C∀ (Cforall) User Manual
Version 1.0

"describe not prescribe"

C∀ Team (past and present)

Andrew Beach, Richard Bilson, Michael Brooks, Peter A. Buhr, Thierry Delisle,
Glen Ditchfield, Rodolfo G. Esteves, Jiada Liang, Aaron Moss, Colby Parsons
Rob Schluntz, Fangren Yu, Mubeen Zulfiqar

April 14, 2025

Contents

1 Introduction

C∀[1] is a modern general-purpose concurrent programming-language, designed as an evolutionary step forward for the C programming language. The syntax of C∀ builds from C and should look immediately familiar to C/C++ programmers. C∀ adds many modern features that directly lead to increased *safety* and *productivity*, while maintaining interoperability with existing C programs and achieving similar performance. Like C, C∀ is a statically typed, procedural (non-object-oriented) language with a low-overhead runtime, meaning there is no global garbage-collection, but regional garbage-collection is possible. The primary new features include polymorphic routines and types, exceptions, concurrency, and modules.

One of the main design philosophies of C∀ is to "describe not prescribe", which means C∀ tries to provide a pathway from low-level C programming to high-level C∀ programming, but it does not force programmers to "do the right thing". Programmers can cautiously add C∀ extensions to their C programs in any order and at any time to incrementally move towards safer, higher-level programming. A programmer is always free to reach back to C from C∀, for any reason, and in many cases, new C∀ features can be locally switched back to their C counterpart. There is no notion or requirement for *rewriting* a legacy C program to C∀; instead, a programmer evolves a legacy program into C∀ by incrementally incorporating C∀ features. As well, new programs can be written in C∀ using a combination of C and C∀ features. In many ways, C∀ is to C as Scala [29] is to Java, providing a vehicle for new typing and control-flow capabilities on top of a highly popular programming language allowing immediate dissemination.

C++ [30] had a similar goal 30 years ago, allowing object-oriented programming to be incrementally added to C. However, C++ currently has the disadvantages of a strong object-oriented bias, multiple legacy design-choices that are difficult to update, and active divergence of the language model from C, requiring significant effort and training to incrementally add C++ to a C code-base. In contrast, C∀ has 30 years of hindsight and a clean starting point.

Like C++, there may be both old and new ways to achieve the same effect. For example, the following programs compare the C, C∀, and C++ I/O mechanisms, where the programs output the same result.

C	C∀	C++
#include <stdio.h>	**#include** <fstream.hfa>	**#include** <iostream>
		using namespace std;
int main(**void**) {	**int** main(**void**) {	**int** main() {
int x = 0, y = 1, z = 2;	**int** x = 0, y = 1, z = 2;	**int** x = 0, y = 1, z = 2;
printf("%d %d %d\n", x, y, z);	sout | x | y | z;	cout << x << ' ' << y << ' ' << z << endl;
}	}	}

While C∀ I/O (see Section 22, p. 52) looks similar to C++, there are important differences, such as automatic spacing between variables and an implicit newline at the end of the expression list, similar to Python [26]. In general, C∀ programs are 10% to 30% shorter than their equivalent C/C++ counterparts.

1.1 Background

This document is a programmer reference-manual for the C∀ programming language. The manual covers the core features of the language and runtime-system, with simple examples illustrating syntax and semantics of features. The manual does not teach programming, *i.e.*, how to combine the new constructs to build complex programs. The reader must have an intermediate knowledge of control flow, data structures, and concurrency issues to understand the ideas presented, as well as some experience programming in C/C++. Implementers should refer to the C∀ Programming Language Specification for details about the language syntax and semantics. Changes to the syntax and additional features are expected to be included in later revisions.

2 Why fix C?

The C programming language is a foundational technology for modern computing with billions of lines of code implementing everything from hobby projects to commercial operating-systems. This installation base and the programmers producing it represent a massive software-engineering investment spanning decades and likely to continue for decades more. Even with all its problems, C continues to be popular because it allows writing software at virtually any level in a computer system without restriction. For system programming, where direct access to hardware, storage management, and real-time issues are a requirement, C is the only language of choice. The TIOBE index [32] for February

---

[1]Pronounced "C-for-all", and written C∀, CFA, or Cforall.

2023 ranks the top six most *popular* programming languages as C 17.4%, Java 12%, Python 12%, C++ 7.6%, C♯ 4%, Visual Basic 3.8% = 56.8%, where the next 50 languages are less than 2% each, with a long tail. The top 4 rankings over the past 35 years are:

	2023	2018	2013	2008	2003	1998	1993	1988
Python	1	4	8	7	12	25	18	-
C	**2**	**2**	**1**	**2**	**2**	**1**	**1**	**1**
C++	3	3	4	4	3	2	2	4
Java	4	1	2	1	1	18	-	-

Hence, C is still an extremely important programming language, with double the usage of C++; in many cases, C++ is often used solely as a better C. Love it or hate it, C has been an important and influential part of computer science for 40 years and its appeal is not diminishing. Nevertheless, C has many problems and omissions that make it an unacceptable programming language for modern needs.

As stated, the goal of the C∀ project is to engineer modern language-features into C in an evolutionary rather than revolutionary way. C++ [22, 12] is an example of a similar project; however, it largely extended the C language, and did not address many of C's existing problems.[2] Fortran [14], Cobol [6], and Ada [1] are examples of programming languages that took an evolutionary approach, where modern language-features (*e.g.*, objects, concurrency) are added and problems fixed within the framework of the existing language. Java [18], Go [19], Rust [28] and D [3] are examples of the revolutionary approach for modernizing C/C++, resulting in a new language rather than an extension of the descendent. These languages have different syntax and semantics from C, do not interoperate directly with C, and are not systems languages because of restrictive memory-management or garbage collection. As a result, there is a significant learning curve to move to these languages, and C legacy-code must be rewritten. These costs can be prohibitive for many companies with a large software-base in C/C++, and a significant number of programmers require retraining in the new programming language.

The result of this project is a language that is largely backwards compatible with C11 [21], but fixes many of the well known C problems while adding modern language-features. To achieve these goals required a significant engineering exercise, *i.e.*, "thinking *inside* the C box". Considering the large body of existing C code and programmers, there is significant impetus to ensure C is transformed into a modern language. While C11 made a few simple extensions to the language, nothing was added to address existing problems in the language or to augment the language with modern language-features. While some may argue that modern language-features may make C complex and inefficient, it is clear a language without modern capabilities is insufficient for the advanced programming problems existing today.

3   History

The C∀ project started with Dave Till's K-W C [5, 31], which extended C with new declaration syntax, multiple return values from routines, and advanced assignment capabilities using the notion of tuples (see [33] for similar work in C++). The first C∀ implementation of these extensions was by Rodolfo Esteves [13].

The signature feature of C∀ is *overloadable* parametric-polymorphic functions [7, 8, 11] with functions generalized using a **forall** clause (giving the language its name):

```
forall( T ) T identity( T val ) { return val; }
int forty_two = identity( 42 );        // T is bound to int, forty_two == 42
```

C∀'s polymorphism was originally formalized by Glen Ditchfield [9], and first implemented by Richard Bilson [2]. However, at that time, there was little interest in extending C, so work did not continue. As the saying goes, "What goes around, comes around.", and there is now renewed interest in the C programming language because of the legacy code-base, so the C∀ project was restarted in 2015.

4   Interoperability

C∀ is designed to integrate directly with existing C programs and libraries. The most important feature of interoperability is using the same calling conventions, so there is no complex interface or overhead to call existing C routines. This feature allows C∀ programmers to take advantage of the existing panoply of C libraries to access thousands of

---

[2]Two important existing problems addressed were changing the type of character literals from **int** to **char** and enumerator from **int** to the type of its enumerators.

software features. Language developers often state that adequate library support takes more work than designing and implementing the language itself. Fortunately, C∀, like C++, starts with immediate access to all exiting C libraries, and in many cases, can easily wrap library routines with simpler and safer interfaces, at zero or very low cost. Hence, C∀ begins by leveraging the large repository of C libraries, and than allows programmers to incrementally augment their C programs with modern backward-compatible features.

However, it is necessary to differentiate between C and C∀ code because of name overloading, as for C++. For example, the C math-library provides the following routines for computing the absolute value of the basic types: abs, labs, llabs, fabs, fabsf, fabsl, cabsf, cabs, and cabsl. Whereas, C∀ wraps these routines into one overloaded name abs:

```
unsigned char abs( signed char );                  // no C equivalent
extern "C" { int abs( int ); }                      // C abs
unsigned long int abs( long int );                  // C labs
unsigned long long int abs( long long int );        // C llabs
float abs( float );                                 // C fabsf
double abs( double );                               // C fabs
long double abs( long double );                     // C fabsl
float _Complex abs( float _Complex );               // C cabsf
double _Complex abs( double _Complex );             // C cabs
long double _Complex abs( long double _Complex );   // C cabsl
```

The problem is a name clash between the C name abs and the C∀ names abs, resulting in two name linkages: **extern "C"** and **extern "Cforall"** (default). Overloaded names must use *name mangling* to create unique names that are different from unmangled C names. Hence, there is the same need as in C++ to know if a name is a C or C∀ name, so it can be correctly formed. The only way around this problem is C's approach of creating unique names for each pairing of operation and type.

This example illustrates a core idea in C∀: *the power of a name*. The name "abs" evokes the notion of absolute value and many mathematical types provide the notion of absolute value. Hence, knowing the name abs is sufficient to apply it to any applicable type. The time savings and safety of using one name uniformly versus *N* unique names cannot be underestimated.

5   C∀ Compilation

C∀ is a *transpiler*, meaning it reads in a programming language (C∀) as input and generates another programming language (C) as output, whereas a *compiler* reads in a programming language and generates assembler/machine code. Hence, C∀ is like the C preprocessor modifying a program and sending it on to another step for further transformation. The order of transformation is C preprocessor, C∀, and finally GNU C compiler, which also has a number of transformation steps, such as assembler and linker.

The command cfa is used to compile a C∀ program and is based on the GNU gcc command, *e.g.*:

cfa [gcc/C∀−options] [C/C∀ source−files] [assembler/loader files]

There is no ordering among options (flags) and files, unless an option has an argument, which must appear immediately after the option possibly with or without a space separating option and argument.

C∀ has the following gcc flags turned on:

−std=gnu11 The 2011 C standard plus GNU extensions.

C∀ has the following new options:

−CFA Only the C preprocessor (flag −E) and the C∀ translator steps are performed and the transformed program is written to standard output, which makes it possible to examine the code generated by the C∀ translator. The generated code starts with the standard C∀ prelude.

−XCFA Pass next flag as-is to the cfa−cpp translator (see details below).

−debug The program is linked with the debugging version of the runtime system. The debug version performs runtime checks to aid the debugging phase of a C∀ program, but can substantially slow program execution. The runtime checks should only be removed after a program is completely debugged. **This option is the default.**

−nodebug The program is linked with the non-debugging version of the runtime system, so the execution of the program is faster. *However, no runtime checks or* asserts *are performed so errors usually result in abnormal program behaviour or termination.*

−help Information about the set of C∀ compilation flags is printed.

1   −nohelp Information about the set of C∀ compilation flags is not printed. **This option is the default.**

2   −quiet The C∀ compilation message is not printed at the beginning of a compilation.

3   −noquiet The C∀ compilation message is printed at the beginning of a compilation. **This option is the default.**

4   The following preprocessor variables are available:

5   __CFA_MAJOR__ is available during preprocessing and its value is the major version number of C∀.[3]

6   __CFA_MINOR__ is available during preprocessing and its value is the minor version number of C∀.

7   __CFA_PATCH__ is available during preprocessing and its value is the patch level number of C∀.

8   __CFA__, __CFORALL__, and __cforall are always available during preprocessing and have no value.

9   These preprocessor variables allow conditional compilation of programs that must work differently in these situations.

10 For example, to toggle between C and C∀ extensions, use the following:

```
11    #ifndef __CFORALL__
12    #include <stdio.h>                    // C header file
13    #else
14    #include <fstream.hfa>                // C∀ header file
15    #endif
```

16 which conditionally includes the correct header file, if the program is compiled using gcc or cfa.

17   The C∀ transpiler has multiple internal steps. The following flags control how the C∀ transpiler works, the stages

18 run, and printing within a stage. The majority of these flags are used by C∀ developers, but some are occasionally

19 useful to programmers. Each option must be escaped with −XCFA to direct it to the C∀ compilation step, similar to the

20 −Xlinker flag for the linker, *e.g.*:

21   cfa test.cfa −CFA −XCFA −p *# print translated code without printing the standard prelude*

22   cfa test.cfa −XCFA −P −XCFA parse −XCFA −n *# show program parse without prelude*

23 Alternatively, multiple flags can be specified separated with commas and *without* spaces.

24   cfa §test§.cfa −XCFA,−Pparse,−n *# show program parse without prelude*

25   −c, −−colors   diagnostic color: never, always, auto
26   −g, −−gdb   wait for gdb to attach
27   −h, −−help   print transpiler help message
28   −i, −−invariant   invariant checking during AST passes
29   −l, −−libcfa   generate libcfa.c
30   −L, −−linemarks   generate line marks
31   −m, −−no−main   do not replace main
32   −N, −−no−linemarks   do not generate line marks
33   −n, −−no−prelude   do not read prelude
34   −p, −−prototypes   do not generate prelude prototypes ⇒ prelude not printed
35   −d, −−deterministic−out   only print deterministic output
36   −P, −−print   one of:
37   ascodegen   print AST as codegen rather than AST
38   asterr   print AST on error
39   declstats   print code property statistics
40   parse   print yacc (parsing) debug information
41   pretty   prettyprint for ascodegen flag
42   rproto   resolver-proto instance
43   rsteps   print resolver steps
44   ast   print AST after parsing
45   excpdecl   print AST after translating exception decls
46   symevt   print AST after symbol table events
47   expralt   print AST after expressions alternatives
48   valdecl   print AST after declaration validation pass
49   bresolver   print AST before resolver step

---

[3]The C preprocessor allows only integer values in a preprocessor variable so a value like "1.0.0 " is not allowed. Hence, the need to have three variables for the major, minor and patch version number.

```
                       // include file uses the CFA keyword "with".
                       #if ! defined( with )                    // nesting ?
                       #define with ``with                      // make keyword an identifier
                       #define __CFA_BFD_H__
                       #endif
                       #include_next <bfdlink.h>                // must have internal check for multiple expansion
                       #if defined( with ) && defined( __CFA_BFD_H__ )  // reset only if set
                       #undef with
                       #undef __CFA_BFD_H__
                       #endif
```

Figure 1: Header-File Interposition

1   expranly   print AST after expression analysis
2   ctordtor   print AST after ctor/dtor are replaced
3   tuple   print AST after tuple expansion
4   instgen   print AST after instantiate generics
5   bbox   print AST before box pass
6   bcodegen   print AST before code generation
7   --prelude-dir <directory> prelude directory for debug/nodebug
8   -S, --statistics <option-list> enable profiling information: counters, heap, time, all, none
9   -t, --tree build in tree

6   Backquote Identifiers

11 C∀ introduces several new keywords (see Section C, p. 76) that can clash with existing C variable-names in legacy
12 code. Keyword clashes are accommodated by syntactic transformations using the C∀ backquote escape-mechanism:

13   **int ``coroutine** = 3;   // make keyword an identifier
14   **double ``forall** = 3.5;

15 Existing C programs with keyword clashes can be converted by prefixing the keyword identifiers with double back-
16 quotes, and eventually the identifier name can be changed to a non-keyword name. Figure 1 shows how clashes in
17 existing C header-files (see Section D, p. 76) can be handled using preprocessor *interposition*: **#include_next** and
18 command-line -I filename. Several common C header-files with keyword clashes are fixed in the standard C∀ header-
19 library, so there is largely a seamless programming-experience.

7   Constant Underscores

21 Numeric constants are extended to allow underscores as a separator, *e.g.*:

22   2_147_483_648;   // decimal constant
23   56_ul;   // decimal unsigned long constant
24   0_377;   // octal constant
25   0x_ff_ff;   // hexadecimal constant
26   0x_ef3d_aa5c;   // hexadecimal constant
27   3.141_592_654;   // floating constant
28   10_e_+1_00;   // floating constant
29   0x_ff_ff_p_3;   // hexadecimal floating
30   0x_1.ffff_ffff_p_128_l;   // hexadecimal floating long constant
31   L_"\x_ff_ee";   // wide character constant

32 The rules for placement of underscores are:

33   1. A sequence of underscores is disallowed, *e.g.*, 12__34 is invalid.

34   2. Underscores may only appear within a sequence of digits (regardless of the digit radix). In other words, an
35   underscore cannot start or end a sequence of digits, *e.g.*, _1, 1_ and _1_ are invalid (actually, the 1st and 3rd
36   examples are identifier names).

3. A numeric prefix may end with an underscore; a numeric infix may begin and/or end with an underscore; a numeric suffix may begin with an underscore. For example, the octal 0 or hexadecimal 0x prefix may end with an underscore 0_377 or 0x_ff; the exponent infix E may start or end with an underscore 1.0_E10, 1.0E_10 or 1.0_E_10; the type suffixes U, L, *etc*. may start with an underscore 1_U, 1_ll or 1.0E10_f.

It is significantly easier to read and enter long constants when they are broken up into smaller groupings (most cultures use comma and/or period among digits for the same purpose). This extension is backwards compatible, matches with the use of underscore in variable names, and appears in Ada and Java 8. C++ uses the single quote (') as a separator, restricted within a sequence of digits, *e.g*., 0xaa'ff, 3.141'592E1'1. However, the drawback of the C++ approach is differentiating between character and numeric constants by IDEs, as quotes are no longer balanced ('x' and 3.14'159).

8 Exponentiation Operator

Exponentiation, x^y, means raise x to the yth power. When y is a positive integer, exponentiation corresponds to $\prod_{i=1}^{y} x$.

C, C++, Java and other programming languages have no exponentiation operator, using a routine like pow(x, y) instead. Ada, Haskell, Python and other programming languages often use operators $\wedge$ or ** for exponentiation. However, neither of these operators work in C as $\wedge$ means exclusive-or and ** means double dereference. Furthermore, using a routine for exponentiation does not match with mathematical expectation, *i.e*., −x**−y becomes pow(−x, −y).

C∀ extends the basic C operator set with symbol \ (backslash) as the exponentiation operator, represented by routines ?\? and ?\=?, respectively. For example, x \ y and x \= y mean x^y and $x \leftarrow x^y$. The priority of the exponentiation operator is between the cast and multiplicative operators, so −f(x) \ −g(y) is parenthesized as (−f(x)) \ (−g(y)). The C pow routines continues to be available for backwards compatibility.

Exponentiation is overloaded for integral and floating types, including the builtin complex types. Integral exponentiation is performed with repeated multiplication ($O(\log y)$) or shifting if the exponent is 2. Overflow for a large exponent or negative exponent returns zero. Floating exponentiation is performed using logarithms, so the exponent cannot be negative.

```
sout | 1 \ 0 | 1 \ 1 | 2 \ 8 | −4 \ 3 | 5 \ 3 | 5 \ 32 | 5L \ 32 | 5L \ 64 | −4 \ −3 | −4.0 \ −3 | 4.0 \ 2.1 | (1.0f+2.0fi) \ (3.0f+2.0fi);
1 1 256 −64 125 0 3273344365508751233 0 0 −0.015625 18.3791736799526 0.264715−1.1922i
```

Note, 5 \ 32 and 5L \ 64 overflow, and −4 \ −3 is a fraction but stored in an integer so all three computations generate an integral zero. Because exponentiation has higher priority than +, parenthesis are required for exponentiation of complex constants or the expression is parsed as 1.0f+(2.0fi \ 3.0f)+2.0fi, requiring (1.0f+2.0fi) \ (3.0f+2.0fi).

The exponentiation operator is available for all the basic types, but for user-defined types, only the integral-computation version is available.

```
forall( T | { void ?{}( T & this, one_t ); T ?*?( T, T ); } )
T ?\?( T ep, unsigned int y );
forall( T | { void ?{}( T & this, one_t ); T ?*?( T, T ); } )
T ?\?( T ep, unsigned long int y );
```

A user type T must define one (1) and multiplication (*) (see Section 26.4, p. 68).

9 Control Structures

C∀ identifies inconsistent, problematic, and missing control structures in C, and extends, modifies, and adds control structures to increase functionality and safety.

9.1 if / while Statement

The **if** and **while** expressions are extended with declarations, similar to the **for** declaration expression.[4]

```
if ( int x = f() ) ...                     // x != 0
if ( int x = f(), y = g() ) ...            // x != 0 && y != 0
if ( int x = f(), y = g(); x < y ) ...     // relational expression
if ( struct S { int i; } x = { f() }; x.i < 4 )  // relational expression

while ( int x = f() ) ...                  // x != 0
while ( int x = f(), y = g() ) ...         // x != 0 && y != 0
```

---

[4]Declarations in the **do**-**while** condition are not useful because they appear after the loop body.

```
1      while ( int x = f(), y = g(); x < y ) ...          // relational expression
2      while ( struct S { int i; } x = { f() }; x.i < 4 ) ...     // relational expression
```

Unless a relational expression is specified, each variable is compared not equal to 0, which is the standard semantics for the **if/while** expression, and the results are combined using the logical && operator. The scope of the declaration(s) is local to the **if/while** statement, *i.e.*, in both *then* and *else* clauses for **if**, and loop body for **while**. C++ only provides a single declaration always compared != to 0.

9.2   case Clause

C restricts the **case** clause in a **switch** statement to a single value. For multiple **case** clauses prefixing a statement within the **switch** statement, it is necessary to have multiple **case** clauses rather than multiple values. Requiring a **case** clause for each value is not in the spirit of brevity normally associated with C. Therefore, the **case** clause is extended with a list of values.

```
              C                              C∀
    switch ( i ) {                switch ( i ) {
      case 1: case 3 : case 5:      case 1, 3, 5:      // odd values
        ...                            ...
      case 2: case 4 : case 6:      case 2, 4, 6:      // even values
        ...                            ...
    }                             }
```

In addition, inclusive ranges are allowed using symbol ~ to specify a contiguous set of case values, both positive and negative.

```
                C                                C∀                      gcc
    switch ( i ) {                      switch ( i ) {           switch ( i ) {
      case −4: case −3: case −2: case −1:   case −4~−1:            case −4 ...−1:    // −4, −3, −2, −1
        ...                                  ...                    ...
      case 10: case 11: case 12: case 13:   case 10~13:            case 10 ...13:    // 10, 11, 12, 13
        ...                                  ...                    ...
    }                                  }                        }
```

While gcc has the same range mechanism, it has an awkward syntax, 2 ...42, because a space is required after the lower bound, otherwise the period is a decimal point.

C∀ also allows lists of subranges.

```
    case −5~−1, 12~21, 35~42:
```

9.3   switch Statement

C allows a number of questionable forms for the **switch** statement:

1. By default, the end of a **case** clause[5] *falls through* to the next **case** clause in the **switch** statement; to exit a **switch** statement from a **case** clause requires explicitly terminating the clause with a transfer statement, most commonly **break**:

   ```
       switch ( i ) {
         case 1:
           ...
         // fall-through
         case 2:
           ...
           break;   // exit switch statement
       }
   ```

   The ability to fall-through to the next clause *is* a useful form of control flow, specifically when a sequence of case actions compound:

---

[5]In this section, the term *case clause* refers to either a **case** or **default** clause.

```
switch ( argc ) {                       if ( argc == 3 ) {
  case 3:                                 // open output file
    // open output file                   // open input file
    // fall−through                     } else if ( argc == 2 ) {
  case 2:                                 // open input file (duplicate)
    // open input file
    break;  // exit switch statement
  default:                              } else {
    // usage message                      // usage message
}                                       }
```

In this example, case 2 is always done if case 3 is done. This control flow is difficult to simulate with **if** statements or a **switch** statement without fall-through as code must be duplicated or placed in a separate routine. C also uses fall-through to handle multiple case-values resulting in the same action:

```
switch ( i ) {
  case 1: case 3: case 5:    // odd values
    // odd action
    break;
  case 2: case 4: case 6:    // even values
    // even action
    break;
}
```

This situation is better handled by a list of case values (see Section 9.2).

While fall-through itself is not a problem, the problem occurs when fall-through is the default, as this semantics is unintuitive for many programmers and is different from most programming languages with a **switch** statement. Hence, default fall-through semantics results in errors as programmers often *forget* the **break** statement at the end of a **case** clause, resulting in inadvertent fall-through.

2. It is possible to place **case** clauses on statements nested *within* the body of the **switch** statement:

```
switch ( i ) {
  case 0:
    if ( j < k ) {
        ...
      case 1:    // transfer into "if" statement
        ...
    } // if
```

This usage branches into control structures, which causes comprehension and technical difficulties. The comprehension problem results from the inability to determine how control reaches a particular point due to the number of branches leading to it. The technical problem results from the inability to ensure declaration and initialization of variables when blocks are not entered at the beginning. There are few arguments for this kind of control flow, and therefore, there is a strong impetus to eliminate it. This C idiom is known as "Duff's device" [10], from the example:

```
register int n = (count + 7) / 8;
switch ( count % 8 ) {
case 0: do{ *to = *from++;
case 7:     *to = *from++;
case 6:     *to = *from++;
case 5:     *to = *from++;
case 4:     *to = *from++;
case 3:     *to = *from++;
case 2:     *to = *from++;
case 1:     *to = *from++;
    } while ( −−n > 0 );
}
```

which unrolls a loop N times (N = 8 above) and uses the **switch** statement to deal with any iterations not a multiple of N. While efficient, this sort of special purpose usage is questionable:

Disgusting, no? But it compiles and runs just fine. I feel a combination of pride and revulsion at this

discovery. [10]

3. It is possible to place the **default** clause anywhere in the list of labelled clauses for a **switch** statement, rather than only at the end. Most programming languages with a **switch** statement require the **default** clause to appear last in the case-clause list. The logic for this semantics is that after checking all the **case** clauses without success, the **default** clause is selected; hence, physically placing the **default** clause at the end of the **case** clause list matches with this semantics. This physical placement can be compared to the physical placement of an **else** clause at the end of a series of connected **if/else** statements.

4. It is possible to place unreachable code at the start of a **switch** statement, as in:

```
switch ( x ) {
    int y = 1;                  // unreachable initialization
    x = 7;                      // unreachable code without label/branch
case 0: ...
    ...
    int z = 0;                  // unreachable initialization, cannot appear after case
    z = 2;
case 1:
    x = z;                      // without fall through, z is uninitialized
}
```

While the declaration of the local variable y is useful with a scope across all **case** clauses, the initialization for such a variable is defined to never be executed because control always transfers over it. Furthermore, any statements before the first **case** clause can only be executed if labelled and transferred to using a **goto**, either from outside or inside of the **switch**, where both are problematic. As well, the declaration of z cannot occur after the **case** because a label can only be attached to a statement, and without a fall-through to case 3, z is uninitialized. The key observation is that the **switch** statement branches into a control structure, *i.e.*, there are multiple entry points into its statement body.

Before discussing potential language changes to deal with these problems, it is worth observing that in a typical C program:

- the number of **switch** statements is small,
- most **switch** statements are well formed (*i.e.*, no Duff's device),
- the **default** clause is usually written as the last case-clause,
- and there is only a medium amount of fall-through from one **case** clause to the next, and most of these result from a list of case values executing common code, rather than a sequence of case actions that compound.

These observations put into perspective the C∀ changes to the **switch** statement.

1. Eliminating default fall-through has the greatest potential for affecting existing code. However, even if fall-through is removed, most **switch** statements would continue to work because of the explicit transfers already present at the end of each **case** clause, the common placement of the **default** clause at the end of the case list, and the most common use of fall-through, *i.e.*, a list of **case** clauses executing common code, *e.g.*:

```
case 1:  case 2:  case 3: ...
```

still works. Nevertheless, reversing the default action would have a non-trivial effect on case actions that compound, such as the above example of processing shell arguments. Therefore, to preserve backwards compatibility, it is necessary to introduce a new kind of **switch** statement, called **choose**, with no implicit fall-through semantics and an explicit fall-through if the last statement of a case-clause ends with the new keyword **fallthrough**, *e.g.*:

```
choose ( i ) {
case 1:  case 2:  case 3:
    ...
    // implicit end of switch (break)
case 5:
    ...
    fallthrough;                // explicit fall through
case 7:
    ...
    break                       // explicit end of switch (redundant)
```

```
1      default:
2         j = 3;
3      }
```

4   Like the **switch** statement, the **choose** statement retains the fall-through semantics for a list of **case** clauses. An
5   implicit **break** is applied only at the end of the *statements* following a **case** clause. An explicit **fallthrough** is
6   retained because it is a C-idiom most C programmers expect, and its absence might discourage programmers
7   from using the **choose** statement. As well, allowing an explicit **break** from the **choose** is a carry over from the
8   **switch** statement, and expected by C programmers.

9   2. Duff's device is eliminated from both **switch** and **choose** statements, and only invalidates a small amount of
10   very questionable code. Hence, the **case** clause must appear at the same nesting level as the **switch/choose**
11   body, as is done in most other programming languages with **switch** statements.

12 3. The issue of **default** at locations other than at the end of the cause clause can be solved by using good program-
13   ming style, and there are a few reasonable situations involving fall-through where the **default** clause needs to
14   appear is locations other than at the end. Therefore, no change is made for this issue.

15 4. Dealing with unreachable code in a **switch/choose** body is solved by restricting declarations and initialization to
16   the start of statement body, which is executed *before* the transfer to the appropriate **case** clause[6] and precluding
17   statements before the first **case** clause. Further declarations at the same nesting level as the statement body are
18   disallowed to ensure every transfer into the body is sound.

```
19      switch ( x ) {
20         int i = 0;                    // allowed only at start
21         case 0:
22           ...
23         int j = 0;                    // disallowed
24         case 1:
25           {
26             int k = 0;                // allowed at different nesting levels
27             ...
28             case 2:                   // disallow case in nested statements
29           }
30         ...
31      }
```

32   9.4   Non-terminating and Labelled **fallthrough**

33 The **fallthrough** clause may be non-terminating within a **case** clause or have a target label to common code from
34 multiple case clauses.

```
choose ( ... ) {                choose ( ... ) {                choose ( ... ) {
  case 3:                         case 3:                         case 3:
    if ( ... ) {                    ... fallthrough common;         choose ( ... ) {
      ... fallthrough; // goto case 4   case 4:                       case 4:
    } else {                        ... fallthrough common;           for ( ... ) {
      ...                                                               // multi−level transfer
    }                             common: // below fallthrough          ... fallthrough common;
    // implicit break                    // at case−clause level      }
  case 4:                           ... // common code for cases 3/4   ...
                                    // implicit break              }
                                  case 4:                           ...
                                                                  common: // below fallthrough
                                                                          // at case−clause level
```

36 The target label must be below the **fallthrough** and may not be nested in a control structure, and the target label must
37 be at the same or higher level as the containing **case** clause and located at the same level as a **case** clause; the target
38 label may be case **default**, but only associated with the current **switch/choose** statement.

---

[6]Essentially, these declarations are hoisted before the **switch/choose** statement and both declarations and statement are surrounded by a compound
statement.

9.5 Loop Control

Looping a predefined number of times, possibly with a loop index, occurs frequently. C∀ condenses writing loops to facilitate coding speed and safety.

for, **while**, and **do** loop-control are extended with an empty conditional, meaning a comparison value of 1 (true).

```
while ( /* empty */ )              // while ( true )
for ( /* empty */ )                // for ( ; true; )
do ... while ( /* empty */ )       // do ... while ( true )
```

The **for** control, *i.e.*, **for** (/* control */), is extended with a range and step. A range is a set of values defined by an optional low value (default to 0), tilde, and high value, L ~ H, with an optional step ~ S (default to 1), which means an ascending set of values from L to H in positive steps of S.

```
0 ~ 5                              // { 0, 1, 2, 3, 4, 5 }
−8 ~ −2 ~ 2                        // { -8. -6, -4, -2 }
−3 ~ 3 ~ 1                         // { -3, -2, -1, 0, 1, 2, 3 }
```

Warning: A range in descending order, *e.g.*, 5 ~ −3 is the null (empty) set, *i.e.*, no values in the set. Warning: A 0 or negative step is undefined. Note, the order of values in a set may not be the order the values are presented during looping.

The range character, '~', is decorated on the left and right to control how the set values are presented in the loop body. The range character can be prefixed with '+' or '−' indicating the *direction* the range is scanned, *i.e.*, from left to right (ascending) or right to left (descending). Ascending stepping uses operator +=; descending stepping uses operator −=. If there is no prefix character, it defaults to '+'.

```
−8 ~ −2                            // ascending, no prefix
0 +~ 5                             // ascending, prefix
−3 −~ 3                            // descending
```

For descending iteration, the L and H values are *implicitly* switched, and the increment/decrement for S is toggled. When changing the iteration direction, this form is faster and safer, *i.e.*, the direction prefix can be added/removed without changing existing (correct) program text. Warning: reversing the range endpoints for descending order results in an empty set.

```
for ( i; 10 −~ 1 )                 // WRONG descending range!
```

Because C uses zero origin, most loops iterate from 0 to $N-1$. Hence, when scanning a range during iteration, the last value is dropped, *e.g.*, 0 ~ 5 is 0, 1, 2, 3, 4, an exclusive range, [L,H). To obtain *all* the values in the range, the range character is postfixed with '=', *e.g.*, 0 ~= 5 is 0, 1, 2, 3, 4, 5, an inclusive range, [L,H].

for control is formalized by the following regular expression:

```
[ L ] [ + | − ] ~ [ = ] H [ ~ S ]
```

where [] denotes optional and | denotes alternative. That is, the optional low set value, the optional scan direction (ascending/descending), the (possibly) required range character, the optional include last-scan value, the required high set value, and the optional range character and step value. Warning: the regular expression allows the form ~H, but this syntax has a preexisting meaning in C: complement the bits of H, *e.g.*, **for** (~5) meaning **for** (−6), as −6 is the complement of 5. This anomaly is unlikely to cause problems because programers should write the shorter **for** (5).

The previous **for** loops have an anonymous loop index in which the range iteration is computed. To access the value of the range iteration in the loop body, a loop index is specified before the range.

```
for ( int i; 0 ~ 10 ~ 2 ) { ... i ... }     // loop index available in loop body
```

Hence, unlike the 3 components in the C **for**-control, there are only two components in the C∀ **for**-control: the optional index variable and the range. The index type is optional (like C++ **auto**), where the type is normally inferred from the low value L because it initializes the index (the type of H can be different from L). When L is omitted, the type of the required high value H is used, as both L and H are the same type in this case.

```
for ( i; 1.5 ~ 5 )                 // typeof(1.5) i; 1.5 is low value
for ( i; 5.5 )                     // typeof(5.5) i; 5.5 is high value
```

The following examples illustrate common C∀ **for**-control combinations, with the C counter-part in the comment.
- H is implicit ascending exclusive range [0,H).

```
    for ( 5 )                      // for ( typeof(5) i; i < 5; i += 1 )
```

- ~= H is implicit ascending inclusive range [0,H].

1   **for** (~= 5)   *// for (typeof(5) i; i <= 5; i += 1)*

2   • L ~ H is explicit ascending exclusive range [L,H).

3   **for** (1 ~ 5)   *// for (typeof(1) i = 1; i < 5; i += 1)*

4   • L ~= H is explicit ascending inclusive range [L,H].

5   **for** (1 ~= 5)   *// for (typeof(1) i = 1; i <= 5; i += 1)*

6   • L –~ H is explicit descending exclusive range (H,L], where L and H are implicitly interchanged to make the range

7   descending.

8   **for** (1 –~ 5)   *// for (typeof(1) i = 5; i > 0; i -= 1)*

9   • L –~= H is explicit descending inclusive range [H,L], where L and H are implicitly interchanged to make the

10   range descending.

11   **for** (1 –~= 5)   *// for (typeof(1) i = 5; i >= 0; i -= 1)*

12   There are situations when the **for**-control actions need to be moved into the loop body, *e.g.*, a mid-loop exit does

13   not need an iteration-completion test in the **for** control. The character ' @ ' indicates that a specific **for**-control action is

14   ignored, *i.e.*, generates no code.

15   **for** (i; @ –~ 10)   *// for (typeof(10) i = 10; /*empty*/; i -= 1)*

16   **for** (i; 1 ~ @ ~ 2)   *// for (typeof(1) i = 1; /* empty */; i += 2)*

17   **for** (i; 1 ~ 10 ~ @)   *// for (typeof(1) i = 1; i < 10; /* empty */)*

18   **for** (i; 1 ~ @ ~ @)   *// for (typeof(1) i = 1; /* empty */; /* empty */)*

19   Warning: L *cannot* be elided for the ascending range, @ ~ 5, nor H for the descending range, 1 –~ @, as the loop index

20   is uninitialized. Warning: H *cannot* be elided in an anonymous loop index, 1 ~ @, as there is no index to stop the loop.

21   There are situations when multiple loop indexes are required. The character ' : ' means add another index, where

22   any number of indices may be chained in a single **for** control.

23   **for** (i; 5 : j; 2 ~ 12 ~ 3)   *// for (typeof(i) i = 1, j = 2; i < 5 && j < 12; i += 1, j += 3)*

24   **for** (i; 5 : j; 2 ~ @ ~ 3)   *// for (typeof(i) i = 1, j = 2; i < 5; i += 1, j += 3)*

25   **for** (i; 5 : j; 2.5 ~ @ ~ 3.5)   *// no C equivalent, without hoisting declaration of floating-point j*

26   Figure 2 shows more complex loop-control examples across all the different options.

27   Finally, any type that satisfies the Iterate trait can be used with **for** control.

28   **forall**(T) **trait** Iterate {

29   **void** ?{}(T & t, **zero_t**);

30   **int** ?<?(T t1, T t2);

31   **int** ?<=?(T t1, T t2);

32   **int** ?>?(T t1, T t2);

33   **int** ?>=?(T t1, T t2);

34   T ?+=?(T & t1, T t2);

35   T ?+=?(T & t, **one_t**);

36   T ?–=?(T & t1, T t2);

37   T ?–=?(T & t, **one_t**);

38   }

39   Figure 3, p. 14 shows an example of a structure using **for** control. Note, the use of (S){0} when implicitly setting the

40   loop-index type, because using 0 incorrect declares the index to **int** rather than S.

41   9.6   Labelled continue / break Statement

42   C **continue** and **break** statements are restricted to one level of nesting for a particular control structure. This restriction

43   forces programmers to use **goto** to achieve the equivalent control-flow for more than one level of nesting. To prevent

44   having to switch to the **goto**, C∀ extends the **continue** and **break** with a target label to support static multi-level

45   exit [4], as in Java. For both **continue** and **break**, the target label must be directly associated with a **for**, **while** or **do**

46   statement; for **break**, the target label can also be associated with a **switch**, **if** or compound ({}) statement. Figure 4,

47   p. 14 shows a comparison between labelled **continue** and **break** and the corresponding C equivalent using **goto** and

48   labels. The innermost loop has 8 exit points, which cause continuation or termination of one or more of the 7 nested

49   control-structures.

50   Both labelled **continue** and **break** are a **goto** restricted in the following ways:

loop control		output
while () { sout | `"empty"`; **break**; }		empty
do { sout | `"empty"`; **break**; } **while** ();		empty
for () { sout | `"empty"`; **break**; }	*sout | nl | nlOff;*	empty
for (0) { sout | `"A"`; }	*sout | nl;*	
for (1) { sout | `"A"`; }	*sout | nl;*	A
for (10) { sout | `"A"`; }	*sout | nl;*	A A A A A A A A A A
for (~= 10) { sout | `"A"`; }	*sout | nl;*	A A A A A A A A A A A
for (1 ~= 10 ~ 2) { sout | `"B"`; }	*sout | nl;*	B B B B B
for (1 −~= 10 ~ 2) { sout | `"C"`; }	*sout | nl;*	C C C C C
for (0.5 ~ 5.5) { sout | `"D"`; }	*sout | nl;*	D D D D D
for (0.5 −~ 5.5) { sout | `"E"`; }	*sout | nl;*	E E E E E
for (i; 10) { sout | i; }	*sout | nl;*	0 1 2 3 4 5 6 7 8 9
for (i; ~= 10) { sout | i; }	*sout | nl;*	0 1 2 3 4 5 6 7 8 9 10
for (i; 1 ~= 10 ~ 2) { sout | i; }	*sout | nl;*	1 3 5 7 9
for (i; 1 −~= 10 ~ 2) { sout | i; }	*sout | nl;*	10 8 6 4 2
for (i; 0.5 ~ 5.5) { sout | i; }	*sout | nl;*	0.5 1.5 2.5 3.5 4.5
for (i; 0.5 −~ 5.5) { sout | i; }	*sout | nl;*	5.5 4.5 3.5 2.5 1.5
for (ui; 2u ~= 10u ~ 2u) { sout | ui; }	*sout | nl;*	2 4 6 8 10
for (ui; 2u −~= 10u ~ 2u) { sout | ui; }	*sout | nl | nl | nl;*	10 8 6 4 2
enum { N = 10 };		
for (N) { sout | `"N"`; }	*sout | nl;*	N N N N N N N N N N
for (i; N) { sout | i; }	*sout | nl;*	0 1 2 3 4 5 6 7 8 9
for (i; −~ N) { sout | i; }	*sout | nl | nl | nl;*	10 9 8 7 6 5 4 3 2 1
const int low = 3, high = 10, inc = 2;		
for (i; low ~ high ~ inc + 1) { sout | i; }	*sout | nl;*	3 6 9
for (i; 1 ~ @) { **if** (i > 10) **break**; sout | i; }	*sout | nl;*	1 2 3 4 5 6 7 8 9 10
for (i; @ −~ 10) { **if** (i < 0) **break**; sout | i; }	*sout | nl;*	10 9 8 7 6 5 4 3 2 1 0
for (i; 2 ~ @ ~ 2) { **if** (i > 10) **break**; sout | i; }	*sout | nl;*	2 4 6 8 10
for (i; 2.1 ~ @ ~ @) { **if** (i > 10.5) **break**; sout | i; i += 1.7; } *sout | nl;*		2.1 3.8 5.5 7.2 8.9
for (i; @ −~ 10 ~ 2) { **if** (i < 0) **break**; sout | i; }	*sout | nl;*	10 8 6 4 2 0
for (i; 12.1 ~ @ ~ @) { **if** (i < 2.5) **break**; sout | i; i −= 1.7; } *sout | nl;*		12.1 10.4 8.7 7. 5.3 3.6
for (i; 5 : j; −5 ~ @) { sout | i | j; }	*sout | nl;*	0 −5 1 −4 2 −3 3 −2 4 −1
for (i; 5 : j; @ −~ −5) { sout | i | j; }	*sout | nl;*	0 −5 1 −6 2 −7 3 −8 4 −9
for (i; 5 : j; −5 ~ @ ~ 2) { sout | i | j; }	*sout | nl;*	0 −5 1 −3 2 −1 3 1 4 3
for (i; 5 : j; @ −~ −5 ~ 2) { sout | i | j; }	*sout | nl;*	0 −5 1 −7 2 −9 3 −11 4 −13
for (i; 5 : j; −5 ~ @) { sout | i | j; }	*sout | nl;*	0 −5 1 −4 2 −3 3 −2 4 −1
for (i; 5 : j; @ −~ −5) { sout | i | j; }	*sout | nl;*	0 −5 1 −6 2 −7 3 −8 4 −9
for (i; 5 : j; −5 ~ @ ~ 2) { sout | i | j; }	*sout | nl;*	0 −5 1 −3 2 −1 3 1 4 3
for (i; 5 : j; @ −~ −5 ~ 2) { sout | i | j; }	*sout | nl;*	0 −5 1 −7 2 −9 3 −11 4 −13
for (i; 5 : j; @ −~ −5 ~ 2 : k; 1.5 ~ @) { sout | i | j | k; } *sout | nl;*		0 −5 1.5 1 −7 2.5 2 −9 3.5 3 −11 4.5 4 −13 5.5
for (i; 5 : j; @ −~ −5 ~ 2 : k; 1.5 ~ @) { sout | i | j | k; } *sout | nl;*		0 −5 1.5 1 −7 2.5 2 −9 3.5 3 −11 4.5 4 −13 5.5
for (i; 5 : k; 1.5 ~ @ : j; @ −~ −5 ~ 2) { sout | i | j | k; } *sout | nl;*		0 −5 1.5 1 −7 2.5 2 −9 3.5 3 −11 4.5 4 −13 5.5

Figure 2: Loop Control Examples

```
struct S { int i, j; };                                              int main() {
void ?{}( S & s, int i = 0, int j = 0 ) { s.[i, j] = [i, j]; }           for ( S i = 0; i < (S){10,10}; i += 1 ) { sout | i; } sout | "A" | nl; // C
void ?{}( S & s, zero_t ) { s.[i, j] = 0; }                             for ( S i; 0 ~ (S){10,10} ) { sout | i; } sout | "B" | nl; // CFA
int ?<?( S t1, S t2 ) { return t1.i < t2.i && t1.j < t2.j; }            for ( i; (S){10,10} ) { sout | i; } sout | "C" | nl;
int ?<=?( S t1, S t2 ) { return t1.i <= t2.i && t1.j <= t2.j; }         for ( i; (S){0} ~ (S){10,10} ) { sout | i; } sout | "D" | nl;
int ?>?( S t1, S t2 ) { return t1.i > t2.i && t1.j > t2.j; }            for ( i; (S){0} ~= (S){10,10} ) { sout | i; } sout | "E" | nl;
int ?>=?( S t1, S t2 ) { return t1.i >= t2.i && t1.j >= t2.j; }         for ( i; (S){0} ~= (S){10,10} ~ (S){2} ) { sout | i; } sout | "F" | nl;
S ?+=?( S & t1, S t2 ) { t1.i += t2.i; t1.j += t2.j; return t1; }       for ( i; (S){0} –~  (S){10,10} ) { sout | i; } sout | "G" | nl;
S ?+=?( S & t, one_t ) { t.i += 1; t.j += 1; return t; }                for ( i; (S){0} –~= (S){10,10} ) { sout | i; } sout | "H" | nl;
S ?–=?( S & t1, S t2 ) { t1.i –= t2.i; t1.j –= t2.j; return t1; }       for ( i; (S){0} –~= (S){10,10} ~ (S){2,1} ) { sout | i; } sout | "I" | nl;
S ?–=?( S & t, one_t ) { t.i –= 1; t.j –= 1; return t; }            }
ofstream & ?|?( ofstream & os, S s ) {                               (0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) A
    return os | "(" | s.i | s.j | ")";                              (0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) B
}                                                                   (0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) C
void & ?|?( ofstream & os, S s ) {                                  (0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) D
    (ofstream &)(os | s); ends( os );                               (0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) (10 10) E
}                                                                   (0 0) (2 0) (4 0) (6 0) (8 0) (10 0) F
                                                                    (10 10) (9 9) (8 8) (7 7) (6 6) (5 5) (4 4) (3 3) (2 2) (1 1) G
                                                                    (10 10) (9 9) (8 8) (7 7) (6 6) (5 5) (4 4) (3 3) (2 2) (1 1) (0 0) H
                                                                    (10 10) (8 9) (6 8) (4 7) (2 6) (0 5) I
```

Figure 3: For Control with Structure Type

```
{                                              Compound: {
                                                 Try: try {
    ForC: for ( ... ) {                              For: for ( ... ) {
      WhileC: while ( ... ) {                          While: while ( ... ) {
        DoC: do {                                        Do: do {
          if ( ... ) {                                     If: if ( ... ) {
            switch ( ... ) {                                 Switch: switch ( ... ) {
              case 3:                                          case 3:
                goto Compound;                                   break Compound;
                goto Try;                                        break Try;
                goto ForB;     /* or */ goto ForC;              break For;     /* or */ continue For;
                goto WhileB; /* or */ goto WhileC;              break While; /* or */ continue While;
                goto DoB;      /* or */ goto DoC;               break Do;      /* or */ continue Do;
                goto If;                                        break If;
                goto Switch;                                    break Switch;
            } Switch: ;                                        } // switch
          } else {                                           } else {
            ... goto If; ...   // terminate if                  ... break If; ...   // terminate if
          } If:;                                              } // if
        } while ( ... ); DoB: ;                             } while ( ... ); // do
      } WhileB: ;                                          } // while
    } ForB: ;                                            } // for
                                                       } finally { // always executed
                                                       } // try
} Compound: ;                                      } // compound
                a) C                                               b) C∀
```

Figure 4: Multi-level Exit

- They cannot create a loop, which means only the looping constructs cause looping. This restriction means all situations resulting in repeated execution are clearly delineated.
- They cannot branch into a control structure. This restriction prevents missing declarations and/or initializations at the start of a control structure resulting in undefined behaviour.

The advantage of the labelled **continue/break** is allowing static multi-level exits without having to use the **goto** statement, and tying control flow to the target control structure rather than an arbitrary point in a program via a label. Furthermore, the location of the label at the *beginning* of the target control structure informs the reader (eye candy) that complex control-flow is occurring in the body of the control structure. With **goto**, the label is at the end of the control structure, which fails to convey this important clue early enough to the reader. Finally, using an explicit target for the transfer, instead of an implicit target, allows new constructs to be added or removed without affecting existing constructs. Otherwise, the implicit targets of the current **continue** and **break**, *i.e.*, the closest enclosing loop or **switch**, change as certain constructs are added or removed.

9.7 Extended else

The **if** statement has an optional **else** clause executed if the conditional is false. This concept is extended to the **while**, **for**, and **do** looping constructs (like Python). Hence, if the loop conditional becomes false, looping stops and the corresponding **else** clause is executed, if present.

The following example is a linear search for the key 3 in an array, where finding the key is handled with a **break** and not finding with the **else** clause on the loop construct.

```
int a[10];
```

```
                                                                        int i = 0;
while ( int i = 0; i < 10 ) {        for ( i; 10 ) {                    do {
  if ( a[i] == 3 ) break; // found     if ( a[i] == 3 ) break; // found   if ( a[i] == 3 ) break; // found
  i += 1;                              i += 1;                            i += 1;
} else { // i == 10                  } else { // i == 10                } while( i < 10 ) else { // i == 10
  sout | "not found";                  sout | "not found";                sout | "not found";
}                                    }                                  }
```

Note, dangling else now occurs with **if**, **while**, **for**, **do**, and **waitfor**.

9.8 with Statement

Grouping heterogeneous data into an *aggregate* (structure/union) is a common programming practice, and aggregates may be nested:

```
struct Person {                     // aggregate
  struct Name {                     // nesting
    char first[20], last[20];
  } name;
  struct Address {                  // nesting
    ...
  } address;
  int sex;
};
```

Functions manipulating aggregates must repeat the aggregate name to access its containing fields.

```
Person p
p.name ...;  p.address ...;  p.sex ...;     // access containing fields
```

which extends to multiple levels of qualification for nested aggregates and multiple aggregates.

```
struct Ticket { ... } t;
p.name.first ...;  p.address.street ...;    // access nested fields
t.departure ...;  t.cost ...;               // access multiple aggregate
```

Repeated aggregate qualification is tedious and makes code difficult to read. Therefore, reducing aggregate qualification is a useful language design goal.

1   C partially addresses the problem by eliminating qualification for enumerated types and unnamed *nested* aggre-
2 gates, which open their scope into the containing aggregate. This feature is used to group fields for attributes and/or
3 with **union** aggregates.

```
4      struct S {
5          struct /* unnamed */ { int g,  h; } __attribute__(( aligned(64) ));
6          int tag;
7          union /* unnamed */ {
8              struct { char c1,  c2; } __attribute__(( aligned(128) ));
9              struct { int i1,  i2; };
10             struct { double d1,  d2; };
11         };
12     } s;
13     enum { R, G, B };
14     s.g; s.h;   s.tag = R;   s.c1; s.c2;   s.i1 = G; s.i2 = B;   s.d1; s.d2;
```

15 Object-oriented languages reduce qualification for class variables within member functions, *e.g.*, C++:

```
16     struct S {
17         char c;   int i;   double d;
18         void f( /* S * this */ ) {              // implicit "this" parameter
19             c;  i;  d;                           // this->c; this->i; this->d;
20         }
21     }
```

22 In general, qualification is elided for the variables and functions in the lexical scopes visible from a member function.
23 However, qualification is necessary for name shadowing and explicit aggregate parameters.

```
24     struct T {
25         char m;   int i;   double n;            // derived class variables
26     };
27     struct S : public T {
28         char c;   int i;   double d;            // class variables
29         void g( double d, T & t ) {
30             d;  t.m;  t.i;  t.n;                 // function parameter
31             c;  i;  this->d;  S::d;              // class S variables
32             m;  T::i;  n;                        // class T variables
33         }
34     };
```

35 Note the three different forms of qualification syntax in C++, ., –>, ::, which is confusing.
36   Since C∀ in not object-oriented, it has no implicit parameter with its implicit qualification. Instead C∀ introduces
37 a general mechanism using the **with** statement (see Pascal [23, § 4.F]) to explicitly elide aggregate qualification by
38 opening a scope containing the field identifiers. Hence, the qualified fields become variables with the side-effect that
39 it is simpler to write, easier to read, and optimize field references in a block.

```
40     void f( S & this ) with ( this ) {          // with statement
41         c;  i;  d;                               // this.c, this.i, this.d
42     }
```

43 with the generality of opening multiple aggregate-parameters:

```
44     void g( S & s, T & t ) with ( s, t ) {      // multiple aggregate parameters
45         c;  s.i;  d;                             // s.c, s.i, s.d
46         m;  t.i;  n;                             // t.m, t.i, t.n
47     }
```

48 where qualification is only necessary to disambiguate the shadowed variable i. In detail, the **with** statement may form
49 a function body or be nested within a function body.
50   The **with** clause takes a list of expressions, where each expression provides an aggregate type and object. (Enu-
51 merations are already opened.) To open a pointer type, the pointer must be dereferenced to obtain a reference to the
52 aggregate type.

```
53     S * sp;
54     with ( *sp ) { ... }
```

55 The expression object is the implicit qualifier for the open structure-fields.

C∀'s ability to overload variables (see Section 26.2, p. 67) and use the left-side of assignment in type resolution means most fields with the same name but different types are automatically disambiguated, eliminating qualification. All expressions in the expression list are open in parallel within the compound statement. This semantic is different from Pascal, which nests the openings from left to right. The difference between parallel and nesting occurs for fields with the same name and type:

```
struct Q { int i; int k; int m; } q, w;
struct R { int i; int j; double m; } r, w;
with ( r, q ) {
    j + k;                    // unambiguous, r.j + q.k
    m = 5.0;                  // unambiguous, q.m = 5.0
    m = 1;                    // unambiguous, r.m = 1
    int a = m;                // unambiguous, a = r.i
    double b = m;             // unambiguous, b = q.m
    int c = r.i + q.i;        // disambiguate with qualification
    (double)m;                // disambiguate with cast
}
```

For parallel semantics, both r.i and q.i are visible, so i is ambiguous without qualification; for nested semantics, q.i hides r.i, so i implies q.i. Pascal nested-semantics is possible by nesting **with** statements.

```
with ( r ) {
    i;                        // unambiguous, r.i
    with ( q ) {
        i;                    // unambiguous, q.i
    }
}
```

A cast or qualification can be used to disambiguate variables within a **with** *statement*. A cast can also be used to disambiguate among overload variables in a **with** *expression*:

```
with ( w ) { ... }                // ambiguous, same name and no context
with ( (Q)w ) { ... }             // unambiguous, cast
```

Because there is no left-side in the **with** expression to implicitly disambiguate between the w variables, it is necessary to explicitly disambiguate by casting w to type Q or R.

Finally, there is an interesting problem between parameters and the function-body **with**, *e.g.*:

```
void f( S & s, char c ) with ( s ) {
    s.c = c;  i = 3;  d = 5.5;        // initialize fields
}
```

Here, the assignment s.c = c means s.c = s.c, which is meaningless, and there is no mechanism to qualify the parameter c, making the assignment impossible using the function-body **with**. To solve this problem, parameters *not* explicitly opened are treated like an initialized aggregate:

```
struct Params {                    // s explicitly opened so S & s elided
    char c;
} params;
```

and implicitly opened *after* a function-body open, to give them higher priority:

```
void f( S & s, char c ) with ( s ) with( params ) { // syntax disallowed, illustration only
    s.c = c;  i = 3;  d = 5.5;
}
```

This implicit semantic matches with programmer expectation.

10   Exception Handling

Exception handling provides two mechanism: change of control flow from a raise to a handler, and communication from the raise to the handler. Transfer of control can be local, within a routine, or non-local, among routines. Non-local transfer can cause stack unwinding, *i.e.*, non-local routine termination, depending on the kind of raise.

Currently, CV uses macros ExceptionDecl and ExceptionInst to declare and instantiate an exception.

```
#include <Exception.hfa>
ExceptionDecl( E,       // must be global scope
```

```
1          ... // exception fields
2       );
3       try {
4          ...
5          if ( ... ) throwResume ExceptionInst( E, /* intialization */ );
6          if ( ... ) throw ExceptionInst( E, /* intialization */ );
7          ...
8       } catchResume( E * ) { // must be pointer
9          ...
10      } catch( E * ) {
11         ...
12      }

13      exception_t E {};                    // exception type
14      void f(...) {
15         ... throw E{}; ...                // termination
16         ... throwResume E{}; ...          // resumption
17      }
18      try {
19         f(...);
20      } catch( E e ; boolean-predicate ) {     // termination handler
21         // recover and continue
22      } catchResume( E e ; boolean-predicate ) { // resumption handler
23         // repair and return
24      } finally {
25         // always executed
26      }
```

27 The kind of raise and handler match: **throw** with **catch** and **throwResume** with **catchResume**. Then the exception type
28 must match along with any additional predicate must be true. The **catch** and **catchResume** handlers may appear in
29 any oder. However, the **finally** clause must appear at the end of the **try** statement.

30 ## 10.1   Non-local Exception

```
31      void main() {
32         try {
33            _Enable {
34               ... resume(); ...
35            }
36         } catchResume( E & ) { // should be reference
37            ...
38         } catch( E & ) {
39            ...
40         }
41      }
```
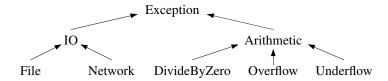
42 ## 10.2   Exception Hierarchy

43 An exception type can be derived from another exception type, just like deriving a subclass from a class, providing
44 a kind of polymorphism among exception types. The exception-type hierarchy that is created is used to organize
45 exception types, similar to a class hierarchy in object-oriented languages, *e.g.*:

47 A programmer can then choose to handle an exception at different degrees of specificity along the hierarchy; derived
48 exception-types support a more flexible programming style. For example, higher-level code should catch general
49 exceptions to reduce coupling to the specific implementation at the lower levels; unnecessary coupling may force

changes in higher-level code when low-level code changes. A consequence of derived exception-types is that multiple exceptions may match, *e.g.*:

catch(Arithmetic)

matches all three derived exception-types: DivideByZero, Overflow, and Underflow. Because the propagation mechanisms perform a simple linear search of the handler clause for a guarded block, and selects the first matching handler, the order of catch clauses in the handler clause becomes important, *e.g.*:

```
try {
    ...
} catch( Overflow ) {   // must appear first
    // handle overflow
} catch( Arithmetic )
    // handle other arithmetic issues
}
```

Multiple derivation among exception is not supported.

11   Alternative Declarations

C declaration syntax is notoriously confusing and error prone. For example, many C programmers are confused by a declaration as simple as:

int ∗ x[5]   x   x

Is this an array of 5 pointers to integers or a pointer to an array of 5 integers? If there is any doubt, it implies productivity and safety issues even for basic programs. Another example of confusion results from the fact that a routine name and its parameters are embedded within the return type, mimicking the way the return value is used at the routine's call site. For example, a routine returning a pointer to an array of integers is defined and used in the following way:

```
int (∗f())[5] {...};                // definition
... (∗f())[3] += 1;                 // usage
```

Essentially, the return type is wrapped around the routine name in successive layers (like an onion). While attempting to make the two contexts consistent is a laudable goal, it has not worked out in practice, even though Dennis Richie believed otherwise:

> In spite of its difficulties, I believe that the C's approach to declarations remains plausible, and am comfortable with it; it is a useful unifying principle. [27, p. 12]

C∀ provides its own type, variable and routine declarations, using a different syntax. The new declarations place qualifiers to the left of the base type, while C declarations place qualifiers to the right of the base type. In the following example, red is the base type and **blue** is qualifiers. The C∀ declarations move the qualifiers to the left of the base type, *i.e.*, move the blue to the left of the red, while the qualifiers have the same meaning but are ordered left to right to specify a variable's type.

C	C∀
int ∗ x1 [5];	[5] ∗ **int** x1;
int (∗x2)[5];	∗ [5] **int** x2;
int (∗f(**int** p))[5];	[∗ [5] **int**] f(**int** p);

The only exception is bit field specification, which always appear to the right of the base type. However, unlike C, C∀ type declaration tokens are distributed across all variables in the declaration list. For instance, variables x and y of type pointer to integer are defined in C∀ as follows:

C	C∀
int ∗x, ∗y;	∗ **int** x, y;

The downside of this semantics is the need to separate regular and pointer declarations:

	C	**C∀**
1	**int** *x, y;	* **int** x;
		int y;

2   which is prescribing a safety benefit. Other examples are:

	C	**C∀**	
	int z[5];	[5] **int** z;	*// array of 5 integers*
	char * w[5];	[5] * **char** w;	*// array of 5 pointers to char*
	double (* v)[5];	* [5] **double** v;	*// pointer to array of 5 doubles*
3	**struct** s {	**struct** s {	
	int f0:3;	**int** f0:3;	*// common bit field syntax*
	int * f1;	* **int** f1;	
	int * f2[5]	[5] * **int** f2;	
	};	};	

4   All type qualifiers, *e.g.*, **const**, **volatile**, *etc.*, are used in the normal way with the new declarations and also appear
5   left to right, *e.g.*:

	C	**C∀**	
6	**int const** * **const** x;	**const** * **const int** x;	*// const pointer to const integer*
	const int (* **const** y)[5]	**const** * [5] **const int** y;	*// const pointer to array of 5 const integers*

7   All declaration qualifiers, *e.g.*, **extern**, **static**, *etc.*, are used in the normal way with the new declarations but can only
8   appear at the start of a C∀ routine declaration,[7] *e.g.*:

	C	**C∀**	
9	**int extern** x[5];	**extern** [5] **int** x;	*// externally visible array of 5 integers*
	const int static * y;	**static** * **const int** y;	*// internally visible pointer to constant int*

10   The new declaration syntax can be used in other contexts where types are required, *e.g.*, casts and the pseudo-
11   routine **sizeof**:

	C	**C∀**
12	y = (**int** *)x;	y = (* **int**)x;
	i = **sizeof**(**int** * [5]);	i = **sizeof**([5] * **int**);

13   Finally, new C∀ declarations may appear together with C declarations in the same program block, but cannot be
14   mixed within a specific declaration. Therefore, a programmer has the option of either continuing to use traditional C
15   declarations or take advantage of the new style. Clearly, both styles need to be supported for some time due to existing
16   C-style header-files, particularly for UNIX-like systems.

17   # 12   Pointer / Reference

18   C provides a *pointer type*; C∀ adds a *reference type*. These types may be derived from an object or routine type, called
19   the *referenced type*. Objects of these types contain an *address*, which is normally a location in memory, but may also
20   address memory-mapped registers in hardware devices. An integer constant expression with the value 0, or such an
21   expression cast to type **void** *, is called a *null-pointer constant*.[8] An address is *sound*, if it points to a valid memory
22   location in scope, *i.e.*, within the program's execution-environment and has not been freed. Dereferencing an *unsound*
23   address, including the null pointer, is undefined, often resulting in a memory fault.

24   A program *object* is a region of data storage in the execution environment, the contents of which can represent
25   values. In most cases, objects are located in memory at an address, and the variable name for an object is an implicit
26   address to the object generated by the compiler and automatically dereferenced, as in:

---

[7] The placement of a storage-class specifier other than at the beginning of the declaration specifiers in a declaration is an obsolescent feature. [21, § 6.11.5(1)]

[8] One way to conceptualize the null pointer is that no variable is placed at this address, so the null-pointer address can be used to denote an uninitialized pointer/reference object; *i.e.*, the null pointer is guaranteed to compare unequal to a pointer to any object or routine. In general, a value with special meaning among a set of values is called a *sentinel value*, *e.g.*, –1 as a return code value.

2   where the right example is how the compiler logically interprets the variables in the left example. Since a variable
3   name only points to one address during its lifetime, it is an immutable pointer; hence, the implicit type of pointer
4   variables x and y are constant pointers in the compiler interpretation. In general, variable addresses are stored in
5   instructions instead of loaded from memory, and hence may not occupy storage. These approaches are contrasted in
6   the following:

explicit variable address	implicit variable address

lda   r1,100 // *load address of x*
 ld   r2,(r1) // *load value of x*
 lda   r3,104 // *load address of y*
 st   r2,(r3) // *store x into y*
	ld r2,(100) // *load value of x*
	
	 st r2,(104) // *store x into y*

8   Finally, the immutable nature of a variable's address and the fact that there is no storage for the variable pointer means
9   pointer assignment is impossible. Therefore, the expression x = y has only one meaning, *x = *y, *i.e.*, manipulate values,
10 which is why explicitly writing the dereferences is unnecessary even though it occurs implicitly as part of instruction
11 decoding.
12   A pointer/reference object is a generalization of an object variable-name, *i.e.*, a mutable address that can point to
13 more than one memory location during its lifetime. (Similarly, an integer variable can contain multiple integer literals
14 during its lifetime versus an integer constant representing a single literal during its lifetime, and like a variable name,
15 may not occupy storage if the literal is embedded directly into instructions.) Hence, a pointer occupies memory to
16 store its current address, and the pointer's value is loaded by dereferencing, *e.g.*:

18   Notice, an address has a duality: a location in memory or the value at that location. In many cases, a compiler
19 might be able to infer the best meaning for these two cases. For example, Algol68 [24] infers pointer dereferencing to
20 select the best meaning for each pointer usage

21   p2 = p1 + x;   // *compiler infers *p2 = *p1 + x;*

22 Algol68 infers the following dereferencing *p2 = *p1 + x, because adding the arbitrary integer value in x to the address
23 of p1 and storing the resulting address into p2 is an unlikely operation. Unfortunately, automatic dereferencing does
24 not work in all cases, and so some mechanism is necessary to fix incorrect choices.
25   Rather than inferring dereference, most programming languages pick one implicit dereferencing semantics, and the
26 programmer explicitly indicates the other to resolve address-duality. In C, objects of pointer type always manipulate
27 the pointer object's address:

28   p1 = p2;   // *p1 = p2 rather than *p1 = *p2*
29   p2 = p1 + x;   // *p2 = p1 + x rather than *p2 = *p1 + x*

30 even though the assignment to p2 is likely incorrect, and the programmer probably meant:

31   p1 = p2;   // *pointer address assignment*
32   *p2 = *p1 + x;   // *pointed-to value assignment / operation*

33 The C semantics work well for situations where manipulation of addresses is the primary meaning and data is rarely
34 accessed, such as storage management (malloc/free).
35   However, in most other situations, the pointed-to value is requested more often than the pointer address.

36   *p2 = ((*p1 + *p2) * (**p3 − *p1)) / (**p3 − 15);

37 In this case, it is tedious to explicitly write the dereferencing, and error prone when pointer arithmetic is allowed. It is
38 better to have the compiler generate the dereferencing and have no implicit pointer arithmetic:

39   p2 = ((p1 + p2) * (p3 − p1)) / (p3 − 15);

40   To support this common case, a reference type is introduced in C∀, denoted by &, which is the opposite dereference

semantics to a pointer type, making the value at the pointed-to location the implicit semantics for dereferencing (similar but not the same as C++ reference types).

```
int x, y, & r1, & r2, && r3;
&r1 = &x;                        // r1 points to x
&r2 = &r1;                       // r2 points to x
&r1 = &y;                        // r1 points to y
&&r3 = &&r2;                     // r3 points to r2
r2 = ((r1 + r2) * (r3 − r1)) / (r3 − 15);   // implicit dereferencing
```

Except for auto-dereferencing by the compiler, this reference example is the same as the previous pointer example. Hence, a reference behaves like the variable name for the current variable it is pointing-to. One way to conceptualize a reference is via a rewrite rule, where the compiler inserts a dereference operator before the reference variable for each reference qualifier in a declaration, so the previous example becomes:

```
*r2 = ((*r1 + *r2) * (**r3 − *r1)) / (**r3 − 15);
```

When a reference operation appears beside a dereference operation, *e.g.*, &*, they cancel out. However, in C, the cancellation always yields a value (rvalue).[9] For a C∀ reference type, the cancellation on the left-hand side of assignment leaves the reference as an address (lvalue):

```
(&*)r1 = &x;                     // (&*) cancel giving address in r1 not variable pointed-to by r1
```

Similarly, the address of a reference can be obtained for assignment or computation (rvalue):

```
(&(&*)*)r3 = &(&*)r2;            // (&*) cancel giving address in r2, (&(&*)*) cancel giving address in r3
```

Cancellation works to arbitrary depth.

Fundamentally, pointer and reference objects are functionally interchangeable because both contain addresses.

```
int x, *p1 = &x, **p2 = &p1, ***p3 = &p2,
        &r1 = x,   &&r2 = r1,   &&&r3 = r2;
***p3 = 3;                       // change x
r3 = 3;                          // change x, ***r3
**p3 = ...;                      // change p1
&r3 = ...;                       // change r1, (&*)**r3, 1 cancellation
*p3 = ...;                       // change p2
&&r3 = ...;                      // change r2, (&(&*)*)*r3, 2 cancellations
&&&r3 = p3;                      // change r3 to p3, (&(&(&*)*)*)r3, 3 cancellations
```

Furthermore, both types are equally performant, as the same amount of dereferencing occurs for both types. Therefore, the choice between them is based solely on whether the address is dereferenced frequently or infrequently, which dictates the amount of implicit dereferencing aid from the compiler.

As for a pointer type, a reference type may have qualifiers:

```
const int cx = 5;                     // cannot change cx;
const int & cr = cx;                  // cannot change what cr points to
&cr = &cx;                            // can change cr
cr = 7;                               // error, cannot change cx
int & const rc = x;                   // must be initialized
&rc = &x;                             // error, cannot change rc
const int & const crc = cx;           // must be initialized
crc = 7;                              // error, cannot change cx
&crc = &cx;                           // error, cannot change crc
```

Hence, for type & **const**, there is no pointer assignment, so &rc = &x is disallowed, and *the address value cannot be the null pointer unless an arbitrary pointer is coerced into the reference*:

```
int & const cr = *0;                  // where 0 is the int * zero
```

Note, constant reference-types do not prevent addressing errors because of explicit storage-management:

```
int & const cr = *malloc();
cr = 5;
free( &cr );
```

---

[9]The unary & operator yields the address of its operand. If the operand has type "type", the result has type "pointer to type". If the operand is the result of a unary * operator, neither that operator nor the & operator is evaluated and the result is as if both were omitted, except that the constraints on the operators still apply and the result is not an lvalue. [21, § 6.5.3.2–3]

1   cr = 7;   // *unsound pointer dereference*

2   The position of the **const** qualifier *after* the pointer/reference qualifier causes confuse for C programmers. The
3   **const** qualifier cannot be moved before the pointer/reference qualifier for C style-declarations; C∀-style declarations
4   (see Section 11, p. 19) attempt to address this issue:

C	C∀
const int * **const** * **const** ccp;	**const** * **const** * **const int** ccp;
	const & **const** & **const int** ccr;

6   where the C∀ declaration is read left-to-right.
7   Finally, like pointers, references are usable and composable with other type operators and generators.

8   **int** w, x, y, z, & ar[3] = { x, y, z };   // *initialize array of references*
9   &ar[1] = &w;   // *change reference array element*
10   **typeof**(ar[1]) p;   // *(gcc) is int, i.e., the type of referenced object*
11   **typeof**(&ar[1]) q;   // *(gcc) is int &, i.e., the type of reference*
12   **sizeof**(ar[1]) == **sizeof**(**int**);   // *is true, i.e., the size of referenced object*
13   **sizeof**(&ar[1]) == **sizeof**(**int** *)   // *is true, i.e., the size of a reference*

14   In contrast to C∀ reference types, C++'s reference types are all **const** references, preventing changes to the reference
15   address, so only value assignment is possible, which eliminates half of the address duality. Also, C++ does not allow
16   arrays of reference[10] Java's reference types to objects (all Java objects are on the heap) are like C pointers, which
17   always manipulate the address, and there is no (bit-wise) object assignment, so objects are explicitly cloned by shallow
18   or deep copying, which eliminates half of the address duality.

19   12.1   Initialization

20   Initialization is different than assignment because initialization occurs on the empty (uninitialized) storage on an
21   object, while assignment occurs on possibly initialized storage of an object. There are three initialization contexts
22   in C∀: declaration initialization, argument/parameter binding, return/temporary binding. Because the object being
23   initialized has no value, there is only one meaningful semantics with respect to address duality: it must mean address
24   as there is no pointed-to value. In contrast, the left-hand side of assignment has an address that has a duality. Therefore,
25   for pointer/reference initialization, the initializing value must be an address not a value.

26   **int** * p = &x;   // *assign address of x*
27   **int** * p = x;   // *assign value of x*
28   **int** & r = x;   // *must have address of x*

29   Like the previous example with C pointer-arithmetic, it is unlikely assigning the value of x into a pointer is meaningful
30   (again, a warning is usually given). Therefore, for safety, this context requires an address, so it is superfluous to
31   require explicitly taking the address of the initialization object, even though the type is incorrect. Note, this is strictly
32   a convenience and safety feature for a programmer. Hence, C∀ allows r to be assigned x because it infers a reference
33   for x, by implicitly inserting a address-of operator, &, and it is an error to put an & because the types no longer match
34   due to the implicit dereference. Unfortunately, C allows p to be assigned with &x (address) or x (value), but most
35   compilers warn about the latter assignment as being potentially incorrect. Similarly, when a reference type is used for
36   a parameter/return type, the call-site argument does not require a reference operator for the same reason.

37   **int** & f(**int** & r);   // *reference parameter and return*
38   z = f(x) + f(y);   // *reference operator added, temporaries needed for call results*

39   Within routine f, it is possible to change the argument by changing the corresponding parameter, and parameter r can
40   be locally reassigned within f. Since operator routine ?+? takes its arguments by value, the references returned from f
41   are used to initialize compiler generated temporaries with value semantics that copy from the references.

42   **int** temp1 = f(x), temp2 = f(y);
43   z = temp1 + temp2;

44   This implicit referencing is crucial for reducing the syntactic burden for programmers when using references; other-
45   wise references have the same syntactic burden as pointers in these contexts.
46   When a pointer/reference parameter has a **const** value (immutable), it is possible to pass literals and expressions.

---

[10]The reason for disallowing arrays of reference is unknown, but possibly comes from references being ethereal (like a textual macro), and hence, replaceable by the referent object.

```
1        void f( const int & cr );
2        void g( const int * cp );
3        f( 3 );          g( &3 );
4        f( x + y );      g( &(x + y) );
```

Here, the compiler passes the address to the literal 3 or the temporary for the expression x + y, knowing the argument cannot be changed through the parameter. The & before the constant/expression for the pointer-type parameter (g) is a C∀ extension necessary to type match and is a common requirement before a variable in C (*e.g.*, scanf). Importantly, &3 may not be equal to &3, where the references occur across calls because the temporaries maybe different on each call.

C∀ *extends* this semantics to a mutable pointer/reference parameter, and the compiler implicitly creates the necessary temporary (copying the argument), which is subsequently pointed-to by the reference parameter and can be changed.[11]

```
13       void f( int & r );
14       void g( int * p );
15       f( 3 );          g( &3 );                 // compiler implicit generates temporaries
16       f( x + y );      g( &(x + y) );           // compiler implicit generates temporaries
```

Essentially, there is an implicit rvalue to lvalue conversion in this case.[12] The implicit conversion allows seamless calls to any routine without having to explicitly name/copy the literal/expression to allow the call.

Finally, C handles routine objects in an inconsistent way. A routine object is both a pointer and a reference (particle and wave).

```
21       void f( int i );
22       void (* fp)( int );                       // routine pointer
23       fp = f;                                    // reference initialization
24       fp = &f;                                   // pointer initialization
25       fp = *f;                                   // reference initialization
26       fp(3);                                     // reference invocation
27       (*fp)(3);                                  // pointer invocation
```

While C's treatment of routine objects has similarity to inferring a reference type in initialization contexts, the examples are assignment not initialization, and all possible forms of assignment are possible (f, &f, *f) without regard for type. Instead, a routine object should be referenced by a **const** reference:

```
31       const void (& fr)( int ) = f;             // routine reference
32       fr = ...;                                  // error, cannot change code
33       &fr = ...;                                 // changing routine reference
34       fr( 3 );                                   // reference call to f
35       (*fr)(3);                                  // error, incorrect type
```

because the value of the routine object is a routine literal, *i.e.*, the routine code is normally immutable during execution.[13] C∀ allows this additional use of references for routine objects in an attempt to give a more consistent meaning for them.

12.2 Address-of Semantics

In C, &E is an rvalue for any expression E. C∀ extends the & (address-of) operator as follows:

- if R is an rvalue of type T $\&_1 \cdots \&_r$, where $r \geq 1$ references (& symbols), than &R has type T $*\&_2 \cdots \&_r$, *i.e.*, T pointer with $r - 1$ references (& symbols).
- if L is an lvalue of type T $\&_1 \cdots \&_l$, where $l \geq 0$ references (& symbols), than &L has type T $*\&_1 \cdots \&_l$, *i.e.*, T pointer with l references (& symbols).

The following example shows the first rule applied to different rvalue contexts:

```
46       int x, * px, ** ppx, *** pppx, **** ppppx;
47       int & rx = x, && rrx = rx, &&& rrrx = rrx ;
48       x = rrrx;                                  // rrrx is an lvalue with type int &&& (equivalent to x)
```

---

[11]If whole program analysis is possible, and shows the parameter is not assigned, *i.e.*, it is **const**, the temporary is unnecessary.

[12]This conversion attempts to address the *const hell* problem, when the innocent addition of a **const** qualifier causes a cascade of type failures, requiring an unknown number of additional **const** qualifiers, until it is discovered a **const** qualifier cannot be added and all the **const** qualifiers must be removed.

[13]Dynamic code rewriting is possible but only in special circumstances.

```
1      px = &rrrx;                          // starting from rrrx, &rrrx is an rvalue with type int *&&& (&x)
2      ppx = &&rrrx;                        // starting from &rrrx, &&rrrx is an rvalue with type int **&& (&rx)
3      pppx = &&&rrrx;                      // starting from &&rrrx, &&&rrrx is an rvalue with type int ***& (&rrx)
4      ppppx = &&&&rrrx;                    // starting from &&&rrrx, &&&&rrrx is an rvalue with type int **** (&rrrx)
```

5 The following example shows the second rule applied to different lvalue contexts:

```
6      int x, * px, ** ppx, *** pppx;
7      int & rx = x, && rrx = rx, &&& rrrx = rrx ;
8      rrrx = 2;                            // rrrx is an lvalue with type int &&& (equivalent to x)
9      &rrrx = px;                          // starting from rrrx, &rrrx is an rvalue with type int *&&& (rx)
10     &&rrrx = ppx;                        // starting from &rrrx, &&rrrx is an rvalue with type int **&& (rrx)
11     &&&rrrx = pppx;                      // starting from &&rrrx, &&&rrrx is an rvalue with type int ***& (rrrx)
```

12.3 Conversions

13 C provides a basic implicit conversion to simplify variable usage:

14   0. lvalue to rvalue conversion: cv T converts to T, which allows implicit variable dereferencing.

```
15         int x;
16         x + 1;                           // lvalue variable (int) converts to rvalue for expression
```

17   An rvalue has no type qualifiers (cv), so the lvalue qualifiers are dropped.

18 C∀ provides three new implicit conversion for reference types to simplify reference usage.

19   1. reference to rvalue conversion: cv T & converts to T, which allows implicit reference dereferencing.

```
20         int x, &r = x, f( int p );
21         x = r + f( r );                  // lvalue reference converts to rvalue
```

22   An rvalue has no type qualifiers (cv), so the reference qualifiers are dropped.

23   2. lvalue to reference conversion: lvalue−type cv1 T converts to cv2 T &, which allows implicitly converting vari-
24   ables to references.

```
25         int x, &r = x, f( int & p );     // lvalue variable (int) convert to reference (int &)
26         f( x );                          // lvalue variable (int) convert to reference (int &)
```

27   Conversion can restrict a type, where cv1 ≤ cv2, *e.g.*, passing an **int** to a **const volatile int** &, which has low cost.
28   Conversion can expand a type, where cv1 > cv2, *e.g.*, passing a **const volatile int** to an **int** &, which has high cost
29   (warning); furthermore, if cv1 has **const** but not cv2, a temporary variable is created to preserve the immutable
30   lvalue.

31   3. rvalue to reference conversion: T converts to cv T &, which allows binding references to temporaries.

```
32         int x, & f( int & p );
33         f( x + 3 );                      // rvalue parameter (int) implicitly converts to lvalue temporary reference (int &)
34         &f(...) = &x;                    // rvalue result (int &) implicitly converts to lvalue temporary reference (int &)
```

35   In both case, modifications to the temporary are inaccessible (warning). Conversion expands the temporary-type
36   with cv, which is low cost since the temporary is inaccessible.

13   string **Type**

38 The C∀ string type is for manipulation of dynamically-size character-strings versus C **char** * type for manipulation of
39 statically-size null-terminated character-strings. That is, the amount of storage for a C∀ string changes dynamically
40 at runtime to fit the string size, whereas the amount of storage for a C string is fixed at compile time. Hence, a
41 string declaration does not specify a maximum length; as a string dynamically grows and shrinks in size, so does its
42 underlying storage. In contrast, a C string also dynamically grows and shrinks is size, but its underlying storage is
43 fixed. The maximum storage for a C∀ string value is size_t characters, which is 2^{32} or 2^{64} respectively. A C∀ string
44 manages its length separately from the string, so there is no null ('\0') terminating value at the end of a string value.
45 Hence, a C∀ string cannot be passed to a C string manipulation routine, such as strcat. Like C strings, the characters
46 in a string are numbered starting from 0.

47   The following operations have been defined to manipulate an instance of type string. The discussion assumes the
48 following declarations and assignment statements are executed.

```
49         #include <string.hfa>
```

```
//  string s = 5;               sout | s;
    string s;
    // conversion of char and char * to string
    s = 'x';                              sout | s;              x
    s = "abc";                            sout | s;              abc
    char cs[5] = "abc";
    s = cs;                               sout | s;              abc
    // conversion of integral, floating−point, and complex to string
    s = 45hh;                             sout | s;              45
    s = 45h;                              sout | s;              45
    s = −(ssize_t)MAX − 1;                sout | s;              −9223372036854775808
    s = (size_t)MAX;                      sout | s;              18446744073709551615
    s = 5.5;                              sout | s;              5.5
    s = 5.5L;                             sout | s;              5.5
    s = 5.5+3.4i;                         sout | s;              5.5+3.4i
    s = 5.5L+3.4Li;                       sout | s;              5.5+3.4i
```

Figure 5: Implicit Conversions to String

```
1       string s, peter, digit, alpha, punctuation, ifstmt;
2       int i;
3       peter  = "PETER";
4       digit  = "0123456789";
5       punctuation = "()., ";
6       ifstmt = "IF (A > B) {";
```

7 Note, the include file string.hfa to access type string.

8   13.1   Implicit String Conversions

9 The types **char**, **char** *, **int**, **double**, **_Complex**, including different signness and sizes, implicitly convert to type string.
10 Figure 5 shows examples of implicit conversions between C strings, integral, floating-point and complex types to
11 string. A conversions can be explicitly specified:

```
12      s = string( "abc" );            // converts char * to string
13      s = string( 5 );                // converts int to string
14      s = string( 5.5 );              // converts double to string
```

15 All conversions from string to **char** *, attempt to be safe: either by requiring the maximum length of the **char** * storage
16 (strncpy) or allocating the **char** * storage for the string characters (ownership), meaning the programmer must free the
17 storage. As well, a string is always null terminates, implying a minimum size of 1 character.

```
        string s = "abcde";
        char cs[3];
        strncpy( cs, s, sizeof(cs) );         sout | cs;         ab
18      char * cp = s;                        sout | cp;         abcde
        delete( cp );
        cp = s + '  ' + s;                    sout | cp;         abcde abcde
        delete( cp );
```

19   13.2   Size (length)

20 The size operation returns the length of a string.

```
21      i = size( "" );                 // i is assigned 0
22      i = size( "abc" );              // i is assigned 3
23      i = size( peter );              // i is assigned 5
```

13.3 Comparison Operators

The binary relational operators, <, <=, >, >=, and equality operators, ==, !=, compare strings using lexicographical ordering, where longer strings are greater than shorter strings.

13.4 Concatenation

The binary operators + and += concatenate two strings, creating the sum of the strings.

```
s = peter + ' ' + digit;          // s is assigned "PETER 0123456789"
s += peter;                       // s is assigned "PETER 0123456789PETER"
```

13.5 Repetition

The binary operators * and *= repeat a string *N* times. If $N = 0$, a zero length string, `""` is returned.

```
s = 'x' * 3;                      // s is assigned "PETER PETER PETER "
s = (peter + ' ') * 3;            // s is assigned "PETER PETER PETER "
```

13.6 Substring

The substring operation returns a subset of the string starting at a position in the string and traversing a length.

```
s = peter( 2, 3 );                // s is assigned "ETE"
s = peter( 4, −3 );               // s is assigned "ETE", length is opposite direction
s = peter( 2, 8 );                // s is assigned "ETER", length is clipped to 4
s = peter( 0, −1 );               // s is assigned "", beyond string so clipped to null
s = peter(−1, −1 );               // s is assigned "R", start and length are negative
```

A negative starting position is a specification from the right end of the string. A negative length means that characters are selected in the opposite (right to left) direction from the starting position. If the substring request extends beyond the beginning or end of the string, it is clipped (shortened) to the bounds of the string. If the substring request is completely outside of the original string, a null string located at the end of the original string is returned. The substring operation can also appear on the left hand side of the assignment operator. The substring is replaced by the value on the right hand side of the assignment. The length of the right-hand-side value may be shorter, the same length, or longer than the length of the substring that is selected on the left hand side of the assignment.

```
digit( 3, 3 ) = "";               // digit is assigned "0156789"
digit( 4, 3 ) = "xyz";            // digit is assigned "015xyz9"
digit( 7, 0 ) = "***";            // digit is assigned "015xyz***9"
digit(−4, 3 ) = "$$$";            // digit is assigned "015xyz$$$9"
```

A substring is treated as a pointer into the base (substringed) string rather than creating a copy of the subtext. As with all pointers, if the item they are pointing at is changed, then the pointer is referring to the changed item. Pointers to the result value of a substring operation are defined to always start at the same location in their base string as long as that starting location exists, independent of changes to themselves or the base string. However, if the base string value changes, this may affect the values of one or more of the substrings to that base string. If the base string value shortens so that its end is before the starting location of a substring, resulting in the substring starting location disappearing, the substring becomes a null string located at the end of the base string.

The following example illustrates passing the results of substring operations by reference and by value to a subprogram. Notice the side-effects to other reference parameters as one is modified.

```
main() {
    string x = "xxxxxxxxxxxxx";
    test( x, x(1,3), x(3,3), x(5,5), x(9,5), x(9,5) );
}

// x, a, b, c, & d are substring results passed by reference
// e is a substring result passed by value
void test(string &x, string &a, string &b, string &c, string &d, string e) {
                                  // x a b c d e
    a( 1, 2 ) = "aaa";            // aaaxxxxxxxxxxx aaax axx xxxxx xxxxx xxxxx
    b( 2, 12 ) = "bbb";           // aaabbbxxxxxxxxx aaab abbb bbxxx xxxxx xxxxx
```

```
c( 4, 5 ) = "ccc";                   // aaabbbxcccxxxxxx aaab abbb bbxccc ccxxx xxxxx
c = "yyy";                           // aaabyyyxxxxxx aaab abyy yyy xxxxx xxxxx
d( 1, 3 ) = "ddd";                   // aaabyyyxdddxx aaab abyy yyy dddxx xxxxx
e( 1, 3 ) = "eee";                   // aaabyyyxdddxx aaab abyy yyy dddxx eeexx
x = e;                               // eeexx eeex exx x eeexx
}
```

There is an assignment form of substring in which only the starting position is specified and the length is assumed to be the remainder of the string.

```
string operator () (int start);
```

For example:

```
s = peter( 2 );                      // s is assigned "ETER"
peter( 2 ) = "IPER";                 // peter is assigned "PIPER"
```

It is also possible to substring using a string as the index for selecting the substring portion of the string.

```
string operator () (const string &index);
```

For example:

```
digit( "xyz$$$" ) = "678";           // digit is assigned "0156789"
digit( "234") = "***";               // digit is assigned "0156789***"
```

13.7   Searching

The index operation

int index(**const** string &key, **int** start = 1, occurrence occ = first);

returns the position of the first or last occurrence of the key (depending on the occurrence indicator occ that is either first or last) in the current string starting the search at position start. If the key does not appear in the current string, the length of the current string plus one is returned. A negative starting position is a specification from the right end of the string.

```
i = digit.index( "567" );            // i is assigned 3
i = digit.index( "567", 7 );         // i is assigned 11
i = digit.index( "567", −1, last );  // i is assigned 3
i = peter.index( "E", 5, last );     // i is assigned 4
```

The next two string operations test a string to see if it is or is not composed completely of a particular class of characters. For example, are the characters of a string all alphabetic or all numeric? Use of these operations involves a two step operation. First, it is necessary to create an instance of type strmask and initialize it to a string containing the characters of the particular character class, as in:

```
strmask digitmask = digit;
strmask alphamask = string( "abcdefghijklmnopqrstuvwxyz" );
```

Second, the character mask is used in the functions include and exclude to check a string for compliance of its characters with the characters indicated by the mask.

The include operation

int include(**const** strmask &, **int** = 1, occurrence occ = first);

returns the position of the first or last character (depending on the occurrence indicator, which is either first or last) in the current string that does not appear in the mask starting the search at position start; hence it skips over characters in the current string that are included (in) the mask. The characters in the current string do not have to be in the same order as the mask. If all the characters in the current string appear in the mask, the length of the current string plus one is returned, regardless of which occurrence is being searched for. A negative starting position is a specification from the right end of the string.

```
i = peter.include( digitmask );      // i is assigned 1
i = peter.include( alphamask );      // i is assigned 6
```

The exclude operation

int exclude(string &mask, **int** start = 1, occurrence occ = first)

returns the position of the first or last character (depending on the occurrence indicator, which is either first or last) in the current string that does appear in the mask string starting the search at position start; hence it skips over characters in the current string that are excluded from (not in) in the mask string. The characters in the current string do not have to be in the same order as the mask string. If all the characters in the current string do NOT appear in the mask string, the length of the current string plus one is returned, regardless of which occurrence is being searched for. A negative starting position is a specification from the right end of the string.

```
i = peter.exclude( digitmask );          // i is assigned 6
i = ifstmt.exclude( strmask( punctuation ) ); // i is assigned 4
```

The includeStr operation:

```
string includeStr( strmask &mask, int start = 1, occurrence occ = first )
```

returns the longest substring of leading or trailing characters (depending on the occurrence indicator, which is either first or last) of the current string that ARE included in the mask string starting the search at position start. A negative starting position is a specification from the right end of the string.

```
s = peter.includeStr( alphamask );        // s is assigned "PETER"
s = ifstmt.includeStr( alphamask );       // s is assigned "IF"
s = peter.includeStr( digitmask );        // s is assigned ""
```

The excludeStr operation:

```
string excludeStr( strmask &mask, int start = 1, occurrence = first )
```

returns the longest substring of leading or trailing characters (depending on the occurrence indicator, which is either first or last) of the current string that are excluded (NOT) in the mask string starting the search at position start. A negative starting position is a specification from the right end of the string.

```
s = peter.excludeStr( digitmask);         // s is assigned "PETER"
s = ifstmt.excludeStr( strmask( punctuation ) ); // s is assigned "IF "
s = peter.excludeStr( alphamask);         // s is assigned ""
```

13.8   Miscellaneous

The trim operation

```
string trim( string &mask, occurrence occ = first )
```

returns a string in that is the longest substring of leading or trailing characters (depending on the occurrence indicator, which is either first or last) which ARE included in the mask are removed.

```
// remove leading blanks
s = string( "    ABC" ).trim( "  " );     // s is assigned "ABC",
// remove trailing blanks
s = string( "ABC    " ).trim( "  ", last ); // s is assigned "ABC",
```

The translate operation

```
string translate( string &from, string &to )
```

returns a string that is the same length as the original string in which all occurrences of the characters that appear in the from string have been translated into their corresponding character in the to string. Translation is done on a character by character basis between the from and to strings; hence these two strings must be the same length. If a character in the original string does not appear in the from string, then it simply appears as is in the resulting string.

```
// upper to lower case
peter = peter.translate( "ABCDEFGHIJKLMNOPQRSTUVWXYZ", "abcdefghijklmnopqrstuvwxyz" );
        // peter is assigned "peter"
s = ifstmt.translate( "ABCDEFGHIJKLMNOPQRSTUVWXYZ", "abcdefghijklmnopqrstuvwxyz" );
        // ifstmt is assigned "if (a > b) {"
// lower to upper case
peter = peter.translate( "abcdefghijklmnopqrstuvwxyz", "ABCDEFGHIJKLMNOPQRSTUVWXYZ" );
        // peter is assigned "PETER"
```

The replace operation

```
string replace( string &from, string &to )
```

returns a string in which all occurrences of the from string in the current string have been replaced by the to string.

char []	string
strcpy, strncpy	=
strcat, strncat	+
strcmp, strncmp	==, !=, <, <=, >, >=
strlen	size
[]	[]
strstr	find
strcspn	find_first_of, find_last_of
strspc	find_fist_not_of, find_last_not_of

Table 1: Companion Routines for C∀ string to C Strings

1   s = peter.replace("E", "XX");   // *s is assigned "PXXTXXR"*

2 The replacement is done left-to-right. When an instance of the from string is found and changed to the to string, it is
3 NOT examined again for further replacement.

13.9   Returning N+1 on Failure

5 Any of the string search routines can fail at some point during the search. When this happens it is necessary to return
6 indicating the failure. Many string types in other languages use some special value to indicate the failure. This value
7 is often 0 or -1 (PL/I returns 0). This section argues that a value of N+1, where N is the length of the base string in the
8 search, is a more useful value to return. The index-of function in APL returns N+1. These are the boundary situations
9 and are often overlooked when designing a string type.
10   The situation that can be optimized by returning N+1 is when a search is performed to find the starting location for
11 a substring operation. For example, in a program that is extracting words from a text file, it is necessary to scan from
12 left to right over whitespace until the first alphabetic character is found.

13   line = line(line.exclude(alpha));

14 If a text line contains all whitespaces, the exclude operation fails to find an alphabetic character. If exclude returns 0 or
15 -1, the result of the substring operation is unclear. Most string types generate an error, or clip the starting value to 1,
16 resulting in the entire whitespace string being selected. If exclude returns N+1, the starting position for the substring
17 operation is beyond the end of the string leaving a null string.
18   The same situation occurs when scanning off a word.

19   start = line.include(alpha);
20   word = line(1, start − 1);

21 If the entire line is composed of a word, the include operation will fail to find a non-alphabetic character. In general,
22 returning 0 or -1 is not an appropriate starting position for the substring, which must substring off the word leaving a
23 null string. However, returning N+1 will substring off the word leaving a null string.

13.10   C Compatibility

25 To ease conversion from C to C∀, there are companion string routines for C strings. Table 1 shows the C routines on
26 the left that also work with string and the rough equivalent string opeation of the right. Hence, it is possible to directly
27 convert a block of C string operations into @string@ just by changing the
28   For example, this block of C code can be converted to C∀ by simply changing the type of variable s from **char** [] to
29 string.

```
30        char s[32];
31        //string s;
32        strcpy( s, "abc" );          PRINT( %s, s );
33        strncpy( s, "abcdef", 3 );      PRINT( %s, s );
34        strcat( s, "xyz" );          PRINT( %s, s );
35        strncat( s, "uvwxyz", 3 );      PRINT( %s, s );
36        PRINT( %zd, strlen( s ) );
37        PRINT( %c, s[3] );
38        PRINT( %s, strstr( s, "yzu" ) ) ;
```

```
1       PRINT( %s, strstr( s, 'y' ) ) ;
```

2 However, the conversion fails with I/O because printf cannot print a string using format code %s because C∀ strings are
3 not null terminated.

13.11   Input/Output Operators

5 Both the C++ operators << and >> are defined on type string. However, input of a string value is different from input of
6 a **char** * value. When a string value is read, *all* input characters from the current point in the input stream to either the
7 end of line ('\n') or the end of file are read.

14   Enumeration

9 An *enumeration* is a compile-time mechanism to alias names to constants, like **typedef** is a mechanism to alias names
10 to types. Its purpose is to define a restricted-value type providing code-readability and maintenance – changing an
11 enum's value automatically updates all name usages during compilation.
12   An enumeration type is a set of names, each called an *enumeration constant* (shortened to *enum*) aliased to a fixed
13 value (constant).

```
14      enum Days { Mon, Tue, Wed, Thu, Fri, Sat, Sun }; // enumeration type definition, set of 7 names & values
15      Days days = Mon; // enumeration type declaration and initialization
```

16 The set of enums is injected into the variable namespace at the definition scope. Hence, enums may be overloaded
17 with variable, enum, and function names.

```
18      int Foo;                                // type/variable separate namespaces
19      enum Foo { Bar };
20      enum Goo { Bar };                       // overload Foo.Bar
21      double Bar;                             // overload Foo.Bar, Goo.Bar
```

22 An anonymous enumeration injects enums with specific values into a scope.

```
23      enum { Prime = 103, BufferSize = 1024 };
```

24 An enumeration is better than using C preprocessor or constant declarations.

```
25      #define Mon 0            const int Mon = 0,
        ...                      ...,
        #define Sun 6            Sun = 6;
```

26 because the enumeration is succinct, has automatic numbering, can appear in **case** labels, does not use storage, and
27 is part of the language type-system. Finally, the type of an enum is implicitly or explicitly specified and the constant
28 value can be implicitly or explicitly specified. Note, enum values may be repeated in an enumeration.

14.1   Enum type

30 The type of enums can be any type, and an enum's value comes from this type. Because an enum is a constant, it
31 cannot appear in a mutable context, *e.g.*, Mon = Sun is disallowed, and has no address (it is an rvalue). Therefore,
32 an enum is automatically converted to its constant's base-type, *e.g.*, comparing/printing an enum compares/prints its
33 value rather than the enum name; there is no mechanism to print the enum name.
34   The default enum type is **int**. Hence, Days is the set type Mon, Tue, ... , Sun, while the type of each enum is **int**
35 and each enum represents a fixed integral value. If no values are specified for an integral enum type, the enums are
36 automatically numbered by one from left to right starting at zero. Hence, the value of enum Mon is 0, Tue is 1, ... , Sun
37 is 6. If an enum value is specified, numbering continues by one from that value for subsequent unnumbered enums. If
38 an enum value is a *constant* expression, the compiler performs constant-folding to obtain a constant value.
39   C∀ allows other integral types with associated values.

```
40      enum( char ) Letter { A = 'A', B, C, I = 'I', J, K };
41      enum( long long int ) BigNum { X = 123_456_789_012_345, Y = 345_012_789_456_123 };
```

42 For enumeration Letter, enum A's value is explicitly set to 'A', with B and C implicitly numbered with increasing
43 values from 'A', and similarly for enums I, J, and K.
44   Non-integral enum types must be explicitly initialized, *e.g.*, **double** is not automatically numbered by one.

```
1      // non −integral numeric
2      enum( double ) Math { PI_2 = 1.570796, PI = 3.141597, E = 2.718282 }
3      // pointer
4      enum( char * ) Name { Fred = "Fred",  Mary = "Mary", Jane = "Jane" };
5      int i, j, k;
6      enum( int * ) ptr { I = &i,  J = &j,  K = &k };
7      enum( int & ) ref { I = i,   J = j,   K = k };
8      // tuple
9      enum( [int, int] ) { T = [ 1, 2 ] };
10     // function
11     void f() {...}   void g() {...}
12     enum( void (*)() ) funs { F = f,  F = g };
13     // aggregate
14     struct S { int i, j; };
15     enum( S ) s { A = { 3,  4 }, B = { 7,  8 } };
16     // enumeration
17     enum( Letter ) Greek { Alph = A, Beta = B, /* more enums */ }; // alphabet intersection
```

Enumeration Greek may have more or less enums than Letter, but the enum values *must* be from Letter. Therefore, Greek enums are a subset of type Letter and are type compatible with enumeration Letter, but Letter enums are not type compatible with enumeration Greek.

The following examples illustrate the difference between the enumeration type and the type of its enums.

```
22     Math m = PI;                          // allowed
23     double d = PI;                        // allowed, conversion to base type
24     m = E;                                // allowed
25     m = Alph;                             // disallowed
26     m = 3.141597;                         // disallowed
27     d = m;                                // allowed
28     d = Alph;                             // disallowed
29     Letter l = A;                         // allowed
30     Greek g = Alph;                       // allowed
31     l = Alph;                             // allowed, conversion to base type
32     g = A;                                // disallowed
```

A constructor *cannot* be used to initialize enums because a constructor executes at runtime. A fallback is explicit C-style initialization using @=.

```
35     enum( struct vec3 ) Axis { Up @= { 1, 0, 0 }, Left @= { 0, 1, 0 }, Front @= { 0, 0, 1 } }
```

Finally, enumeration variables are assignable and comparable only if the appropriate operators are defined for its enum type.

14.2   Inheritance

Plan-9 inheritance may be used with enumerations.

```
40     enum( char * ) Name2 { inline Name, Jack = "Jack", Jill = "Jill" };
41     enum /* inferred */ Name3 { inline Name2, Sue = "Sue", Tom = "Tom" };
```

Enumeration Name2 inherits all the enums and their values from enumeration Name by containment, and a Name enumeration is a subtype of enumeration Name2. Note, enums must be unique in inheritance but enum values may be repeated. The enum type for the inheriting type must be the same as the inherited type; hence the enum type may be omitted for the inheriting enumeration and it is inferred from the inherited enumeration, as for Name3. When inheriting from integral types, automatic numbering may be used, so the inheritance placement left to right is important, *e.g.*, the placement of Sue and Tom before or after **inline** Name2.

Specifically, the inheritance relationship for Names is:

Name ⊆ Name2 ⊆ Name3 ⊆ **const char** * // enum type of Name

Hence, given

```
51     void f( Name );
52     void g( Name2 );
53     void h( Name3 );
```

1 **void** j(**const char** *);

2 the following calls are valid

3 f(Fred);

4 g(Fred);   g(Jill);

5 h(Fred);   h(Jill);   h(Sue);

6 j(Fred);   j(Jill);   j(Sue);   j('W');

7 Note, the validity of calls is the same for call-by-reference as for call-by-value, and **const** restrictions are the same as

8 for other types.

9   Enums cannot be created at runtime, so inheritence problems, such as contra-variance do not apply. Only instances

10 of the enum base-type may be created at runtime.

11 15   Routine Definition

12 C∀ supports a new syntax for routine definition, as well as C11 and K&R routine syntax. The point of the new syntax

13 is to allow returning multiple values from a routine [16, 25], *e.g.*:

14   [**int** o1, **int** o2, **char** o3] f(**int** i1, **char** i2, **char** i3) {

15   *routine body*

16   }

17 where routine f has three output (return values) and three input parameters. Existing C syntax cannot be extended with

18 multiple return types because it is impossible to embed a single routine name within multiple return type-specifications.

19   In detail, the brackets, [], enclose the result type, where each return value is named and that name is a local variable

20 of the particular return type.[14] The value of each local return variable is automatically returned at routine termination.

21 Declaration qualifiers can only appear at the start of a routine definition, *e.g.*:

22   **extern** [**int** x] g(**int** y) { }

23 Lastly, if there are no output parameters or input parameters, the brackets and/or parentheses must still be specified; in

24 both cases the type is assumed to be void as opposed to old style C defaults of int return type and unknown parameter

25 types, respectively, as in:

26   [] g();   *// no input or output parameters*

27   [**void**] g(**void**);   *// no input or output parameters*

28   Routine f is called as follows:

29   [i, j, ch] = f(3, 'a', ch);

30 The list of return values from f and the grouping on the left-hand side of the assignment is called a *return list* and

31 discussed in Section 12.

32   C∀ style declarations cannot be used to declare parameters for K&R style routine definitions because of the fol-

33 lowing ambiguity:

34   **int** (*f(x))[5] **int** x; {}

35 The string "**int** (*f(x))[5]" declares a K&R style routine of type returning a pointer to an array of 5 integers, while the

36 string "[5] **int** x" declares a C∀ style parameter x of type array of 5 integers. Since the strings overlap starting with the

37 open bracket, [, there is an ambiguous interpretation for the string.

38   As well, C∀-style declarations cannot be used to declare parameters for C-style routine-definitions because of the

39 following ambiguity:

40   **typedef int** foo;

41   **int** f(**int** (* foo));   *// foo is redefined as a parameter name*

42 The string "**int** (* foo)" declares a C-style named-parameter of type pointer to an integer (the parenthesis are superflu-

43 ous), while the same string declares a C∀ style unnamed parameter of type routine returning integer with unnamed

44 parameter of type pointer to foo. The redefinition of a type name in a parameter list is the only context in C where the

45 character * can appear to the left of a type name, and C∀ relies on all type qualifier characters appearing to the right of

46 the type name. The inability to use C∀ declarations in these two contexts is probably a blessing because it precludes

47 programmers from arbitrarily switching between declarations forms within a declaration contexts.

48   C-style declarations can be used to declare parameters for C∀ style routine definitions, *e.g.*:

---

[14]Michael Tiemann, with help from Doug Lea, provided named return values in g++, circa 1989.

1   [**int**] f(* **int**, **int** *);   *// returns an integer, accepts 2 pointers to integers*
2   [* **int**, **int** *] f(**int**);   *// returns 2 pointers to integers, accepts an integer*

3 The reason for allowing both declaration styles in the new context is for backwards compatibility with existing pre-
4 processor macros that generate C-style declaration-syntax, as in:

5   **#define** ptoa(n, d) **int** (*n)[d]
6   **int** f(ptoa(p, 5)) ...   *// expands to int f(int (*p)[5])*
7   [**int**] f(ptoa(p, 5)) ...   *// expands to [int] f(int (*p)[5])*

8 Again, programmers are highly encouraged to use one declaration form or the other, rather than mixing the forms.

9 ## 15.1   Named Return Values

10 Named return values handle the case where it is necessary to define a local variable whose value is then returned in a
11 **return** statement, as in:

12   **int** f() {
13   **int** x;
14   ... x = 0; ... x = y; ...
15   **return** x;
16   }

17 Because the value in the return variable is automatically returned when a C∀ routine terminates, the **return** statement
18 *does not* contain an expression, as in:

   [**int** x, **int** y] f() {
      int z;
19   ... x = 0; ... y = z; ...
      return;   *// implicitly return x, y*
   }

20 When the return is encountered, the current values of x and y are returned to the calling routine. As well, "falling off
21 the end" of a routine without a **return** statement is permitted, as in:

22   [**int** x, **int** y] f() {
23   ...
24   }   *// implicitly return x, y*

25 In this case, the current values of x and y are returned to the calling routine just as if a **return** had been encountered.
26   Named return values may be used in conjunction with named parameter values; specifically, a return and parameter
27 can have the same name.

28   [**int** x, **int** y] f(**int**, x, **int** y) {
29   ...
30   }   *// implicitly return x, y*

31 This notation allows the compiler to eliminate temporary variables in nested routine calls.

32   [**int** x, **int** y] f(**int**, x, **int** y);   *// prototype declaration*
33   **int** a, b;
34   [a, b] = f(f(f(a, b)));

35 While the compiler normally ignores parameters names in prototype declarations, here they are used to eliminate
36 temporary return-values by inferring that the results of each call are the inputs of the next call, and ultimately, the
37 left-hand side of the assignment. Hence, even without the body of routine f (separate compilation), it is possible to
38 perform a global optimization across routine calls. The compiler warns about naming inconsistencies between routine
39 prototype and definition in this case, and behaviour is undefined if the programmer is inconsistent.

40 ## 15.2   Routine Prototype

41 The syntax of the new routine prototype declaration follows directly from the new routine definition syntax; as well,
42 parameter names are optional, *e.g.*:

43   [**int** x] f ();   *// returning int with no parameters*
44   [* **int**] g (**int** y);   *// returning pointer to int with int parameter*
45   [ ] h (**int**, **char**);   *// returning no result with int and char parameters*
46   [* **int**, **int**] j (**int**);   *// returning pointer to int and int, with int parameter*

1 This syntax allows a prototype declaration to be created by cutting and pasting source text from the routine definition
2 header (or vice versa). Like C, it is possible to declare multiple routine-prototypes in a single declaration, where the
3 return type is distributed across *all* routine names in the declaration list (see Section 11, p. 19), *e.g.*:

4   C :   **const double** bar1(), bar2(**int**), bar3(**double**);
5   C∀ :   [**const double**] foo(), foo(**int**), foo(**double**) { **return** 3.0; }

6 C∀ allows the last routine in the list to define its body.

7   Declaration qualifiers can only appear at the start of a C∀ routine declaration,[7] *e.g.*:

8   **extern** [**int**] f (**int**);
9   **static** [**int**] g (**int**);

10   15.3 Postfix Function

11 C∀ provides an alternative call syntax where the argument appears before the function name. The syntax uses the
12 backquote ` to separate the parameters/arguments and function name: ?` denotes a postfix-function name, *e.g.*,
13 **int** ?`h( **int** s ) and ` denotes a postfix-function call, *e.g.*, 0`h meaning h(0).

postfix function	constant argument call	variable argument call	postfix function pointer
int ?`h(**int** s);	0`h;		**int** (* ?`p)(**int** i);
int ?`h(**double** s);	3.5`h;	**int** i = 7;	?`p = ?`h;
int ?`m(**char** c);	'1'`m;	i`h;	3`p;
int ?`m(**const char** * s);	"123" "456"`m;	(i + 3)`h;	i`p;
int ?`t(**int** a, **int** b, **int** c);	[1, 2, 3]`t;	(i + 3.5)`h;	(i + 3)`p;

15 Note, to pass *multiple* arguments to a postfix function requires a tuple, *e.g.*, [1, 2, 3]`t, which forms a single argument
16 that is flattened into the multiple arguments (see Section 20, p. 39). Similarly, if the argument is an expression, it must
17 be parenthesized, *e.g.*, (i + 3)`h, or only the last operand of the expression is the argument, *e.g.*, i + (3`h).

18   Figure 6 shows a common example for postfix functions: converting basic literals into user literals. (See Sec-
19 tion E.1, p. 77 for other uses for postfix functions.) The C∀ example (left) stores a mass in units of stones (1 stone
20 = 14 lb or 6.35 kg) and provides an addition operator ?+? (imagine a full set of arithmetic operators). The arithmetic
21 operators manipulate stones and the postfix operations convert to/from different units. The three postfixing function
22 names st, lb, and kg, represent units stones, pounds, and kilograms, respectively. Each name has two forms that bidi-
23 rectional convert: a value of a specified unit to stones, *e.g.*, w = 14`lb ⇒ w == 1 stone or a Weight from stones back to
24 specific units, *e.g.*, w`lb (1 stone) to 14. A similar group of postfix functions provide user constants for converting time
25 units into nanoseconds, which is the basic time unit, *e.g.*, ns, us, ms, s, m, h, d, and w, for nanosecond, microsecond,
26 millisecond, second, minute, hour, day, and week, respectively. (Note, month is not a fixed period of time nor is year
27 because of leap years.)

28   The C++ example (right) provides a *restricted* capability via user literals. The **operator""** only takes a constant
29 argument (*i.e.*, no variable as an argument), and the constant type must be the highest-level constant-type, *e.g.*,
30 **long double** for all floating-point constants. As well, there is no constant conversion, *i.e.*, **int** to **double** constants,
31 so integral constants are handled by a separate set of routines, with maximal integral type **unsigned long long int**.
32 Finally, there is no mechanism to use this syntax for a bidirectional conversion because **operator""** only accepts a
33 constant argument.

34   16   Routine Pointers

35 The syntax for pointers to C∀ routines specifies the pointer name on the right, *e.g.*:

36   * [**int** x] () fp;   *// pointer to routine returning int with no parameters*
37   * [* **int**] (**int** y) gp;   *// pointer to routine returning pointer to int with int parameter*
38   * [ ] (**int**,**char**) hp;   *// pointer to routine returning no result with int and char parameters*
39   * [* **int**,**int**] (**int**) jp;   *// pointer to routine returning pointer to int and int, with int parameter*

40 While parameter names are optional, *a routine name cannot be specified*; for example, the following is incorrect:

41   * [**int** x] f () fp;   *// routine name "f" is disallowed*

C∀ Postfix Routine	C++ User Literals
struct Weight { **double** stones; }; Weight ?+?(Weight l, Weight r) { **return** l.stones + r.stones; } Weight ?\`st( **double** w ) { **return** w; } **double** ?\`st(Weight w) { **return** w.stones; } Weight ?\`lb( **double** w ) { **return** w / 14.0; } **double** ?\`lb(Weight w) { **return** w.stones * 14.0; } Weight ?\`kg( **double** w ) { **return** w / 6.35; } **double** ?\`kg(Weight w) { **return** w.stones * 6.35; } **int** main() { Weight w, heavy = { 20 }; *// stones* w = 155\`lb;     w = 0b_1111\`st; w = 0_233\`lb;     w = 0x_9b\`kg; w = 5.5\`st + 8\`kg + 25.01\`lb + heavy; }	**struct** Weight { **double** stones; Weight() {} Weight(**double** w) { stones = w; } }; Weight **operator**+(Weight l, Weight r) { **return** l.stones + r.stones; } Weight **operator" "**_st(**long double** w) { **return** w; } Weight **operator" "**_lb(**long double** w) { **return** w / 14.0; } Weight **operator" "**_kg(**long double** w) { **return** w / 6.35; } Weight **operator" "**_st(**unsigned long long int** w) { **return** w; } Weight **operator" "**_lb(**unsigned long long int** w) { **return** w / 14.0; } Weight **operator" "**_kg(**unsigned long long int** w) { **return** w / 6.35; } **int** main() { Weight w, heavy = { 20 }; *// stones* w = 155_lb; w = 0b1111_st; w = 0'233®_lb®; // quote separator w = 0x9b®_kg®; w = 5.5®_st® + 8®_kg® + 25.01®_lb® + heavy; }

Figure 6: Units: Stone, Pound, Kilogram Comparison

17   Default and Named Parameter

Default and named parameters [20][15] are two mechanisms to simplify routine call.

17.1   Default

A default parameter associates a default value with a parameter so it can be optionally specified in the argument list. For example, given the routine prototype:

   void f(**int** x = 1, **int** y = 2, **int** z = 3);

allowable calls are:

positional arguments		**empty arguments**	
f();	*// rewrite ⇒ f(1, 2, 3)*	f(?, 4, 4);	*// rewrite ⇒ f(1, 4, 4)*
f(4);	*// rewrite ⇒ f(4, 2, 3)*	f(4, ?, 4);	*// rewrite ⇒ f(4, 2, 4)*
f(4, 4);	*// rewrite ⇒ f(4, 4, 3)*	f(4, 4, ?);	*// rewrite ⇒ f(4, 4, 3)*
f(4, 4, 4);	*// rewrite ⇒ f(4, 4, 4)*	f(4, ?, ?);	*// rewrite ⇒ f(4, 2, 3)*
		f(?, 4, ?);	*// rewrite ⇒ f(1, 4, 3)*
		f(?, ?, 4);	*// rewrite ⇒ f(1, 2, 4)*
		f(?, ?, ?);	*// rewrite ⇒ f(1, 2, 3)*

where the ? selects the default value as the argument. Here the missing arguments are inserted from the default values in the parameter list. The compiler rewrites missing default values into explicit positional arguments. The advantages of default values are:

  • Routines with a large number of parameters are often very generalized, giving a programmer a number of different options on how a computation is performed. For many of these routines, there are standard or default settings that work for the majority of computations. Without default values for parameters, a programmer is forced to specify these common values all the time, resulting in long argument lists that are error prone.

  • When a routine's interface is augmented with new parameters, it extends the interface providing generalizability[16] (somewhat like the generalization provided by inheritance for classes). That is, all existing calls are still

---

[15]Francez [15] proposed a further extension to the named-parameter passing style, which specifies what type of communication (by value, by reference, by name) the argument is passed to the routine.

[16]"It should be possible for the implementor of an abstraction to increase its generality. So long as the modified abstraction is a generalization of

valid, although the call must still be recompiled.

The only disadvantage of default arguments is that unintentional omission of an argument may not result in a compiler-time error. Instead, a default value is used, which may not be the programmer's intent.

Default parameters may only appear in a prototype versus definition context:

```
void f( int x, int y = 2, int z = 3 );        // prototype: allowed
void f( int, int = 2, int = 3 );              // prototype: allowed
void f( int x, int y = 2, int z = 3 ) {}      // definition: disallowed
```

The reason for this restriction is to allow separate compilation. Multiple prototypes with different default values is undefined.

Default arguments and overloading (see Section 26, p. 65) are complementary. While in theory default arguments can be simulated with overloading, as in:

default arguments	**overloading**
void f(int x, int y = 2, int z = 3) {...}	void f(int x, int y, int z) {...}
	void f(int x) { f(x, 2, 3); }
	void f(int x, int y) { f(x, y, 3); }

the number of required overloaded routines is linear in the number of default values, which is unacceptable growth. In general, overloading is used over default parameters, if the body of the routine is significantly different. Furthermore, overloading cannot handle accessing default arguments in the middle of a positional list.

```
f( 1, ?, 5 );                                 // rewrite ⇒ f( 1, 2, 5 )
```

17.2 Named (or Keyword)

A named (keyword) parameter provides the ability to specify an argument to a routine call using the parameter name rather than the position of the parameter. For example, given the routine prototype:

```
void f( int ?x, int ?y, int ?z );
```

allowable calls are:

```
f( ?x = 3, ?y = 4, ?z = 5 );                  // rewrite ⇒ f( 3, 4, 5 )
f( ?y = 4, ?z = 5, ?x = 3 );                  // rewrite ⇒ f( 3, 4, 5 )
f( ?z = 5, ?x = 3, ?y = 4 );                  // rewrite ⇒ f( 3, 4, 5 )
f( ?x = 3, ?z = 5, ?y = 4 );                  // rewrite ⇒ f( 3, 4, 5 )
```

Here the ordering of the the parameters and arguments is unimportant, and the names of the parameters are used to associate argument values with the corresponding parameters. The compiler rewrites a named call into a positional call. Note, the syntax ?x = 3 is necessary for the argument, because x = 3 has an existing meaning, *i.e.*, assign 3 to x and pass the value of x. The advantages of named parameters are:

- Remembering the names of the parameters may be easier than the order in the routine definition.

- Parameter names provide documentation at the call site (assuming the names are descriptive).

- Changes can be made to the order or number of parameters without affecting the call (although the call must still be recompiled).

Named parameters may only appear in a prototype versus definition context:

```
void f( int  x, int ?y, int ?z );             // prototype: allowed
void f( int ?x, int , int ?z );               // prototype: allowed
void f( int x, int ?y, int ?z ) {}            // definition: disallowed
```

The reason for this restriction is to allow separate compilation. Multiple prototypes with different positional parameter names is an error.

The named parameter is not part of type resolution; only the type of the expression assigned to the named parameter affects type resolution.

```
int f( int ?i, int ?j );
int f( int ?i, double ?j );
```

---

the original, existing uses of the abstraction will not require change. It might be possible to modify an abstraction in a manner which is not a generalization without affecting existing uses, but, without inspecting the modules in which the uses occur, this possibility cannot be determined. This criterion precludes the addition of parameters, unless these parameters have default or inferred values that are valid for all possible existing applications." [8, p. 128]

```
1      f( ?j = 3, ?i = 4 );                              // 1st f
2      f( ?i = 7, ?j = 8.1 );                            // 2nd f
```

17.3   Mixed Default/Named

Default and named parameters can be intermixed and named parameters can have a default value. For example, given the routine prototype:

```
6      void f( int x, int y = 1, int ?z = 2 );
```

allowable calls are:

```
8      f( 3 );                                           // rewrite ⇒ f( 3, 1, 2 )
9      f( 3, 4 );                                        // rewrite ⇒ f( 3, 4, 2 )
10     f( 3, ?z = 5 );                                   // rewrite ⇒ f( 3, 1, 5 )
11     f( 3, 4, ?z = 5 );                                // rewrite ⇒ f( 3, 4, 5 )
12     f( ?z = 5, 3 );                                   // rewrite ⇒ f( 3, 1, 5 )
13     f( 3, ?z = 5, 4 );                                // rewrite ⇒ f( 3, 4, 5 )
```

Finally, the ellipse ("...") parameter must appear after positional and named parameters in a routine prototype.

```
15     void f( int i = 1, int ?j = 2, ... );
```

C∀ named and default arguments are backwards compatible with C. C⧺ only supports default parameters; Ada supports both named and default parameters.

18   Unnamed Structure Fields

C requires each field of a structure to have a name, except for a bit field associated with a basic type, *e.g.*:

```
20     struct {
21        int f1;                                        // named field
22        int f2 : 4;                                    // named field with bit field size
23        int : 3;                                       // unnamed field for basic type with bit field size
24        int ;                                          // disallowed, unnamed field
25        int *;                                         // disallowed, unnamed field
26        int (*)( int );                                // disallowed, unnamed field
27     };
```

This requirement is relaxed by making the field name optional for all field declarations; therefore, all the field declarations in the example are allowed. As for unnamed bit fields, an unnamed field is used for padding a structure to a particular size. A list of unnamed fields is also supported, *e.g.*:

```
31     struct {
32        int , , ;                                      // 3 unnamed fields
33     }
```

19   Nesting

Nesting of types and routines is useful for controlling name visibility (*name hiding*).

19.1   Type Nesting

C∀ allows type nesting, and type qualification of the nested types (see Figure 7), where as C hoists (refactors) nested types into the enclosing scope and has no type qualification. In the left example in C, types C, U and T are implicitly hoisted outside of type S into the containing block scope. In the right example in C∀, the types are not hoisted and accessed using the field-selection operator "." for type qualification, as does Java, rather than the C⧺ type-selection operator "::".

19.2   Routine Nesting

While C∀ does not provide object programming by putting routines into structures, it does rely heavily on locally nested routines to redefine operations at or close to a call site. For example, the C quick-sort is wrapped into the following polymorphic C∀ routine:

C Type Nesting	C Implicit Hoisting	C∀
| ```
struct S {
    enum C { R, G, B };
    struct T {
        union U { int i, j; };
        enum C c;
        short int i, j;
    };
    struct T t;
} s;

int fred() {
    s.t.c = R;
    struct T t = { R, 1, 2 };
    enum C c;
    union U u;
}
``` | ```
enum C { R, G, B };
union U { int i, j; };
struct T {
    enum C c;
    short int i, j;
};
struct S {
    struct T t;
} s;
``` | ```
struct S {
    enum C { R, G, B };
    struct T {
        union U { int i, j; };
        enum C c;
        short int i, j;
    };
    struct T t;
} s;

int fred() {
    s.t.c = S.R;   // type qualification
    struct S.T t = { S.R, 1, 2 };
    enum S.C c;
    union S.T.U u;
}
``` |

Figure 7: Type Nesting / Qualification

```
1   forall( T | { int ?<?( T, T ); } )
2   void qsort( const T * arr, size_t dimension );
```

3 which can be used to sort in ascending and descending order by locally redefining the less-than operator into greater-
4 than.

```
5    const unsigned int size = 5;
6    int ia[size];
7    ...                                    // assign values to array ia
8    qsort( ia, size );                     // sort ascending order using builtin ?<?
9    {
10       int ?<?( int x, int y ) { return x > y; } // nested routine
11       qsort( ia, size );                 // sort descending order by local redefinition
12   }
```

13 Nested routines are not first-class, meaning a nested routine cannot be returned if it has references to variables in
14 its enclosing blocks; the only exception is references to the external block of the translation unit, as these variables
15 persist for the duration of the program. The following program in undefined in C∀ (and Indexcgcc)

```
16   [* [int]( int )] foo() {               // int (* foo())( int )
17       int i = 7;
18       int bar( int p ) {
19           i += 1;                        // dependent on local variable
20           sout | i;
21       }
22       return bar;                        // undefined because of local dependence
23   }
24   int main() {
25       * [int]( int ) fp = foo();         // int (* fp)( int )
26       sout | fp( 3 );
27   }
```

28 because

29 Currently, there are no lambda expressions, *i.e.*, unnamed routines because routine names are very important to
30 properly select the correct routine.

20   Tuple

32 In C and C∀, lists of elements appear in several contexts, such as the parameter list of a routine call.

```
33   f( 2, x, 3 + i );                      // element list
```

A list of elements is called a *tuple*, and is different from a comma expression.

20.1   Multiple-Return-Value Functions

In C and most programming languages, functions return at most one value; however, many operations have multiple outcomes, some exceptional (see Section 10, p. 17). To emulate functions with multiple return values, *aggregation* and/or *aliasing* is used.

In the former approach, a record type is created combining all of the return values. For example, consider C's div function, which returns the quotient and remainder for a division of an integer value.

```
typedef struct { int quot, rem; } div_t; // from include stdlib.h
div_t div( int num, int den );
div_t qr = div( 13, 5 );                  // return quotient/remainder aggregate
printf( "%d %d\n", qr.quot, qr.rem );  // print quotient/remainder
```

This approach requires a name for the return type and fields, where naming is a common programming-language issue. That is, naming creates an association that must be managed when reading and writing code. While effective when used sparingly, this approach does not scale when functions need to return multiple combinations of types.

In the latter approach, additional return values are passed as pointer parameters. A pointer parameter is assigned inside the routine to emulate a return. For example, consider C's modf function, which returns the integral and fractional part of a floating value.

```
double modf( double x, double * i );   // from include math.h
double intp, frac = modf( 13.5, &intp ); // return integral and fractional components
printf( "%g %g\n", intp, frac );        // print integral/fractional components
```

This approach requires allocating storage for the return values, which complicates the call site with a sequence of variable declarations leading to the call. Also, while a disciplined use of **const** can give clues about whether a pointer parameter is used as an out parameter, it is not obvious from the routine signature whether the callee expects such a parameter to be initialized before the call. Furthermore, while many C routines that accept pointers are safe for a NULL argument, there are many C routines that are not null-safe. Finally, C does not provide a mechanism to state that a parameter is going to be used as an additional return value, which makes the job of ensuring that a value is returned more difficult for the compiler. Still, not every routine with multiple return values should be required to return an error code, and error codes are easily ignored, so this is not a satisfying solution. As with the previous approach, this technique can simulate multiple return values, but in practice it is verbose and error prone.

C∀ allows functions to return multiple values by extending the function declaration syntax. Multiple return values are declared as a comma-separated list of types in square brackets in the same location that the return type appears in standard C function declarations.

```
[ char, int, double ] f( ... );
```

The ability to return multiple values from a function requires a new syntax for the return statement. For consistency, the return statement in C∀ accepts a comma-separated list of expressions in square brackets.

```
return [ c, i, d ];
```

The expression resolution ensures the correct form is used depending on the values being returned and the return type of the current function. A multiple-returning function with return type T can return any expression that is implicitly convertible to T.

A common use of a function's output is input to another function. C∀ allows this case, without any new syntax; a multiple-returning function can be used in any of the contexts where an expression is allowed. When a function call is passed as an argument to another call, the best match of actual arguments to formal parameters is evaluated given all possible expression interpretations in the current scope.

```
void g( int, int );                  // 1
void g( double, double );            // 2
g( div( 13, 5 ) );                   // select 1
g( modf( 13.5 ) );                   // select 2
```

In this case, there are two overloaded g routines. Both calls to g expect two arguments that are matched by the two return values from div and modf. respectively, which are fed directly to the first and second parameters of g. As well, both calls to g have exact type matches for the two different versions of g, so these exact matches are chosen. When type matches are not exact, conversions are used to find a best match.

1   The previous examples can be rewritten passing the multiple returned-values directly to the printf function call.

2   [**int**, **int**] div(**int** x, **int** y);   *// from include stdlib*
3   printf(`"%d %d\n"`, div(13, 5));   *// print quotient/remainder*

4

5   [**double**, **double**] modf(**double** x);   *// from include math*
6   printf(`"%g %g\n"`, modf(13.5));   *// print integral/fractional components*

7   This approach provides the benefits of compile-time checking for appropriate return statements as in aggregation, but
8   without the required verbosity of declaring a new named type.

9   Finally, the addition of multiple-return-value functions necessitates a syntax for retaining the multiple values at the
10 call-site versus their temporary existence during a call. The simplest mechanism for retaining a return value in C is
11 variable assignment. By assigning the multiple return-values into multiple variables, the values can be retrieved later.
12 As such, C∀ allows assigning multiple values from a function into multiple variables, using a square-bracketed list of
13 lvalue expressions on the left side.

14   **int** quot, rem;
15   [quot, rem] = div(13, 5);   *// assign multiple variables*
16   printf(`"%d %d\n"`, quot, rem);   *// print quotient/remainder*

17 Here, the multiple return-values are matched in much the same way as passing multiple return-values to multiple
18 parameters in a call.

19   20.2   Expressions

20 Multiple-return-value functions provide C∀ with a new syntax for expressing a combination of expressions in the
21 return statement and a combination of types in a function signature. These notions are generalized to provide C∀ with
22 *tuple expression*s and *tuple type*s. A tuple expression is an expression producing a fixed-size, ordered list of values of
23 heterogeneous types. The type of a tuple expression is the tuple of the subexpression types, or a tuple type.

24   In C∀, a tuple expression is denoted by a comma-separated list of expressions enclosed in square brackets. For
25 example, the expression [5, `'x'`, 10.5] has type [**int**, **char**, **double**]. The previous expression has 3 *components*. Each
26 component in a tuple expression can be any C∀ expression, including another tuple expression. The order of evalu-
27 ation of the components in a tuple expression is unspecified, to allow a compiler the greatest flexibility for program
28 optimization. It is, however, guaranteed that each component of a tuple expression is evaluated for side-effects, even
29 if the result is not used. Multiple-return-value functions can equivalently be called *tuple-returning functions*.

30   20.3   Variables

31 The previous call of div still requires the preallocation of multiple return-variables in a manner similar to the aliasing
32 example. In C∀, it is possible to overcome this restriction by declaring a *tuple variable*.

33   [**int**, **int**] qr = div(13, 5);   *// initialize tuple variable*
34   printf(`"%d %d\n"`, qr);   *// print quotient/remainder*

35 It is now possible to match the multiple return-values to a single variable, in much the same way as aggregation. As
36 well, the components of the tuple value are passed as separate parameters to printf, allowing direct printing of tuple
37 variables. One way to access the individual components of a tuple variable is with assignment.

38   [quot, rem] = qr;   *// assign multiple variables*

39   In addition to variables of tuple type, it is also possible to have pointers to tuples, and arrays of tuples. Tuple
40 types can be composed of any types, except for array types, since array assignment is disallowed, which makes tuple
41 assignment difficult when a tuple contains an array.

42   [**double**, **int**] di;
43   [**double**, **int**] * pdi
44   [**double**, **int**] adi[10];

45 This examples declares a variable of type [**double**, **int**], a variable of type pointer to [**double**, **int**], and an array of ten
46 [**double**, **int**].

20.4   Indexing

It is also possible to access a single component of a tuple-valued expression without creating temporary variables. Given a tuple-valued expression e np and a compile-time constant integer i where $0 \leq i < n$, where n is the number of components in e, $e.i$ accesses the i^{th} component of e, *e.g.*:

```
[int, double] x;
[char *, int] f();
void g(double, int);
[int, double] * p;

int y = x.0;                  // access int component of x
y = f().1;                    // access int component of f
p->0 = 5;                     // access int component of tuple pointed-to by p
g( x.1, x.0 );                // rearrange x to pass to g
double z = [ x, f() ].0.1;    // access second component of first component of tuple expression
```

Tuple-index expressions can occur on any tuple-typed expression, including tuple-returning functions, square-bracketed tuple expressions, and other tuple-index expressions, provided the retrieved component is also a tuple. This feature was proposed for K-W C but never implemented [31, p. 45].

20.5   Flattening and Structuring

As evident in previous examples, tuples in C∀ do not have a rigid structure. In function call contexts, tuples support implicit flattening and restructuring conversions. Tuple flattening recursively expands a tuple into the list of its basic components. Tuple structuring packages a list of expressions into a value of tuple type.

```
int f(int, int);
int g([int, int]);
int h(int, [int, int]);
[int, int] x;
int y;

f(x);     // flatten
g(y, 10); // structure
h(x, y);  // flatten & structure
```

In C∀, each of these calls is valid. In the call to f, x is implicitly flattened so that the components of x are passed as the two arguments to f. For the call to g, the values y and 10 are structured into a single argument of type [**int**, **int**] to match the type of the parameter of g. Finally, in the call to h, x is flattened to yield an argument list of length 3, of which the first component of x is passed as the first parameter of h, and the second component of x and y are structured into the second argument of type [**int**, **int**]. The flexible structure of tuples permits a simple and expressive function-call syntax to work seamlessly with both single- and multiple-return-value functions, and with any number of arguments of arbitrarily complex structure.

In K-W C [5, 31], there were 4 tuple coercions: opening, closing, flattening, and structuring. Opening coerces a tuple value into a tuple of values, while closing converts a tuple of values into a single tuple value. Flattening coerces a nested tuple into a flat tuple, *i.e.*, it takes a tuple with tuple components and expands it into a tuple with only non-tuple components. Structuring moves in the opposite direction, *i.e.*, it takes a flat tuple value and provides structure by introducing nested tuple components.

In C∀, the design has been simplified to require only the two conversions previously described, which trigger only in function call and return situations. This simplification is a primary contribution of this thesis to the design of tuples in C∀. Specifically, the expression resolution algorithm examines all of the possible alternatives for an expression to determine the best match. In resolving a function call expression, each combination of function value and list of argument alternatives is examined. Given a particular argument list and function value, the list of argument alternatives is flattened to produce a list of non-tuple valued expressions. Then the flattened list of expressions is compared with each value in the function's parameter list. If the parameter's type is not a tuple type, then the current argument value is unified with the parameter type, and on success the next argument and parameter are examined. If the parameter's type is a tuple type, then the structuring conversion takes effect, recursively applying the parameter matching algorithm using the tuple's component types as the parameter list types. Assuming a successful unification, eventually the algorithm gets to the end of the tuple type, which causes all of the matching expressions to be consumed

1 and structured into a tuple expression. For example, in

2   **int** f(**int**, [**double**, **int**]);
3   f([5, 10.2], 4);

4 There is only a single definition of f, and 3 arguments with only single interpretations. First, the argument alternative
5 list [5, 10.2], 4 is flattened to produce the argument list 5, 10.2, 4. Next, the parameter matching algorithm begins,
6 with $P =$ **int** and $A =$ **int**, which unifies exactly. Moving to the next parameter and argument, $P =$ [**double**, **int**] and
7 $A =$ **double**. This time, the parameter is a tuple type, so the algorithm applies recursively with $P' =$ **double** and
8 $A =$ **double**, which unifies exactly. Then $P' =$ **int** and $A =$ **double**, which again unifies exactly. At this point, the end
9 of P' has been reached, so the arguments 10.2, 4 are structured into the tuple expression [10.2, 4]. Finally, the end of
10 the parameter list P has also been reached, so the final expression is f(5, [10.2, 4]).

11 ## 20.6 Assignment

12 An assignment where the left side of the assignment operator has a tuple type is called *tuple assignment*. There are
13 two kinds of tuple assignment depending on whether the right side of the assignment operator has a non-tuple or tuple
14 type, called *mass* and *multiple* assignment, respectively.

15   **int** x;
16   **double** y;
17   [**int**, **double**] z;
18   [y, x] = 3.14;   *// mass assignment*
19   [x, y] = z;   *// multiple assignment*
20   z = 10;   *// mass assignment*
21   z = [x, y];   *// multiple assignment*

22 Let L_i for i in $[0,n)$ represent each component of the flattened left side, R_i represent each component of the flattened
23 right side of a multiple assignment, and R represent the right side of a mass assignment.

24   For a multiple assignment to be valid, both tuples must have the same number of elements when flattened. For
25 example, the following is invalid because the number of components on the left does not match the number of compo-
26 nents on the right.

27   [**int**, **int**] x, y, z;
28   [x, y] = z;   *// multiple assignment, invalid 4 != 2*

29 Multiple assignment assigns R_i to L_i for each i. That is, ?=?(&L_i, R_i) must be a well-typed expression. In the
30 previous example, [x, y] = z, z is flattened into z.0, z.1, and the assignments x = z.0 and y = z.1 happen.

31   A mass assignment assigns the value R to each L_i. For a mass assignment to be valid, ?=?(&L_i, R) must be
32 a well-typed expression. These semantics differ from C cascading assignment (*e.g.*, a=b=c) in that conversions are
33 applied to R in each individual assignment, which prevents data loss from the chain of conversions that can happen
34 during a cascading assignment. For example, [y, x] = 3.14 performs the assignments y = 3.14 and x = 3.14, which results
35 in the value 3.14 in y and the value 3 in x. On the other hand, the C cascading assignment y = x = 3.14 performs the
36 assignments x = 3.14 and y = x, which results in the value 3 in x, and as a result the value 3 in y as well.

37   Both kinds of tuple assignment have parallel semantics, such that each value on the left side and right side is
38 evaluated *before* any assignments occur. As a result, it is possible to swap the values in two variables without explicitly
39 creating any temporary variables or calling a function.

40   **int** x = 10, y = 20;
41   [x, y] = [y, x];

42 After executing this code, x has the value 20 and y has the value 10.

43   In C∀, tuple assignment is an expression where the result type is the type of the left side of the assignment, as
44 in normal assignment. That is, a tuple assignment produces the value of the left-hand side after assignment. These
45 semantics allow cascading tuple assignment to work out naturally in any context where a tuple is permitted. These
46 semantics are a change from the original tuple design in K-W C [31], wherein tuple assignment was a statement that
47 allows cascading assignments as a special case. Restricting tuple assignment to statements was an attempt to to fix
48 what was seen as a problem with side-effects, wherein assignment can be used in many different locations, such as in
49 function-call argument position. While permitting assignment as an expression does introduce the potential for subtle
50 complexities, it is impossible to remove assignment expressions from C∀ without affecting backwards compatibility.
51 Furthermore, there are situations where permitting assignment as an expression improves readability by keeping code
52 succinct and reducing repetition, and complicating the definition of tuple assignment puts a greater cognitive burden

on the user. In another language, tuple assignment as a statement could be reasonable, but it would be inconsistent for tuple assignment to be the only kind of assignment that is not an expression. In addition, K-W C permits the compiler to optimize tuple assignment as a block copy, since it does not support user-defined assignment operators. This optimization could be implemented in C∀, but it requires the compiler to verify that the selected assignment operator is trivial.

The following example shows multiple, mass, and cascading assignment used in one expression

```
int a, b;
double c, d;
[ void ] f( [ int, int ] );
f( [ c, a ] = [ b, d ] = 1.5 );            // assignments in parameter list
```

The tuple expression begins with a mass assignment of 1.5 into [b, d], which assigns 1.5 into b, which is truncated to 1, and 1.5 into d, producing the tuple [1, 1.5] as a result. That tuple is used as the right side of the multiple assignment (*i.e.*, [c, a] = [1, 1.5]) that assigns 1 into c and 1.5 into a, which is truncated to 1, producing the result [1, 1]. Finally, the tuple [1, 1] is used as an expression in the call to f.

20.7   Construction

Tuple construction and destruction follow the same rules and semantics as tuple assignment, except that in the case where there is no right side, the default constructor or destructor is called on each component of the tuple. As constructors and destructors did not exist in previous versions of C∀ or in K-W C, this is a primary contribution of this thesis to the design of tuples.

```
struct S;
void ?{}(S *);                        // (1)
void ?{}(S *, int);                   // (2)
void ?{}(S * double);                 // (3)
void ?{}(S *, S);                     // (4)

[S, S] x = [3, 6.28];                 // uses (2), (3), specialized constructors
[S, S] y;                             // uses (1), (1), default constructor
[S, S] z = x.0;                       // uses (4), (4), copy constructor
```

In this example, x is initialized by the multiple constructor calls ?{}(&x.0, 3) and ?{}(&x.1, 6.28), while y is initialized by two default constructor calls ?{}(&y.0) and ?{}(&y.1). z is initialized by mass copy constructor calls ?{}(&z.0, x.0) and ?{}(&z.1, x.0). Finally, x, y, and z are destructed, *i.e.*, the calls ^?{}(&x.0), ^?{}(&x.1), ^?{}(&y.0), ^?{}(&y.1), ^?{}(&z.0), and ^?{}(&z.1).

It is possible to define constructors and assignment functions for tuple types that provide new semantics, if the existing semantics do not fit the needs of an application. For example, the function **void** ?{}([T, U] *, S); can be defined to allow a tuple variable to be constructed from a value of type S.

```
struct S { int x; double y; };
void ?{}([int, double] * this, S s) {
    this−>0 = s.x;
    this−>1 = s.y;
}
```

Due to the structure of generated constructors, it is possible to pass a tuple to a generated constructor for a type with a member prefix that matches the type of the tuple. For example,

```
struct S { int x; double y; int z };
[int, double] t;
S s = t;
```

The initialization of s with t works by default because t is flattened into its components, which satisfies the generated field constructor ?{}(S *, **int**, **double**) to initialize the first two values.

20.8   Member-Access Expression

Tuples may be used to select multiple fields of a record by field name. The result is a single tuple-valued expression whose type is the tuple of the types of the members. For example,

```
struct S { char x; int y; double z; } s;
```

1   s.[x, y, z];

2   Here, the type of s.[x, y, z] is [**char**, **int**, **double**]. A member tuple expression has the form *e*.[x, y, z]; where *e* is an
3   expression with type T, where T supports member access expressions, and x, y, z are all members of T with types
4   T$_x$, T$_y$, and T$_z$ respectively. Then the type of *e*.[x, y, z] is [T$_x$, T$_y$, T$_z$].

5   A member-access tuple may be used anywhere a tuple can be used, *e.g.*:

6   s.[y, z, x] = [3, 3.2, 'x'];   // *equivalent to s.x = 'x', s.y = 3, s.z = 3.2*
7   f(s.[y, z]);   // *equivalent to f(s.y, s.z)*

8   Note, the fields appearing in a record-field tuple may be specified in any order; also, it is unnecessary to specify all the
9   fields of a struct in a multiple record-field tuple.

10 Since tuple-index expressions are a form of member-access expression, it is possible to use tuple-index expressions
11 in conjunction with member-access expressions to restructure a tuple (*e.g.*, rearrange components, drop components,
12 duplicate components, *etc.*).

13 [**int**, **int**, **long**, **double**] x;
14 **void** f(**double**, **long**);
15
16 f(x.[0, 3]);   // *f(x.0, x.3)*
17 x.[0, 1] = x.[1, 0];   // *[x.0, x.1] = [x.1, x.0]*
18 [**long**, **int**, **long**] y = x.[2, 0, 2];

19 It is possible for a member tuple expression to contain other member access expressions, *e.g.*:

20 **struct** A { **double** i; **int** j; };
21 **struct** B { **int** * k; **short** l; };
22 **struct** C { **int** x; A y; B z; } v;
23 v.[x, y.[i, j], z.k];

24 This expression is equivalent to [v.x, [v.y.i, v.y.j], v.z.k]. That is, the aggregate expression is effectively distributed
25 across the tuple allowing simple and easy access to multiple components in an aggregate without repetition. It is
26 guaranteed that the aggregate expression to the left of the . in a member tuple expression is evaluated exactly once. As
27 such, it is safe to use member tuple expressions on the result of a function with side-effects.

28 [**int**, **float**, **double**] f();
29 [**double**, **float**] x = f().[2, 1];   // *f() called once*

30 In K-W C, member tuple expressions are known as *record field tuples* [31]. Since C∀ permits these tuple-access
31 expressions using structures, unions, and tuples, *member tuple expression* or *field tuple expression* is more appropriate.

32 ## 20.9   Casting

33 Casting is a mechanism to explicitly change the type and representation of a value. If the type and representation are
34 changed, the cast is a *conversion*; if only the type is changed but not the value representation, the cast is a *coercion*.
35 For example, in:

36 **int** i, *ip;
37 **double** d;
38 d = (**double**)i;   // *conversion*
39 ip = (**int** *)d;   // *coercion*

40 the conversion cast implicitly runs code that transforms an integer representation into the best-effort floating-point
41 representation. Another conversion case exists in object-oriented programming-languages to walk an inheritance
42 hierarchy looking for specific types along the path. The coercion cast lies about the representation of the value as
43 the integer point is actually pointing at a floating-point value; indirect operations through ip are as odds with direct
44 operations on d. In general, coercion casts are only necessary for systems programming, like building a memory
45 allocator, where raw storage is typed and returned for use by the language or the runtime system to access storage in
46 special ways.

47 For coercion casts, there are often fine-grain variations to precisely expalin how the storage is to be typed. C++ and
48 C∀ have a number specialized casts., there are four types of explicit casting operators.

49 1. dynamic_cast Used for conversion of polymorphic types.

50 2. static_cast Used for conversion of nonpolymorphic types.

51 3. const_cast Used to remove the type qualifiers and possibly attributes.

1   4. reinterpret_cast Used for simple reinterpretation of bits.

2   `'(' type_no_function ')'` cast_expression
3   `'(' aggregate_control '&' ')'` cast_expression   *// CFA*
4   `'(' aggregate_control '*' ')'` cast_expression   *// CFA*
5   `'(' VIRTUAL ')'` cast_expression   *// CFA*
6   `'(' VIRTUAL type_no_function ')'` cast_expression   *// CFA*
7   `'(' RETURN type_no_function ')'` cast_expression   *// CFA (ASCRIPTION)*
8   `'(' COERCE type_no_function ')'` cast_expression   *// CFA (COERCION)*
9   `'(' qualifier_cast_list ')'` cast_expression   *// CFA, (modify CVs of cast_expression)*

10   Specialized Casts

11   There is some use in Cforall for cast operators with semantics other than the standard C cast. To make these
12 alternate casts look like the familiar C cast, this proposal follows the example of the virtual proposal's virtual cast
13 '(virtual Foo)x' and uses an added (pseudo-)keyword inside the cast parens.

14   C (Conversion) Cast

15   The standard C cast performs conversions, transformations between types which may make a new object with a
16 different in-memory representation. Cforall maintains these semantics in a backward-compatible way while account-
17 ing for name overloading by choosing the lowest-cost interpretation of the argument expression which is convertable
18 to the target type, breaking ties by conversion cost.

19   The C cast must be maintained for backward-compatibility, and developing a second cast operator with identical
20 semantics seems an undesirable multiplication of language features, but '(convert Foo)' or '(to Foo)' would be reason-
21 able options for a keyword. An alternate semantics for a Cforall-specific conversion cast would be to choose the cast
22 interpretation with the lowest sum of conversion cost and interpretation cost, which aligns better with Cforall function
23 call resolution algorithm.

24   Ascription Cast

25   Using casts in Cforall for type ascription ("select the interpretation of this type") works by the conversion-cost
26 tiebreaker behaviour of the cast operator. However, the ascription interpretation of casts is prioritized less than the
27 conversion interpretation of casts, sometimes resulting in some surprising results, as in the following example:

28   **int** f(**int**);   *// f1*
29   **int** f(**double**); *// f2*
30   **int** g(**int**);   *// g1*
31   **double** g(**long**); *// g2*
32
33   f((**double**)42); *// selects f2 by cast on argument*
34   (**double**)g(42); *// does NOT select g2, argument conversion cost results in g1*

35 An ascription cast which reversed the priorities of the C cast would be useful for selecting expressions based on
36 their return type; a reversal of the priorities of the standard C cast would work for this (that is, select the lowest-
37 cost conversion, breaking ties based on argument cost). A plausible stricter semantics would be to select the cheapest
38 interpretation with a zero-cost conversion to the target type, reporting a compiler error otherwise (this semantics would
39 make ascription a solely compile-time phenomenon, rather than relying on possible runtime conversions). A resonable
40 keyword would be '(as Foo)', which is short, evocative, and echos "ascription"; '(return Foo)' would not introduce
41 new keywords, and speaks to its use in return-type selection, as in the following corrected version of the example
42 above:

43   (as **double**)g(42); *// selects g2, as expected (under either presented ascription semantics)*

44   Coercion Cast

45   Some of the explicit conversions in C are defined to be a coercions (reinterpret the bits of this value as another
46 type). Use of coercions often relies on non-standard implementation details of the provided environment, and as such
47 is discouraged, but is sometimes necessary. Since all explicit pointer casts in C are coercions, any lvalue x in C/Cforall
48 can be coerced with the pattern *(Foo*)&x, but this is complex and doesn't extend to rvalues.

49   **int** i = 5;
50   **double** d = *(**double***)&i; *// value coercion*
51   printf(`"%g %g %x\n"`, d, *(**double** *)&i, *(**int** *)&d);
52
53   **int** i = 5; *// pointer coercion*
54   **double** d = *(**double***)&i; *// value coercion*

A dedicated coercion cast would solve these issues; (reinterpret Foo) (from C++), (transmute Foo) (from Rust), or (**coerce** Foo) would be reasonable keywords.

Qualifier Cast

A more restricted (and thus safer) form of coercion is modifiying the qualifiers of a type; C++ has const_cast for this purpose, and a similar feature would be useful for Cforall. With regard to syntax, (requalify **const** Foo)/(requalify Foo) to add/strip **const** would echo C++, but given that the vast majority of uses are stripping const-qualfiers, (non **const**) would be shorter, clearer, easily searchable, and not require the programmer to exactly match the argument type. In this syntax, coercion casts could be used to add qualifiers, or another cast type (say (**with const**)) could be introduced to add qualfiers.

Virtual Cast see virtual.txt; semantics equivalent to C++ dynamic cast

In C∀, the cast operator has a secondary use, which is type ascription, since it forces the expression resolution algorithm to choose the lowest cost conversion to the target type. That is, a cast can be used to select the type of an expression when it is ambiguous, as in the call to an overloaded function.

```
int f();                            // (1)
double f();                         // (2)

f();                                // ambiguous - (1),(2) both equally viable
(int)f();                           // choose (2)
```

Since casting is a fundamental operation in C∀, casts need to be given a meaningful interpretation in the context of tuples. Taking a look at standard C provides some guidance with respect to the way casts should work with tuples.

```
1   int f();
2   void g();

2
3   (void)f();                      // valid, ignore results
4   (int)g();                       // invalid, void cannot be converted to int

4
5   struct A { int x; };
6   (struct A)f();                  // invalid, int cannot be converted to A
```

In C, line 4 is a valid cast, which calls f and discards its result. On the other hand, line 5 is invalid, because g does not produce a result, so requesting an **int** to materialize from nothing is nonsensical. Finally, line 8 is also invalid, because in C casts only provide conversion between scalar types [21, p. 91]. For consistency, this implies that any case wherein the number of components increases as a result of the cast is invalid, while casts that have the same or fewer number of components may be valid.

Formally, a cast to tuple type is valid when $T_n \leq S_m$, where T_n is the number of components in the target type and S_m is the number of components in the source type, and for each i in $[0,n)$, S_i can be cast to T_i. Excess elements (S_j for all j in $[n,m)$) are evaluated, but their values are discarded so that they are not included in the result expression. This discarding naturally follows the way that a cast to void works in C.

For example,

```
[int, int, int] f();
[int, [int, int], int] g();

([int, double])f();                 // (1) valid
([int, int, int])g();               // (2) valid
([void, [int, int]])g();            // (3) valid
([int, int, int, int])g();          // (4) invalid
([int, [int, int, int]])g();        // (5) invalid
```

(1) discards the last element of the return value and converts the second element to type double. Since **int** is effectively a 1-element tuple, (2) discards the second component of the second element of the return value of g. If g is free of side effects, this is equivalent to [(**int**)(g().0), (**int**)(g().1.0), (**int**)(g().2)]. Since **void** is effectively a 0-element tuple, (3) discards the first and third return values, which is effectively equivalent to [(**int**)(g().1.0), (**int**)(g().1.1)]). if casts become function calls, what would they look like? would need a way to specify the target type, which seems awkward. Also, C++ basically only has this because classes are closed to extension, while we don't have that problem (can have floating constructors for any type). Note that a cast is not a function call in C∀, so flattening and structuring conversions do not occur for cast expressions. As such, (4) is invalid because the cast target type contains 4 components, while the

1   source type contains only 3. Similarly, (5) is invalid because the cast ([**int**, **int**, **int**])(g().1) is invalid. That is, it is invalid
2   to cast [**int**, **int**] to [**int**, **int**, **int**].

3   20.10   Polymorphism

4   Due to the implicit flattening and structuring conversions involved in argument passing, object and opaque param-
5   eters are restricted to matching only with non-tuple types. The integration of polymorphism, type assertions, and
6   monomorphic specialization of tuple-assertions are a primary contribution of this thesis to the design of tuples.

```
7       forall(T, U &)
8       void f(T x, U * y);
9
10      f([5, "hello"]);
```

11 In this example, [5, "hello"] is flattened, so that the argument list appears as 5, "hello". The argument matching
12 algorithm binds T to **int** and U to **const char**, and calls the function as normal.
13   Tuples can contain polymorphic types. For example, a plus operator can be written to add two triples of a type
14 together.

```
15      forall(T | { T ?+?(T, T); })
16      [T, T, T] ?+?([T, T, T] x, [T, T, T] y) {
17          return [x.0+y.0, x.1+y.1, x.2+y.2];
18      }
19      [int, int, int] x;
20      int i1, i2, i3;
21      [i1, i2, i3] = x + ([10, 20, 30]);
```

22 Note that due to the implicit tuple conversions, this function is not restricted to the addition of two triples. A call to
23 this plus operator type checks as long as a total of 6 non-tuple arguments are passed after flattening, and all of the
24 arguments have a common type that can bind to T, with a pairwise ?+? over T. For example, these expressions also
25 succeed and produce the same value.

```
26      ([x.0, x.1]) + ([x.2, 10, 20, 30]);  // x + ([10, 20, 30])
27      x.0 + ([x.1, x.2, 10, 20, 30]);      // x + ([10, 20, 30])
```

28 This presents a potential problem if structure is important, as these three expressions look like they should have
29 different meanings. Furthermore, these calls can be made ambiguous by introducing seemingly different functions.

```
30      forall(T | { T ?+?(T, T); })
31      [T, T, T] ?+?([T, T] x, [T, T, T, T]);
32      forall(T | { T ?+?(T, T); })
33      [T, T, T] ?+?(T x, [T, T, T, T, T]);
```

34 It is also important to note that these calls could be disambiguated if the function return types were different, as they
35 likely would be for a reasonable implementation of ?+?, since the return type is used in overload resolution. Still,
36 these semantics are a deficiency of the current argument matching algorithm, and depending on the function, differing
37 return values may not always be appropriate. These issues could be rectified by applying an appropriate conversion
38 cost to the structuring and flattening conversions, which are currently 0-cost conversions in the expression resolver.
39 Care would be needed in this case to ensure that exact matches do not incur such a cost.

```
40      void f([int, int], int, int);
41
42      f([0, 0], 0, 0);                      // no cost
43      f(0, 0, 0, 0);                        // cost for structuring
44      f([0, 0,], [0, 0]);                   // cost for flattening
45      f([0, 0, 0], 0);                      // cost for flattening and structuring
```

46   Until this point, it has been assumed that assertion arguments must match the parameter type exactly, modulo
47 polymorphic specialization (*i.e.*, no implicit conversions are applied to assertion arguments). This decision presents a
48 conflict with the flexibility of tuples.

49   20.10.1   Assertion Inference

```
50      int f([int, double], double);
51      forall(T, U | { T f(T, U, U); })
```

1   **void** g(T, U);
2   g(5, 10.21);

3 If assertion arguments must match exactly, then the call to g cannot be resolved, since the expected type of f is flat,
4 while the only f in scope requires a tuple type. Since tuples are fluid, this requirement reduces the usability of tuples
5 in polymorphic code. To ease this pain point, function parameter and return lists are flattened for the purposes of type
6 unification, which allows the previous example to pass expression resolution.

7   This relaxation is made possible by extending the existing thunk generation scheme, as described by Bilson [2].
8 Now, whenever a candidate's parameter structure does not exactly match the formal parameter's structure, a thunk is
9 generated to specialize calls to the actual function.

10   **int** _thunk(**int** _p0, **double** _p1, **double** _p2) {
11   **return** f([_p0, _p1], _p2);
12   }

13 Essentially, this provides flattening and structuring conversions to inferred functions, improving the compatibility of
14 tuples and polymorphism.

21   Tuples

16 In C and C∀, lists of elements appear in several contexts, such as the parameter list for a routine call. (More contexts
17 are added shortly.) A list of such elements is called a *lexical list*. The general syntax of a lexical list is:

18   [*exprlist*]

19 where *exprlist* is a list of one or more expressions separated by commas. The brackets, [], allow differentiating between
20 lexical lists and expressions containing the C comma operator. The following are examples of lexical lists:

21   [x, y, z]
22   [2]
23   [v + w, x * y, 3.14159, f()]

24 Tuples are permitted to contain sub-tuples (*i.e.*, nesting), such as [[14, 21], 9], which is a 2-element tuple whose first
25 element is itself a tuple. Note, a tuple is not a record (structure); a record denotes a single value with substructure,
26 whereas a tuple is multiple values with no substructure (see flattening coercion in Section 20.5, p. 42). In essence,
27 tuples are largely a compile time phenomenon, having little or no runtime presence.

28   Tuples can be organized into compile-time tuple variables; these variables are of *tuple type*. Tuple variables and
29 types can be used anywhere lists of conventional variables and types can be used. The general syntax of a tuple type
30 is:

31   [*typelist*]

32 where *typelist* is a list of one or more legal C∀ or C type specifications separated by commas, which may include other
33 tuple type specifications. Examples of tuple types include:

34   [**unsigned int**, **char**]
35   [**double**, **double**, **double**]
36   [* **int**, **int** *]   // mix of CFA and ANSI
37   [* [5] **int**, * * **char**, * [[**int**, **int**]] (**int**, **int**)]

38 Like tuples, tuple types may be nested, such as [[**int**, **int**], **int**], which is a 2-element tuple type whose first element is
39 itself a tuple type.

40   Examples of declarations using tuple types are:

41   [**int**, **int**] x;   // 2 element tuple, each element of type int
42   * [**char**, **char**] y;   // pointer to a 2 element tuple
43   [[**int**, **int**]] z ([**int**, **int**]);

44 The last example declares an external routine that expects a 2 element tuple as an input parameter and returns a 2
45 element tuple as its result.

46   As mentioned, tuples can appear in contexts requiring a list of value, such as an argument list of a routine call.
47 In unambiguous situations, the tuple brackets may be omitted, *e.g.*, a tuple that appears as an argument may have its
48 square brackets omitted for convenience; therefore, the following routine invocations are equivalent:

49   f([1, x+2, fred()]);
50   f(1, x+2, fred());

Also, a tuple or a tuple variable may be used to supply all or part of an argument list for a routine expecting multiple input parameters or for a routine expecting a tuple as an input parameter. For example, the following are all legal:

```
[ int, int ] w1;
[ int, int, int ] w2;
[ void ] f (int, int, int);          // three input parameters of type int
[ void ] g ([ int, int, int ]);      // 3 element tuple as input
f( [ 1, 2, 3 ] );
f( w1, 3 );
f( 1, w1 );
f( w2 );
g( [ 1, 2, 3 ] );
g( w1, 3 );
g( 1, w1 );
g( w2 );
```

Note, in all cases 3 arguments are supplied even though the syntax may appear to supply less than 3. As mentioned, a tuple does not have structure like a record; a tuple is simply converted into a list of components.

□   The present implementation of C∀ does not support nested routine calls when the inner routine returns multiple values; *i.e.*, a statement such as g(f()) is not supported. Using a temporary variable to store the results of the inner routine and then passing this variable to the outer routine works, however.   □

A tuple can contain a C comma expression, provided the expression containing the comma operator is enclosed in parentheses. For instance, the following tuples are equivalent:

```
[ 1, 3, 5 ]
[ 1, (2, 3), 5 ]
```

The second element of the second tuple is the expression (2, 3), which yields the result 3. This requirement is the same as for comma expressions in argument lists.

Type qualifiers, *i.e.*, **const** and **volatile**, may modify a tuple type. The meaning is to distribute the qualifier across all of the types in the tuple, *e.g.*:

```
const volatile [ int, float, const int ] x;
```

is equivalent to:

```
[ const volatile int, const volatile float, const volatile int ] x;
```

Declaration qualifiers can only appear at the start of a CV tuple declaration4, *e.g.*:

```
extern [ int, int ] w1;
static [ int, int, int ] w2;
```

□   Unfortunately, C's syntax for subscripts precluded treating them as tuples. The C subscript list has the form [i][j]... and not [i, j, ...]. Therefore, there is no syntactic way for a routine returning multiple values to specify the different subscript values, *e.g.*, f[g()] always means a single subscript value because there is only one set of brackets. Fixing this requires a major change to C because the syntactic form M[i, j, k] already has a particular meaning: i, j, k is a comma expression.   □

21.1   Tuple Coercions

There are four coercions that can be performed on tuples and tuple variables: closing, opening, flattening and structuring. In addition, the coercion of dereferencing can be performed on a tuple variable to yield its value(s), as for other variables. A *closing coercion* takes a set of values and converts it into a tuple value, which is a contiguous set of values, as in:

```
[ int, int, int, int ] w;
w = [ 1, 2, 3, 4 ];
```

First the right-hand tuple is closed into a tuple value and then the tuple value is assigned.

An *opening coercion* is the opposite of closing; a tuple value is converted into a tuple of values, as in:

```
[ a, b, c, d ] = w
```

w is implicitly opened to yield a tuple of four values, which are then assigned individually.

A *flattening coercion* coerces a nested tuple, *i.e.*, a tuple with one or more components, which are themselves tuples, into a flattened tuple, which is a tuple whose components are not tuples, as in:

    [a, b, c, d] = [1, [2, 3], 4];

First the right-hand tuple is flattened and then the values are assigned individually. Flattening is also performed on tuple types. For example, the type [**int**, [**int**, **int**], **int**] can be coerced, using flattening, into the type [**int**, **int**, **int**, **int**]. A *structuring coercion* is the opposite of flattening; a tuple is structured into a more complex nested tuple. For example, structuring the tuple [1, 2, 3, 4] into the tuple [1, [2, 3], 4] or the tuple type [**int**, **int**, **int**, **int**] into the tuple type [**int**, [**int**, **int**], **int**]. In the following example, the last assignment illustrates all the tuple coercions:

    [**int**, **int**, **int**, **int**] w = [1, 2, 3, 4];
    int x = 5;
    [x, w] = [w, x];   // all four tuple coercions

Starting on the right-hand tuple in the last assignment statement, w is opened, producing a tuple of four values; therefore, the right-hand tuple is now the tuple [[1, 2, 3, 4], 5]. This tuple is then flattened, yielding [1, 2, 3, 4, 5], which is structured into [1, [2, 3, 4, 5]] to match the tuple type of the left-hand side. The tuple [2, 3, 4, 5] is then closed to create a tuple value. Finally, x is assigned 1 and w is assigned the tuple value using multiple assignment (see Section 20.6, p. 43).

☐   A possible additional language extension is to use the structuring coercion for tuples to initialize a complex record with a tuple.   ☐

21.2   Mass Assignment

C∀ permits assignment to several variables at once using mass assignment [25]. Mass assignment has the following form:

    [*lvalue*, ... , *lvalue*] = *expr*;

The left-hand side is a tuple of *lvalue*s, which is a list of expressions each yielding an address, *i.e.*, any data object that can appear on the left-hand side of a conventional assignment statement. *expr* is any standard arithmetic expression. Clearly, the types of the entities being assigned must be type compatible with the value of the expression.

Mass assignment has parallel semantics, *e.g.*, the statement:

    [x, y, z] = 1.5;

is equivalent to:

    x = 1.5; y = 1.5; z = 1.5;

This semantics is not the same as the following in C:

    x = y = z = 1.5;

as conversions between intermediate assignments may lose information. A more complex example is:

    [i, y[i], z] = a + b;

which is equivalent to:

    t = a + b;
    a1 = &i; a2 = &y[i]; a3 = &z;
    *a1 = t; *a2 = t; *a3 = t;

The temporary t is necessary to store the value of the expression to eliminate conversion issues. The temporaries for the addresses are needed so that locations on the left-hand side do not change as the values are assigned. In this case, y[i] uses the previous value of i and not the new value set at the beginning of the mass assignment.

21.3   Multiple Assignment

C∀ also supports the assignment of several values at once, known as multiple assignment [25, 16]. Multiple assignment has the following form:

    [*lvalue*, ... , *lvalue*] = [*expr*, ... , *expr*];

The left-hand side is a tuple of *lvalue*s, and the right-hand side is a tuple of *expr*s. Each *expr* appearing on the right-hand side of a multiple assignment statement is assigned to the corresponding *lvalue*s on the left-hand side of the statement using parallel semantics for each assignment. An example of multiple assignment is:

1   [x, y, z] = [1, 2, 3];

2 Here, the values 1, 2 and 3 are assigned, respectively, to the variables x, y and z. A more complex example is:

3   [i, y[i], z] = [1, i, a + b];

4 Here, the values 1, i and a + b are assigned to the variables i, y[i] and z, respectively. Note, the parallel semantics of
5 multiple assignment ensures:

6   [x, y] = [y, x];

7 correctly interchanges (swaps) the values stored in x and y. The following cases are errors:

8   [a, b, c] = [1, 2, 3, 4];
9   [a, b, c] = [1, 2];

10 because the number of entities in the left-hand tuple is unequal with the right-hand tuple.

11   As for all tuple contexts in C, side effects should not be used because C does not define an ordering for the
12 evaluation of the elements of a tuple; both these examples produce indeterminate results:

13   f(x++, x++);   // C routine call with side effects in arguments
14   [v1, v2] = [x++, x++];   // side effects in right-hand side of multiple assignment

15 **21.4   Cascade Assignment**

16 As in C, C∀ mass and multiple assignments can be cascaded, producing cascade assignment. Cascade assignment has
17 the following form:

18   *tuple* = *tuple* = ... = *tuple*;

19 and it has the same parallel semantics as for mass and multiple assignment. Some examples of cascade assignment
20 are:

21   x1 = y1 = x2 = y2 = 0;
22   [x1, y1] = [x2, y2] = [x3, y3];
23   [x1, y1] = [x2, y2] = 0;
24   [x1, y1] = z = 0;

25 As in C, the rightmost assignment is performed first, *i.e.*, assignment parses right to left.

26 # 22   Stream I/O Library

27 The goal of C∀ stream input/output (I/O) is to simplify the common cases, while fully supporting polymorphism
28 and user defined types in a consistent way. Stream I/O can be implicitly or explicitly formatted. Implicit formatting
29 means C∀ selects an I/O format for values that matches a variable's type. Explicit formatting means additional I/O
30 information is specified to control how a value is interpreted.

31   C∀ formatting incorporates ideas from C printf, C++ stream manipulators, and Python implicit spacing and newline.
32 Specifically:

33   • printf/Python format codes are dense, making them difficult to read and remember. C∀/C++ format manipulators
34   are named, making them easier to read and remember.

35   • printf/Python separate format codes from associated variables, making it difficult to match codes with variables.
36   C∀/C++ co-locate codes with associated variables, where C∀ has the tighter binding.

37   • Format manipulators in printf/Python/C∀ have local effect, whereas C++ have global effect, except setw. Hence,
38   it is common C++ programming practice to toggle manipulators on and then back to the default to prevent
39   downstream side-effects. Without this programming style, errors occur when moving prints, as manipulator
40   effects incorrectly flow into the new location. Furthermore, to guarantee no side-effects, manipulator values
41   must be saved and restored across function calls. C++ programers never do any of this.

42   • C∀ has more sophisticated implicit value spacing than Python, plus implicit newline at the end of a print.

43 **22.1   Basic I/O**

44 The standard polymorphic I/O streams are stdin/sin (input), stdout/sout, and stderr/serr (output) (like C++ cin/cout/cerr).
45 The standard I/O operator is the bit-wise (or) operator, ' | ', which is used to cascade multiple I/O operations. The C∀
46 header file for the I/O library is fstream.hfa.

22.1.1   Stream Output

For implicit formatted output, the common case is printing a series of variables separated by whitespace.

C∀	C++	Python
int x = 1, y = 2, z = 3;		x = 1; y = 2; z = 3
sout | x | y | z;	cout << x << " " << y << " " << z << endl;	print(x, y, z)
1␣2␣3	1␣2␣3	1␣2␣3

The C∀ form has half the characters of the C++ form, and is similar to Python I/O with respect to implicit separators and newline. Similar simplification occurs for tuple I/O, which flattens the tuple and prints each value separated by "␣" (comma space).

```
[int, [ int, int ] ] t1 = [ 1, [ 2, 3 ] ], t2 = [ 4, [ 5, 6 ] ];
sout | t1 | t2;                    // print tuples
1,␣2,␣3␣4,␣5,␣6
```

The bit-wise | operator is used for I/O, rather C++ shift-operators, << and >>, as it is the lowest-priority *overloadable* operator, other than assignment. (Operators || and && are not overloadable in C∀.) Therefore, fewer output expressions require parenthesis.

C∀:   sout | x * 3 | y + 1 | z << 2 | x == y | (x | y) | (x || y) | (x > z ? 1 : 2);

C++:   cout << x * 3 << y + 1 << (z << 2) << (x == y) << (x | y) << (x || y) << (x > z ? 1 : 2) << endl;

3␣3␣12␣0␣3␣1␣2

There is a weak similarity between the C∀ logical-or operator and the Shell pipe-operator for moving data, where data flows in the correct direction for input but the opposite direction for output. Input and output use a uniform operator, |, rather than C++'s << and >> input/output operators to prevent this common error in C++:

cin << i; // *why is this generating a lot of error messages?*

Streams exit and abort provide output with immediate program termination without and with generating a stack trace and core file. Stream exit implicitly returns EXIT_FAILURE to the shell.

```
exit  | "x (" | x | ") negative value.";   // print, terminate, and return EXIT_FAILURE to shell
abort | "x (" | x | ") negative value.";   // print, terminate, and generate stack trace and core file
```

Note, C∀ stream variables stdin, stdout, stderr, exit, and abort overload C variables stdin, stdout, stderr, and functions exit and abort, respectively.

22.1.2   Stream Input

For implicit formatted input, the common case is reading a sequence of values separated by whitespace, where the type of an input constant must match with the type of the input variable.

char c;   **int** i;   **double** d

C∀	C++	Python
sin | c | i | d;	cin >> c >> i >> d;	c = input(); i = **int**(input()); d = **float**(input());
A␣1␣2.5	A␣1␣2.5	A
		1
		2.5

The format of numeric input values in the same as C constants without a trailing type suffix, as the input value-type is denoted by the input variable. For **bool** type, the constants are true and false. For integral types, any number of digits, optionally preceded by a sign (+ or –), where a

- 1-9 prefix introduces a decimal value (0-9),
- 0 prefix introduces an octal value (0-7), and
- 0x or 0X prefix introduces a hexadecimal value (0-f) with lower or upper case letters.

For floating-point types, any number of decimal digits, optionally preceded by a sign (+ or –), optionally containing a decimal point, and optionally followed by an exponent, e or E, with signed (optional) decimal digits. Floating-point values can also be written in hexadecimal format preceded by 0x or 0X with hexadecimal digits and exponent denoted by p or P. In all cases, whitespace characters are skipped until an appropriate value is found.

char ch;   **int** i;   **float** f; **double** d;   **_Complex double** cxd;

```
1    sin | ch | i | f | d | cxd;
2    X  42  1234.5   0xfffp−2   3.5+7.1i
```

3 It is also possible to scan and ignore specific strings and whitespace using a string format.

```
4    sin | "abc def";                    // space matches arbitrary whitespace (2 blanks, 2 tabs)
5    abc␣␣␣⟶·⟶def
```

6 A non-whitespace format character reads the next input character, compares the format and input characters, and if
7 equal, the input character is discarded and the next format character is tested. Note, a single whitespace in the format
8 string matches **any** quantity of whitespace characters from the stream (including none).

9   For the C-string type, the default input format is any number of **non-whitespace** characters. There is no escape
10 character supported in an input string, but any Latin-1 character can be typed directly in the input string. For example,
11 if the following non-whitespace output is redirected into a file by the shell:

```
12   sout | "\n\t\f\0234\x23";
```

13 it can be read back from the file by redirecting the file as input using:

```
14   char s[64];
15   sin | wdi( sizeof(s), s );             // must specify string size
```

16 The input string is always null terminated '\0' in the input variable. Because of potential buffer overrun when reading
17 C strings, strings are restricted to work with input manipulators (see Section 22.6, p. 60). As well, there are multiple
18 input-manipulators for scanning complex input string formats, *e.g.*, a quoted character or string.

19   **In all cases, if an invalid data value is not found for a type or format string, the exception** missing_data **is**
20 **raised and the input variable is unchanged.** For example, when reading an integer and the string "abc" is found,
21 the exception missing_data is raised to ensure the program does not proceed erroneously. If a valid data value is found,
22 but it is larger than the capacity of the input variable, such reads are undefined.

23 ### 22.1.3   Stream Files

24 Figure 8 shows the I/O stream operations for interacting with files other than sin, sout, and cerr.

25   • fail tests the stream error-indicator, returning nonzero if it is set.

26   • clear resets the stream error-indicator.

27   • flush (ofstream only) causes any unwritten data for a stream to be written to the file.

28   • eof (ifstream only) tests the end-of-file indicator for the stream pointed to by stream. Returns true if the end-of-
29   file indicator is set, otherwise false.

30   • open binds the file with name to a stream accessed with mode (see fopen).

31   • close flushes the stream and closes the file.

32   • write (ofstream only) writes size bytes to the stream. The bytes are written lazily when an internal buffer fills.
33   Eager buffer writes are done with flush

34   • read (ifstream only) reads size bytes from the stream.

35   • ungetc (ifstream only) pushes the character back to the input stream. Pushed-back characters returned by subse-
36   quent reads in the reverse order of pushing.

37 The constructor functions:

38   • create an unbound stream, which is subsequently bound to a file with open.

39   • create a bound stream to the associated file with given mode.

40 The destructor closes the stream.

41   Figure 9, p. 56 demonstrates the file operations by showing the idiomatic C∀ command-line processing and copying
42 an input file to an output file. Note, a stream variable may be copied because it is a reference to an underlying stream
43 data-structures. **All unusual I/O cases are handled as exceptions, including end-of-file.**

44 ## 22.2   Implicit Separator

45 The implicit separator character (space/blank) is a separator not a terminator for output. The rules for implicitly adding
46 a separator are:

47   1. A separator does not appear at the start or end of a line.

```
48       sout | 1 | 2 | 3;
```

```
// ******************************** ofstream ************************************
bool fail( ofstream & );
void clearerr( ofstream & );
int flush( ofstream & );
void open( ofstream &, const char name[], const char mode[] = "w" );
void close( ofstream & );
ofstream & write( ofstream &, const char data[], size_t size );
void ?{}( ofstream & );
void ?{}( ofstream &, const char name[], const char mode[] = "w" );
void ^?{}( ofstream & );

// ******************************** ifstream ************************************
bool fail( ifstream & is );
void clearerr( ifstream & );
bool eof( ifstream & is );
void open( ifstream & is, const char name[], const char mode[] = "r" );
void close( ifstream & is );
ifstream & read( ifstream & is, char data[], size_t size );
ifstream & ungetc( ifstream & is, char c );
void ?{}( ifstream & is );
void ?{}( ifstream & is, const char name[], const char mode[] = "r" );
void ^?{}( ifstream & is );
```

Figure 8: I/O Stream Functions

1. 1␣2␣3

2. A separator does not appear before or after a character literal or variable.

    ```
    sout | '1' | '2' | '3';
    123
    ```

3. A separator does not appear before or after a null (empty) C string, which is a local mechanism to disable insertion of the separator character.

    ```
    sout | 1 | "" | 2 | "" | 3;
    123
    ```

4. A separator does not appear before a C string starting with the (extended) ASCII characters: , . ; ! ?)] } %¢», where » is a closing citation mark.

    ```
    sout | 1 | ", x" | 2 | ". x" | 3 | "; x" | 4 | "! x" | 5 | "? x" | 6 | "% x"
         | 7 | "¢ x" | 8 | "» x" | 9 | ") x" | 10 | "] x" | 11 | "} x";
    ```

    Input1,␣x␣2.␣x␣3;␣x␣4!␣x␣5?␣x␣6%␣x␣7¢␣x␣8»␣x␣9)␣x␣10]␣x␣11}␣x

5. A separator does not appear after a C string ending with the (extended) ASCII characters: ([{ = $ £ ¥ ¡ ¿ «, where ¡ ¿ are inverted opening exclamation and question marks, and « is an opening citation mark.

    ```
    sout | "x (" | 1 | "x [" | 2 | "x {" | 3 | "x =" | 4 | "x $" | 5 | "x £" | 6 | "x ¥"
         | 7 | "x ¡" | 8 | "x ¿" | 9 | "x «" | 10;
    ```

    x␣(1␣x␣[2␣x␣{3␣x␣=4␣x␣$5␣x␣£6␣x␣¥7␣x␣¡8␣x␣¿9␣x␣«10

6. A separator does not appear before/after a C string starting/ending with the ASCII quote or whitespace characters: ` ' " : ␣ \t \v \f \r \n

    ```
    sout | "x`" | 1 | "`x'" | 2 | "'x\"" | 3 | "\"x:" | 4 | ":x " | 5 | " x\t" | 6 | "\tx";
    ```

    x`1`x'2'x"3"x:4:x␣5␣x␣␣␣6␣␣␣x

7. If a space is desired before or after one of the special string start/end characters, explicitly insert a space.

    ```
    sout | "x (␣" | 1 | "␣) x" | 2 | "␣, x" | 3 | "␣:x:␣" | 4;
    ```

    x␣(␣1␣)␣x␣2␣␣x␣3␣:x:␣4

```
#include <fstream.hfa>

int main( int argc, char * argv[] ) {
    ifstream in  = stdin;                  // copy default files
    ofstream out = stdout;

    try {
        choose ( argc ) {
          case 3, 2:
            open( in, argv[1] );           // open input file first as output creates file
            if ( argc == 3 ) open( out, argv[2] ); // do not create output unless input opens
          case 1: ;                        // use default files
          default:
            exit | "Usage" | argv[0] | "[ input−file (default stdin) "
                "[ output−file (default stdout) ] ]";
        } // choose
    } catch( open_failure * ex; ex−>istream == &in ) { // input file errors
        exit | "Unable to open input file" | argv[1];
    } catch( open_failure * ex; ex−>ostream == &out ) { // output file errors
        close( in );                       // optional
        exit | "Unable to open output file" | argv[2];
    } // try

    out | nlOff;                           // turn off auto newline
    in | nlOn;                             // turn on reading newline
    char ch;
    try {
        for () {                           // read/write characters
            in | ch;
            out | ch;
        } // for
    } catch( end_of_file * ) {             // end-of-file raised
    } // try
} // main
```

Figure 9: C∀ Command-Line Processing

22.3   Separation Manipulators

The following manipulators control implicit output separation. The effect of these manipulators is global for an output stream (except sep and nosep).

1. sepSet and sepVal/sepGet set and get the separator string. The separator string can be at most 16 characters including the `'\0'` string terminator (15 printable characters).

```
sepSet( sout, ",  $" );          // set separator from " " to ", $"
sout | 1 | 2 | 3 | " \"" | sepVal | "\"";
1,␣$2,␣$3␣",␣$"

sepSet( sout, "  " );            // reset separator to " "
sout | 1 | 2 | 3 | " \"" | sepGet( sout ) | "\"";
1␣2␣3␣"␣"
```

sepGet can be used to store a separator and then restore it:

```
char store[sepSize];             // sepSize is the maximum separator size
strcpy( store, sepGet( sout ) ); // copy current separator
sepSet( sout, "_" );             // change separator to underscore
sout | 1 | 2 | 3;
1_2_3

sepSet( sout, store );           // change separator back to original
```

```
1        sout | 1 | 2 | 3;
2        1␣2␣3
```

2. sepSetTuple and sepTupleVal/sepGetTuple get and set the tuple separator-string. The tuple separator-string can be at most 16 characters including the `'\0'` string terminator (15 printable characters).

```
5        sepSetTuple( sout, "  " );         // set tuple separator from ", " to " "
6        sout | t1 | t2 | "  \"" | sepTupleVal | "\"";
7        1␣2␣3␣4␣5␣6␣"␣"
8        sepSetTuple( sout, ",  " );        // reset tuple separator to ", "
9        sout | t1 | t2 | "  \"" | sepGetTuple( sout ) | "\"";
10       1,␣2,␣3␣4,␣5,␣6␣",␣"
```

As for sepGet, sepGetTuple can be use to store a tuple separator and then restore it.

3. sepOff and sepOn globally toggle printing the separator.

```
13       sout | sepOff | 1 | 2 | 3;          // turn off implicit separator
14       123
15       sout | sepOn | 1 | 2 | 3;           // turn on implicit separator
16       1␣2␣3
```

4. sep and nosep locally toggle printing the separator with respect to the next printed item, and then return to the global separator setting.

```
19       sout | 1 | nosep | 2 | 3;           // turn off implicit separator for the next item
20       12␣3
21       sout | sepOff | 1 | sep | 2 | 3;    // turn on implicit separator for the next item
22       1␣23
```

The tuple separator also responses to being turned on and off.

```
24       sout | t1 | nosep | t2;             // turn off implicit separator for the next item
25       1,␣2,␣34,␣5,␣6
```

sep *cannot* be used to start/end a line with a separator because separators do not appear at the start/end of a line. Use sep to accomplish this functionality.

```
28       sout | sep | 1 | 2 | 3 | sep;       // sep does nothing at start/end of line
29       1␣2␣3
30       sout | sepVal | 1 | 2 | 3 | sepVal ; // use sepVal to print separator at start/end of line
31       ␣1␣2␣3␣
```

22.4   Newline Manipulators

The following manipulators control newline separation for input and output.

For input:

1. nlOn reads the newline character, when reading single characters.

2. nlOff does *not* read the newline character, when reading single characters.

3. nl scans characters until the next newline character, *i.e.*, ignore the remaining characters in the line. If nlOn is enabled, the nl is also consumed.

For example, in:

```
40       int i, j;
41       sin | i | nl | j;
42       1 2
43       3
```

variable i is assigned 1, the 2 is skipped, and variable j is assigned 3. For example, in:

```
         char ch
                                      sin | nlOn; // enable reading newlines
45       sin | ch; // read X          sin | ch; // read newline

         X
```

the left example skips the newline and reads `'X'` into ch, while the right example reads the newline into ch.

For output:

1. nlOn implicitly prints a newline at the end of each output expression.
2. nlOff does *not* implicitly print a newline at the end of each output expression.
3. nl inserts a newline.

```
sout | nl;                      // only print newline
sout | 2;                       // implicit newline
sout | 3 | nl | 4 | nl;         // terminating nl merged with implicit newline
sout | 5 | nl | nl;             // again terminating nl merged with implicit newline
sout | 6;                       // implicit newline


2

3

4

5


6
```

Note, a terminating nl is merged with (overrides) the implicit newline at the end of the sout expression, otherwise it is impossible to print a single newline

22.5   Output Manipulators

The following manipulators control formatting (printing) of the argument output values.

1. bin(integer) print value in base 2 preceded by 0b/0B.

```
sout | bin( 0 ) | bin( 27HH ) | bin( 27H ) | bin( 27 ) | bin( 27L );
0b0 0b11011 0b11011 0b11011 0b11011
sout | bin( −27HH ) | bin( −27H ) | bin( −27 ) | bin( −27L );
0b11100101 0b1111111111100101 0b11111111111111111111111111100101 0b(58 1s)100101
```

2. oct(integer) print value in base 8 preceded by 0.

```
sout | oct( 0 ) | oct( 27HH ) | oct( 27H ) | oct( 27 ) | oct( 27L );
0 033 033 033 033
sout | oct( −27HH ) | oct( −27H ) | oct( −27 ) | oct( −27L );
0345 0177745 037777777745 01777777777777777777745
```

Note, octal 0 is *not* preceded by 0 to prevent confusion.

3. hex(integer / floating-point) print value in base 16 preceded by 0x/0X.

```
sout | hex( 0 ) | hex( 27HH ) | hex( 27H ) | hex( 27 ) | hex( 27L );
0x0 0x1b 0x1b 0x1b 0x1b
sout | hex( −27HH ) | hex( −27H ) | hex( −27 ) | hex( −27L );
0xe5 0xffe5 0xffffffe5 0xffffffffffffffe5

sout | hex( 0.0 ) | hex( 27.5F ) | hex( 27.5 ) | hex( 27.5L );
0x0p+0 0x1.b8p+4 0x1.b8p+4 0xd.cp+1
sout | hex( −27.5F ) | hex( −27.5 ) | hex( −27.5L );
−0x1.b8p+4 −0x1.b8p+4 −0xd.cp+1
```

4. sci(floating-point) print value in scientific notation with exponent. Default is 6 digits of precision.

```
sout | sci( 0.0 ) | sci( 27.5 ) | sci( −27.5 );
0.000000e+00 2.750000e+01 −2.750000e+01
```

5. eng(floating-point) print value in engineering notation with exponent, which means the exponent is adjusted to a multiple of 3.

```
sout | eng( 0.0 ) | eng( 27000.5 ) | eng( −27.5e7 );
0e0 27.0005e3 −275e6
```

6. unit(engineering-notation) print engineering exponent as a letter between the range 10^{-24} and 10^{24}:

y $\Rightarrow 10^{-24}$, z $\Rightarrow 10^{-21}$, a $\Rightarrow 10^{-18}$, f $\Rightarrow 10^{-15}$, p $\Rightarrow 10^{-12}$, n $\Rightarrow 10^{-9}$, u $\Rightarrow 10^{-6}$, m $\Rightarrow 10^{-3}$, K $\Rightarrow 10^{3}$, M $\Rightarrow 10^{6}$, G $\Rightarrow 10^{9}$, T $\Rightarrow 10^{12}$, P $\Rightarrow 10^{15}$, E $\Rightarrow 10^{18}$, Z $\Rightarrow 10^{21}$, Y $\Rightarrow 10^{24}$.

For exponent 10^{0}, no decimal point or letter is printed.

```
sout | unit(eng( 0.0 )) | unit(eng( 27000.5 )) | unit(eng( −27.5e7 ));
0 27.0005K −275M
```

7. upcase(bin / hex / floating-point) print letters in a value in upper case. Lower case is the default.

```
sout | upcase( bin( 27 ) ) | upcase( hex( 27 ) ) | upcase( 27.5e−10 ) | upcase( hex( 27.5 ) );
0B11011 0X1B 2.75E−09 0X1.B8P+4
```

8. nobase(integer) do not precede bin, oct, hex with 0b/0B, 0, or 0x/0X. Printing the base is the default.

```
sout | nobase( bin( 27 ) ) | nobase( oct( 27 ) ) | nobase( hex( 27 ) );
11011 33 1b
```

9. nodp(floating-point) do not print a decimal point if there are no fractional digits. Printing a decimal point is the default, if there are no fractional digits.

```
sout | 0. | nodp( 0. ) | 27.0 | nodp( 27.0 ) | nodp( 27.5 );
0.0 0 27.0 27 27.5
```

10. sign(integer / floating-point) prefix with plus or minus sign (+ or −). Only printing the minus sign is the default.

```
sout | sign( 27 ) | sign( −27 ) | sign( 27. ) | sign( −27. ) | sign( 27.5 ) | sign( −27.5 );
+27 −27 +27.0 −27.0 +27.5 −27.5
```

11. wd(minimum, value), wd(minimum, precision, value) For all types, minimum is the number of printed characters. If the value is shorter than the minimum, it is padded on the right with spaces.

```
sout | wd( 4, 34) | wd( 3, 34 ) | wd( 2, 34 );
sout | wd( 10, 4.) | wd( 9, 4. ) | wd( 8, 4. );
sout | wd( 4, "ab" ) | wd( 3, "ab" ) | wd( 2, "ab" );
␣␣34␣␣34␣34
␣␣4.000000␣␣4.000000␣4.000000
␣␣ab␣␣ab␣ab
```

If the value is larger, it is printed without truncation, ignoring the minimum.

```
sout | wd( 4, 34567 ) | wd( 3, 34567 ) | wd( 2, 34567 );
sout | wd( 4, 3456. ) | wd( 3, 3456. ) | wd( 2, 3456. );
sout | wd( 4, "abcde" ) | wd( 3, "abcde" ) | wd( 2,"abcde" );
34567␣34567␣34567
3456.␣3456.␣3456.
abcde␣abcde␣abcde
```

For integer types, precision is the minimum number of printed digits. If the value is shorter, it is padded on the left with leading zeros.

```
sout | wd( 4,3, 34 ) | wd( 8,4, 34 ) | wd( 10,10, 34 );
␣034␣␣␣␣␣0034␣0000000034
```

If the value is larger, it is printed without truncation, ignoring the precision.

```
sout | wd( 4,1, 3456 ) | wd( 8,2, 3456 ) | wd( 10,3, 3456 );
3456␣␣␣␣␣3456␣␣␣␣␣␣3456
```

If precision is 0, nothing is printed for zero. If precision is greater than the minimum, it becomes the minimum.

```
sout | wd( 4,0, 0 ) | wd( 3,10, 34 );
␣␣␣␣␣0000000034
```

For floating-point types, precision is the minimum number of digits after the decimal point.

```
sout | wd( 6,3, 27.5 ) | wd( 8,1, 27.5 ) | wd( 8,0, 27.5 ) | wd( 3,8, 27.5 );
27.500␣␣␣␣␣27.5␣␣␣␣␣␣28.␣27.50000000
```

For the C-string type, precision is the maximum number of printed characters, so the string is truncated if it exceeds the maximum.

```
sout | wd( 6,8, "abcd" ) | wd( 6,8, "abcdefghijk" ) | wd( 6,3, "abcd" ) | wd( 10, "" ) | 'X';
␣␣abcd␣abcdefgh␣␣␣␣␣abc␣␣␣␣␣␣␣␣␣␣X
```

Note, printing the null string with minimum width L pads with L spaces.

12. ws(minimum, significant, floating-point) For floating-point types, minimum is the same as for manipulator wd, but significant is the maximum number of significant digits to be printed for both the integer and fractions (versus only the fraction for wd). If a value's significant digits is greater than significant, the last significant digit

is rounded up.

```
sout | ws(6,6, 234.567) | ws(6,5, 234.567) | ws(6,4, 234.567) | ws(6,3, 234.567);
234.567␣234.57 7␣␣234.6␣␣␣␣235
```

If a value's magnitude is greater than significant, the value is printed in scientific notation with the specified number of significant digits.

```
sout | ws(6,6, 234567.) | ws(6,5, 234567.) | ws(6,4, 234567.) | ws(6,3, 234567.);
234567.␣2.3457e+05␣2.346e+05␣2.35e+05
```

If significant is greater than minimum, it defines the number of printed characters.

```
sout | ws(3,6, 234567.) | ws(4,6, 234567.) | ws(5,6, 234567.) | ws(6,6, 234567.);
234567.␣234567.␣234567.␣234567.
```

13. left(field-width) left justify within the given field.

```
sout | left(wd(4, 27)) | left(wd(10, 27.)) | left(wd(10, 27.5)) | left(wd(4,3, 27)) | left(wd(10,3, 27.5));
27␣␣27.000000␣␣27.500000␣␣027␣␣27.500␣␣␣␣
```

14. pad0(field-width) left pad with zeroes (0).

```
sout | pad0( wd( 4, 27 ) ) | pad0( wd( 4,3, 27 ) ) | pad0( wd( 8,3, 27.5 ) );
0027  027 0027.500
```

22.6   Input Manipulators

A string variable *must* be large enough to contain the input sequence. To force programmers to consider buffer overruns for C-string input, C-strings may only be read with a width field, which should specify a size less than or equal to the C-string size, *e.g.*:

```
char line[64];
sin | wdi( sizeof(line), line );              // must specify string size
```

Certain input manipulators support a *scanset*, which is a simple regular expression, where the matching set contains any Latin-1 character (8-bits) or character ranges using minus. For example, the scanset "a-zA-Z -/?§" matches any number of characters between 'a' and 'z', between 'A' and 'Z', between space and '/', and characters '?' and (Latin-1) '§'. The following string is matched by this scanset:

```
!&%$ abAA () ZZZ ??§ xx§§
```

To match a minus, make it the first character in the set, *e.g.*, "-0-9". Other complex forms of regular-expression matching are unsupported.

The following manipulators control scanning of input values (reading) and only affect the format of the argument.

1. skip(*scanset*), skip(*N*) consumes either the *scanset* or the next *N* characters, including newlines. If the match successes, the input characters are ignored, and input continues with the next character. If the match fails, the input characters are left unread.

```
char scanset[ ] = "abc";
sin | "abc␣" | skip( scanset ) | skip( 5 ); // match and skip input sequence
abc␣␣␣abc␣␣xxx
```

Again, the blank in the format string "abc␣" matches any number of whitespace characters.

2. wdi(*maximum*, T & v) For all types except char *, whitespace is skipped and the longest sequence of non-whitespace characters matching an appropriate typed (T) value is read, converted into its corresponding internal form, and written into the T variable. *maximum* is the maximum number of characters read for the current value rather than the longest sequence.

```
char ch;  char ca[3];  int i;  double d;
sin | wdi( sizeof(ch), ch ) | wdi( sizeof(ca), ca[0] ) | wdi( 3, i ) | wdi( 8, d ); // c == 'a', ca == "bcd", i == 123, d == 345.6
abcd1233.456E+2
```

Here, ca[0] is type char, so the width reads 3 characters **without** a null terminator. If an input value is not found for a variable, the exception missing_data is raised, and the input variable is unchanged.

Note, input wdi cannot be overloaded with output wd because both have the same parameters but return different types. Currently, C∀ cannot distinguish between these two manipulators in the middle of an sout/sin expression based on return type.

3. wdi(*maximum size*, **char** s[]) For type **char** *, whitespace is skippped and the longest sequence of non-whitespace characters is read, without conversion, and written into the string variable (null terminated). *maximum size* is the maximum number of characters in the string variable. If the non-whitespace sequence of input characters is greater than *maximum size* − 1 (null termination), the exception cstring_length is raised.

```
char cs[10];
sin | wdi( sizeof(cs), cs );
012345678
```

Nine non-whitespace character are read and the null character is added to make ten.

4. wdi(*maximum size*, *maximum read*, **char** s[]) This manipulator is the same as the previous one, except *maximum read* is the maximum number of characters read for the current value rather than the longest sequence, where *maximum read* ≤ *maximum size*.

```
char cs[10];
sin | wdi( sizeof(cs), 9, cs );
0123456789
```

The exception cstring_length is not raised, because the read stops reading after nine characters.

5. getline(*wdi manipulator*, **const char** delimiter = '\n') consumes the scanset "[^D]D", where D is the delimiter character, which reads all characters from the current input position to the delimiter character into the string (null terminated), and consumes and ignores the delimiter. If the delimiter character is omitted, it defaults to '\n' (newline).

```
char cs[10];
sin | getline( wdi( sizeof(cs), cs ) );
sin | getline( wdi( sizeof(cs), cs ), 'X' ); // X is the line delimiter
abc␣??␣#@%
abc␣??␣#@%X␣w
```

The same value is read for both input strings.

6. quote(**char** & ch, **const char** Ldelimiter = '\'', **const char** Rdelimiter = '\0') consumes the string "LCR", where L is the left delimiter character, C is the value in ch, and R is the right delimiter character, which skips whitespace, consumes and ignores the left delimiter, reads a single character into ch, and consumes and ignores the right delimiter (3 characters). If the delimit character is omitted, it defaults to '\'' (single quote).

```
char ch;
sin | quote( ch );  sin | quote( ch, '"' );  sin | quote( ch, '[', ']' );
␣␣␣'a'␣␣"a"[a]
```

7. quote(*wdi manipulator*, **const char** Ldelimiter = '\'', **const char** Rdelimiter = '\0') consumes the scanset "L[^R]R", where L is the left delimiter character and R is the right delimiter character, which skips whitespace, consumes and ignores the left delimiter, reads characters until the right-delimiter into the string variable (null terminated), and consumes and ignores the right delimiter. If the delimit character is omitted, it defaults to '"' (double quote).

```
char cs[10];
sin | quote( wdi( sizeof(cs), cs ) );              // " is the start/end delimiter
sin | quote( wdi( sizeof(cs), cs ), '\'' );        // ' is the start/end delimiter
sin | quote( wdi( sizeof(cs), cs ), '[', ']' );    // [ is the start and ] is the end delimiter
␣␣␣"abc"␣␣'abc'[abc]
```

8. incl(scanset, *wdi manipulator*) consumes the scanset, which reads all the scanned characters into the string variable (null terminated).

```
char cs[10];
sin | incl( "abc", cs );
bcaxyz
```

9. excl(scanset, *wdi manipulator*) consumes the *not* scanset, which reads all the scanned characters into the string variable (null terminated).

```
char cs[10];
sin | excl( "abc", cs );
xyzbca
```

10. ignore(T & v or **const char** cs[] or *string manipulator*) consumes the appropriate characters for the type and ignores them, so the input variable is unchanged.

```
double d;
char cs[10];
sin | ignore( d );                    // d is unchanged
sin | ignore( cs );                   // cs is unchanged, no wdi required
sin | ignore( quote( wdi( sizeof(cs), cs ) ) ); // cs is unchanged
␣␣−75.35e−4␣25␣"abc"
```

22.7   Concurrent Stream Access

When a stream is shared by multiple threads, input or output characters can be intermixed or cause failure. For example, if two threads execute the following:

thread$_1$: sout | "abc " | "def ";
thread$_2$: sout | "uvw " | "xyz ";

possible outputs are:

abc def	abc uvw xyz	uvw abc xyz def	abuvwc dexf	uvw abc def
uvw xyz	def		yz	xyz

Concurrent operations can even corrupt the internal state of the stream resulting in failure. As a result, some form of mutual exclusion is required for concurrent stream access.

A coarse-grained solution is to perform all stream operations via a single thread or within a monitor providing the necessary mutual exclusion for the stream. A fine-grained solution is to have a lock for each stream, which is acquired and released around stream operations by each thread. C∀ provides a fine-grained solution where a recursive lock is acquired and released indirectly via a manipulator acquire or instantiating an RAII type specific for the kind of stream: osacquire for output streams and isacquire for input streams.

The common usage is the short form of the mutex statement to lock a stream during a single cascaded I/O expression, *e.g.*:

thread$_1$: **mutex**(sout) sout | "abc " | "def ";
thread$_2$: **mutex**(sout) sout | "uvw " | "xyz ";

Now, the order of the thread execution is still non-deterministic, but the output is constrained to two possible lines in either order.

abc def	uvw xyz
uvw xyz	abc def

In summary, the stream lock is acquired by the acquire manipulator and implicitly released at the end of the cascaded I/O expression ensuring all operations in the expression occur atomically.

To lock a stream across multiple I/O operations, he long form of the mutex statement is used, *e.g.*:

```
mutex( sout ) {
    sout | 1;
    mutex( sout ) sout | 2 | 3;        // unnecessary, but ok because of recursive lock
    sout | 4;
} // implicitly release sout lock
```

Note, the unnecessary **mutex** in the middle of the mutex statement, works because the recursive stream-lock can be acquired/released multiple times by the owner thread. Hence, calls to functions that also acquire a stream lock for their output do not result in deadlock.

The previous values written by threads 1 and 2 can be read in concurrently:

```
mutex( sin ) {
    int x, y, z, w;
    sin | x;
    mutex( sin ) sin | y | z;          // unnecessary, but ok because of recursive lock
    sin | w;
} // implicitly release sin lock
```

Again, the order of the reading threads is non-deterministic. Note, non-deterministic reading is rare.

WARNING: The general problem of nested locking can occur if routines are called in an I/O sequence that block, *e.g.*:

```
mutex( sout ) sout | "data:" | rtn( mon );   // mutex call on monitor
```

If the thread executing the I/O expression blocks in the monitor with the sout lock, other threads writing to sout also block until the thread holding the lock is unblocked and releases it. This scenario can lead to deadlock, if the thread that is going to unblock the thread waiting in the monitor first writes to sout (deadly embrace). To prevent nested locking, a simple precaution is to factor out the blocking call from the expression, *e.g.*:

```
int data = rtn( mon );
mutex( sout ) sout | "data:" | data;
```

22.8 Locale

Cultures use different syntax, called a *locale*, for printing numbers so they are easier to read, *e.g.*:

```
12,345.123          // comma separator, period decimal-point
12.345,123          // period separator, comma decimal-point
12␣345,123.         // space separator, comma decimal-point, period terminator
```

A locale is selected with function setlocale, and the corresponding locale package *must* be installed on the underlying system; setlocale returns 0p if the requested locale is unavailable. Furthermore, a locale covers the syntax for many cultural items, *e.g.*, address, measurement, money, etc. This discussion applies to item LC_NUMERIC for formatting non-monetary integral and floating-point values. Figure 10 shows selecting different cultural syntax, which may be associated with one or more countries.

23 String Stream

The stream types ostrstream and istrstream provide all the stream formatting capabilities to/from a C string rather than a stream file. Figure 11, p. 65 shows writing (output) to and reading (input) from a C string. The only string stream operations different from a file stream are:

- constructors to create a stream that writes to a write buffer (ostrstream) of size, or reads from a read buffer (istrstream) containing a C string terminated with '\0'.

    ```
    void ?{}( ostrstream &, char buf[], size_t size );
    void ?{}( istrstream & is, char buf[] );
    ```

- write (ostrstream only) writes all the buffered characters to the specified stream (stdout default).

    ```
    ostrstream & write( ostrstream & os, FILE * stream = stdout );
    ```

    There is no read for istrstream.

24 Structures

Structures in C∀ are basically the same as structures in C. A structure is defined with the same syntax as in C. When referring to a structure in C∀, users may omit the struct keyword.

```
struct Point {
    double x;
    double y;
};

Point p = {0.0, 0.0};
```

C∀ does not support inheritance among types, but instead uses composition to enable reuse of structure fields. Composition is achieved by embedding one type into another. When type A is embedded in type B, an object with type B may be used as an object of type A, and the fields of type A are directly accessible. Embedding types is achieved using anonymous members. For example, using Point from above:

```
void foo(Point p);

struct ColoredPoint {
```

```
#include <fstream.hfa>
#include <locale.h>                     // setlocale
#include <stdlib.h>                     // getenv

int main() {
    void print() {
        sout | 12 | 123 | 1234 | 12345 | 123456 | 1234567;
        sout | 12. | 123.1 | 1234.12 | 12345.123 | 123456.1234 | 1234567.12345;
        sout | nl;
    }
    sout | "Default locale off";
    print();
    sout | "Locale on" | setlocale( LC_NUMERIC, getenv( "LANG" ) );  // enable local locale
    print();
    sout | "German" | setlocale( LC_NUMERIC, "de_DE.UTF-8" );  // enable German locale
    print();
    sout | "Ukraine" | setlocale( LC_NUMERIC, "uk_UA.utf8" );  // enable Ukraine locale
    print();
    sout | "Default locale off" | setlocale( LC_NUMERIC, "C" );  // disable locale
    print();
}
```

```
Default locale off
12 123 1234 12345 123456 1234567
12. 123.1 1234.12 12345.123 123456.1234 1234567.12345

Locale on en_US.UTF−8
12 123 1,234 12,345 123,456 1,234,567
12. 123.1 1,234.12 12,345.123 123,456.1234 1,234,567.12345

German de_DE.UTF−8
12 123 1.234 12.345 123.456 1.234.567
12. 123,1. 1.234,12 12.345,123 123.456,1234 1.234.567,12345

Ukraine uk_UA.utf8
12 123 1 234 12 345 123 456 1 234 567
12. 123,1. 1 234,12. 12 345,123. 123 456,1234. 1 234 567,12345.

Default locale off C
12 123 1234 12345 123456 1234567
12. 123.1 1234.12 12345.123 123456.1234 1234567.12345
```

Figure 10: Stream Locale

```
1        Point; // anonymous member (no identifier)
2        int Color;
3    };
4    ...
5        ColoredPoint cp = ...;
6        cp.x = 10.3; // x from Point is accessed directly
7        cp.color = 0x33aaff; // color is accessed normally
8        foo(cp); // cp can be used directly as a Point
```

9 25 Constructors and Destructors

10 C∀ supports C initialization of structures, but it also adds constructors for more advanced initialization. Additionally,
11 C∀ adds destructors that are called when a variable is deallocated (variable goes out of scope or object is deleted).
12 These functions take a reference to the structure as a parameter (see Section for more information).

```
#include <fstream.hfa>
#include <strstream.hfa>

int main() {
    enum { size = 256 };
    char buf[size];                        // output buffer
    ostrstream osstr = { buf, size };      // bind output buffer/size
    int i = 3, j = 5, k = 7;
    double x = 12345678.9, y = 98765.4321e−11;

    osstr | i | hex(j) | wd(10, k) | sci(x) | unit(eng(y)) | "abc";
    write( osstr );                        // write string to stdout
    printf( "%s", buf );                   // same lines of output
    sout | i | hex(j) | wd(10, k) | sci(x) | unit(eng(y)) | "abc";

    char buf2[] = "12 14 15 3.5 7e4 abc"; // input buffer
    istrstream isstr = { buf2 };
    char s[10];
    isstr | i | j | k | x | y | s;
    sout  | i | j | k | x | y | s;
}

3␣0x5␣␣␣␣␣␣␣␣␣␣␣7␣1.234568e+07␣987.654n␣abc
3␣0x5␣␣␣␣␣␣␣␣␣␣␣7␣1.234568e+07␣987.654n␣abc
3␣0x5␣␣␣␣␣␣␣␣␣␣␣7␣1.234568e+07␣987.654n␣abc
12␣14␣15␣3.5␣70000.␣abc
```

Figure 11: String Stream Processing

26   Overloading

Overloading refers to the capability of a programmer to define and use multiple objects in a program with the same name. In C∀, a declaration may overload declarations from outer scopes with the same name, instead of hiding them as is the case in C. This may cause identical C and C∀ programs to behave differently. The compiler selects the appropriate object (overload resolution) based on context information at the place where it is used. Overloading allows programmers to give functions with different signatures but similar semantics the same name, simplifying the interface to users. Disadvantages of overloading are that it can be used to give functions with different semantics the same name, causing confusion, or that the compiler may resolve to a different function from what the programmer expected. C∀ allows overloading of functions, operators, variables, and even the constants 0 and 1.

The compiler follows some overload resolution rules to determine the best interpretation of all of these overloads. 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. For details about type inference and overload resolution, please see the C∀ Language Specification.

```
int foo(int a, int b) {
    float sum = 0.0;
    float special = 1.0;
    {
        int sum = 0;
        // both the float and int versions of sum are available
        float special = 4.0;
        // this inner special hides the outer version
        ...
    }
    ...
```

```
struct Widget {
    int id;
    float size;
    int * optionalint;
};

// ?{} is the constructor operator identifier
// The first argument is a reference to the type to initialize
// Subsequent arguments can be specified for initialization

void ?{}( Widget & w ) {                    // default constructor
    w.id = −1;
    w.size = 0.0;
    w.optionalint = 0p;
}

// constructor with values (does not need to include all fields)
void ?{}( Widget & w, int id, float size ) {
    w.id = id;
    w.size = size;
    w.optionalint = 0p;
}

// ^ ?{} is the destructor operator identifier
void ^?{}( Widget & w ) {                    // destructor
    w.id = 0;
    w.size = 0.0;
    if ( w.optionalint != 0p ) {
        free( w.optionalint );
        w.optionalint = 0p;
    }
}

Widget baz;                          // reserve space only
Widget foo{};                        // calls default constructor
Widget bar{ 23, 2.45 };              // calls constructor with values
baz{ 24, 0.91 };                     // calls constructor with values
?{}( baz, 24, 0.91 );                // explicit call to constructor
^?{} (bar );                         // explicit call to destructor
```

Figure 12: Constructors and Destructors

1   }

2 26.1   Constant

3 The constants 0 and 1 have special meaning. In C∀, as in C, all scalar types can be incremented and decremented,
4 which is defined in terms of adding or subtracting 1. The operations &&, ||, and ! can be applied to any scalar arguments
5 and are defined in terms of comparison against 0 (*e.g.*, (a && b) becomes (a != 0 && b != 0)).

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

12   Why just 0 and 1? Why not other integers? No other integers have special status in C. A facility that let program-
13 mers 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 does not seem worth the gain.

For example, to define the constants for a complex type, the programmer would define the following:

```
struct Complex {
    double real;
    double imaginary;
}

const Complex 0 = {0, 0};
const Complex 1 = {1, 0};
...

    Complex a = 0;
...

    a++;
...
    if (a) { // same as if (a == 0)
...
    }
```

26.2   Variable

The overload rules of C∀ allow a programmer to define multiple variables with the same name, but different types. Allowing overloading of variable names enables programmers to use the same name across multiple types, simplifying naming conventions and is compatible with the other overloading that is allowed. For example, a developer may want to do the following:

```
int pi = 3;
float pi = 3.14;
char pi = .p.;
```

26.3   Function Overloading

Overloaded functions in C∀ are resolved based on the number and type of arguments, type of return value, and the level of specialization required (specialized functions are preferred over generic).

The examples below give some basic intuition about how the resolution works.

```
// Choose the one with less conversions
int doSomething(int value) {...} // option 1
int doSomething(short value) {...} // option 2

int a, b = 4;
short c = 2;

a = doSomething(b); // chooses option 1
a = doSomething(c); // chooses option 2

// Choose the specialized version over the generic

generic(type T)
T bar(T rhs, T lhs) {...} // option 3
float bar(float rhs, float lhs){...} // option 4
float a, b, c;
double d, e, f;
c = bar(a, b); // chooses option 4

// specialization is preferred over unsafe conversions

f = bar(d, e); // chooses option 5
```

26.4   Operator

C∀ also allows operators to be overloaded, to simplify the use of user-defined types. Overloading the operators allows the users to use the same syntax for their custom types that they use for built-in types, increasing readability and improving productivity. C∀ uses the following special identifiers to name overloaded operators:

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

Table 2: Operator Identifiers

These identifiers are defined such that the question marks in the name identify the location of the operands. These operands represent the parameters to the functions, and define how the operands are mapped to the function call. For example, a + b becomes ?+?(a, b).

In the example below, a new type, myComplex, is defined with an overloaded constructor, + operator, and string operator. These operators are called using the normal C syntax.

```
type Complex = struct { // define a Complex type
    double real;
    double imag;
}

// Constructor with default values

void ?{}(Complex &c, double real = 0.0, double imag = 0.0) {
    c.real = real;
    c.imag = imag;
}

Complex ?+?(Complex lhs, Complex rhs) {
    Complex sum;
    sum.real = lhs.real + rhs.real;
    sum.imag = lhs.imag + rhs.imag;
    return sum;
}

String ()?(const Complex c) {
    // use the string conversions for the structure members
    return (String)c.real + . + . + (String)c.imag + .i.;
}
...

Complex a, b, c = {1.0}; // constructor for c w/ default imag
...
c = a + b;
```

```
print(.sum = . + c);
```

27   Auto Type-Inferencing

Auto type-inferencing occurs in a declaration where a variable's type is inferred from its initialization expression type.

C++	gcc	
	#define expr 3.0 * i	
auto j = 3.0 * 4;	**typeof**(expr) j = expr;	*// use type of initialization expression*
int i;	**int** i;	
auto k = i;	**typeof**(i) k = i;	*// use type of primary variable*

The two important capabilities are:

- not determining or writing long generic types,

- ensuring secondary variables, related to a primary variable, always have the same type.

In C∀, **typedef** provides a mechanism to alias long type names with short ones, both globally and locally, but not eliminate the use of the short name. gcc provides **typeof** to declare a secondary variable from a primary variable. C∀ also relies heavily on the specification of the left-hand side of assignment for type inferencing, so in many cases it is crucial to specify the type of the left-hand side to select the correct type of the right-hand expression. Only for overloaded routines *with the same return type* is variable type-inferencing possible. Finally, **auto** presents the programming problem of tracking down a type when the type is actually needed. For example, given

```
auto j = ...
```

and the need to write a routine to compute using j

```
void rtn( ... parm );
rtn( j );
```

A programmer must work backwards to determine the type of j's initialization expression, reconstructing the possibly long generic type-name. In this situation, having the type name or a short alias is very useful.

There is also the conundrum in type inferencing of when to *brand* a type. That is, when is the type of the variable more important than the type of its initialization expression. For example, if a change is made in an initialization expression, it can cause cascading type changes and/or errors. At some point, a variable type needs to remain constant and the expression to be in error when it changes.

Given **typedef** and **typeof** in C∀, and the strong need to use the type of left-hand side in inferencing, auto type-inferencing is not supported at this time. Should a significant need arise, this feature can be revisited.

28   Concurrency

Concurrency support in C∀ is implemented on top of a highly efficient runtime system of light-weight, M:N, user level threads. The model integrates concurrency features into the language by making the structure type the core unit of concurrency. All communication occurs through method calls, where data is sent via method arguments, and received via the return value. This enables a very familiar interface to all programmers, even those with no parallel programming experience. It also allows the compiler to do static type checking of all communication, a very important safety feature. This controlled communication with type safety has some similarities with channels in Go, and can actually implement channels exactly, as well as create additional communication patterns that channels cannot. Mutex objects, monitors, are used to contain mutual exclusion within an object and synchronization across concurrent threads.

28.1   Coroutine

Coroutines are the precursor to threads. Figure 13 shows a coroutine that computes the Fibonacci numbers.

28.2   Monitors

A monitor is a structure in C∀ which includes implicit locking of its fields. Users of a monitor interact with it just like any structure, but the compiler handles code as needed to ensure mutual exclusion. An example of the definition of a monitor is shown here:

```
#include <fstream.hfa>
#include <coroutine.hfa>

coroutine Fibonacci {
    int fn;                              // used for communication
};

void main( Fibonacci & fib ) with( fib ) { // called on first resume
    int fn1, fn2;                        // retained between resumes
    fn = 0;  fn1 = fn;                   // 1st case
    suspend;                             // restart last resume
    fn = 1;  fn2 = fn1;  fn1 = fn;       // 2nd case
    suspend;                             // restart last resume
    for () {
        fn = fn1 + fn2;  fn2 = fn1;  fn1 = fn; // general case
        suspend;                         // restart last resume
    }
}
int next( Fibonacci & fib ) with( fib ) {
    resume( fib );                       // restart last suspend
    return fn;
}
int main() {
    Fibonacci f1, f2;
    for ( 10 ) {                         // print N Fibonacci values
        sout | next( f1 ) | next( f2 );
    }
}
```

Figure 13: Fibonacci Coroutine

```
1    type Account = monitor {
2        const unsigned long number; // account number
3        float balance; // account balance
4    };
```

28.3   Threads

C∀ also provides a simple mechanism for creating and utilizing user level threads. A thread provides mutual exclusion like a monitor, and also has its own execution state and a thread of control. Similar to a monitor, a thread is defined like a structure:

29   Language Comparisons

C∀ is one of many languages that attempts to improve upon C. In developing C∀, many other languages were consulted for ideas, constructs, and syntax. Therefore, it is important to show how these languages each compare with Do. In this section, C∀ is compared with what the writers of this document consider to be the closest competitors of Do: C++, Go, Rust, and D.

29.1   C++

C++ is a general-purpose programming language. It is an imperative, object-oriented and generic programming language, while also providing facilities for low-level memory manipulation. The primary focus of C++ was adding object-oriented programming to C, and this is the primary difference between C++ and C∀. C++ uses classes to encapsulate data and the functions that operate on that data, and to hide the internal representation of the data. C∀ uses modules instead to perform these same tasks. Classes in C++ also enable inheritance among types. Instead of inheritance, C∀ embraces composition and interfaces to achieve the same goals with more flexibility. There are many studies

```
#include <fstream.hfa>
#include <thread.hfa>

monitor AtomicCnt { int counter; };
void ?{}( AtomicCnt & c, int init = 0 ) with(c) { counter = init; }
int inc( AtomicCnt & mutex c, int inc = 1 ) with(c) { return counter += inc; }
int dec( AtomicCnt & mutex c, int dec = 1 ) with(c) { return counter −= dec; }
forall( ostype & | ostream( ostype ) ) {    // print any stream
    ostype & ?|?( ostype & os, AtomicCnt c ) { return os | c.counter; }
    void ?|?( ostype & os, AtomicCnt c ) { (ostype &)(os | c.counter); ends( os ); }
}

AtomicCnt global;                          // shared

thread MyThread {};
void main( MyThread & ) {
    for ( i; 100_000 ) {
        inc( global );
        dec( global );
    }
}
int main() {
    enum { Threads = 4 };
    processor p[Threads − 1];               // + starting processor
    {
        MyThread t[Threads];
    }
    sout | global;                          // print 0
}
```

Figure 14: Atomic-Counter Monitor

1 and articles comparing inheritance and composition (or is-a versus has-a relationships), so we will not go into more
2 detail here (Venners, 1998) (Pike, Go at Google: Language Design in the Service of Software Engineering , 2012).
3   Overloading in C∀ is very similar to overloading in C++, with the exception of the additional use, in C∀, of the
4 return type to differentiate between overloaded functions. References and exceptions in C∀ are heavily based on the
5 same features from C++. The mechanism for interoperating with C code in C∀ is also borrowed from C++.
6   Both C∀ and C++ provide generics, and the syntax is quite similar. The key difference between the two, is that in
7 C++ templates are expanded at compile time for each type for which the template is instantiated, while in C∀, function
8 pointers are used to make the generic fully compilable. This means that a generic function can be defined in a compiled
9 library, and still be used as expected from source.

10   29.2   Go

11 Go, also commonly referred to as golang, is a programming language developed at Google in 2007 [19]. It is a
12 statically typed language with syntax loosely derived from that of C, adding garbage collection, type safety, some
13 structural typing capabilities, additional built-in types such as variable-length arrays and key-value maps, and a large
14 standard library. (Wikipedia)
15   Go and C∀ differ significantly in syntax and implementation, but the underlying core concepts of the two languages
16 are aligned. Both Go and C∀ use composition and interfaces as opposed to inheritance to enable encapsulation and
17 abstraction. Both languages (along with their tooling ecosystem) provide a simple packaging mechanism for building
18 units of code for easy sharing and reuse. Both languages also include built-in light weight, user level threading
19 concurrency features that attempt to simplify the effort and thought process required for writing parallel programs
20 while maintaining high performance.
21   Go has a significant runtime which handles the scheduling of its light weight threads, and performs garbage col-
22 lection, among other tasks. C∀ uses a cooperative scheduling algorithm for its tasks, and uses automatic reference

```
#include <fstream.hfa>
#include <kernel>
#include <stdlib>
#include <thread>

thread First  { signal_once * lock; };
thread Second { signal_once * lock; };

void ?{}( First * this, signal_once* lock ) { this->lock = lock; }
void ?{}( Second * this, signal_once* lock ) { this->lock = lock; }

void main( First * this ) {
    for ( int i = 0; i < 10; i += 1 ) {
        sout | "First : Suspend No." | i + 1;
        yield();
    }
    signal( this->lock );
}

void main( Second * this ) {
    wait( this->lock );
    for ( int i = 0; i < 10; i += 1 ) {
        sout | "Second : Suspend No." | i + 1;
        yield();
    }
}

int main( void ) {
    signal_once lock;
    sout | "User main begin";
    {
        processor p;
        {
            First  f = { &lock };
            Second s = { &lock };
        }
    }
    sout | "User main end";
}
```

Figure 15: Simple Threads

1 counting to enable advanced memory management without garbage collection. This results in Go requiring significant
2 overhead to interface with C libraries while C∀ has no overhead.

3 29.3   Rust

4 Rust is a general-purpose, multi-paradigm, compiled programming language developed by Mozilla Research. It is
5 designed to be a "safe, concurrent, practical language", supporting pure-functional, concurrent-actor[dubious . dis-
6 cuss][citation needed], imperative-procedural, and object-oriented styles.

7   The primary focus of Rust is in safety, especially in concurrent programs. To enforce a high level of safety, Rust
8 has added ownership as a core feature of the language to guarantee memory safety. This safety comes at the cost of a
9 difficult learning curve, a change in the thought model of the program, and often some runtime overhead.

10   Aside from those key differences, Rust and C∀ also have several similarities. Both languages support no overhead
11 interoperability with C and have minimal runtimes. Both languages support inheritance and polymorphism through
12 the use of interfaces (traits).

29.4 D

The D programming language is an object-oriented, imperative, multi-paradigm system programming language created by Walter Bright of Digital Mars and released in 2001. [.] Though it originated as a re-engineering of C++, D is a distinct language, having redesigned some core C++ features while also taking inspiration from other languages, notably Java, Python, Ruby, C#, and Eiffel.

D and C∀ both start with C and add productivity features. The obvious difference is that D uses classes and inheritance while C∀ uses composition and interfaces. D is closer to C∀ than C++ since it is limited to single inheritance and also supports interfaces. Like C++, and unlike C∀, D uses garbage collection and has compile-time expanded templates. D does not have any built-in concurrency constructs in the language, though it does have a standard library for concurrency which includes the low-level primitives for concurrency.

A Syntax Ambiguities

C has a number of syntax ambiguities, which are resolved by taking the longest sequence of overlapping characters that constitute a token. For example, the program fragment x+++++y is parsed as x␣++␣++␣+␣y because operator tokens ++ and + overlap. Unfortunately, the longest sequence violates a constraint on increment operators, even though the parse x␣++␣+␣++␣y might yield a correct expression. Hence, C programmers are aware that spaces have to added to disambiguate certain syntactic cases.

In C∀, there are ambiguous cases with dereference and operator identifiers, *e.g.*, **int** *?*?(), where the string *?*? can be interpreted as:

```
*?␣*?                              // dereference operator, dereference operator
*␣?*?                              // dereference, multiplication operator
```

By default, the first interpretation is selected, which does not yield a meaningful parse. Therefore, C∀ does a lexical look-ahead for the second case, and backtracks to return the leading unary operator and reparses the trailing operator identifier. Otherwise a space is needed between the unary operator and operator identifier to disambiguate this common case.

A similar issue occurs with the dereference, *?(...), and routine-call, ?()(...) identifiers. The ambiguity occurs when the deference operator has no parameters:

```
*?()␣... ;
*?()␣...(...) ;
```

requiring arbitrary whitespace look-ahead for the routine-call parameter-list to disambiguate. However, the dereference operator *must* have a parameter/argument to dereference *?(...). Hence, always interpreting the string *?() as *␣?() does not preclude any meaningful program.

The remaining cases are with the increment/decrement operators and conditional expression, *e.g.*:

```
i++?␣...(...);
i?++␣...(...);
```

requiring arbitrary whitespace look-ahead for the operator parameter-list, even though that interpretation is an incorrect expression (juxtaposed identifiers). Therefore, it is necessary to disambiguate these cases with a space:

```
i++␣? i : 0;
i?␣++i : 0;
```

B C Incompatibles

The following incompatibles exist between C∀ and C, and are similar to Annex C for C++ [22].

1. **Change:** add new keywords
   New keywords are added to C∀ (see Section C, p. 76).

   Rationale: keywords added to implement new semantics of C∀.

   Effect on original feature: change to semantics of well-defined feature.
   Any C11 programs using these keywords as identifiers are invalid C∀ programs.

   Difficulty of converting: keyword clashes are accommodated by syntactic transformations using the C∀ back-quote escape-mechanism (see Section 6, p. 5).

   How widely used: clashes among new C∀ keywords and existing identifiers are rare.

2. **Change:** drop K&R C declarations

   K&R declarations allow an implicit base-type of **int**, if no type is specified, plus an alternate syntax for declaring parameters. *e.g.*:

   ```
   x;                        // int x
   *y;                       // int * y
   f( p1, p2 );              // int f( int p1, int p2 );
   g( p1, p2 ) int p1, p2;   // int g( int p1, int p2 );
   ```

   C∀ continues to support K&R routine definitions:

   ```
   f( a, b, c )              // default int return
       int a, b; char c;     // K&R parameter declarations
   {
       ...
   }
   ```

   Rationale: dropped from C11 standard.[17]

   Effect on original feature: original feature is deprecated.

   Any old C programs using these K&R declarations are invalid C∀ programs.

   Difficulty of converting: trivial to convert to C∀.

   How widely used: existing usages are rare.

3. **Change:** type of character literal **int** to **char** to allow more intuitive overloading:

   ```
   int rtn( int i );
   int rtn( char c );
   rtn( 'x' );                       // programmer expects 2nd rtn to be called
   ```

   Rationale: it is more intuitive for the call to rtn to match the second version of definition of rtn rather than the first. In particular, output of **char** variable now print a character rather than the decimal ASCII value of the character.

   ```
   sout | 'x' | "  " | (int)'x';
   x 120
   ```

   Having to cast 'x' to **char** is non-intuitive.

   Effect on original feature: change to semantics of well-defined feature that depend on:

   ```
   sizeof( 'x' ) == sizeof( int )
   ```

   no long work the same in C∀ programs.

   Difficulty of converting: simple

   How widely used: programs that depend upon **sizeof**('x') are rare and can be changed to **sizeof**(**char**).

4. **Change:** make string literals **const**:

   ```
   char * p = "abc";          // valid in C, deprecated in C∀
   char * q = expr ? "abc" : "de";   // valid in C, invalid in C∀
   ```

   The type of a string literal is changed from [] **char** to **const** [] **char**. Similarly, the type of a wide string literal is changed from [] wchar_t to **const** [] wchar_t.

   Rationale: This change is a safety issue:

   ```
   char * p = "abc";
   p[0] = 'w';                        // segment fault or change constant literal
   ```

   The same problem occurs when passing a string literal to a routine that changes its argument.

   Effect on original feature: change to semantics of well-defined feature.

   Difficulty of converting: simple syntactic transformation, because string literals can be converted to **char** *.

   How widely used: programs that have a legitimate reason to treat string literals as pointers to potentially modifiable memory are rare.

5. **Change:** remove *tentative definitions*, which only occurs at file scope:

---

[17]At least one type specifier shall be given in the declaration specifiers in each declaration, and in the specifier-qualifier list in each structure declaration and type name [21, § 6.7.2(2)]

```
1      int i;                          // forward definition
2      int *j = &i;                    // forward reference, valid in C, invalid in C∀
3      int i = 0;                      // definition
```

is valid in C, and invalid in C∀ because duplicate overloaded object definitions at the same scope level are disallowed. This change makes it impossible to define mutually referential file-local static objects, if initializers are restricted to the syntactic forms of C. For example,

```
7      struct X { int i; struct X *next; };
8      static struct X a;              // forward definition
9      static struct X b = { 0, &a }; // forward reference, valid in C, invalid in C∀
10     static struct X a = { 1, &b };    // definition
```

Rationale: avoids having different initialization rules for builtin types and user-defined types.

Effect on original feature: change to semantics of well-defined feature.

Difficulty of converting: the initializer for one of a set of mutually-referential file-local static objects must invoke a routine call to achieve the initialization.

How widely used: seldom

6. **Change:** have **struct** introduce a scope for nested types:

```
17     enum Colour { R, G, B, Y, C, M };
18     struct Person {
19        enum Colour { R, G, B }; // nested type
20        struct Face {            // nested type
21           Colour Eyes, Hair;    // type defined outside (1 level)
22        };
23        .Colour shirt;           // type defined outside (top level)
24        Colour pants;            // type defined same level
25        Face looks[10];           // type defined same level
26     };
27     Colour c = R;               // type/enum defined same level
28     Person.Colour pc = Person.R;  // type/enum defined inside
29     Person.Face pretty;         // type defined inside
```

In C, the name of the nested types belongs to the same scope as the name of the outermost enclosing structure, *i.e.*, the nested types are hoisted to the scope of the outer-most type, which is not useful and confusing. C∀ is C *incompatible* on this issue, and provides semantics similar to C++. Nested types are not hoisted and can be referenced using the field selection operator ".", unlike the C++ scope-resolution operator "::".

Rationale: **struct** scope is crucial to C∀ as an information structuring and hiding mechanism.

Effect on original feature: change to semantics of well-defined feature.

Difficulty of converting: Semantic transformation.

How widely used: C programs rarely have nest types because they are equivalent to the hoisted version.

7. **Change:** In C++, the name of a nested class is local to its enclosing class.

Rationale: C++ classes have member functions which require that classes establish scopes.

Difficulty of converting: Semantic transformation. To make the struct type name visible in the scope of the enclosing struct, the struct tag could be declared in the scope of the enclosing struct, before the enclosing struct is defined. Example:

```
43     struct Y;                        // struct Y and struct X are at the same scope
44     struct X {
45        struct Y { /* ... */ } y;
46     };
```

All the definitions of C struct types enclosed in other struct definitions and accessed outside the scope of the enclosing struct could be exported to the scope of the enclosing struct. Note: this is a consequence of the difference in scope rules, which is documented in 3.3.

How widely used: Seldom.

8. **Change:** remove implicit conversion of **void** * to or from any T * pointer:

```
52     void foo() {
```

```
int * b = malloc( sizeof(int) ); // implicitly convert void * to int *
char * c = b;                    // implicitly convert int * to void *, and then void * to char *
}
```

Rationale: increase type safety

Effect on original feature: deletion of semantically well-defined feature.

Difficulty of converting: requires adding a cast (see Section E.1 for better alternatives):

```
int * b = (int *)malloc( sizeof(int) );
char * c = (char *)b;
```

How widely used: Significant. Some C translators already give a warning if the cast is not used.

9. **Change:** Types must be declared in declarations, not in expressions In C, a sizeof expression or cast expression may create a new type. For example,

```
p = (void*)(struct x {int i;} *)0;
```

declares a new type, struct x .

Rationale: This prohibition helps to clarify the location of declarations in the source code.

Effect on original feature: Deletion of a semantically welldefined feature.

Difficulty of converting: Syntactic transformation.

How widely used: Seldom.

10. **Change:** comma expression is disallowed as subscript

Rationale: safety issue to prevent subscripting error for multidimensional arrays: x[i,j] instead of x[i][j], and this syntactic form then taken by C∀ for new style arrays.

Effect on original feature: change to semantics of well-defined feature.

Difficulty of converting: semantic transformation of x[i,j] to x[(i,j)]

How widely used: Seldom.

C   C∀ **Keywords**

C∀ introduces the following new keywords, which cannot be used as identifiers.

basetypeof, **choose**, **coroutine**, **disable**, **enable**, **exception**, **fallthrough**, **fallthrough**, **finally**, **fixup**, **forall,generator**,
int128, **monitor**, **mutex**, **one_t**, **report**, **suspend**, **throw**, **throwResume**, **trait**, **try**, **virtual**, **waitfor**, **when**, **with**, **zero_t**

C∀ introduces the following new quasi-keywords, which can be used as identifiers.

catch, **catchResume**, **finally**, **fixup**, **or**, **timeout**

D   **Standard Headers**

C11 prescribes the following standard header-files [21, § 7.1.2]:

assert.h, complex.h, ctype.h, errno.h, fenv.h, float.h, inttypes.h, iso646.h, limits.h, locale.h, math.h, setjmp.h, signal.h, stdalign.h, stdarg.h, stdatomic.h, stdbool.h, stddef.h, stdint.h, stdio.h, stdlib.h, stdnoreturn.h, string.h, tgmath.h, threads.h, time.h, uchar.h, wchar.h, wctype.h

and C∀ adds to this list:

gmp.h, malloc.h, unistd.h

For the prescribed head-files, C∀ uses header interposition to wraps these includes in an **extern** "C"; hence, names in these include files are not mangled (see Section 4, p. 2). All other C header files must be explicitly wrapped in **extern** "C" to prevent name mangling. This approach is different from C++ where the name-mangling issue is handled internally in C header-files through checks for preprocessor variable __cplusplus, which adds appropriate **extern** "C" qualifiers.

E   **Standard Library**

The C∀ standard-library extends existing C library routines by adding new function, wrapping existing explicitly-polymorphic C routines into implicitly-polymorphic versions, and adding new C∀ extensions.

Table 3: Allocation Routines versus Storage-Management Properties

	routine	**fill**	**alignment**	**scale**	**resize**
C	malloc	no	no	no	no
	calloc	yes (0 only)	no	yes	no
	realloc	copy	no	no	yes
	reallocarray	copy	no	yes	yes
	memalign	no	yes	no	no
	aligned_alloc[a]	no	yes	no	no
	posix_memalign	no	yes	no	no
	valloc	no	yes (page size)	no	no
	pvalloc[b]	no	yes (page size)	no	no
C∀	cmemalign	yes (0 only)	yes	yes	no
	resize	no copy	yes	no	yes
	realloc	copy	yes	no	yes
	alloc[c]	yes	yes	yes	yes

[a]Same as memalign but size is an integral multiple of alignment.
[b]Same as valloc but rounds size to multiple of page size.
[c]Multiple overloads with different parameters.

E.1 Dynamic Storage-Management

Dynamic storage-management in C is based on explicit allocation and deallocation (malloc/free). Programmer's must manage all allocated storage via its address (pointer) and subsequently deallocate the storage via this address. Storage that is not deallocated becomes inaccessible, called a *memory leak*, which can only be detected at program termination. Storage freed twice is an error, called a *duplicate free*, which can sometimes be detected. Storage used after it is deallocated is an error, called using a *dangling pointer*, which can sometimes be detected.

E.1.1 C Interface

C dynamic storage-management provides the following properties.

fill storage after an allocation with a specified character or value.

align an allocation on a specified memory boundary, *e.g.*, an address multiple of 64 or 128 for cache-line purposes.

scale an allocation size to the specified number of array elements. An array may be filled, resized, or aligned.

resize an existing allocation to decreased or increased its size. In either direction, new storage may or may not be allocated, but if there is a new allocation, as much data from the existing allocation is copied into the new allocation. When new storage is allocated, it may be aligned and storage after copied data may be filled.

Table 3 shows different combinations of storage-management properties provided by the C and C∀ allocation routines.

E.1.2 C∀ Interface

C∀ dynamic memory management:

1. extends type safety of all allocation routines by using the left-hand assignment type to determine the allocation size and alignment, and return a matching type for the new storage, which removes many common allocation errors.

```
int * ip = (int *)malloc( sizeof(int) ); // C
int * ip = malloc();              // C∀ type-safe call of C malloc
int * ip = calloc();              // C∀ type-safe call of C calloc
struct __attribute__(( aligned(128) )) spinlock { ... }; // cache alignment
spinlock * slp = malloc();        // correct size, alignment, and return type
```

Here, the alignment of the ip storage is 16 (default) and 128 for slp.

2. introduces the notion of *sticky properties* used in resizing. All initial allocation properties are remembered and maintained for use should resize require new storage. For example, the initial alignment and fill properties in the initial allocation

```
struct __attribute__(( aligned(4096) )) S { ... };
S * sp = calloc( 10 );            // align 4K and zero fill
```

```
1        sp = reallocarray( sp, 100 );          // preserve 4K alignment and zero fill new storage
```

2   are preserved in the resize so the new storage has the same alignment and extra storage after the data copy is zero

3   filled. Without sticky properties it is dangerous to resize, resulting in the C idiom of manually performing the

4   reallocation to maintain correctness, which is error prone.

5   3. provides resizing without data copying, which is useful to repurpose an existing block of storage, rather than

6   freeing the old storage and performing a new allocation. A resize can take advantage of unused storage after the

7   data to preventing a free/reallocation step altogether.

8   4. provides free/delete functions that delete a variable number of allocations.

```
9        int * ip = malloc(), * jp = malloc(), * kp = malloc();
10       double * xp = malloc(), * yp = malloc(), * zp = malloc();
11       free( ip, jp, kp, xp, yp, zp );        // multiple deallocations
```

12 5. supports constructors for initialization of allocated storage and destructors for deallocation (like C++).

```
13       struct S { int v; };                   // default constructors
14       void ^?{}( S & ) { ... }               // destructor
15       S & sp = *new( 3 );                    // allocate and call constructor
16       sout | sp.v;
17       delete( &sp );                         // call destructor
18       S * spa1 = anew( 10, 5 ), * spa2 = anew( 10, 8 ); // allocate array and call constructor for each array element
19       for ( i; 10 ) sout | spa1[i].v | spa2[i].v | nonl; sout | nl;
20       adelete( spa1, spa2 );                 // call destructors on all array objects
21
22       3
23       5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8
```

24 Allocation routines new/anew allocate a variable/array and initialize storage using the allocated type's constructor.

25 Note, the matching deallocation routines delete/adelete. C++ only supports the default constructor for intializing

26 array elements.

```
27       S * sp = new S[10]{5};                 // disallowed
```

28 In addition, C∀ provides a new allocator interface to further increase orthogonality and usability of dynamic-

29 memory allocation. This interface helps programmers in three ways.

30 1. naming: C∀ regular and **ttype** polymorphism (similar to C++ variadic templates) is used to encapsulate a wide

31 range of allocation functionality into a single routine name, so programmers do not have to remember multiple

32 routine names for different kinds of dynamic allocations.

33 2. named arguments: individual allocation properties are specified using postfix function call (see Section 15.3,

34 p. 35), so programmers do not have to remember parameter positions in allocation calls.

35 3. safe usage: like the C∀'s C-interface, programmers do not have to specify object size or cast allocation results.

36 The polymorphic functions

```
37       T * alloc( ... );
38       T * alloc( size_t dim, ... );
```

39 are overloaded with a variable number of allocation properties. These allocation properties can be passed as named

40 arguments when calling the alloc routine. A call without parameters returns an uninitialized dynamically allocated

41 object of type T (malloc). A call with only the dimension (dim) parameter returns an uninitialized dynamically allocated

42 array of objects with type T (aalloc). The variable number of arguments consist of allocation properties to specialize

43 the allocation. The properties resize and realloc are associated with an existing allocation variable indicating how its

44 storage is modified.

45 The following allocation property functions may be combined and appear in any order as arguments to alloc,

46 • T_align ?`align(size_t alignment) to align an allocation. The alignment parameter must be ≥ the default align-

47 ment (libAlign() in C∀) and a power of two, *e.g.*, the following return a dynamic object and object array aligned

48 on a 256 and 4096-byte boundary.

```
49       int * i0 = alloc( 256`align );  sout | i0 | nl;
50       int * i1 = alloc( 3, 4096`align );  for (i; 3 ) sout | &i1[i] | nonl;  sout | nl;
51       free( i0, i1 );
52
```

```
      0x555555656900  // 256 alignment
      0x55555656c000 0x5656c004 0x5656c008  // 4K array alignment
```

- T_fill(T) ?`fill(/* *various types* */) to initialize storage. There are three ways to fill storage:

     1. A **char** fills every byte of each object.

     2. An object of the returned type fills each object.

     3. An object array pointer fills some or all of the corresponding object array.

  For example:

```
1   int * i0 = alloc( 0n`fill );  sout | *i0 | nl;  // 0n disambiguates 0p
2   int * i1 = alloc( 5`fill );  sout | *i1 | nl;
3   int * i2 = alloc( '\xfe'`fill );  sout | hex( *i2 ) | nl;
4   int * i3 = alloc( 5, 5`fill );  for ( i; 5 ) sout | i3[i] | nonl; sout | nl;
5   int * i4 = alloc( 5, 0xdeadbeefN`fill );  for ( i; 5 ) sout | hex( i4[i] ) | nonl; sout | nl;
6   int * i5 = alloc( 5, i3`fill );  for ( i; 5 ) sout | i5[i] | nonl; sout | nl; // completely fill from i3
7   int * i6 = alloc( 5, [i3, 3]`fill );  for ( i; 5 ) sout | i6[i] | nonl; sout | nl; // partial fill from i3
8   free( i0, i1, i2, i3, i4, i5, i6 );
```

```
1   0
2   5
3   0xfefefefe
4   5 5 5 5 5
5   0xdeadbeef 0xdeadbeef 0xdeadbeef 0xdeadbeef 0xdeadbeef
6   5 5 5 5 5
7   5 5 5 −555819298 −555819298  // two undefined values
```

  Examples 1 to 3 fill an object with a value or characters. Examples 4 to 7 fill an array of objects with values, another array, or part of an array.

- S_resize(T) ?`resize(**void** * oaddr) used to resize, realign, and fill, where the old object data is not copied to the new object. The old object type may be different from the new object type, since the values are not used. For example:

```
1   int * ip = alloc( 5`fill );  sout | ip | *ip;
2   ip = alloc( ip`resize, 256`align, 7`fill );  sout | ip | *ip;
3   double * dp = alloc( ip`resize, 4096`align, 13.5`fill );  sout | dp | *dp;
4   free( dp );  // DO NOT FREE ip AS ITS STORAGE IS MOVED TO dp
```

```
1   0x555555580a80 5
2   0x555555581100 7
3   0x555555587000 13.5
```

  Examples 2 to 3 change the alignment, fill, and size for the initial storage of i.

```
1   int * ia = alloc( 5, 5`fill );  sout | ia | nonl;  for ( i; 5 ) sout | ia[i] | nonl; sout | nl;
2   ia = alloc( 10, ia`resize, 7`fill );  sout | ia | nonl;  for ( i; 10 ) sout | ia[i] | nonl; sout | nl;
3   ia = alloc( 5, ia`resize, 512`align, 13`fill );  sout | ia | nonl;  for ( i; 5 ) sout | ia[i] | nonl; sout | nl;;
4   ia = alloc( 3, ia`resize, 4096`align, 2`fill );  for ( i; 3 ) sout | &ia[i] | ia[i] | nonl; sout | nl;
5   free( ia );
```

```
1   0x55555656d540 5 5 5 5 5
2   0x55555656d480 7 7 7 7 7 7 7 7 7 7
3   0x55555656fe00 13 13 13 13 13
4   0x555556570000 2 0x555556570004 2 0x555556570008 2
```

  Examples 2 to 4 change the array size, alignment, and fill initializes all storage because no data is copied.

- S_realloc(T) ?`realloc(T * a)) used to resize, realign, and fill, where the old object data is copied to the new object. The old object type must be the same as the new object type, since the value is used. Note, for fill, only the extra space after copying the data from the old object is filled with the given parameter. For example:

```
1   int * ip = alloc( 5`fill );  sout | ip | *ip;
2   ip = alloc( ip`realloc, 256`align );  sout | ip | *ip;
3   ip = alloc( ip`realloc, 4096`align, 13`fill );  sout | ip | *ip;
4   free( ip );
```

```
1      0x55555556d5c0 5
2      0x555555570000 5
3      0x555555571000 5
```

Examples 2 to 3 change the alignment for the initial storage of i. The 13`fill in example 3 does nothing because no new storage is added.

```
1    int * ia = alloc( 5, 5`fill );  sout | ia | nonl;  for ( i; 5 ) sout | ia[i] | nonl; sout | nl;
2    ia = alloc( 10, ia`realloc, 7`fill );  sout | ia | nonl;  for ( i; 10 ) sout | ia[i] | nonl; sout | nl;
3    ia = alloc( 5, ia`realloc, 512`align, 13`fill );  sout | ia | nonl; for ( i; 5 ) sout | ia[i] | nonl; sout | nl;;
4    ia = alloc( 3, ia`realloc, 4096`align, 2`fill );  for ( i; 3 ) sout | &ia[i] | ia[i] | nonl; sout | nl;
5    free( ia );
```

```
1    0x55555656d540 5 5 5 5 5
2    0x55555656d480 7 7 7 7 7 7 7 7 7 7
3    0x555556570e00 5 5 5 5 5
4    0x5555556571000 5 0x555556571004 5 0x555556571008 5
```

Examples 2 to 4 change the array size, alignment, and fill does no initialization after the copied data, as no new storage is added.

```
extern "C" {
    // New C allocation operations.
    void * aalloc( size_t dim, size_t elemSize );
    void * resize( void * oaddr, size_t size );
    void * amemalign( size_t align, size_t dim, size_t elemSize );
    void * cmemalign( size_t align, size_t dim, size_t elemSize );
    size_t malloc_alignment( void * addr );
    bool malloc_zero_fill( void * addr );
    size_t malloc_size( void * addr );
    int malloc_stats_fd( int fd );
    size_t malloc_expansion();          // heap expansion size (bytes)
    size_t malloc_mmap_start();         // crossover allocation size from sbrk to mmap
    size_t malloc_unfreed();            // heap unfreed size (bytes)
    void malloc_stats_clear();          // clear heap statistics
}

// New allocation operations.
void * resize( void * oaddr, size_t alignment, size_t size );
void * realloc( void * oaddr, size_t alignment, size_t size );
void * reallocarray( void * oaddr, size_t nalign, size_t dim, size_t elemSize );

forall( T & | sized(T) ) {
    // C∀  safe equivalents, i.e., implicit size specification, eliminate return−type cast
    T * malloc( void );
    T * aalloc( size_t dim );
    T * calloc( size_t dim );
    T * resize( T * ptr, size_t size );
    T * resize( T * ptr, size_t alignment, size_t size );
    T * realloc( T * ptr, size_t size );
    T * realloc( T * ptr, size_t alignment, size_t size );
    T * reallocarray( T * ptr, size_t dim );
    T * reallocarray( T * ptr, size_t alignment, size_t dim );
    T * memalign( size_t align );
    T * amemalign( size_t align, size_t dim );
    T * cmemalign( size_t align, size_t dim );
    T * aligned_alloc( size_t align );
    int posix_memalign( T ** ptr, size_t align );
    T * valloc( void );
    T * pvalloc( void );

    // C∀  safe general allocation, fill, resize, alignment, array
```

```
1    T * alloc( ... );                        // variable, T size
2    T * alloc( size_t dim, ... );
3    T_align ?`align( size_t alignment );
4    T_fill(T) ?`fill( /* various types */ );
5    T_resize ?`resize( void * oaddr );
6    T_realloc ?`realloc( void * oaddr ));
7    }
8
9    forall( T &, List ... ) void free( T * ptr, ... ) // deallocation list
10
11   // C∀  allocation/deallocation and constructor/destructor, non−array types
12   forall( T &, Parms ... | { void ?{}( T &, Parms ); } ) T * new( Parms ... );
13   forall( T &, List ... | { void ^?{}( T & ); void delete( List ... ); } );
14   // C∀  allocation/deallocation and constructor/destructor, array types
15   forall( T & | sized(T), Parms ... | { void ?{}( T &, Parms ); } ) T * anew( size_t dim, Parms ... );
16   forall( T & | sized(T) | { void ^?{}( T & ); }, List ... } ) void adelete( T arr[], List ... );
```

17 E.2 Memory Set and Copy

18 Like safe memory allocation, C∀ provides safe block initialization and copy. While objects should be initialized/copied
19 with constructors/assignment, block operations can be very performant. In certain cases the compiler generates block
20 copy operations, such as assigning structures s = t, however C arrays cannot be assigned.

```
21
22   struct S { int i, j, k; };
23   S s, t, *sp = &s, * tp = &t, sa[10], ta[10];
```

C∀	C
memset(s, '\0');	memset(&s, '\0', **sizeof**(s));
memset(sp, '\0');	memset(sp, '\0', **sizeof**(s));
24 memcpy(s, t);	memcpy(&s, &t, **sizeof**(s));
memcpy(sp, tp);	memcpy(sp, tp, **sizeof**(s));
amemset(sa, '\0', 10);	memset(sa, '\0', **sizeof**(sa));
amemcpy(sa, ta, 10);	memcpy(sa, ta, **sizeof**(sa));

25 These operations provide uniformity between reference and pointer, so object dereferencing, '&', is unnecessary.

```
26   static inline forall( T & | sized(T) ) {
27       // CFA safe initialization/copy, i.e., implicit size specification, non−array types
28       T * memset( T * dest, char fill );        // all combinations of pointer/reference
29       T * memset( T & dest, char fill );
30
31       T * memcpy( T * dest, const T * src );  // all combinations of pointer/reference
32       T * memcpy( T & dest, const T & src );
33       T * memcpy( T * dest, const T & src );
34       T * memcpy( T & dest, const T * src );
35
36       // CFA safe initialization/copy, i.e., implicit size specification, array types
37       T * amemset( T dest[], char fill, size_t dim );
38       T * amemcpy( T dest[], const T src[], size_t dim );
39   }
```

40 E.3 String to Value Conversion

```
41   int ato( const char * ptr );
42   unsigned int ato( const char * ptr );
43   long int ato( const char * ptr );
44   unsigned long int ato( const char * ptr );
45   long long int ato( const char * ptr );
46   unsigned long long int ato( const char * ptr );
```

```
1    float ato( const char * ptr );
2    double ato( const char * ptr );
3    long double ato( const char * ptr );
4    float _Complex ato( const char * ptr );
5    double _Complex ato( const char * ptr );
6    long double _Complex ato( const char * ptr );
7
8    int strto( const char * sptr, char ** eptr, int base );
9    unsigned int strto( const char * sptr, char ** eptr, int base );
10   long int strto( const char * sptr, char ** eptr, int base );
11   unsigned long int strto( const char * sptr, char ** eptr, int base );
12   long long int strto( const char * sptr, char ** eptr, int base );
13   unsigned long long int strto( const char * sptr, char ** eptr, int base );
14   float strto( const char * sptr, char ** eptr );
15   double strto( const char * sptr, char ** eptr );
16   long double strto( const char * sptr, char ** eptr );
17   float _Complex strto( const char * sptr, char ** eptr );
18   double _Complex strto( const char * sptr, char ** eptr );
19   long double _Complex strto( const char * sptr, char ** eptr );
```

20   E.4   Search / Sort

```
21   forall( T | { int ?<?( T, T ); } )              // location
22   T * bsearch( T key, const T * arr, size_t dim );
23
24   forall( T | { int ?<?( T, T ); } )              // position
25   unsigned int bsearch( T key, const T * arr, size_t dim );
26
27   forall( T | { int ?<?( T, T ); } )
28   void qsort( const T * arr, size_t dim );
29
30   forall( E | { int ?<?( E, E ); } ) {
31      E * bsearch( E key, const E * vals, size_t dim ); // location
32      size_t bsearch( E key, const E * vals, size_t dim );// position
33      E * bsearchl( E key, const E * vals, size_t dim );
34      size_t bsearchl( E key, const E * vals, size_t dim );
35      E * bsearchu( E key, const E * vals, size_t dim );
36      size_t bsearchu( E key, const E * vals, size_t dim );
37   }
38
39   forall( K, E | { int ?<?( K, K ); K getKey( const E & ); } ) {
40      E * bsearch( K key, const E * vals, size_t dim );
41      size_t bsearch( K key, const E * vals, size_t dim );
42      E * bsearchl( K key, const E * vals, size_t dim );
43      size_t bsearchl( K key, const E * vals, size_t dim );
44      E * bsearchu( K key, const E * vals, size_t dim );
45      size_t bsearchu( K key, const E * vals, size_t dim );
46   }
47
48   forall( E | { int ?<?( E, E ); } ) {
49      void qsort( E * vals, size_t dim );
50   }
```

51   E.5   Absolute Value

```
52   unsigned char abs( signed char );
53   int abs( int );
54   unsigned long int abs( long int );
55   unsigned long long int abs( long long int );
56   float abs( float );
```

1   **double** abs(**double**);
2   **long double** abs(**long double**);
3   **float** abs(**float _Complex**);
4   **double** abs(**double _Complex**);
5   **long double** abs(**long double _Complex**);
6   **forall**(T | { **void** ?{}(T *, **zero_t**); **int** ?<?(T, T); T −?(T); })
7   T abs(T);

8   E.6   C Random Numbers

9   **void** srandom(**unsigned int** seed);
10   **char** random(**void**);
11   **char** random(**char** u);   *// [0,u)*
12   **char** random(**char** l, **char** u);   *// [l,u]*
13   **int** random(**void**);
14   **int** random(**int** u);   *// [0,u)*
15   **int** random(**int** l, **int** u);   *// [l,u]*
16   **unsigned int** random(**void**);
17   **unsigned int** random(**unsigned int** u); *// [0,u)*
18   **unsigned int** random(**unsigned int** l, **unsigned int** u); *// [l,u]*
19   **long int** random(**void**);
20   **long int** random(**long int** u);   *// [0,u)*
21   **long int** random(**long int** l, **long int** u); *// [l,u]*
22   **unsigned long int** random(**void**);
23   **unsigned long int** random(**unsigned long int** u); *// [0,u)*
24   **unsigned long int** random(**unsigned long int** l, **unsigned long int** u); *// [l,u]*
25   **float** random(**void**);   *// [0.0, 1.0)*
26   **double** random(**void**);   *// [0.0, 1.0)*
27   **float _Complex** random(**void**);   *// [0.0, 1.0)+[0.0, 1.0)i*
28   **double _Complex** random(**void**);   *// [0.0, 1.0)+[0.0, 1.0)i*
29   **long double _Complex** random(**void**); *// [0.0, 1.0)+[0.0, 1.0)i*

30   E.7   Algorithms

31   **forall**(T | { **int** ?<?(T, T); }) T min(T t1, T t2);
32   **forall**(T | { **int** ?>?(T, T); }) T max(T t1, T t2);
33   **forall**(T | { T min(T, T); T max(T, T); }) T clamp(T value, T min_val, T max_val);
34   **forall**(T) **void** swap(T * t1, T * t2);

35   F   Math Library

36   The C∀ math-library wraps explicitly-polymorphic C math-routines into implicitly-polymorphic versions.

37   F.1   General

38   **float** ?%?(**float**, **float**);
39   **float** fmod(**float**, **float**);
40   **double** ?%?(**double**, **double**);
41   **double** fmod(**double**, **double**);
42   **long double** ?%?(**long double**, **long double**);
43   **long double** fmod(**long double**, **long double**);
44
45   **float** remainder(**float**, **float**);
46   **double** remainder(**double**, **double**);
47   **long double** remainder(**long double**, **long double**);
48
49   **float** remquo(**float**, **float**, **int** *);
50   **double** remquo(**double**, **double**, **int** *);
51   **long double** remquo(**long double**, **long double**, **int** *);
52   [**int**, **float**] remquo(**float**, **float**);

```
1        [ int, double ] remquo( double, double );
2        [ int, long double ] remquo( long double, long double );
3
4        [ int, float ] div( float, float );
5        [ int, double ] div( double, double );
6        [ int, long double ] div( long double, long double );
7
8        float fma( float, float, float );
9        double fma( double, double, double );
10       long double fma( long double, long double, long double );
11
12       float fdim( float, float );
13       double fdim( double, double );
14       long double fdim( long double, long double );
15
16       float nan( const char * );
17       double nan( const char * );
18       long double nan( const char * );
```

19   F.2   Exponential

```
20       float exp( float );
21       double exp( double );
22       long double exp( long double );
23       float _Complex exp( float _Complex );
24       double _Complex exp( double _Complex );
25       long double _Complex exp( long double _Complex );
26
27       float exp2( float );
28       double exp2( double );
29       long double exp2( long double );
30       // float _Complex exp2( float _Complex );
31       // double _Complex exp2( double _Complex );
32       // long double _Complex exp2( long double _Complex );
33
34       float expm1( float );
35       double expm1( double );
36       long double expm1( long double );
37
38       float pow( float, float );
39       double pow( double, double );
40       long double pow( long double, long double );
41       float _Complex pow( float _Complex, float _Complex );
42       double _Complex pow( double _Complex, double _Complex );
43       long double _Complex pow( long double _Complex, long double _Complex );
```

44   F.3   Logarithm

```
45       float log( float );
46       double log( double );
47       long double log( long double );
48       float _Complex log( float _Complex );
49       double _Complex log( double _Complex );
50       long double _Complex log( long double _Complex );
51
52       int log2( unsigned int );
53       long int log2( unsigned long int );
54       long long int log2( unsigned long long int )
55       float log2( float );
56       double log2( double );
```

```
1    long double log2( long double );
2    // float _Complex log2( float _Complex );
3    // double _Complex log2( double _Complex );
4    // long double _Complex log2( long double _Complex );
5
6    float log10( float );
7    double log10( double );
8    long double log10( long double );
9    // float _Complex log10( float _Complex );
10   // double _Complex log10( double _Complex );
11   // long double _Complex log10( long double _Complex );
12
13   float log1p( float );
14   double log1p( double );
15   long double log1p( long double );
16
17   int ilogb( float );
18   int ilogb( double );
19   int ilogb( long double );
20
21   float logb( float );
22   double logb( double );
23   long double logb( long double );
24
25   float sqrt( float );
26   double sqrt( double );
27   long double sqrt( long double );
28   float _Complex sqrt( float _Complex );
29   double _Complex sqrt( double _Complex );
30   long double _Complex sqrt( long double _Complex );
31
32   float cbrt( float );
33   double cbrt( double );
34   long double cbrt( long double );
35
36   float hypot( float, float );
37   double hypot( double, double );
38   long double hypot( long double, long double );
```

39   F.4   Trigonometric

```
40   float sin( float );
41   double sin( double );
42   long double sin( long double );
43   float _Complex sin( float _Complex );
44   double _Complex sin( double _Complex );
45   long double _Complex sin( long double _Complex );
46
47   float cos( float );
48   double cos( double );
49   long double cos( long double );
50   float _Complex cos( float _Complex );
51   double _Complex cos( double _Complex );
52   long double _Complex cos( long double _Complex );
53
54   float tan( float );
55   double tan( double );
56   long double tan( long double );
57   float _Complex tan( float _Complex );
58   double _Complex tan( double _Complex );
```

```
1      long double _Complex tan( long double _Complex );
2
3      float asin( float );
4      double asin( double );
5      long double asin( long double );
6      float _Complex asin( float _Complex );
7      double _Complex asin( double _Complex );
8      long double _Complex asin( long double _Complex );
9
10     float acos( float );
11     double acos( double );
12     long double acos( long double );
13     float _Complex acos( float _Complex );
14     double _Complex acos( double _Complex );
15     long double _Complex acos( long double _Complex );
16
17     float atan( float );
18     double atan( double );
19     long double atan( long double );
20     float _Complex atan( float _Complex );
21     double _Complex atan( double _Complex );
22     long double _Complex atan( long double _Complex );
23
24     float atan2( float, float );
25     double atan2( double, double );
26     long double atan2( long double, long double );
27
28     float atan( float, float );               // alternative name for atan2
29     double atan( double, double );
30     long double atan( long double, long double );
```

31 F.5   Hyperbolic

```
32     float sinh( float );
33     double sinh( double );
34     long double sinh( long double );
35     float _Complex sinh( float _Complex );
36     double _Complex sinh( double _Complex );
37     long double _Complex sinh( long double _Complex );
38
39     float cosh( float );
40     double cosh( double );
41     long double cosh( long double );
42     float _Complex cosh( float _Complex );
43     double _Complex cosh( double _Complex );
44     long double _Complex cosh( long double _Complex );
45
46     float tanh( float );
47     double tanh( double );
48     long double tanh( long double );
49     float _Complex tanh( float _Complex );
50     double _Complex tanh( double _Complex );
51     long double _Complex tanh( long double _Complex );
52
53     float asinh( float );
54     double asinh( double );
55     long double asinh( long double );
56     float _Complex asinh( float _Complex );
57     double _Complex asinh( double _Complex );
58     long double _Complex asinh( long double _Complex );
```

```
1
2       float acosh( float );
3       double acosh( double );
4       long double acosh( long double );
5       float _Complex acosh( float _Complex );
6       double _Complex acosh( double _Complex );
7       long double _Complex acosh( long double _Complex );
8
9       float atanh( float );
10      double atanh( double );
11      long double atanh( long double );
12      float _Complex atanh( float _Complex );
13      double _Complex atanh( double _Complex );
14      long double _Complex atanh( long double _Complex );
```

15   F.6   Error / Gamma

```
16      float erf( float );
17      double erf( double );
18      long double erf( long double );
19      float _Complex erf( float _Complex );
20      double _Complex erf( double _Complex );
21      long double _Complex erf( long double _Complex );
22
23      float erfc( float );
24      double erfc( double );
25      long double erfc( long double );
26      float _Complex erfc( float _Complex );
27      double _Complex erfc( double _Complex );
28      long double _Complex erfc( long double _Complex );
29
30      float lgamma( float );
31      double lgamma( double );
32      long double lgamma( long double );
33      float lgamma( float, int * );
34      double lgamma( double, int * );
35      long double lgamma( long double, int * );
36
37      float tgamma( float );
38      double tgamma( double );
39      long double tgamma( long double );
```

40   F.7   Nearest Integer

```
41      // n / align * align
42      signed char floor( signed char n, signed char align );
43      unsigned char floor( unsigned char n, unsigned char align );
44      short int floor( short int n, short int align );
45      unsigned short int floor( unsigned short int n, unsigned short int align );
46      int floor( int n, int align );
47      unsigned int floor( unsigned int n, unsigned int align );
48      long int floor( long int n, long int align );
49      unsigned long int floor( unsigned long int n, unsigned long int align );
50      long long int floor( long long int n, long long int align );
51      unsigned long long int floor( unsigned long long int n, unsigned long long int align );
52
53      // (n + (align − 1)) / align
54      signed char ceiling_div( signed char n, char align );
55      unsigned char ceiling_div( unsigned char n, unsigned char align );
56      short int ceiling_div( short int n, short int align );
```

```
1     unsigned short int ceiling_div( unsigned short int n, unsigned short int align );
2     int ceiling_div( int n, int align );
3     unsigned int ceiling_div( unsigned int n, unsigned int align );
4     long int ceiling_div( long int n, long int align );
5     unsigned long int ceiling_div( unsigned long int n, unsigned long int align );
6     long long int ceiling_div( long long int n, long long int align );
7     unsigned long long int ceiling_div( unsigned long long int n, unsigned long long int align );
8
9     // floor( n + (n % align != 0 ? align − 1 : 0), align )
10    signed char ceiling( signed char n, signed char align );
11    unsigned char ceiling( unsigned char n, unsigned char align );
12    short int ceiling( short int n, short int align );
13    unsigned short int ceiling( unsigned short int n, unsigned short int align );
14    int ceiling( int n, int align );
15    unsigned int ceiling( unsigned int n, unsigned int align );
16    long int ceiling( long int n, long int align );
17    unsigned long int ceiling( unsigned long int n, unsigned long int align );
18    long long int ceiling( long long int n, long long int align );
19    unsigned long long int ceiling( unsigned long long int n, unsigned long long int align );
20
21    float floor( float );
22    double floor( double );
23    long double floor( long double );
24
25    float ceil( float );
26    double ceil( double );
27    long double ceil( long double );
28
29    float trunc( float );
30    double trunc( double );
31    long double trunc( long double );
32
33    float rint( float );
34    long double rint( long double );
35    long int rint( float );
36    long int rint( double );
37    long int rint( long double );
38    long long int rint( float );
39    long long int rint( double );
40    long long int rint( long double );
41
42    long int lrint( float );
43    long int lrint( double );
44    long int lrint( long double );
45    long long int llrint( float );
46    long long int llrint( double );
47    long long int llrint( long double );
48
49    float nearbyint( float );
50    double nearbyint( double );
51    long double nearbyint( long double );
52
53    float round( float );
54    long double round( long double );
55    long int round( float );
56    long int round( double );
57    long int round( long double );
58    long long int round( float );
59    long long int round( double );
60    long long int round( long double );
```

1
2   **long int** lround(**float**);
3   **long int** lround(**double**);
4   **long int** lround(**long double**);
5   **long long int** llround(**float**);
6   **long long int** llround(**double**);
7   **long long int** llround(**long double**);

8   F.8   Manipulation

9   **float** copysign(**float**, **float**);
10   **double** copysign(**double**, **double**);
11   **long double** copysign(**long double**, **long double**);
12
13   **float** frexp(**float**, **int** *);
14   **double** frexp(**double**, **int** *);
15   **long double** frexp(**long double**, **int** *);
16
17   **float** ldexp(**float**, **int**);
18   **double** ldexp(**double**, **int**);
19   **long double** ldexp(**long double**, **int**);
20
21   [**float**, **float**] modf(**float**);
22   **float** modf(**float**, **float** *);
23   [**double**, **double**] modf(**double**);
24   **double** modf(**double**, **double** *);
25   [**long double**, **long double**] modf(**long double**);
26   **long double** modf(**long double**, **long double** *);
27
28   **float** nextafter(**float**, **float**);
29   **double** nextafter(**double**, **double**);
30   **long double** nextafter(**long double**, **long double**);
31
32   **float** nexttoward(**float**, **long double**);
33   **double** nexttoward(**double**, **long double**);
34   **long double** nexttoward(**long double**, **long double**);
35
36   **float** scalbn(**float**, **int**);
37   **double** scalbn(**double**, **int**);
38   **long double** scalbn(**long double**, **int**);
39
40   **float** scalbln(**float**, **long int**);
41   **double** scalbln(**double**, **long int**);
42   **long double** scalbln(**long double**, **long int**);

43   G   Time Keeping

44   G.1   Duration

45   **struct** Duration {
46   int64_t tn;   *// nanoseconds*
47   };
48
49   **void** ?{}(Duration & dur);
50   **void** ?{}(Duration & dur, **zero_t**);
51   **void** ?{}(Duration & dur, timeval t)
52   **void** ?{}(Duration & dur, timespec t)
53
54   Duration ?=?(Duration & dur, **zero_t**);
55   Duration ?=?(Duration & dur, timeval t)

```
1       Duration ?=?( Duration & dur, timespec t )
2
3       Duration +?( Duration rhs );
4       Duration ?+?( Duration & lhs, Duration rhs );
5       Duration ?+=?( Duration & lhs, Duration rhs );
6
7       Duration −?( Duration rhs );
8       Duration ?−?( Duration & lhs, Duration rhs );
9       Duration ?−=?( Duration & lhs, Duration rhs );
10
11      Duration ?*?( Duration lhs, int64_t rhs );
12      Duration ?*?( int64_t lhs, Duration rhs );
13      Duration ?*=?( Duration & lhs, int64_t rhs );
14
15      int64_t ?/?( Duration lhs, Duration rhs );
16      Duration ?/?( Duration lhs, int64_t rhs );
17      Duration ?/=?( Duration & lhs, int64_t rhs );
18      double div( Duration lhs, Duration rhs );
19
20      Duration ?%?( Duration lhs, Duration rhs );
21      Duration ?%=?( Duration & lhs, Duration rhs );
22
23      bool ?==?( Duration lhs, zero_t );
24      bool ?!=?( Duration lhs, zero_t );
25      bool ?<? ( Duration lhs, zero_t );
26      bool ?<=?( Duration lhs, zero_t );
27      bool ?>? ( Duration lhs, zero_t );
28      bool ?>=?( Duration lhs, zero_t );
29
30      bool ?==?( Duration lhs, Duration rhs );
31      bool ?!=?( Duration lhs, Duration rhs );
32      bool ?<? ( Duration lhs, Duration rhs );
33      bool ?<=?( Duration lhs, Duration rhs );
34      bool ?>? ( Duration lhs, Duration rhs );
35      bool ?>=?( Duration lhs, Duration rhs );
36
37      Duration abs( Duration rhs );
38
39      Duration ?`ns( int64_t nsec );
40      Duration ?`us( int64_t usec );
41      Duration ?`ms( int64_t msec );
42      Duration ?`s( int64_t sec );
43      Duration ?`s( double sec );
44      Duration ?`m( int64_t min );
45      Duration ?`m( double min );
46      Duration ?`h( int64_t hours );
47      Duration ?`h( double hours );
48      Duration ?`d( int64_t days );
49      Duration ?`d( double days );
50      Duration ?`w( int64_t weeks );
51      Duration ?`w( double weeks );
52
53      int64_t ?`ns( Duration dur );
54      int64_t ?`us( Duration dur );
55      int64_t ?`ms( Duration dur );
56      int64_t ?`s( Duration dur );
57      int64_t ?`m( Duration dur );
58      int64_t ?`h( Duration dur );
59      int64_t ?`d( Duration dur );
60      int64_t ?`w( Duration dur );
```

1
2   **double** ?`dns(Duration dur);
3   **double** ?`dus(Duration dur);
4   **double** ?`dms(Duration dur);
5   **double** ?`ds(Duration dur);
6   **double** ?`dm(Duration dur);
7   **double** ?`dh(Duration dur);
8   **double** ?`dd(Duration dur);
9   **double** ?`dw(Duration dur);
10
11   Duration max(Duration lhs, Duration rhs);
12   Duration min(Duration lhs, Duration rhs);
13
14   **forall**(ostype & | ostream(ostype)) ostype & ?|?(ostype & os, Duration dur);

15   G.2   timeval

16   **void** ?{}(timeval & t);
17   **void** ?{}(timeval & t, **zero_t**);
18   **void** ?{}(timeval & t, time_t sec, suseconds_t usec);
19   **void** ?{}(timeval & t, time_t sec);
20   **void** ?{}(timeval & t, Time time);
21
22   timeval ?=?(timeval & t, **zero_t**);
23   timeval ?+?(timeval & lhs, timeval rhs);
24   timeval ?−?(timeval & lhs, timeval rhs);
25   **bool** ?==?(timeval lhs, timeval rhs);
26   **bool** ?!=?(timeval lhs, timeval rhs);

27   G.3   timespec

28   **void** ?{}(timespec & t);
29   **void** ?{}(timespec & t, **zero_t**);
30   **void** ?{}(timespec & t, time_t sec, __syscall_slong_t nsec);
31   **void** ?{}(timespec & t, time_t sec);
32   **void** ?{}(timespec & t, Time time);
33
34   timespec ?=?(timespec & t, **zero_t**);
35   timespec ?+?(timespec & lhs, timespec rhs);
36   timespec ?−?(timespec & lhs, timespec rhs);
37   **bool** ?==?(timespec lhs, timespec rhs);
38   **bool** ?!=?(timespec lhs, timespec rhs);

39   G.4   itimerval

40   **void** ?{}(itimerval & itv, Duration alarm);
41   **void** ?{}(itimerval & itv, Duration alarm, Duration interval);

42   G.5   Time

43   **struct** Time {
44   uint64_t tn;   *// nanoseconds since UNIX epoch*
45   };
46
47   **void** ?{}(Time & time);
48   **void** ?{}(Time & time, **zero_t**);
49   **void** ?{}(Time & time, timeval t);
50   **void** ?{}(Time & time, timespec t);
51
52   Time ?=?(Time & time, **zero_t**);
53   Time ?=?(Time & time, timeval t);

```
1     Time ?=?( Time & time, timespec t );
2
3     Time ?+?( Time & lhs, Duration rhs );
4     Time ?+?( Duration lhs, Time rhs );
5     Time ?+=?( Time & lhs, Duration rhs );
6
7     Duration ?−?( Time lhs, Time rhs );
8     Time ?−?( Time lhs, Duration rhs );
9     Time ?−=?( Time & lhs, Duration rhs );
10    bool ?==?( Time lhs, Time rhs );
11    bool ?!=?( Time lhs, Time rhs );
12    bool ?<?( Time lhs, Time rhs );
13    bool ?<=?( Time lhs, Time rhs );
14    bool ?>?( Time lhs, Time rhs );
15    bool ?>=?( Time lhs, Time rhs );
16
17    int64_t ?`ns( Time t );
18
19    char * yy_mm_dd( Time time, char * buf );
20    char * ?`ymd( Time time, char * buf ); // short form
21
22    char * mm_dd_yy( Time time, char * buf );
23    char * ?`mdy( Time time, char * buf ); // short form
24
25    char * dd_mm_yy( Time time, char * buf );
26    char * ?`dmy( Time time, char * buf ); // short form
27
28    size_t strftime( char * buf, size_t size, const char * fmt, Time time );
29
30    forall( ostype & | ostream( ostype ) ) ostype & ?|?( ostype & os, Time time );
```

H   Clock

H.1   C time

```
33    char * ctime( time_t tp );
34    char * ctime_r( time_t tp, char * buf );
35    tm * gmtime( time_t tp );
36    tm * gmtime_r( time_t tp, tm * result );
37    tm * localtime( time_t tp );
38    tm * localtime_r( time_t tp, tm * result );
```

H.2   Clock

```
40    struct Clock {                           // virtual clock
41        Duration offset;                     // offset from computer real-time
42    };
43
44    void ?{}( Clock & clk );                  // create no offset
45    void ?{}( Clock & clk, Duration adj );    // create with offset
46    void reset( Clock & clk, Duration adj );  // change offset
47
48    Duration resolutionHi();                  // clock resolution in nanoseconds (fine)
49    Duration resolution();                    // clock resolution without nanoseconds (coarse)
50
51    Time timeHiRes();                         // real time with nanoseconds
52    Time time();                              // real time without nanoseconds
53    Time time( Clock & clk );                 // real time for given clock
54    Time ?()( Clock & clk );                  //   alternative syntax
55    timeval time( Clock & clk );              // convert to C time format
```

```
1   tm time( Clock & clk );
2   Duration processor();              // non-monotonic duration of kernel thread
3   Duration program();               // non-monotonic duration of program CPU
4   Duration boot();                  // monotonic duration since computer boot
```

I   Pseudo Random Number Generator

Random numbers are values generated independently, i.e., new values do not depend on previous values (independent trials), *e.g.*, lottery numbers, shuffled cards, dice roll, coin flip. While a primary goal of programming is computing values that are *not* random, random values are useful in simulation, cryptography, games, etc. A random-number generator is an algorithm that computes independent values. If the algorithm uses deterministic computation (a predictable sequence of values), it generates *pseudo* random numbers versus *true* random numbers.

All *pseudo random-number generators* (*PRNG*) involve some technique to scramble bits of a value, *e.g.*, multiplicative recurrence:

```
13  rand = 33967 * (rand + 1063); // scramble bits
```

Multiplication of large values adds new least-significant bits and drops most-significant bits.

bits 63–32 (most)	bits 31–0 (least)
0x0	0x3e8e36
0x5f	0x718c25e1
0xad3e	0x7b5f1dbe
0xbc3b	0xac69ff19
0x1070f	0x2d258dc6

By dropping bits 63–32, bits 31–0 become scrambled after each multiply. The least-significant bits *appear* random but the same bits are always generated given a fixed starting value, called the *seed* (value 0x3e8e36 above). Hence, if a program uses the same seed, the same sequence of pseudo-random values is generated from the PRNG. Often the seed is set to another random value like a program's process identifier (getpid) or time when the program is run; hence, one random value bootstraps another. Finally, a PRNG usually generates a range of large values, *e.g.*, [0, UINT_MAX], which are scaled using the modulus operator, *e.g.*, prng() % 5 produces random values in the range 0–4.

C∀ provides 32/64-bit sequential PRNG type only accessible by a single thread (not thread-safe) and a set of global routines and companion thread PRNG functions accessible by multiple threads without contention. To use the PRNG interface **requires including** stdlib.hfa.

- The PRNG types for sequential programs, including coroutining, are:

```
26      struct PRNG32 {};                                       // opaque type, no copy or assignment
27      void ?{}( PRNG32 & prng, uint32_t seed );               // fixed seed
28      void ?{}( PRNG32 & prng );                              // random seed
29      void set_seed( PRNG32 & prng, uint32_t seed );          // set seed
30      uint32_t get_seed( PRNG32 & prng );                     // get seed
31      uint32_t prng( PRNG32 & prng );                         // [0,UINT_MAX]
32      uint32_t prng( PRNG32 & prng, uint32_t u );             // [0,u)
33      uint32_t prng( PRNG32 & prng, uint32_t l, uint32_t u ); // [l,u]
34      uint32_t calls( PRNG32 & prng );                        // number of calls
35      void copy( PRNG32 & dst, PRNG32 & src );                // checkpoint PRNG state

36      struct PRNG64 {};                                       // opaque type, no copy or assignment
37      void ?{}( PRNG64 & prng, uint64_t seed );               // fixed seed
38      void ?{}( PRNG64 & prng );                              // random seed
39      void set_seed( PRNG64 & prng, uint64_t seed );          // set seed
40      uint64_t get_seed( PRNG64 & prng );                     // get seed
41      uint64_t prng( PRNG64 & prng );                         // [0,UINT_MAX]
42      uint64_t prng( PRNG64 & prng, uint64_t u );             // [0,u)
43      uint64_t prng( PRNG64 & prng, uint64_t l, uint64_t u ); // [l,u]
44      uint64_t calls( PRNG64 & prng );                        // number of calls
45      void copy( PRNG64 & dst, PRNG64 & src );                // checkpoint PRNG state
```

The type PRNG is aliased to PRNG64 on 64-bit architectures and PRNG32 on 32-bit architectures. A PRNG object is used to randomize behaviour or values during execution, *e.g.*, in games, a character makes a random

```
PRNG sprng1, sprng2;                        // select appropriate 32/64-bit PRNG
set_seed( sprng1, 1009 );  set_seed( sprng2, 1009 );
for ( 10 ) {
    // Do not cascade prng calls because side−effect functions called in arbitrary order.
    sout | nlOff | prng( sprng1 ); sout | prng( sprng1, 5 ); sout | prng( sprng1, 0, 5 ) | '\t';
    sout | prng( sprng2 ); sout | prng( sprng2, 5 ); sout | prng( sprng2, 0, 5 ) | nlOn;
}

37301721 2 2       37301721 2 2
1681308562 1 3     1681308562 1 3
290112364 3 2      290112364 3 2
1852700364 4 3     1852700364 4 3
733221210 1 3      733221210 1 3
1775396023 2 3     1775396023 2 3
123981445 2 3      123981445 2 3
2062557687 2 0     2062557687 2 0
283934808 1 0      283934808 1 0
672325890 1 3      672325890 1 3
```

Figure 16: Sequential PRNG

move or an object takes on a random value. In this scenario, it is useful to have multiple PRNG objects, *e.g.*, one per player or object. However, sequential execution is still repeatable given the same starting seeds for all PRNGs. Figure 16 shows an example that creates two sequential PRNGs, sets both to the same seed (1009), and illustrates the three forms for generating random values, where both PRNGs generate the same sequence of values. Note, to prevent accidental PRNG copying, the copy constructor and assignment are hidden. To copy a PRNG for checkpointing, use the explicit copy member.

• The PRNG global and companion thread functions are for concurrent programming, such as randomizing execution in short-running programs, *e.g.*, yield(prng() % 5).

```
void set_seed( size_t seed );                   // set global seed
size_t get_seed();                              // get global seed
// SLOWER, global routines
size_t prng( void );                            // [0,UINT_MAX]
size_t prng( size_t u );                        // [0,u)
size_t prng( size_t l, size_t u );              // [l,u]
// FASTER, thread members
size_t prng( thread$ & th );                    // [0,UINT_MAX]
size_t prng( thread$ & th, size_t u );          // [0,u)
size_t prng( thread$ & th, size_t l, size_t u ); // [l,u]
```

The only difference between the two sets of prng routines is performance.

Because concurrent execution is non-deterministic, seeding the concurrent PRNG is less important, as repeatable execution is impossible. Hence, there is one system-wide PRNG (global seed) but each C∀ thread has its own non-contended PRNG state. If the global seed is set, threads start with this seed, until it is reset and then threads start with the reset seed. Hence, these threads generate the same sequence of random numbers from their specific starting seed. If the global seed is *not* set, threads start with a random seed, until the global seed is set. Hence, these threads generate different sequences of random numbers. If each thread needs its own seed, use a sequential PRNG in each thread. The slower prng global functions, *i.e.*, *without* a thread argument, call active_thread internally to indirectly access the current thread's PRNG state, while the faster prng functions, *i.e.*, *with* a thread argument, directly access the thread through the thread parameter. If a thread pointer is available, *e.g.*, in thread main, eliminating the call to active_thread significantly reduces the cost of accessing the thread's PRNG state. Figure 17 shows an example using the slower/faster concurrent PRNG in the program main and a thread.

```
thread T {};
void main( T & th ) {  // thread address
    for ( i; 10 ) {
        sout | nlOff | prng();  sout | prng( 5 );  sout | prng( 0, 5 ) | '\t';  // SLOWER
        sout | nlOff | prng( th );  sout | prng( th, 5 );  sout | prng( th, 0, 5 ) | nlOn;  // FASTER
    }
}
int main() {
    set_seed( 1009 );
    thread$ & th = *active_thread();  // program−main thread−address
    for ( i; 10 ) {
        sout | nlOff | prng();  sout | prng( 5 );  sout | prng( 0, 5 ) | '\t';  // SLOWER
        sout | nlOff | prng( th );  sout | prng( th, 5 );  sout | prng( th, 0, 5 ) | nlOn;  // FASTER
    }
    sout | nl;
    T t;  // run thread
}

37301721 2 2     1681308562 1 3
290112364 3 2    1852700364 4 3
733221210 1 3    1775396023 2 3
123981445 2 3    2062557687 2 0
283934808 1 0    672325890 1 3
1414344101 1 3   873424536 3 4
871831898 3 4    866783532 0 1
2142057611 4 4   17310256 2 5
802117363 0 4    492964499 0 0
2346353643 1 3   2143013105 3 2

// same output as above from thread t
```

Figure 17: Concurrent PRNG

J   Multi-precision Integers

C∀ has an interface to the GMP multi-precision signed-integers [17], similar to the C++ interface provided by GMP. The C∀ interface wraps GMP routines into operator routines to make programming with multi-precision integers identical to using fixed-sized integers. The C∀ type name for multi-precision signed-integers is Int and the header file is gmp.

```
void ?{}( Int * this );                          // constructor/destructor
void ?{}( Int * this, Int init );
void ?{}( Int * this, zero_t );
void ?{}( Int * this, one_t );
void ?{}( Int * this, signed long int init );
void ?{}( Int * this, unsigned long int init );
void ?{}( Int * this, const char * val );
void ^?{}( Int * this );

Int ?=?( Int * lhs, Int rhs );                   // assignment
Int ?=?( Int * lhs, long int rhs );
Int ?=?( Int * lhs, unsigned long int rhs );
Int ?=?( Int * lhs, const char * rhs );

char ?=?( char * lhs, Int rhs );
short int ?=?( short int * lhs, Int rhs );
int ?=?( int * lhs, Int rhs );
long int ?=?( long int * lhs, Int rhs );
unsigned char ?=?( unsigned char * lhs, Int rhs );
unsigned short int ?=?( unsigned short int * lhs, Int rhs );
unsigned int ?=?( unsigned int * lhs, Int rhs );
```

```
1      unsigned long int ?=?( unsigned long int * lhs, Int rhs );
2
3      long int narrow( Int val );
4      unsigned long int narrow( Int val );
5
6      int ?==?( Int oper1, Int oper2 );              // comparison
7      int ?==?( Int oper1, long int oper2 );
8      int ?==?( long int oper2, Int oper1 );
9      int ?==?( Int oper1, unsigned long int oper2 );
10     int ?==?( unsigned long int oper2, Int oper1 );
11
12     int ?!=?( Int oper1, Int oper2 );
13     int ?!=?( Int oper1, long int oper2 );
14     int ?!=?( long int oper1, Int oper2 );
15     int ?!=?( Int oper1, unsigned long int oper2 );
16     int ?!=?( unsigned long int oper1, Int oper2 );
17
18     int ?<?( Int oper1, Int oper2 );
19     int ?<?( Int oper1, long int oper2 );
20     int ?<?( long int oper2, Int oper1 );
21     int ?<?( Int oper1, unsigned long int oper2 );
22     int ?<?( unsigned long int oper2, Int oper1 );
23
24     int ?<=?( Int oper1, Int oper2 );
25     int ?<=?( Int oper1, long int oper2 );
26     int ?<=?( long int oper2, Int oper1 );
27     int ?<=?( Int oper1, unsigned long int oper2 );
28     int ?<=?( unsigned long int oper2, Int oper1 );
29
30     int ?>?( Int oper1, Int oper2 );
31     int ?>?( Int oper1, long int oper2 );
32     int ?>?( long int oper1, Int oper2 );
33     int ?>?( Int oper1, unsigned long int oper2 );
34     int ?>?( unsigned long int oper1, Int oper2 );
35
36     int ?>=?( Int oper1, Int oper2 );
37     int ?>=?( Int oper1, long int oper2 );
38     int ?>=?( long int oper1, Int oper2 );
39     int ?>=?( Int oper1, unsigned long int oper2 );
40     int ?>=?( unsigned long int oper1, Int oper2 );
41
42     Int +?( Int oper );                            // arithmetic
43     Int −?( Int oper );
44     Int ~?( Int oper );
45
46     Int ?&?( Int oper1, Int oper2 );
47     Int ?&?( Int oper1, long int oper2 );
48     Int ?&?( long int oper1, Int oper2 );
49     Int ?&?( Int oper1, unsigned long int oper2 );
50     Int ?&?( unsigned long int oper1, Int oper2 );
51     Int ?&=?( Int * lhs, Int rhs );
52
53     Int ?|?( Int oper1, Int oper2 );
54     Int ?|?( Int oper1, long int oper2 );
55     Int ?|?( long int oper1, Int oper2 );
56     Int ?|?( Int oper1, unsigned long int oper2 );
57     Int ?|?( unsigned long int oper1, Int oper2 );
58     Int ?|=?( Int * lhs, Int rhs );
59
60     Int ?^?( Int oper1, Int oper2 );
```

```
1    Int ?^?( Int oper1, long int oper2 );
2    Int ?^?( long int oper1, Int oper2 );
3    Int ?^?( Int oper1, unsigned long int oper2 );
4    Int ?^?( unsigned long int oper1, Int oper2 );
5    Int ?^=?( Int * lhs, Int rhs );
6
7    Int ?+?( Int addend1, Int addend2 );
8    Int ?+?( Int addend1, long int addend2 );
9    Int ?+?( long int addend2, Int addend1 );
10   Int ?+?( Int addend1, unsigned long int addend2 );
11   Int ?+?( unsigned long int addend2, Int addend1 );
12   Int ?+=?( Int * lhs, Int rhs );
13   Int ?+=?( Int * lhs, long int rhs );
14   Int ?+=?( Int * lhs, unsigned long int rhs );
15   Int ++?( Int * lhs );
16   Int ?++( Int * lhs );
17
18   Int ?−?( Int minuend, Int subtrahend );
19   Int ?−?( Int minuend, long int subtrahend );
20   Int ?−?( long int minuend, Int subtrahend );
21   Int ?−?( Int minuend, unsigned long int subtrahend );
22   Int ?−?( unsigned long int minuend, Int subtrahend );
23   Int ?−=?( Int * lhs, Int rhs );
24   Int ?−=?( Int * lhs, long int rhs );
25   Int ?−=?( Int * lhs, unsigned long int rhs );
26   Int −−?( Int * lhs );
27   Int ?−−( Int * lhs );
28
29   Int ?*?( Int multiplicator, Int multiplicand );
30   Int ?*?( Int multiplicator, long int multiplicand );
31   Int ?*?( long int multiplicand, Int multiplicator );
32   Int ?*?( Int multiplicator, unsigned long int multiplicand );
33   Int ?*?( unsigned long int multiplicand, Int multiplicator );
34   Int ?*=?( Int * lhs, Int rhs );
35   Int ?*=?( Int * lhs, long int rhs );
36   Int ?*=?( Int * lhs, unsigned long int rhs );
37
38   Int ?/?( Int dividend, Int divisor );
39   Int ?/?( Int dividend, unsigned long int divisor );
40   Int ?/?( unsigned long int dividend, Int divisor );
41   Int ?/?( Int dividend, long int divisor );
42   Int ?/?( long int dividend, Int divisor );
43   Int ?/=?( Int * lhs, Int rhs );
44   Int ?/=?( Int * lhs, long int rhs );
45   Int ?/=?( Int * lhs, unsigned long int rhs );
46
47   [ Int, Int ] div( Int dividend, Int divisor );
48   [ Int, Int ] div( Int dividend, unsigned long int divisor );
49
50   Int ?%?( Int dividend, Int divisor );
51   Int ?%?( Int dividend, unsigned long int divisor );
52   Int ?%?( unsigned long int dividend, Int divisor );
53   Int ?%?( Int dividend, long int divisor );
54   Int ?%?( long int dividend, Int divisor );
55   Int ?%=?( Int * lhs, Int rhs );
56   Int ?%=?( Int * lhs, long int rhs );
57   Int ?%=?( Int * lhs, unsigned long int rhs );
58
59   Int ?<<?( Int shiften, mp_bitcnt_t shift );
60   Int ?<<=?( Int * lhs, mp_bitcnt_t shift );
```

```
1     Int ?>>?( Int shiften, mp_bitcnt_t shift );
2     Int ?>>=?( Int * lhs, mp_bitcnt_t shift );
3
4     Int abs( Int oper );                        // number functions
5     Int fact( unsigned long int N );
6     Int gcd( Int oper1, Int oper2 );
7     Int pow( Int base, unsigned long int exponent );
8     Int pow( unsigned long int base, unsigned long int exponent );
9     void srandom( gmp_randstate_t state );
10    Int random( gmp_randstate_t state, mp_bitcnt_t n );
11    Int random( gmp_randstate_t state, Int n );
12    Int random( gmp_randstate_t state, mp_size_t max_size );
13    int sgn( Int oper );
14    Int sqrt( Int oper );
15
16    forall( istype & | istream( istype ) ) istype * ?|?( istype * is, Int * mp );  // I/O
17    forall( ostype & | ostream( ostype ) ) ostype * ?|?( ostype * os, Int mp );
```

18 Figure 18 shows C∀ and C factorial programs using the GMP interfaces. (Compile with flag –lgmp to link with the
19 GMP library.)

K   Rational Numbers

21 Rational numbers are numbers written as a ratio, *i.e.*, as a fraction, where the numerator (top number) and the denomi-
22 nator (bottom number) are whole numbers. When creating and computing with rational numbers, results are constantly
23 reduced to keep the numerator and denominator as small as possible.

```
24    // implementation
25    struct Rational {
26        long int numerator, denominator;   // invariant: denominator > 0
27    }; // Rational
28
29    Rational rational();                    // constructors
30    Rational rational( long int n );
31    Rational rational( long int n, long int d );
32    void ?{}( Rational * r, zero_t );
33    void ?{}( Rational * r, one_t );
34
35    long int numerator( Rational r );       // numerator/denominator getter/setter
36    long int numerator( Rational r, long int n );
37    long int denominator( Rational r );
38    long int denominator( Rational r, long int d );
39
40    int ?==?( Rational l, Rational r );     // comparison
41    int ?!=?( Rational l, Rational r );
42    int ?<?( Rational l, Rational r );
43    int ?<=?( Rational l, Rational r );
44    int ?>?( Rational l, Rational r );
45    int ?>=?( Rational l, Rational r );
46
47    Rational −?( Rational r );              // arithmetic
48    Rational ?+?( Rational l, Rational r );
49    Rational ?−?( Rational l, Rational r );
50    Rational ?*?( Rational l, Rational r );
51    Rational ?/?( Rational l, Rational r );
52
53    double widen( Rational r );             // conversion
54    Rational narrow( double f, long int md );
55
56    forall( istype & | istream( istype ) ) istype * ?|?( istype *, Rational * ); // I/O
57    forall( ostype & | ostream( ostype ) ) ostype * ?|?( ostype *, Rational );
```

C∀	C
| ```c
#include <gmp.hfa>
int main( void ) {
    sout | "Factorial Numbers";
    Int fact = 1;

    sout | 0 | fact;
    for ( i; 40 ) {
        fact *= i;
        sout | i | fact;
    }
}
``` | ```c
#include <gmp.h>
int main( void ) {
    gmp_printf( "Factorial Numbers\n" );
    mpz_t fact;
    mpz_init_set_ui( fact, 1 );
    gmp_printf( "%d %Zd\n", 0, fact );
    for ( unsigned int i = 1; i <= 40; i += 1 ) {
        mpz_mul_ui( fact, fact, i );
        gmp_printf( "%d %Zd\n", i, fact );
    }
}
``` |

```
Factorial Numbers
0 1
1 1
2 2
3 6
4 24
5 120
6 720
7 5040
8 40320
9 362880
10 3628800
11 39916800
12 479001600
13 6227020800
14 87178291200
15 1307674368000
16 20922789888000
17 355687428096000
18 6402373705728000
19 121645100408832000
20 2432902008176640000
21 51090942171709440000
22 1124000727777607680000
23 25852016738884976640000
24 620448401733239439360000
25 15511210043330985984000000
26 403291461126605635584000000
27 10888869450418352160768000000
28 304888344611713860501504000000
29 8841761993739701954543616000000
30 265252859812191058636308480000000
31 8222838654177922817725562880000000
32 263130836933693530167218012160000000
33 8683317618811886495518194401280000000
34 295232799039604140847618609643520000000
35 10333147966386144929666651337523200000000
36 371993326789901217467999448150835200000000
37 13763753091226345046315979581580902400000000
38 523022617466601111760007224100074291200000000
39 20397882081197443358640281739902897356800000000
40 815915283247897734345611269596115894272000000000
```

Figure 18: Multi-precision Factorials

References

[1] Ada16. *Ada Reference Manual ISO/IEC 8652:2012(E) with COR.1:2016*. AXE Consultants, Madison WI, USA, 3rd with technical corrigendum 1 for ada 2012 edition, 2016. https://docs.adacore.com/live/wave/arm12/pdf/arm12/arm-12.pdf. 2

[2] Richard C. Bilson. Implementing overloading and polymorphism in C∀. Master's thesis, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 2003. http://plg.uwaterloo.ca/theses/BilsonThesis.pdf. 2, 49

[3] Walter Bright and Andrei Alexandrescu. *D Programming Language*. Digital Mars, Vienna Virginia, U.S.A., 2016. http://dlang.org/spec/spec.html. 2

[4] P. A. Buhr. A case for teaching multi-exit loops to beginning programmers. *SIGPLAN Not.*, 20(11):14–22, November 1985. 12

[5] P. A. Buhr, David Till, and C. R. Zarnke. Assignment as the sole means of updating objects. *Softw. Pract. Exper.*, 24(9):835–870, September 1994. 2, 42

[6] Cobol14. *Programming Languages – Cobol ISO/IEC 1989:2014*. International Standard Organization, Geneva, Switzerland, 2nd edition, 2014. https://www.iso.org/standard/51416.html. 2

[7] G. V. Cormack and A. K. Wright. Polymorphism in the compiled language ForceOne. In *Proceedings of the 20th Hawaii International Conference on Systems Sciences*, pages 284–292, January 1987. 2

[8] G. V. Cormack and A. K. Wright. Type-dependent parameter inference. *SIGPLAN Not.*, 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. 2, 37

[9] Glen Jeffrey Ditchfield. *Contextual Polymorphism*. PhD thesis, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 1992. http://plg.uwaterloo.ca/theses/DitchfieldThesis.pdf. 2

[10] Tom Duff. Duff's device. http://www.lysator.liu.se/c/duffs-device.html, November 1983. 8, 9

[11] Dominic Duggan, Gordon V. Cormack, and John Ophel. Kinded type inference for parametric overloading. *Acta Infomatica*, 33(1):21–68, 1996. 2

[12] Margaret A. Ellis and Bjarne Stroustrup. *The Annotated C++ Reference Manual*. Addison-Wesley, Boston, 1st edition, 1990. 2

[13] Rodolfo Gabriel Esteves. C∀, a study in evolutionary design in programming languages. Master's thesis, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 2004. http://plg.uwaterloo.ca/theses/EstevesThesis.pdf. 2

[14] Fortran08. *Programming Languages – Fortran Part 1:Base Language ISO/IEC 1539-1:2010*. International Standard Organization, Geneva, Switzerland, 3rd edition, 2010. https://www.iso.org/standard/50459.html. 2

[15] Nissim Francez. Another advantage of keyword notation for parameter communication with subprograms. *Communications of the ACM*, 20(8):604–605, August 1977. 36

[16] John Galletly. *OCCAM 2: Including OCCAM 2.1*. UCL (University College London) Press, London, 2nd edition, 1996. 33, 51

[17] *GNU Multiple Precision Arithmetic Library*. GNU, 2016. https://gmplib.org. 95

[18] James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. *Java Language Specification*, Java SE 8 edition, 2015. 2

[19] Robert Griesemer, Rob Pike, and Ken Thompson. *Go Programming Language*. Google, Mountain View, CA, USA, 2009. https://go.dev/ref/spec. 2, 71

[20] W. T. Hardgrave. Positional versus keyword parameter communication in programming languages. *SIGPLAN Not.*, 11(5):52–58, May 1976. 36

[21] International Standard Organization, Geneva, Switzerland. *C Programming Language ISO/IEC 9889:2011-12*, 3rd edition, December 2011. https://www.iso.org/standard/57853.html. 2, 20, 22, 47, 74, 76

[22] International Standard Organization, Geneva, Switzerland. *C++ Programming Language ISO/IEC 14882:2014*, 4th edition, 2014. https://www.iso.org/standard/64029.html. 2, 73

[23] Kathleen Jensen and Niklaus Wirth. *Pascal User Manual and Report, ISO Pascal Standard*. Springer–Verlag, 4th edition, 1991. Revised by Andrew B. Mickel and James F. Miner. 16

[24] C. H. Lindsey and S. G. van der Meulen. *Informal Introduction to ALGOL 68*. North-Holland, London, 1977. 21

[25] 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, New York, 1981. 33, 51

[26] Python. *Python Language Reference, Release 3.7.2*. Python Software Foundation, https://docs.python.org/3/reference/index.html, 2018. 1

[27] Dennis M. Ritchie. The development of the C language. *SIGPLAN Not.*, 28(3):201–208, March 1993. 19

[28] *Rust Programming Language*. Rust Project Developers, 2015. https://doc.rust-lang.org/reference.html. 2

[29] *Scala Language Specification, Version 2.11*. École Polytechnique Fédérale de Lausanne, 2016. http://www.scala-lang.org/files/archive/spec/2.11. 1

[30] Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley Series in Computer Science. Addison-Wesley, Boston, 1st edition, 1986. 1

[31] David W. Till. Tuples in imperative programming languages. Master's thesis, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 1989. 2, 42, 43, 45

[32] TIOBE Index. https://www.tiobe.com/tiobe-index. 1

[33] Ben Werther and Damian Conway. A modest proposal: C++ resyntaxed. *SIGPLAN Not.*, 31(11):74–82, November 1996. 2

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.