

1 CV (Cforall) User Manual
2 Version 1.0

3 “describe not prescribe”

4 CV 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

5 January 15, 2026

6 © 2016, 2018, 2021, 2024 CV Project

7

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

| | | |
|----|---|-----------|
| 1 | Contents | |
| 2 | 1 Introduction | 1 |
| 3 | 1.1 Background | 1 |
| 4 | 2 Why fix C? | 1 |
| 5 | 3 History | 2 |
| 6 | 4 Interoperability | 2 |
| 7 | 5 CV Compilation | 3 |
| 8 | 6 Backquote Identifiers | 5 |
| 9 | 7 Constant Underscores | 5 |
| 10 | 8 Exponentiation Operator | 6 |
| 11 | 9 Control Structures | 6 |
| 12 | 9.1 if / while Statement | 6 |
| 13 | 9.2 case Clause | 7 |
| 14 | 9.3 switch Statement | 7 |
| 15 | 9.4 Non-terminating and Labelled fallthrough | 10 |
| 16 | 9.5 Loop Control | 10 |
| 17 | 9.6 Labelled continue / break Statement | 12 |
| 18 | 9.7 Extended else | 15 |
| 19 | 9.8 with Statement | 15 |
| 20 | 10 Exception Handling | 17 |
| 21 | 10.1 Non-local Exception | 18 |
| 22 | 10.2 Exception Hierarchy | 18 |
| 23 | 11 Alternative Declarations | 19 |
| 24 | 12 Pointer / Reference | 20 |
| 25 | 12.1 Initialization | 23 |
| 26 | 12.2 Address-of Semantics | 24 |
| 27 | 12.3 Conversions | 25 |
| 28 | 13 string Type | 25 |
| 29 | 13.1 Implicit String Conversions | 26 |
| 30 | 13.2 Length | 26 |
| 31 | 13.3 Comparison Operators | 26 |
| 32 | 13.4 Concatenation | 26 |
| 33 | 13.5 Repetition | 27 |
| 34 | 13.6 Substring | 27 |
| 35 | 13.7 Searching | 28 |
| 36 | 13.8 Returning N on Search Failure | 29 |
| 37 | 13.9 C Compatibility | 29 |
| 38 | 13.10 I/O Operators | 30 |
| 39 | 14 Enumeration | 31 |
| 40 | 14.1 Enum type | 31 |
| 41 | 14.2 Inheritance | 32 |

| | | |
|----|--|-----------|
| 1 | 15 Routine Definition | 33 |
| 2 | 15.1 Named Return Values | 34 |
| 3 | 15.2 Routine Prototype | 35 |
| 4 | 15.3 Postfix Function | 35 |
| 5 | 16 Routine Pointers | 36 |
| 6 | 17 Default and Named Parameter | 36 |
| 7 | 17.1 Default | 36 |
| 8 | 17.2 Named (or Keyword) | 37 |
| 9 | 17.3 Mixed Default/Named | 38 |
| 10 | 18 Unnamed Structure Fields | 38 |
| 11 | 19 Nesting | 38 |
| 12 | 19.1 Type Nesting | 38 |
| 13 | 19.2 Routine Nesting | 39 |
| 14 | 20 Tuple | 40 |
| 15 | 20.1 Multiple-Return-Value Functions | 40 |
| 16 | 20.2 Expressions | 41 |
| 17 | 20.3 Variables | 41 |
| 18 | 20.4 Indexing | 42 |
| 19 | 20.5 Flattening and Structuring | 42 |
| 20 | 20.6 Assignment | 43 |
| 21 | 20.7 Construction | 44 |
| 22 | 20.8 Member-Access Expression | 45 |
| 23 | 20.9 Casting | 45 |
| 24 | 20.10 Polymorphism | 48 |
| 25 | 20.10.1 Assertion Inference | 49 |
| 26 | 21 Tuples | 49 |
| 27 | 21.1 Tuple Coercions | 50 |
| 28 | 21.2 Mass Assignment | 51 |
| 29 | 21.3 Multiple Assignment | 51 |
| 30 | 21.4 Cascade Assignment | 52 |
| 31 | 22 Stream I/O Library | 52 |
| 32 | 22.1 Basic I/O | 53 |
| 33 | 22.1.1 Stream Output | 53 |
| 34 | 22.1.2 Stream Input | 53 |
| 35 | 22.1.3 Stream Files | 54 |
| 36 | 22.2 Implicit Separator | 55 |
| 37 | 22.3 Separation Manipulators | 56 |
| 38 | 22.4 Newline Manipulators | 57 |
| 39 | 22.5 Output Manipulators | 58 |
| 40 | 22.6 Input Manipulators | 60 |
| 41 | 22.7 Concurrent Stream Access | 62 |
| 42 | 22.8 Locale | 63 |
| 43 | 23 String Stream | 63 |
| 44 | 24 Structures | 63 |
| 45 | 25 Constructors and Destructors | 65 |

| | | |
|----|--|-----------|
| 1 | 26 Overloading | 65 |
| 2 | 26.1 Constant | 66 |
| 3 | 26.2 Variable | 67 |
| 4 | 26.3 Function Overloading | 67 |
| 5 | 26.4 Operator | 68 |
| 6 | 27 Auto Type-Inferencing | 69 |
| 7 | 28 Concurrency | 69 |
| 8 | 28.1 Coroutine | 69 |
| 9 | 28.2 Monitors | 70 |
| 10 | 28.3 Threads | 70 |
| 11 | 29 Language Comparisons | 70 |
| 12 | 29.1 C++ | 70 |
| 13 | 29.2 Go | 71 |
| 14 | 29.3 Rust | 72 |
| 15 | 29.4 D | 73 |
| 16 | A Syntax Ambiguities | 73 |
| 17 | B C Incompatibles | 73 |
| 18 | C CV Keywords | 76 |
| 19 | D Standard Headers | 76 |
| 20 | E Standard Library | 77 |
| 21 | E.1 Dynamic Storage-Management | 77 |
| 22 | E.1.1 C Interface | 77 |
| 23 | E.1.2 CV Interface | 77 |
| 24 | E.2 Memory Set and Copy | 81 |
| 25 | E.3 String to Value Conversion | 82 |
| 26 | E.4 Search / Sort | 82 |
| 27 | E.5 Absolute Value | 83 |
| 28 | E.6 C Random Numbers | 83 |
| 29 | E.7 Algorithms | 83 |
| 30 | F Math Library | 83 |
| 31 | F.1 General | 83 |
| 32 | F.2 Exponential | 84 |
| 33 | F.3 Logarithm | 84 |
| 34 | F.4 Trigonometric | 85 |
| 35 | F.5 Hyperbolic | 86 |
| 36 | F.6 Error / Gamma | 87 |
| 37 | F.7 Nearest Integer | 87 |
| 38 | F.8 Manipulation | 89 |
| 39 | G Time Keeping | 89 |
| 40 | G.1 Duration | 89 |
| 41 | G.2 timeval | 91 |
| 42 | G.3 timespec | 91 |
| 43 | G.4 itimerval | 91 |
| 44 | G.5 Time | 91 |

| | | |
|---|---|------------|
| 1 | H Clock | 92 |
| 2 | H.1 C time | 92 |
| 3 | H.2 Clock | 92 |
| 4 | I Pseudo Random Number Generator | 93 |
| 5 | J Multi-precision Integers | 95 |
| 6 | K Rational Numbers | 98 |
| 7 | Index | 101 |

1 Introduction

`CV`¹ is a modern general-purpose concurrent programming-language, designed as an evolutionary step forward for the C programming language. The syntax of `CV` builds from C and should look immediately familiar to C/C++ programmers. `CV` 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, `CV` 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 `CV` is to “describe not prescribe”, which means `CV` tries to provide a pathway from low-level C programming to high-level `CV` programming, but it does not force programmers to “do the right thing”. Programmers can cautiously add `CV` 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 `CV`, for any reason, and in many cases, new `CV` features can be locally switched back to their C counterpart. There is no notion or requirement for *rewriting* a legacy C program to `CV`; instead, a programmer evolves a legacy program into `CV` by incrementally incorporating `CV` features. As well, new programs can be written in `CV` using a combination of C and `CV` features. In many ways, `CV` is to C as Scala [31] 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++ [32] 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, `CV` 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, `CV`, and C++ I/O mechanisms, where the programs output the same result.

| | C | CV | C++ |
|----|---|--|---|
| | <pre>#include <stdio.h></pre> | <pre>#include <fstream.hfa></pre> | <pre>#include <iostream></pre> |
| | | | <pre>using namespace std;</pre> |
| 24 | <pre>int main(void) { int x = 0, y = 1, z = 2; printf("%d %d %d\n", x, y, z); }</pre> | <pre>int main(void) { int x = 0, y = 1, z = 2; sout x y z; }</pre> | <pre>int main() { int x = 0, y = 1, z = 2; cout << x << ' ' << y << ' ' << z << endl; }</pre> |

While `CV` 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 [28]. In general, `CV` programs are 10% to 30% shorter than their equivalent C/C++ counterparts.

1.1 Background

This document is a programmer reference-manual for the `CV` 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 `CV` 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 [34] for February

¹Pronounced “C-for-all”, and written `CV`, `CFA`, or `Cforall`.

1 2023 ranks the top six most *popular* programming languages as C 17.4%, Java 12%, Python 12%, C++ 7.6%, C# 4%,
 2 Visual Basic 3.8% = 56.8%, where the next 50 languages are less than 2% each, with a long tail. The top 4 rankings
 3 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 | - | - |

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

9 As stated, the goal of the CV project is to engineer modern language-features into C in an evolutionary rather than
 10 revolutionary way. C++ [23, 12] is an example of a similar project; however, it largely extended the C language, and
 11 did not address many of C’s existing problems.² Fortran [14], Cobol [6], and Ada [1] are examples of programming
 12 languages that took an evolutionary approach, where modern language-features (*e.g.*, objects, concurrency) are added
 13 and problems fixed within the framework of the existing language. Java [18], Go [19], Rust [30] and D [3] are
 14 examples of the revolutionary approach for modernizing C/C++, resulting in a new language rather than an extension
 15 of the descendent. These languages have different syntax and semantics from C, do not interoperate directly with C,
 16 and are not systems languages because of restrictive memory-management or garbage collection. As a result, there
 17 is a significant learning curve to move to these languages, and C legacy-code must be rewritten. These costs can be
 18 prohibitive for many companies with a large software-base in C/C++, and a significant number of programmers require
 19 retraining in the new programming language.

20 The result of this project is a language that is largely backwards compatible with C11 [22], but fixes many of the
 21 well known C problems while adding modern language-features. To achieve these goals required a significant engi-
 22 neering exercise, *i.e.*, “thinking *inside* the C box”. Considering the large body of existing C code and programmers,
 23 there is significant impetus to ensure C is transformed into a modern language. While C11 made a few simple exten-
 24 sions to the language, nothing was added to address existing problems in the language or to augment the language
 25 with modern language-features. While some may argue that modern language-features may make C complex and
 26 inefficient, it is clear a language without modern capabilities is insufficient for the advanced programming problems
 27 existing today.

28 3 History

29 The CV project started with Dave Till’s K-W C [5, 33], which extended C with new declaration syntax, multiple return
 30 values from routines, and advanced assignment capabilities using the notion of tuples (see [35] for similar work in
 31 C++). The first CV implementation of these extensions was by Rodolfo Esteves [13].

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

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

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

40 4 Interoperability

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

²Two 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.

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

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

```

9   unsigned char abs( signed char );           // no C equivalent
10  extern "C" { int abs( int ); }             // C abs
11  unsigned long int abs( long int );         // C labs
12  unsigned long long int abs( long long int ); // C llabs
13  float abs( float );                       // C fabsf
14  double abs( double );                    // C fabs
15  long double abs( long double );          // C fabsl
16  float _Complex abs( float _Complex );    // C cabsf
17  double _Complex abs( double _Complex );  // C cabs
18  long double _Complex abs( long double _Complex ); // C cabsl

```

19 The problem is a name clash between the C name abs and the CV names abs, resulting in two name linkages:
 20 `extern "C"` and `extern "Cforall"` (default). Overloaded names must use *name mangling* to create unique names
 21 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 CV
 22 name, so it can be correctly formed. The only way around this problem is C's approach of creating unique names for
 23 each pairing of operation and type.

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

28 5 CV Compilation

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

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

```

35 cfa [ gcc/CV-options ] [ C/CV source-files ] [ assembler/loader files ]

```

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

38 CV has the following `gcc` flags turned on:

39 `-std=gnu11` The 2011 C standard plus GNU extensions.

40 CV has the following new options:

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

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

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

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

51 `-help` Information about the set of CV compilation flags is printed.

- 1 `-nohelp` Information about the set of CV compilation flags is not printed. **This option is the default.**
- 2 `-quiet` The CV compilation message is not printed at the beginning of a compilation.
- 3 `-noquiet` The CV 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 CV.³
- 6 `__CFA_MINOR__` is available during preprocessing and its value is the minor version number of CV.
- 7 `__CFA_PATCH__` is available during preprocessing and its value is the patch level number of CV.
- 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 CV extensions, use the following:

```

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

```

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

17 The CV transpiler has multiple internal steps. The following flags control how the CV transpiler works, the stages run, and printing within a stage. The majority of these flags are used by CV developers, but some are occasionally useful to programmers. Each option must be escaped with `-XCFA` to direct it to the CV compilation step, similar to the `-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

³The 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      ctorctor  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

```

10 6 Backquote Identifiers

11 CV 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 CV 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 CV header-
 19 library, so there is largely a seamless programming-experience.

20 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_L 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 y th 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 have an exponentiation operator often using operators `^` or `**`. However, neither of these operators work in C as `^` 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)`.

CV 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

CV 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
```

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

```

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

```

4 Unless a relational expression is specified, each variable is compared not equal to 0, which is the standard semantics
5 for the **if/while** expression, and the results are combined using the logical && operator. The scope of the declaration(s)
6 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
7 single declaration always compared != to 0.

8 9.2 case Clause

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

| | C | Cv | |
|----|--|-----------------------|----------------|
| | switch (i) { | switch (i) { | |
| 13 | case 1: case 3: case 5: | case 1, 3, 5: | // odd values |
| | ... | ... | |
| | case 2: case 4: case 6: | case 2, 4, 6: | // even values |
| | ... | ... | |
| | } | } | |

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

| | C | Cv | gcc | |
|----|---|-----------------------|-----------------------|-------------------|
| | switch (i) { | switch (i) { | switch (i) { | |
| 16 | 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 |
| | ... | ... | ... | |
| | } | } | } | |

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

19 Cv also allows lists of subranges.

```

20  case -5~-1, 12~21, 35~42:

```

21 9.3 switch Statement

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

- 23 1. By default, the end of a **case** clause⁵ *falls through* to the next **case** clause in the **switch** statement; to exit a
24 **switch** statement from a **case** clause requires explicitly terminating the clause with a transfer statement, most
25 commonly **break**:

```

26  switch ( i ) {
27  case 1:
28  ...
29  // fall-through
30  case 2:
31  ...
32  break; // exit switch statement
33  }

```

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

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

```

switch ( argc ) {
  case 3:
    // open output file
    // fall-through
  case 2:
    // open input file
    break; // exit switch statement
  default:
    // usage message
}

```

```

if ( argc == 3 ) {
    // open output file
    // open input file
} else if ( argc == 2 ) {
    // open input file (duplicate)
} else {
    // 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.

- 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

1 discovery. [10]

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

```

9     switch ( x ) {
10         int y = 1;           // unreachable initialization
11         x = 7;             // unreachable code without label/branch
12     case 0: ...
13         ...
14         int z = 0;         // unreachable initialization, cannot appear after case
15         z = 2;
16     case 1:
17         x = z;             // without fall through, z is uninitialized
18     }
```

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

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

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

32 These observations put into perspective the C99 changes to the **switch** statement.

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

```

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

38 still works. Nevertheless, reversing the default action would have a non-trivial effect on case actions that com-
 39 pound, such as the above example of processing command-line arguments. To preserve backwards compatibil-
 40 ity, a new kind of **switch** statement, called **choose** is introduced, with no implicit fall-through semantics and an
 41 explicit fall-through if the last statement of a case-clause ends with the new keyword **fallthrough**, *e.g.*:

```

42     choose ( i ) {
43         case 1: case 2: case 3:
44             ...
45             // implicit end of switch (break)
46         case 5:
47             ...
48             fallthrough;           // explicit fall through
49         case 7:
50             ...
51             break                 // explicit end of switch (redundant)
52         default:
53             j = 3;
54     }
```

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

2. Duff's device is eliminated from both **switch** and **choose** statements, and only invalidates a small amount of very questionable code. Hence, the **case** clause must appear at the same nesting level as the **switch/choose** body, as is done in most other programming languages with **switch** statements.
3. The issue of **default** at locations other than at the end of the cause clause can be solved by using good programming style, and there are a few reasonable situations involving fall-through where the **default** clause needs to appear is locations other than at the end. Therefore, no change is made for this issue.
4. Dealing with unreachable code in a **switch/choose** body is solved by restricting declarations and initialization to the start of statement body, which is executed *before* the transfer to the appropriate **case** clause⁶ and precluding statements before the first **case** clause. Further declarations at the same nesting level as the statement body are disallowed to ensure every transfer into the body is sound.

```

switch ( x ) {
  int i = 0;           // allowed only at start
  case 0:
  ...
  int j = 0;         // disallowed
  case 1:
  {
    int k = 0;       // allowed at different nesting levels
    ...
  case 2:           // disallow case in nested statements
  }
  ...
}

```

9.4 Non-terminating and Labelled fallthrough

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

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

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

9.5 Loop Control

CV condenses writing loops to facilitate coding speed and safety.

⁶These declarations are hoisted before the **switch/choose** statement and both declarations and statement are surrounded by a compound statement.

1 To simplify creating an infinite loop, the **for**, **while**, and **do** loop-predicate is extended with an empty conditional,
2 meaning a comparison value of 1 (true).

```
3 while ( ) // while ( true )
4 for ( ) // for ( ; true; )
5 do ... while ( ) // do ... while ( true )
```

6 Looping a predefined number of times, possibly with a loop index, occurs frequently. The **for** control, *i.e.*,
7 **for** (*/* control */*), is extended with a range and step. A range is a set of values defined by an optional low value
8 (default to 0), tilde, and high value, $L \sim H$, with an optional step $\sim S$ (default to 1), which means an ascending set of
9 values from L to H in positive steps of S.

```
10 0 ~ 5 // { 0, 1, 2, 3, 4, 5 }
11 -8 ~ -2 ~ 2 // { -8, -6, -4, -2 }
12 -3 ~ 3 ~ 1 // { -3, -2, -1, 0, 1, 2, 3 }
```

13 **Warning:** A range in descending order, *e.g.*, $5 \sim -3$ is the null (empty) set, *i.e.*, no values in the set. As well, a 0 or
14 negative step is undefined.

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

```
19 -8 ~ -2 // ascending, no prefix
20 0 + ~ 5 // ascending, prefix
21 -3 - ~ 3 // descending, prefix
```

22 For descending iteration, the L and H values are *implicitly* switched, and the increment/decrement for S is toggled.
23 Hence, the order of values in a set may not be the order the values are presented during looping. Changing the
24 iteration direction is faster and safer because the direction prefix can be added/removed without changing existing
25 (correct) range information. **Warning:** reversing the range endpoints for descending order results in an empty set.

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

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

30 **for** control is formalized by the following regular expression:

```
31 [ L ] [ + | - ] ~ [ = ] H [ ~ S ]
```

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

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

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

40 Hence, unlike the 3 components in the C **for**-control, there are only two components in the C \forall **for**-control: the optional
41 index variable and the range. The index type is optional (like C++ **auto**), where the type is normally inferred from the
42 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
43 required high value H is used, as both L and H are the same type in this case.

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

46 The following examples illustrate common C \forall **for**-control combinations, with the C counter-part in the comment.

- H is implicit ascending exclusive range [0,H).

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

- $\sim = H$ is implicit ascending inclusive range [0,H].

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

- $L \sim H$ is explicit ascending exclusive range [L,H).

```

1      for ( 1 ~ 5 )                // for ( typeof(1) i = 1; i < 5; i += 1 )
2      • L ~ H is explicit ascending inclusive range [L,H].
3      for ( 1 ~ = 5 )              // for ( typeof(1) i = 1; i <= 5; i += 1 )
4      • L ~ H is explicit descending exclusive range (H,L), where L and H are implicitly
5      descending.
6      for ( 1 ~ ~ 5 )              // for ( typeof(1) i = 5; i > 0; i -= 1 )
7      • L ~ ~ H is explicit descending inclusive range [H,L], where L and H are implicitly
8      range descending.
9      for ( 1 ~ ~ = 5 )            // for ( typeof(1) i = 5; i >= 0; i -= 1 )

```

There are situations when the **for**-control actions need to be moved into the loop body, *e.g.*, a mid-loop exit does not need an iteration-completion test in the **for** control. The character '@' indicates that a specific **for**-control action is ignored, *i.e.*, generates no code.

```

13     for ( i; @ ~ ~ 10 )           // for ( typeof(10) i = 10; /*empty*/; i -= 1 )
14     for ( i; 1 ~ @ ~ 2 )          // for ( typeof(1) i = 1; /* empty */; i += 2 )
15     for ( i; 1 ~ 10 ~ @ )         // for ( typeof(1) i = 1; i < 10; /* empty */ )
16     for ( i; 1 ~ @ ~ @ )          // for ( typeof(1) i = 1; /* empty */; /* empty */ )

```

Warning: L *cannot* be elided for the ascending range, @ ~ 5, nor H for the descending range, 1 ~ ~ @, as the loop index is uninitialized. **Warning:** H *cannot* be elided in an anonymous loop index, 1 ~ @, as there is no index to stop the loop.

There are situations when multiple loop indexes are required. The character ':' means add another index, where any number of indices may be chained in a single **for** control.

```

21     for ( i; 5 : j; 2 ~ 12 ~ 3 )   // for ( typeof(i) i = 1, j = 2; i < 5 && j < 12; i += 1, j += 3 )
22     for ( i; 5 : j; 2 ~ @ ~ 3 )     // for ( typeof(i) i = 1, j = 2; i < 5; i += 1, j += 3 )
23     for ( i; 5 : j; 2.5 ~ @ ~ 3.5 ) // no C equivalent, without hoisting declaration of floating-point j

```

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

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

```

26     forall( T ) trait Iterate {
27         void ?{}( T & t, zero_t );
28         int ?<?( T t1, T t2 );
29         int ?<=?( T t1, T t2 );
30         int ?>?( T t1, T t2 );
31         int ?>=?( T t1, T t2 );
32         T ?+=?( T & t1, T t2 );
33         T ?+=?( T & t, one_t );
34         T ?-=?( T & t1, T t2 );
35         T ?-=?( T & t, one_t );
36     }

```

Figure 3, p. 14 shows an example of a structure using **for** control. Note, the use of (S){0} when implicitly setting the loop-index type, because using 0 incorrect declares the index to **int** rather than S.

39 9.6 Labelled **continue** / **break** Statement

40 C **continue** and **break** statements are restricted to one level of nesting for a particular control structure. This restriction
41 forces programmers to use **goto** to achieve the equivalent control-flow for more than one level of nesting. To prevent
42 having to switch to the **goto**, C \forall extends the **continue** and **break** with a target label to support static multi-level
43 exit [4], as in Java. For both **continue** and **break**, the target label must be directly associated with a **for**, **while** or **do**
44 statement; for **break**, the target label can also be associated with a **switch**, **if** or compound ({}) statement. Figure 4,
45 p. 14 shows a comparison between labelled **continue** and **break** and the corresponding C equivalent using **goto** and
46 labels. The innermost loop has 8 exit points, which cause continuation or termination of one or more of the 7 nested
47 control-structures.

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

- 49 • They cannot create a loop, which means only the looping constructs cause looping. This restriction means all
50 situations resulting in repeated execution are clearly delineated.


```

struct S { int i, j; };
void ?{ ( S & s, int i = 0, int j = 0 ) { s.[i, j] = [i, j]; }
void ?{ ( S & s, zero_t ) { s.[i, j] = 0; }
int ?<?( S t1, S t2 ) { return t1.i < t2.i && t1.j < t2.j; }
int ?<=?( S t1, S t2 ) { return t1.i <= t2.i && t1.j <= t2.j; }
int ?>?( S t1, S t2 ) { return t1.i > t2.i && t1.j > t2.j; }
int ?>=?( S t1, S t2 ) { return t1.i >= t2.i && t1.j >= t2.j; }
S ?+=?( S & t1, S t2 ) { t1.i += t2.i; t1.j += t2.j; return t1; }
S ?+=?( S & t, one_t ) { t.i += 1; t.j += 1; return t; }
S ?-=?( S & t1, S t2 ) { t1.i -= t2.i; t1.j -= t2.j; return t1; }
S ?-=?( S & t, one_t ) { t.i -= 1; t.j -= 1; return t; }
ofstream & ?{ ( ofstream & os, S s ) {
    return os | " (" | s.i | s.j | " ) ";
}
void & ?{ ( ofstream & os, S s ) {
    ( ofstream & )( os | s ); ends ( os );
}

```

```

int main() {
    sout | nloff;
    for ( S i = 0; i < (S){10,10}; i += 1 ) { sout | i; } sout | "A" | nl; // C
    for ( S i; 0 ~ (S){10,10} ) { sout | i; } sout | "B" | nl; // CFA
    for ( i; (S){10,10} ) { sout | i; } sout | "C" | nl;
    for ( i; (S){0} ~ (S){10,10} ) { sout | i; } sout | "D" | nl;
    for ( i; (S){0} ~ (S){10,10} ) { sout | i; } sout | "E" | nl;
    for ( i; (S){0} ~ (S){10,10} ~ (S){2} ) { sout | i; } sout | "F" | nl;
    for ( i; (S){0} ~ (S){10,10} ) { sout | i; } sout | "G" | nl;
    for ( i; (S){0} ~ (S){10,10} ) { sout | i; } sout | "H" | nl;
    for ( i; (S){0} ~ (S){10,10} ~ (S){2,1} ) { sout | i; } sout | "I" | nl;
}
(0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) A
(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
(0 0) (1 1) (2 2) (3 3) (4 4) (5 5) (6 6) (7 7) (8 8) (9 9) D
(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

| | |
|--|--|
| <pre> { ForC: for (...) { WhileC: while (...) { DoC: do { if (...) { switch (...) { case 3: goto Compound; goto Try; goto ForB; /* or */ goto ForC; goto WhileB; /* or */ goto WhileC; goto DoB; /* or */ goto DoC; goto If; goto Switch; } Switch; } else { ... goto If; ... // terminate if } if; } while (...); DoB; } WhileB; } ForB; } Compound; </pre> <p style="text-align: center;">a) C</p> | <pre> Compound: { Try: try { For: for (...) { While: while (...) { Do: do { if: if (...) { Switch: switch (...) { case 3: break Compound; break Try; break For; /* or */ continue For; break While; /* or */ continue While; break Do; /* or */ continue Do; break If; break Switch; } // switch } else { ... break If; ... // terminate if } // if } while (...); // do } // while } // for } finally { // always executed } // try } // compound </pre> <p style="text-align: center;">b) CV</p> |
|--|--|

Figure 4: Multi-level Exit

1 • They cannot branch into a control structure. This restriction prevents missing declarations and/or initializations
 2 at the start of a control structure resulting in undefined behaviour.
 3 The advantage of the labelled **continue/break** is allowing static multi-level exits without having to use the **goto** state-
 4 ment, and tying control flow to the target control structure rather than an arbitrary point in a program via a label.
 5 Furthermore, the location of the label at the *beginning* of the target control structure informs the reader (eye candy)
 6 that complex control-flow is occurring in the body of the control structure. With **goto**, the label is at the end of the
 7 control structure, which fails to convey this important clue early enough to the reader. Finally, using an explicit target
 8 for the transfer, instead of an implicit target, allows new constructs to be added or removed without affecting existing
 9 constructs. Otherwise, the implicit targets of the current **continue** and **break**, *i.e.*, the closest enclosing loop or **switch**,
 10 change as certain constructs are added or removed.

11 9.7 Extended else

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

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

```

17 int a[10]

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

```

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

19 9.8 with Statement

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

```

22 struct Person {                       // aggregate
23     struct Name {                     // nesting
24         char first[20], last[20];
25     } name;
26     struct Address {                 // nesting
27         ...
28     } address;
29     int sex;
30 };

```

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

```

32 Person p
33 p.name ...; p.address ...; p.sex ...; // access containing fields

```

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

```

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

```

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

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

```

43 struct S {
44     struct /* unnamed */ { int g, h; } __attribute__(( aligned(64) ));

```

```

1   int tag;
2   union /* unnamed */ {
3       struct { char c1, c2; } __attribute__(( aligned(128) ));
4       struct { int i1, i2; };
5       struct { double d1, d2; };
6   };
7   } s;
8   enum { R, G, B };
9   s.g; s.h; s.tag = R; s.c1; s.c2; s.i1 = G; s.i2 = B; s.d1; s.d2;
10  Object-oriented languages reduce qualification for class variables within member functions, e.g., C++:
11  struct S {
12      char c; int i; double d;
13      void f( /* S * this */ ) {           // implicit "this" parameter
14          c; i; d;                         // this->c; this->i; this->d;
15      }
16  }

```

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

```

19  struct T {
20      char m; int i; double n;           // derived class variables
21  };
22  struct S : public T {
23      char c; int i; double d;           // class variables
24      void g( double d, T & t ) {
25          d; t.m; t.i; t.n;             // function parameter
26          c; i; this->d; S::d;           // class S variables
27          m; T::i; n;                   // class T variables
28      }
29  };

```

Note the three different forms of qualification syntax in C++, `.`, `->`, `::`, which is confusing.

Since CV is not object-oriented, it has no implicit parameter with its implicit qualification. Instead CV introduces a general mechanism using the **with** statement (see Pascal [25, § 4.F]) to explicitly elide aggregate qualification by opening a scope containing the field identifiers. Hence, the qualified fields become variables with the side-effect that it is simpler to write, easier to read, and optimize field references in a block.

```

35  void f( S & this ) with ( this ) {     // with statement
36      c; i; d;                           // this.c, this.i, this.d
37  }

```

with the generality of opening multiple aggregate-parameters:

```

39  void g( S & s, T & t ) with ( s, t ) { // multiple aggregate parameters
40      c; s.i; d;                           // s.c, s.i, s.d
41      m; t.i; n;                           // t.m, t.i, t.n
42  }

```

where qualification is only necessary to disambiguate the shadowed variable `i`. In detail, the **with** statement may form a function body or be nested within a function body.

The **with** clause takes a list of expressions, where each expression provides an aggregate type and object. (Enumerations are already opened.) To open a pointer type, the pointer must be dereferenced to obtain a reference to the aggregate type.

```

48  S * sp;
49  with ( *sp ) { ... }

```

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

CV'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:

```

1  struct Q { int i; int k; int m; } q, w;
2  struct R { int i; int j; double m; } r, w;
3  with ( r, q ) {
4      j + k;           // unambiguous, r.j + q.k
5      m = 5.0;        // unambiguous, q.m = 5.0
6      m = 1;          // unambiguous, r.m = 1
7      int a = m;      // unambiguous, a = r.i
8      double b = m;   // unambiguous, b = q.m
9      int c = r.i + q.i; // disambiguate with qualification
10     (double)m;      // disambiguate with cast
11 }

```

12 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.

```

14  with ( r ) {
15      i;               // unambiguous, r.i
16      with ( q ) {
17          i;           // unambiguous, q.i
18      }
19  }

```

20 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:

```

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

```

24 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`.

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

```

27  void f( S & s, char c ) with ( s ) {
28      s.c = c; i = 3; d = 5.5; // initialize fields
29  }

```

30 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:

```

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

```

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

```

37  void f( S & s, char c ) with ( s ) with( params ) { // syntax disallowed, illustration only
38      s.c = c; i = 3; d = 5.5;
39  }

```

40 This implicit semantic matches with programmer expectation.

41 10 Exception Handling

42 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.

45 Currently, C_V uses macros `ExceptionDecl` and `ExceptionInst` to declare and instantiate an exception.

```

46  #include <Exception.hfa>
47  ExceptionDecl( E, // must be global scope
48      ... // exception fields
49  );
50  try {
51      ...
52      if ( ... ) throwResume ExceptionInst( E, /* initialization */ );
53      if ( ... ) throw ExceptionInst( E, /* initialization */ );

```

```

1  ...
2  } catchResume( E * ) { // must be pointer
3  ...
4  } catch( E * ) {
5  ...
6  }
7  exception_t E {}; // exception type
8  void f(...) {
9  ... throw E{}; ... // termination
10 ... throwResume E{}; ... // resumption
11 }
12 try {
13 f(...);
14 } catch( E e ; boolean-predicate ) { // termination handler
15 // recover and continue
16 } catchResume( E e ; boolean-predicate ) { // resumption handler
17 // repair and return
18 } finally {
19 // always executed
20 }

```

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

24 10.1 Non-local Exception

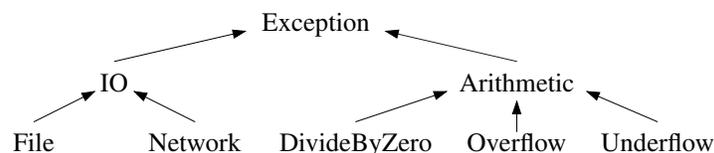
```

25 void main() {
26     try {
27         _Enable {
28             ... resume(); ...
29         }
30     } catchResume( E & ) { // should be reference
31     ...
32     } catch( E & ) {
33     ...
34     }
35 }

```

36 10.2 Exception Hierarchy

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



41 A programmer can then choose to handle an exception at different degrees of specificity along the hierarchy; derived
 42 exception-types support a more flexible programming style. For example, higher-level code should catch general
 43 exceptions to reduce coupling to the specific implementation at the lower levels; unnecessary coupling may force
 44 changes in higher-level code when low-level code changes. A consequence of derived exception-types is that multiple
 45 exceptions may match, *e.g.*:

```
46 catch( Arithmetic )
```

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

```

1  try {
2  ...
3  } catch( Overflow ) { // must appear first
4  // handle overflow
5  } catch( Arithmetic )
6  // handle other arithmetic issues
7  }
8  Multiple derivation among exception is not supported.

```

9 11 Alternative Declarations

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

```

12  int * x[5]      x [ ] [ ] [ ] [ ] [ ]      x [ ] → 0 1 2 3 4
                   ↓ ↓ ↓ ↓ ↓
                   0 1 2 3 4

```

13 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
 14 productivity and safety issues even for basic programs. Another example of confusion results from the fact that a
 15 routine name and its parameters are embedded within the return type, mimicking the way the return value is used at
 16 the routine's call site. For example, a routine returning a pointer to an array of integers is defined and used in the
 17 following way:

```

18  int (*f())[5] {...};           // definition
19  ... (*f())[3] += 1;           // usage

```

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

23 In spite of its difficulties, I believe that the C's approach to declarations remains plausible, and am com-
 24 fortable with it; it is a useful unifying principle. [29, p. 12]

25 CV provides its own type, variable and routine declarations, using a different syntax. The new declarations place
 26 qualifiers to the left of the base type, while C declarations place qualifiers to the right of the base type. In the following
 27 example, **red** is the base type and **blue** is qualifiers. The CV declarations move the qualifiers to the left of the base
 28 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
 29 to specify a variable's type.

| | C | CV |
|----|------------------------------------|------------------------------------|
| 30 | int * x1 [5]; | [5] * int x1; |
| | int (*x2)[5]; | * [5] int x2; |
| | int (*f(int p))[5]; | * [5] int f(int p); |

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

| | C | CV |
|----|--------------------|--------------------|
| 34 | int *x, *y; | * int x, y; |

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

| | C | CV |
|----|-------------------|----------------------------------|
| 36 | int *x, y; | * int x; int y; |

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

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

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

| C | CV | |
|---|---|--|
| <code>int const * const x;</code> | <code>const * const int x;</code> | <i>// const pointer to const integer</i> |
| <code>const int (* const y)[5]</code> | <code>const * [5] const int y;</code> | <i>// const pointer to array of 5 const integers</i> |

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

| C | CV | |
|------------------------------------|------------------------------------|--|
| <code>int extern x[5];</code> | <code>extern [5] int x;</code> | <i>// externally visible array of 5 integers</i> |
| <code>const int static * y;</code> | <code>static * const int y;</code> | <i>// internally visible pointer to constant int</i> |

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

| C | CV |
|---------------------------------------|---------------------------------------|
| <code>y = (int *)x;</code> | <code>y = (* int)x;</code> |
| <code>i = sizeof(int * [5]);</code> | <code>i = sizeof([5] * int);</code> |

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

12 Pointer / Reference

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

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

| | | | | |
|---------------------|-----|---|-----|--|
| <code>int x;</code> | x | 3 | int | <code>int * const x = (int *)100</code> |
| <code>x = 3;</code> | 100 | | | <code>*x = 3;</code> <i>// implicit dereference</i> |
| <code>int y;</code> | y | 3 | int | <code>int * const y = (int *)104;</code> |
| <code>y = x;</code> | 104 | | | <code>*y = *x;</code> <i>// implicit dereference</i> |

where the right example is how the compiler logically interprets the variables in the left example. Since a variable name only points to one address during its lifetime, it is an immutable pointer; hence, the implicit type of pointer variables `x` and `y` are constant pointers in the compiler interpretation. In general, variable addresses are stored 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. [22, § 6.11.5(1)]

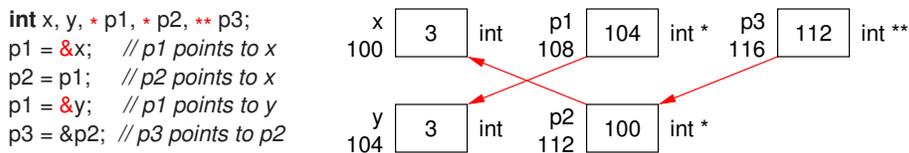
⁸One way to conceptualize the null pointer is that no variable is placed at this address, so the null-pointer address can be used to denote an uninitialized pointer/reference object; *i.e.*, the null pointer is guaranteed to compare unequal to a pointer to any object or routine. In general, a value with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

1 instructions instead of loaded from memory, and hence may not occupy storage. These approaches are contrasted in
2 the following:

| | explicit variable address | | implicit variable address |
|---|---------------------------------|----|-----------------------------|
| 3 | lda r1,100 // load address of x | | |
| | ld r2,(r1) // load value of x | ld | r2,(100) // load value of x |
| | lda r3,104 // load address of y | | |
| | st r2,(r3) // store x into y | st | r2,(104) // store x into y |

4 Finally, the immutable nature of a variable's address and the fact that there is no storage for the variable pointer means
5 pointer assignment is impossible. Therefore, the expression $x = y$ has only one meaning, $*x = *y$, *i.e.*, manipulate values,
6 which is why explicitly writing the dereferences is unnecessary even though it occurs implicitly as part of instruction
7 decoding.

8 A pointer/reference object is a generalization of an object variable-name, *i.e.*, a mutable address that can point to
9 more than one memory location during its lifetime. (Similarly, an integer variable can contain multiple integer literals
10 during its lifetime versus an integer constant representing a single literal during its lifetime, and like a variable name,
11 may not occupy storage if the literal is embedded directly into instructions.) Hence, a pointer occupies memory to
12 store its current address, and the pointer's value is loaded by dereferencing, *e.g.*:



14 Notice, an address has a duality: a location in memory or the value at that location. In many cases, a compiler
15 might be able to infer the best meaning for these two cases. For example, Algol68 [26] infers pointer dereferencing to
16 select the best meaning for each pointer usage

```
17 p2 = p1 + x; // compiler infers *p2 = *p1 + x;
```

18 Algol68 infers the following dereferencing $*p2 = *p1 + x$, because adding the arbitrary integer value in x to the address
19 of $p1$ and storing the resulting address into $p2$ is an unlikely operation. Unfortunately, automatic dereferencing does
20 not work in all cases, and so some mechanism is necessary to fix incorrect choices.

21 Rather than inferring dereference, most programming languages pick one implicit dereferencing semantics, and the
22 programmer explicitly indicates the other to resolve address-duality. In C, objects of pointer type always manipulate
23 the pointer object's address:

```
24 p1 = p2; // p1 = p2 rather than *p1 = *p2
25 p2 = p1 + x; // p2 = p1 + x rather than *p2 = *p1 + x
```

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

```
27 p1 = p2; // pointer address assignment
28 *p2 = *p1 + x; // pointed-to value assignment / operation
```

29 The C semantics work well for situations where manipulation of addresses is the primary meaning and data is rarely
30 accessed, such as storage management (malloc/free).

31 However, in most other situations, the pointed-to value is requested more often than the pointer address.

```
32 *p2 = ((*p1 + *p2) * (**p3 - *p1)) / (**p3 - 15);
```

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

```
35 p2 = ((p1 + p2) * (p3 - p1)) / (p3 - 15);
```

36 To support this common case, a reference type is introduced in CV, denoted by $\&$, which is the opposite dereference
37 semantics to a pointer type, making the value at the pointed-to location the implicit semantics for dereferencing (similar
38 but not the same as C++ reference types).

```
39 int x, y, & r1, & r2, && r3;
40 &r1 = &x; // r1 points to x
41 &r2 = &r1; // r2 points to x
42 &r1 = &y; // r1 points to y
43 &&r3 = &&r2; // r3 points to r2
```

```
1 r2 = ((r1 + r2) * (r3 - r1)) / (r3 - 15); // implicit dereferencing
```

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

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

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

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

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

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

13 Cancellation works to arbitrary depth.

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

```
15 int x, *p1 = &x, **p2 = &p1, ***p3 = &p2,
16     &r1 = x, &&r2 = r1, &&&r3 = r2;
17 ***p3 = 3; // change x
18 r3 = 3; // change x, ***r3
19 **p3 = ...; // change p1
20 &r3 = ...; // change r1, (&*)**r3, 1 cancellation
21 *p3 = ...; // change p2
22 &&r3 = ...; // change r2, (&(&*)*)r3, 2 cancellations
23 &&&r3 = p3; // change r3 to p3, (&(&(&*)*)r3, 3 cancellations
```

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

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

```
28 const int cx = 5; // cannot change cx;
29 const int & cr = cx; // cannot change what cr points to
30 &cr = &cx; // can change cr
31 cr = 7; // error, cannot change cx
32 int & const rc = x; // must be initialized
33 &rc = &x; // error, cannot change rc
34 const int & const crc = cx; // must be initialized
35 crc = 7; // error, cannot change cx
36 &crc = &cx; // error, cannot change crc
```

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

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

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

```
41 int & const cr = *malloc();
42 cr = 5;
43 free( &cr );
44 cr = 7; // unsound pointer dereference
```

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

| C | CV |
|--|--|
| 48 <code>const int * const * const ccp;</code> | <code>const * const * const int ccp;</code> <code>const & const & const int ccr;</code> |

⁹The 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. [22, § 6.5.3.2–3]

1 where the `CV` declaration is read left-to-right.

2 Finally, like pointers, references are usable and composable with other type operators and generators.

```
3 int w, x, y, z, & ar[3] = { x, y, z }; // initialize array of references
4 &ar[1] = &w; // change reference array element
5 typedef( ar[1] ) p; // (gcc) is int, i.e., the type of referenced object
6 typedef( &ar[1] ) q; // (gcc) is int &, i.e., the type of reference
7 sizeof( ar[1] ) == sizeof( int ); // is true, i.e., the size of referenced object
8 sizeof( &ar[1] ) == sizeof( int * ); // is true, i.e., the size of a reference
```

9 In contrast to `CV` reference types, `C++`'s reference types are all **const** references, preventing changes to the reference address, so only value assignment is possible, which eliminates half of the address duality. Also, `C++` does not allow arrays of reference¹⁰ Java's reference types to objects (all Java objects are on the heap) are like `C` pointers, which always manipulate the address, and there is no (bit-wise) object assignment, so objects are explicitly cloned by shallow or deep copying, which eliminates half of the address duality.

14 12.1 Initialization

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

```
21 int * p = &x; // assign address of x
22 int * p = x; // assign value of x
23 int & r = x; // must have address of x
```

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

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

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

```
37 int temp1 = f( x ), temp2 = f( y );
38 z = temp1 + temp2;
```

39 This implicit referencing is crucial for reducing the syntactic burden for programmers when using references; otherwise references have the same syntactic burden as pointers in these contexts.

41 When a pointer/reference parameter has a **const** value (immutable), it is possible to pass literals and expressions.

```
42 void f( const int & cr );
43 void g( const int * cp );
44 f( 3 ); g( &3 );
45 f( x + y ); g( &(x + y) );
```

46 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 `CV` extension necessary to type match and is a common requirement before a variable in `C` (e.g., `scanf`). Importantly,

¹⁰The 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 &3 may not be equal to &3, where the references occur across calls because the temporaries maybe different on each
2 call.

3 CV extends this semantics to a mutable pointer/reference parameter, and the compiler implicitly creates the nec-
4 essary temporary (copying the argument), which is subsequently pointed-to by the reference parameter and can be
5 changed.¹¹

```
6 void f( int & r );
7 void g( int * p );
8 f( 3 );      g( &3 );      // compiler implicit generates temporaries
9 f( x + y );  g( &(x + y) ); // compiler implicit generates temporaries
```

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

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

```
14 void f( int i );
15 void (* fp)( int );      // routine pointer
16 fp = f;                 // reference initialization
17 fp = &f;                // pointer initialization
18 fp = *f;                // reference initialization
19 fp(3);                  // reference invocation
20 (*fp)(3);               // pointer invocation
```

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

```
24 const void (& fr)( int ) = f;      // routine reference
25 fr = ...;                          // error, cannot change code
26 &fr = ...;                          // changing routine reference
27 fr( 3 );                            // reference call to f
28 (*fr)(3);                          // error, incorrect type
```

29 because the value of the routine object is a routine literal, *i.e.*, the routine code is normally immutable during execu-
30 tion.¹³ CV allows this additional use of references for routine objects in an attempt to give a more consistent meaning
31 for them.

32 12.2 Address-of Semantics

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

- 34 • if R is an rvalue of type T &₁ ··· &_r, where $r \geq 1$ references (& symbols), than &R has type T *&₂ ··· &_r, *i.e.*, T
35 pointer with $r - 1$ references (& symbols).
- 36 • if L is an lvalue of type T &₁ ··· &_l, where $l \geq 0$ references (& symbols), than &L has type T *&₁ ··· &_l, *i.e.*, T
37 pointer with l references (& symbols).

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

```
39 int x, * px, ** ppx, *** pppx, **** ppppx;
40 int & rx = x, && rrx = rx, &&& rrrx = rrx ;
41 x = rrrx;      // rrrx is an lvalue with type int &&& (equivalent to x)
42 px = &rrrx;   // starting from rrrx, &rrrx is an rvalue with type int *&&& (&x)
43 ppx = &&rrrx; // starting from &rrrx, &&rrrx is an rvalue with type int **&& (&rx)
44 pppx = &&&rrrx; // starting from &&rrrx, &&&rrrx is an rvalue with type int ***& (&rx)
45 ppppx = &&&&rrrx; // starting from &&&rrrx, &&&&rrrx is an rvalue with type int **** (&rrrx)
```

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

```
47 int x, * px, ** ppx, *** pppx;
48 int & rx = x, && rrx = rx, &&& rrrx = rrx ;
```

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

¹²This 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.

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

```

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

```

12.3 Conversions

C provides a basic implicit conversion to simplify variable usage:

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

```

8   int x;
9   x + 1;           // lvalue variable (int) converts to rvalue for expression

```

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

CV provides three new implicit conversion for reference types to simplify reference usage.

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

```

13  int x, &r = x, f( int p );
14  x = r + f( r );   // lvalue reference converts to rvalue

```

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

2. lvalue to reference conversion: lvalue-type cv1 T converts to cv2 T &, which allows implicitly converting variables to references.

```

18  int x, &r = x, f( int & p ); // lvalue variable (int) convert to reference (int &)
19  f( x );                     // lvalue variable (int) convert to reference (int &)

```

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

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

```

25  int x, &f( int & p );
26  f( x + 3 );         // rvalue parameter (int) implicitly converts to lvalue temporary reference (int &)
27  &f(...) = &x;      // rvalue result (int &) implicitly converts to lvalue temporary reference (int &)

```

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

13 string Type

A string is a sequence of symbols, where the form of a symbol can vary significantly: regular 7/8-bit ASCII/Latin-1, or wide 2/4/8-byte UNICODE or variable length UTF-8/16/32. A C character string is zero or more regular, wide, or escape characters enclosed in double-quotes "xyz\n". Currently, CV strings only support regular characters.

A string type is designed to operate on groups of characters for assigning, copying, scanning, and updating, rather than working with individual characters. The CV string type is for manipulation of dynamically-sized strings versus C **char *** type for manipulation of statically-sized null-terminated strings. Therefore, the amount of storage for a CV string changes dynamically at runtime to fit the string size, whereas the amount of storage for a C string is fixed at compile time. As a result, a string declaration does not specify a maximum length, where a C string array does. For CV, as a string dynamically grows and shrinks in size, so does its underlying storage. For C, as a string dynamically grows and shrinks in size, but its underlying storage does not. The maximum storage for a CV string value is `size_t` characters, which is 2^{32} or 2^{64} respectively. A CV string manages its length separately from the string, so there is no null (' \0 ') terminating value at the end of a string value. Hence, a CV string cannot be passed to a C string manipulation function, such as `strcat`. Like C strings, characters in a string are numbered from the left starting at 0 (because subscripting is zero-origin), and in CV numbered from the right starting at -1.

| | | | | | |
|----|----|----|----|----|---------------------|
| "a | b | c | d | e" | |
| 0 | 1 | 2 | 3 | 4 | left to right index |
| -5 | -4 | -3 | -2 | -1 | right to left index |

The include file `string.hfa` is necessary to access type string.

1 13.1 Implicit String Conversions

2 The ability to convert from internal (machine) to external (human) format is useful in situations other than I/O. Hence,
3 the basic types **char**, **char ***, **int**, **double**, **_Complex**, including any signness and size variations, implicitly convert to
4 type string (as in Java).

| | | | | |
|---|---|--|--|--|
| 5 | <pre>string s = 5; s = 'x'; s = "abc"; s = 42hh; /* signed char */ s = 42h; /* short int */ s = 0xff;</pre> | <pre>"5" "x" "abc" "42" "42" "255"</pre> | <pre>s = (ssize_t)MIN; s = (size_t)MAX; s = 5.5; s = 5.5L; s = 5.5+3.4i; s = 5.5L+3.4Li;</pre> | <pre>"-9223372036854775808" "18446744073709551615" "5.5" "5.5" "5.5+3.4i" "5.5+3.4i"</pre> |
|---|---|--|--|--|

6 Conversions can be explicitly specified using a compound literal.

```
7 s = (string){ 5 }; s = (string){ "abc" }; s = (string){ 5.5 };
```

8 Conversions from string to **char *** attempt to be safe. The `strncpy` conversion requires the maximum length for
9 the pointer's target buffer. The assignment operator and constructor both allocate the buffer and return its address,
10 meaning the programmer must free it. Note, a C string is always null terminated, implying storage is always necessary
11 for the null.

| | | |
|----|---|--|
| 12 | <pre>string s = "abcde"; char cs[4]; strncpy(cs, s, sizeof(cs)); char * cp = s; // ownership delete(cp); cp = s + ' ' + s; // ownership delete(cp);</pre> | <pre>"abc\0", in place "abcde\0", malloc "abcde abcde\0", malloc</pre> |
|----|---|--|

13 13.2 Length

14 The `len` operation (short for `strlen`) returns the length of a C or C \forall string. For compatibility, `strlen` works with C \forall
15 strings.

| | | |
|----|---|------------------------|
| 16 | <pre>i = len(""); i = len("abc"); i = len(cs); i = strlen(cs); i = len(name); i = strlen(name);</pre> | <pre>0 3 3 3 4 4</pre> |
|----|---|------------------------|

17 13.3 Comparison Operators

18 The binary relational, `<`, `<=`, `>`, `>=`, and equality, `==`, `!=`, operators compare C \forall strings using lexicographical ordering,
19 where longer strings are greater than shorter strings. In C, these operators compare the C string pointer not its value,
20 which does not match programmer expectation. C strings use function `strcmp` to lexicographically compare the string
21 value. Java has the same issue with `==` and `.equals`.

22 13.4 Concatenation

23 The binary operators `+` and `+=` concatenate C **char**, **char *** and C \forall strings, creating the sum of the characters.

| | | | | | | |
|----|--|-------------------------------------|--|--|--|--|
| 24 | <pre>s = ""; s = 'a' + 'b'; s = 'a' + "b"; s = "a" + 'b'; s = "a" + "b";</pre> | <pre>"ab" "ab" "ab" "ab" "ab"</pre> | <pre>s = ""; s = 'a' + 'b' + s; s = 'a' + 'b' + s; s = 'a' + "b" + s; s = "a" + 'b' + s;</pre> | <pre>"ab" "abab" "ababab" "abababab"</pre> | <pre>s = ""; s = s + 'a' + 'b'; s = s + 'a' + "b"; s = s + "a" + 'b'; s = s + "a" + "b";</pre> | <pre>"ab" "abab" "ababab" "abababab"</pre> |
|----|--|-------------------------------------|--|--|--|--|

25 However, including `<string.h>` can result in ambiguous uses of the overloaded `+` operator.¹⁴ For example, subtracting
26 characters or pointers has valid use-cases:

¹⁴Combining multiple packages in any programming language can result in name clashes or ambiguities.

| | length | | pattern | |
|---|------------------------|---|-------------------|-----------------------|
| | string name = "PETER"; | | | |
| | s = name(0, 4); | "PETE" | s = name("ET"); | "ET" |
| | s = name(1, 4); | "ETER" | s = name("WW"); | " " // does not occur |
| 1 | s = name(2, 4); | "TER" // clip length to 3 | | |
| | s = name(4, -2); | "ER" | | |
| | s = name(8, 2); | " " // beyond string to right, clip to null | | |
| | s = name(0, -2); | " " // beyond string to left, clip to null | | |
| | s = name(-1, -2); | "ER" | | |
| | s = name(-3); | "TER" // to end of string | | |

2 For the length form, a negative starting position is a specification from the right end of the string. A negative length
3 means that characters are selected in the opposite (right to left) direction from the starting position. If the substring
4 request extends beyond the beginning or end of the string, it is clipped (shortened) to the bounds of the string. If the
5 substring request is completely outside of the original string, a null string is returned. For the pattern-form, it returns
6 the pattern string if the pattern matches or a null string if the pattern does not match. The usefulness of this mechanism
7 is discussed next.

8 The substring operation can appear on the left side of assignment, where it defines a replacement substring. The
9 length of the right string may be shorter, the same, or longer than the length of left string. Hence, the left string may
10 decrease, stay the same, or increase in length.

```

string digit = "0123456789";
digit( 3, 3 ) = " ";           "0126789"
11 digit( 4, 3 ) = "xyz";      "0126xyz"
digit( 7, 0 ) = "***";       "0126xyz"
digit(-4, 3) = "$$$";        "012$$$$"
digit( 5 ) = "LLL";          "012$$LLL"

```

12 Now substring pattern matching is useful on the left-hand side of assignment.

```

13 digit( "$$" ) = "345";      "012345LLL"
digit( "LLL" ) = "6789";     "0123456789"

```

14 The replace operation extends substring to substitute all occurrences.

```

15 s = replace( "PETER", "E", "XX" );   "PXXTXXR"
s = replace( "PETER", "ET", "XX" );   "PXXER"
s = replace( "PETER", "W", "XX" );     "PETER"

```

16 The replacement is done left-to-right and substituted text is not examined for replacement.

17 13.7 Searching

18 The find operation returns the position of the first occurrence of a key in a string. If the key does not appear in the
19 string, the length of the string is returned.

```

string digit = "0123456789";
20 i = find( digit, '3' );      3
i = find( digit, "45" );      4
i = find( digit, "abc" );     10

```

21 A character-class operation indicates if a string is composed completely of a particular class of characters, *e.g.*,
22 alphabetic, numeric, vowels, *etc.*

```

charclass vowels{ "aeiouy" };
23 i = include( "aaeiuyoo", vowels );  8 // compliant
i = include( "aabiuyoo", vowels );   2 // b non-compliant

```

24 vowels defines a character class and function include checks if all characters in the string appear in the class (compli-
25 ance). The position of the last character is returned if the string is compliant or the position of the first non-compliant
26 character. There is no relationship between the order of characters in the two strings. Function exclude is the reverse
27 of include, checking if all characters in the string are excluded from the class (compliance).

```

1   i = exclude( "cdbfghmk", vowels ); | 8 // compliant
   i = exclude( "cdy fghmk", vowels ); | 2 // y non-compliant

```

2 Both forms can return the longest substring of compliant characters.

```

3   s = include( "aaeiuyoo", vowels ); | "aaeiuyoo"
   s = include( "aabiuyoo", vowels ); | "aa"
   s = exclude( "cdbfghmk", vowels ); | "cdbfghmk"
   s = exclude( "cdy fghmk", vowels ); | "cd"

```

4 There are also versions of include and exclude, returning a position or string, taking a validation function, like one
5 of the C character-class functions.¹⁵

```

6   i = include( "1FeC34aB", isxdigit ); | 8 // compliant
   i = include( ". , ; ! \ " , ispunct ); | 6 // compliant
   i = include( "XXXx", isupper ); | 3 // non-compliant

```

7 These operations perform an *apply* of the validation function to each character, where the function returns a boolean
8 indicating a stopping condition for the search. The position of the last character is returned if the string is compliant
9 or the position of the first non-compliant character.

10 The translate operation returns a string with each character transformed by one of the C character transformation
11 functions.

```

12  s = translate( "abc", toupper ); | "ABC"
   s = translate( "ABC", tolower ); | "abc"
   int tospace( int c ) { return isspace( c ) ? ' ' : c; }
   s = translate( "X X\tX\nX", tospace ); | "X X X X"

```

13 13.8 Returning N on Search Failure

14 Some of the prior string operations are composite, *e.g.*, string operations returning the longest substring of compliant
15 characters (include) are built using a search and then substring the appropriate text. However, string search can fail,
16 which is reported as an alternate search outcome, possibly an exception. Many string libraries use a return code to
17 indicate search failure, with a failure value of 0 or -1 (PL/I [21] returns 0). This semantics leads to the awkward
18 pattern, which can appear many times in a string library or user code.

```

19   i = exclude( s, alpha );
20   if ( i != -1 ) return s( 0, i );
21   else return "";

```

22 C \forall adopts a return code but the failure value is taken from the index-of function in APL [24], which returns the
23 length of the target string N (or $N + 1$ for 1 origin). This semantics allows many search and substring functions to be
24 written without conditions, *e.g.*:

```

25   string include( const string & s, int (*f)( int ) ) { return s( 0, include( s, f ) ); }
26   string exclude( const string & s, int (*f)( int ) ) { return s( 0, exclude( s, f ) ); }

```

27 In string systems with an $O(1)$ length operator, checking for failure is low cost.

```

28   if ( include( line, alpha ) == len( line ) ) ... // not found, 0 origin

```

29 Figure 5 compares C++ and C \forall string code for extracting words from a line of text, repeatedly removing non-word
30 text and then a word until the line is empty. The C \forall code is simpler solely because of the choice for indicating search
31 failure. (A simplification of the C++ version is to concatenate a sentinel character at the end of the line so the call to
32 find_first_not_of does not fail.)

33 13.9 C Compatibility

34 To ease conversion from C to C \forall , C \forall provides companion C string functions. Hence, it is possible to convert a block
35 of C string operations to C \forall strings just by changing the type `char *` to `string`.

¹⁵It is part of the hereditary of C that these function take and return an `int` rather than a `bool`, which affects the function type.

| C++ | CV |
|---|--|
| <pre> for (;;) { string::size_type posn = line.find_first_of(alpha); if (posn == string::npos) break; line = line.substr(posn); posn = line.find_first_not_of(alpha); if (posn != string::npos) { cout << line.substr(0, posn) << endl; line = line.substr(posn); } else { cout << line << endl; line = ""; } } </pre> | <pre> for () { size_t posn = exclude(line, alpha); if (posn == len(line)) break; line = line(posn); posn = include(line, alpha); sout line(0, posn); line = line(posn); } </pre> |

Figure 5: Extracting Words from Line of Text

```

char s[32]; // string s;
strlen( s );          strcpy( s, "abc" );
1  strlen( s, 3 );    strncpy( s, "abcdef", 3 );
  strcmp( s, "abc" );  strcat( s, "xyz" );
  strncmp( s, "abc", 3 );  strncat( s, "uvwxyz", 3 );

```

2 However, the conversion fails with I/O because printf cannot print a string using format code %s as CV strings are not
3 null terminated. Nevertheless, this capability does provide a useful starting point for conversion to safer CV strings.

4 13.10 I/O Operators

5 The ability to input and output strings is as essential as for any other type. The goal for character I/O is to also work
6 with groups rather than individual characters. A comparison with C++ string I/O is presented as a counterpoint to CV
7 string I/O.

8 The C++ output << and input >> operators are defined on type string. C++ output for **char**, **char ***, and string are
9 similar. The C++ manipulators are setw, and its associated width controls left, right and setfill.

```

10 string s = "abc";
   cout << setw(10) << left << setfill( 'x' ) << s << endl; | "abcxxxxxxxx"

```

11 The CV input/output operator | is defined on type string. CV output for **char**, **char ***, and string are similar. The CV
12 manipulators are bin, oct, hex, wd, and its associated width control and left.

```

string s = "abc";
13 sout | bin( s ) | nl | "0b1100001 0b1100010 0b1100011"
   | oct( s ) | nl | "0141 0142 0143"
   | hex( s ) | nl | "0x61 0x62 0x63"
   | wd( 10, s ) | nl | "          abc"
   | wd( 10, 2, s ) | nl | "         ab"
   | left( wd( 10, s ) ); | "abc      "

```

14 C++ setfill is not considered an important string manipulator.

15 C++ input matching for **char**, **char ***, and string are similar, where *all* input characters are read from the current
16 point in the input stream to the end of the type size, format width, whitespace, end of line ('\n'), or end of file. The
17 C++ manipulator is setw to restrict the size. Reading into a **char** is safe as the size is 1, **char *** is unsafe without using
18 setw to constraint the length (which includes '\0'), string is safe as its grows dynamically as characters are read.

```

char ch, c[10];
string s;
19 cin >> ch >> setw( 5 ) >> c >> s; | 'a' "bcde" "fg"
   abcde fg

```

20 Input text can be *gulped*, including whitespace, from the current point to an arbitrary delimiter character using getline.

1 The CV philosophy for input is that, for every constant type in C, these constants should be usable as input. For
 2 example, the complex constant 3.5+4.1i can appear as input to a complex variable. CV input matching for **char**, **char ***,
 3 and string are similar. C-strings may only be read with a width field, which should match the string size. Certain input
 4 manipulators support a scanset, which is a simple regular expression from printf. The CV manipulators for these types
 5 are wdi,¹⁶ and its associated width control and left, quote, incl, excl, and getline.

```

char ch, c[10];
string s;

sin | ch | wdi( 5, c ) | s;
abcde fg
6 sin | quote( ch ) | quote( wdi( sizeof(c), c ) ) | quote( s, ' [', ' ] ' ) | nl;
'a' "bcde" [fg]
sin | incl( "a-zA-Z0-9 ?!\&\n", s ) | nl;
x?&000xyz TOM !.
sin | excl( "a-zA-Z0-9 ?!\&\n", s );
<>{}STOP
'a' "bcde" "fg"
'a' "bcde" "fg"
"x?&000xyz TOM !"
"<>{}{}"

```

7 Note, the ability to read in quoted strings with whitespace to match with program string constants. The nl at the end of
 8 an input ignores the rest of the line.

9 14 Enumeration

10 An *enumeration* is a compile-time mechanism to alias names to constants, like **typedef** is a mechanism to alias names
 11 to types. Its purpose is to define a restricted-value type providing code-readability and maintenance – changing an
 12 enum's value automatically updates all name usages during compilation.

13 An enumeration type is a set of names, each called an *enumeration constant* (shortened to *enum*) aliased to a fixed
 14 value (constant).

```

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

```

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

```

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

```

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

```

24 enum { Prime = 103, BufferSize = 1024 };

```

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

```

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

```

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

30 14.1 Enum type

31 The type of enums can be any type, and an enum's value comes from this type. Because an enum is a constant, it
 32 cannot appear in a mutable context, *e.g.*, Mon = Sun is disallowed, and has no address (it is an rvalue). Therefore,
 33 an enum is automatically converted to its constant's base-type, *e.g.*, comparing/printing an enum compares/prints its
 34 value rather than the enum name; there is no mechanism to print the enum name.

¹⁶Due to an overloading issue in the type-resolver, the input width name must be temporarily different from the output, wdi versus wd.

1 The default enum type is **int**. Hence, Days is the set type Mon, Tue, ..., Sun, while the type of each enum is **int**
 2 and each enum represents a fixed integral value. If no values are specified for an integral enum type, the enums are
 3 automatically numbered by one from left to right starting at zero. Hence, the value of enum Mon is 0, Tue is 1, ..., Sun
 4 is 6. If an enum value is specified, numbering continues by one from that value for subsequent unnumbered enums. If
 5 an enum value is a *constant* expression, the compiler performs constant-folding to obtain a constant value.

6 CV allows other integral types with associated values.

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

9 For enumeration Letter, enum A's value is explicitly set to 'A', with B and C implicitly numbered with increasing
 10 values from 'A', and similarly for enums I, J, and K.

11 Non-integral enum types must be explicitly initialized, *e.g.*, **double** is not automatically numbered by one.

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

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

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

```
33 Math m = PI; // allowed
34 double d = PI; // allowed, conversion to base type
35 m = E; // allowed
36 m = Alph; // disallowed
37 m = 3.141597; // disallowed
38 d = m; // allowed
39 d = Alph; // disallowed
40 Letter l = A; // allowed
41 Greek g = Alph; // allowed
42 l = Alph; // allowed, conversion to base type
43 g = A; // disallowed
```

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

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

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

49 14.2 Inheritance

50 Plan-9 inheritance may be used with enumerations.

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

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

7 Specifically, the inheritance relationship for Names is:

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

9 Hence, given

```
10 void f( Name );
11 void g( Name2 );
12 void h( Name3 );
13 void j( const char * );
```

14 the following calls are valid

```
15 f( Fred );
16 g( Fred ); g( Jill );
17 h( Fred ); h( Jill ); h( Sue );
18 j( Fred ); j( Jill ); j( Sue ); j( 'W' );
```

19 Note, the validity of calls is the same for call-by-reference as for call-by-value, and **const** restrictions are the same as
 20 for other types.

21 Enums cannot be created at runtime, so inheritance problems, such as contra-variance do not apply. Only instances
 22 of the enum base-type may be created at runtime.

23 15 Routine Definition

24 C \forall supports a new syntax for routine definition, as well as C11 and K&R routine syntax. The point of the new syntax
 25 is to allow returning multiple values from a routine [16, 27], *e.g.*:

```
26 [ int o1, int o2, char o3 ] f( int i1, char i2, char i3 ) {
27     routine body
28 }
```

29 where routine f has three output (return values) and three input parameters. Existing C syntax cannot be extended with
 30 multiple return types because it is impossible to embed a single routine name within multiple return type-specifications.

31 In detail, the brackets, [], enclose the result type, where each return value is named and that name is a local variable
 32 of the particular return type.¹⁷ The value of each local return variable is automatically returned at routine termination.
 33 Declaration qualifiers can only appear at the start of a routine definition, *e.g.*:

```
34 extern [ int x ] g( int y ) {}
```

35 Lastly, if there are no output parameters or input parameters, the brackets and/or parentheses must still be specified; in
 36 both cases the type is assumed to be void as opposed to old style C defaults of int return type and unknown parameter
 37 types, respectively, as in:

```
38 [] g(); // no input or output parameters
39 [ void ] g( void ); // no input or output parameters
```

40 Routine f is called as follows:

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

42 The list of return values from f and the grouping on the left-hand side of the assignment is called a *return list* and
 43 discussed in Section 12.

44 C \forall style declarations cannot be used to declare parameters for K&R style routine definitions because of the fol-
 45 lowing ambiguity:

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

47 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
 48 string “[5] **int** x” declares a C \forall style parameter x of type array of 5 integers. Since the strings overlap starting with the

¹⁷Michael Tiemann, with help from Doug Lea, provided named return values in g++, circa 1989.

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

2 As well, C \forall -style declarations cannot be used to declare parameters for C-style routine-definitions because of the
3 following ambiguity:

```
4 typedef int foo;
5 int f( int (* foo) ); // foo is redefined as a parameter name
```

6 The string “**int** (* foo)” declares a C-style named-parameter of type pointer to an integer (the parenthesis are superfluous), while the same string declares a C \forall style unnamed parameter of type routine returning integer with unnamed
7 parameter of type pointer to foo. The redefinition of a type name in a parameter list is the only context in C where the
8 character * can appear to the left of a type name, and C \forall relies on all type qualifier characters appearing to the right of
9 the type name. The inability to use C \forall declarations in these two contexts is probably a blessing because it precludes
10 programmers from arbitrarily switching between declarations forms within a declaration contexts.

11 C-style declarations can be used to declare parameters for C \forall style routine definitions, e.g.:

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

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

```
17 #define ptoa( n, d ) int (*n)[ d ]
18 int f( ptoa( p, 5 ) ) ... // expands to int f( int (*p)[ 5 ] )
19 [ int ] f( ptoa( p, 5 ) ) ... // expands to [ int ] f( int (*p)[ 5 ] )
```

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

21 15.1 Named Return Values

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

```
24 int f() {
25     int x;
26     ... x = 0; ... x = y; ...
27     return x;
28 }
```

29 Because the value in the return variable is automatically returned when a C \forall routine terminates, the **return** statement
30 *does not* contain an expression, as in:

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

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

```
34 [ int x, int y ] f() {
35     ...
36 } // implicitly return x, y
```

37 In this case, the current values of x and y are returned to the calling routine just as if a **return** had been encountered.

38 Named return values may be used in conjunction with named parameter values; specifically, a return and parameter
39 can have the same name.

```
40 [ int x, int y ] f( int, x, int y ) {
41     ...
42 } // implicitly return x, y
```

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

```
44 [ int x, int y ] f( int, x, int y ); // prototype declaration
45 int a, b;
46 [a, b] = f( f( a, b ) );
```

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

6 15.2 Routine Prototype

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

```
9 [ int x ] f (); // returning int with no parameters
10 [ * int ] g ( int y ); // returning pointer to int with int parameter
11 [ ] h ( int, char ); // returning no result with int and char parameters
12 [ * int, int ] j ( int ); // returning pointer to int and int, with int parameter
```

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

```
16 C : const double bar1(), bar2( int ), bar3( double );
17 CV : [const double] foo(), foo( int ), foo( double ) { return 3.0; }
```

18 CV allows the last routine in the list to define its body.

19 Declaration qualifiers can only appear at the start of a CV routine declaration,⁷ *e.g.*:

```
20 extern [ int ] f ( int );
21 static [ int ] g ( int );
```

22 15.3 Postfix Function

23 CV provides an alternative call syntax where the argument appears before the function name. The syntax uses the
 24 backquote ``` to separate the parameters/arguments and function name: `?`` denotes a postfix-function name, *e.g.*,
 25 `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 |
|----|--|------------------------------|---------------------------|------------------------------------|
| | <code>int ?`h(int s);</code> | <code>0`h;</code> | | <code>int (* ?`p)(int i);</code> |
| | <code>int ?`h(double s);</code> | <code>3.5`h;</code> | <code>int i = 7;</code> | <code>?`p = ?`h;</code> |
| 26 | <code>int ?`m(char c);</code> | <code>'1'`m;</code> | <code>i`h;</code> | <code>3`p;</code> |
| | <code>int ?`m(const char * s);</code> | <code>"123" "456" `m;</code> | <code>(i + 3)`h;</code> | <code>i`p;</code> |
| | <code>int ?`t(int a, int b, int c);</code> | <code>[1, 2, 3]`t;</code> | <code>(i + 3.5)`h;</code> | <code>(i + 3)`p;</code> |

27 Note, to pass *multiple* arguments to a postfix function requires a tuple, *e.g.*, `[1, 2, 3]`t`, which forms a single argument
 28 that is flattened into the multiple arguments (see Section 20, p. 40). Similarly, if the argument is an expression, it must
 29 be parenthesized, *e.g.*, `(i + 3)`h`, or only the last operand of the expression is the argument, *e.g.*, `i + (3`h)`.

30 Figure 6 shows a common example for postfix functions: converting basic literals into user literals. (See Sec-
 31 tion E.1, p. 77 for other uses for postfix functions.) The CV example (left) stores a mass in units of stones (1 stone
 32 = 14 lb or 6.35 kg) and provides an addition operator `?+?` (imagine a full set of arithmetic operators). The arithmetic
 33 operators manipulate stones and the postfix operations convert to/from different units. The three postfixing function
 34 names `st`, `lb`, and `kg`, represent units stones, pounds, and kilograms, respectively. Each name has two forms that bidi-
 35 rectional convert: a value of a specified unit to stones, *e.g.*, `w = 14`lb` \Rightarrow `w == 1` stone or a Weight from stones back to
 36 specific units, *e.g.*, `w`lb` (1 stone) to 14. A similar group of postfix functions provide user constants for converting time
 37 units into nanoseconds, which is the basic time unit, *e.g.*, `ns`, `us`, `ms`, `s`, `m`, `h`, `d`, and `w`, for nanosecond, microsecond,
 38 millisecond, second, minute, hour, day, and week, respectively. (Note, month is not a fixed period of time nor is year
 39 because of leap years.)

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

| CV Postfix Routine | C++ User Literals |
|---|---|
| <pre> 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_11111`st; w = 0_233`lb; w = 0x_9b`kg; w = 5.5`st + 8`kg + 25.01`lb + heavy; } </pre> | <pre> 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 = 0b11111`st; w = 0'233@_lb@; // quote separator w = 0x9b@_kg@; w = 5.5@_st@ + 8@_kg@ + 25.01@_lb@ + heavy; } </pre> |

Figure 6: Units: Stone, Pound, Kilogram Comparison

16 Routine Pointers

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

```

3 * [ int x ] () fp; // pointer to routine returning int with no parameters
4 * [ * int ] ( int y ) gp; // pointer to routine returning pointer to int with int parameter
5 * [ ] ( int, char ) hp; // pointer to routine returning no result with int and char parameters
6 * [ * int, int ] ( int ) jp; // pointer to routine returning pointer to int and int, with int parameter

```

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

```

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

```

17 Default and Named Parameter

Default and named parameters [20]¹⁸ 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:

```

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

```

allowable calls are:

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

¹⁸Francez [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.

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

- 4 • Routines with a large number of parameters are often very generalized, giving a programmer a number of
 5 different options on how a computation is performed. For many of these routines, there are standard or default
 6 settings that work for the majority of computations. Without default values for parameters, a programmer is
 7 forced to specify these common values all the time, resulting in long argument lists that are error prone.
- 8 • When a routine's interface is augmented with new parameters, it extends the interface providing generalizabil-
 9 ity¹⁹ (somewhat like the generalization provided by inheritance for classes). That is, all existing calls are still
 10 valid, although the call must still be recompiled.

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

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

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

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

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

| | default arguments | overloading |
|----|--|---|
| 21 | <code>void f(int x, int y = 2, int z = 3) {...}</code> | <code>void f(int x, int y, int z) {...}</code> <code>void f(int x) { f(x, 2, 3); }</code> <code>void f(int x, int y) { f(x, y, 3); }</code> |

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

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

26 17.2 Named (or Keyword)

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

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

30 allowable calls are:

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

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

- 39 • Remembering the names of the parameters may be easier than the order in the routine definition.
- 40 • Parameter names provide documentation at the call site (assuming the names are descriptive).
- 41 • Changes can be made to the order or number of parameters without affecting the call (although the call must
 42 still be recompiled).

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

¹⁹“It should be possible for the implementor of an abstraction to increase its generality. So long as the modified abstraction is a generalization of 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 void f( int x, int ?y, int ?z ); // prototype: allowed
2 void f( int ?x, int , int ?z ); // prototype: allowed
3 void f( int x, int ?y, int ?z ) {} // definition: disallowed

```

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

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

```

8 int f( int ?i, int ?j );
9 int f( int ?i, double ?j );
10 f( ?j = 3, ?i = 4 ); // 1st f
11 f( ?i = 7, ?j = 8.1 ); // 2nd f

```

12 17.3 Mixed Default/Named

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

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

16 allowable calls are:

```

17 f( 3 ); // rewrite => f( 3, 1, 2 )
18 f( 3, 4 ); // rewrite => f( 3, 4, 2 )
19 f( 3, ?z = 5 ); // rewrite => f( 3, 1, 5 )
20 f( 3, 4, ?z = 5 ); // rewrite => f( 3, 4, 5 )
21 f( ?z = 5, 3 ); // rewrite => f( 3, 1, 5 )
22 f( 3, ?z = 5, 4 ); // rewrite => f( 3, 4, 5 )

```

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

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

25 C# named and default arguments are backwards compatible with C. C++ only supports default parameters; Ada supports both named and default parameters.

27 18 Unnamed Structure Fields

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

```

29 struct {
30     int f1; // named field
31     int f2 : 4; // named field with bit field size
32     int : 3; // unnamed field for basic type with bit field size
33     int ; // disallowed, unnamed field
34     int *; // disallowed, unnamed field
35     int (*)( int ); // disallowed, unnamed field
36 };

```

37 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.:

```

40 struct {
41     int , , ; // 3 unnamed fields
42 }

```

43 19 Nesting

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

45 19.1 Type Nesting

46 C# allows type nesting, and type qualification of the nested types (see Figure 7), whereas 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

| C Type Nesting | C Implicit Hoisting | CV |
|--|--|---|
| <pre> 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; } </pre> | <pre> 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; </pre> | <pre> 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; } </pre> |

Figure 7: Type Nesting / Qualification

1 hoisted outside of type S into the containing block scope. In the right example in CV, the types are not hoisted and
 2 accessed using the field-selection operator “.” for type qualification, as does Java, rather than the C++ type-selection
 3 operator “::”.

4 19.2 Routine Nesting

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

```

8 forall( T | { int ?<( T, T ); } )
9 void qsort( const T * arr, size_t dimension );
10 which can be used to sort in ascending and descending order by locally redefining the less-than operator into greater-
11 than.
12 const unsigned int size = 5;
13 int ia[size];
14 ... // assign values to array ia
15 qsort( ia, size ); // sort ascending order using builtin ?<?
16 {
17 int ?<( int x, int y ) { return x > y; } // nested routine
18 qsort( ia, size ); // sort descending order by local redefinition
19 }

```

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

```

23 [* int]( int ) foo() { // int (*foo())( int )
24 int i = 7;
25 int bar( int p ) {
26 i += 1; // dependent on local variable
27 sout | i;
28 }
29 return bar; // undefined because of local dependence
30 }
31 int main() {
32 [* int]( int ) fp = foo(); // int (*fp)( int )
33 sout | fp( 3 );
34 }

```

1 because

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

4 20 Tuple

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

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

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

8 20.1 Multiple-Return-Value Functions

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

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

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

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

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

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

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

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

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

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

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

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

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

```

1  void g( int, int );           // 1
2  void g( double, double );   // 2
3  g( div( 13, 5 ) );         // select 1
4  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.

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

```

10 [ int, int ] div( int x, int y ); // from include stdlib
11 printf( "%d %d\n", div( 13, 5 ) ); // print quotient/remainder
12
13 [ double, double ] modf( double x ); // from include math
14 printf( "%g %g\n", modf( 13.5 ) ); // print integral/fractional components

```

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

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

```

22 int quot, rem;
23 [ quot, rem ] = div( 13, 5 ); // assign multiple variables
24 printf( "%d %d\n", quot, rem ); // print quotient/remainder

```

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

20.2 Expressions

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

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

20.3 Variables

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

```

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

```

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

```

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

```

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

```

50 [ double, int ] di;

```

```

1  [ double, int ] * pdi
2  [ double, int ] adi[10];
3  This examples declares a variable of type [double, int], a variable of type pointer to [double, int], and an array of ten
4  [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 enp 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.*:

```

9  [int, double] x;
10 [char *, int] f();
11 void g(double, int);
12 [int, double] * p;
13
14 int y = x.0;           // access int component of x
15 y = f().1;           // access int component of f
16 p->0 = 5;             // access int component of tuple pointed-to by p
17 g( x.1, x.0 );       // rearrange x to pass to g
18 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 [33, p. 45].

20.5 Flattening and Structuring

As evident in previous examples, tuples in CV 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.

```

26 int f(int, int);
27 int g([int, int]);
28 int h(int, [int, int]);
29 [int, int] x;
30 int y;
31
32 f(x); // flatten
33 g(y, 10); // structure
34 h(x, y); // flatten & structure

```

In CV, 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, 33], 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 CV, 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 CV. 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

1 compared with each value in the function's parameter list. If the parameter's type is not a tuple type, then the current
 2 argument value is unified with the parameter type, and on success the next argument and parameter are examined. If
 3 the parameter's type is a tuple type, then the structuring conversion takes effect, recursively applying the parameter
 4 matching algorithm using the tuple's component types as the parameter list types. Assuming a successful unification,
 5 eventually the algorithm gets to the end of the tuple type, which causes all of the matching expressions to be consumed
 6 and structured into a tuple expression. For example, in

```
7   int f(int, [double, int]);
8   f([5, 10.2], 4);
```

9 There is only a single definition of `f`, and 3 arguments with only single interpretations. First, the argument alternative
 10 list `[5, 10.2], 4` is flattened to produce the argument list `5, 10.2, 4`. Next, the parameter matching algorithm begins,
 11 with $P = \text{int}$ and $A = \text{int}$, which unifies exactly. Moving to the next parameter and argument, $P = [\text{double}, \text{int}]$ and
 12 $A = \text{double}$. This time, the parameter is a tuple type, so the algorithm applies recursively with $P' = \text{double}$ and
 13 $A = \text{double}$, which unifies exactly. Then $P' = \text{int}$ and $A = \text{double}$, which again unifies exactly. At this point, the end
 14 of P' has been reached, so the arguments `10.2, 4` are structured into the tuple expression `[10.2, 4]`. Finally, the end of
 15 the parameter list P has also been reached, so the final expression is `f(5, [10.2, 4])`.

16 20.6 Assignment

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

```
20   int x;
21   double y;
22   [int, double] z;
23   [y, x] = 3.14;           // mass assignment
24   [x, y] = z;             // multiple assignment
25   z = 10;                 // mass assignment
26   z = [x, y];            // multiple assignment
```

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

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

```
32   [ int, int ] x, y, z;
33   [ x, y ] = z;           // multiple assignment, invalid 4 != 2
```

34 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
 35 previous example, `[x, y] = z`, `z` is flattened into `z.0, z.1`, and the assignments `x = z.0` and `y = z.1` happen.

36 A mass assignment assigns the value R to each L_i . For a mass assignment to be valid, $?=(\&\$L_i\$, \&\$R\$)$ must be
 37 a well-typed expression. These semantics differ from C cascading assignment (*e.g.*, `a=b=c`) in that conversions are
 38 applied to R in each individual assignment, which prevents data loss from the chain of conversions that can happen
 39 during a cascading assignment. For example, `[y, x] = 3.14` performs the assignments `y = 3.14` and `x = 3.14`, which results
 40 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
 41 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.

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

```
45   int x = 10, y = 20;
46   [ x, y ] = [ y, x ];
```

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

48 In CV, tuple assignment is an expression where the result type is the type of the left side of the assignment, as
 49 in normal assignment. That is, a tuple assignment produces the value of the left-hand side after assignment. These
 50 semantics allow cascading tuple assignment to work out naturally in any context where a tuple is permitted. These
 51 semantics are a change from the original tuple design in K-W C [33], wherein tuple assignment was a statement that
 52 allows cascading assignments as a special case. Restricting tuple assignment to statements was an attempt to fix

1 what was seen as a problem with side-effects, wherein assignment can be used in many different locations, such as in
 2 function-call argument position. While permitting assignment as an expression does introduce the potential for subtle
 3 complexities, it is impossible to remove assignment expressions from \mathcal{CV} without affecting backwards compatibility.
 4 Furthermore, there are situations where permitting assignment as an expression improves readability by keeping code
 5 succinct and reducing repetition, and complicating the definition of tuple assignment puts a greater cognitive burden
 6 on the user. In another language, tuple assignment as a statement could be reasonable, but it would be inconsistent
 7 for tuple assignment to be the only kind of assignment that is not an expression. In addition, K-W C permits the
 8 compiler to optimize tuple assignment as a block copy, since it does not support user-defined assignment operators.
 9 This optimization could be implemented in \mathcal{CV} , but it requires the compiler to verify that the selected assignment
 10 operator is trivial.

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

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

16 The tuple expression begins with a mass assignment of 1.5 into [b, d], which assigns 1.5 into b, which is truncated to
 17 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
 18 (*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
 19 tuple [1, 1] is used as an expression in the call to f.

20 20.7 Construction

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

```
25     struct S;
26     void ?{S *};           // (1)
27     void ?{S *, int};     // (2)
28     void ?{S * double};  // (3)
29     void ?{S *, S};      // (4)
30
31     [S, S] x = [3, 6.28]; // uses (2), (3), specialized constructors
32     [S, S] y;           // uses (1), (1), default constructor
33     [S, S] z = x.0;     // uses (4), (4), copy constructor
```

34 In this example, x is initialized by the multiple constructor calls $?\{&x.0, 3\}$ and $?\{&x.1, 6.28\}$, while y is initialized
 35 by two default constructor calls $?\{&y.0\}$ and $?\{&y.1\}$. z is initialized by mass copy constructor calls $?\{&z.0, x.0\}$ and
 36 $?\{&z.1, x.0\}$. Finally, x, y, and z are destructed, *i.e.*, the calls $\wedge?\{&x.0\}$, $\wedge?\{&x.1\}$, $\wedge?\{&y.0\}$, $\wedge?\{&y.1\}$, $\wedge?\{&z.0\}$,
 37 and $\wedge?\{&z.1\}$.

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

```
41     struct S { int x; double y; };
42     void ?{([int, double] * this, S s) {
43         this->0 = s.x;
44         this->1 = s.y;
45     }
```

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

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

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

1 20.8 Member-Access Expression

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

```
4 struct S { char x; int y; double z; } s;  
5 s.[x, y, z];
```

6 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
7 expression with type `T`, where `T` supports member access expressions, and `x, y, z` are all members of `T` with types
8 `T$_x$, T$_y$, and T$_z$` respectively. Then the type of `e.[x, y, z]` is `[T$_x$, T$_y$, T$_z$]`.

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

```
10 s.[y, z, x] = [3, 3.2, 'x']; // equivalent to s.x = 'x', s.y = 3, s.z = 3.2  
11 f(s.[y, z]); // equivalent to f(s.y, s.z)
```

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

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

```
17 [int, int, long, double] x;  
18 void f(double, long);  
19  
20 f(x.[0, 3]); // f(x.0, x.3)  
21 x.[0, 1] = x.[1, 0]; // [x.0, x.1] = [x.1, x.0]  
22 [long, int, long] y = x.[2, 0, 2];
```

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

```
24 struct A { double i; int j; };  
25 struct B { int * k; short l; };  
26 struct C { int x; A y; B z; } v;  
27 v.[x, y.[i, j], z.k];
```

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

```
32 [int, float, double] f();  
33 [double, float] x = f().[2, 1]; // f() called once
```

34 In K-W C, member tuple expressions are known as *record field tuples* [33]. Since `Cv` permits these tuple-access
35 expressions using structures, unions, and tuples, *member tuple expression* or *field tuple expression* is more appropriate.

36 20.9 Casting

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

```
40 int i, *ip;  
41 double d;  
42 d = (double)i; // conversion  
43 ip = (int *)d; // coercion
```

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

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

1. `dynamic_cast` Used for conversion of polymorphic types.
2. `static_cast` Used for conversion of nonpolymorphic types.
3. `const_cast` Used to remove the type qualifiers and possibly attributes.
4. `reinterpret_cast` Used for simple reinterpretation of bits.

```

7      ' ( ' type_no_function ' ) ' cast_expression
8      ' ( ' aggregate_control '&' ' ) ' cast_expression // CFA
9      ' ( ' aggregate_control '*' ' ) ' cast_expression // CFA
10     ' ( ' VIRTUAL ' ) ' cast_expression // CFA
11     ' ( ' VIRTUAL type_no_function ' ) ' cast_expression // CFA
12     ' ( ' RETURN type_no_function ' ) ' cast_expression // CFA (ASCRPTION)
13     ' ( ' COERCE type_no_function ' ) ' cast_expression // CFA (COERCION)
14     ' ( ' qualifier_cast_list ' ) ' cast_expression // CFA, (modify CVs of cast_expression)

```

Specialized Casts

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

C (Conversion) Cast

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

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

Ascription Cast

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

```

33     int f(int); // f1
34     int f(double); // f2
35     int g(int); // g1
36     double g(long); // g2
37
38     f((double)42); // selects f2 by cast on argument
39     (double)g(42); // does NOT select g2, argument conversion cost results in g1

```

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

```

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

```

Coercion Cast

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

```

1   int i = 5;
2   double d = *(double*)&i; // value coercion
3   printf( "%g %g %x\n", d, *(double *)&i, *(int *)&d );
4
5   int i = 5; // pointer coercion
6   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 modifying 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`-qualifiers, (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 qualifiers.

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

In CV, 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.

```

20  int f(); // (1)
21  double f(); // (2)
22
23  f(); // ambiguous - (1),(2) both equally viable
24  (int)f(); // choose (2)

```

Since casting is a fundamental operation in CV, 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.

```

27  1  int f();
28  2  void g();
29  2
30  3  (void)f(); // valid, ignore results
31  4  (int)g(); // invalid, void cannot be converted to int
32  4
33  5  struct A { int x; };
34  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 [22, 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,

```

45  [int, int, int] f();
46  [int, [int, int], int] g();
47
48  ([int, double])f(); // (1) valid
49  ([int, int, int])g(); // (2) valid
50  ([void, [int, int]])g(); // (3) valid
51  ([int, int, int, int])g(); // (4) invalid
52  ([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

1 tuple, (3) discards the first and third return values, which is effectively equivalent to `[(int)(g().1.0), (int)(g().1.1)]`. if casts
 2 become function calls, what would they look like? would need a way to specify the target type, which seems awkward.
 3 Also, C++ basically only has this because classes are closed to extension, while we don't have that problem (can have
 4 floating constructors for any type). Note that a cast is not a function call in C \forall , so flattening and structuring conversions
 5 do not occur for cast expressions. As such, (4) is invalid because the cast target type contains 4 components, while the
 6 source type contains only 3. Similarly, (5) is invalid because the cast `[(int, int, int)(g().1)]` is invalid. That is, it is invalid
 7 to cast `[int, int]` to `[int, int, int]`.

8 20.10 Polymorphism

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

```
12 forall(T, U &)  
13 void f(T x, U * y);  
14  
15 f([5, "hello"]);
```

16 In this example, `[5, "hello"]` is flattened, so that the argument list appears as `5, "hello"`. The argument matching
 17 algorithm binds `T` to `int` and `U` to `const char`, and calls the function as normal.

18 Tuples can contain polymorphic types. For example, a plus operator can be written to add two triples of a type
 19 together.

```
20 forall(T | { T ?+?(T, T); })  
21 [T, T, T] ?+?([T, T, T] x, [T, T, T] y) {  
22     return [x.0+y.0, x.1+y.1, x.2+y.2];  
23 }  
24 [int, int, int] x;  
25 int i1, i2, i3;  
26 [i1, i2, i3] = x + ([10, 20, 30]);
```

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

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

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

```
35 forall(T | { T ?+?(T, T); })  
36 [T, T, T] ?+?([T, T] x, [T, T, T, T]);  
37 forall(T | { T ?+?(T, T); })  
38 [T, T, T] ?+?(T x, [T, T, T, T, T]);
```

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

```
45 void f([int, int], int, int);  
46  
47 f([0, 0], 0, 0); // no cost  
48 f(0, 0, 0, 0); // cost for structuring  
49 f([0, 0], [0, 0]); // cost for flattening  
50 f([0, 0, 0], 0); // cost for flattening and structuring
```

51 Until this point, it has been assumed that assertion arguments must match the parameter type exactly, modulo
 52 polymorphic specialization (*i.e.*, no implicit conversions are applied to assertion arguments). This decision presents a

1 conflict with the flexibility of tuples.

2 20.10.1 Assertion Inference

```
3 int f([int, double], double);
4 forall(T, U | { T f(T, U, U); })
5 void g(T, U);
6 g(5, 10.21);
```

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

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

```
14 int _thunk(int _p0, double _p1, double _p2) {
15     return f([_p0, _p1], _p2);
16 }
```

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

19 21 Tuples

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

```
22 [ exprlist ]
```

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

```
25 [ x, y, z ]
26 [ 2 ]
27 [ v + w, x * y, 3.14159, f() ]
```

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

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

```
35 [ typelist ]
```

36 where *typelist* is a list of one or more legal `CV` or `C` type specifications separated by commas, which may include other
37 tuple type specifications. Examples of tuple types include:

```
38 [ unsigned int, char ]
39 [ double, double, double ]
40 [ * int, int * ] // mix of CFA and ANSI
41 [ * [ 5 ] int, * * char, * [ [ int, int ] ] (int, int) ]
```

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

44 Examples of declarations using tuple types are:

```
45 [ int, int ] x; // 2 element tuple, each element of type int
46 * [ char, char ] y; // pointer to a 2 element tuple
47 [ [ int, int ] ] z ([ int, int ]);
```

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

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

```
4 f( [ 1, x+2, fred() ] );
5 f( 1, x+2, fred() );
```

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

```
8 [ int, int ] w1;
9 [ int, int, int ] w2;
10 [ void ] f( int, int, int );           // three input parameters of type int
11 [ void ] g( [ int, int, int ] );     // 3 element tuple as input
12 f( [ 1, 2, 3 ] );
13 f( w1, 3 );
14 f( 1, w1 );
15 f( w2 );
16 g( [ 1, 2, 3 ] );
17 g( w1, 3 );
18 g( 1, w1 );
19 g( w2 );
```

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

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

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

```
27 [ 1, 3, 5 ]
28 [ 1, (2, 3), 5 ]
```

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

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

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

34 is equivalent to:

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

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

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

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

44 21.1 Tuple Coercions

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

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

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

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

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

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

5 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:

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

8 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].

10 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:

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

```
14 int x = 5;
```

```
15 [ x, w ] = [ w, x ]; // all four tuple coercions
```

16 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).

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

23 21.2 Mass Assignment

24 CV permits assignment to several variables at once using mass assignment [27]. Mass assignment has the following form:

```
26 [ lvalue, ... , lvalue ] = expr;
```

27 The left-hand side is a tuple of *lvalues*, 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.

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

```
31 [ x, y, z ] = 1.5;
```

32 is equivalent to:

```
33 x = 1.5; y = 1.5; z = 1.5;
```

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

```
35 x = y = z = 1.5;
```

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

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

38 which is equivalent to:

```
39 t = a + b;
```

```
40 a1 = &i; a2 = &y[i]; a3 = &z;
```

```
41 *a1 = t; *a2 = t; *a3 = t;
```

42 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.

45 21.3 Multiple Assignment

46 CV also supports the assignment of several values at once, known as multiple assignment [27, 16]. Multiple assignment has the following form:

47

```
1 [ lvalue, ... , lvalue ] = [ expr, ... , expr ];
```

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

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

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

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

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

```
10 [ x, y ] = [ y, x ];
```

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

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

```
13 [ a, b, c ] = [ 1, 2 ];
```

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

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

```
17 f( x++, x++ );
```

```
// C routine call with side effects in arguments
```

```
18 [ v1, v2 ] = [ x++, x++ ];
```

```
// side effects in right-hand side of multiple assignment
```

19 21.4 Cascade Assignment

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

```
22 tuple = tuple = ... = tuple;
```

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

```
25 x1 = y1 = x2 = y2 = 0;
```

```
26 [ x1, y1 ] = [ x2, y2 ] = [ x3, y3 ];
```

```
27 [ x1, y1 ] = [ x2, y2 ] = 0;
```

```
28 [ x1, y1 ] = z = 0;
```

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

30 22 Stream I/O Library

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

35 CV formatting incorporates ideas from C printf, C++ stream manipulators, and Python implicit spacing and newline. Specifically:

- 37 • printf/Python format codes are dense, making them difficult to read and remember. CV/C++ format manipulators are named, making them easier to read and remember.
- 39 • printf/Python separate format codes from associated variables, making it difficult to match codes with variables. CV/C++ co-locate codes with associated variables, where CV has the tighter binding.
- 41 • Format manipulators in printf/Python/CV have local effect, whereas C++ have global effect, except setw. Hence, it is common C++ programming practice to toggle manipulators on and then back to the default to prevent downstream side-effects. Without this programming style, errors occur when moving prints, as manipulator effects incorrectly flow into the new location. Furthermore, to guarantee no side-effects, manipulator values must be saved and restored across function calls. C++ programmers never do any of this.
- 46 • CV has more sophisticated implicit value spacing than Python, plus implicit newline at the end of a print.

1 22.1 Basic I/O

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

5 22.1.1 Stream Output

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

| | CV | C++ | Python |
|---|---------------------------------------|---|----------------------------------|
| 7 | <code>int x = 1, y = 2, z = 3;</code> | | <code>x = 1; y = 2; z = 3</code> |
| | <code>sout x y z;</code> | <code>cout << x << " " << y << " " << z << endl;</code> | <code>print(x, y, z)</code> |
| | <code>1_2_3</code> | <code>1_2_3</code> | <code>1_2_3</code> |

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

| | | |
|----|---|------------------------------|
| 11 | <code>[int, [int, int]] t1 = [1, [2, 3]], t2 = [4, [5, 6]];</code> | |
| 12 | <code>sout t1 t2;</code> | <code>// print tuples</code> |
| 13 | <code>1_2_3_4_5_6</code> | |

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

| | | |
|----|-------------|---|
| | CV: | <code>sout x * 3 y + 1 z << 2 x == y (x y) (x y) (x > z ? 1 : 2);</code> |
| 17 | C++: | <code>cout << x * 3 << y + 1 << (z << 2) << (x == y) << (x y) << (x y) << (x > z ? 1 : 2) << endl;</code> |
| | | <code>3_3_12_0_3_1_2</code> |

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

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

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

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

26 Note, CV stream variables `stdin`, `stdout`, `stderr`, `exit`, and `abort` overload C variables `stdin`, `stdout`, `stderr`, and functions `exit`
 27 and `abort`, respectively.

28 22.1.2 Stream Input

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

| | CV | C++ | Python |
|----|--------------------------------------|--|---|
| | <code>char c; int i; double d</code> | | |
| 31 | <code>sin c i d;</code> | <code>cin >> c >> i >> d;</code> | <code>c = input(); i = int(input()); d = float(input());</code> |
| | <code>A_1_2.5</code> | <code>A_1_2.5</code> | <code>A</code> <code>1</code> <code>2.5</code> |

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

- 35 • 1-9 prefix introduces a decimal value (0-9),
- 36 • 0 prefix introduces an octal value (0-7), and
- 37 • 0x or 0X prefix introduces a hexadecimal value (0-f) with lower or upper case letters.

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

```
5 char ch; int i; float f; double d; _Complex double cxd;
6 sin | ch | i | f | d | cxd;
7 X 42 1234.5 0xffff-2 3.5+7.1i
```

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

```
9 sin | "abc def"; // space matches arbitrary whitespace (2 blanks, 2 tabs)
10 abc_ _ _ _ _ def
```

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

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

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

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

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

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

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

28 22.1.3 Stream Files

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

- 30 • `fail` tests the stream error-indicator, returning nonzero if it is set.
- 31 • `clear` resets the stream error-indicator.
- 32 • `flush` (ofstream only) causes any unwritten data for a stream to be written to the file.
- 33 • `eof` (ifstream only) tests the end-of-file indicator for the stream pointed to by stream. Returns true if the end-of-
- 34 file indicator is set, otherwise false.
- 35 • `open` binds the file with name to a stream accessed with mode (see `fopen`).
- 36 • `close` flushes the stream and closes the file.
- 37 • `write` (ofstream only) writes size bytes to the stream. The bytes are written lazily when an internal buffer fills.
- 38 Eager buffer writes are done with `flush`
- 39 • `read` (ifstream only) reads size bytes from the stream.
- 40 • `ungetc` (ifstream only) pushes the character back to the input stream. Pushed-back characters returned by subse-
- 41 quent reads in the reverse order of pushing.

42 The constructor functions:

- 43 • create an unbound stream, which is subsequently bound to a file with `open`.
- 44 • create a bound stream to the associated file with given mode.

45 The destructor closes the stream.

46 Figure 9, p. 56 demonstrates the file operations by showing the idiomatic C++ command-line processing and copying
 47 an input file to an output file. Note, a stream variable may be copied because it is a reference to an underlying stream
 48 data-structures. **All unusual I/O cases are handled as exceptions, including end-of-file.**

```

// ***** 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 22.2 Implicit Separator

2 The implicit separator character (space/blank) is a separator not a terminator for output. The rules for implicitly adding
3 a separator are:

4 1. A separator does not appear at the start or end of a line.

```
5     sout | 1 | 2 | 3;
6     1_2_3
```

7 2. A separator does not appear before or after a character literal or variable.

```
8     sout | '1' | '2' | '3';
9     123
```

10 3. A separator does not appear before or after a null (empty) C string, which is a local mechanism to disable
11 insertion of the separator character.

```
12     sout | 1 | "" | 2 | "" | 3;
13     123
```

14 4. A separator does not appear before a C string starting with the (extended) ASCII characters: , . ; ! ?)] } % < > ,
15 where » is a closing citation mark.

```
16     sout | 1 | ", x" | 2 | ". x" | 3 | "; x" | 4 | "! x" | 5 | "? x" | 6 | "% x"
17     | 7 | "ç x" | 8 | "» x" | 9 | ")" x" | 10 | "]" x" | 11 | "}" x";
```

```
18     Input1_2_3_4!_5?_6%_7ç_8»_9)_10]_11}_x
```

19 5. A separator does not appear after a C string ending with the (extended) ASCII characters: ([{ = \$ £ ¥ ; ¿ «, where
20 ¿ ¿ are inverted opening exclamation and question marks, and « is an opening citation mark.

```
21     sout | "x (" | 1 | "x [" | 2 | "x {" | 3 | "x =" | 4 | "x $" | 5 | "x £" | 6 | "x ¥"
22     | 7 | "x ;" | 8 | "x ¿" | 9 | "x «" | 10;
```

```
23     x_(1_x_[2_x_{3_x_=4_x_$5_x_£6_x_¥7_x_;8_x_¿9_x_«10
```

24 6. A separator does not appear before/after a C string starting/ending with the ASCII quote or whitespace charac-
25 ters: ` ' " : _ \ t \ v \ f \ r \ n

```
26     sout | "x`" | 1 | "x'" | 2 | "x\"" | 3 | "x:" | 4 | "x_" | 5 | "x\t" | 6 | "x\n";
```

```

#include <fstream.h>

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

```
1      x`1`x' 2' x"3"x:4;x_5x_6__x
```

2 7. If a space is desired before or after one of the special string start/end characters, explicitly insert a space.

```
3      sout | "x ( _ | | " _ ) x" | 2 | " _ , x" | 3 | " _ : x : _ " | 4;
```

```
4      x_( _ 1 _ ) _ x_2_ _ x_3_ : x : _ 4
```

5 22.3 Separation Manipulators

6 The following manipulators control implicit output separation. The effect of these manipulators is global for an output
7 stream (except sep and nosep).

8 1. sepSet and sepVal/sepGet set and get the separator string. The separator string can be at most 16 characters
9 including the '\0' string terminator (15 printable characters).

```
10      sepSet( sout, " , $" );                // set separator from " " to " , $"
11      sout | 1 | 2 | 3 | " \ " | sepVal | " \ " ;
12      1_ _ $ 2_ _ $ 3_ _ " , _ $ "

13      sepSet( sout, " " );                  // reset separator to " "
14      sout | 1 | 2 | 3 | " \ " | sepGet( sout ) | " \ " ;
15      1_ _ 2_ _ 3_ _ " _ "
```

16 sepGet can be used to store a separator and then restore it:

```

1      char store[sepSize];           // sepSize is the maximum separator size
2      strcpy( store, sepGet( sout ) ); // copy current separator
3      sepSet( sout, "_" );           // change separator to underscore
4      sout | 1 | 2 | 3;
5      1_2_3
6
6      sepSet( sout, store );         // change separator back to original
7      sout | 1 | 2 | 3;
8      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).

```

11     sepSetTuple( sout, " " );       // set tuple separator from ", " to " "
12     sout | t1 | t2 | " \ " | sepTupleVal | " \ " ;
13     1_2_3_4_5_6_ " \ "
14
14     sepSetTuple( sout, ", " );     // reset tuple separator to ", "
15     sout | t1 | t2 | " \ " | sepGetTuple( sout ) | " \ " ;
16     1_2_3_4_5_6_ ", "

```

As for `sepGet`, `sepGetTuple` can be used to store a tuple separator and then restore it.

3. `sepOff` and `sepOn` globally toggle printing the separator.

```

19     sout | sepOff | 1 | 2 | 3;      // turn off implicit separator
20     123
21
21     sout | sepOn | 1 | 2 | 3;       // turn on implicit separator
22     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.

```

25     sout | 1 | nosep | 2 | 3;       // turn off implicit separator for the next item
26     12_3
27
27     sout | sepOff | 1 | sep | 2 | 3; // turn on implicit separator for the next item
28     1_23

```

The tuple separator also responds to being turned on and off.

```

30     sout | t1 | nosep | t2;         // turn off implicit separator for the next item
31     1_2_3_4_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.

```

34     sout | sep | 1 | 2 | 3 | sep;    // sep does nothing at start/end of line
35     1_2_3
36
36     sout | sepVal | 1 | 2 | 3 | sepVal ; // use sepVal to print separator at start/end of line
37     _1_2_3_

```

38 22.4 Newline Manipulators

39 The following manipulators control newline separation for input and output.

40 For input:

- 41 1. `nlOn` reads the newline character, when reading single characters.
- 42 2. `nlOff` does *not* read the newline character, when reading single characters.
- 43 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.

45 For example, in:

```

46     int i, j;
47     sin | i | nl | j;
48     1 2
49     3

```

50 variable `i` is assigned 1, the 2 is skipped, and variable `j` is assigned 3. For example, in:

```

char ch
1   sin | ch; // read X      sin | nlOn; // enable reading newlines
                               sin | ch; // read newline
X

```

2 the left example skips the newline and reads 'x' into ch, while the right example reads the newline into ch.

3 For output:

- 4 1. nlOn implicitly prints a newline at the end of each output expression.
- 5 2. nlOff does *not* implicitly print a newline at the end of each output expression.
- 6 3. nl inserts a newline.

```

7   sout | nl;                // only print newline
8   sout | 2;                 // implicit newline
9   sout | 3 | nl | 4 | nl;    // terminating nl merged with implicit newline
10  sout | 5 | nl | nl;        // again terminating nl merged with implicit newline
11  sout | 6;                 // implicit newline
12
13  2
14  3
15  4
16  5
17
18  6

```

19 Note, a terminating nl is merged with (overrides) the implicit newline at the end of the sout expression, otherwise
20 it is impossible to print a single newline

21 22.5 Output Manipulators

22 The following manipulators control formatting (printing) of the argument output values.

- 23 1. bin(integer) print value in base 2 preceded by 0b/0B.

```

24   sout | bin( 0 ) | bin( 27HH ) | bin( 27H ) | bin( 27 ) | bin( 27L );
25   0b0 0b11011 0b11011 0b11011 0b11011
26   sout | bin( -27HH ) | bin( -27H ) | bin( -27 ) | bin( -27L );
27   0b11100101 0b11111111111100101 0b1111111111111111111111111111111100101 0b(58 1s)100101

```

- 28 2. oct(integer) print value in base 8 preceded by 0.

```

29   sout | oct( 0 ) | oct( 27HH ) | oct( 27H ) | oct( 27 ) | oct( 27L );
30   0 033 033 033 033
31   sout | oct( -27HH ) | oct( -27H ) | oct( -27 ) | oct( -27L );
32   0345 01777745 037777777745 0177777777777777777777777777777745

```

33 Note, octal 0 is *not* preceded by 0 to prevent confusion.

- 34 3. hex(integer / floating-point) print value in base 16 preceded by 0x/0X.

```

35   sout | hex( 0 ) | hex( 27HH ) | hex( 27H ) | hex( 27 ) | hex( 27L );
36   0x0 0x1b 0x1b 0x1b 0x1b
37   sout | hex( -27HH ) | hex( -27H ) | hex( -27 ) | hex( -27L );
38   0xe5 0xffe5 0xffffffe5 0xfffffffffffe5
39
40   sout | hex( 0.0 ) | hex( 27.5F ) | hex( 27.5 ) | hex( 27.5L );
41   0x0p+0 0x1.b8p+4 0x1.b8p+4 0xd.cp+1
42   sout | hex( -27.5F ) | hex( -27.5 ) | hex( -27.5L );
43   -0x1.b8p+4 -0x1.b8p+4 -0xd.cp+1

```

- 44 4. sci(floating-point) print value in scientific notation with exponent. Default is 6 digits of precision.

```

45   sout | sci( 0.0 ) | sci( 27.5 ) | sci( -27.5 );
46   0.000000e+00 2.750000e+01 -2.750000e+01

```

- 47 5. eng(floating-point) print value in engineering notation with exponent, which means the exponent is adjusted to
48 a multiple of 3.

```

49   sout | eng( 0.0 ) | eng( 27000.5 ) | eng( -27.5e7 );

```

```

1      0e0 27.0005e3 -275e6
2  6. unit( engineering-notation ) print engineering exponent as a letter between the range 10-24 and 1024:
3      y ⇒ 10-24, z ⇒ 10-21, a ⇒ 10-18, f ⇒ 10-15, p ⇒ 10-12, n ⇒ 10-9, u ⇒ 10-6, m ⇒ 10-3, K ⇒ 103, M
4      ⇒ 106, G ⇒ 109, T ⇒ 1012, P ⇒ 1015, E ⇒ 1018, Z ⇒ 1021, Y ⇒ 1024.
5      For exponent 100, no decimal point or letter is printed.
6      sout | unit(eng( 0.0 )) | unit(eng( 27000.5 )) | unit(eng( -27.5e7 ));
7      0 27.0005K -275M
8  7. upcase( bin / hex / floating-point ) print letters in a value in upper case. Lower case is the default.
9      sout | upcase( bin( 27 ) ) | upcase( hex( 27 ) ) | upcase( 27.5e-10 ) | upcase( hex( 27.5 ) );
10     0B11011 0X1B 2.75E-09 0X1.B8P+4
11  8. nobase( integer ) do not precede bin, oct, hex with 0b/0B, 0, or 0x/0X. Printing the base is the default.
12     sout | nobase( bin( 27 ) ) | nobase( oct( 27 ) ) | nobase( hex( 27 ) );
13     11011 33 1b
14  9. nodp( floating-point ) do not print a decimal point if there are no fractional digits. Printing a decimal point is the
15     default, if there are no fractional digits.
16     sout | 0. | nodp( 0. ) | 27.0 | nodp( 27.0 ) | nodp( 27.5 );
17     0.0 0 27.0 27.5
18  10. sign( integer / floating-point ) prefix with plus or minus sign (+ or -). Only printing the minus sign is the default.
19     sout | sign( 27 ) | sign( -27 ) | sign( 27. ) | sign( -27. ) | sign( 27.5 ) | sign( -27.5 );
20     +27 -27 +27.0 -27.0 +27.5 -27.5
21  11. wd( minimum, value ), wd( minimum, precision, value ) For all types, minimum is the number of printed char-
22     acters. If the value is shorter than the minimum, it is padded on the right with spaces.
23     sout | wd( 4, 34 ) | wd( 3, 34 ) | wd( 2, 34 );
24     sout | wd( 10, 4. ) | wd( 9, 4. ) | wd( 8, 4. );
25     sout | wd( 4, "ab" ) | wd( 3, "ab" ) | wd( 2, "ab" );
26     _34_34_34
27     _4.000000_4.000000_4.000000
28     _ab_ab_ab
29     If the value is larger, it is printed without truncation, ignoring the minimum.
30     sout | wd( 4, 34567 ) | wd( 3, 34567 ) | wd( 2, 34567 );
31     sout | wd( 4, 3456. ) | wd( 3, 3456. ) | wd( 2, 3456. );
32     sout | wd( 4, "abcde" ) | wd( 3, "abcde" ) | wd( 2, "abcde" );
33     34567_34567_34567
34     3456_3456_3456.
35     abcde_abcde_abcde
36     For integer types, precision is the minimum number of printed digits. If the value is shorter, it is padded on
37     the left with leading zeros.
38     sout | wd( 4,3, 34 ) | wd( 8,4, 34 ) | wd( 10,10, 34 );
39     _034_0034_000000034
40     If the value is larger, it is printed without truncation, ignoring the precision.
41     sout | wd( 4,1, 3456 ) | wd( 8,2, 3456 ) | wd( 10,3, 3456 );
42     3456_3456_3456
43     If precision is 0, nothing is printed for zero. If precision is greater than the minimum, it becomes the minimum.
44     sout | wd( 4,0, 0 ) | wd( 3,10, 34 );
45     _000000034
46     For floating-point types, precision is the minimum number of digits after the decimal point.
47     sout | wd( 6,3, 27.5 ) | wd( 8,1, 27.5 ) | wd( 8,0, 27.5 ) | wd( 3,8, 27.5 );
48     27.500_27.5_28_27.5000000
49     For the C-string type, precision is the maximum number of printed characters, so the string is truncated if it
50     exceeds the maximum.

```


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, `CV` 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 skipped 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.

```

10 char cs[10];
11 sin | wdi( sizeof(cs), cs );
12 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* \leq *maximum size*.

```

17 char cs[10];
18 sin | wdi( sizeof(cs), 9, cs );
19 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).

```

25 char cs[10];
26 sin | getline( wdi( sizeof(cs), cs ) );
27 sin | getline( wdi( sizeof(cs), cs ), 'X' ); // X is the line delimiter
28 abc_??_#@%
29 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 delimiter character is omitted, it defaults to `'\ '` (single quote).

```

35 char ch;
36 sin | quote( ch ); sin | quote( ch, ' ' ); sin | quote( ch, '[', ']' );
37 _ _ _ '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 delimiter character is omitted, it defaults to `' ' ' ' (double quote).`

```

43 char cs[10];
44 sin | quote( wdi( sizeof(cs), cs ) ); // " is the start/end delimiter
45 sin | quote( wdi( sizeof(cs), cs ), '\ ' ); // ' is the start/end delimiter
46 sin | quote( wdi( sizeof(cs), cs ), '[', ']' ); // [ is the start and ] is the end delimiter
47 _ _ _ "abc" _ _ 'abc'[abc]

```

8. `incl(scanset, wdi manipulator)` consumes the scanset, which reads all the scanned characters into the string variable (null terminated).

```

50 char cs[10];
51 sin | incl( "abc", cs );
52 bcxyz

```

```

1  9. excl( scanset, wdi manipulator ) consumes the not scanset, which reads all the scanned characters into the string
2  variable (null terminated).
3      char cs[10];
4      sin | excl( "abc", cs );
5      xyzbca
6  10. ignore( T & v or const char cs[] or string manipulator ) consumes the appropriate characters for the type and
7  ignores them, so the input variable is unchanged.
8      double d;
9      char cs[10];
10     sin | ignore( d );           // d is unchanged
11     sin | ignore( cs );        // cs is unchanged, no wdi required
12     sin | ignore( quote( wdi( sizeof(cs), cs ) ) ); // cs is unchanged
13     _-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:

```

17  thread1 : sout | "abc " | "def ";
18  thread2 : sout | "uvw " | "xyz ";

```

possible outputs are:

```

20  abc def   |   abc uvw xyz   |   uvw abc xyz def   |   abuvwc dext   |   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_V 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.*:

```

30  thread1 : mutex( sout ) sout | "abc " | "def ";
31  thread2 : 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.

```

34  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, the long form of the `mutex` statement is used, *e.g.*:

```

38  mutex( sout ) {
39     sout | 1;
40     mutex( sout ) sout | 2 | 3;           // unnecessary, but ok because of recursive lock
41     sout | 4;
42 } // 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:

```

47  mutex( sin ) {
48     int x, y, z, w;
49     sin | x;

```

```

1      mutex( sin ) sin | y | z;           // unnecessary, but ok because of recursive lock
2      sin | w;
3      } // implicitly release sin lock

```

4 Again, the order of the reading threads is non-deterministic. Note, non-deterministic reading is rare.

5 **WARNING:** The general problem of nested locking can occur if routines are called in an I/O sequence that block,
6 *e.g.*:

```

7      mutex( sout ) sout | "data: " | rtn( mon ); // mutex call on monitor

```

8 If the thread executing the I/O expression blocks in the monitor with the sout lock, other threads writing to sout also
9 block until the thread holding the lock is unblocked and releases it. This scenario can lead to deadlock, if the thread
10 that is going to unblock the thread waiting in the monitor first writes to sout (deadly embrace). To prevent nested
11 locking, a simple precaution is to factor out the blocking call from the expression, *e.g.*:

```

12     int data = rtn( mon );
13     mutex( sout ) sout | "data: " | data;

```

14 22.8 Locale

15 Cultures use different syntax, called a *locale*, for printing numbers so they are easier to read, *e.g.*:

```

16     12,345.123      // comma separator, period decimal-point
17     12.345,123     // period separator, comma decimal-point
18     12_345,123.    // space separator, comma decimal-point, period terminator

```

19 A locale is selected with function `setlocale`, and the corresponding locale package *must* be installed on the underlying
20 system; `setlocale` returns 0p if the requested locale is unavailable. Furthermore, a locale covers the syntax for many
21 cultural items, *e.g.*, address, measurement, money, etc. This discussion applies to item LC_NUMERIC for formatting
22 non-monetary integral and floating-point values. Figure 10 shows selecting different cultural syntax, which may be
23 associated with one or more countries.

24 23 String Stream

25 The stream types `ostrstream` and `istrstream` provide all the stream formatting capabilities to/from a C string rather than
26 a stream file. Figure 11, p. 65 shows writing (output) to and reading (input) from a C string. The only string stream
27 operations different from a file stream are:

- 28 • constructors to create a stream that writes to a write buffer (`ostrstream`) of size, or reads from a read buffer
29 (`istrstream`) containing a C string terminated with `'\0'`.

```

30     void ?{?( ostrstream &, char buf[], size_t size );
31     void ?{?( istrstream & is, char buf[] );

```

- 32 • write (`ostrstream` only) writes all the buffered characters to the specified stream (stdout default).

```

33     ostrstream & write( ostrstream & os, FILE * stream = stdout );

```

34 There is no read for `istrstream`.

35 24 Structures

36 Structures in C++ are basically the same as structures in C. A structure is defined with the same syntax as in C. When
37 referring to a structure in C++, users may omit the `struct` keyword.

```

38     struct Point {
39         double x;
40         double y;
41     };
42
43     Point p = {0.0, 0.0};

```

44 C++ does not support inheritance among types, but instead uses composition to enable reuse of structure fields.
45 Composition is achieved by embedding one type into another. When type A is embedded in type B, an object with
46 type B may be used as an object of type A, and the fields of type A are directly accessible. Embedding types is
47 achieved using anonymous members. For example, using `Point` from above:

```

#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 void foo(Point p);
2
3 struct ColoredPoint {
4     Point; // anonymous member (no identifier)
5     int Color;
6 };
7 ...
8 ColoredPoint cp = ...;
9 cp.x = 10.3; // x from Point is accessed directly
10 cp.color = 0x33aaff; // color is accessed normally
11 foo(cp); // cp can be used directly as a Point

```

```

#include <fstream.h>
#include <sstream.h>

int main() {
    enum { size = 256 };
    char buf[size];           // output buffer
    ostream 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
    istream isstr = { buf2 };
    char s[10];
    isstr | i | j | k | x | y | s;
    sout | i | j | k | x | y | s;
}

3_0x5_12_14_15_3.5_70000_abc
3_0x5_12_14_15_3.5_70000_abc
3_0x5_12_14_15_3.5_70000_abc
12_14_15_3.5_70000_abc

```

Figure 11: String Stream Processing

1 25 Constructors and Destructors

- 2 C++ supports C initialization of structures, but it also adds constructors for more advanced initialization. Additionally,
- 3 C++ adds destructors that are called when a variable is deallocated (variable goes out of scope or object is deleted).
- 4 These functions take a reference to the structure as a parameter (see Section 12, p. 20 for more information).

5 26 Overloading

- 6 Overloading refers to the capability of a programmer to define and use multiple objects in a program with the same
- 7 name. In C++, a declaration may overload declarations from outer scopes with the same name, instead of hiding
- 8 them as is the case in C. This may cause identical C and C++ programs to behave differently. The compiler selects
- 9 the appropriate object (overload resolution) based on context information at the place where it is used. Overloading
- 10 allows programmers to give functions with different signatures but similar semantics the same name, simplifying the
- 11 interface to users. Disadvantages of overloading are that it can be used to give functions with different semantics the
- 12 same name, causing confusion, or that the compiler may resolve to a different function from what the programmer
- 13 expected. C++ allows overloading of functions, operators, variables, and even the constants 0 and 1.

- 14 The compiler follows some overload resolution rules to determine the best interpretation of all of these overloads.
- 15 The best valid interpretations are the valid interpretations that use the fewest unsafe conversions. Of these, the best
- 16 are those where the functions and objects involved are the least polymorphic. Of these, the best have the lowest total
- 17 conversion cost, including all implicit conversions in the argument expressions. Of these, the best have the highest
- 18 total conversion cost for the implicit conversions (if any) applied to the argument expressions. If there is no single best
- 19 valid interpretation, or if the best valid interpretation is ambiguous, then the resulting interpretation is ambiguous. For
- 20 details about type inference and overload resolution, please see the C++ Language Specification.

```

21 int foo(int a, int b) {
22     float sum = 0.0;
23     float special = 1.0;
24     {
25         int sum = 0;

```

```

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           // both the float and int versions of sum are available
2           float special = 4.0;
3           // this inner special hides the outer version
4           ...
5           }
6           ...
7           }

```

8 26.1 Constant

9 The constants 0 and 1 have special meaning. In C#, as in C, all scalar types can be incremented and decremented,
10 which is defined in terms of adding or subtracting 1. The operations &&, ||, and ! can be applied to any scalar arguments
11 and are defined in terms of comparison against 0 (e.g., (a && b) becomes (a != 0 && b != 0)).

12 In C, the integer constants 0 and 1 suffice because the integer promotion rules can convert them to any arithmetic
13 type, and the rules for pointer expressions treat constant expressions evaluating to 0 as a special case. However, user-
14 defined arithmetic types often need the equivalent of a 1 or 0 for their functions or operators, polymorphic functions

1 often need 0 and 1 constants of a type matching their polymorphic parameters, and user-defined pointer-like types may
 2 need a null value. Defining special constants for a user-defined type is more efficient than defining a conversion to the
 3 type from **bool**.

4 Why just 0 and 1? Why not other integers? No other integers have special status in C. A facility that let program-
 5 mers declare specific constants **const** Rational 12, for instance. would not be much of an improvement. Some facility
 6 for defining the creation of values of programmer-defined types from arbitrary integer tokens would be needed. The
 7 complexity of such a feature does not seem worth the gain.

8 For example, to define the constants for a complex type, the programmer would define the following:

```

9  struct Complex {
10     double real;
11     double imaginary;
12 }
13
14 const Complex 0 = {0, 0};
15 const Complex 1 = {1, 0};
16 ...
17
18     Complex a = 0;
19 ...
20
21     a++;
22 ...
23     if (a) { // same as if (a == 0)
24 ...
25 }
```

26 26.2 Variable

27 The overload rules of CV allow a programmer to define multiple variables with the same name, but different types.
 28 Allowing overloading of variable names enables programmers to use the same name across multiple types, simplifying
 29 naming conventions and is compatible with the other overloading that is allowed. For example, a developer may want
 30 to do the following:

```

31 int pi = 3;
32 float pi = 3.14;
33 char pi = .p.;
```

34 26.3 Function Overloading

35 Overloaded functions in CV are resolved based on the number and type of arguments, type of return value, and the
 36 level of specialization required (specialized functions are preferred over generic).

37 The examples below give some basic intuition about how the resolution works.

```

38 // Choose the one with less conversions
39 int doSomething(int value) {...} // option 1
40 int doSomething(short value) {...} // option 2
41
42 int a, b = 4;
43 short c = 2;
44
45 a = doSomething(b); // chooses option 1
46 a = doSomething(c); // chooses option 2
47
48 // Choose the specialized version over the generic
49
50 generic(type T)
51 T bar(T rhs, T lhs) {...} // option 3
52 float bar(float rhs, float lhs){...} // option 4
53 float a, b, c;
```

```

1  double d, e, f;
2  c = bar(a, b); // chooses option 4
3
4  // specialization is preferred over unsafe conversions
5
6  f = bar(d, e); // chooses option 5

```

26.4 Operator

CV 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. CV 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 1: 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.

```

16  type Complex = struct { // define a Complex type
17      double real;
18      double imag;
19  }
20
21  // Constructor with default values
22
23  void ?{}(Complex &c, double real = 0.0, double imag = 0.0) {
24      c.real = real;
25      c.imag = imag;
26  }
27
28  Complex ?+?(Complex lhs, Complex rhs) {
29      Complex sum;
30      sum.real = lhs.real + rhs.real;
31      sum.imag = lhs.imag + rhs.imag;
32      return sum;
33  }
34
35  String ()?(const Complex c) {
36      // use the string conversions for the structure members

```

```

1     return (String)c.real + . + . + (String)c.imag + .i.;
2   }
3   ...
4
5   Complex a, b, c = {1.0}; // constructor for c w/ default imag
6   ...
7   c = a + b;
8   print(.sum = . + c);

```

9 27 Auto Type-Inferencing

10 Auto type-inferencing occurs in a declaration where a variable's type is inferred from its initialization expression type.

| | C++ | gcc | |
|----|--------------------------|--------------------------------|--|
| 11 | auto j = 3.0 * 4; | #define expr 3.0 * i | |
| | int i; | typeof (expr) j = expr; | // use type of initialization expression |
| | auto k = i; | typeof (i) k = i; | // use type of primary variable |

12 The two important capabilities are:

- 13 • not determining or writing long generic types,
- 14 • ensuring secondary variables, related to a primary variable, always have the same type.

15 In C \forall , **typedef** provides a mechanism to alias long type names with short ones, both globally and locally, but not
 16 eliminate the use of the short name. gcc provides **typeof** to declare a secondary variable from a primary variable.
 17 C \forall also relies heavily on the specification of the left-hand side of assignment for type inferencing, so in many cases
 18 it is crucial to specify the type of the left-hand side to select the correct type of the right-hand expression. Only
 19 for overloaded routines *with the same return type* is variable type-inferencing possible. Finally, **auto** presents the
 20 programming problem of tracking down a type when the type is actually needed. For example, given

```
21 auto j = ...
```

22 and the need to write a routine to compute using j

```
23 void rtn( ... parm );
24 rtn(j);
```

25 A programmer must work backwards to determine the type of j's initialization expression, reconstructing the possibly
 26 long generic type-name. In this situation, having the type name or a short alias is very useful.

27 There is also the conundrum in type inferencing of when to *brand* a type. That is, when is the type of the variable
 28 more important than the type of its initialization expression. For example, if a change is made in an initialization
 29 expression, it can cause cascading type changes and/or errors. At some point, a variable type needs to remain constant
 30 and the expression to be in error when it changes.

31 Given **typedef** and **typeof** in C \forall , and the strong need to use the type of left-hand side in inferencing, auto type-
 32 inferencing is not supported at this time. Should a significant need arise, this feature can be revisited.

33 28 Concurrency

34 Concurrency support in C \forall is implemented on top of a highly efficient runtime system of light-weight, M:N, user
 35 level threads. The model integrates concurrency features into the language by making the structure type the core
 36 unit of concurrency. All communication occurs through method calls, where data is sent via method arguments, and
 37 received via the return value. This enables a very familiar interface to all programmers, even those with no parallel
 38 programming experience. It also allows the compiler to do static type checking of all communication, a very important
 39 safety feature. This controlled communication with type safety has some similarities with channels in Go, and can
 40 actually implement channels exactly, as well as create additional communication patterns that channels cannot. Mutex
 41 objects, monitors, are used to contain mutual exclusion within an object and synchronization across concurrent threads.

42 28.1 Coroutine

43 Coroutines are the precursor to threads. Figure 13 shows a coroutine that computes the Fibonacci numbers.

```

#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 28.2 Monitors

2 A monitor is a structure in CV which includes implicit locking of its fields. Users of a monitor interact with it just like
3 any structure, but the compiler handles code as needed to ensure mutual exclusion. An example of the definition of a
4 monitor is shown here:

```

5 type Account = monitor {
6     const unsigned long number; // account number
7     float balance; // account balance
8 };

```

9 28.3 Threads

10 CV also provides a simple mechanism for creating and utilizing user level threads. A thread provides mutual exclusion
11 like a monitor, and also has its own execution state and a thread of control. Similar to a monitor, a thread is defined
12 like a structure:

13 29 Language Comparisons

14 CV is one of many languages that attempts to improve upon C. In developing CV, many other languages were consulted
15 for ideas, constructs, and syntax. Therefore, it is important to show how these languages each compare with Do. In
16 this section, CV is compared with what the writers of this document consider to be the closest competitors of Do: C++,
17 Go, Rust, and D.

18 29.1 C++

19 C++ is a general-purpose programming language. It is an imperative, object-oriented and generic programming lan-
20 guage, while also providing facilities for low-level memory manipulation. The primary focus of C++ was adding

```

#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 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 and articles comparing inheritance and composition (or is-a versus has-a relationships), so we will not go into more detail here (Venners, 1998) (Pike, Go at Google: Language Design in the Service of Software Engineering, 2012).

7 Overloading in C# is very similar to overloading in C++, with the exception of the additional use, in C#, of the return type to differentiate between overloaded functions. References and exceptions in C# are heavily based on the same features from C++. The mechanism for interoperating with C code in C# is also borrowed from C++.

10 Both C# and C++ provide generics, and the syntax is quite similar. The key difference between the two, is that in C++ templates are expanded at compile time for each type for which the template is instantiated, while in C#, function pointers are used to make the generic fully compilable. This means that a generic function can be defined in a compiled library, and still be used as expected from source.

14 29.2 Go

15 Go, also commonly referred to as golang, is a programming language developed at Google in 2007 [19]. It is a statically typed language with syntax loosely derived from that of C, adding garbage collection, type safety, some structural typing capabilities, additional built-in types such as variable-length arrays and key-value maps, and a large standard library. (Wikipedia)

19 Go and C# differ significantly in syntax and implementation, but the underlying core concepts of the two languages are aligned. Both Go and C# use composition and interfaces as opposed to inheritance to enable encapsulation and abstraction. Both languages (along with their tooling ecosystem) provide a simple packaging mechanism for building units of code for easy sharing and reuse. Both languages also include built-in light weight, user level threading

```

#include <fstream.h>
#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 concurrency features that attempt to simplify the effort and thought process required for writing parallel programs
 2 while maintaining high performance.

3 Go has a significant runtime which handles the scheduling of its light weight threads, and performs garbage col-
 4 lection, among other tasks. C \forall uses a cooperative scheduling algorithm for its tasks, and uses automatic reference
 5 counting to enable advanced memory management without garbage collection. This results in Go requiring significant
 6 overhead to interface with C libraries while C \forall has no overhead.

7 29.3 Rust

8 Rust is a general-purpose, multi-paradigm, compiled programming language developed by Mozilla Research. It is
 9 designed to be a "safe, concurrent, practical language", supporting pure-functional, concurrent-actor[dubious . dis-
 10 cuss][citation needed], imperative-procedural, and object-oriented styles.

11 The primary focus of Rust is in safety, especially in concurrent programs. To enforce a high level of safety, Rust
 12 has added ownership as a core feature of the language to guarantee memory safety. This safety comes at the cost of a
 13 difficult learning curve, a change in the thought model of the program, and often some runtime overhead.

1 Aside from those key differences, Rust and C_V also have several similarities. Both languages support no overhead
 2 interoperability with C and have minimal runtimes. Both languages support inheritance and polymorphism through
 3 the use of interfaces (traits).

4 29.4 D

5 The D programming language is an object-oriented, imperative, multi-paradigm system programming language cre-
 6 ated by Walter Bright of Digital Mars and released in 2001. [.] Though it originated as a re-engineering of C++, D
 7 is a distinct language, having redesigned some core C++ features while also taking inspiration from other languages,
 8 notably Java, Python, Ruby, C#, and Eiffel.

9 D and C_V both start with C and add productivity features. The obvious difference is that D uses classes and
 10 inheritance while C_V uses composition and interfaces. D is closer to C_V than C++ since it is limited to single inheritance
 11 and also supports interfaces. Like C++, and unlike C_V, D uses garbage collection and has compile-time expanded
 12 templates. D does not have any built-in concurrency constructs in the language, though it does have a standard library
 13 for concurrency which includes the low-level primitives for concurrency.

14 A Syntax Ambiguities

15 C has a number of syntax ambiguities, which are resolved by taking the longest sequence of overlapping characters
 16 that constitute a token. For example, the program fragment `x+++++y` is parsed as `x_++_+_+_+_y` because operator tokens
 17 `++` and `+` overlap. Unfortunately, the longest sequence violates a constraint on increment operators, even though the
 18 parse `x_+_+_+_+_+_y` might yield a correct expression. Hence, C programmers are aware that spaces have to added to
 19 disambiguate certain syntactic cases.

20 In C_V, there are ambiguous cases with dereference and operator identifiers, *e.g.*, `int *?*?()`, where the string `*?*?`
 21 can be interpreted as:

```
22 *?_ *?           // dereference operator, dereference operator
23 *_?*?          // dereference, multiplication operator
```

24 By default, the first interpretation is selected, which does not yield a meaningful parse. Therefore, C_V does a lexical
 25 look-ahead for the second case, and backtracks to return the leading unary operator and reparses the trailing operator
 26 identifier. Otherwise a space is needed between the unary operator and operator identifier to disambiguate this common
 27 case.

28 A similar issue occurs with the dereference, `*?(...)`, and routine-call, `?()(...)` identifiers. The ambiguity occurs when
 29 the dereference operator has no parameters:

```
30 *?()_... ;
31 *?()_...(...);
```

32 requiring arbitrary whitespace look-ahead for the routine-call parameter-list to disambiguate. However, the dereference
 33 operator *must* have a parameter/argument to dereference `*?(...)`. Hence, always interpreting the string `*?()` as `*_?()` does
 34 not preclude any meaningful program.

35 The remaining cases are with the increment/decrement operators and conditional expression, *e.g.*:

```
36 i++?_...(...);
37 i?+_...(...);
```

38 requiring arbitrary whitespace look-ahead for the operator parameter-list, even though that interpretation is an incorrect
 39 expression (juxtaposed identifiers). Therefore, it is necessary to disambiguate these cases with a space:

```
40 i++_? i : 0;
41 i?_++i : 0;
```

42 B C Incompatibles

43 The following incompatibles exist between C_V and C, and are similar to Annex C for C++ [23].

44 1. **Change:** add new keywords

45 New keywords are added to C_V (see Section C, p. 76).

46 **Rationale:** keywords added to implement new semantics of C_V.

Effect on original feature: change to semantics of well-defined feature.

Any C11 programs using these keywords as identifiers are invalid C \forall programs.

Difficulty of converting: keyword clashes are accommodated by syntactic transformations using the C \forall back-quote escape-mechanism (see Section 6, p. 5).

How widely used: clashes among new C \forall 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.*:

```

9         x;                // int x
10        *y;               // int *y
11        f( p1, p2 );      // int f( int p1, int p2 );
12        g( p1, p2 ) int p1, p2; // int g( int p1, int p2 );
```

C \forall continues to support K&R routine definitions:

```

14        f( a, b, c )      // default int return
15        int a, b; char c; // K&R parameter declarations
16        {
17        ...
18        }
```

Rationale: dropped from C11 standard.²⁰

Effect on original feature: original feature is deprecated.

Any old C programs using these K&R declarations are invalid C \forall programs.

Difficulty of converting: trivial to convert to C \forall .

How widely used: existing usages are rare.

3. **Change:** type of character literal **int** to **char** to allow more intuitive overloading:

```

25        int rtn( int i );
26        int rtn( char c );
27        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.

```

31        sout | 'x' | " " | (int) 'x';
32        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 longer work the same in C \forall 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**:

```

40        char * p = "abc"; // valid in C, deprecated in C $\forall$ 
41        char * q = expr ? "abc" : "de"; // valid in C, invalid in C $\forall$ 
```

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:

```

45        char * p = "abc";
46        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.

²⁰At 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 [22, § 6.7.2(2)]

1 **Effect on original feature:** change to semantics of well-defined feature.

2 **Difficulty of converting:** simple syntactic transformation, because string literals can be converted to `char *`.

3 **How widely used:** programs that have a legitimate reason to treat string literals as pointers to potentially mod-
4 ifiable memory are rare.

5 5. **Change:** remove *tentative definitions*, which only occurs at file scope:

```
6     int i;                // forward definition
7     int *j = &i;         // forward reference, valid in C, invalid in C\
8     int i = 0;           // definition
```

9 is valid in C, and invalid in C\ because duplicate overloaded object definitions at the same scope level
10 are disallowed. This change makes it impossible to define mutually referential file-local static objects, if
11 initializers are restricted to the syntactic forms of C. For example,

```
12     struct X { int i; struct X *next; };
13     static struct X a;    // forward definition
14     static struct X b = { 0, &a }; // forward reference, valid in C, invalid in C\
15     static struct X a = { 1, &b }; // definition
```

16 **Rationale:** avoids having different initialization rules for builtin types and user-defined types.

17 **Effect on original feature:** change to semantics of well-defined feature.

18 **Difficulty of converting:** the initializer for one of a set of mutually-referential file-local static objects must
19 invoke a routine call to achieve the initialization.

20 **How widely used:** seldom

21 6. **Change:** have `struct` introduce a scope for nested types:

```
22     enum Colour { R, G, B, Y, C, M };
23     struct Person {
24         enum Colour { R, G, B }; // nested type
25         struct Face {           // nested type
26             Colour Eyes, Hair; // type defined outside (1 level)
27         };
28         Colour shirt;          // type defined outside (top level)
29         Colour pants;          // type defined same level
30         Face looks[10];        // type defined same level
31     };
32     Colour c = R;              // type/enum defined same level
33     Person.Colour pc = Person.R; // type/enum defined inside
34     Person.Face pretty;        // type defined inside
```

35 In C, the name of the nested types belongs to the same scope as the name of the outermost enclosing structure,
36 *i.e.*, the nested types are hoisted to the scope of the outer-most type, which is not useful and confusing. C\
37 is C *incompatible* on this issue, and provides semantics similar to C++. Nested types are not hoisted and can
38 be referenced using the field selection operator “.”, unlike the C++ scope-resolution operator “::”.

39 **Rationale:** `struct` scope is crucial to C\ as an information structuring and hiding mechanism.

40 **Effect on original feature:** change to semantics of well-defined feature.

41 **Difficulty of converting:** Semantic transformation.

42 **How widely used:** C programs rarely have nest types because they are equivalent to the hoisted version.

43 7. **Change:** In C++, the name of a nested class is local to its enclosing class.

44 **Rationale:** C++ classes have member functions which require that classes establish scopes.

45 **Difficulty of converting:** Semantic transformation. To make the struct type name visible in the scope of the
46 enclosing struct, the struct tag could be declared in the scope of the enclosing struct, before the enclosing
47 struct is defined. Example:

```
48     struct Y;                // struct Y and struct X are at the same scope
49     struct X {
50         struct Y { /* ... */ } y;
51     };
```

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:

```
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 \forall 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 \forall Keywords

C \forall 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 \forall 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 [22, § 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 \forall adds to this list:

`gmp.h, malloc.h, unistd.h`

For the prescribed head-files, C \forall 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

Table 2: 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 |
| Cv | cmemalign | yes (0 only) | yes | yes | no |
| | resize | no copy | yes | no | yes |
| | realloc | copy | yes | no | yes |
| | alloc ^c | yes | yes | yes | yes |

^aSame as memalign but size is an integral multiple of alignment.

^bSame as valloc but rounds size to multiple of page size.

^cMultiple overloads with different parameters.

1 **extern "C"** to prevent name mangling. This approach is different from C++ where the name-mangling issue is handled
 2 internally in C header-files through checks for preprocessor variable `__cplusplus`, which adds appropriate **extern "C"**
 3 qualifiers.

4 E Standard Library

5 The Cv standard-library extends existing C library routines by adding new function, wrapping existing explicitly-
 6 polymorphic C routines into implicitly-polymorphic versions, and adding new Cv extensions.

7 E.1 Dynamic Storage-Management

8 Dynamic storage-management in C is based on explicit allocation and deallocation (`malloc/free`). Programmer's must
 9 manage all allocated storage via its address (pointer) and subsequently deallocate the storage via this address. Storage
 10 that is not deallocated becomes inaccessible, called a *memory leak*, which can only be detected at program termination.
 11 Storage freed twice is an error, called a *duplicate free*, which can sometimes be detected. Storage used after it is
 12 deallocated is an error, called using a *dangling pointer*, which can sometimes be detected.

13 E.1.1 C Interface

14 C dynamic storage-management provides the following properties.

15 **fill** storage after an allocation with a specified character or value.

16 **align** an allocation on a specified memory boundary, *e.g.*, an address multiple of 64 or 128 for cache-line purposes.

17 **scale** an allocation size to the specified number of array elements. An array may be filled, resized, or aligned.

18 **resize** an existing allocation to decreased or increased its size. In either direction, new storage may or may not be
 19 allocated, but if there is a new allocation, as much data from the existing allocation is copied into the new allocation.
 20 When new storage is allocated, it may be aligned and storage after copied data may be filled.

21 Table 2 shows different combinations of storage-management properties provided by the C and Cv allocation routines.

22 E.1.2 Cv Interface

23 Cv dynamic memory management:

24 1. extends type safety of all allocation routines by using the left-hand assignment type to determine the allocation size
 25 and alignment, and return a matching type for the new storage, which removes many common allocation errors.

26 `int * ip = (int *)malloc(sizeof(int)); // C`

27 `int * ip = malloc(); // Cv type-safe call of C malloc`

```

1   int * ip = calloc();           // CV type-safe call of C calloc
2   struct __attribute__(( aligned(128) )) spinlock { ... }; // cache alignment
3   spinlock * slp = malloc();    // correct size, alignment, and return type

```

4 Here, the alignment of the ip storage is 16 (default) and 128 for slp.

5 2. introduces the notion of *sticky properties* used in resizing. All initial allocation properties are remembered and
6 maintained for use should resize require new storage. For example, the initial alignment and fill properties in the
7 initial allocation

```

8   struct __attribute__(( aligned(4096) )) S { ... };
9   S * sp = calloc( 10 );        // align 4K and zero fill
10  sp = reallocarray( sp, 100 ); // preserve 4K alignment and zero fill new storage

```

11 are preserved in the resize so the new storage has the same alignment and extra storage after the data copy is zero
12 filled. Without sticky properties it is dangerous to resize, resulting in the C idiom of manually performing the
13 reallocation to maintain correctness, which is error prone.

14 3. provides resizing without data copying, which is useful to repurpose an existing block of storage, rather than
15 freeing the old storage and performing a new allocation. A resize can take advantage of unused storage after the
16 data to preventing a free/reallocation step altogether.

17 4. provides free/delete functions that delete a variable number of allocations.

```

18  int * ip = malloc(), * jp = malloc(), * kp = malloc();
19  double * xp = malloc(), * yp = malloc(), * zp = malloc();
20  free( ip, jp, kp, xp, yp, zp ); // multiple deallocations

```

21 5. supports constructors for initialization of allocated storage and destructors for deallocation (like C++).

```

22  struct S { int v; };           // default constructors
23  void ^?{}( S & ) { ... }      // destructor
24  S & sp = *new( 3 );           // allocate and call constructor
25  sout | sp.v;
26  delete( &sp );               // call destructor
27  S * spa1 = anew( 10, 5 ), * spa2 = anew( 10, 8 ); // allocate array and call constructor for each array element
28  for ( i; 10 ) sout | spa1[i].v | spa2[i].v | nonl; sout | nl;
29  adelete( spa1, spa2 );        // call destructors on all array objects
30
31  3
32  5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8 5 8

```

33 Allocation routines new/anew allocate a variable/array and initialize storage using the allocated type's constructor.
34 Note, the matching deallocation routines delete/adelete. C++ only supports the default constructor for initializing
35 array elements.

```

36  S * sp = new S[10][5];        // disallowed

```

37 In addition, CV provides a new allocator interface to further increase orthogonality and usability of dynamic-
38 memory allocation. This interface helps programmers in three ways.

39 1. naming: CV regular and **ttype** polymorphism (similar to C++ variadic templates) is used to encapsulate a wide
40 range of allocation functionality into a single routine name, so programmers do not have to remember multiple
41 routine names for different kinds of dynamic allocations.

42 2. named arguments: individual allocation properties are specified using postfix function call (see Section 15.3,
43 p. 35), so programmers do not have to remember parameter positions in allocation calls.

44 3. safe usage: like the CV's C-interface, programmers do not have to specify object size or cast allocation results.

45 The polymorphic functions

```

46  T * alloc( ... );
47  T * alloc( size_t dim, ... );

```

48 are overloaded with a variable number of allocation properties. These allocation properties can be passed as named
49 arguments when calling the alloc routine. A call without parameters returns an uninitialized dynamically allocated
50 object of type T (malloc). A call with only the dimension (dim) parameter returns an uninitialized dynamically allocated
51 array of objects with type T (aalloc). The variable number of arguments consist of allocation properties to specialize
52 the allocation. The properties resize and realloc are associated with an existing allocation variable indicating how its

1 storage is modified.

2 The following allocation property functions may be combined and appear in any order as arguments to alloc,

- 3 • T_align ?`align(size_t alignment) to align an allocation. The alignment parameter must be \geq the default alignment (libAlign() in CV) and a power of two, e.g., the following return a dynamic object and object array aligned on a 256 and 4096-byte boundary.

```
6 int * i0 = alloc( 256`align ); sout | i0 | nl;
7 int * i1 = alloc( 3, 4096`align ); for ( i; 3 ) sout | &i1[i] | nonl; sout | nl;
8 free( i0, i1 );
```

```
9
10 0x5555556569900 // 256 alignment
11 0x555555656c000 0x5656c004 0x5656c008 // 4K array alignment
```

- 12 • T_fill(T) ?`fill(/* various types */) to initialize storage. There are three ways to fill storage:

- 13 1. A char fills every byte of each object.
- 14 2. An object of the returned type fills each object.
- 15 3. An object array pointer fills some or all of the corresponding object array.

16 For example:

```
17 1 int * i0 = alloc( 0n`fill ); sout | *i0 | nl; // 0n disambiguates 0p
18 2 int * i1 = alloc( 5`fill ); sout | *i1 | nl;
19 3 int * i2 = alloc( 'x`fill ); sout | hex( *i2 ) | nl;
20 4 int * i3 = alloc( 5, 5`fill ); for ( i; 5 ) sout | i3[i] | nonl; sout | nl;
21 5 int * i4 = alloc( 5, 0xdeadbeefN`fill ); for ( i; 5 ) sout | hex( i4[i] ) | nonl; sout | nl;
22 6 int * i5 = alloc( 5, i3`fill ); for ( i; 5 ) sout | i5[i] | nonl; sout | nl; // completely fill from i3
23 7 int * i6 = alloc( 5, [i3, 3]`fill ); for ( i; 5 ) sout | i6[i] | nonl; sout | nl; // partial fill from i3
24 8 free( i0, i1, i2, i3, i4, i5, i6 );
```

```
25 1 0
26 2 5
27 3 0xfefefefe
28 4 5 5 5 5
29 5 0xdeadbeef 0xdeadbeef 0xdeadbeef 0xdeadbeef 0xdeadbeef
30 6 5 5 5 5
31 7 5 5 5 -555819298 -555819298 // two undefined values
```

32 Examples 1 to 3 fill an object with a value or characters. Examples 4 to 7 fill an array of objects with values,
33 another array, or part of an array.

- 34 • S_resize(T) ?`resize(void * oaddr) used to resize, realign, and fill, where the old object data is not copied to the
35 new object. The old object type may be different from the new object type, since the values are not used. For
36 example:

```
37 1 int * ip = alloc( 5`fill ); sout | ip | *ip;
38 2 ip = alloc( ip`resize, 256`align, 7`fill ); sout | ip | *ip;
39 3 double * dp = alloc( ip`resize, 4096`align, 13.5`fill ); sout | dp | *dp;
40 4 free( dp ); // DO NOT FREE ip AS ITS STORAGE IS MOVED TO dp
```

```
41 1 0x555555580a80 5
42 2 0x555555581100 7
43 3 0x555555587000 13.5
```

44 Examples 2 to 3 change the alignment, fill, and size for the initial storage of i.

```
45 1 int * ia = alloc( 5, 5`fill ); sout | ia | nonl; for ( i; 5 ) sout | ia[i] | nonl; sout | nl;
46 2 ia = alloc( 10, ia`resize, 7`fill ); sout | ia | nonl; for ( i; 10 ) sout | ia[i] | nonl; sout | nl;
47 3 ia = alloc( 5, ia`resize, 512`align, 13`fill ); sout | ia | nonl; for ( i; 5 ) sout | ia[i] | nonl; sout | nl;;
48 4 ia = alloc( 3, ia`resize, 4096`align, 2`fill ); for ( i; 3 ) sout | &ia[i] | ia[i] | nonl; sout | nl;
49 5 free( ia );
50 1 0x555555656d540 5 5 5 5 5
51 2 0x555555656d480 7 7 7 7 7 7 7 7 7
52 3 0x555555656fe00 13 13 13 13 13
53 4 0x5555556570000 2 0x5555556570004 2 0x5555556570008 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  0x555555556d5c0 5
2  0x5555555570000 5
3  0x5555555571000 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  0x555555656d540 5 5 5 5 5
2  0x555555656d480 7 7 7 7 7 7 7 7 7 7
3  0x5555556570e00 5 5 5 5 5
4  0x5555556571000 5 0x5555556571004 5 0x5555556571008 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.

```

25  extern "C" {
26      // New C allocation operations.
27      void * aalloc( size_t dim, size_t elemSize );
28      void * resize( void * oaddr, size_t size );
29      void * amemalign( size_t align, size_t dim, size_t elemSize );
30      void * cmemalign( size_t align, size_t dim, size_t elemSize );
31      size_t malloc_alignment( void * addr );
32      bool malloc_zero_fill( void * addr );
33      size_t malloc_size( void * addr );
34      int malloc_stats_fd( int fd );
35      size_t malloc_expansion();           // heap expansion size (bytes)
36      size_t malloc_mmap_start();         // crossover allocation size from sbrk to mmap
37      size_t malloc_unfreed();           // heap unfreed size (bytes)
38      void malloc_stats_clear();         // clear heap statistics
39  }

```

```

41  // New allocation operations.
42  void * resize( void * oaddr, size_t alignment, size_t size );
43  void * realloc( void * oaddr, size_t alignment, size_t size );
44  void * reallocarray( void * oaddr, size_t nalign, size_t dim, size_t elemSize );

```

```

46  forall( T & | sized(T) ) {
47      //CV safe equivalents, i.e., implicit size specification, eliminate return-type cast
48      T * malloc( void );
49      T * aalloc( size_t dim );
50      T * calloc( size_t dim );
51      T * resize( T * ptr, size_t size );
52      T * resize( T * ptr, size_t alignment, size_t size );
53      T * realloc( T * ptr, size_t size );
54      T * realloc( T * ptr, size_t alignment, size_t size );
55      T * reallocarray( T * ptr, size_t dim );
56      T * reallocarray( T * ptr, size_t alignment, size_t dim );

```

```

1      T * memalign( size_t align );
2      T * amemalign( size_t align, size_t dim );
3      T * cmemalign( size_t align, size_t dim );
4      T * aligned_alloc( size_t align );
5      int posix_memalign( T ** ptr, size_t align );
6      T * valloc( void );
7      T * pvalloc( void );
8
9      //CV safe general allocation, fill, resize, alignment, array
10     T * alloc( ... ); // variable, T size
11     T * alloc( size_t dim, ... );
12     T_align ?` align( size_t alignment );
13     T_fill(T) ?` fill( /* various types */ );
14     T_resize ?` resize( void * oaddr );
15     T_realloc ?` realloc( void * oaddr );
16 }
17
18 forall( T &, List ... ) void free( T * ptr, ... ) // deallocation list
19
20 //CV allocation/deallocation and constructor/destructor, non-array types
21 forall( T &, Params ... | { void ?{}( T &, Params ); } ) T * new( Params ... );
22 forall( T &, List ... | { void ^?{}( T & ); void delete( List ... ); } );
23 //CV allocation/deallocation and constructor/destructor, array types
24 forall( T & | sized(T), Params ... | { void ?{}( T &, Params ); } ) T * anew( size_t dim, Params ... );
25 forall( T & | sized(T) | { void ^?{}( T & ); }, List ... ) void adelete( T arr[], List ... );

```

26 E.2 Memory Set and Copy

27 Like safe memory allocation, CV provides safe block initialization and copy. While objects should be initialized/copied
 28 with constructors/assignment, block operations can be very performant. In certain cases the compiler generates block
 29 copy operations, such as assigning structures `s = t`, however C arrays cannot be assigned.

| CV | C |
|---|---|
| 30 struct S { int i, j, k; }; 31 S s, t, *sp = &s, *tp = &t, sa[10], ta[10]; | |
| 32 | |
| 33 | |
| 34 | |
| 35 | |
| 36 | |
| 37 | |
| 38 | |
| 39 | |
| 40 | |
| 41 | |
| 42 | |
| 43 | |
| 44 | |
| 45 | |
| 46 | |
| 47 | |
| 48 | |

34 These operations provide uniformity between reference and pointer, so object dereferencing, '&', is unnecessary.

```

35 static inline forall( T & | sized(T) ) {
36     //CFA safe initialization/copy, i.e., implicit size specification, non-array types
37     T * memset( T * dest, char fill ); // all combinations of pointer/reference
38     T * memset( T & dest, char fill );
39
40     T * memcpy( T * dest, const T * src ); // all combinations of pointer/reference
41     T * memcpy( T & dest, const T & src );
42     T * memcpy( T * dest, const T & src );
43     T * memcpy( T & dest, const T * src );
44
45     //CFA safe initialization/copy, i.e., implicit size specification, array types
46     T * amemset( T dest[], char fill, size_t dim );
47     T * amemcpy( T dest[], const T src[], size_t dim );
48 }

```

1 E.3 String to Value Conversion

```

2   int ato( const char * ptr );
3   unsigned int ato( const char * ptr );
4   long int ato( const char * ptr );
5   unsigned long int ato( const char * ptr );
6   long long int ato( const char * ptr );
7   unsigned long long int ato( const char * ptr );
8   float ato( const char * ptr );
9   double ato( const char * ptr );
10  long double ato( const char * ptr );
11  float _Complex ato( const char * ptr );
12  double _Complex ato( const char * ptr );
13  long double _Complex ato( const char * ptr );
14
15  int strtol( const char * s, char ** eptr, int base );
16  unsigned int strtoul( const char * s, char ** eptr, int base );
17  long int strtol( const char * s, char ** eptr, int base );
18  unsigned long int strtoul( const char * s, char ** eptr, int base );
19  long long int strtoll( const char * s, char ** eptr, int base );
20  unsigned long long int strtoull( const char * s, char ** eptr, int base );
21  float strtod( const char * s, char ** eptr );
22  double strtod( const char * s, char ** eptr );
23  long double strtold( const char * s, char ** eptr );
24  float _Complex strtoc( const char * s, char ** eptr );
25  double _Complex strtod( const char * s, char ** eptr );
26  long double _Complex strtold( const char * s, char ** eptr );

```

27 E.4 Search / Sort

```

28  forall( T | { int ?<?( T, T ); } )           // location
29  T * bsearch( T key, const T * arr, size_t dim );
30
31  forall( T | { int ?<?( T, T ); } )           // position
32  unsigned int bsearch( T key, const T * arr, size_t dim );
33
34  forall( T | { int ?<?( T, T ); } )
35  void qsort( const T * arr, size_t dim );
36
37  forall( E | { int ?<?( E, E ); } ) {
38      E * bsearch( E key, const E * vals, size_t dim ); // location
39      size_t bsearch( E key, const E * vals, size_t dim ); // position
40      E * bsearchl( E key, const E * vals, size_t dim );
41      size_t bsearchl( E key, const E * vals, size_t dim );
42      E * bsearchu( E key, const E * vals, size_t dim );
43      size_t bsearchu( E key, const E * vals, size_t dim );
44  }
45
46  forall( K, E | { int ?<?( K, K ); K getKey( const E & ); } ) {
47      E * bsearch( K key, const E * vals, size_t dim );
48      size_t bsearch( K key, const E * vals, size_t dim );
49      E * bsearchl( K key, const E * vals, size_t dim );
50      size_t bsearchl( K key, const E * vals, size_t dim );
51      E * bsearchu( K key, const E * vals, size_t dim );
52      size_t bsearchu( K key, const E * vals, size_t dim );
53  }
54
55  forall( E | { int ?<?( E, E ); } ) {
56      void qsort( E * vals, size_t dim );
57  }

```

1 **E.5 Absolute Value**

```

2   unsigned char abs( signed char );
3   int abs( int );
4   unsigned long int abs( long int );
5   unsigned long long int abs( long long int );
6   float abs( float );
7   double abs( double );
8   long double abs( long double );
9   float abs( float _Complex );
10  double abs( double _Complex );
11  long double abs( long double _Complex );
12  forall( T | { void ?{}( T *, zero_t ); int ?<?( T, T ); T -?( T ); } )
13  T abs( T );

```

14 **E.6 C Random Numbers**

```

15  void srand( unsigned int seed );
16  char random( void );
17  char random( char u ); // [0,u)
18  char random( char l, char u ); // [l,u]
19  int random( void );
20  int random( int u ); // [0,u)
21  int random( int l, int u ); // [l,u]
22  unsigned int random( void );
23  unsigned int random( unsigned int u ); // [0,u)
24  unsigned int random( unsigned int l, unsigned int u ); // [l,u]
25  long int random( void );
26  long int random( long int u ); // [0,u)
27  long int random( long int l, long int u ); // [l,u]
28  unsigned long int random( void );
29  unsigned long int random( unsigned long int u ); // [0,u)
30  unsigned long int random( unsigned long int l, unsigned long int u ); // [l,u]
31  float random( void ); // [0.0, 1.0)
32  double random( void ); // [0.0, 1.0)
33  float _Complex random( void ); // [0.0, 1.0)+[0.0, 1.0)i
34  double _Complex random( void ); // [0.0, 1.0)+[0.0, 1.0)i
35  long double _Complex random( void ); // [0.0, 1.0)+[0.0, 1.0)i

```

36 **E.7 Algorithms**

```

37  forall( T | { int ?<?( T, T ); } ) T min( T t1, T t2 );
38  forall( T | { int ?>?( T, T ); } ) T max( T t1, T t2 );
39  forall( T | { T min( T, T ); T max( T, T ); } ) T clamp( T value, T min_val, T max_val );
40  forall( T ) void swap( T * t1, T * t2 );

```

41 **F Math Library**

42 The CV math-library wraps explicitly-polymorphic C math-routines into implicitly-polymorphic versions.

43 **F.1 General**

```

44  float ?%?( float, float );
45  float fmod( float, float );
46  double ?%?( double, double );
47  double fmod( double, double );
48  long double ?%?( long double, long double );
49  long double fmod( long double, long double );
50
51  float remainder( float, float );
52  double remainder( double, double );

```

```

1   long double remainder( long double, long double );
2
3   float remquo( float, float, int * );
4   double remquo( double, double, int * );
5   long double remquo( long double, long double, int * );
6   [ int, float ] remquo( float, float );
7   [ int, double ] remquo( double, double );
8   [ int, long double ] remquo( long double, long double );
9
10  [ int, float ] div( float, float );
11  [ int, double ] div( double, double );
12  [ int, long double ] div( long double, long double );
13
14  float fma( float, float, float );
15  double fma( double, double, double );
16  long double fma( long double, long double, long double );
17
18  float fdim( float, float );
19  double fdim( double, double );
20  long double fdim( long double, long double );
21
22  float nan( const char * );
23  double nan( const char * );
24  long double nan( const char * );

```

25 F.2 Exponential

```

26  float exp( float );
27  double exp( double );
28  long double exp( long double );
29  float _Complex exp( float _Complex );
30  double _Complex exp( double _Complex );
31  long double _Complex exp( long double _Complex );
32
33  float exp2( float );
34  double exp2( double );
35  long double exp2( long double );
36  // float _Complex exp2( float _Complex );
37  // double _Complex exp2( double _Complex );
38  // long double _Complex exp2( long double _Complex );
39
40  float expm1( float );
41  double expm1( double );
42  long double expm1( long double );
43
44  float pow( float, float );
45  double pow( double, double );
46  long double pow( long double, long double );
47  float _Complex pow( float _Complex, float _Complex );
48  double _Complex pow( double _Complex, double _Complex );
49  long double _Complex pow( long double _Complex, long double _Complex );

```

50 F.3 Logarithm

```

51  float log( float );
52  double log( double );
53  long double log( long double );
54  float _Complex log( float _Complex );
55  double _Complex log( double _Complex );
56  long double _Complex log( long double _Complex );

```

```

1
2  int log2( unsigned int );
3  long int log2( unsigned long int );
4  long long int log2( unsigned long long int )
5  float log2( float );
6  double log2( double );
7  long double log2( long double );
8  // float _Complex log2( float _Complex );
9  // double _Complex log2( double _Complex );
10 // long double _Complex log2( long double _Complex );
11
12 float log10( float );
13 double log10( double );
14 long double log10( long double );
15 // float _Complex log10( float _Complex );
16 // double _Complex log10( double _Complex );
17 // long double _Complex log10( long double _Complex );
18
19 float log1p( float );
20 double log1p( double );
21 long double log1p( long double );
22
23 int ilogb( float );
24 int ilogb( double );
25 int ilogb( long double );
26
27 float logb( float );
28 double logb( double );
29 long double logb( long double );
30
31 float sqrt( float );
32 double sqrt( double );
33 long double sqrt( long double );
34 float _Complex sqrt( float _Complex );
35 double _Complex sqrt( double _Complex );
36 long double _Complex sqrt( long double _Complex );
37
38 float cbrt( float );
39 double cbrt( double );
40 long double cbrt( long double );
41
42 float hypot( float, float );
43 double hypot( double, double );
44 long double hypot( long double, long double );

```

45 F.4 Trigonometric

```

46 float sin( float );
47 double sin( double );
48 long double sin( long double );
49 float _Complex sin( float _Complex );
50 double _Complex sin( double _Complex );
51 long double _Complex sin( long double _Complex );
52
53 float cos( float );
54 double cos( double );
55 long double cos( long double );
56 float _Complex cos( float _Complex );
57 double _Complex cos( double _Complex );
58 long double _Complex cos( long double _Complex );

```

```

1
2  float tan( float );
3  double tan( double );
4  long double tan( long double );
5  float _Complex tan( float _Complex );
6  double _Complex tan( double _Complex );
7  long double _Complex tan( long double _Complex );
8
9  float asin( float );
10 double asin( double );
11 long double asin( long double );
12 float _Complex asin( float _Complex );
13 double _Complex asin( double _Complex );
14 long double _Complex asin( long double _Complex );
15
16 float acos( float );
17 double acos( double );
18 long double acos( long double );
19 float _Complex acos( float _Complex );
20 double _Complex acos( double _Complex );
21 long double _Complex acos( long double _Complex );
22
23 float atan( float );
24 double atan( double );
25 long double atan( long double );
26 float _Complex atan( float _Complex );
27 double _Complex atan( double _Complex );
28 long double _Complex atan( long double _Complex );
29
30 float atan2( float, float );
31 double atan2( double, double );
32 long double atan2( long double, long double );
33
34 float atan( float, float );           // alternative name for atan2
35 double atan( double, double );
36 long double atan( long double, long double );

```

37 F.5 Hyperbolic

```

38 float sinh( float );
39 double sinh( double );
40 long double sinh( long double );
41 float _Complex sinh( float _Complex );
42 double _Complex sinh( double _Complex );
43 long double _Complex sinh( long double _Complex );
44
45 float cosh( float );
46 double cosh( double );
47 long double cosh( long double );
48 float _Complex cosh( float _Complex );
49 double _Complex cosh( double _Complex );
50 long double _Complex cosh( long double _Complex );
51
52 float tanh( float );
53 double tanh( double );
54 long double tanh( long double );
55 float _Complex tanh( float _Complex );
56 double _Complex tanh( double _Complex );
57 long double _Complex tanh( long double _Complex );
58

```

```

1  float asinh( float );
2  double asinh( double );
3  long double asinh( long double );
4  float _Complex asinh( float _Complex );
5  double _Complex asinh( double _Complex );
6  long double _Complex asinh( long double _Complex );
7
8  float acosh( float );
9  double acosh( double );
10 long double acosh( long double );
11 float _Complex acosh( float _Complex );
12 double _Complex acosh( double _Complex );
13 long double _Complex acosh( long double _Complex );
14
15 float atanh( float );
16 double atanh( double );
17 long double atanh( long double );
18 float _Complex atanh( float _Complex );
19 double _Complex atanh( double _Complex );
20 long double _Complex atanh( long double _Complex );

```

21 F.6 Error / Gamma

```

22 float erf( float );
23 double erf( double );
24 long double erf( long double );
25 float _Complex erf( float _Complex );
26 double _Complex erf( double _Complex );
27 long double _Complex erf( long double _Complex );
28
29 float erfc( float );
30 double erfc( double );
31 long double erfc( long double );
32 float _Complex erfc( float _Complex );
33 double _Complex erfc( double _Complex );
34 long double _Complex erfc( long double _Complex );
35
36 float lgamma( float );
37 double lgamma( double );
38 long double lgamma( long double );
39 float lgamma( float, int * );
40 double lgamma( double, int * );
41 long double lgamma( long double, int * );
42
43 float tgamma( float );
44 double tgamma( double );
45 long double tgamma( long double );

```

46 F.7 Nearest Integer

```

47 // n / align * align
48 signed char floor( signed char n, signed char align );
49 unsigned char floor( unsigned char n, unsigned char align );
50 short int floor( short int n, short int align );
51 unsigned short int floor( unsigned short int n, unsigned short int align );
52 int floor( int n, int align );
53 unsigned int floor( unsigned int n, unsigned int align );
54 long int floor( long int n, long int align );
55 unsigned long int floor( unsigned long int n, unsigned long int align );
56 long long int floor( long long int n, long long int align );

```

```

1  unsigned long long int floor( unsigned long long int n, unsigned long long int align );
2
3  // (n + (align - 1)) / align
4  signed char ceiling_div( signed char n, char align );
5  unsigned char ceiling_div( unsigned char n, unsigned char align );
6  short int ceiling_div( short int n, short int align );
7  unsigned short int ceiling_div( unsigned short int n, unsigned short int align );
8  int ceiling_div( int n, int align );
9  unsigned int ceiling_div( unsigned int n, unsigned int align );
10 long int ceiling_div( long int n, long int align );
11 unsigned long int ceiling_div( unsigned long int n, unsigned long int align );
12 long long int ceiling_div( long long int n, long long int align );
13 unsigned long long int ceiling_div( unsigned long long int n, unsigned long long int align );
14
15 // floor( n + (n % align != 0 ? align - 1 : 0), align )
16 signed char ceiling( signed char n, signed char align );
17 unsigned char ceiling( unsigned char n, unsigned char align );
18 short int ceiling( short int n, short int align );
19 unsigned short int ceiling( unsigned short int n, unsigned short int align );
20 int ceiling( int n, int align );
21 unsigned int ceiling( unsigned int n, unsigned int align );
22 long int ceiling( long int n, long int align );
23 unsigned long int ceiling( unsigned long int n, unsigned long int align );
24 long long int ceiling( long long int n, long long int align );
25 unsigned long long int ceiling( unsigned long long int n, unsigned long long int align );
26
27 float floor( float );
28 double floor( double );
29 long double floor( long double );
30
31 float ceil( float );
32 double ceil( double );
33 long double ceil( long double );
34
35 float trunc( float );
36 double trunc( double );
37 long double trunc( long double );
38
39 float rint( float );
40 long double rint( long double );
41 long int rint( float );
42 long int rint( double );
43 long int rint( long double );
44 long long int rint( float );
45 long long int rint( double );
46 long long int rint( long double );
47
48 long int lrint( float );
49 long int lrint( double );
50 long int lrint( long double );
51 long long int llrint( float );
52 long long int llrint( double );
53 long long int llrint( long double );
54
55 float nearbyint( float );
56 double nearbyint( double );
57 long double nearbyint( long double );
58
59 float round( float );
60 long double round( long double );

```

```

1   long int round( float );
2   long int round( double );
3   long int round( long double );
4   long long int round( float );
5   long long int round( double );
6   long long int round( long double );
7
8   long int lround( float );
9   long int lround( double );
10  long int lround( long double );
11  long long int llround( float );
12  long long int llround( double );
13  long long int llround( long double );

```

14 F.8 Manipulation

```

15  float copysign( float, float );
16  double copysign( double, double );
17  long double copysign( long double, long double );
18
19  float frexp( float, int * );
20  double frexp( double, int * );
21  long double frexp( long double, int * );
22
23  float ldexp( float, int );
24  double ldexp( double, int );
25  long double ldexp( long double, int );
26
27  [ float, float ] modf( float );
28  float modf( float, float * );
29  [ double, double ] modf( double );
30  double modf( double, double * );
31  [ long double, long double ] modf( long double );
32  long double modf( long double, long double * );
33
34  float nextafter( float, float );
35  double nextafter( double, double );
36  long double nextafter( long double, long double );
37
38  float nexttoward( float, long double );
39  double nexttoward( double, long double );
40  long double nexttoward( long double, long double );
41
42  float scalbn( float, int );
43  double scalbn( double, int );
44  long double scalbn( long double, int );
45
46  float scalbln( float, long int );
47  double scalbln( double, long int );
48  long double scalbln( long double, long int );

```

49 G Time Keeping

50 G.1 Duration

```

51  struct Duration {
52      int64_t tn;           // nanoseconds
53  };
54
55  void ?{}( Duration & dur );

```

```

1  void ?{}( Duration & dur, zero_t );
2  void ?{}( Duration & dur, timeval t )
3  void ?{}( Duration & dur, timespec t )
4
5  Duration ?=( Duration & dur, zero_t );
6  Duration ?=( Duration & dur, timeval t )
7  Duration ?=( Duration & dur, timespec t )
8
9  Duration +( Duration rhs );
10 Duration ?+( Duration & lhs, Duration rhs );
11 Duration ?+=( Duration & lhs, Duration rhs );
12
13 Duration -( Duration rhs );
14 Duration ?-( Duration & lhs, Duration rhs );
15 Duration ?-=( Duration & lhs, Duration rhs );
16
17 Duration ?*( Duration lhs, int64_t rhs );
18 Duration ?*( int64_t lhs, Duration rhs );
19 Duration ?*=( Duration & lhs, int64_t rhs );
20
21 int64_t ?/( Duration lhs, Duration rhs );
22 Duration ?/( Duration lhs, int64_t rhs );
23 Duration ?/=( Duration & lhs, int64_t rhs );
24 double div( Duration lhs, Duration rhs );
25
26 Duration ?%( Duration lhs, Duration rhs );
27 Duration ?%=( Duration & lhs, Duration rhs );
28
29 bool ?==( Duration lhs, zero_t );
30 bool ?!==( Duration lhs, zero_t );
31 bool ?<? ( Duration lhs, zero_t );
32 bool ?<=( Duration lhs, zero_t );
33 bool ?>? ( Duration lhs, zero_t );
34 bool ?>=( Duration lhs, zero_t );
35
36 bool ?==( Duration lhs, Duration rhs );
37 bool ?!==( Duration lhs, Duration rhs );
38 bool ?<? ( Duration lhs, Duration rhs );
39 bool ?<=( Duration lhs, Duration rhs );
40 bool ?>? ( Duration lhs, Duration rhs );
41 bool ?>=( Duration lhs, Duration rhs );
42
43 Duration abs( Duration rhs );
44
45 Duration ?`ns( int64_t nsec );
46 Duration ?`us( int64_t usec );
47 Duration ?`ms( int64_t msec );
48 Duration ?`s( int64_t sec );
49 Duration ?`s( double sec );
50 Duration ?`m( int64_t min );
51 Duration ?`m( double min );
52 Duration ?`h( int64_t hours );
53 Duration ?`h( double hours );
54 Duration ?`d( int64_t days );
55 Duration ?`d( double days );
56 Duration ?`w( int64_t weeks );
57 Duration ?`w( double weeks );
58
59 int64_t ?`ns( Duration dur );
60 int64_t ?`us( Duration dur );

```

```

1   int64_t ?`ms( Duration dur );
2   int64_t ?`s( Duration dur );
3   int64_t ?`m( Duration dur );
4   int64_t ?`h( Duration dur );
5   int64_t ?`d( Duration dur );
6   int64_t ?`w( Duration dur );
7
8   double ?`dns( Duration dur );
9   double ?`dus( Duration dur );
10  double ?`dms( Duration dur );
11  double ?`ds( Duration dur );
12  double ?`dm( Duration dur );
13  double ?`dh( Duration dur );
14  double ?`dd( Duration dur );
15  double ?`dw( Duration dur );
16
17  Duration max( Duration lhs, Duration rhs );
18  Duration min( Duration lhs, Duration rhs );
19
20  forall( ostype & | ostream( ostype ) ) ostype & ?|( ostype & os, Duration dur );

```

21 G.2 timeval

```

22  void ?{}( timeval & t );
23  void ?{}( timeval & t, zero_t );
24  void ?{}( timeval & t, time_t sec, suseconds_t usec );
25  void ?{}( timeval & t, time_t sec );
26  void ?{}( timeval & t, Time time );
27
28  timeval ?=? ( timeval & t, zero_t );
29  timeval ?+?( timeval & lhs, timeval rhs );
30  timeval ?-?( timeval & lhs, timeval rhs );
31  bool ?==?( timeval lhs, timeval rhs );
32  bool ?!=?( timeval lhs, timeval rhs );

```

33 G.3 timespec

```

34  void ?{}( timespec & t );
35  void ?{}( timespec & t, zero_t );
36  void ?{}( timespec & t, time_t sec, __syscall_slong_t nsec );
37  void ?{}( timespec & t, time_t sec );
38  void ?{}( timespec & t, Time time );
39
40  timespec ?=? ( timespec & t, zero_t );
41  timespec ?+?( timespec & lhs, timespec rhs );
42  timespec ?-?( timespec & lhs, timespec rhs );
43  bool ?==?( timespec lhs, timespec rhs );
44  bool ?!=?( timespec lhs, timespec rhs );

```

45 G.4 itimerval

```

46  void ?{}( itimerval & itv, Duration alarm );
47  void ?{}( itimerval & itv, Duration alarm, Duration interval );

```

48 G.5 Time

```

49  struct Time {
50      uint64_t tn;           // nanoseconds since UNIX epoch
51  };
52
53  void ?{}( Time & time );

```

```

1  void ?{}( Time & time, zero_t );
2  void ?{}( Time & time, timeval t );
3  void ?{}( Time & time, timespec t );
4
5  Time ?=? ( Time & time, zero_t );
6  Time ?=? ( Time & time, timeval t );
7  Time ?=? ( Time & time, timespec t );
8
9  Time ?+?( Time & lhs, Duration rhs );
10 Time ?+?( Duration lhs, Time rhs );
11 Time ?+=( Time & lhs, Duration rhs );
12
13 Duration ?-?( Time lhs, Time rhs );
14 Time ?-?( Time lhs, Duration rhs );
15 Time ?-=( Time & lhs, Duration rhs );
16 bool ?==( Time lhs, Time rhs );
17 bool ?!=( Time lhs, Time rhs );
18 bool ?<?( Time lhs, Time rhs );
19 bool ?<=( Time lhs, Time rhs );
20 bool ?>?( Time lhs, Time rhs );
21 bool ?>=( Time lhs, Time rhs );
22
23 int64_t ?`ns( Time t );
24
25 char * yy_mm_dd( Time time, char * buf );
26 char * ?`ymd( Time time, char * buf ); // short form
27
28 char * mm_dd_yy( Time time, char * buf );
29 char * ?`mdy( Time time, char * buf ); // short form
30
31 char * dd_mm_yy( Time time, char * buf );
32 char * ?`dmy( Time time, char * buf ); // short form
33
34 size_t strftime( char * buf, size_t size, const char * fmt, Time time );
35
36 forall( ostype & | ostream( ostype ) ) ostype & ?{?( ostype & os, Time time );

```

37 H Clock

38 H.1 C time

```

39 char * ctime( time_t tp );
40 char * ctime_r( time_t tp, char * buf );
41 tm * gmtime( time_t tp );
42 tm * gmtime_r( time_t tp, tm * result );
43 tm * localtime( time_t tp );
44 tm * localtime_r( time_t tp, tm * result );

```

45 H.2 Clock

```

46 struct Clock { // virtual clock
47     Duration offset; // offset from computer real-time
48 };
49
50 void ?{}( Clock & clk ); // create no offset
51 void ?{}( Clock & clk, Duration adj ); // create with offset
52 void reset( Clock & clk, Duration adj ); // change offset
53
54 Duration resolutionHi(); // clock resolution in nanoseconds (fine)
55 Duration resolution(); // clock resolution without nanoseconds (coarse)

```

```

1
2   Time timeHiRes();           // real time with nanoseconds
3   Time time();               // real time without nanoseconds
4   Time time( Clock & clk );  // real time for given clock
5   Time ?()( Clock & clk );   // alternative syntax
6   timeval time( Clock & clk ); // convert to C time format
7   tm time( Clock & clk );
8   Duration processor();      // non-monotonic duration of kernel thread
9   Duration program();        // non-monotonic duration of program CPU
10  Duration boot();           // monotonic duration since computer boot

```

11 I Pseudo Random Number Generator

12 Random numbers are values generated independently, i.e., new values do not depend on previous values (independent
13 trials), e.g., lottery numbers, shuffled cards, dice roll, coin flip. While a primary goal of programming is computing
14 values that are *not* random, random values are useful in simulation, cryptography, games, etc. A random-number gen-
15 erator is an algorithm that computes independent values. If the algorithm uses deterministic computation (a predictable
16 sequence of values), it generates *pseudo* random numbers versus *true* random numbers.

17 All *pseudo random-number generators (PRNG)* involve some technique to scramble bits of a value, e.g., multi-
18 plicative recurrence:

```
19   rand = 33967 * (rand + 1063); // scramble bits
```

20 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 |
| 21 | 0xad3e | 0x7b5f1dbe |
| | 0xbc3b | 0xac69ff19 |
| | 0x1070f | 0x2d258dc6 |

22 By dropping bits 63–32, bits 31–0 become scrambled after each multiply. The least-significant bits *appear* random
23 but the same bits are always generated given a fixed starting value, called the *seed* (value 0x3e8e36 above). Hence, if
24 a program uses the same seed, the same sequence of pseudo-random values is generated from the PRNG. Often the
25 seed is set to another random value like a program’s process identifier (getpid) or time when the program is run; hence,
26 one random value bootstraps another. Finally, a PRNG usually generates a range of large values, e.g., [0, UINT_MAX],
27 which are scaled using the modulus operator, e.g., prng() % 5 produces random values in the range 0–4.

28 CV provides 32/64-bit sequential PRNG type only accessible by a single thread (not thread-safe) and a set of global
29 routines and companion thread PRNG functions accessible by multiple threads without contention. To use the PRNG
30 interface **requires including `stdlib.h`**.

31 • The PRNG types for sequential programs, including coroutines, are:

```

32   struct PRNG32 {}; // opaque type, no copy or assignment
33   void ?{}( PRNG32 & prng, uint32_t seed ); // fixed seed
34   void ?{}( PRNG32 & prng ); // random seed
35   void set_seed( PRNG32 & prng, uint32_t seed ); // set seed
36   uint32_t get_seed( PRNG32 & prng ); // get seed
37   uint32_t prng( PRNG32 & prng ); // [0,UINT_MAX]
38   uint32_t prng( PRNG32 & prng, uint32_t u ); // [0,u]
39   uint32_t prng( PRNG32 & prng, uint32_t l, uint32_t u ); // [l,u]
40   uint32_t calls( PRNG32 & prng ); // number of calls
41   void copy( PRNG32 & dst, PRNG32 & src ); // checkpoint PRNG state

42   struct PRNG64 {}; // opaque type, no copy or assignment
43   void ?{}( PRNG64 & prng, uint64_t seed ); // fixed seed
44   void ?{}( PRNG64 & prng ); // random seed
45   void set_seed( PRNG64 & prng, uint64_t seed ); // set seed
46   uint64_t get_seed( PRNG64 & prng ); // get seed
47   uint64_t prng( PRNG64 & prng ); // [0,UINT_MAX]

```

```

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

```

1      uint64_t prng( PRNG64 & prng, uint64_t u );           // [0,u]
2      uint64_t prng( PRNG64 & prng, uint64_t l, uint64_t u ); // [l,u]
3      uint64_t calls( PRNG64 & prng );                   // number of calls
4      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 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)`.

```

15      void set_seed( size_t seed );                       // set global seed
16      size_t get_seed();                                 // get global seed
17      // SLOWER, global routines
18      size_t prng( void );                               // [0,UINT_MAX]
19      size_t prng( size_t u );                           // [0,u]
20      size_t prng( size_t l, size_t u );                 // [l,u]
21      // FASTER, thread members
22      size_t prng( thread$ & th );                       // [0,UINT_MAX]
23      size_t prng( thread$ & th, size_t u );             // [0,u]
24      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 CV 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,

```

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

1 e.g., in thread main, eliminating the call to active_thread significantly reduces the cost of accessing the thread's
2 PRNG state. Figure 17 shows an example using the slower/faster concurrent PRNG in the program main and a
3 thread.

4 J Multi-precision Integers

5 CV has an interface to the GMP multi-precision signed-integers [17], similar to the C++ interface provided by GMP. The
6 CV interface wraps GMP routines into operator routines to make programming with multi-precision integers identical
7 to using fixed-sized integers. The CV type name for multi-precision signed-integers is lnt and the header file is gmp.

```

8 void ?{}( lnt * this ); // constructor/destructor
9 void ?{}( lnt * this, lnt init );
10 void ?{}( lnt * this, zero_t );
11 void ?{}( lnt * this, one_t );
12 void ?{}( lnt * this, signed long int init );
13 void ?{}( lnt * this, unsigned long int init );
14 void ?{}( lnt * this, const char * val );
15 void ^?{}( lnt * this );
16
17 lnt ?=? ( lnt * lhs, lnt rhs ); // assignment
18 lnt ?=? ( lnt * lhs, long int rhs );
19 lnt ?=? ( lnt * lhs, unsigned long int rhs );
20 lnt ?=? ( lnt * lhs, const char * rhs );
21
22 char ?=? ( char * lhs, lnt rhs );
23 short int ?=? ( short int * lhs, lnt rhs );

```

```

1  int ?=( int * lhs, Int rhs );
2  long int ?=( long int * lhs, Int rhs );
3  unsigned char ?=( unsigned char * lhs, Int rhs );
4  unsigned short int ?=( unsigned short int * lhs, Int rhs );
5  unsigned int ?=( unsigned int * lhs, Int rhs );
6  unsigned long int ?=( unsigned long int * lhs, Int rhs );
7
8  long int narrow( Int val );
9  unsigned long int narrow( Int val );
10
11 int ?==( Int oper1, Int oper2 );           // comparison
12 int ?==( Int oper1, long int oper2 );
13 int ?==( long int oper2, Int oper1 );
14 int ?==( Int oper1, unsigned long int oper2 );
15 int ?==( unsigned long int oper2, Int oper1 );
16
17 int ?!==( Int oper1, Int oper2 );
18 int ?!==( Int oper1, long int oper2 );
19 int ?!==( long int oper1, Int oper2 );
20 int ?!==( Int oper1, unsigned long int oper2 );
21 int ?!==( unsigned long int oper1, Int oper2 );
22
23 int ?<?( Int oper1, Int oper2 );
24 int ?<?( Int oper1, long int oper2 );
25 int ?<?( long int oper2, Int oper1 );
26 int ?<?( Int oper1, unsigned long int oper2 );
27 int ?<?( unsigned long int oper2, Int oper1 );
28
29 int ?<=( Int oper1, Int oper2 );
30 int ?<=( Int oper1, long int oper2 );
31 int ?<=( long int oper2, Int oper1 );
32 int ?<=( Int oper1, unsigned long int oper2 );
33 int ?<=( unsigned long int oper2, Int oper1 );
34
35 int ?>?( Int oper1, Int oper2 );
36 int ?>?( Int oper1, long int oper2 );
37 int ?>?( long int oper1, Int oper2 );
38 int ?>?( Int oper1, unsigned long int oper2 );
39 int ?>?( unsigned long int oper1, Int oper2 );
40
41 int ?>=( Int oper1, Int oper2 );
42 int ?>=( Int oper1, long int oper2 );
43 int ?>=( long int oper1, Int oper2 );
44 int ?>=( Int oper1, unsigned long int oper2 );
45 int ?>=( unsigned long int oper1, Int oper2 );
46
47 Int +?( Int oper );                       // arithmetic
48 Int -( Int oper );
49 Int ~( Int oper );
50
51 Int ?&?( Int oper1, Int oper2 );
52 Int ?&?( Int oper1, long int oper2 );
53 Int ?&?( long int oper1, Int oper2 );
54 Int ?&?( Int oper1, unsigned long int oper2 );
55 Int ?&?( unsigned long int oper1, Int oper2 );
56 Int ?&=( Int * lhs, Int rhs );
57
58 Int ?|?( Int oper1, Int oper2 );
59 Int ?|?( Int oper1, long int oper2 );
60 Int ?|?( long int oper1, Int oper2 );

```

```

1   Int ?|( Int oper1, unsigned long int oper2 );
2   Int ?|?( unsigned long int oper1, Int oper2 );
3   Int ?|=( Int * lhs, Int rhs );
4
5   Int ?^( Int oper1, Int oper2 );
6   Int ?^( Int oper1, long int oper2 );
7   Int ?^( long int oper1, Int oper2 );
8   Int ?^( Int oper1, unsigned long int oper2 );
9   Int ?^( unsigned long int oper1, Int oper2 );
10  Int ?^=( Int * lhs, Int rhs );
11
12  Int ?+?( Int addend1, Int addend2 );
13  Int ?+?( Int addend1, long int addend2 );
14  Int ?+?( long int addend2, Int addend1 );
15  Int ?+?( Int addend1, unsigned long int addend2 );
16  Int ?+?( unsigned long int addend2, Int addend1 );
17  Int ?+=( Int * lhs, Int rhs );
18  Int ?+=( Int * lhs, long int rhs );
19  Int ?+=( Int * lhs, unsigned long int rhs );
20  Int ++?( Int * lhs );
21  Int ?++( Int * lhs );
22
23  Int ?-( Int minuend, Int subtrahend );
24  Int ?-( Int minuend, long int subtrahend );
25  Int ?-( long int minuend, Int subtrahend );
26  Int ?-( Int minuend, unsigned long int subtrahend );
27  Int ?-( unsigned long int minuend, Int subtrahend );
28  Int ?-=( Int * lhs, Int rhs );
29  Int ?-=( Int * lhs, long int rhs );
30  Int ?-=( Int * lhs, unsigned long int rhs );
31  Int --?( Int * lhs );
32  Int ?--( Int * lhs );
33
34  Int ?*( Int multiplicator, Int multiplicand );
35  Int ?*( Int multiplicator, long int multiplicand );
36  Int ?*( long int multiplicand, Int multiplicator );
37  Int ?*( Int multiplicator, unsigned long int multiplicand );
38  Int ?*( unsigned long int multiplicand, Int multiplicator );
39  Int ?*=( Int * lhs, Int rhs );
40  Int ?*=( Int * lhs, long int rhs );
41  Int ?*=( Int * lhs, unsigned long int rhs );
42
43  Int ?/( Int dividend, Int divisor );
44  Int ?/( Int dividend, unsigned long int divisor );
45  Int ?/( unsigned long int dividend, Int divisor );
46  Int ?/( Int dividend, long int divisor );
47  Int ?/( long int dividend, Int divisor );
48  Int ?/=( Int * lhs, Int rhs );
49  Int ?/=( Int * lhs, long int rhs );
50  Int ?/=( Int * lhs, unsigned long int rhs );
51
52  [ Int, Int ] div( Int dividend, Int divisor );
53  [ Int, Int ] div( Int dividend, unsigned long int divisor );
54
55  Int ?%( Int dividend, Int divisor );
56  Int ?%( Int dividend, unsigned long int divisor );
57  Int ?%( unsigned long int dividend, Int divisor );
58  Int ?%( Int dividend, long int divisor );
59  Int ?%( long int dividend, Int divisor );
60  Int ?%=( Int * lhs, Int rhs );

```

```

1  Int ?%=( Int * lhs, long int rhs );
2  Int ?%=( Int * lhs, unsigned long int rhs );
3
4  Int ?<<?( Int shiften, mp_bitcnt_t shift );
5  Int ?<=<?( Int * lhs, mp_bitcnt_t shift );
6  Int ?>>?( Int shiften, mp_bitcnt_t shift );
7  Int ?>=>?( Int * lhs, mp_bitcnt_t shift );
8
9  Int abs( Int oper );           // number functions
10 Int fact( unsigned long int N );
11 Int gcd( Int oper1, Int oper2 );
12 Int pow( Int base, unsigned long int exponent );
13 Int pow( unsigned long int base, unsigned long int exponent );
14 void srandom( gmp_randstate_t state );
15 Int random( gmp_randstate_t state, mp_bitcnt_t n );
16 Int random( gmp_randstate_t state, Int n );
17 Int random( gmp_randstate_t state, mp_size_t max_size );
18 int sgn( Int oper );
19 Int sqrt( Int oper );
20
21 forall( istype & | istream( istype ) ) istype * ?|( istype * is, Int * mp ); // I/O
22 forall( ostype & | ostream( ostype ) ) ostype * ?|( ostype * os, Int mp );

```

23 Figure 18 shows C \forall and C factorial programs using the GMP interfaces. (Compile with flag `-lgmp` to link with the
24 GMP library.)

25 K Rational Numbers

26 Rational numbers are numbers written as a ratio, *i.e.*, as a fraction, where the numerator (top number) and the denomi-
27 nator (bottom number) are whole numbers. When creating and computing with rational numbers, results are constantly
28 reduced to keep the numerator and denominator as small as possible.

```

29 // implementation
30 struct Rational {
31     long int numerator, denominator; // invariant: denominator > 0
32 }; // Rational
33
34 Rational rational();           // constructors
35 Rational rational( long int n );
36 Rational rational( long int n, long int d );
37 void ?{}( Rational * r, zero_t );
38 void ?{}( Rational * r, one_t );
39
40 long int numerator( Rational r ); // numerator/denominator getter/setter
41 long int numerator( Rational r, long int n );
42 long int denominator( Rational r );
43 long int denominator( Rational r, long int d );
44
45 int ?==( Rational l, Rational r ); // comparison
46 int ?!==( Rational l, Rational r );
47 int ?<?( Rational l, Rational r );
48 int ?<=( Rational l, Rational r );
49 int ?>?( Rational l, Rational r );
50 int ?>=( Rational l, Rational r );
51
52 Rational -( Rational r ); // arithmetic
53 Rational ?+?( Rational l, Rational r );
54 Rational ?-?( Rational l, Rational r );
55 Rational ?*( Rational l, Rational r );
56 Rational ?/( Rational l, Rational r );
57

```

| CV | C |
|---|---|
| <pre> #include <gmp.h> int main(void) { sout "Factorial Numbers"; int fact = 1; sout 0 fact; for (i; 40) { fact *= i; sout i fact; } } </pre> | <pre> #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); } } </pre> |
| <pre> 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 112400072777607680000 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 8159152832478977343456112695961158942720000000000 </pre> | |

Figure 18: Multi-precision Factorials

```

1  double widen( Rational r );           // conversion
2  Rational narrow( double f, long int md );
3
4  forall( istype & | istream( istype ) ) istype * ?|( istype *, Rational * ); // I/O
5  forall( ostype & | ostream( ostype ) ) ostype * ?|( ostype *, Rational );

```

6 References

- 7 [1] Ada16. *Ada Reference Manual ISO/IEC 8652:2012(E) with COR.1:2016*. AXE Consultants, Madison WI, USA,
8 3rd with technical corrigendum 1 for ada 2012 edition, 2016. <https://docs.adacore.com/live/wave/arm12/pdf/arm12/arm-12.pdf>. 2
- 10 [2] Richard C. Bilson. Implementing overloading and polymorphism in C \forall . Master's thesis, University of Waterloo,
11 Waterloo, Ontario, Canada, 2003. <http://plg.uwaterloo.ca/theses/BilsonThesis.pdf>. 2, 49
- 12 [3] Walter Bright and Andrei Alexandrescu. *D Programming Language*. Digital Mars, Vienna Virginia, USA, 2016.
13 <http://dlang.org/spec/spec.html>. 2
- 14 [4] P. A. Buhr. A case for teaching multi-exit loops to beginning programmers. *SIGPLAN Not.*, 20(11):14–22,
15 November 1985. 12
- 16 [5] P. A. Buhr, David Till, and C. R. Zarnke. Assignment as the sole means of updating objects. *Softw. Pract. Exper.*,
17 24(9):835–870, September 1994. 2, 42
- 18 [6] Cobol14. *Programming Languages – Cobol ISO/IEC 1989:2014*. International Standard Organization, Geneva,
19 Switzerland, 2nd edition, 2014. <https://www.iso.org/standard/51416.html>. 2
- 20 [7] G. V. Cormack and A. K. Wright. Polymorphism in the compiled language ForceOne. In *Proceedings of the*
21 *20th Hawaii International Conference on Systems Sciences*, pages 284–292, January 1987. 2
- 22 [8] G. V. Cormack and A. K. Wright. Type-dependent parameter inference. *SIGPLAN Not.*, 25(6):127–136, June
23 1990. Proceedings of the ACM Sigplan'90 Conference on Programming Language Design and Implementation
24 June 20–22, 1990, White Plains, New York, USA. 2, 37
- 25 [9] Glen Jeffrey Ditchfield. *Contextual Polymorphism*. PhD thesis, Department of Computer Science, University of
26 Waterloo, Waterloo, Ontario, Canada, 1992. <http://plg.uwaterloo.ca/theses/DitchfieldThesis.pdf>. 2
- 27 [10] Tom Duff. Duff's device. <http://www.lysator.liu.se/c/duffs-device.html>, November 1983. 8, 9
- 28 [11] Dominic Duggan, Gordon V. Cormack, and John Ophel. Kindred type inference for parametric overloading. *Acta*
29 *Infomatica*, 33(1):21–68, 1996. 2
- 30 [12] Margaret A. Ellis and Bjarne Stroustrup. *The Annotated C++ Reference Manual*. Addison-Wesley, Boston, 1st
31 edition, 1990. 2
- 32 [13] Rodolfo Gabriel Esteves. C \forall , a study in evolutionary design in programming languages. Master's thesis, School
33 of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2004. [http://plg.uwaterloo.ca/theses/](http://plg.uwaterloo.ca/theses/EstevesThesis.pdf)
34 [EstevesThesis.pdf](http://plg.uwaterloo.ca/theses/EstevesThesis.pdf). 2
- 35 [14] Fortran08. *Programming Languages – Fortran Part 1:Base Language ISO/IEC 1539-1:2010*. International
36 Standard Organization, Geneva, Switzerland, 3rd edition, 2010. <https://www.iso.org/standard/50459.html>. 2
- 37 [15] Nissim Francez. Another advantage of keyword notation for parameter communication with subprograms. *Com-*
38 *munications of the ACM*, 20(8):604–605, August 1977. 36
- 39 [16] John Galletly. *OCCAM 2: Including OCCAM 2.1*. UCL (University College London) Press, London, 2nd edition,
40 1996. 33, 51
- 41 [17] *GNU Multiple Precision Arithmetic Library*. GNU, 2016. <https://gmp1b.org>. 95
- 42 [18] James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. *Java Language Specification*, Java SE 8
43 edition, 2015. 2
- 44 [19] Robert Griesemer, Rob Pike, and Ken Thompson. *Go Programming Language*. Google, Mountain View, CA,
1 USA, 2009. <https://go.dev/ref/spec>. 2, 71

- 2 [20] W. T. Hardgrave. Positional versus keyword parameter communication in programming languages. *SIGPLAN*
3 *Not.*, 11(5):52–58, May 1976. 36
- 4 [21] International Business Machines. *OS and DOS PL/I Reference Manual*, 1st edition, September 1981. Manual
5 GC26-3977-0. 29
- 6 [22] International Standard Organization, Geneva, Switzerland. *C Programming Language ISO/IEC 9889:2011-12*,
7 3rd edition, December 2011. <https://www.iso.org/standard/57853.html>. 2, 20, 22, 47, 74, 76
- 8 [23] International Standard Organization, Geneva, Switzerland. *C++ Programming Language ISO/IEC 14882:2014*,
9 4th edition, 2014. <https://www.iso.org/standard/64029.html>. 2, 73
- 10 [24] K. E. Iverson. *A Programming Language*. Wiley, New York, 1962. 29
- 11 [25] Kathleen Jensen and Niklaus Wirth. *Pascal User Manual and Report, ISO Pascal Standard*. Springer-Verlag,
12 4th edition, 1991. Revised by Andrew B. Mickel and James F. Miner. 16
- 13 [26] C. H. Lindsey and S. G. van der Meulen. *Informal Introduction to ALGOL 68*. North-Holland, London, 1977.
14 21
- 15 [27] Barbara Liskov, Russell Atkinson, Toby Bloom, Eliot Moss, J. Craig Schaffert, Robert Scheifler, and Alan Sny-
16 der. *CLU Reference Manual*, volume 114 of *Lecture Notes in Computer Science*. Springer, New York, 1981. 33,
17 51
- 18 [28] Python. *Python Language Reference, Release 3.7.2*. Python Software Foundation, [https://docs.python.org/3/](https://docs.python.org/3/reference/index.html)
19 [reference/index.html](https://docs.python.org/3/reference/index.html), 2018. 1
- 20 [29] Dennis M. Ritchie. The development of the C language. *SIGPLAN Not.*, 28(3):201–208, March 1993. 19
- 21 [30] *The Rust Reference*. Rust Developers, 2015. <https://doc.rust-lang.org/stable/reference>. 2
- 22 [31] *Scala Language Specification, Version 2.11*. École Polytechnique Fédérale de Lausanne, 2016. [http://www.](http://www.scala-lang.org/files/archive/spec/2.11)
23 [scala-lang.org/files/archive/spec/2.11](http://www.scala-lang.org/files/archive/spec/2.11). 1
- 24 [32] Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley Series in Computer Science. Addison-
25 Wesley, Boston, 1st edition, 1986. 1
- 26 [33] David W. Till. Tuples in imperative programming languages. Master’s thesis, Department of Computer Science,
27 University of Waterloo, Waterloo, Ontario, Canada, 1989. 2, 42, 43, 45
- 28 [34] TIOBE Index. <https://www.tiobe.com/tiobe-index>. 1
- 29 [35] Ben Werther and Damian Conway. A modest proposal: C++ resyntaxed. *SIGPLAN Not.*, 31(11):74–82, Novem-
30 ber 1996. 2

Index

- 31 *Italic page numbers give the location of the main entry for*
1 *the referenced term. Plain page numbers denote uses of*
2 *the indexed term. Entries for grammar non-terminals are*
3 *italicized. A typewriter font is used for grammar terminals,*
4 *and program identifiers.*
- 5 *!?*, **68**
6 *&&*, **7**
7 ***, **27**
8 **=*, **27**
9 **?*, **68**
10 *+*, **26**
11 *++?*, **68**
12 *+=*, **11, 26**
13 *+?*, **68**
14 *--?*, **68**
15 *--colors*, **4**
16 *--deterministic-out*, **4**
17 *--gdb*, **4**
18 *--help*, **4**
19 *--invariant*, **4**
20 *--libcfa*, **4**
21 *--linemarks*, **4**
22 *--no-linemarks*, **4**
23 *--no-main*, **4**
24 *--no-prelude*, **4**
25 *--prelude-dir*, **5**
26 *--print*, **4**
27 *--prototypes*, **4**
28 *--statistics*, **5**
29 *--tree*, **5**
30 *--*, **11**
31 *-?*, **68**
32 *-CFA*, **3**
33 *-L*, **4**
34 *-N*, **4**
35 *-P*, **4**
36 *-S*, **5**
37 *-XCFA*, **3, 4**
38 *-̄*, **11**
39 *-̄=*, **11**
40 *-c*, **4**
41 *-d*, **4**
42 *-debug*, **3**
43 *-g*, **4**
44 *-h*, **4**
45 *-help*, **3**
46 *-i*, **4**
47 *-l*, **4**
48 *-lgmp*, **98**
49 *-m*, **4**
50 *-n*, **4**
- nodebug*, **3**
-nohelp, **4**
-noquiet, **4**
-p, **4**
-quiet, **4**
-std=gnu11, **3**
-t, **5**
-~, **12**
-~=, **12**
<, **26**
<=, **26**
=, **26**
>, **26**
>=, **26**
?!=?, **68**
?%=?, **68**
?%?, **68**
?&=?, **68**
?&?, **68**
?(), **68**
?=?*, **68**
??*, **68**
?++, **68**
?+=?, **68**
?+?, **68**
?--, **68**
?-=?, **68**
?-?, **68**
?/=?, **68**
?/?, **68**
?<<=?, **68**
?<<?, **68**
?<=?, **68**
?<?, **68**
?==?, **68**
?=?, **68**
?>=?, **68**
?>>=?, **68**
?>>?, **68**
?>?, **68**
?[?], **68**
?=?**, **68**
??**, **68**
?^=?, **68**
?^?, **68**
?|=?, **68**
?|?, **68**
̄, **11**
̄=, **11**
__CFA_MAJOR__, **4**
__CFA_MINOR__, **4**

- 102 `__CFA_PATCH__`, 4
- 1 `__CFA__`, 4
- 2 `__CFORALL__`, 4
- 3 `__cforall`, 4
- 4 `~`, 11
- 5 `~=`, 12
- 6 `~?`, 68

- 7 `aalloc`, 78, 80, 81
- 8 `abs`, 83
- 9 `acos`, 86
- 10 `acosh`, 87
- 11 `address`
 - 12 `duality`, 21
- 13 `adelete`, 81
- 14 `align`, 81
- 15 `aligned_alloc`, 81
- 16 `alloc`, 78, 81
- 17 `amemalign`, 80, 81
- 18 `anew`, 81
- 19 `array`
 - 20 `reference`, 23
- 21 `ascending exclusive range`, 11
- 22 `ascending inclusive range`, 11
- 23 `ASCII`
 - 24 `extended`, 55
- 25 `ascodegen`, 4
- 26 `asin`, 86
- 27 `asinh`, 87
- 28 `assert.h`, 76
- 29 `assignment`
 - 30 `cascade`, 52
 - 31 `pointer`, 21
- 32 `ast`, 4
- 33 `asterr`, 4
- 34 `atan`, 86
- 35 `atan2`, 86
- 36 `atanh`, 87
- 37 `ato`, 82

- 38 `basetypeof`, 76
- 39 `bbox`, 5
- 40 `bcodegen`, 5
- 41 `Bilson, Richard`, 2
- 42 `bin`, 58
- 43 `break`, 7, 12
- 44 `break`
 - 45 `labelled`, 12
- 46 `bresolver`, 4
- 47 `bsearch`, 82
- 48 `bsearchl`, 82
- 49 `bsearchu`, 82

- 50 `C linkage`, 3
- 51 `C++`, 1–3, 23, 38, 70, 75, 77

- 52 `C11`, 2
- 53 `calloc`, 80
- 54 `cancellation`
 - 55 `pointer/reference`, 22
- 56 `cascade assignment`, 52
- 57 `case`, 7, 10
- 58 `catch`, 76
- 59 `catchResume`, 76
- 60 `cbrt`, 85
- 61 `ceil`, 88
- 62 `cfa`, 3, 4
- 63 `CV`, 1
- 64 `char *`, 25
- 65 `choose`, 9, 10, 76
- 66 `clamp`, 83
- 67 `clear`, 54
- 68 `clearerr`, 55
- 69 `close`, 54, 55
- 70 `cmemalign`, 80
- 71 `coercion`, 22
- 72 `compilation`
 - 73 `cfa`, 3
- 74 `compilation option`
 - 75 `-CFA`, 3
 - 76 `-XCFA`, 3
 - 77 `-debug`, 3
 - 78 `-help`, 3
 - 79 `-nodebug`, 3
 - 80 `-nohelp`, 4
 - 81 `-noquiet`, 4
 - 82 `-quiet`, 4
 - 83 `-std=gnu11`, 3
- 84 `complex.h`, 76
- 85 `constant`
 - 86 `underscore`, 5
- 87 `continue`, 12
- 88 `continue`
 - 89 `labelled`, 12
- 90 `copysign`, 89
- 91 `coroutine`, 76
- 92 `cos`, 85
- 93 `cosh`, 86
- 94 `ctordtor`, 5
- 95 `ctype.h`, 76

- 96 `declstats`, 4
- 97 `default`, 10
- 98 `default parameter`, 36
- 99 `delete`, 81
- 100 `descending exclusive range`, 11
- 101 `descending inclusive range`, 11
- 102 `disable`, 76
- 103 `Ditchfield, Glen`, 2
- 104 `div`, 40

- 105 **do**, 6, 11, 12, 15
- 1 dynamic storage-management, 77
- 2 **enable**, 76
- 3 eng, 58
- 4 eof, 54, 55
- 5 erf, 87
- 6 erfc, 87
- 7 errno.h, 76
- 8 Esteves, Rodolfo, 2
- 9 **exception**, 76
- 10 excl, 62
- 11 excpdecl, 4
- 12 exp, 84
- 13 exp2, 84
- 14 expm1, 84
- 15 exponentiation
 - 16 floating, 6
 - 17 logarithm, 6
 - 18 operator, 6
 - 19 unsigned integral, 6
- 20 expralt, 4
- 21 expranly, 5
- 22 extended **else**, 15
- 23 fail, 54, 55
- 24 **fallthrough**, 9, 10, 76
- 25 fdim, 84
- 26 fenv.h, 76
- 27 fill, 81
- 28 **finally**, 76
- 29 **fixup**, 76
- 30 float.h, 76
- 31 floor, 88
- 32 flush, 54, 55
- 33 fma, 84
- 34 fmod, 83
- 35 **for**, 6, 11, 12, 15
- 36 for control, 11
- 37 **forall**, 76
- 38 frexp, 89
- 39 fstream.hfa, 1, 4, 53
- 40 garbage-collection
 - 41 regional, 1
- 42 gcc, 3, 4, 69
- 43 **generator**, 76
- 44 getline, 61
- 45 getpid, 93
- 46 gmp, 95, 99
- 47 gmp.h, 76, 99
- 48 **goto**, 12
- 49 **goto**
 - 50 restricted, 12
- 51 hex, 58
- 52 hypot, 85
- 53 I/O
 - 54 common case, 52
 - 55 separator, 55
- 56 **if**, 6, 12, 15
- 57 ifstream
 - 58 `{}`, 55
 - 59 `^{}` , 55
 - 60 clearerr, 55
 - 61 close, 55
 - 62 eof, 55
 - 63 fail, 55
 - 64 open, 55
 - 65 read, 55
 - 66 unget, 55
- 67 ignore, 62
- 68 ilogb, 85
- 69 incl, 61
- 70 inheritance
 - 71 enumeration, 32
- 72 input, 52
- 73 instgen, 5
- 74 lnt, 95
- 75 **int128**, 76
- 76 inttypes.h, 76
- 77 iostream, 1
- 78 isacquire, 62
- 79 iso646.h, 76
- 80 istream
 - 81 isacquire, 62
- 82 labelled
 - 83 **break**, 12
 - 84 **continue**, 12
- 85 ldexp, 89
- 86 Lea, Doug, 33
- 87 left, 60
- 88 lgamma, 87
- 89 library
 - 90 stream, 52
- 91 limits.h, 76
- 92 locale
 - 93 stream, 63
- 94 locale.h, 76
- 95 log, 84
- 96 log10, 85
- 97 log1p, 85
- 98 log2, 85
- 99 logb, 85
- 100 loop predicate, 11
- 101 lrint, 88
- 102 lround, 89

- 103 malloc, 78, 80
- 1 malloc.h, 76
- 2 malloc_alignment, 80
- 3 malloc_expansion, 80
- 4 malloc_mmap_start, 80
- 5 malloc_size, 80
- 6 malloc_stats_clear, 80
- 7 malloc_stats_fd, 80
- 8 malloc_unfreed(), 80
- 9 malloc_zero_fill, 80
- 10 mangling
 - 11 name, 3, 76
- 12 manipulator
 - 13 bin, 58
 - 14 eng, 58
 - 15 excl, 62
 - 16 getline, 61
 - 17 hex, 58
 - 18 ignore, 62
 - 19 incl, 61
 - 20 left, 60
 - 21 nl, 57, 58
 - 22 nlOff, 57, 58
 - 23 nlOn, 57, 58
 - 24 nobase, 59
 - 25 nodp, 59
 - 26 nosep, 57
 - 27 oct, 58
 - 28 pad0, 60
 - 29 quote, 61
 - 30 sci, 58
 - 31 sep, 57
 - 32 sepGet, 56
 - 33 sepGetTuple, 57
 - 34 sepOff, 57
 - 35 sepOn, 57
 - 36 sepSet, 56
 - 37 sepSetTuple, 57
 - 38 sepTupleVal, 57
 - 39 sepVal, 56
 - 40 sign, 59
 - 41 skip, 60
 - 42 unit, 59
 - 43 upcase, 59
 - 44 wd, 59
 - 45 wdi, 60, 61
 - 46 ws, 60
- 47 mass, 43
- 48 math.h, 76
- 49 max, 83
- 50 memalign, 81
- 51 memcpy, 81
- 52 memset, 81
- 53 min, 83
- 54 modf, 40, 89
- 55 **monitor**, 76
- 56 multi-level exit, 12
- 57 multiple, 43
- 58 **mutex**, 76
- 59 named parameter, 36
- 60 nan, 84
- 61 nearbyint, 88
- 62 new, 81
- 63 nextafter, 89
- 64 nexttoward, 89
- 65 nl, 57, 58
- 66 nlOff, 57, 58
- 67 nlOn, 57, 58
- 68 nobase, 59
- 69 nodp, 59
- 70 nosep, 57
- 71 oct, 58
- 72 ofstream
 - 73 ?{}, 55
 - 74 ^?{}, 55
 - 75 clearerr, 55
 - 76 close, 55
 - 77 fail, 55
 - 78 flush, 55
 - 79 open, 55
 - 80 write, 55
- 81 **one_t**, 76
- 82 open, 54, 55
- 83 operator
 - 84 exponentiation, 6
- 85 **or**, 76
- 86 ostream
 - 87 **mutex**, 62
 - 88 osacquire, 62
- 89 output, 52
- 90 pad0, 60
- 91 parameter
 - 92 default, 36
 - 93 named, 36
- 94 parse, 4
- 95 pointer
 - 96 assignment, 21
 - 97 cancellation, 22
- 98 posix_memalign, 81
- 99 postfix function, 35
- 100 pow, 84
- 101 pow(x, y), 6
- 102 preprocessor variables
 - 103 __CFA_MINOR__, 4
 - 104 __CFA_PATCH__, 4
 - 105 __CFA__, 4

- 106 `__CFA__`, 4
- 1 `__CFORALL__`, 4
- 2 `__cforall`, 4
- 3 `pretty`, 4
- 4 `pvalloc`, 81
- 5 `qsort`, 82
- 6 `quote`, 61
- 7 `random`, 83
- 8 `Rational`, 98
- 9 `read`, 54, 55
- 10 `realloc`, 80, 81
- 11 `reallocarray`, 80
- 12 `remainder`, 83
- 13 `remquo`, 84
- 14 **report**, 76
- 15 `resize`, 80, 81
- 16 `rint`, 88
- 17 `round`, 88
- 18 `rproto`, 4
- 19 `rsteps`, 4
- 20 `scalbn`, 89
- 21 `scalbn`, 89
- 22 `sci`, 58
- 23 `sep`, 57
- 24 `sepGet`, 56
- 25 `sepGetTuple`, 57
- 26 `sepOff`, 57
- 27 `sepOn`, 57
- 28 `sepSet`, 56
- 29 `sepSetTuple`, 57
- 30 `sepTupleVal`, 57
- 31 `sepVal`, 56
- 32 `setjmp.h`, 76
- 33 `sign`, 59
- 34 `signal.h`, 76
- 35 `sin`, 85
- 36 `sinh`, 86
- 37 `skip`, 60
- 38 `sout`, 1
- 39 `sqrt`, 85
- 40 `srandom`, 83
- 41 `standard library`, 77
- 42 `static multi-level exit`, 12
- 43 `stdalign.h`, 76
- 44 `stdarg.h`, 76
- 45 `stdatomic.h`, 76
- 46 `stdbool.h`, 76
- 47 `stddef.h`, 76
- 48 `stdint.h`, 76
- 49 `stdio.h`, 1, 4, 76
- 50 `stdlib.h`, 76
- 51 `stdlib.hfa`, 93
- 52 `stdnoreturn.h`, 76
- 53 `storage management`, 77
- 54 `stream`
- 55 `locale`, 63
- 56 `stream library`, 52
- 57 `string`
- 58 `equality operators`, 26
- 59 `relational operators`, 26
- 60 `string`, 25
- 61 `string.h`, 76
- 62 `string.hfa`, 25
- 63 **suspend**, 76
- 64 `swap`, 83
- 65 **switch**, 7, 10, 12
- 66 `symevt`, 4
- 67 `tan`, 86
- 68 `tanh`, 86
- 69 `tgamma`, 87
- 70 `tgmath.h`, 76
- 71 `threads.h`, 76
- 72 **throw**, 76
- 73 **throwResume**, 76
- 74 `Tiemann, Michael`, 33
- 75 `Till, Dave`, 2
- 76 `time.h`, 76
- 77 **timeout**, 76
- 78 **trait**, 76
- 79 `transpiler option`
- 80 `--colors`, 4
- 81 `--deterministic-out`, 4
- 82 `--gdb`, 4
- 83 `--help`, 4
- 84 `--invariant`, 4
- 85 `--libcfa`, 4
- 86 `--linemarks`, 4
- 87 `--no-linemarks`, 4
- 88 `--no-main`, 4
- 89 `--no-prelude`, 4
- 90 `--print`, 4
- 91 `-print`
- 92 `ascodegen`, 4
- 93 `asterr`, 4
- 94 `ast`, 4
- 95 `bbox`, 5
- 96 `bcodegen`, 5
- 97 `bresolver`, 4
- 98 `ctordtor`, 5
- 99 `declstats`, 4
- 100 `excpdecl`, 4
- 101 `expralt`, 4
- 102 `expranly`, 5
- 103 `instgen`, 5
- 104 `parse`, 4

- 105 pretty, 4
- 1 rproto, 4
- 2 rsteps, 4
- 3 symevt, 4
- 4 tuple, 5
- 5 valdecl, 4
- 6 --prototypes, 4
- 7 --statistics
- 8 counters,heap,time,all,none, 5
- 9 --tree, 5
- 10 -L, 4
- 11 -N, 4
- 12 -P, 4
- 13 ascodegen, 4
- 14 asterr, 4
- 15 ast, 4
- 16 bbox, 5
- 17 bcodegen, 5
- 18 bresolver, 4
- 19 ctordtor, 5
- 20 declstats, 4
- 21 excpdecl, 4
- 22 expralt, 4
- 23 expranly, 5
- 24 instgen, 5
- 25 parse, 4
- 26 pretty, 4
- 27 rproto, 4
- 28 rsteps, 4
- 29 symevt, 4
- 30 tuple, 5
- 31 valdecl, 4
- 32 -S
- 33 counters,heap,time,all,none, 5
- 34 -XCFA, 4
- 35 -c, 4
- 36 -d, 4
- 37 -g, 4
- 38 -h, 4
- 39 -i, 4
- 40 -l, 4
- 41 -m, 4
- 42 -n, 4
- 43 -p, 4
- 44 -t, 5
- 45 trunc, 88
- 46 try, 76
- 47 tuple, 5
- 48 type hoisting, 38
- 49 uchar.h, 76
- 50 unget, 55
- 51 ungetc, 54
- 52 unistd.h, 76
- 53 unit, 59
- 54 upcase, 59
- 55 valdecl, 4
- 56 valloc, 81
- 57 **virtual**, 76
- 58 **waitfor**, 15, 76
- 59 wchar.h, 76
- 60 wctype.h, 76
- 61 wd, 59
- 62 wdi, 60, 61
- 63 **when**, 76
- 64 **while**, 6, 11, 12, 15
- 65 **with**, 76
- 66 write, 54, 55, 63
- 67 ws, 60
- 68 **zero_t**, 76