Atomic DT* restrict volatile*, DT _Atomic *),
    * ?=?( _Atomic DT* restrict volatile*, void *),
    * ?=?( _Atomic DT* restrict volatile*, forall(dtype D) D*),
    * ?=?( _Atomic DT* _Atomic restrict volatile*, DT _Atomic *),
    * ?=?( _Atomic DT* _Atomic restrict volatile*, void *),
    * ?=?( _Atomic DT* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT const
    * ?=(DT const* restrict volatile*, DT const *),
    * ?=(DT const* restrict volatile*, void const *),
    * ?=(DT const* restrict volatile*, forall(dtype D) D*),
    * ?=(DT const* _Atomic restrict volatile*, DT const *),
    * ?=(DT const* _Atomic restrict volatile*, void const *),
    * ?=(DT const* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT restrict
    * ?=(restrict DT* restrict volatile*, DT restrict *),
    * ?=(restrict DT* restrict volatile*, void *),

```

```

* ??(restrict DT* restrict volatile*, forall(dtype D) D*),
* ??(restrict DT* _Atomic restrict volatile*, DT restrict *),
* ??(restrict DT* _Atomic restrict volatile*, void *),
* ??(restrict DT* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT volatile
* ??(DT volatile* restrict volatile*, DT volatile*),
* ??(DT volatile* restrict volatile*, void volatile *),
* ??(DT volatile* restrict volatile*, forall(dtype D) D*),
* ??(DT volatile* _Atomic restrict volatile*, DT volatile*),
* ??(DT volatile* _Atomic restrict volatile*, void volatile *),
* ??(DT volatile* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT _Atomic const
* ??(DT _Atomic const* restrict volatile*, DT _Atomic const *),
* ??(DT _Atomic const* restrict volatile*, void const *),
* ??(DT _Atomic const* restrict volatile*, forall(dtype D) D*),
* ??(DT _Atomic const* _Atomic restrict volatile*, DT _Atomic const *),
* ??(DT _Atomic const* _Atomic restrict volatile*, void const *),
* ??(DT _Atomic const* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT _Atomic restrict
* ??(_Atomic restrict DT* restrict volatile*, DT _Atomic restrict *),
* ??(_Atomic restrict DT* restrict volatile*, void *),
* ??(_Atomic restrict DT* restrict volatile*, forall(dtype D) D*),
* ??(_Atomic restrict DT* _Atomic restrict volatile*, DT _Atomic restrict *),
* ??(_Atomic restrict DT* _Atomic restrict volatile*, void *),
* ??(_Atomic restrict DT* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT _Atomic volatile
* ??(DT _Atomic volatile* restrict volatile*, DT _Atomic volatile*),
* ??(DT _Atomic volatile* restrict volatile*, void volatile *),
* ??(DT _Atomic volatile* restrict volatile*, forall(dtype D) D*),
* ??(DT _Atomic volatile* _Atomic restrict volatile*, DT _Atomic volatile*),
* ??(DT _Atomic volatile* _Atomic restrict volatile*, void volatile *),
* ??(DT _Atomic volatile* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT const restrict
* ??(DT const restrict* restrict volatile*, DT const restrict *),
* ??(DT const restrict* restrict volatile*, void const *),
* ??(DT const restrict* restrict volatile*, forall(dtype D) D*),
* ??(DT const restrict* _Atomic restrict volatile*, DT const restrict *),
* ??(DT const restrict* _Atomic restrict volatile*, void const *),
* ??(DT const restrict* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT const volatile
* ??(DT const volatile* restrict volatile*, DT const volatile*),
* ??(DT const volatile* restrict volatile*, void const volatile*),
* ??(DT const volatile* restrict volatile*, forall(dtype D) D*),

```

```

* ??(DT const volatile* _Atomic restrict volatile*, DT const volatile*),
* ??(DT const volatile* _Atomic restrict volatile*, void const volatile *),
* ??(DT const volatile* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT restrict volatile
* ??(DT restrict volatile* restrict volatile*, DT restrict volatile*),
* ??(DT restrict volatile* restrict volatile*, void volatile *),
* ??(DT restrict volatile* restrict volatile*, forall(dtype D) D*),
* ??(DT restrict volatile* _Atomic restrict volatile*, DT restrict volatile*),
* ??(DT restrict volatile* _Atomic restrict volatile*, void volatile *),
* ??(DT restrict volatile* _Atomic restrict volatile*, forall(dtype D) D*);

forall(dtype DT) DT _Atomic const restrict
* ??(DT _Atomic const restrict* restrict volatile*,
    DT _Atomic const restrict *),
* ??(DT _Atomic const restrict* restrict volatile*,
    void const *),
* ??(DT _Atomic const restrict* restrict volatile*,
    forall(dtype D) D*),
* ??(DT _Atomic const restrict * _Atomic restrict volatile*,
    DT _Atomic const restrict *),
* ??(DT _Atomic const restrict * _Atomic restrict volatile*,
    void const *),
* ??(DT _Atomic const restrict * _Atomic restrict volatile*,
    forall(dtype D) D*);

forall(dtype DT) DT _Atomic const volatile
* ??(DT _Atomic const volatile* restrict volatile*,
    DT _Atomic const volatile*),
* ??(DT _Atomic const volatile* restrict volatile*,
    void const volatile *),
* ??(DT _Atomic const volatile* restrict volatile*,
    forall(dtype D) D*),
* ??(DT _Atomic const volatile* _Atomic restrict volatile*,
    DT _Atomic const volatile *),
* ??(DT _Atomic const volatile* _Atomic restrict volatile*,
    void const volatile*),
* ??(DT _Atomic const volatile* _Atomic restrict volatile*,
    forall(dtype D) D*);

forall(dtype DT) DT _Atomic restrict volatile
* ??(DT _Atomic restrict volatile* restrict volatile*,
    DT _Atomic restrict volatile*),
* ??(DT _Atomic restrict volatile* restrict volatile*,
    void volatile *),
* ??(DT _Atomic restrict volatile* restrict volatile*,
    forall(dtype D) D*),
* ??(DT _Atomic restrict volatile* _Atomic restrict volatile*,
    DT _Atomic restrict volatile*),

```

```

* ?=(DT _Atomic restrict volatile* _Atomic restrict volatile*,
    void volatile *),
* ?=(DT _Atomic restrict volatile* _Atomic restrict volatile*,
    forall(dtype D) D*);

forall(dtype DT) DT const restrict volatile
* ?=(DT const restrict volatile* restrict volatile*,
    DT const restrict volatile*),
* ?=(DT const restrict volatile* restrict volatile*,
    void const volatile *),
* ?=(DT const restrict volatile* restrict volatile*,
    forall(dtype D) D*),
* ?=(DT const restrict volatile* _Atomic restrict volatile*,
    DT const restrict volatile*),
* ?=(DT const restrict volatile* _Atomic restrict volatile*,
    void const volatile*),
* ?=(DT const restrict volatile* _Atomic restrict volatile*,
    forall(dtype D) D*);

forall(dtype DT) DT _Atomic const restrict volatile
* ?=(DT _Atomic const restrict volatile* restrict volatile*,
    DT _Atomic const restrict volatile*),
* ?=(DT _Atomic const restrict volatile* restrict volatile*,
    void const volatile*),
* ?=(DT _Atomic const restrict volatile* restrict volatile*,
    forall(dtype D) D*),
* ?=(DT _Atomic const restrict volatile* _Atomic restrict volatile*,
    DT _Atomic const restrict volatile*),
* ?=(DT _Atomic const restrict volatile* _Atomic restrict volatile*,
    void const volatile*),
* ?=(DT _Atomic const restrict volatile* _Atomic restrict volatile*,
    forall(dtype D) D*);

forall(dtype DT) void
* ?=(void* restrict volatile*, DT *);

forall(dtype DT) void const
* ?=(void const* restrict volatile*, DT const*);

forall(dtype DT) void volatile
* ?=(void volatile* restrict volatile*, DT volatile*);

forall(dtype DT) void const volatile
* ?=(void const volatile* restrict volatile*, DT const volatile*);

```

- The pattern of overloadings for simple assignment resembles that of pointer increment and decrement, except that the polymorphic pointer assignment functions declare a `dtype` parameter, instead of a `cdtype` parameter, because the left operand may be a pointer to an incomplete type.
-

For every complete structure or union type S there exist

```
 $S$  ?=( $S$  volatile*,  $S$ ),
    ?=( $S$  _Atomic volatile*,  $S$ );
```

For every extended integer type X there exist

```
 $X$  ?=( $X$  volatile* ,  $X$ ),
    ?=( $X$  _Atomic volatile* ,  $X$ );
```

For every complete enumerated type E there exist

```
 $E$  ?=( $E$  volatile*, int),
    ?=( $E$  _Atomic volatile*, int);
```

□ The right-hand argument is `int` because enumeration constants have type `int`. □

Semantics

The structure assignment functions provide member-wise assignment; each non-array member and each element of each array member of the right argument is assigned to the corresponding member or element of the left argument using the assignment function defined for its type. All other assignment functions have the same effect as the corresponding C assignment expression.

□ Note that, by default, union assignment uses C semantics—that is, bitwise copy—even if some of the union members have programmer-defined assignment functions. □

6.5.16.2 Compound assignment

Predefined Identifiers

```
forall(type T) T
    * ?+=(T* restrict volatile*, ptrdiff_t),
    * ?-=(T* restrict volatile*, ptrdiff_t),
    * ?+=(T* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic
    * ?+=(T _Atomic* restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic* restrict volatile*, ptrdiff_t),
    * ?+=(T _Atomic* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T const
    * ?+=(T const* restrict volatile*, ptrdiff_t),
    * ?-=(T const* restrict volatile*, ptrdiff_t),
    * ?+=(T const* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T const* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T restrict
    * ?+=(T restrict* restrict volatile*, ptrdiff_t),
    * ?-=(T restrict* restrict volatile*, ptrdiff_t),
    * ?+=(T restrict* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T restrict* _Atomic restrict volatile*, ptrdiff_t);
```

```

forall(type T) T volatile
    * ?+=(T volatile* restrict volatile*, ptrdiff_t),
    * ?-=(T volatile* restrict volatile*, ptrdiff_t),
    * ?+=(T volatile* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic const
    * ?+=(T _Atomic const restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic const restrict volatile*, ptrdiff_t),
    * ?+=(T _Atomic const _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic const _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic restrict
    * ?+=(T _Atomic restrict* restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic restrict* restrict volatile*, ptrdiff_t),
    * ?+=(T _Atomic restrict* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic restrict* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic volatile
    * ?+=(T _Atomic volatile* restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic volatile* restrict volatile*, ptrdiff_t),
    * ?+=(T _Atomic volatile* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T const restrict
    * ?+=(T const restrict* restrict volatile*, ptrdiff_t),
    * ?-=(T const restrict* restrict volatile*, ptrdiff_t),
    * ?+=(T const restrict* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T const restrict* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T const volatile
    * ?+=(T const volatile* restrict volatile*, ptrdiff_t),
    * ?-=(T const volatile* restrict volatile*, ptrdiff_t),
    * ?+=(T const volatile* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T const volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T restrict volatile
    * ?+=(T restrict volatile* restrict volatile*, ptrdiff_t),
    * ?-=(T restrict volatile* restrict volatile*, ptrdiff_t),
    * ?+=(T restrict volatile* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T restrict volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic const restrict
    * ?+=(T _Atomic const restrict* restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic const restrict* restrict volatile*, ptrdiff_t),
    * ?+=(T _Atomic const restrict* _Atomic restrict volatile*, ptrdiff_t),
    * ?-=(T _Atomic const restrict* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic const volatile

```

```

* ?+?(T _Atomic const volatile* restrict volatile*, ptrdiff_t),
* ?-?(T _Atomic const volatile* restrict volatile*, ptrdiff_t),
* ?+?(T _Atomic const volatile* _Atomic restrict volatile*, ptrdiff_t),
* ?-?(T _Atomic const volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic restrict volatile
* ?+?(T _Atomic restrict volatile* restrict volatile*, ptrdiff_t),
* ?-?(T _Atomic restrict volatile* restrict volatile*, ptrdiff_t),
* ?+?(T _Atomic restrict volatile* _Atomic restrict volatile*, ptrdiff_t),
* ?-?(T _Atomic restrict volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T const restrict volatile
* ?+?(T const restrict volatile* restrict volatile*, ptrdiff_t),
* ?-?(T const restrict volatile* restrict volatile*, ptrdiff_t),
* ?+?(T const restrict volatile* _Atomic restrict volatile*, ptrdiff_t),
* ?-?(T const restrict volatile* _Atomic restrict volatile*, ptrdiff_t);

forall(type T) T _Atomic const restrict volatile
* ?+?(T _Atomic const restrict volatile* restrict volatile*, ptrdiff_t),
* ?-?(T _Atomic const restrict volatile* restrict volatile*, ptrdiff_t),
* ?+?(T _Atomic const restrict volatile* _Atomic restrict volatile*, ptrdiff_t),
* ?-?(T _Atomic const restrict volatile* _Atomic restrict volatile*, ptrdiff_t);

_Bool
?*=?(_Bool volatile*, _Bool),
?/?=?(_Bool volatile*, _Bool),
?+=?(_Bool volatile*, _Bool),
?-=?(_Bool volatile*, _Bool),
?%=?(_Bool volatile*, _Bool),
?<<=?(_Bool volatile*, int),
?>>=?(_Bool volatile*, int),
?&=?(_Bool volatile*, _Bool),
?^=?(_Bool volatile*, _Bool),
?|=?(_Bool volatile*, _Bool);

char
?*=? (char volatile*, char),
?/?=? (char volatile*, char),
?+=? (char volatile*, char),
?-=? (char volatile*, char),
?%=? (char volatile*, char),
?<<=? (char volatile*, int),
?>>=? (char volatile*, int),
?&=? (char volatile*, char),
?^=? (char volatile*, char),
?|=? (char volatile*, char);

unsigned char
?*=? (unsigned char volatile*, unsigned char),
?/?=? (unsigned char volatile*, unsigned char),
?+=? (unsigned char volatile*, unsigned char),

```

```

?-=?(unsigned char volatile*, unsigned char),
?%=?(unsigned char volatile*, unsigned char),
?<<=?(unsigned char volatile*, int),
?>>=?(unsigned char volatile*, int),
?&=?(unsigned char volatile*, unsigned char),
?^=?(unsigned char volatile*, unsigned char),
?|=?(unsigned char volatile*, unsigned char);
signed char
?*=?(signed char volatile*, signed char),
?/=?(signed char volatile*, signed char),
?+=?(signed char volatile*, signed char),
?-=?(signed char volatile*, signed char),
?%=?(signed char volatile*, signed char),
?<<=?(signed char volatile*, int),
?>>=?(signed char volatile*, int),
?&=?(signed char volatile*, signed char),
?^=?(signed char volatile*, signed char),
?|=?(signed char volatile*, signed char);
short int
?*=?(short int volatile*, short int),
?/=?(short int volatile*, short int),
?+=?(short int volatile*, short int),
?-=?(short int volatile*, short int),
?%=?(short int volatile*, short int),
?<<=?(short int volatile*, int),
?>>=?(short int volatile*, int),
?&=?(short int volatile*, short int),
?^=?(short int volatile*, short int),
?|=?(short int volatile*, short int);
unsigned short int
?*=?(unsigned short int volatile*, unsigned short int),
?/=?(unsigned short int volatile*, unsigned short int),
?+=?(unsigned short int volatile*, unsigned short int),
?-=?(unsigned short int volatile*, unsigned short int),
?%=?(unsigned short int volatile*, unsigned short int),
?<<=?(unsigned short int volatile*, int),
?>>=?(unsigned short int volatile*, int),
?&=?(unsigned short int volatile*, unsigned short int),
?^=?(unsigned short int volatile*, unsigned short int),
?|=?(unsigned short int volatile*, unsigned short int);
int
?*=?(int volatile*, int),
?/=?(int volatile*, int),
?+=?(int volatile*, int),
?-=?(int volatile*, int),
?%=?(int volatile*, int),
?<<=?(int volatile*, int),
?>>=?(int volatile*, int),
?&=?(int volatile*, int),

```

```

    ^=?(int volatile*, int),
    |=?(int volatile*, int);
unsigned int
    *=?(unsigned int volatile*, unsigned int),
    /=?(unsigned int volatile*, unsigned int),
    +=?(unsigned int volatile*, unsigned int),
    -=?(unsigned int volatile*, unsigned int),
    %=?(unsigned int volatile*, unsigned int),
    <<=?(unsigned int volatile*, int),
    >>=?(unsigned int volatile*, int),
    &=?(unsigned int volatile*, unsigned int),
    ^=?(unsigned int volatile*, unsigned int),
    |=?(unsigned int volatile*, unsigned int);
long int
    *=?(long int volatile*, long int),
    /=?(long int volatile*, long int),
    +=?(long int volatile*, long int),
    -=?(long int volatile*, long int),
    %=?(long int volatile*, long int),
    <<=?(long int volatile*, int),
    >>=?(long int volatile*, int),
    &=?(long int volatile*, long int),
    ^=?(long int volatile*, long int),
    |=?(long int volatile*, long int);
unsigned long int
    *=?(unsigned long int volatile*, unsigned long int),
    /=?(unsigned long int volatile*, unsigned long int),
    +=?(unsigned long int volatile*, unsigned long int),
    -=?(unsigned long int volatile*, unsigned long int),
    %=?(unsigned long int volatile*, unsigned long int),
    <<=?(unsigned long int volatile*, int),
    >>=?(unsigned long int volatile*, int),
    &=?(unsigned long int volatile*, unsigned long int),
    ^=?(unsigned long int volatile*, unsigned long int),
    |=?(unsigned long int volatile*, unsigned long int);
long long int
    *=?(long long int volatile*, long long int),
    /=?(long long int volatile*, long long int),
    +=?(long long int volatile*, long long int),
    -=?(long long int volatile*, long long int),
    %=?(long long int volatile*, long long int),
    <<=?(long long int volatile*, int),
    >>=?(long long int volatile*, int),
    &=?(long long int volatile*, long long int),
    ^=?(long long int volatile*, long long int),
    |=?(long long int volatile*, long long int);
unsigned long long int
    *=?(unsigned long long int volatile*, unsigned long long int),
    /=?(unsigned long long int volatile*, unsigned long long int),

```

```

?+=?(unsigned long long int volatile*, unsigned long long int),
?-=? (unsigned long long int volatile*, unsigned long long int),
?%=?(unsigned long long int volatile*, unsigned long long int),
?<<=? (unsigned long long int volatile*, int),
?>>=? (unsigned long long int volatile*, int),
?&=? (unsigned long long int volatile*, unsigned long long int),
?^=? (unsigned long long int volatile*, unsigned long long int),
?!=? (unsigned long long int volatile*, unsigned long long int);
float
?*=? (float volatile*, float),
?/=?(float volatile*, float),
?+=?(float volatile*, float),
?-=? (float volatile*, float);
double
?*=? (double volatile*, double),
?/=?(double volatile*, double),
?+=?(double volatile*, double),
?-=? (double volatile*, double);
long double
?*=? (long double volatile*, long double),
?/=?(long double volatile*, long double),
?+=?(long double volatile*, long double),
?-=? (long double volatile*, long double);
_Complex float
?*=? (_Complex float volatile*, _Complex float),
?/=?( _Complex float volatile*, _Complex float),
?+=?( _Complex float volatile*, _Complex float),
?-=? (_Complex float volatile*, _Complex float);
_Complex double
?*=? (_Complex double volatile*, _Complex double),
?/=?( _Complex double volatile*, _Complex double),
?+=?( _Complex double volatile*, _Complex double),
?-=? (_Complex double volatile*, _Complex double);
_Complex long double
?*=? (_Complex long double volatile*, _Complex long double),
?/=?( _Complex long double volatile*, _Complex long double),
?+=?( _Complex long double volatile*, _Complex long double),
?-=? (_Complex long double volatile*, _Complex long double);

```

For every extended integer type X there exist

```

X
?*=? (X volatile*, X),
?/=?(X volatile*, X),
?+=?(X volatile*, X),
?-=? (X volatile*, X),
?%=?(X volatile*, X),
?<<=? (X volatile*, int),
?>>=? (X volatile*, int),
?&=? (X volatile*, X),

```

```
?^=?(X volatile*, X),
?|=?(X volatile*, X);
```

For every complete enumerated type E there exist

```
E
?*=?(E volatile*, E),
?/=?(E volatile*, E),
?+=?(E volatile*, E),
?-=?(E volatile*, E),
?%=?(E volatile*, E),
?<<=?(E volatile*, int),
?>>=?(E volatile*, int),
?&=?(E volatile*, E),
?^=?(E volatile*, E),
?|=?(E volatile*, E);
```

6.5.17 Comma operator

Syntax

expression:

```
assignment-expression
expression , assignment-expression
```

Semantics

In the comma expression “ a, b ”, the first operand is interpreted as “(void) (a)”, which shall be unambiguous. The interpretations of the expression are the interpretations of the second operand.

6.6 Constant expressions

6.7 Declarations

Syntax

declaration: ...

```
type-declaration
spec-definition
```

Constraints

If an identifier has no linkage, there shall be no more than one declaration of the identifier (in a declarator or type specifier) with compatible types in the same scope and in the same name space, except that:

- a typedef name may be redefined to denote the same type as it currently does, provided that type is not a variably modified type;
 - tags may be redeclared as specified in section 6.7.2.3 of the ISO/IEC C standard.
- This constraint adds the phrase “with compatible types” to the ISO/IEC C constraint, to allow overloading. □

An identifier declared by a type declaration shall not be redeclared as a parameter in a function definition whose declarator includes an identifier list.

- This restriction echos ISO/IEC C's ban on the redeclaration of typedef names as parameters. This avoids an ambiguity between old-style function declarations and new-style function prototypes:

```
void f(Complex, /* ... 3000 characters ... */);
void g(Complex, /* ... 3000 characters ... */)
    int Complex; { /* ... */ }
```

Without the rule, `Complex` would be a type in the first case, and a parameter name in the second.

□

6.7.2 Type specifiers

Syntax

type-specifier: ...
forall-specifier

Semantics

Forall specifiers are discussed in 6.7.2.5.

6.7.2.1 Structure and union specifiers

Semantics

Cforall extends the ISO/IEC C definition of *anonymous structure* to include structure specifiers with tags, and extends the ISO/IEC C definition of *anonymous union* to include union specifiers with tags.

- This extension imitates an extension in the Plan 9 C compiler [8].

□

Examples

```
struct point {
    int x, y;
};
struct color_point {
    enum { RED, BLUE, GREEN } color;
    struct point;
};
struct color_point cp;
cp.x = 0;
cp.color = RED;

struct literal {
    enum { NUMBER, STRING } tag;
    union {
        double n;
        char *s;
    };
};
```

```

};
struct literal *next;
int          length;
extern int   strlen(const char*);
/* ... */
if (next->tag == STRING) length = strlen(next->s);

```

6.7.2.5 Forall specifiers

Syntax

```

forall-specifier:
    forall ( type-parameter-list )

```

Constraints

If the *declaration-specifiers* of a declaration that contains a *forall-specifier* declares a structure or union tag, the types of the members of the structure or union shall not use any of the type identifiers declared by the *type-parameter-list*.

- This sort of declaration is illegal because the scope of the type identifiers ends at the end of the declaration, but the scope of the structure tag does not.

```
forall(type T) struct Pair { T a,b; } mkPair(T, T); /* illegal */
```

If an instance of `struct Pair` was declared later in the current scope, what would the members' type be? □

Semantics

The *type-parameter-lists* and assertions of the *forall-specifiers* declare type identifiers, function and object identifiers with no linkage.

If, in the declaration “T D1”, T contains *forall-specifiers* and D1 has the form

```
D(parameter-type-list)
```

then a type identifier declared by one of the *forall-specifiers* is an *inferred parameter* of the function declarator if and only if it is not an inferred parameter of a function declarator in D, and it is used in the type of a parameter in the following *type-parameter-list* or it and an inferred parameter are used as arguments of a specification in one of the *forall-specifiers*. The identifiers declared by assertions that use an inferred parameter of a function declarator are assertion parameters of that function declarator.

- Since every inferred parameter is used by some parameter, inference can be understood as a single bottom-up pass over the expression tree, that only needs to apply local reasoning at each node.

If this restriction were lifted, it would be possible to write

```
forall(type T) T* alloc(void);
int *p = alloc();
```

Here `alloc()` would receive `int` as an inferred argument, and return an `int*`. In general, if a call to `alloc()` is a subexpression of an expression involving polymorphic functions and overloaded identifiers, there could be considerable distance between the call and the subexpression that causes T to be bound.

With the current restriction, `alloc()` must be given an argument that determines T:

```
forall(type T) T* alloc(T initial_value);
```

□

If a function declarator is part of a function definition, its inferred parameters and assertion parameters have block scope; otherwise, identifiers declared by assertions have a *declaration scope*, which terminates at the end of the *declaration*.

A function type that has at least one inferred parameter is a *polymorphic function* type. Function types with no inferred parameters are *monomorphic function* types. One function type is *less polymorphic* than another if it has fewer inferred parameters, or if it has the same number of inferred parameters and fewer of its explicit parameters have types that depend on an inferred parameter.

The names of inferred parameters and the order of identifiers in forall specifiers are not relevant to polymorphic function type compatibility. Let f and g be two polymorphic function types with the same number of inferred parameters, and let f_i and g_i be the inferred parameters of f and g in their order of occurrence in the function types' *parameter-type-lists*. Let f' be f with every occurrence of f_i replaced by g_i , for all i . Then f and g are compatible types if f' 's and g 's return types and parameter lists are compatible, and if for every assertion parameter of f' there is an assertion parameter in g with the same identifier and compatible type, and vice versa.

Examples

Consider these analogous monomorphic and polymorphic declarations.

```
int fi(int);
forall(type T) T fT(T);
```

`fi()` takes an `int` and returns an `int`. `fT()` takes a `T` and returns a `T`, for any type `T`.

```
int (*pfi)(int) = fi;
forall(type T) T (*pfT)(T) = fT;
```

`pfi` and `pfT` are pointers to functions. `pfT` is not polymorphic, but the function it points at is.

```
int (*fvpfi(void))(int) {
    return pfi;
}
forall(type T) T (*fvpfT(void))(T) {
    return pfT;
}
```

`fvpfi()` and `fvpfT()` are functions taking no arguments and returning pointers to functions. `fvpfT()` is monomorphic, but the function that its return value points at is polymorphic.

```
forall(type T)          int (*fTpfT(T))(int);
forall(type T)          T  (*fTpfT(T))(T);
forall(type T, type U)  U  (*fTpfU(T))(U);
```

`fTpfT()` is a polymorphic function that returns a pointer to a monomorphic function taking an integer and returning an integer. It could return `pfi`. `fTpfT()` is subtle: it is a polymorphic function returning a *monomorphic* function taking and returning `T`, where `T` is an inferred parameter of `fTpfT()`. For instance, in the expression “`fTpfT(17)`”, `T` is inferred to be `int`, and the returned value would have type `int (*) (int)`. “`fTpfT(17)(13)`” and “`fTpfT("yes")("no")`” are legal, but “`fTpfT(17)("no")`” is illegal. `fTpfU()` is polymorphic (in type `T`), and returns a pointer to a function that is polymorphic (in type `U`). “`f5(17)("no")`” is a legal expression of type `char*`.

```
forall(type T, type U, type V) U* f(T*, U, V* const);
forall(type U, type V, type W) U* g(V*, U, W* const);
```

The functions $f()$ and $g()$ have compatible types. Let f and g be their types; then $f_1 = T, f_2 = U, f_3 = V, g_1 = V, g_2 = U,$, and $g_3 = W$. Replacing every f_i by g_i in f gives

```
forall(type V, type U, type W) U* f(V*, U, W* const);
```

which has a return type and parameter list that is compatible with g .

□ The word “**type**” in a forall specifier is redundant at the moment, but I want to leave room for inferred parameters of ordinary types in case parameterized types get added one day.

Even without parameterized types, I might try to allow

```
forall(int n) int sum(int vector[n]);
```

but C currently rewrites array parameters as pointer parameters, so the effects of such a change require more thought. □

□ A polymorphic declaration must do two things: it must introduce type parameters, and it must apply assertions to those types. Adding this to existing C declaration syntax and semantics was delicate, and not entirely successful.

C depends on declaration-before-use, so a forall specifier must introduce type names before they can be used in the declaration specifiers. This could be done by making the forall specifier part of the declaration specifiers, or by making it a new introductory clause of declarations.

Assertions are also part of polymorphic function types, because it must be clear which functions have access to the assertion parameters declared by the assertions. All attempts to put assertions inside an introductory clause produced complex semantics and confusing code. Building them into the declaration specifiers could be done by placing them in the function’s parameter list, or in a forall specifier that is a declaration specifier. Assertions are also used with type parameters of specifications, and by type declarations. For consistency’s sake it seems best to attach assertions to the type declarations in forall specifiers, which means that forall specifiers must be declaration specifiers. □

6.7.3 Type qualifiers

Cforall defines a new type qualifier `lvalue`.

Syntax

```
type-qualifier: ...
               lvalue
```

Constraints

Types other than type parameters and pointer types whose referenced type is an object type shall not be restrict-qualified.

Semantics

An object’s type may be a restrict-qualified type parameter. `restrict` does not establish any special semantics in that case.

□ Cforall loosens the constraint on the restrict qualifier so that restrict-qualified pointers may be passed to polymorphic functions. □

`lvalue` may be used to qualify the return type of a function type. Let `T` be an unqualified version of a type; then the result of calling a function with return type `lvalue T` is a modifiable lvalue of type `T`. `const` and `volatile` qualifiers may also be added to indicate that the function result is a constant or volatile lvalue.

- The `const` and `volatile` qualifiers can only be sensibly used to qualify the return type of a function if the `lvalue` qualifier is also used. □

An lvalue-qualified type may be used in a cast expression if the operand is an lvalue; the result of the expression is an lvalue.

- `lvalue` provides some of the functionality of C++’s “`T&`” (reference to object of type `T`) type. Reference types have four uses in C++.

- They are necessary for user-defined operators that return lvalues, such as “subscript” and “dereference”.
- A reference can be used to define an alias for a complicated lvalue expression, as a way of getting some of the functionality of the Pascal `with` statement. The following C++ code gives an example.

```
{ char& code = long_name.some_field[i].data->code;
  code = toupper(code);
}
```

This is not very useful.

- A reference parameter can be used to allow a function to modify an argument without forcing the caller to pass the address of the argument. This is most useful for user-defined assignment operators. In C++, plain assignment is done by a function called “`operator=`”, and the two expressions

```
a = b;
operator=(a,b);
```

are equivalent. If `a` and `b` are of type `T`, then the first parameter of `operator=` must have type “`T&`”. It cannot have type `T`, because then assignment couldn’t alter the variable, and it can’t have type “`T*`”, because the assignment would have to be written “`&a = b;`”.

In the case of user-defined operators, this could just as well be handled by using pointer types and by changing the rewrite rules so that “`a = b;`” is equivalent to “`operator=(&(a),b)`”. Reference parameters of “normal” functions are Bad Things, because they remove a useful property of C function calls: an argument can only be modified by a function if it is preceded by “`&`”.

- References to `const`-qualified types can be used instead of value parameters. Given the C++ function call “`fiddle(a_thing)`”, where the type of `a_thing` is `Thing`, the type of `fiddle` could be either of

```
void fiddle(Thing);
void fiddle(const Thing&);
```

If the second form is used, then constructors and destructors are not invoked to create a temporary variable at the call site (and it is bad style for the caller to make any assumptions about such things), and within `fiddle` the parameter is subject to the usual problems caused by aliases. The reference form might be chosen for efficiency’s sake if `Things` are too large or their constructors or destructors are too expensive. An implementation may switch between them without causing trouble for well-behaved clients. This leaves the implementor to define “too large” and “too expensive”.

I propose to push this job onto the compiler by allowing it to implement

```
void fiddle(const volatile Thing);
```

with call-by-reference. Since it knows all about the size of `Things` and the parameter passing mechanism, it should be able to come up with a better definition of “too large”, and may be able to make a good guess at “too expensive”.

In summary, since references are only really necessary for returning lvalues, I’ll only provide lvalue functions. □

6.7.9 Initialization

An expression that is used as an *initializer* is treated as being cast to the type of the object being initialized. An expression used in an *initializer-list* is treated as being cast to the type of the aggregate member that it initializes. In either case the cast must have a single unambiguous interpretation.

6.7.11 Specification definitions

Syntax

spec-definition:

```
spec identifier ( type-parameter-list ) { spec-declaration-listopt }
```

spec-declaration-list:

```
spec-declaration ;
spec-declaration-list spec-declaration ;
```

spec-declaration:

```
specifier-qualifier-list declarator-list
```

declarator-list:

```
declarator
declarator-list , declarator
```

□ The declarations allowed in a specification are much the same as those allowed in a structure, except that bit fields are not allowed, and incomplete types and function types are allowed. □

Semantics

A *specification definition* defines a name for a *specification*: a parameterized collection of object and function declarations.

The declarations in a specification consist of the declarations in the *spec-declaration-list* and declarations produced by any assertions in the *spec-parameter-list*. If the collection contains two declarations that declare the same identifier and have compatible types, they are combined into one declaration with the composite type constructed from the two types.

6.7.11.1 Assertions

Syntax

assertion-list:

```
assertion
assertion-list assertion
```

assertion:

```
| identifier ( type-name-list )
| spec-declaration
```

```

type-name-list:
    type-name
    type-name-list , type-name

```

Constraints

The *identifier* in an assertion that is not a *spec-declaration* shall be the name of a specification. The *type-name-list* shall contain one *type-name* argument for each *type-parameter* in that specification's *spec-parameter-list*. If the *type-parameter* uses type-class **type**, the argument shall be the type name of an object type; if it uses **dtype**, the argument shall be the type name of an object type or an incomplete type; and if it uses **ftype**, the argument shall be the type name of a function type.

Semantics

An *assertion* is a declaration of a collection of objects and functions, called *assertion parameters*.

The assertion parameters produced by an assertion that applies the name of a specification to type arguments are found by taking the declarations specified in the specification and treating each of the specification's parameters as a synonym for the corresponding *type-name* argument.

The collection of assertion parameters produced by the *assertion-list* are found by combining the declarations produced by each assertion. If the collection contains two declarations that declare the same identifier and have compatible types, they are combined into one declaration with the composite type constructed from the two types.

Examples

```

forall(type T | T ?*(T,T))
T square(T val) {
    return val + val;
}

spec summable(type T) {
    T      ?+?(T*, T);
    const T 0;
};

spec list_of(type List, type Element) {
    Element car(List);
    List    cdr(List);
    List    cons(Element, List);
    List    nil;
    int     is_nil(List);
};

spec sum_list(type List,
              type Element | summable(Element) | list_of(List, Element)) {};

```

`sum_list` contains seven declarations, which describe a list whose elements can be added up. The assertion “`|sum_list(i_list, int)`” produces the assertion parameters

```

int      ?+?(int*, int);
const int 0;
int      car(i_list);
i_list   cdr(i_list);

```

```

i_list    cons(int, i_list);
i_list    nil;
int       is_nil;

```

6.7.12 Type declarations

Syntax

```

type-parameter-list:
    type-parameter
    type-parameter-list , type-parameter
type-parameter:
    type-class identifier assertion-listopt
type-class:
    type
    dtype
    ftype
type-declaration:
    storage-class-specifieropt type type-declarator-list ;
type-declarator-list:
    type-declarator
    type-declarator-list , type-declarator
type-declarator:
    identifier assertion-listopt = type-name
    identifier assertion-listopt

```

Constraints

If a type declaration has block scope, and the declared identifier has external or internal linkage, the declaration shall have no initializer for the identifier.

Semantics

A *type-parameter* or a *type-declarator* declares an identifier to be a type name for a type distinct from all other types.

An identifier declared by a *type-parameter* has no linkage. Identifiers declared with type-class **type** are object types; those declared with type-class **dtype** are incomplete types; and those declared with type-class **ftype** are function types. The identifier has block scope that terminates at the end of the *spec-declaration-list* or polymorphic function that contains the *type-parameter*.

A *type-declarator* with an initializer is a *type definition*. The declared identifier is an incomplete type within the initializer, and an object type after the end of the initializer. The type in the initializer is called the *implementation type*. Within the scope of the declaration, implicit conversions can be performed between the defined type and the implementation type, and between pointers to the defined type and pointers to the implementation type.

A type declaration without an initializer and without a storage-class specifier or with storage-class specifier **static** defines an incomplete type. If a translation unit or block contains one or more such declarations for an identifier, it must contain exactly one definition of the identifier (but not in an enclosed block, which would define a new type known only within that block).

- Incomplete type declarations allow compact mutually-recursive types.

```

type t1;                                /* Incomplete type declaration. */
type t2 = struct { t1* p; /* ... */ };
type t1 = struct { t2* p; /* ... */ };

```

Without them, mutual recursion could be handled by declaring mutually recursive structures, then initializing the types to those structures.

```

struct s1;
type t2 = struct s2 { struct s1* p; /* ... */ };
type t1 = struct s1 { struct s2* p; /* ... */ };

```

This introduces extra names, and may force the programmer to cast between the types and their implementations. □

A type declaration without an initializer and with storage-class specifier `extern` is an *opaque type declaration*. Opaque types are object types. An opaque type is not a *constant-expression*; neither is a structure or union that has a member whose type is not a *constant-expression*. Every other object type is a *constant-expression*. Objects with static storage duration shall be declared with a type that is a *constant-expression*.

□ Type declarations can declare identifiers with external linkage, whereas typedef declarations declare identifiers that only exist within a translation unit. These opaque types can be used in declarations, but the implementation of the type is not visible.

Static objects can not have opaque types because space for them would have to be allocated at program start-up. This is a deficiency, but I don't want to deal with "module initialization" code just now. □

An incomplete type which is not a qualified version of a type is a value of type-class `dtype`. An object type which is not a qualified version of a type is a value of type-classes `type` and `dtype`. A function type is a value of type-class `ftype`.

□ Syntactically, a type value is a *type-name*, which is a declaration for an object which omits the identifier being declared.

Object types are precisely the types that can be instantiated. Type qualifiers are not included in type values because the compiler needs the information they provide at compile time to detect illegal statements or to produce efficient machine instructions. For instance, the code that a compiler must generate to manipulate an object that has volatile-qualified type may be different from the code to manipulate an ordinary object.

Type qualifiers are a weak point of C's type system. Consider the standard library function `strchr()` which, given a string and a character, returns a pointer to the first occurrence of the character in the string.

```

char *strchr(const char *s, int c) {
    char real_c = c; /* done because c was declared as int. */
    for (; *s != real_c; s++)
        if (*s == '\0') return NULL;
    return (char*)s;
}

```

The parameter `s` must be `const char*`, because `strchr()` might be used to search a constant string, but the return type must be `char*`, because the result might be used to modify a non-constant string. Hence the body must perform a cast, and (even worse) `strchr()` provides a type-safe way to attempt to modify constant strings. What is needed is some way to say that `s`'s type might contain qualifiers, and the result type has exactly the same qualifiers. Polymorphic functions do not provide a fix for this deficiency, because type qualifiers are not part of type

values. Instead, overloading can be used to define `strchr()` for each combination of qualifiers.

□

□ Since incomplete types are not type values, they can not be used as the initializer in a type declaration, or as the type of a structure or union member. This prevents the declaration of types that contain each other.

```
type t1;
type t2 = t1;                /* illegal: incomplete type 't1'. */
type t1 = t2;
```

The initializer in a file-scope declaration must be a constant expression. This means type declarations can not build on opaque types, which is a deficiency.

```
extern type Huge; /* extended-precision integer type. */
type Rational = struct {
    Huge numerator, denominator; /* illegal */
};
struct Pair {
    Huge first, second;          /* legal */
};
```

Without this restriction, Cforall might require “module initialization” code (since `Rational` has external linkage, it must be created before any other translation unit instantiates it), and would force an ordering on the initialization of the translation unit that defines `Huge` and the translation that declares `Rational`.

A benefit of the restriction is that it prevents the declaration in separate translation units of types that contain each other, which would be hard to prevent otherwise.

File a.c:

```
extern type t1;
type t2 = struct { t1 f1; /* ... */ } /* illegal */
```

File b.c:

```
extern type t2;
type t1 = struct { t2 f2; /* ... */ } /* illegal */
```

□

□ Since a *type-declaration* is a *declaration* and not a *struct-declaration*, type declarations can not be structure members. The form of *type-declaration* forbids arrays of, pointers to, and functions returning `type`. Hence the syntax of *type-specifier* does not have to be extended to allow type-valued expressions. It also side-steps the problem of type-valued expressions producing different values in different declarations.

Since a type declaration is not a *parameter-declaration*, functions can not have explicit type parameters. This may be too restrictive, but it attempts to make compilation simpler. Recall that when traditional C scanners read in an identifier, they look it up in the symbol table to determine whether or not it is a typedef name, and return a “type” or “identifier” token depending on what they find. A type parameter would add a type name to the current scope. The scope manipulations involved in parsing the declaration of a function that takes function pointer parameters and returns a function pointer may just be too complicated.

Explicit type parameters don’t seem to be very useful, anyway, because their scope would not include the return type of the function. Consider the following attempt to define a type-safe memory allocation function.

```

#include <stdlib.h>
T* new(type T) { return (T*) malloc(sizeof(T)); };
...
int* ip = new(int);

```

This looks sensible, but Cforall’s declaration-before-use rules mean that “T” in the function body refers to the parameter, but the “T” in the return type refers to the meaning of T in the scope that contains `new`; it could be undefined, or a type name, or a function or variable name. Nothing good can result from such a situation. \square

Examples

Since type declarations create new types, instances of types are always passed by value.

```

type A1 = int[2];
void f1(A1 a) { a[0] = 0; };
typedef int A2[2];
void f2(A2 a) { a[0] = 0; };
A1 v1;
A2 v2;
f1(v1);
f2(v2);

```

V1 is passed by value, so `f1()`’s assignment to `a[0]` does not modify `v1`. V2 is converted to a pointer, so `f2()` modifies `v2[0]`.

A translation unit containing the declarations

```

extern type Complex; /* opaque type declaration. */
extern float abs(Complex);

```

can contain declarations of complex numbers, which can be passed to `abs`. Some other translation unit must implement `Complex` and `abs`. That unit might contain the declarations

```

type Complex = struct { float re, im; };
Complex cplx_i = {0.0, 1.0};
float abs(Complex c) {
    return sqrt(c.re*c.re + c.im*c.im);
}

```

Note that `c` is implicitly converted to a `struct` so that its components can be retrieved.

```

type Time_of_day = int; /* seconds since midnight. */
Time_of_day ?+?(Time_of_day t1, int seconds) {
    return ((int)t1 + seconds)%86400;
}

```

`t1` must be cast to its implementation type to prevent infinite recursion.

\square Within the scope of a type definition, an instance of the type can be viewed as having that type or as having the implementation type. In the `Time_of_day` example, the difference is important. Different languages have treated the distinction between the abstraction and the implementation in different ways.

- Inside a Clu cluster [6], the declaration of an instance states which view applies. Two primitives called `up` and `down` can be used to convert between the views.

- The Simula class [7] is essentially a record type. Since the only operations on a record are member selection and assignment, which can not be overloaded, there is never any ambiguity as to whether the abstraction or the implementation view is being used. In C++ [5], operations on class instances include assignment and “&”, which can be overloaded. A “scope resolution” operator can be used inside the class to specify whether the abstract or implementation version of the operation should be used.
- An Ada derived type definition [9] creates a new type from an old type, and also implicitly declares derived subprograms that correspond to the existing subprograms that use the old type as a parameter type or result type. The derived subprograms are clones of the existing subprograms with the old type replaced by the derived type. Literals and aggregates of the old type are also cloned. In other words, the abstract view provides exactly the same operations as the implementation view. This allows the abstract view to be used in all cases. The derived subprograms can be replaced by programmer-specified subprograms. This is an exception to the normal scope rules, which forbid duplicate definitions of a subprogram in a scope. In this case, explicit conversions between the derived type and the old type can be used.

Cforall’s rules are like Clu’s, except that implicit conversions and conversion costs allow it to do away with most uses of `up` and `down`. □

6.7.12.1 Default functions and objects

A declaration of a type identifier `T` with type-class `type` implicitly declares a *default assignment* function `T` \rightarrow `?(T*, T)`, with the same scope and linkage as the identifier `T`.

□ Assignment is central to C’s imperative programming style, and every existing C object type has assignment defined for it (except for array types, which are treated as pointer types for purposes of assignment). Without this rule, nearly every inferred type parameter would need an accompanying assignment assertion parameter. If a type parameter should not have an assignment operation, `dtype` should be used. If a type should not have assignment defined, the user can define an assignment function that causes a run-time error, or provide an external declaration but no definition and thus cause a link-time error. □

A definition of a type identifier `T` with implementation type `I` and type-class `type` implicitly defines a default assignment function. A definition of a type identifier `T` with implementation type `I` and an assertion list implicitly defines *default functions* and *default objects* as declared by the assertion declarations. The default objects and functions have the same scope and linkage as the identifier `T`. Their values are determined as follows:

- If at the definition of `T` there is visible a declaration of an object with the same name as the default object, and if the type of that object with all occurrence of `I` replaced by `T` is compatible with the type of the default object, then the default object is initialized with that object. Otherwise the scope of the declaration of `T` must contain a definition of the default object.
- If at the definition of `T` there is visible a declaration of a function with the same name as the default function, and if the type of that function with all occurrence of `I` replaced by `T` is compatible with the type of the default function, then the default function calls that function after converting its arguments and returns the converted result.

Otherwise, if `I` contains exactly one anonymous member such that at the definition of `T` there is visible a declaration of a function with the same name as the default function, and the type of that function with all occurrences of the anonymous member’s type in its parameter list replaced by `T` is compatible

with the type of the default function, then the default function calls that function after converting its arguments and returns the result.

Otherwise the scope of the declaration of `T` must contain a definition of the default function.

□ Note that a pointer to a default function will not compare as equal to a pointer to the inherited function. □

A function or object with the same type and name as a default function or object that is declared within the scope of the definition of `T` replaces the default function or object.

Examples

```
spec s(type T) {
    T a, b;
}
struct impl { int left, right; } a = { 0, 0 };
type Pair | s(Pair) = struct impl;
Pair b = { 1, 1 };
```

The definition of `Pair` implicitly defines two objects `a` and `b`. `Pair a` inherits its value from the `struct impl a`. The definition of `Pair b` is compulsory because there is no `struct impl b` to construct a value from.

```
spec ss(type T) {
    T clone(T);
    void munge(T*);
}
type Whatsit | ss(Whatsit);
type Doodad | ss(Doodad) = struct doodad {
    Whatsit; /* anonymous member */
    int extra;
};
Doodad clone(Doodad) { /* ... */ }
```

The definition of `Doodad` implicitly defines three functions:

```
Doodad ?=(Doodad*, Doodad);
Doodad clone(Doodad);
void munge(Doodad*);
```

The assignment function inherits `struct doodad`'s assignment function because the types match when `struct doodad` is replaced by `Doodad` throughout. `munge()` inherits `Whatsit`'s `munge()` because the types match when `Whatsit` is replaced by `Doodad` in the parameter list. `clone()` does *not* inherit `Whatsit`'s `clone()`: replacement in the parameter list yields "`Whatsit clone(Doodad)`", which is not compatible with `Doodad`'s `clone()`'s type. Hence the definition of "`Doodad clone(Doodad)`" is necessary.

Default functions and objects are subject to the normal scope rules.

```
type T = ...;
T a_T = ...; /* Default assignment used. */
T ?=(T*, T);
T a_T = ...; /* Programmer-defined assignment called. */
```

- A compiler warning would be helpful in this situation. □
- The *class* construct of object-oriented programming languages performs three independent functions. It *encapsulates* a data structure; it defines a *subtype* relationship, whereby instances of one class may be used in contexts that require instances of another; and it allows one class to *inherit* the implementation of another.
 - In Cforall, encapsulation is provided by opaque types and the scope rules, and subtyping is provided by specifications and assertions. Inheritance is provided by default functions and objects. □

6.8 Statements and blocks

Many statements contain expressions, which may have more than one interpretation. The following sections describe how the Cforall translator selects an interpretation. In all cases the result of the selection shall be a single unambiguous interpretation.

6.8.3 Expression and null statements

The expression in an expression statement is treated as being cast to void.

6.8.4 Selection statements

The controlling expression e in the switch statement

```
switch ( e ) ...
```

may have more than one interpretation, but it shall have only one interpretation with an integral type. An integer promotion is performed on the expression if necessary. The constant expressions in **case** statements with the switch are converted to the promoted type.

6.8.5 Iteration statements

The controlling expression e in the loops

```
if ( e ) ...
while ( e ) ...
do ... while ( e );
```

is treated as “(int)((e)!=0)”.

The statement

```
for ( a; b; c ) ...
```

is treated as

```
for ((void)(a); (int)((b)!=0); (void)(c)) ...
```

6.8.6 Jump statements

An expression in a **return** statement is treated as being cast to the result type of the function.

6.10 Preprocessing directives

6.10.8 Predefined macro names

The implementation shall define the macro names `__LINE__`, `__FILE__`, `__DATE__`, and `__TIME__`, as in the ISO/IEC C standard. It shall not define the macro name `__STDC__`.

In addition, the implementation shall define the macro name `__CFORALL__` to be the decimal constant 1.

Appendix A

Examples

A.1 C types

This section gives example specifications for some groups of types that are important in the C language, in terms of the predefined operations that can be applied to those types.

A.1.1 Scalar, arithmetic, and integral types

The pointer, integral, and floating-point types are all *scalar types*. All of these types can be logically negated and compared. The assertion “`scalar(Complex)`” should be read as “type `Complex` is scalar”.

```
spec scalar(type T) {
    int !?(T);
    int ?<?(T, T), ?<=? (T, T), ?==?(T, T), ?>=? (T, T), ?>?(T, T), ?!=?(T, T);
};
```

The integral and floating-point types are *arithmetic types*, which support the basic arithmetic operators. The use of an assertion in the *spec-parameter-list* declares that, in order to be arithmetic, a type must also be scalar (and hence that scalar operations are available). This is equivalent to inheritance of specifications.

```
spec arithmetic(type T | scalar(T) ) {
    T +?(T),      -?(T);
    T ?*?(T, T), ?/?(T, T),  ?+?(T, T),  ?-?(T, T);
};
```

The various flavors of `char` and `int` and the enumerated types make up the *integral types*.

```
spec integral(type T | arithmetic(T) ) {
    T ~?(T);
    T ?&?(T, T),  ?|?(T, T),  ?^?(T, T);
    T ?%?(T, T);
    T ?<<?(T, T), ?>>?(T, T);
};
```

A.1.2 Modifiable types

The only operation that can be applied to all modifiable lvalues is simple assignment.

```
spec m_lvalue(type T) {
    T ?=(T*, T);
};
```

Modifiable scalar lvalues are scalars and are modifiable lvalues, and assertions in the *spec-parameter-list* reflect those relationships. This is equivalent to multiple inheritance of specifications. Scalars can also be incremented and decremented.

```
spec m_l_scalar(type T | scalar(T) | m_lvalue(T) ) {
    T ?++(T*), ?--(T*);
    T ++?(T*), --?(T*);
};
```

Modifiable arithmetic lvalues are both modifiable scalar lvalues and arithmetic. Note that this results in the “inheritance” of `scalar` along both paths.

```
spec m_l_arithmetic(type T | m_l_scalar(T) | arithmetic(T) ) {
    T ?/=(T*, T), ?*=(T*, T);
    T ?+=(T*, T), ?-=(T*, T);
};

spec m_l_integral(type T | m_l_arithmetic(T) | integral(T) ) {
    T ?&=(T*, T), ?|=(T*, T), ?^=(T*, T);
    T ?%=(T*, T), ?<=(T*, T), ?>=(T*, T);
};
```

A.1.3 Pointer and array types

Array types can barely be said to exist in ISO/IEC C, since in most cases an array name is treated as a constant pointer to the first element of the array, and the subscript expression “`a[i]`” is equivalent to the dereferencing expression “`*(a+i)`”. Technically, pointer arithmetic and pointer comparisons other than “`==`” and “`!=`” are only defined for pointers to array elements, but the type system does not enforce those restrictions. Consequently, there is no need for a separate “array type” specification.

Pointer types are scalar types. Like other scalar types, they have “`+`” and “`-`” operators, but the types do not match the types of the operations in `arithmetic`, so these operators cannot be consolidated in `scalar`.

```
spec pointer(type P | scalar(P) ) {
    P      ?+(P, long int), ?+(long int, P), ?-(P, long int);
    ptrdiff_t ?-(P, P);
};

spec m_l_pointer(type P | pointer(P) | m_l_scalar(P) ) {
    P      ?+=(P*, long int), ?-=(P*, long int);
    P      ?=(P*, void*);
    void* ?=(void**, P);
};
```

Specifications that define the dereference operator (or subscript operator) require two parameters, one for the pointer type and one for the pointed-at (or element) type. Different specifications are needed for each set of type qualifiers, because qualifiers are not included in types. The assertion “`|ptr_to(Safe_pointer, int)`” should be read as “`Safe_pointer` acts like a pointer to `int`”.

```

spec ptr_to(type P | pointer(P), type T) {
    lvalue T *(P); lvalue T ?[?](P, long int);
};

spec ptr_to_const(type P | pointer(P), type T) {
    const lvalue T *(P); const lvalue T ?[?](P, long int);
};

spec ptr_to_volatile(type P | pointer(P), type T) {
    volatile lvalue T *(P); volatile lvalue T ?[?](P, long int);
};

spec ptr_to_const_volatile(type P | pointer(P), type T) {
    const volatile lvalue T *(P);
    const volatile lvalue T ?[?](P, long int);
};

```

Assignment to pointers is more complicated than is the case with other types, because the target's type can have extra type qualifiers in the pointed-at type: a “T*” can be assigned to a “const T*”, a “volatile T*”, and a “const volatile T*”. Again, the pointed-at type is passed in, so that assertions can connect these specifications to the “ptr_to” specifications.

```

spec m_l_ptr_to(type P | m_l_pointer(P),
                type T | ptr_to(P,T) {
    P    ?=(P*, T*);
    T*   ?=(T**, P);
};

spec m_l_ptr_to_const(type P | m_l_pointer(P),
                      type T | ptr_to_const(P,T) {
    P    ?=(P*, const T*);
    const T* ?=(const T**, P);
};

spec m_l_ptr_to_volatile(type P | m_l_pointer(P),
                         type T | ptr_to_volatile(P,T) {
    P    ?=(P*, volatile T*);
    volatile T* ?=(volatile T**, P);
};

spec m_l_ptr_to_const_volatile(
    type P | ptr_to_const_volatile(P),
    type T | m_l_ptr_to_volatile(P,T) | m_l_ptr_to_const(P)) {
    P    ?=(P*, const volatile T*);
    const volatile T* ?=(const volatile T**, P);
};

```

Note the regular manner in which type qualifiers appear in those specifications. An alternative specification can make use of the fact that qualification of the pointed-at type is part of a pointer type to capture that regularity.

```

spec m_l_ptr_like(type MyP | m_l_pointer(MyP),
                  type CP | m_l_pointer(CP) ) {
    MyP ??(MyP*, CP);
    CP  ??(CP*, MyP);
};

```

The assertion “`| m_l_ptr_like(Safe_ptr, const int*)`” should be read as “`Safe_ptr` is a pointer type like `const int*`”. This specification has two defects, compared to the original four: there is no automatic assertion that dereferencing a `MyP` produces an lvalue of the type that `CP` points at, and the “`|m_l_pointer(CP)`” assertion provides only a weak assurance that the argument passed to `CP` really is a pointer type.

A.2 Relationships between operations

Different operators often have related meanings; for instance, in C, “+”, “+=”, and the two versions of “++” perform variations of addition. Languages like C++ and Ada allow programmers to define operators for new types, but do not require that these relationships be preserved, or even that all of the operators be implemented. Completeness and consistency is left to the good taste and discretion of the programmer. It is possible to encourage these attributes by providing generic operator functions, or member functions of abstract classes, that are defined in terms of other, related operators.

In Cforall, polymorphic functions provide the equivalent of these generic operators, and specifications explicitly define the minimal implementation that a programmer should provide. This section shows a few examples.

A.2.1 Relational and equality operators

The different comparison operators have obvious relationships, but there is no obvious subset of the operations to use in the implementation of the others. However, it is usually convenient to implement a single comparison function that returns a negative integer, 0, or a positive integer if its first argument is respectively less than, equal to, or greater than its second argument; the library function `strcmp` is an example.

C and Cforall have an extra, non-obvious comparison operator: “!”, logical negation, returns 1 if its operand compares equal to 0, and 0 otherwise.

```

spec comparable(type T) {
    const T 0;
    int      compare(T, T);
}

forall(type T | comparable(T)) int ?<?(T l, T r) {
    return compare(l,r) < 0;
}
/* ... similarly for <=, ==, >=, >,
   and !=. */

forall(type T | comparable(T)) int !?(T operand) {
    return !compare(operand, 0);
}

```

A.2.2 Arithmetic and integer operations

A complete arithmetic type would provide the arithmetic operators and the corresponding assignment operators. Of these, the assignment operators are more likely to be implemented directly, because it is usually more efficient to alter the contents of an existing object than to create and return a new one. Similarly, a complete integral type would provide integral operations based on integral assignment operations.

```

spec arith_base(type T) {
    const T 1;
    T      ?+=?(T*,T), ?-=?(T*,T), ?*=? (T*,T), ?/=?(T*,T);
}

forall(type T | arith_base(T)) T ?+?(T l, T r) {
    return l += r;
}

forall(type T | arith_base(T)) T ?++(T* operand) {
    T temporary = *operand;
    *operand += 1;
    return temporary;
}

forall(type T | arith_base(T)) T ++?(T* operand) {
    return *operand += 1;
}
/* ... similarly for -, --, *,and /. */

spec int_base(type T) {
    T ?&=? (T*, T), ?|=?(T*, T), ?^=? (T*, T);
    T ?%=? (T*, T), ?<<=? (T*, T), ?>>=? (T*, T);
}

forall(type T | int_base(T)) T ?&&?(T l, T r) {
    return l &= r;
}
/* ... similarly for |, ^, %, <<,and >>. */

```

Note that, although an arithmetic type would certainly provide comparison functions, and an integral type would provide arithmetic operations, there does not have to be any relationship among `int_base`, `arith_base` and `comparable`. Note also that these declarations provide guidance and assistance, but they do not define an absolutely minimal set of requirements. A truly minimal implementation of an arithmetic type might only provide 0, 1, and `?-=?`, which would be used by polymorphic `?+=?`, `?*=?`, and `?/=?` functions.

Note also that `short` is an integer type in C11 terms, but has no operations!

Appendix B

TODO

Review index entries.

Restrict allowed to qualify anything, or type/dtype parameters, but only affects pointers. This gets into `noalias` territory. Qualifying anything (“`short restrict rs`”) means pointer parameters of `?++`, etc, would need restrict qualifiers.

Enumerated types. Constants are not ints. Overloading. Definition should be “representable as an integer type”, not “as an int”. C11 usual conversions freely convert to and from ordinary integer types via assignment, which works between any integer types. Does `enum Color ?*(enum Color, enum Color)` really make sense? `?++` does, but it adds `(int)1`.

Operators on `,signed,unsigned char` and other small types. `?|?` harmless; `?*?` questionable for chars. Generic selections make these choices visible. Safe conversion operators? Predefined “promotion” function? `register` assignment might be handled as assignment to a temporary with copying back and forth, but copying must not be done by assignment.

Don’t use `ptrdiff_t` by name in the predefineds.

Bibliography

- [1] American National Standards Institute, 25 West 43rd Street, New York, New York 10036. *American National Standard Information Systems – Programming languages – C*, May 2012. INCITS/ISO/IEC 9899-2011[2012].
- [2] T. P. Baker. A one-pass algorithm for overload resolution in Ada. *ACM Transactions on Programming Languages and Systems*, 4(4):601–614, October 1982.
- [3] G. V. Cormack and A. K. Wright. Type-dependent parameter inference. *SIGPLAN Notices*, 25(6):127–136, June 1990. Proceedings of the ACM Sigplan’90 Conference on Programming Language Design and Implementation June 20-22, 1990, White Plains, New York, U.S.A.
- [4] Glen Ditchfield. An overview of cforall. in preparation, 1996.
- [5] Margaret A. Ellis and Bjarne Stroustrup. *The Annotated C++ Reference Manual*. Addison Wesley, first edition, 1990.
- [6] Barbara Liskov, Russell Atkinson, Toby Bloom, Eliot Moss, J. Craig Schaffert, Robert Scheifler, and Alan Snyder. *CLU Reference Manual*, volume 114 of *Lecture Notes in Computer Science*. Springer-Verlag, 1981.
- [7] Standardiseringskommissionen i Sverige. *Databehandling – Programspråk – SIMULA*, 1987. Svensk Standard SS 63 61 14.
- [8] Ken Thompson. A new c compiler. In *Proceedings of the Summer 1990 UKUUG Conference*, pages 41–51, 1990.
- [9] United States Department of Defense. *The Programming Language Ada: Reference Manual*, ANSI/MIL-STD-1815A-1983 edition, February 1983. Published by Springer-Verlag.

Index

Italic page numbers give the location of the main entry for the referenced term. Plain page numbers denote uses of the indexed term. Entries for grammar non-terminals are italicized. A typewriter font is used for grammar terminals and program identifiers.

- !?, 8, 22
- *?, 8, 22
- ++?, 8, 22
- +?, 8, 22
- ?, 8, 22
- ?, 8, 22
- ?!=?, 8, 38
- ?%=?, 8, 45
- ?%?, 8, 31
- ?&=?, 8, 45
- ?&?, 8, 40
- ?(), 8, 12
- ?*=?, 8, 45
- ?*?, 8, 31, 64
- ?++, 8, 12
- ?+=?, 8, 45, 64
- ?+?, 8, 33, 68
- ?--, 8, 12
- ?-=?, 8, 45
- ?-?, 8, 33
- ?/=?, 8, 45
- ?/? , 8, 31
- ?:, 43
- ?<<=?, 8
- ?<<?, 8
- ?<=?, 8
- ?<?, 8
- ?==?, 8, 38
- ?=?, 8, 45, 69
- ?>=?, 8, 37
- ?>>=?, 8, 45
- ?>>?, 8, 36
- ?>?, 8, 37
- ?[?] , 8, 12
- ?^=?, 8, 45
- ?~?, 8, 41
- ?|=?, 8, 45
- ?|? , 8
- restrict, 61
- _Alignof, 30
- register, 20
- ~?, 8, 22

- 0, 7, 9--11, 15, 42, 64
- 1, 7, 10

- abs, 68
- abs(Complex), 68
- additive-expression*, 33, 33, 36
- alloc, 61, 61
- ambiguous interpretation, 9, 9, 11, 13, 15, 27, 30, 57
- AND-expression*, 40, 40, 41
- anonymous member, 69
- anonymous structure, 4, 6, 58
- anonymous union, 4, 6, 58
- argument-expression-list*, 11, 11
- arithmetic, 72, 72, 73
- arithmetic types, 72
- assertion, 63
- assertion*, 63, 63
- assertion parameters, 5, 13, 59, 63
- assertion-list*, 63, 63, 64, 65
- assignment-expression*, 11, 45, 45, 57
- assignment-operator*, 45, 45

- best valid interpretations, 9, 11, 13
- block, 65
- block scope, 59, 65

- cast expression, 9, 30, 61
- cast-expression*, 22, 30, 30, 31
- color_point, 58
- comma expression, 8
- compatible type, 3, 5, 13, 59
- Complex, 68, 68
- complex type, 4, 32

- composite type, 13, 64
- conditional-expression*, 42, 42, 45
- const**, 61
- const-qualified, 62
- constant*, 10, 11
- constant identifiers, 7
- constant-expression*, 7, 65, 66
- conversion cost, 6, 9, 32
- cplx_i*, 68
- declaration*, 57, 59, 67
- declaration scope, 59
- declaration-specifiers*, 58
- declarator*, 63
- declarator-list*, 63, 63
- default argument promotions, 13
- default assignment, 69
- default functions, 69
- default objects, 69
- deficiencies
 - generalizability, 14
 - nesting opaque types, 66
 - pointers to bit-fields, 21
 - pointers to qualified types, 12, 21, 66
 - static opaque objects, 66
 - union assignment, 51
 - void* conversion, 6
- direct safe conversion, 6
- Doodad, 70
- dtype**, 65
- eat_double**, 29, 32
- equality-expression*, 38, 38, 40
- exclusive-OR-expression*, 41, 41
- explicit conversions, 4
- expression*, 10, 11, 42, 57, 57
- extern**, 65
- external linkage, 3
- file scope, 10
- forall-specifier*, 57, 58, 58, 59
- ftype**, 65
- function designator, 13
- function types, 63, 65, 66
- generalizability, 13
- generic-selection*, 10
- identifier*, 7, 10, 11, 63–65
- identifiers
 - for constants, 7
 - for operators, 8
- implementation type, 65, 69
- implicit arguments, 13
- implicit conversions, 4, 4, 65
- inclusive-OR-expression*, 41, 41, 42
- incomplete types, 63, 65, 66
- inferred parameter, 5, 30, 59
- initializer, 65
- initializer*, 62
- initializer-list*, 11, 62
- integer conversion rank, 4, 28, 32, 35, 36, 38, 40, 41
- integer promotion, 36, 71
- integer promotions, 4, 32
- integral**, 72, 73
- integral types, 72
- internal linkage, 3, 10
- interpretations, 9, 27, 30, 62, 71
- keyword*, 7
- less polymorphic, 5, 9, 15, 59
- linkage, 3, 69
- list_of*, 64
- literal**, 58
- logical-AND-expression*, 42, 42
- logical-OR-expression*, 42, 42
- lvalue, 61
- lvalue*, 61
- m_l_arithmetic*, 73, 73
- m_l_integral*, 73
- m_l_pointer*, 73, 74, 75
- m_l_ptr_like*, 75
- m_l_ptr_to*, 74
- m_l_ptr_to_const*, 74, 74
- m_l_ptr_to_const_volatile*, 74
- m_l_ptr_to_volatile*, 74, 74
- m_l_scalar*, 73, 73
- m_lvalue*, 73, 73
- min_max**, 15
- modifiable lvalue, 61, 72
- monomorphic function, 13, 59
- move_by**, 5
- multiplicative-expression*, 31, 31, 33
- name spaces, 3
- no linkage, 3, 57, 59, 65
- null pointer, 9, 15, 40
- object types, 63, 65, 66

- opaque type declaration, 65
- operator identifiers, 8, 10
- overloading, 3, 7
- parameter-declaration*, 67
- parameter-type-list*, 59
- point, 58
- pointer, 73, 74
- polymorphic function, 13, 59
- postfix-expression*, 11, 11, 13, 22
- primary-expression*, 10, 11
- ptr_to*, 74, 74
- ptr_to_const*, 74, 74
- ptr_to_const_volatile*, 74, 74
- ptr_to_volatile*, 74, 74
- ptrdiff_t*, 12
- qualified type, 6, 66
- rank, *see* integer conversion rank
- Rational**, 30, 42
- real type, 4, 32
- relational-expression*, 36, 36, 38
- rewrite rules, 10
- safe arithmetic conversions, 4
- safe conversions, 6, 13
- scalar**, 72, 72, 73
- scalar types, 72
- scopes, 3, 69
- shift-expression*, 36, 36
- sizeof**, 30
- spec-declaration*, 63, 63
- spec-declaration-list*, 63, 63, 65
- spec-definition*, 57, 63
- spec-parameter-list*, 63, 72, 73
- specialization, 11, 15
- specification, 59, 63
- specification definition, 63
- specifier-qualifier-list*, 63
- square**, 64
- static**, 65
- storage-class specifier, 20
- storage-class specifiers, 65
- storage-class-specifier*, 65
- strchr**, 66
- string-literal*, 10, 11
- struct-declaration*, 67
- sum_list**, 64
- summable**, 64
- Time_of_day**, 68
- translation unit, 65
- type**, 3, 13, 63, 65
- type declaration, 6, 69
- type definition, 65, 69
- type names, 65
- type qualifiers, 73
- type-class, 13, 66
- type-class*, 64, 65
- type-declaration*, 30, 57, 65, 67
- type-declarator*, 65, 65
- type-declarator-list*, 65, 65
- type-name*, 11, 22, 30, 63–66
- type-name-list*, 63, 63
- type-parameter*, 30, 63, 64, 64, 65
- type-parameter-list*, 58, 59, 63, 64, 64
- type-qualifier*, 61
- type-specifier*, 57, 67
- typedef**, 3
- unary-expression*, 22, 22, 30, 45
- unary-operator*, 22, 22
- unsafe conversions, 6, 6, 9, 14
- usual arithmetic conversions, 4, 32, 44
- valid interpretations, 9, 11, 13
- visible, 3, 11
- volatile**, 61
- Whatsit**, 70