

5

Complex Type Simplified WG14/N450 (X3J11/95-051)

Tom MacDonald

David Knaak

Cray Research, Inc.

655F Lone Oak Drive

Eagan, MN 55121

tam@cray.com knaak@cray.com

23 August 1995

10

Introduction

15 This document is a full specification for a proposed extension to Standard C that adds complex data types. The specification is given as modifications to document **WG14/N412** (a.k.a. **X3J11/95-013**) **C9X Draft 2**, and will track the progress of the C9X Draft specification.

This document also includes rationale, summarizing some of the discussions related to this proposed extension.

20 This document, "Complex Type Simplified," focuses on basic language issues for complex types and the basics of the complex elementary math functions. It does not specify behavior for the special values of signed zeros, NaNs, and infinities for those implementations that support them.

25 This document does not include imaginary types as part of the proposed extension. The reasons for this are discussed in the rationale part of section **6.1.2.5**.

Sections that can be regarded as extending or modifying a particular section of C9X Draft 2 are marked with the section numbers.

Changes

The following major changes appear in this document:

- 30
- Complex constants can only occur after **<complex.h>** is included.
 - **6.3.9** description changed slightly.

5.2.4.2 Numerical limits

5.2.4.2.2 Characteristics of floating types `<float.h>`

Add after the first sentence:

- 5 The characteristics of each of the real and imaginary parts of `float complex`, `double complex`, and `long double complex` numbers are the same as for the corresponding real floating types.

RATIONALE:

- 10 The specification that the parts of the complex floating types have the same characteristics as the corresponding real floating types, including their numerical limits, implies that the underlying implementation of the complex floating types is to be Cartesian. With this requirement, the limits for the complex plane are clearly and easily defined in terms of the limits for real floating types. See rationale with section 6.1.2.5 for further discussion.

6.1.1 Keywords

Add to the list of keywords:

- 15 **complex**

Semantics

The token `complex` becomes a keyword when the header `<complex.h>` is included, and not before.

Add to the list of forward references:

- 20 **Forward references:** complex mathematics (7.14)

RATIONALE:

- 25 Since some C programs already define the type `complex` as a named type, this identifier only becomes a keyword when the `<complex.h>` header is included. This avoids conflicts with existing usage, and allows implementations supporting freestanding environments to avoid the burden of implementing a complex type. One simple implementation that avoids this conflict is to define a macro inside the header such as:

```
#define complex _Complex
```

- 30 which uses an identifier `_Complex` that is in the implementor's name space. The implementation recognizes `_Complex` as a type. It is possible to move `complex` into the regular list of standard keywords if desired.

6.1.2.5 Types

Replace first sentence of sixth paragraph with:

There are three *real floating types*, designated as `float`, `double`, and `long double`.

Add the following two paragraphs:

- 35 There are three *complex floating types*, designated as `float complex`, `double complex`, and `long double complex`. The real and imaginary parts of the complex floating types each have the same representation as the corresponding real floating types. The set of values of the type `float complex` is a subset of the set of

values of the type **double complex**, the set of values of the type **double complex** is a subset of the set of values of the type **long double complex**.

The integral types and the real floating types are collectively called the *real types*. The real floating types and the complex floating types are collectively called the *floating types*.

5

RATIONALE:

The complex extension to C is intended to provide as direct a mapping as possible of complex mathematics (for example, complex algebra or complex analysis) onto the C language. This mapping should be as close as possible in both how it is expressed and the results that are obtained. Because of the inherent characteristics of the C language and of digital computer systems, this mapping may be less than perfect, but this mapping should still be the guiding principle for the extension.

10

The statement "The real and imaginary parts of the complex floating types each have the same representation as the corresponding real floating types" is the same as saying that the underlying implementation of the complex types is Cartesian. This implementation is explicitly stated so that behaviors can be defined simply and unambiguously. If the underlying representation is not explicit, the description of the underlying implementation, limits, promotion rules, and behavior of math functions is very abstract. Therefore, if one representation is to be selected, Cartesian is the most natural for virtually all modern computers. This specification is consistent with the Fortran 90 specification of complex types. Other topologies, such as polar, are useful for many problems but are not defined by this extension. When needed, users can define additional functions or macros in terms of the underlying Cartesian implementation.

15

20

This extension does not include integral types. However, the addition of complex integral types at some future date would be upward compatible with the current proposal.

25

This extension does not include imaginary types. Inclusion of imaginary types is also being proposed and discussed by the committee. The justification for imaginary types is that these are necessary in order to avoid generation of undesired NaNs, to avoid getting the sign of zero wrong when infinities and minus zeros are part of a complex operand, and to get efficient arithmetic when a binary operator has both a real and a complex operand.

30

35

The imaginary types are not needed because, in the vast majority of complex arithmetic cases, the sign of zero does not matter, and NaNs and infinities provide equal information. Signed zeros typically do not matter because they do not provide the same information with the complex plane that they do with the real number line (i.e., the direction of underflow). NaNs and infinities provide equal information with the complex plane because both indicate something exceptional happen. Since there are many infinities present in the complex plane, infinities do not provide the same information with the complex plane as with the real number line. (i.e., the direction of overflow). Acknowledgment that NaN and infinity values give equal information, and that +0 and -0 provide equal information, allows an implementation to provide efficient operations through standard optimizations. For example:

40

```
float * complex → float*real + float*imag
```

and not necessarily:

```
float * complex → (float*real - 0.0*imag) + (float*imag + 0.0*real)
```

In summary, the costs associated with adding imaginary types are not justified by the benefits gained.

6.1.3 Constants

6.1.3.1 Floating constants

Syntax

- floating-constant*:
- 5 *real-floating-constant*
 complex-floating-constant
- real-floating-constant*:
- fractional-constant* *exponent-part* *real-suffix*_{opt}
 digit-sequence *exponent-part* *real-suffix*_{opt}
- 10 *complex-floating-constant*:
- fractional-constant* *exponent-part* *complex-suffix*_{opt}
 digit-sequence *exponent-part* *complex-suffix*_{opt}
 digit-sequence *imaginary-suffix*
- fractional-constant*:
- 15 *digit-sequence*_{opt} . *digit-sequence*
 digit-sequence .
- exponent-part*:
- e* *sign*_{opt} *digit-sequence*
 E *sign*_{opt} *digit-sequence*
- 20 *sign*: one of
 + -
- digit-sequence*:
- digit*
 digit-sequence *digit*
- 25 *complex-suffix*:
- float-suffix* *imaginary-suffix*
 imaginary-suffix *float-suffix*
 long-suffix *imaginary-suffix*
 imaginary-suffix *long-suffix*
- 30 *imaginary-suffix*
- real-suffix*:
- float-suffix*
 long-suffix
- float-suffix*: one of
- 35 **f** **F**
- long-suffix*: one of
- l** **L**
- imaginary-suffix*:
- i**

Semantics

5 An unsuffixed floating constant has type **double**. If suffixed by the letter **f** or **F**, it has type **float**. If suffixed by the letter **l** or **L**, it has type **long double**. If suffixed by the letter **i** it has type **double complex**. If suffixed by both the letters **i** and **f** or **F**, it has type **float complex**. If suffixed by both the letters **i** and **l** or **L**, it has type **long double complex**. A floating constant that has a complex type specifies the value of the imaginary part, and the "real" part has the value zero. Any floating constant suffixed with **i** is called a *complex floating constant*. Any floating constant not suffixed with **i** is called a *real floating constant*.

A complex floating constant shall not occur before the `<complex.h>` header is included.

10 RATIONALE:

Either **i** or **j** could be selected for the imaginary suffix used to form complex constants. And in fact, both could be allowed. However, allowing both would add unnecessary duplication to the language.

15 The production rule *digit-sequence imaginary-suffix* is purely for notational convenience and is not necessary for completeness. It allows the form **1i** to be used rather than the equivalent **1.i** or **1.0i**.

20 Inclusion of the `<complex.h>` header is not really necessary to define a complex constant because these constants are upward compatible. However, as mentioned earlier, this allows implementations supporting freestanding environments to avoid the burden of implementing a complex type.

6.2 Conversions

6.2.1 Arithmetic operands

Add this section:

6.2.1.4.1 Complex types

25 When a value of complex type is promoted or demoted to another complex type, both parts follow the same promotion and demotion rules as for the corresponding real types.

When a value of real type is promoted to a complex type, the real part of the complex value gets the same value as if the promotion was to the corresponding real type and the imaginary part of the complex value gets the value of positive zero (+0).

30 When a value of complex type is demoted to a real type, the value of the imaginary part of the complex value is discarded and the value of the real part gets demoted according to the demotion rules for the corresponding real type.

6.2.1.5 Usual arithmetic conversions

Replace this section with:

35 All types have three type attributes called the *dimension*, the *format*, and the *length*. The dimension attribute specifies whether the values of the type can be represented on a one dimensional line, (i.e., real numbers) or on a two dimensional plane, (i.e., complex numbers). The format attribute specifies whether the values of the type are represented with an exponent part, (i.e., floating numbers) or without an exponent part, (i.e., integral numbers). The length attribute specifies how many bits are used to represent the magnitude and precision of the type. The values for each of these attributes are ranked, from highest to lowest, as shown below. For example, complex ranks higher than real for the dimension attribute.

40

dimension	format	floating length	integral length
5 complex	floating	long double	unsigned long
10 real	integral	double	signed long
		float	unsigned int
			signed int

15 Many binary operators that have operands of arithmetic types cause implicit conversions of one or both operands. The purpose of the conversions is to yield a common format and length for the two operands which is the type of the result. These implicit conversions of the operands are called the *usual arithmetic conversions*.

20 The conversions shall preserve the original magnitude and precision of both operands except that precision may be lost when an integral type is converted to a floating type. This will occur if the magnitude of the integer is too great for the mantissa of the floating type to represent it exactly.

The rules for the usual arithmetic conversions are:

- 25 1) The dimension of the result type is that of the higher ranking dimension of the operands.
- 2) The format of the result type is that of the higher ranking format of the operands.
- 3) If the format of the result type is floating then:
 - the length of the result type is that of the higher ranking floating length of the operands.
 - 30 else the format of the result type is integral and:
 - the integral promotions are performed on both operands, and the length of the result type is that of the higher ranking integral length of the promoted operands with one exception. The exception is that if one operand has type **signed long** and the other has type **unsigned int** and if a **signed long** cannot represent all the values of an **unsigned int**, the length of the result is **unsigned long**.
- 4) After the result type is determined and before the binary operation is performed, any operand that does not already have the same type as the result type is promoted to the result type.

40 RATIONALE:

The conversion rules spell out an unambiguous path for determining the result type of a binary operator and for specifying any conversions to be done on the operands. However, if it is more efficient in a given implementation to take a different path to get the same result, that is allowed. The primary requirement is that the result must be the same "as if" all the arithmetic conversion rules were followed.

The standard specifies a set of promotion rules such that if the user writes code that assumes these promotions will be done, then the code will be portable across conforming implementations.

50 An implementation is not required to actually convert the operands as long as the result of the binary operation is the same as the result obtained if all the conversion rules were followed. Similarly, operands may be converted when not required as long as the result of the binary operation is the same as the result obtained if all the conversion rules were followed. Since an implementation can assume that NaN and infinity values provide equal information for complex types, optimizations can be performed without regard to these exceptional values.

6.3 Expressions**6.3.2 Postfix operators****6.3.3 Unary operators****6.3.4 Cast operators**5 **6.3.5 Multiplicative operators****6.3.6 Additive operators****6.3.7 Bitwise shift operators****6.3.10 Bitwise AND operator****6.3.11 Bitwise exclusive OR operator**10 **6.3.12 Bitwise inclusive OR operator****6.3.13 Logical AND operator****6.3.14 Logical OR operator****6.3.15 Conditional operator**

No change to the above sections.

15 **RATIONALE:**

A complex number is an arithmetic type and a scalar type. Therefore, all operators that apply to arithmetic and scalar types also apply to complex types. If complex integral types are ever added to the language, then the constraints for some, but not all, of the operators must change ‘integral type’ to ‘real integral type’.

20 The following are mathematical formulas for the basic complex operations:

$$z_1 + z_2 = (a + bi) + (c + di) = (a + c) + (b + d)i$$

$$z_1 - z_2 = (a + bi) - (c + di) = (a - c) + (b - d)i$$

$$z_1 * z_2 = (a + bi) * (c + di) = (ac - bd) + (bc + ad)i$$

$$z_1 \div z_2 = (a + bi) \div (c + di) = [(ac + bd) \div (c^2 + d^2)] + [(bc - bd) \div (c^2 + d^2)]i$$

25 The standard specifies what the result of an operation must be but does not specify how an implementation is to perform the operation. On digital computers with finite range of representable floating point values, operations such as complex division can lead to overflow or underflow of intermediate values even when the mathematical result is within the range of representable values. Unless the standard specifies an algorithm for all implementations to use, whether and when overflow or underflow occurs and the speed of the operation may vary from implementation to implementation.

30 If code is written for implementations that support special values such as signed zeros, NaNs, and infinities, and if the code is written with the assumption that these values could arise during execution of the program, then the user needs to know how that implementation performs these operations, and it is up to the user to decide the meaning of these special values in the context of the program.

6.3.8 Relational operators

Change the first constraint to:

both operands have real type;

40 *Change the first sentence of Semantics to:*

If both of the operands have real type, the usual arithmetic conversions are performed.

RATIONALE:

While equality and inequality have clear mathematical meaning for complex numbers, greater than or less than do not have any standard meaning. A user who wants to design an ordering for complex values, can write a macro or function, similar to the following:

```
5      #define C_LE(z1,z2) (cabs(z1) <= cabs(z2))
```

This gives the user control over the meaning of relational operations with complex operands.

6.3.9 Equality operators*Add to Semantics:*

10 For IEEE implementations, if a NaN appears in either operand when comparing two complex numbers, they compare unequal.

6.5 Declarations**6.5.2 Type specifiers**

Add to the list of type-specifiers:

complex

15 *Add to the set of type-specifiers:*

- **float complex**
- **double complex**
- **long double complex**

6.5.6 Type definitions

20 *Remove or modify the example that gives the name **complex** to a structure type.*

7 Library

Add the following sections:

7.14 Complex mathematics <complex.h>

The header <complex.h> defines four macros and declares several mathematical functions.

- 5 These functions take **double complex** arguments and return **double** or **double complex** values as specified below. This header (as currently defined) must be included to declare complex types or to create a complex constant.

The functions are:

```

10      csin
      ccos
      cexp
      clog
      cpow
15      csqrt
      cabs
      cimag
      conj
      creal

```

The macros are:

```

20      complex

```

which expands into implementation-defined spelling for declaring complex types, and

```

      CMPLXF
      CMPLX
      CMPLXL

```

- 25 which are described in 7.14.1.

7.14.1 Complex operators

7.14.1.1 The **CMPLXF** macro

Synopsis

```

30      #include <complex.h>
      float complex CMPLXF(float x, float y);

```

Description

The **CMPLXF** macro behaves like an operator that creates a value with type **float complex**. If either argument is not of type **float**, it is first converted to type **float**.

Returns

- 35 The **CMPLXF** macro returns a value with type **float complex** whose real part is **x** and whose imaginary part is **y**.

7.14.1.2 The `CMPLEX` macro

Synopsis

```
#include <complex.h>
double complex CMPLEX(double x, double y);
```

5 Description

The `CMPLEX` macro behaves like an operator that creates a value with type `double complex`. If either argument is not of type `double`, it is first converted to type `double`.

Returns

- 10 The `CMPLEX` macro returns a value with type `double complex` whose real part is `x` and whose imaginary part is `y`.

7.14.1.3 The `CMPLEXL` macro

Synopsis

```
15 #include <complex.h>
long double complex CMPLEXL(long double x, long double y);
```

Description

The `CMPLEXL` macro behaves like an operator that creates a value with type `long double complex`. If either argument is not of type `long double`, it is first converted to type `long double`.

20 Returns

The `CMPLEXL` macro returns a value with type `long double complex` whose real part is `x` and whose imaginary part is `y`.

RATIONALE:

- 25 Various methods have been proposed for creating a complex number from two real numbers. One proposed method is to use imaginary constants as in the expression:

$$x + y * 1.0i$$

- 30 This method has the advantage of adding only a small amount of syntax to the language and is a compact form in many expressions. But IEEE implementations may require a way to create complex values where either the real or imaginary parts are special values such as `+`, `-`, and `NaN`, and this method does not always produce the desired complex result. For example, if `y` has the value `+`, and the appropriate optimizations are not performed, then:

```

y * 1.0i
=> (+) * (0.0, 1.0)
=> [(+ * 0.0), (+ * 1.0)]
35 => (NaN, +)
```

when the desired result is:

```
(0.0, +∞)
```

Another proposed method is to use something similar to compound literals as the expression:

```
(double complex){ ._real = x, ._imag = y }
```

5 The full specification of compound literals is defined in document WG14/N357 and X3J11/94-042, "X3J11 Technical Report – Compound Literals." This approach builds upon that specification and assumes that complex numbers are similar to structures with the real and imaginary parts having the names `_real` and `_imag` respectively (although the order is still unimportant). This method could be exploited to replace the `CMPLX*` macros.

10 Yet a third method is an infix operator as in the expression:

```
x ⊗ y
```

Choosing the suitable character(s) for the `⊗` operator is difficult.

By specifying macros to create complex values from real values, each implementation can choose a method most appropriate for that implementation.

15 7.14.2 Complex math functions

Unless otherwise specified, the domain (or region) and range for all complex math functions is the complex plane with the real and imaginary parts defined for all representable values.

An implementation may set `errno` but is not required to.

20 Whenever there are multiple mathematical values for a complex function, (multi-valued complex functions) the return value of the function shall be the principal value of the function. Additional rules may be given for a specific function.

For IEEE implementations, the behavior of these functions, if the input contains a NaN, an ∞ , or a negative zero, is not specified.

The `csin` function

25 Synopsis

```
#include <complex.h>
double complex csin(double complex z);
```

Description

30 The `csin` function selects the sine of `z`, where the real part of `z` is regarded as a value in radians.

Returns

The `csin` function returns the sine value.

The ccos function

Synopsis

```
#include <complex.h>
double complex ccos(double complex z);
```

5 Description

The **ccos** function computes the cosine of **z**, where the real part of **z** is regarded as a value in radians.

Returns

The **ccos** function returns the cosine value.

10 The cexp function

Synopsis

```
#include <complex.h>
double complex cexp(double complex z);
```

Description

- 15 The **cexp** function computes the exponential function of **z**, where the imaginary part is regarded as a value in radians. An output error occurs if the magnitude of **z** is too large.

Returns

The **cexp** function returns the exponential value.

20 The clog function

Synopsis

```
#include <complex.h>
double complex clog(double complex z);
```

Description

- 25 The **clog** function computes the natural logarithm of **z**. An output error occurs if the argument is zero.

Returns

- 30 The **clog** function returns the principal value of the natural logarithm with the imaginary part **w** of the result in the range $-\pi < w \leq \pi$. The imaginary part of the result is π only when the real part of the argument is less than zero and the imaginary part of the argument is zero. The sign of the imaginary part of the log is the same as the sign of the imaginary part of **z**.

The cpow function

Synopsis

```
#include <complex.h>
double complex cpow(double complex z1, double complex z2);
```

5 Description

The **cpow** function computes **z1** raised to the power **z2**. An input error occurs if **z1** is zero and **z2** is not an integral value. An output error occurs if the magnitude of **z1** or **z2** is too large.

Returns

- 10 The **cpow** function returns the value of **z1** raised to the power **z2**. For **cpow(0.0,0.0)** the return value is **1.0 + 0.0i**.

The csqrt function

Synopsis

```
#include <complex.h>
15 double complex csqrt(double complex z);
```

Description

The **csqrt** function computes the square root of **z**.

Returns

- 20 The **csqrt** function returns the principal value with the real part greater than or equal to zero. If the real part of the result is not zero, then the sign of the imaginary part of the root is the same as the sign of the imaginary part of **z**. If the real part of the result is zero, then the imaginary part is greater than or equal to zero.

The cabs function

Synopsis

```
25 #include <complex.h>
double cabs(double complex z);
```

Description

The **cabs** function computes the magnitude of a double complex number **z**.

Returns

- 30 The **cabs** function returns the value of the magnitude of **z**.

The cimag function

Synopsis

```
#include <complex.h>
double cimag(double complex z);
```

5 Description

The **cimag** function computes the imaginary part of **z**.

Returns

The **cimag** function returns the value of the imaginary part of **z**.

The conj function

10 Synopsis

```
#include <complex.h>
double complex conj(double complex z);
```

Description

The **conj** function computes the conjugate of **z**.

15 Returns

If **z** has the value of $x+yi$, the **conj** function returns the value of $x-yi$.

The creal function

Synopsis

```
20 #include <complex.h>
double creal(double complex z);
```

Description

The **creal** function selects the real part of **z**.

Returns

The **creal** function returns the value of the real part of **z**.

25 RATIONALE:

There is no explicit mention of **errno** in the descriptions of the complex math functions. Section 4.1.3 of the standard states:

30 The value of **errno** is zero at program startup, but is never set to zero by any library function. The value of **errno** may be set to nonzero by a library function call whether or not there is an error, provided the use of **errno** is not documented in the description of the function in the standard.

None of the complex math functions are required to use **errno** as an error indicator. Therefore, whether **errno** is changed by any of these functions, is unspecified.

35 The terms “input” and “output” errors are used instead of “domain” and “range” errors so as to *not* imply that **errno** is set to either EDOM or ERANGE.

The type **double complex** was chosen for the complex functions to correspond to type **double** for the functions in `<math.h>`. More functions can be added to the list if there is sufficient demand.

- 5 No specific proposal is made here for handling complex numbers by the **printf** and **scanf** function families. Explicit use of **cimag** and **creal** are sufficient to make these functions perform with complex numbers.