

# CV (Cforall) User Manual

## Version 1.0

“describe not prescribe”

### 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

February 13, 2024

© 2016, 2018, 2021, 2024 CV Project

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

1	<b>Contents</b>	
2	<b>1 Introduction</b>	<b>1</b>
3	1.1 Background . . . . .	1
4	<b>2 Why fix C?</b>	<b>1</b>
5	<b>3 History</b>	<b>2</b>
6	<b>4 Interoperability</b>	<b>2</b>
7	<b>5 CV Compilation</b>	<b>3</b>
8	<b>6 Backquote Identifiers</b>	<b>5</b>
9	<b>7 Constant Underscores</b>	<b>5</b>
10	<b>8 Exponentiation Operator</b>	<b>6</b>
11	<b>9 Control Structures</b>	<b>6</b>
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 . . . . .	11
17	9.6 Labelled continue / break Statement . . . . .	13
18	9.7 Extended else . . . . .	14
19	9.8 with Statement . . . . .	14
20	<b>10 Exception Handling</b>	<b>16</b>
21	10.1 Non-local Exception . . . . .	17
22	10.2 Exception Hierarchy . . . . .	17
23	<b>11 Alternative Declarations</b>	<b>18</b>
24	<b>12 Pointer / Reference</b>	<b>19</b>
25	12.1 Initialization . . . . .	22
26	12.2 Address-of Semantics . . . . .	23
27	12.3 Conversions . . . . .	24
28	<b>13 string Type</b>	<b>24</b>
29	13.1 Implicit String Conversions . . . . .	25
30	13.2 Size (length) . . . . .	25
31	13.3 Comparison Operators . . . . .	25
32	13.4 Concatenation . . . . .	25
33	13.5 Repetition . . . . .	26
34	13.6 Substring . . . . .	26
35	13.7 Searching . . . . .	27
36	13.8 Miscellaneous . . . . .	28
37	13.9 Returning N+1 on Failure . . . . .	28
38	13.10C Compatibility . . . . .	29
39	13.11 Input/Output Operators . . . . .	29
40	<b>14 Enumeration</b>	<b>30</b>
41	14.1 Enum type . . . . .	30
42	14.2 Inheritance . . . . .	31

1	<b>15 Routine Definition</b>	<b>32</b>
2	15.1 Named Return Values . . . . .	32
3	15.2 Routine Prototype . . . . .	33
4	<b>16 Routine Pointers</b>	<b>33</b>
5	<b>17 Named and Default Arguments</b>	<b>34</b>
6	<b>18 Unnamed Structure Fields</b>	<b>36</b>
7	<b>19 Nesting</b>	<b>36</b>
8	19.1 Type Nesting . . . . .	36
9	19.2 Routine Nesting . . . . .	37
10	<b>20 Tuple</b>	<b>38</b>
11	20.1 Multiple-Return-Value Functions . . . . .	38
12	20.2 Expressions . . . . .	39
13	20.3 Variables . . . . .	39
14	20.4 Indexing . . . . .	40
15	20.5 Flattening and Structuring . . . . .	40
16	20.6 Assignment . . . . .	41
17	20.7 Construction . . . . .	42
18	20.8 Member-Access Expression . . . . .	43
19	20.9 Casting . . . . .	43
20	20.10 Polymorphism . . . . .	44
21	20.10.1 Assertion Inference . . . . .	45
22	<b>21 Tuples</b>	<b>45</b>
23	21.1 Tuple Coercions . . . . .	47
24	21.2 Mass Assignment . . . . .	47
25	21.3 Multiple Assignment . . . . .	48
26	21.4 Cascade Assignment . . . . .	48
27	<b>22 Stream I/O Library</b>	<b>49</b>
28	22.1 Basic I/O . . . . .	49
29	22.1.1 Stream Output . . . . .	49
30	22.1.2 Stream Input . . . . .	50
31	22.1.3 Stream Files . . . . .	50
32	22.2 Implicit Separator . . . . .	51
33	22.3 Separation Manipulators . . . . .	53
34	22.4 Newline Manipulators . . . . .	54
35	22.5 Output Manipulators . . . . .	54
36	22.6 Input Manipulators . . . . .	56
37	22.7 Concurrent Stream Access . . . . .	58
38	22.8 Locale . . . . .	59
39	<b>23 String Stream</b>	<b>59</b>
40	<b>24 Structures</b>	<b>60</b>
41	<b>25 Constructors and Destructors</b>	<b>61</b>

1	<b>26 Overloading</b>	<b>62</b>
2	26.1 Constant . . . . .	63
3	26.2 Variable . . . . .	63
4	26.3 Function Overloading . . . . .	64
5	26.4 Operator . . . . .	64
6	<b>27 Auto Type-Inferencing</b>	<b>65</b>
7	<b>28 Concurrency</b>	<b>66</b>
8	28.1 Coroutine . . . . .	66
9	28.2 Monitors . . . . .	66
10	28.3 Threads . . . . .	67
11	<b>29 Language Comparisons</b>	<b>67</b>
12	29.1 C++ . . . . .	67
13	29.2 Go . . . . .	69
14	29.3 Rust . . . . .	69
15	29.4 D . . . . .	69
16	<b>A Syntax Ambiguities</b>	<b>69</b>
17	<b>B C Incompatibles</b>	<b>70</b>
18	<b>C <math>\mathbf{cv}</math> Keywords</b>	<b>73</b>
19	<b>D Standard Headers</b>	<b>73</b>
20	<b>E Standard Library</b>	<b>73</b>
21	E.1 Storage Management . . . . .	73
22	E.2 String to Value Conversion . . . . .	76
23	E.3 Search / Sort . . . . .	76
24	E.4 Absolute Value . . . . .	77
25	E.5 C Random Numbers . . . . .	77
26	E.6 Algorithms . . . . .	77
27	<b>F Math Library</b>	<b>77</b>
28	F.1 General . . . . .	77
29	F.2 Exponential . . . . .	78
30	F.3 Logarithm . . . . .	78
31	F.4 Trigonometric . . . . .	79
32	F.5 Hyperbolic . . . . .	80
33	F.6 Error / Gamma . . . . .	81
34	F.7 Nearest Integer . . . . .	81
35	F.8 Manipulation . . . . .	83
36	<b>G Time Keeping</b>	<b>83</b>
37	G.1 Duration . . . . .	83
38	G.2 timeval . . . . .	85
39	G.3 timespec . . . . .	85
40	G.4 itimerval . . . . .	85
41	G.5 Time . . . . .	85
42	<b>H Clock</b>	<b>86</b>
43	H.1 C time . . . . .	86
44	H.2 Clock . . . . .	86

1	<b>I Pseudo Random Number Generator</b>	<b>87</b>
2	<b>J Multi-precision Integers</b>	<b>89</b>
3	<b>K Rational Numbers</b>	<b>92</b>
4	<b>Index</b>	<b>95</b>

# 1 Introduction

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

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

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

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

	C	C <sup>V</sup>	C++
	<code>#include &lt;stdio.h&gt;</code>	<code>#include &lt;fstream&gt;</code>	<code>#include &lt;iostream&gt;</code> <code>using namespace std;</code>
24	<code>int main( void ) {</code> <code>    int x = 0, y = 1, z = 2;</code> <code>    printf( "%d %d %d\n", x, y, z );</code> <code>}</code>	<code>int main( void ) {</code> <code>    int x = 0, y = 1, z = 2;</code> <code>    sout   x   y   z;</code> <code>}</code>	<code>int main() {</code> <code>    int x = 0, y = 1, z = 2;</code> <code>    cout &lt;&lt; x &lt;&lt; ' ' &lt;&lt; y &lt;&lt; ' ' &lt;&lt; z &lt;&lt; endl;</code> <code>}</code>

While C<sup>V</sup> I/O (see Section 22, p. 49) looks similar to C++, there are important differences, such as automatic spacing between variables and an implicit newline at the end of the expression list, similar to Python [26]. In general, C<sup>V</sup> programs are 10% to 30% shorter than their equivalent C/C++ counterparts.

## 1.1 Background

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

## 2 Why fix C?

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

<sup>1</sup>Pronounced “C-for-all”, and written C<sup>V</sup>, CFA, or Cforall.

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

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

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

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

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

### 3 History

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

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

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

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

### 4 Interoperability

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

<sup>2</sup>Two important existing problems addressed were changing the type of character literals from **int** to **char** and enumerator from **int** to the type of its enumerators.

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

However, it is necessary to differentiate between C and C<sub>V</sub> code because of name overloading, as for C++. For example, the C math-library provides the following routines for computing the absolute value of the basic types: `abs`, `labs`, `llabs`, `fabs`, `fabsf`, `fabsl`, `cabsf`, `cabs`, and `cabsl`. Whereas, C<sub>V</sub> 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

```

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

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

## 5 C<sub>V</sub> Compilation

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

The command `cfa` is used to compile a C<sub>V</sub> program and is based on the GNU `gcc` command, *e.g.*:

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

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

C<sub>V</sub> has the following `gcc` flags turned on:

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

–fgnu89-inline Use the traditional GNU semantics for inline routines in C11 mode, which allows inline routines in header files.

C<sub>V</sub> has the following new options:

–CFA Only the C preprocessor (flag `–E`) and the C<sub>V</sub> translator steps are performed and the transformed program is written to standard output, which makes it possible to examine the code generated by the C<sub>V</sub> translator. The generated code starts with the standard C<sub>V</sub> prelude.

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

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

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



### *program behaviour or termination.*

`-help` Information about the set of C<sub>V</sub> compilation flags is printed.

`-nohelp` Information about the set of C<sub>V</sub> compilation flags is not printed. **This option is the default.**

`-quiet` The C<sub>V</sub> compilation message is not printed at the beginning of a compilation.

`-noquiet` The C<sub>V</sub> compilation message is printed at the beginning of a compilation. **This option is the default.**

The following preprocessor variables are available:

`__CFA_MAJOR__` is available during preprocessing and its value is the major version number of C<sub>V</sub>.<sup>3</sup>

`__CFA_MINOR__` is available during preprocessing and its value is the minor version number of C<sub>V</sub>.

`__CFA_PATCH__` is available during preprocessing and its value is the patch level number of C<sub>V</sub>.

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

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

For example, to toggle between C and C<sub>V</sub> extensions, use the following:

```
#ifndef __CFORALL__
#include <stdio.h>           // C header file
#else
#include <fstream>          // CV header file
#endif
```

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

The C<sub>V</sub> transpiler has multiple internal steps. The following flags control how the C<sub>V</sub> transpiler works, the stages run, and printing within a stage. The majority of these flags are used by C<sub>V</sub> developers, but some are occasionally useful to programmers. Each option must be escaped with `-XCFA` to direct it to the C<sub>V</sub> compilation step, similar to the `-Xlinker` flag for the linker, *e.g.*:

```
cfa test.cfa -CFA -XCFA -p # print translated code without printing the standard prelude
cfa test.cfa -XCFA -P -XCFA parse -XCFA -n # show program parse without prelude
```

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

```
cfa test.cfa -XCFA,-Pparse,-n # show program parse without prelude
```

`-c`, `--colors` diagnostic color: never, always, auto

`-g`, `--gdb` wait for gdb to attach

`-h`, `--help` print transpiler help message

`-i`, `--invariant` invariant checking during AST passes

`-l`, `--libcfa` generate libcfa.c

`-L`, `--linemarks` generate line marks

`-m`, `--no-main` do not replace main

`-N`, `--no-linemarks` do not generate line marks

`-n`, `--no-prelude` do not read prelude

`-p`, `--prototypes` do not generate prelude prototypes ⇒ prelude not printed

`-d`, `--deterministic-out` only print deterministic output

`-P`, `--print` one of:

`ascodegen` print AST as codegen rather than AST

`asterr` print AST on error

`declstats` print code property statistics

`parse` print yacc (parsing) debug information

`pretty` prettyprint for ascodegen flag

`rproto` resolver-`proto` instance

`rsteps` print resolver steps

`ast` print AST after parsing

`excpdecl` print AST after translating exception decls

`symevt` print AST after symbol table events

<sup>3</sup>The C preprocessor allows only integer values in a preprocessor variable so a value like “1.0.0” is not allowed. Hence, the need to have three variables for the major, minor and patch version number.

```

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

```

Figure 1: Header-File Interposition

```

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

```

## 6 Backquote Identifiers

CV introduces several new keywords (see Section C, p. 73) that can clash with existing C variable-names in legacy code. Keyword clashes are accommodated by syntactic transformations using the CV backquote escape-mechanism:

```

16     int ``coroutine = 3;           // make keyword an identifier
17     double ``forall = 3.5;

```

Existing C programs with keyword clashes can be converted by prefixing the keyword identifiers with double back-quotes, and eventually the identifier name can be changed to a non-keyword name. Figure 1 shows how clashes in existing C header-files (see Section D, p. 73) can be handled using preprocessor *interposition*: `#include_next` and command-line `-I filename`. Several common C header-files with keyword clashes are fixed in the standard CV header-library, so there is largely a seamless programming-experience.

## 7 Constant Underscores

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

```

25     2_147_483_648;                // decimal constant
26     56_ul;                        // decimal unsigned long constant
27     0_377;                        // octal constant
28     0x_ff_ff;                     // hexadecimal constant
29     0x_ef3d_aa5c;                  // hexadecimal constant
30     3.141_592_654;                 // floating constant
31     10_e+1_00;                     // floating constant
32     0x_ff_ff_p_3;                  // hexadecimal floating
33     0x_1.ffff_ffff_p_128_l;        // hexadecimal floating long constant
34     L_"\x_ff_ee";                  // wide character constant

```

The rules for placement of underscores are:

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

2. Underscores may only appear within a sequence of digits (regardless of the digit radix). In other words, an underscore cannot start or end a sequence of digits, *e.g.*, `_1`, `1_` and `_1_` are invalid (actually, the 1st and 3rd 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 difficulties parsing for IDEs between character and numeric constants, as quotes are no longer balanced (`'x'` and `3.14'159`).

## 8 Exponentiation Operator

C, C++, and Java (and other programming languages) have *no* exponentiation operator, *i.e.*,  $x^y$ , and instead use a routine, like `pow(x,y)`, to perform the exponentiation operation. CV extends the basic operators with the exponentiation operator `??` and `??=`, as in, `x ?? y` and `x ??= y`, which means  $x^y$  and  $x \leftarrow x^y$ . The priority of the exponentiation operator is between the cast and multiplicative operators, so that `w * (int)x \ (int)y * z` is parenthesized as `(w * (((int)x) \ ((int)y))) * z`.

There are exponentiation operators 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 necessary 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. 64).

## 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.<sup>4</sup>

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

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

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

```

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

```

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

## 7 9.2 case Clause

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

	C	CV
12	<pre> switch ( i ) {   case 1: case 3 : case 5:     ...   case 2: case 4 : case 6:     ... } </pre>	<pre> switch ( i ) {   case 1, 3, 5: // odd values     ...   case 2, 4, 6: // even values     ... } </pre>

13 In addition, subranges are allowed to specify a contiguous set of case values.

	C	CV	gcc
14	<pre> switch ( i ) {   case 1: case 2: case 3: case 4:     ...   case 10: case 11: case 12: case 13:     ... } </pre>	<pre> switch ( i ) {   case 1~4:     ...   case 10~13:     ... } </pre>	<pre> switch ( i ) {   case 1_...4: // 1, 2, 3, 4     ...   case 10_...13: // 10, 11, 12, 13     ... } </pre>

15 While gcc has the same range mechanism, it has an awkward syntax, 2\_...42, because a space is required after a  
 16 number, otherwise the period is a decimal point.

17 CV also allows lists of subranges.

```

18 case 1~5, 12~21, 35~42:

```

## 19 9.3 switch Statement

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

21 1. By default, the end of a **case** clause<sup>5</sup> *falls through* to the next **case** clause in the **switch** statement; to exit a  
 22 **switch** statement from a **case** clause requires explicitly terminating the clause with a transfer statement, most  
 23 commonly **break**:

```

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

```

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

<sup>5</sup>In 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 programming errors as programmers often *forget* the **break** statement at the end of a **case** clause, resulting in inadvertent fall-through.

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

```

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

```

This usage branches into control structures, which is known to cause both 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 this example:

```

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

```

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

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

discovery. [10]

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

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

```

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

```

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

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

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

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

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

```

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

```

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

```

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

```

```

1      default:
2          j = 3;
3      }

```

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 **fallthru** 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** clause<sup>6</sup> 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.

```

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

```

## 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> 35      choose ( ... ) { 36          case 3: 37              if ( ... ) { 38                  ... fallthru; // goto case 4 39              } else { 40                  ... 41              } 42              // implicit break 43          case 4: </pre>	<pre> 35      choose ( ... ) { 36          case 3: 37              ... fallthrough common; 38          case 4: 39              ... fallthrough common; 40          common: // below fallthrough 41                  // at case-clause level 42              ... // common code for cases 3/4 43              // implicit break 44          case 4: </pre>	<pre> 35      choose ( ... ) { 36          case 3: 37              choose ( ... ) { 38                  case 4: 39                      for ( ... ) { 40                          // multi-level transfer 41                          ... fallthru common; 42                      } 43                  ... 44              } 45          common: // below fallthrough 46                  // 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.

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



## 1 9.5 Loop Control

2 Looping a fixed number of times, possibly with a loop index, occurs frequently. CV condenses simply looping to  
3 facilitate coding speed and safety (see examples in Figure 2).

4 The **for**, **while**, and **do** loop-control are extended to allow an empty conditional, which implies a comparison value  
5 of 1 (true).

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

9 The **for** control is extended with 4 new loop-control operators, which are not overloadable:

```
10 ~ (tilde) up-to exclusive range,
11 ~= (tilde equal) up-to inclusive range,
12 -~ (minus tilde) down-to exclusive range,
13 -~= (minus tilde equal) down-to inclusive range.
```

14 The **for** index components, low (L), high (H), and increment (I), are optional. If missing:

- 15 • 0 is the implicit low value.
- 16 • 1 is the implicit increment value.
- 17 • The up-to range uses operator += for increment.
- 18 • The down-to range uses operator -= for decrement.

19 If no type is specified for the loop index, it is the type of the high value H (when the low value is implicit) or the  
20 low value L.

```
21 for ( 5 )                       // typeof(5) anonymous-index; 5 is high value
22 for ( i; 1.5 ~ 5.5 )            // typeof(1.5) i; 1.5 is low value
23 for ( int i; 0 ~ 10 ~ 2 )       // int i; type is explicit
```

24 The following are examples for constructing different for-control:

- 25 • H is implicit up-to exclusive range [0,H).

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

- 27 • ~= H is implicit up-to inclusive range [0,H].

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

- 29 • L ~ H is explicit up-to exclusive range [L,H).

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

- 31 • L ~= H is explicit up-to inclusive range [L,H].

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

- 33 • L -~ H is explicit down-to exclusive range [H,L), where L and H are implicitly interchanged to make the range  
34 down-to.

```
35 for ( 1 -~ 5 )                 // for ( typeof(1) i = 5; i > 0; i -= 1 )
```

- 36 • L -~= H is explicit down-to inclusive range [H,L], where L and H are implicitly interchanged to make the range  
37 down-to.

```
38 for ( 1 -~= 5 )                // for ( typeof(1) i = 5; i >= 0; i -= 1 )
```

- 39 • @ means put nothing in this field.

```
40 for ( i; 1 ~ @ ~ 2 )           // for ( typeof(1) i = 1; /* empty */; i += 2 )
41 for ( i; 1 ~ 10 ~ @ )          // for ( typeof(1) i = 1; i < 10; /* empty */ )
42 for ( i; 1 ~ @ ~ @ )           // for ( typeof(1) i = 1; /*empty*/; /* empty */ )
```

43 L cannot be elided for the up-to range, @ ~ 5, and H for the down-to range, 1 -~ @, because then the loop index is  
44 uninitialized. Similarly, the high value cannot be elided is an anonymous loop index (1 ~ @), as there is no index  
45 to stop the loop.

- 46 • : means add another index.

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

48 **Warning:** specifying the down-to range maybe unexpected because the loop control *implicitly* switches the L and H



loop control		output
<b>while</b> () { <b>sout</b>   "empty"; <b>break</b> ; }		empty
<b>do</b> { <b>sout</b>   "empty"; <b>break</b> ; } <b>while</b> ();		empty
<b>for</b> () { <b>sout</b>   "empty"; <b>break</b> ; }	<i>sout / nl / nlOff;</i>	empty
<b>for</b> ( 0 ) { <b>sout</b>   "A"; } <b>sout</b>   "zero";	<i>sout / nl;</i>	zero
<b>for</b> ( 1 ) { <b>sout</b>   "A"; }	<i>sout / nl;</i>	A
<b>for</b> ( 10 ) { <b>sout</b>   "A"; }	<i>sout / nl;</i>	A A A A A A A A A A
<b>for</b> ( ~ = 10 ) { <b>sout</b>   "A"; }	<i>sout / nl;</i>	A A A A A A A A A A A
<b>for</b> ( 1 ~ = 10 ~ 2 ) { <b>sout</b>   "B"; }	<i>sout / nl;</i>	B B B B B
<b>for</b> ( 1 ~ = 10 ~ 2 ) { <b>sout</b>   "C"; }	<i>sout / nl;</i>	C C C C C
<b>for</b> ( 0.5 ~ 5.5 ) { <b>sout</b>   "D"; }	<i>sout / nl;</i>	D D D D D
<b>for</b> ( 0.5 ~ 5.5 ) { <b>sout</b>   "E"; }	<i>sout / nl;</i>	E E E E E
<b>for</b> ( i; 10 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	0 1 2 3 4 5 6 7 8 9
<b>for</b> ( i; ~ = 10 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	0 1 2 3 4 5 6 7 8 9 10
<b>for</b> ( i; 1 ~ = 10 ~ 2 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	1 3 5 7 9
<b>for</b> ( i; 1 ~ = 10 ~ 2 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	10 8 6 4 2
<b>for</b> ( i; 0.5 ~ 5.5 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	0.5 1.5 2.5 3.5 4.5
<b>for</b> ( i; 0.5 ~ 5.5 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	5.5 4.5 3.5 2.5 1.5
<b>for</b> ( ui; 2u ~ = 10u ~ 2u ) { <b>sout</b>   ui; }	<i>sout / nl;</i>	2 4 6 8 10
<b>for</b> ( ui; 2u ~ = 10u ~ 2u ) { <b>sout</b>   ui; }	<i>sout / nl / nl / nl;</i>	10 8 6 4 2
<b>enum</b> { N = 10 }; <b>for</b> ( N ) { <b>sout</b>   "N"; }	<i>sout / nl;</i>	N N N N N N N N N N
<b>for</b> ( i; N ) { <b>sout</b>   i; }	<i>sout / nl;</i>	0 1 2 3 4 5 6 7 8 9
<b>for</b> ( i; ~ N ) { <b>sout</b>   i; }	<i>sout / nl / nl / nl;</i>	10 9 8 7 6 5 4 3 2 1
<b>const int</b> low = 3, high = 10, inc = 2; <b>for</b> ( i; low ~ high ~ inc + 1 ) { <b>sout</b>   i; }	<i>sout / nl;</i>	3 6 9
<b>for</b> ( i; 1 ~ @ ) { <b>if</b> ( i > 10 ) <b>break</b> ; <b>sout</b>   i; }	<i>sout / nl;</i>	1 2 3 4 5 6 7 8 9 10
<b>for</b> ( i; @ ~ 10 ) { <b>if</b> ( i < 0 ) <b>break</b> ; <b>sout</b>   i; }	<i>sout / nl;</i>	10 9 8 7 6 5 4 3 2 1 0
<b>for</b> ( i; 2 ~ @ ~ 2 ) { <b>if</b> ( i > 10 ) <b>break</b> ; <b>sout</b>   i; }	<i>sout / nl;</i>	2 4 6 8 10
<b>for</b> ( i; 2.1 ~ @ ~ @ ) { <b>if</b> ( i > 10.5 ) <b>break</b> ; <b>sout</b>   i; i += 1.7; } <i>sout / nl;</i>		2.1 3.8 5.5 7.2 8.9
<b>for</b> ( i; @ ~ 10 ~ 2 ) { <b>if</b> ( i < 0 ) <b>break</b> ; <b>sout</b>   i; }	<i>sout / nl;</i>	10 8 6 4 2 0
<b>for</b> ( i; 12.1 ~ @ ~ @ ) { <b>if</b> ( i < 2.5 ) <b>break</b> ; <b>sout</b>   i; i -= 1.7; } <i>sout / nl;</i>		12.1 10.4 8.7 7.5 3.6
<b>for</b> ( i; 5 : j; -5 ~ @ ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -4 2 -3 3 -2 4 -1
<b>for</b> ( i; 5 : j; @ ~ -5 ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -6 2 -7 3 -8 4 -9
<b>for</b> ( i; 5 : j; -5 ~ @ ~ 2 ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -3 2 -1 3 1 4 3
<b>for</b> ( i; 5 : j; @ ~ -5 ~ 2 ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -7 2 -9 3 -11 4 -13
<b>for</b> ( i; 5 : j; -5 ~ @ ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -4 2 -3 3 -2 4 -1
<b>for</b> ( i; 5 : j; @ ~ -5 ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -6 2 -7 3 -8 4 -9
<b>for</b> ( i; 5 : j; -5 ~ @ ~ 2 ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -3 2 -1 3 1 4 3
<b>for</b> ( i; 5 : j; @ ~ -5 ~ 2 ) { <b>sout</b>   i   j; }	<i>sout / nl;</i>	0 -5 1 -7 2 -9 3 -11 4 -13
<b>for</b> ( i; 5 : j; @ ~ -5 ~ 2 : k; 1.5 ~ @ ) { <b>sout</b>   i   j   k; } <i>sout / nl;</i>		0 -5 1.5 1 -7 2.5 2 -9 3.5 3 -11 4.5 4 -13 5.5
<b>for</b> ( i; 5 : j; @ ~ -5 ~ 2 : k; 1.5 ~ @ ) { <b>sout</b>   i   j   k; } <i>sout / nl;</i>		0 -5 1.5 1 -7 2.5 2 -9 3.5 3 -11 4.5 4 -13 5.5
<b>for</b> ( i; 5 : k; 1.5 ~ @ : j; @ ~ -5 ~ 2 ) { <b>sout</b>   i   j   k; } <i>sout / nl;</i>		0 -5 1.5 1 -7 2.5 2 -9 3.5 3 -11 4.5 4 -13 5.5

Figure 2: Loop Control Examples




Figure 3: Multi-level Exit

1 values (and toggles the increment/decrement for I):

```

2     for ( i; 1 ~ 10 )           // up range
3     for ( i; 1 ~ 10 )           // down range
4     for ( i; 10 ~ 1 )           // WRONG down range!

```

5 The reason for this semantics is that the range direction can be toggled by adding/removing the minus, '-', versus  
6 interchanging the L and H expressions, which has a greater chance of introducing errors.

## 7 9.6 Labelled **continue** / **break** Statement

8 **C continue** and **break** statements are restricted to one level of nesting for a particular control structure. This restriction  
9 forces programmers to use **goto** to achieve the equivalent control-flow for more than one level of nesting. To prevent  
10 having to switch to the **goto**, CV extends the **continue** and **break** with a target label to support static multi-level exit [4],  
11 as in Java. For both **continue** and **break**, the target label must be directly associated with a **for**, **while** or **do** statement; for  
12 **break**, the target label can also be associated with a **switch**, **if** or compound ({} statement. Figure 3 shows a comparison  
13 between labelled **continue** and **break** and the corresponding C equivalent using **goto** and labels. The innermost loop  
14 has 8 exit points, which cause continuation or termination of one or more of the 7 nested control-structures.

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

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

20 The advantage of the labelled **continue/break** is allowing static multi-level exits without having to use the **goto** state-  
21 ment, and tying control flow to the target control structure rather than an arbitrary point in a program via a label.  
22 Furthermore, the location of the label at the *beginning* of the target control structure informs the reader (eye candy)  
23 that complex control-flow is occurring in the body of the control structure. With **goto**, the label is at the end of the  
24 control structure, which fails to convey this important clue early enough to the reader. Finally, using an explicit target

for the transfer, instead of an implicit target, allows new constructs to be added or removed without affecting existing constructs. Otherwise, the implicit targets of the current **continue** and **break**, *i.e.*, the closest enclosing loop or **switch**, change as certain constructs are added or removed.

## 9.7 Extended else

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

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

```

11  int a[10];

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

13  for ( i; 10 ) {
    if ( a[i] == 3 ) break; // found
  } else { // i == 10
    sout | "not found";
  }

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

```

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

## 9.8 with Statement

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

```

17  struct Person {           // aggregate
18      struct Name {         // nesting
19          char first[20], last[20];
20      } name;
21      struct Address {      // nesting
22          ...
23      } address;
24      int sex;
25  };

```

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

```

27  Person p
28  p.name ...; p.address ...; p.sex ...; // access containing fields

```

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

```

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

```

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

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

```

38  struct S {
39      struct /* unnamed */ { int g, h; } __attribute__(( aligned(64) ));
40      int tag;
41      union /* unnamed */ {
42          struct { char c1, c2; } __attribute__(( aligned(128) ));
43          struct { int i1, i2; };
44          struct { double d1, d2; };
45      };

```

```

} s;
enum { R, G, B };
s.g; s.h; s.tag = R; s.c1; s.c2; s.i1 = G; s.i2 = B; s.d1; s.d2;

```

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

```
struct S {
    char c; int i; double d;
    void f( /* S * this */ ) {
        c; i; d;
    }
}
```

// implicit "this" parameter  
// this->c; this->i; this->d;

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.

```
struct T {
    char m; int i; double n;           // derived class variables
};

struct S : public T {
    char c; int i; double d;           // class variables
    void g( double d, T & t ) {
        d; t.m; t.i; t.n;             // function parameter
        c; i; this->d; S::d;           // class S variables
        m; T::i; n;                   // class T variables
    }
};
```

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 [23, § 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.

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

with the generality of opening multiple aggregate-parameters:

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

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.

```
S * sp;
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. 63) and use the left-side of assignment in type resolution means most fields with the same name but different types are automatically disambiguated, eliminating qualification. All expressions in the expression list are open in parallel within the compound statement. This semantic is different from Pascal, which nests the openings from left to right. The difference between parallel and nesting occurs for fields with the same name and type:

```

struct Q { int i; int k; int m; } q, w;
struct R { int i; int j; double m; } r, w;
with ( r, q ) {
    j + k;           // unambiguous, r.j + q.k
    m = 5.0;         // unambiguous, q.m = 5.0
}

```

```

1      m = 1;                      // unambiguous, r.m = 1
2      int a = m;                  // unambiguous, a = r.i
3      double b = m;              // unambiguous, b = q.m
4      int c = r.i + q.i;         // disambiguate with qualification
5      (double)m;                // disambiguate with cast
6  }
```

7 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.

```

9      with ( r ) {
10         i;                      // unambiguous, r.i
11         with ( q ) {
12             i;                  // unambiguous, q.i
13         }
14     }
```

15 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:

```

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

19 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*.

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

```

22     void f( S & s, char c ) with ( s ) {
23         s.c = c; i = 3; d = 5.5;    // initialize fields
24     }
```

25 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:

```

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

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

```

32     void f( S & s, char c ) with ( s ) with( params ) { // syntax not allowed, illustration only
33         s.c = c; i = 3; d = 5.5;
34     }
```

35 This implicit semantic matches with programmer expectation.

## 36 10 Exception Handling

37 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.

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

```

41     #include <Exception.hfa>
42     ExceptionDecl( E,           // must be global scope
43         ... // exception fields
44     );
45     try {
46         ...
47         if ( ... ) throwResume ExceptionInst( E, /* initialization */);
48         if ( ... ) throw ExceptionInst( E, /* initialization */);
49         ...
50     } catchResume( E * ) { // must be pointer
51         ...
52     } catch( E * ) {
53         ...
```

```

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

```

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

## 19 10.1 Non-local Exception


```

20 void main() {
21     try {
22         _Enable {
23             ... resume(); ...
24         }
25     } catchResume( E & ) { // should be reference
26         ...
27     } catch( E & ) {
28         ...
29     }
30 }

```

## 31 10.2 Exception Hierarchy

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



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

```

41 catch( Arithmetic )
42 matches all three derived exception-types: DivideByZero, Overflow, and Underflow. Because the propagation mechanisms perform a simple linear search of the handler clause for a guarded block, and selects the first matching handler, the order of catch clauses in the handler clause becomes important, e.g.:
43
44 try {
45     ...
46 } catch( Overflow ) { // must appear first
47     // handle overflow
48 }

```

```

1      } catch( Arithmetic )
2          // handle other arithmetic issues
3      }
4  Multiple derivation among exception is not supported.

```

## 11 Alternative Declarations

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

```

8      int * x[5]

```

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

```

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

```

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

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

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

C	Cv
<b>int</b> * x1 [5];	[5] * <b>int</b> x1;
<b>int</b> (*x2)[5];	* [5] <b>int</b> x2;
<b>int</b> (*f( <b>int</b> p )) [5];	[* [5] <b>int</b> ] f( <b>int</b> p );

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

C	Cv
<b>int</b> *x, *y;	* <b>int</b> x, y;

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

C	Cv
<b>int</b> *x, y;	* <b>int</b> x;
	<b>int</b> y;

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,<sup>7</sup> *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**:

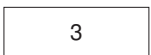
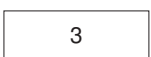
C	Cv
<code>y = (int *)x;</code>	<code>y = (* int)x;</code>
<code>i = sizeof(int * [ 5 ]);</code>	<code>i = sizeof([ 5 ] * int);</code>

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*.<sup>8</sup> 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>	$x$		<code>int</code>	<code>int * const x = (int *)100</code>
<code>x = 3;</code>	100			<code>*x = 3;</code> <i>// implicit dereference</i>
<code>int y;</code>	$y$		<code>int</code>	<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

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

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




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

	explicit variable address		implicit variable address
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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

```
int x, y, & r1, & r2, && r3;
&r1 = &x; // r1 points to x
&r2 = &r1; // r2 points to x
&r1 = &y; // r1 points to y
&&r3 = &&r2; // r3 points to r2
```

```
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 cancel-  
 8 lation always yields a value (rvalue).<sup>9</sup> For a C $\forall$  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; C $\forall$ -style declarations  
 47 (see Section 11, p. 18) attempt to address this issue:

C

C $\forall$ 

```
48   const int * const * const ccp;    const * const * const int ccp;
    const & const & const int ccr;    const & const & const int ccr;
```

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

where the **CV** declaration is read left-to-right.

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

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

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 reference<sup>10</sup> 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.

## 12.1 Initialization

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, for pointer/reference initialization, the initializing value must be an address not a value.

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

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 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.

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

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.

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

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.

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

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

Here, the compiler passes the address to the literal **3** or the temporary for the expression **x + y**, knowing the argument cannot be changed through the parameter. The **&** before the constant/expression for the pointer-type parameter (**g**) is a **CV** extension necessary to type match and is a common requirement before a variable in **C** (e.g., **scanf**). Importantly,

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

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

`CV` extends this semantics to a mutable pointer/reference parameter, and the compiler implicitly creates the necessary temporary (copying the argument), which is subsequently pointed-to by the reference parameter and can be changed.<sup>11</sup>

```

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

```

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

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

```

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

```

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

```

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

```

because the value of the routine object is a routine literal, *i.e.*, the routine code is normally immutable during execution.<sup>13</sup> `CV` allows this additional use of references for routine objects in an attempt to give a more consistent meaning for them.

## 12.2 Address-of Semantics

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

- if `R` is an rvalue of type `T &1 ··· &r`, where  $r \geq 1$  references (& symbols), then `&R` has type `T *&2 ··· &r`, *i.e.*, `T` pointer with  $r - 1$  references (& symbols).
- if `L` is an lvalue of type `T &1 ··· &l`, where  $l \geq 0$  references (& symbols), then `&L` has type `T *&1 ··· &l`, *i.e.*, `T` pointer with  $l$  references (& symbols).

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 ***& (&rrx)
45 ppppx = &&&&rrrx; // starting from &&&rrrx, &&&&rrrx is an rvalue with type int **** (&rrrx)

```

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

```

47 int x, * px, ** ppx, *** pppx;

```

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

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

<sup>13</sup>Dynamic code rewriting is possible but only in special circumstances.

```

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

```

9  int x;
10 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.

```

14 int x, &r = x, f( int p );
15 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.

```

19 int x, &r = x, f( int & p );          // lvalue variable (int) convert to reference (int &)
20 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.

```

26 int x, & f( int & p );
27 f( x + 3 );                          // rvalue parameter (int) implicitly converts to lvalue temporary reference (int &)
28 &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

The CV string type is for manipulation of dynamically-size character-strings versus C **char \*** type for manipulation of statically-size null-terminated character-strings. That is, 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. Hence, a string declaration does not specify a maximum length; as a string dynamically grows and shrinks in size, so does its underlying storage. In contrast, a C string also dynamically grows and shrinks in size, but its underlying storage is fixed. 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 routine, such as `strcat`. Like C strings, the characters in a string are numbered starting from 0.

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

```

43 #include <string.h>
44 string s, peter, digit, alpha, punctuation, ifstmt;
45 int i;
46 peter = "PETER";
47 digit = "0123456789";
48 punctuation = " ( ) . , " ;
49 ifstmt = "IF (A > B) {";

```

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

Figure 4: Implicit Conversions to String

1 Note, the include file `string.h` to access type `string`.

## 2 13.1 Implicit String Conversions

3 The types **char**, **char \***, **int**, **double**, **\_Complex**, including different signness and sizes, implicitly convert to type `string`.  
 4 Figure 4 shows examples of implicit conversions between C strings, integral, floating-point and complex types to  
 5 `string`. A conversions can be explicitly specified:

```
6 s = string( "abc" );           // converts char * to string
7 s = string( 5 );              // converts int to string
8 s = string( 5.5 );            // converts double to string
```

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

<code>string s = "abcde";</code>		
<code>char cs[3];</code>		
<code>strncpy( cs, s, sizeof(cs) );</code>	<code>sout   cs;</code>	ab
<code>char * cp = s;</code>	<code>sout   cp;</code>	abcde
<code>delete( cp );</code>		
<code>cp = s + ' ' + s;</code>	<code>sout   cp;</code>	abcde abcde
<code>delete( cp );</code>		

## 13 13.2 Size (length)

14 The size operation returns the length of a string.

```
15 i = size( "" );               // i is assigned 0
16 i = size( "abc" );            // i is assigned 3
17 i = size( peter );            // i is assigned 5
```

## 18 13.3 Comparison Operators

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

## 21 13.4 Concatenation

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



```

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

```

### 3 13.5 Repetition

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

```

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

```

### 7 13.6 Substring

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

```

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

```

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

```

21     digit( 3, 3 ) = " ";               // digit is assigned "0156789"
22     digit( 4, 3 ) = "xyz";             // digit is assigned "015xyz9"
23     digit( 7, 0 ) = "***";             // digit is assigned "015xyz***9"
24     digit(-4, 3 ) = "$$$";             // digit is assigned "015xyz$$$9"

```

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

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

```

34     main() {
35         string x = "xxxxxxxxxxxxxxxx";
36         test( x, x(1,3), x(3,3), x(5,5), x(9,5), x(9,5) );
37     }
38
39     // x, a, b, c, & d are substring results passed by reference
40     // e is a substring result passed by value
41     void test(string &x, string &a, string &b, string &c, string &d, string e) {
42         // x a b c d e
43         a( 1, 2 ) = "aaa";               // aaaxxxxxxxxxx aaax axx xxxxx xxxxx xxxxx
44         b( 2, 12 ) = "bbb";              // aaabbbxxxxxxxxx aaab abbb bbxxx xxxxx xxxxx
45         c( 4, 5 ) = "ccc";               // aaabbbccccxxxxxx aaab abbb bbxccc ccxxx xxxxx
46         c = "yyy";                      // aaabyyyxxxxxxxx aaab aby yyy xxxxx xxxxx
47         d( 1, 3 ) = "ddd";              // aaabyyyxxxxdddx aaab aby yyy dddxx xxxxx
48         e( 1, 3 ) = "eee";              // aaabyyyxxxxdddx aaab aby yyy dddxx eeexx
49         x = e;                          // eeexx eeex exx x eeexx
50     }

```

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

```

1      string operator () (int start);
2  For example:
3      s = peter( 2 );           // s is assigned "ETER"
4      peter( 2 ) = "IPER";     // peter is assigned "PIPER"
5  It is also possible to substring using a string as the index for selecting the substring portion of the string.
6      string operator () (const string &index);
7  For example:
8      digit( "xyz$$$$" ) = "678"; // digit is assigned "0156789"
9      digit( "234" ) = "***";     // digit is assigned "0156789****"

```

## 10 13.7 Searching

11 The index operation

```

12      int index( const string &key, int start = 1, occurrence occ = first );
13  returns the position of the first or last occurrence of the key (depending on the occurrence indicator occ that is either
14  first or last) in the current string starting the search at position start. If the key does not appear in the current string, the
15  length of the current string plus one is returned. A negative starting position is a specification from the right end of the
16  string.
17      i = digit.index( "567" ); // i is assigned 3
18      i = digit.index( "567", 7 ); // i is assigned 11
19      i = digit.index( "567", -1, last ); // i is assigned 3
20      i = peter.index( "E", 5, last ); // i is assigned 4

```

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

```

25      strmask digitmask = digit;
26      strmask alphamask = string( "abcdefghijklmnopqrstuvwxyz" );

```

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

29 The include operation

```

30      int include( const strmask &, int = 1, occurrence occ = first );
31  returns the position of the first or last character (depending on the occurrence indicator, which is either first or last) in
32  the current string that does not appear in the mask starting the search at position start; hence it skips over characters
33  in the current string that are included (in) the mask. The characters in the current string do not have to be in the same
34  order as the mask. If all the characters in the current string appear in the mask, the length of the current string plus one
35  is returned, regardless of which occurrence is being searched for. A negative starting position is a specification from
36  the right end of the string.

```

```

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

```

39 The exclude operation

```

40      int exclude( string &mask, int start = 1, occurrence occ = first )
41  returns the position of the first or last character (depending on the occurrence indicator, which is either first or last) in
42  the current string that does appear in the mask string starting the search at position start; hence it skips over characters
43  in the current string that are excluded from (not in) in the mask string. The characters in the current string do not have
44  to be in the same order as the mask string. If all the characters in the current string do NOT appear in the mask string,
45  the length of the current string plus one is returned, regardless of which occurrence is being searched for. A negative
46  starting position is a specification from the right end of the string.

```

```

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

```

49 The includeStr operation:



```

1  string includeStr( strmask &mask, int start = 1, occurrence occ = first )
2  returns the longest substring of leading or trailing characters (depending on the occurrence indicator, which is either
3  first or last) of the current string that ARE included in the mask string starting the search at position start. A negative
4  starting position is a specification from the right end of the string.
5      s = peter.includeStr( alphamask );      // s is assigned "PETER"
6      s = ifstmt.includeStr( alphamask );      // s is assigned "IF"
7      s = peter.includeStr( digitmask );      // s is assigned ""
8  The excludeStr operation:
9      string excludeStr( strmask &mask, int start = 1, occurrence = first )
10 returns the longest substring of leading or trailing characters (depending on the occurrence indicator, which is either
11 first or last) of the current string that are excluded (NOT) in the mask string starting the search at position start. A
12 negative starting position is a specification from the right end of the string.
13      s = peter.excludeStr( digitmask );      // s is assigned "PETER"
14      s = ifstmt.excludeStr( strmask( punctuation ) ); // s is assigned "IF "
15      s = peter.excludeStr( alphamask );      // s is assigned ""

```

## 16 13.8 Miscellaneous

```

17 The trim operation
18      string trim( string &mask, occurrence occ = first )
19 returns a string in that is the longest substring of leading or trailing characters (depending on the occurrence indicator,
20 which is either first or last) which ARE included in the mask are removed.
21      // remove leading blanks
22      s = string( "   ABC" ).trim( " " );      // s is assigned "ABC",
23      // remove trailing blanks
24      s = string( "ABC   " ).trim( " ", last ); // s is assigned "ABC",
25 The translate operation
26      string translate( string &from, string &to )
27 returns a string that is the same length as the original string in which all occurrences of the characters that appear in the
28 from string have been translated into their corresponding character in the to string. Translation is done on a character
29 by character basis between the from and to strings; hence these two strings must be the same length. If a character in
30 the original string does not appear in the from string, then it simply appears as is in the resulting string.
31      // upper to lower case
32      peter = peter.translate( "ABCDEFGHIJKLMNOPQRSTUVWXYZ", "abcdefghijklmnopqrstuvwxyz" );
33      // peter is assigned "peter"
34      s = ifstmt.translate( "ABCDEFGHIJKLMNOPQRSTUVWXYZ", "abcdefghijklmnopqrstuvwxyz" );
35      // ifstmt is assigned "if (a > b) {"
36      // lower to upper case
37      peter = peter.translate( "abcdefghijklmnopqrstuvwxyz", "ABCDEFGHIJKLMNOPQRSTUVWXYZ" );
38      // peter is assigned "PETER"
39 The replace operation
40      string replace( string &from, string &to )
41 returns a string in which all occurrences of the from string in the current string have been replaced by the to string.
42      s = peter.replace( "E", "XX" );          // s is assigned "PXXTXXR"
43 The replacement is done left-to-right. When an instance of the from string is found and changed to the to string, it is
44 NOT examined again for further replacement.

```

## 45 13.9 Returning N+1 on Failure

```

46 Any of the string search routines can fail at some point during the search. When this happens it is necessary to return
47 indicating the failure. Many string types in other languages use some special value to indicate the failure. This value
48 is often 0 or -1 (PL/I returns 0). This section argues that a value of N+1, where N is the length of the base string in the

```

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

Table 1: Companion Routines for C $\forall$  string to C Strings

search, is a more useful value to return. The index-of function in APL returns N+1. These are the boundary situations and are often overlooked when designing a string type.

The situation that can be optimized by returning N+1 is when a search is performed to find the starting location for a substring operation. For example, in a program that is extracting words from a text file, it is necessary to scan from left to right over whitespace until the first alphabetic character is found.

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

If a text line contains all whitespaces, the exclude operation fails to find an alphabetic character. If exclude returns 0 or -1, the result of the substring operation is unclear. Most string types generate an error, or clip the starting value to 1, resulting in the entire whitespace string being selected. If exclude returns N+1, the starting position for the substring operation is beyond the end of the string leaving a null string.

The same situation occurs when scanning off a word.

```
start = line.include(alpha);
word = line(1, start - 1);
```

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

### 13.10 C Compatibility

To ease conversion from C to C $\forall$ , there are companion string routines for C strings. Table 1 shows the C routines on the left that also work with string and the rough equivalent string operation of the right. Hence, it is possible to directly convert a block of C string operations into @string@ just by changing the

For example, this block of C code can be converted to C $\forall$  by simply changing the type of variable s from char [] to string.

```
char s[32];
//string s;
strcpy( s, "abc" );      PRINT( %s, s );
strncpy( s, "abcdef", 3 ); PRINT( %s, s );
strcat( s, "xyz" );      PRINT( %s, s );
strncat( s, "vwxyz", 3 ); PRINT( %s, s );
PRINT( %zd, strlen( s ) );
PRINT( %c, s[3] );
PRINT( %s, strstr( s, "yzu" ) );
PRINT( %s, strstr( s, 'y' ) );
```

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

### 13.11 Input/Output Operators

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

## 14 Enumeration

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

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

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

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

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

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

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

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

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

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

### 14.1 Enum type

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

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

CV allows other integral types with associated values.

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

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

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

```
// non-integral numeric
enum( double ) Math { PI_2 = 1.570796, PI = 3.141597, E = 2.718282 }
// pointer
enum( char * ) Name { Fred = "Fred", Mary = "Mary", Jane = "Jane" };
int i, j, k;
enum( int * ) ptr { I = &i, J = &j, K = &k };
enum( int & ) ref { I = i, J = j, K = k };
// tuple
enum( [int, int] ) { T = [ 1, 2 ] };
// function
void f() {...} void g() {...}
enum( void (*)() ) funs { F = f, F = g };
```

```

1 // aggregate
2 struct S { int i, j; };
3 enum( S ) s { A = { 3, 4 }, B = { 7, 8 } };
4 // enumeration
5 enum( Letter ) Greek { Alph = A, Beta = B, /* more enums */ }; // alphabet intersection

```

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

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

```

10 Math m = PI; // allowed
11 double d = PI; // allowed, conversion to base type
12 m = E; // allowed
13 m = Alph; // disallowed
14 m = 3.141597; // disallowed
15 d = m; // allowed
16 d = Alph; // disallowed
17 Letter l = A; // allowed
18 Greek g = Alph; // allowed
19 l = Alph; // allowed, conversion to base type
20 g = A; // disallowed

```

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

```

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

```

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

## 14.2 Inheritance

Plan-9 inheritance may be used with enumerations.

```

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

```

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

Specifically, the inheritance relationship for Names is:

```

37 Name  $\subseteq$  Name2  $\subseteq$  Name3  $\subseteq$  const char * // enum type of Name

```

Hence, given

```

39 void f( Name );
40 void g( Name2 );
41 void h( Name3 );
42 void j( const char * );

```

the following calls are valid

```

44 f( Fred );
45 g( Fred ); g( Jill );
46 h( Fred ); h( Jill ); h( Sue );
47 j( Fred ); j( Jill ); j( Sue ); j( 'w' );

```

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

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

## 15 Routine Definition

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

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

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

In detail, the brackets, [], enclose the result type, where each return value is named and that name is a local variable of the particular return type.<sup>14</sup> The value of each local return variable is automatically returned at routine termination. Declaration qualifiers can only appear at the start of a routine definition, e.g.:

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

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

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

Routine *f* is called as follows:

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

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

CV style declarations cannot be used to declare parameters for K&R style routine definitions because of the following ambiguity:

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

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 string “*[ 5 ] int x*” declares a CV style parameter *x* of type array of 5 integers. Since the strings overlap starting with the open bracket, [, there is an ambiguous interpretation for the string. As well, CV-style declarations cannot be used to declare parameters for C-style routine-definitions because of the following ambiguity:

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

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 CV style unnamed parameter of type routine returning integer with unnamed parameter of type pointer to *foo*. The redefinition of a type name in a parameter list is the only context in C where the character *\** can appear to the left of a type name, and CV relies on all type qualifier characters appearing to the right of the type name. The inability to use CV declarations in these two contexts is probably a blessing because it precludes programmers from arbitrarily switching between declarations forms within a declaration contexts.

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

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

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

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

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

### 15.1 Named Return Values

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

<sup>14</sup>Michael Tiemann, with help from Doug Lea, provided named return values in g++, circa 1989.

```

1  int f() {
2      int x;
3      ... x = 0; ... x = y; ...
4      return x;
5  }

```

6 Because the value in the return variable is automatically returned when a CV routine terminates, the **return** statement  
7 *does not* contain an expression, as in:

```

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

```

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

```

11 [ int x, int y ] f() {
12     ...
13 }                      // implicitly return x, y

```

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

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

```

17 [ int x, int y ] f( int, x, int y ) {
18     ...
19 }                      // implicitly return x, y

```

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

```

21 [ int x, int y ] f( int, x, int y );      // prototype declaration
22 int a, b;
23 [a, b] = f( f( a, b ) );

```

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

## 29 15.2 Routine Prototype

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

```

32 [ int x ] f();           // returning int with no parameters
33 [ *int ] g( int y);      // returning pointer to int with int parameter
34 [ ] h( int, char );      // returning no result with int and char parameters
35 [ *int, int ] j( int );  // returning pointer to int and int, with int parameter

```

36 This syntax allows a prototype declaration to be created by cutting and pasting source text from the routine definition  
37 header (or vice versa). Like C, it is possible to declare multiple routine-prototypes in a single declaration, where the  
38 return type is distributed across *all* routine names in the declaration list (see Section 11, p. 18), *e.g.*:

```

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

```

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

42 Declaration qualifiers can only appear at the start of a CV routine declaration,<sup>7</sup> *e.g.*:

```

43 extern [ int ] f( int );
44 static [ int ] g( int );

```

## 45 16 Routine Pointers

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

```

1  * [ int x ] () fp;           // pointer to routine returning int with no parameters
2  * [ * int ] (int y) gp;      // pointer to routine returning pointer to int with int parameter
3  * [ ] (int, char) hp;        // pointer to routine returning no result with int and char parameters
4  * [ * int, int ] ( int ) jp; // pointer to routine returning pointer to int and int, with int parameter
5  While parameter names are optional, a routine name cannot be specified; for example, the following is incorrect:
6  * [ int x ] f () fp;         // routine name "f" is not allowed

```

## 17 Named and Default Arguments

Named and default arguments [20]<sup>15</sup> are two mechanisms to simplify routine call. Both mechanisms are discussed with respect to CV.

**Named (or Keyword) Arguments:** provide the ability to specify an argument to a routine call using the parameter name rather than the position of the parameter. For example, given the routine:

```
void p( int x, int y, int z ) {...}
```

a positional call is:

```
p( 4, 7, 3 );
```

whereas a named (keyword) call may be:

```
p( z : 3, x : 4, y : 7 );           // rewrite ⇒ p( 4, 7, 3 )
```

Here the order of the arguments is unimportant, and the names of the parameters are used to associate argument values with the corresponding parameters. The compiler rewrites a named call into a positional call. The advantages of named parameters are:

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

Unfortunately, named arguments do not work in C-style programming-languages because a routine prototype is not required to specify parameter names, nor do the names in the prototype have to match with the actual definition. For example, the following routine prototypes and definition are all valid.

```

27  void p( int, int, int );           // equivalent prototypes
28  void p( int x, int y, int z );
29  void p( int y, int x, int z );
30  void p( int z, int y, int x );
31  void p( int q, int r, int s ) {}    // match with this definition

```

Forcing matching parameter names in routine prototypes with corresponding routine definitions is possible, but goes against a strong tradition in C programming. Alternatively, prototype definitions can be eliminated by using a two-pass compilation, and implicitly creating header files for exports. The former is easy to do, while the latter is more complex.

Furthermore, named arguments do not work well in a CV-style programming-languages because they potentially introduces a new criteria for type matching. For example, it is technically possible to disambiguate between these two overloaded definitions of f based on named arguments at the call site:

```

39  int f( int i, int j );
40  int f( int x, double y );
41
42  f( j : 3, i : 4 );           // 1st f
43  f( x : 7, y : 8.1 );         // 2nd f
44  f( 4, 5 );                   // ambiguous call

```

However, named arguments compound routine resolution in conjunction with conversions:

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



```
1      f( i : 3, 5.7 );           // ambiguous call ?
```

2 Depending on the cost associated with named arguments, this call could be resolvable or ambiguous. Adding  
 3 named argument into the routine resolution algorithm does not seem worth the complexity. Therefore, C<sub>V</sub> does *not*  
 4 attempt to support named arguments.

5 **Default Arguments** provide the ability to associate a default value with a parameter so it can be optionally specified  
 6 in the argument list. For example, given the routine:

```
7      void p( int x = 1, int y = 2, int z = 3 ) { ... }
```

8 the allowable positional calls are:

```
9      p();                      // rewrite ⇒ p( 1, 2, 3 )
10     p( 4 );                   // rewrite ⇒ p( 4, 2, 3 )
11     p( 4, 4 );               // rewrite ⇒ p( 4, 4, 3 )
12     p( 4, 4, 4 );           // rewrite ⇒ p( 4, 4, 4 )
13     // empty arguments
14     p( , 4, 4 );             // rewrite ⇒ p( 1, 4, 4 )
15     p( 4, , 4 );            // rewrite ⇒ p( 4, 2, 4 )
16     p( 4, 4, );             // rewrite ⇒ p( 4, 4, 3 )
17     p( 4, , );              // rewrite ⇒ p( 4, 2, 3 )
18     p( , 4, );              // rewrite ⇒ p( 1, 4, 3 )
19     p( , , 4 );             // rewrite ⇒ p( 1, 2, 4 )
20     p( , , );               // rewrite ⇒ p( 1, 2, 3 )
```

21 Here the missing arguments are inserted from the default values in the parameter list. The compiler rewrites  
 22 missing default values into explicit positional arguments. The advantages of default values are:

- 23 • Routines with a large number of parameters are often very generalized, giving a programmer a number of  
 24 different options on how a computation is performed. For many of these kinds of routines, there are stan-  
 25 dard or default settings that work for the majority of computations. Without default values for parameters, a  
 26 programmer is forced to specify these common values all the time, resulting in long argument lists that are  
 27 error prone.
- 28 • When a routine's interface is augmented with new parameters, it extends the interface providing generaliz-  
 29 ability<sup>16</sup> (somewhat like the generalization provided by inheritance for classes). That is, all existing calls are  
 30 still valid, although the call must still be recompiled.

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

33 Default values may only appear in a prototype versus definition context:

```
34     void p( int x, int y = 2, int z = 3 ); // prototype: allowed
35     void p( int, int = 2, int = 3 );      // prototype: allowed
36     void p( int x, int y = 2, int z = 3 ) {} // definition: not allowed
```

37 The reason for this restriction is to allow separate compilation. Multiple prototypes with different default values is  
 38 an error.

39 Ellipse ("...") arguments present problems when used with default arguments. The conflict occurs because both  
 40 named and ellipse arguments must appear after positional arguments, giving two possibilities:

```
41     p( /* positional */, ... , /* named */ );
42     p( /* positional */, /* named */, ... );
```

43 While it is possible to implement both approaches, the first possibly is more complex than the second, *e.g.*:

```
44     p( int x, int y, int z, ... );
45     p( 1, 4, 5, 6, z : 3, y : 2 );           // assume p( /* positional */, ... , /* named */ );
46     p( 1, z : 3, y : 2, 4, 5, 6 );           // assume p( /* positional */, /* named */, ... );
```

<sup>16</sup>"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]

In the first call, it is necessary for the programmer to conceptually rewrite the call, changing named arguments into positional, before knowing where the ellipse arguments begin. Hence, this approach seems significantly more difficult, and hence, confusing and error prone. In the second call, the named arguments separate the positional and ellipse arguments, making it trivial to read the call.

The problem is exacerbated with default arguments, *e.g.*:

```
void p( int x, int y = 2, int z = 3... );
p( 1, 4, 5, 6, z : 3 );           // assume p( /* positional */, ... , /* named */ );
p( 1, z : 3, 4, 5, 6 );           // assume p( /* positional */, /* named */, ... );
```

The first call is an error because arguments 4 and 5 are actually positional not ellipse arguments; therefore, argument 5 subsequently conflicts with the named argument *z* : 3. In the second call, the default value for *y* is implicitly inserted after argument 1 and the named arguments separate the positional and ellipse arguments, making it trivial to read the call. For these reasons, C<sub>V</sub> requires named arguments before ellipse arguments. Finally, while ellipse arguments are needed for a small set of existing C routines, like `printf`, the extended C<sub>V</sub> type system largely eliminates the need for ellipse arguments (see Section 26, p. 62), making much of this discussion moot.

Default arguments and overloading (see Section 26, p. 62) are complementary. While in theory default arguments can be simulated with overloading, as in:

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

the number of required overloaded routines is linear in the number of default values, which is unacceptable growth. In general, overloading should only be used over default arguments if the body of the routine is significantly different. Furthermore, overloading cannot handle accessing default arguments in the middle of a positional list, via a missing argument, such as:

```
p( 1, /* default */, 5 );           // rewrite ⇒ p( 1, 2, 5 )
```

Given the C<sub>V</sub> restrictions above, both named and default arguments are backwards compatible. C++ only supports default arguments; Ada supports both named and default arguments.

## 18 Unnamed Structure Fields

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

```
struct {
    int f1;           // named field
    int f2 : 4;       // named field with bit field size
    int : 3;          // unnamed field for basic type with bit field size
    int ;             // disallowed, unnamed field
    int *;            // disallowed, unnamed field
    int (*)( int );   // disallowed, unnamed field
};
```

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.*:

```
struct {
    int , , ;         // 3 unnamed fields
}
```

## 19 Nesting

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

### 19.1 Type Nesting

C<sub>V</sub> allows type nesting, and type qualification of the nested types (see Figure 5), where as C hoists (refactors) nested types into the enclosing scope and has no type qualification. In the left example in C, types C, U and T are implicitly

C Type Nesting	C Implicit Hoisting	C $\forall$
<pre> <b>struct</b> S {   <b>enum</b> C { R, G, B };   <b>struct</b> T {     <b>union</b> U { <b>int</b> i, j };     <b>enum</b> C c;     <b>short int</b> i, j;   };   <b>struct</b> T t; } s;  <b>int</b> fred() {   s.t.c = R;   <b>struct</b> T t = { R, 1, 2 };   <b>enum</b> C c;   <b>union</b> U u; } </pre>	<pre> <b>enum</b> C { R, G, B }; <b>union</b> U { <b>int</b> i, j }; <b>struct</b> T {   <b>enum</b> C c;   <b>short int</b> i, j; }; <b>struct</b> S {   <b>struct</b> T t; } s; </pre>	<pre> <b>struct</b> S {   <b>enum</b> C { R, G, B };   <b>struct</b> T {     <b>union</b> U { <b>int</b> i, j };     <b>enum</b> C c;     <b>short int</b> i, j;   };   <b>struct</b> T t; } s;  <b>int</b> fred() {   s.t.c = <b>S</b>.R; // type qualification   <b>struct</b> <b>S</b>.T t = { <b>S</b>.R, 1, 2 };   <b>enum</b> <b>S</b>.C c;   <b>union</b> <b>S</b>.T.U u; } </pre>

Figure 5: Type Nesting / Qualification

1 hoisted outside of type S into the containing block scope. In the right example in C $\forall$ , 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 C $\forall$  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 C $\forall$  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 C $\forall$  (and Indexcgc) because of local dependence

```

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  }

```

because

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

## 20 Tuple

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

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

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

### 20.1 Multiple-Return-Value Functions

In C and most programming languages, functions return at most one value; however, many operations have multiple outcomes, some exceptional (see Section 10, p. 16). To emulate functions with multiple return values, *aggregation* and/or *aliasing* is used.

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

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

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

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

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

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

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

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

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

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

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

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

```

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

In 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



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

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

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

## 20.6 Assignment

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

```
int x;
double y;
[int, double] z;
[y, x] = 3.14;           // mass assignment
[x, y] = z;               // multiple assignment
z = 10;                   // mass assignment
z = [x, y];               // multiple assignment
```

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

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

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

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

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

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

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

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

In CV, tuple assignment is an expression where the result type is the type of the left side of the assignment, as in normal assignment. That is, a tuple assignment produces the value of the left-hand side after assignment. These semantics allow cascading tuple assignment to work out naturally in any context where a tuple is permitted. These semantics are a change from the original tuple design in K-W C [31], wherein tuple assignment was a statement that allows cascading assignments as a special case. Restricting tuple assignment to statements was an attempt to fix



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

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

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

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

## 20.7 Construction

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

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

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

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

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

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

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

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

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

## 20.8 Member-Access Expression

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

```

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

```

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

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

```

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

```

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

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

```

[ int, int, long, double ] x;
void f( double, long );

f( x.[ 0, 3 ] );           // f( x.0, x.3 )
x.[ 0, 1 ] = x.[ 1, 0 ];   // [ x.0, x.1 ] = [ x.1, x.0 ]
[ long, int, long ] y = x.[ 2, 0, 2 ];

```

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

```

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

```

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

```

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

```

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

## 20.9 Casting

In C, the cast operator is used to explicitly convert between types. 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.

```

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

f();               // ambiguous - (1),(2) both equally viable
(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.

```

1 int f();
2 void g();
3
4 (void)f();           // valid, ignore results

```

```

1  5  (int)g();                      // invalid, void cannot be converted to int
2  6
3  7  struct A { int x; };
4  8  (struct A)f();                 // invalid, int cannot be converted to A

```

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

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

For example,

```

15  [int, int, int] f();
16  [int, [int, int], int] g();
17
18  ([int, double])f();              // (1) valid
19  ([int, int, int])g();            // (2) valid
20  ([void, [int, int]])g();         // (3) valid
21  ([int, int, int, int])g();       // (4) invalid
22  ([int, [int, int, int]])g();     // (5) invalid

```

(1) discards the last element of the return value and converts the second element to type `double`. Since `int` is effectively a 1-element tuple, (2) discards the second component of the second element of the return value of `g`. If `g` is free of side effects, this is equivalent to `[(int)(g().0), (int)(g().1.0), (int)(g().2)]`. Since `void` is effectively a 0-element tuple, (3) discards the first and third return values, which is effectively equivalent to `[(int)(g().1.0), (int)(g().1.1)]`. Note that a cast is not a function call in C, so flattening and structuring conversions do not occur for cast expressions. As such, (4) is invalid because the cast target type contains 4 components, while the source type contains only 3. Similarly, (5) is invalid because the cast `([int, int, int])(g().1)` is invalid. That is, it is invalid to cast `[int, int]` to `[int, int, int]`.

## 30 20.10 Polymorphism

Due to the implicit flattening and structuring conversions involved in argument passing, `otype` and `dtype` parameters are restricted to matching only with non-tuple types. The integration of polymorphism, type assertions, and monomorphic specialization of tuple-assertions are a primary contribution of this thesis to the design of tuples.

```

34  forall(T, dtype U)
35  void f(T x, U * y);
36
37  f([5, "hello"]);

```

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

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

```

42  forall(T | { T ?+?(T, T); })
43  [T, T, T] ?+?([T, T, T] x, [T, T, T] y) {
44      return [x.0+y.0, x.1+y.1, x.2+y.2];
45  }
46  [int, int, int] x;
47  int i1, i2, i3;
48  [i1, i2, i3] = x + ([10, 20, 30]);

```

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

```

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

```

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

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

```
4    forall(T | { T ?+?(T, T); })
5    [T, T, T] ?+?([T, T] x, [T, T, T, T]);
6    forall(T | { T ?+?(T, T); })
7    [T, T, T] ?+?(T x, [T, T, T, T, T]);
```

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

```
14    void f([int, int], int, int);
15
16    f([0, 0], 0, 0);                // no cost
17    f(0, 0, 0, 0);                 // cost for structuring
18    f([0, 0], [0, 0]);             // cost for flattening
19    f([0, 0, 0], 0);               // cost for flattening and structuring
```

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

### 23 20.10.1 Assertion Inference

```
24    int f([int, double], double);
25    forall(T, U | { T f(T, U, U); })
26    void g(T, U);
27    g(5, 10.21);
```

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

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

```
35    int _thunk(int _p0, double _p1, double _p2) {
36        return f([_p0, _p1], _p2);
37    }
```

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

## 40 21 Tuples

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

```
43    [ exprlist ]
```

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

```
46    [ x, y, z ]
47    [ 2 ]
48    [ v + w, x * y, 3.14159, f() ]
```

49 Tuples are permitted to contain sub-tuples (*i.e.*, nesting), such as `[ [ 14, 21 ], 9 ]`, which is a 2-element tuple whose first  
50 element is itself a tuple. Note, a tuple is not a record (structure); a record denotes a single value with substructure,

whereas a tuple is multiple values with no substructure (see flattening coercion in Section 20.5, p. 40). In essence, tuples are largely a compile time phenomenon, having little or no runtime presence.

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

```
[ typelist ]
```

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

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

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

Examples of declarations using tuple types are:

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

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

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

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

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

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

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

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

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

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

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

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

1     **const volatile** [ **int**, **float**, **const int** ] x;

2 is equivalent to:

3     [ **const volatile int**, **const volatile float**, **const volatile int** ] x;

4 Declaration qualifiers can only appear at the start of a CV tuple declaration<sup>4</sup>, *e.g.*:

5     **extern** [ **int**, **int** ] w1;

6     **static** [ **int**, **int**, **int** ] w2;

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

## 12 21.1 Tuple Coercions

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

17     [ **int**, **int**, **int**, **int** ] w;

18     w = [ 1, 2, 3, 4 ];

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

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

21     [ a, b, c, d ] = w

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

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

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

26 First the right-hand tuple is flattened and then the values are assigned individually. Flattening is also performed on  
27 tuple types. For example, the type [ **int**, [ **int**, **int** ], **int** ] can be coerced, using flattening, into the type [ **int**, **int**, **int**, **int** ].

28 A *structuring coercion* is the opposite of flattening; a tuple is structured into a more complex nested tuple. For  
29 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  
30 type [ **int**, [ **int**, **int** ], **int** ]. In the following example, the last assignment illustrates all the tuple coercions:

31     [ **int**, **int**, **int**, **int** ] w = [ 1, 2, 3, 4 ];

32     **int** x = 5;

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

34 Starting on the right-hand tuple in the last assignment statement, w is opened, producing a tuple of four values;  
35 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 ],  
36 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  
37 closed to create a tuple value. Finally, x is assigned 1 and w is assigned the tuple value using multiple assignment (see  
38 Section 20.6, p. 41).

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

## 41 21.2 Mass Assignment

42 CV permits assignment to several variables at once using mass assignment [25]. Mass assignment has the following  
43 form:

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

45 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  
46 can appear on the left-hand side of a conventional assignment statement. *expr* is any standard arithmetic expression.

47 Clearly, the types of the entities being assigned must be type compatible with the value of the expression.

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





## 22 Stream I/O Library

The goal of C<sub>V</sub> 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 C<sub>V</sub> 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.

C<sub>V</sub> formatting incorporates ideas from C printf, C++ stream manipulators, and Python implicit spacing and newline. Specifically:

- printf/Python format codes are dense, making them difficult to read and remember. C<sub>V</sub>/C++ format manipulators are named, making them easier to read and remember.
- printf/Python separate format codes from associated variables, making it difficult to match codes with variables. C<sub>V</sub>/C++ co-locate codes with associated variables, where C<sub>V</sub> has the tighter binding.
- Format manipulators in printf/Python/C<sub>V</sub> 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.
- C<sub>V</sub> has more sophisticated implicit value spacing than Python, plus implicit newline at the end of a print.

### 22.1 Basic I/O

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

#### 22.1.1 Stream Output

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

	C <sub>V</sub>	C++	Python
24	<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 &lt;&lt; x &lt;&lt; " " &lt;&lt; y &lt;&lt; " " &lt;&lt; z &lt;&lt; endl;</code>	<code>print( x, y, z )</code>
	<code>1_2_3</code>	<code>1_2_3</code>	<code>1_2_3</code>

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

28	<code>[int, [ int, int ] ] t1 = [ 1, [ 2, 3 ] ], t2 = [ 4, [ 5, 6 ] ];</code>	
29	<code>sout   t1   t2;</code>	<code>// print tuples</code>
30	<code>1_2_3_4_5_6</code>	

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

34	<b>C<sub>V</sub>:</b> <code>sout   x * 3   y + 1   z &lt;&lt; 2   x == y   (x   y)   (x    y)   (x &gt; z ? 1 : 2);</code>
	<b>C++:</b> <code>cout &lt;&lt; x * 3 &lt;&lt; y + 1 &lt;&lt; (z &lt;&lt; 2) &lt;&lt; (x == y) &lt;&lt; (x   y) &lt;&lt; (x    y) &lt;&lt; (x &gt; z ? 1 : 2) &lt;&lt; endl;</code>
	<code>3_3_12_0_3_1_2</code>

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

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

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

41	<code>exit   "x ("   x   ")" negative value."; // print, terminate, and return EXIT_FAILURE to shell</code>
42	<code>abort   "x ("   x   ")" negative value."; // print, terminate, and generate stack trace and core file</code>

1 Note, C stream variables `stdin`, `stdout`, `stderr`, `exit`, and `abort` overload C variables `stdin`, `stdout`, `stderr`, and functions `exit`  
 2 and `abort`, respectively.

### 3 22.1.2 Stream Input

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

	<b>C</b>	<b>C++</b>	<b>Python</b>
6	<code>char c; int i; double d</code> <code>sin   c   i   d;</code>	<code>cin &gt;&gt; c &gt;&gt; i &gt;&gt; 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>

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

- 10 • 1-9 prefix introduces a decimal value (0-9),
- 11 • 0 prefix introduces an octal value (0-7), and
- 12 • 0x or 0X prefix introduces a hexadecimal value (0-f) with lower or upper case letters.

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

```
17 char ch; int i; float f; double d; _Complex double cxd;
18 sin | ch | i | f | d | cxd;
19 X 42 1234.5 0xffffp-2 3.5+7.1i
```

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

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

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

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

```
29 sout | "\n\t\tf\0234\x23";
```

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

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

33 The input string is always null terminated '\0' in the input variable. Because of potential buffer overrun when reading  
 34 C strings, strings are restricted to work with input manipulators (see Section 22.6, p. 56). As well, there are multiple  
 35 input-manipulators for scanning complex input string formats, e.g., a quoted character or string.

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

### 40 22.1.3 Stream Files

41 Figure 6 shows the I/O stream operations for interacting with files other than `sin`, `sout`, and `cerr`.

- 42 • `fail` tests the stream error-indicator, returning nonzero if it is set.
- 43 • `clear` resets the stream error-indicator.

```
// ***** ofstream *****
bool fail( ofstream & );
void clear( 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 clear( 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 6: I/O Stream Functions

- flush (ofstream only) causes any unwritten data for a stream to be written to the file.
  - eof (ifstream only) tests the end-of-file indicator for the stream pointed to by stream. Returns true if the end-of-file indicator is set, otherwise false.
  - open binds the file with name to a stream accessed with mode (see fopen).
  - close flushes the stream and closes the file.
  - write (ofstream only) writes size bytes to the stream. The bytes are written lazily when an internal buffer fills. Eager buffer writes are done with flush
  - read (ifstream only) reads size bytes from the stream.
  - ungetc (ifstream only) pushes the character back to the input stream. Pushed-back characters returned by subsequent reads in the reverse order of pushing.
- The constructor functions:
- create an unbound stream, which is subsequently bound to a file with open.
  - create a bound stream to the associated file with given mode.

The destructor closes the stream.

Figure 7 demonstrates the file operations by showing the idiomatic C++ command-line processing and copying an input file to an output file. Note, a stream variable may be copied because it is a reference to an underlying stream data-structures. **All I/O errors are handled as exceptions**, but end-of-file is not an exception as C programmers are use to explicitly checking for it.

## 22.2 Implicit Separator

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

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

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

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

```
sout | '1' | '2' | '3';
```

```

#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;
    for () {                                  // read/write characters
        in | ch;
        if ( eof( in ) ) break;               // eof ?
        out | ch;
    } // for
} // main

```

Figure 7: C\&amp; Command-Line Processing

```

1      123
2      3. A separator does not appear before or after a null (empty) C string, which is a local mechanism to disable
3      insertion of the separator character.
4      sout | 1 | " " | 2 | " " | 3;
5      123
6      4. A separator does not appear before a C string starting with the (extended) ASCII characters: , . ; ! ? ) ] } % < »,
7      where » is a closing citation mark.
8      sout | 1 | " , x " | 2 | " . x " | 3 | " ; x " | 4 | " ! x " | 5 | " ? x " | 6 | " % x "
9      | 7 | " < x " | 8 | " » x " | 9 | " ) x " | 10 | " ] x " | 11 | " } x ";
10     Input1,2,3,4,5,6%,7<,8»,9),10],11}x
11     5. A separator does not appear after a C string ending with the (extended) ASCII characters: ( [ { = $ £ ¥ ¤ ¨, where
12     ¤ ¨ are inverted opening exclamation and question marks, and « is an opening citation mark.
13     sout | "x ( " | 1 | "x [ " | 2 | "x { " | 3 | "x = " | 4 | "x $ " | 5 | "x £ " | 6 | "x ¥ "
14     | 7 | "x ¤ " | 8 | "x ¨ " | 9 | "x « " | 10;
15     x(1,x[2,x{3,x=4,x$5,x£6,x¥7,x¤8,x¨9,x«10
16     6. A separator does not appear before/after a C string starting/ending with the ASCII quote or whitespace charac-
17     ters: ` ' " : _ \ t \ v \ f \ r \ n
18     sout | "x ` " | 1 | "x ' " | 2 | "x \" " | 3 | "x : " | 4 | "x _ " | 5 | "x \t " | 6 | "x \t x ";

```

```
x`1`x' 2' x"3"x:4:x_5_x_6_x
```

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

```
sout | "x ( _ | 1 | " _ ) x" | 2 | " _ , x" | 3 | " _ : x : _ " | 4;
x_( _1 _ )_x_2_ _x_3_:_x:_4
```

## 22.3 Separation Manipulators

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

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

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

sepSet( sout, " " ); // reset separator to " "
sout | 1 | 2 | 3 | " \ " " | sepGet( sout ) | " \ " ";
1_2_3_ " "
```

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

```
char store[sepSize]; // sepSize is the maximum separator size
strcpy( store, sepGet( sout ) ); // copy current separator
sepSet( sout, " _ " ); // change separator to underscore
sout | 1 | 2 | 3;
1_2_3

sepSet( sout, store ); // change separator back to original
sout | 1 | 2 | 3;
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).

```
sepSetTuple( sout, " , " ); // set tuple separator from " , " to " "
sout | t1 | t2 | " \ " " | sepTupleVal | " \ " ";
1_2_3_4_5_6_ " "

sepSetTuple( sout, " , " ); // reset tuple separator to " , "
sout | t1 | t2 | " \ " " | sepGetTuple( sout ) | " \ " ";
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.

```
sout | sepOff | 1 | 2 | 3; // turn off implicit separator
123

sout | sepOn | 1 | 2 | 3; // turn on implicit separator
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.

```
sout | 1 | nosep | 2 | 3; // turn off implicit separator for the next item
12_3

sout | sepOff | 1 | sep | 2 | 3; // turn on implicit separator for the next item
1_23
```

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

```
sout | t1 | nosep | t2; // turn off implicit separator for the next item
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.

```

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

```

## 22.4 Newline Manipulators

The following manipulators control newline separation for input and output.

For input:

1. `nlOn` reads the newline character, when reading single characters.
2. `nlOff` does *not* read the newline character, when reading single characters.
3. `nl` scans characters until the next newline character, *i.e.*, ignore the remaining characters in the line. If `nlOn` is enabled, the `nl` is also consumed.

For example, in:

```

13      int i, j;
14      sin | i | nl | j;
15      1 2
16      3

```

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

```

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

```

the left example skips the newline and reads 'X' into `ch`, while the right example reads the newline into `ch`.

For output:

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

```

24      sout | nl;      // only print newline
25      sout | 2;      // implicit newline
26      sout | 3 | nl | 4 | nl;      // terminating nl merged with implicit newline
27      sout | 5 | nl | nl;      // again terminating nl merged with implicit newline
28      sout | 6;      // implicit newline
29
30      2
31      3
32      4
33      5
34
35      6

```

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

## 22.5 Output Manipulators

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

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

```

41      sout | bin( 0 ) | bin( 27HH ) | bin( 27H ) | bin( 27 ) | bin( 27L );
42      0b0 0b11011 0b11011 0b11011 0b11011
43      sout | bin( -27HH ) | bin( -27H ) | bin( -27 ) | bin( -27L );
44      0b11100101 0b11111111111100101 0b1111111111111111111111111111111100101 0b(58 1s)100101

```

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

```
sout | oct( 0 ) | oct( 27HH ) | oct( 27H ) | oct( 27 ) | oct( 27L );
0 033 033 033 033
sout | oct( -27HH ) | oct( -27H ) | oct( -27 ) | oct( -27L );
0345 0177745 03777777745 017777777777777777745
```

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

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

```
sout | hex( 0 ) | hex( 27HH ) | hex( 27H ) | hex( 27 ) | hex( 27L );
0x0 0x1b 0x1b 0x1b 0x1b
sout | hex( -27HH ) | hex( -27H ) | hex( -27 ) | hex( -27L );
0xe5 0xffe5 0xffffffffe5 0xffffffffffffe5
```

```
sout | hex( 0.0 ) | hex( 27.5F ) | hex( 27.5 ) | hex( 27.5L );
0x0p+0 0x1.b8p+4 0x1.b8p+4 0xd.cp+1
sout | hex( -27.5F ) | hex( -27.5 ) | hex( -27.5L );
-0x1.b8p+4 -0x1.b8p+4 -0xd.cp+1
```

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

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

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

```
sout | eng( 0.0 ) | eng( 27000.5 ) | eng( -27.5e7 );
0e0 27.0005e3 -275e6
```

6. unit( engineering-notation ) print engineering exponent as a letter between the range  $10^{-24}$  and  $10^{24}$ :

$y \Rightarrow 10^{-24}$ ,  $z \Rightarrow 10^{-21}$ ,  $a \Rightarrow 10^{-18}$ ,  $f \Rightarrow 10^{-15}$ ,  $p \Rightarrow 10^{-12}$ ,  $n \Rightarrow 10^{-9}$ ,  $u \Rightarrow 10^{-6}$ ,  $m \Rightarrow 10^{-3}$ ,  $K \Rightarrow 10^3$ ,  
 $M \Rightarrow 10^6$ ,  $G \Rightarrow 10^9$ ,  $T \Rightarrow 10^{12}$ ,  $P \Rightarrow 10^{15}$ ,  $E \Rightarrow 10^{18}$ ,  $Z \Rightarrow 10^{21}$ ,  $Y \Rightarrow 10^{24}$ .

For exponent  $10^0$ , no decimal point or letter is printed.

```
sout | unit(eng( 0.0 )) | unit(eng( 27000.5 )) | unit(eng( -27.5e7 ));
0 27.0005K -275M
```

7. upcase( bin / hex / floating-point ) print letters in a value in upper case. Lower case is the default.

```
sout | upcase( bin( 27 ) ) | upcase( hex( 27 ) ) | upcase( 27.5e-10 ) | upcase( hex( 27.5 ) );
0B11011 0X1B 2.75E-09 0X1.B8P+4
```

8. nobase( integer ) do not precede bin, oct, hex with 0b/0B, 0, or 0x/0X. Printing the base is the default.

```
sout | nobase( bin( 27 ) ) | nobase( oct( 27 ) ) | nobase( hex( 27 ) );
11011 33 1b
```

9. nodp( floating-point ) do not print a decimal point if there are no fractional digits. Printing a decimal point is the default, if there are no fractional digits.

```
sout | 0. | nodp( 0. ) | 27.0 | nodp( 27.0 ) | nodp( 27.5 );
0.0 0 27.0 27.5
```

10. sign( integer / floating-point ) prefix with plus or minus sign (+ or -). Only printing the minus sign is the default.

```
sout | sign( 27 ) | sign( -27 ) | sign( 27. ) | sign( -27. ) | sign( 27.5 ) | sign( -27.5 );
+27 -27 +27.0 -27.0 +27.5 -27.5
```

11. wd( minimum, value ), wd( minimum, precision, value ) For all types, minimum is the number of printed characters. If the value is shorter than the minimum, it is padded on the right with spaces.

```
sout | wd( 4, 34 ) | wd( 3, 34 ) | wd( 2, 34 );
34 34 34
sout | wd( 10, 4. ) | wd( 9, 4. ) | wd( 8, 4. );
4.000000 4.000000 4.000000
sout | wd( 4, "ab" ) | wd( 3, "ab" ) | wd( 2, "ab" );
ab ab ab
```

If the value is larger, it is printed without truncation, ignoring the minimum.

```
sout | wd( 4, 34567 ) | wd( 3, 34567 ) | wd( 2, 34567 );
```



```

1      sout | wd( 4, 3456. ) | wd( 3, 3456. ) | wd( 2, 3456. );
2      sout | wd( 4, "abcde" ) | wd( 3, "abcde" ) | wd( 2, "abcde" );
3      34567_34567_34567
4      3456_3456_3456.
5      abcde_abcde_abcde

```

For integer types, precision is the minimum number of printed digits. If the value is shorter, it is padded on the left with leading zeros.

```

8      sout | wd( 4,3, 34 ) | wd( 8,4, 34 ) | wd( 10,10, 34 );
9      _034_000000034_0000000034

```

If the value is larger, it is printed without truncation, ignoring the precision.

```

11     sout | wd( 4,1, 3456 ) | wd( 8,2, 3456 ) | wd( 10,3, 3456 );
12     3456_3456_3456_3456

```

If precision is 0, nothing is printed for zero. If precision is greater than the minimum, it becomes the minimum.

```

14     sout | wd( 4,0, 0 ) | wd( 3,10, 34 );
15     _0000000034

```

For floating-point types, precision is the minimum number of digits after the decimal point.

```

17     sout | wd( 6,3, 27.5 ) | wd( 8,1, 27.5 ) | wd( 8,0, 27.5 ) | wd( 3,8, 27.5 );
18     27.500_27.5_28.27.50000000

```

For the C-string type, precision is the maximum number of printed characters, so the string is truncated if it exceeds the maximum.

```

21     sout | wd( 6,8, "abcd" ) | wd( 6,8, "abcdefghijk" ) | wd( 6,3, "abcd" ) | wd( 10, " " ) | 'X';
22     _abcd_abcdefgh_abc_X

```

Note, printing the null string with minimum width L pads with L spaces.

12. `ws( minimum, significant, floating-point )` For floating-point types, minimum is the same as for manipulator `wd`, but significant is the maximum number of significant digits to be printed for both the integer and fractions (versus only the fraction for `wd`). If a value's significant digits is greater than significant, the last significant digit is rounded up.

```

28     sout | ws(6,6, 234.567) | ws(6,5, 234.567) | ws(6,4, 234.567) | ws(6,3, 234.567);
29     234.567_234.57_234.6_235

```

If a value's magnitude is greater than significant, the value is printed in scientific notation with the specified number of significant digits.

```

32     sout | ws(6,6, 234567.) | ws(6,5, 234567.) | ws(6,4, 234567.) | ws(6,3, 234567.);
33     234567._2.3457e+05_2.346e+05_2.35e+05

```

If significant is greater than minimum, it defines the number of printed characters.

```

35     sout | ws(3,6, 234567.) | ws(4,6, 234567.) | ws(5,6, 234567.) | ws(6,6, 234567.);
36     234567._234567._234567._234567.

```

13. `left( field-width )` left justify within the given field.

```

38     sout | left(wd(4, 27)) | left(wd(10, 27.)) | left(wd(10, 27.5)) | left(wd(4,3, 27)) | left(wd(10,3, 27.5));
39     27_27.000000_27.500000_027_27.500

```

14. `pad0( field-width )` left pad with zeroes (0).

```

41     sout | pad0( wd( 4, 27 ) ) | pad0( wd( 4,3, 27 ) ) | pad0( wd( 8,3, 27.5 ) );
42     0027 027 0027.500

```

## 22.6 Input Manipulators

A string variable *must* be large enough to contain the input sequence. To force programmers to consider buffer overruns for C-string input, C-strings may only be read with a width field, which should specify a size less than or equal to the C-string size, e.g.:

```

47     char line[64];
48     sin | wdi( sizeof(line), line );           // must specify string size

```

Certain input manipulators support a *scanset*, which is a simple regular expression, where the matching set contains any Latin-1 character (8-bits) or character ranges using minus. For example, the scanset `"a-zA-Z -/?$"` matches

any number of characters between 'a' and 'z', between 'A' and 'Z', between space and '/', and characters '?' and (Latin-1) '§'. The following string is matched by this scanset:

```
!&%%$ abAA () ZZZ ??§ xx§§
```

To match a minus, make it the first character in the set, e.g., "-0-9". Other complex forms of regular-expression matching are unsupported.

The following manipulators control scanning of input values (reading) and only affect the format of the argument.

1. `skip( scanset )`, `skip( N )` consumes either the *scanset* or the next *N* characters, including newlines. If the match succeeds, the input characters are ignored, and input continues with the next character. If the match fails, the input characters are left unread.

```
char scanset[] = "abc";
sin | "abc_" | skip( scanset ) | skip( 5 ); // match and skip input sequence
abc_abc_abc_xxx
```

Again, the blank in the format string "abc\_" matches any number of whitespace characters.

2. `wdi( maximum, T & v )` For all types except **char** \*, whitespace is skipped and the longest sequence of non-whitespace characters matching an appropriate typed (*T*) value is read, converted into its corresponding internal form, and written into the *T* variable. *maximum* is the maximum number of characters read for the current value rather than the longest sequence.

```
char ch; char ca[3]; int i; double d;
sin | wdi( sizeof(ch), ch ) | wdi( sizeof(ca), ca[0] ) | wdi( 3, i ) | wdi( 8, d ); // c == 'a', ca == "bcd", i == 123, d == 345.6
abcd1233.456E+2
```

Here, `ca[0]` is type **char**, so the width reads 3 characters **without** a null terminator. If an input value is not found for a variable, the exception `missing_data` is raised, and the input variable is unchanged.

Note, input `wdi` cannot be overloaded with output `wd` because both have the same parameters but return different types. Currently, C++ cannot distinguish between these two manipulators in the middle of an `sout/sin` expression based on return type.

3. `wdi( maximum size, char s[] )` For type **char** \*, whitespace is 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.

```
char cs[10];
sin | wdi( sizeof(cs), cs );
012345678
```

Nine non-whitespace character are read and the null character is added to make ten.

4. `wdi( maximum size, maximum read, char s[] )` This manipulator is the same as the previous one, except *maximum read* is the maximum number of characters read for the current value rather than the longest sequence, where *maximum read* ≤ *maximum size*.

```
char cs[10];
sin | wdi( sizeof(cs), 9, cs );
0123456789
```

The exception `cstring_length` is not raised, because the read stops reading after nine characters.

5. `getline( wdi manipulator, const char delimiter = '\n' )` consumes the scanset "[^D]D", where *D* is the delimiter character, which reads all characters from the current input position to the delimiter character into the string (null terminated), and consumes and ignores the delimiter. If the delimiter character is omitted, it defaults to '\n' (newline).

```
char cs[10];
sin | getline( wdi( sizeof(cs), cs ) );
sin | getline( wdi( sizeof(cs), cs ), 'X' ); // X is the line delimiter
abc_??_#@%
abc_??_#@%X_w
```

The same value is read for both input strings.

6. `quoted( 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 whites-

pace, consumes and ignores the left delimiter, reads a single character into `ch`, and consumes and ignores the right delimiter (3 characters). If the delimit character is omitted, it defaults to `'\''` (single quote).

```
char ch;
sin | quoted( ch ); sin | quoted( ch, ' "' ); sin | quoted( ch, ' [ ', ' ] ' );
    _ _ _ 'a' _ _ _ "a"[a]
```

7. `quoted( wdi manipulator, const char Ldelimiter = '\'', const char Rdelimiter = '\0' )` consumes the scanset `"L[^R]R"`, where `L` is the left delimiter character and `R` is the right delimiter character, which skips whitespace, consumes and ignores the left delimiter, reads characters until the right-delimiter into the string variable (null terminated), and consumes and ignores the right delimiter. If the delimit character is omitted, it defaults to `'\''` (single quote).

```
char cs[10];
sin | quoted( wdi( sizeof(cs), cs ) ); // " is the start/end delimiter
sin | quoted( wdi( sizeof(cs), cs ), '\ ' ); // ' is the start/end delimiter
sin | quoted( wdi( sizeof(cs), cs ), ' [ ', ' ] ' ); // [ is the start and ] is the end delimiter
    _ _ _ "abc" _ _ _ 'abc'[abc]
```

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

```
char cs[10];
sin | incl( "abc", cs );
    bcxyz
```

9. `excl( scanset, wdi manipulator )` consumes the *not* scanset, which reads all the scanned characters into the string variable (null terminated).

```
char cs[10];
sin | excl( "abc", cs );
    xyzbca
```

10. `ignore( T & v or const char cs[] or string manipulator )` consumes the appropriate characters for the type and ignores them, so the input variable is unchanged.

```
double d;
char cs[10];
sin | ignore( d ); // d is unchanged
sin | ignore( cs ); // cs is unchanged, no wdi required
sin | ignore( quoted( wdi( sizeof(cs), cs ) ) ); // cs is unchanged
    _ _ -75.35e-4 _ _ 25 _ _ "abc"
```

## 34 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:

```
thread1 : sout | "abc " | "def ";
thread2 : sout | "uvw " | "xyz ";
```

possible outputs are:

abc def	abc uvw xyz	uvw abc xyz def	abuvwcdexf	uvw abc def
uvw xyz	def	yz		xyz

Concurrent operations can even corrupt the internal state of the stream resulting in failure. As a result, some form of mutual exclusion is required for concurrent stream access.

A coarse-grained solution is to perform all stream operations via a single thread or within a monitor providing the necessary mutual exclusion for the stream. A fine-grained solution is to have a lock for each stream, which is acquired and released around stream operations by each thread. C++ provides a fine-grained solution where a recursive lock is acquired and released indirectly via a manipulator `acquire` or instantiating an RAII type specific for the kind of stream: `osacquire` for output streams and `isacquire` for input streams.

The common usage is the short form of the mutex statement to lock a stream during a single cascaded I/O expression, *e.g.*:

```

1  thread1 : mutex( sout ) sout | "abc " | "def ";
2  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.

```

5  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.*:

```

9  mutex( sout ) {
10     sout | 1;
11     mutex( sout ) sout | 2 | 3;      // unnecessary, but ok because of recursive lock
12     sout | 4;
13 } // 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:

```

18 mutex( sin ) {
19     int x, y, z, w;
20     sin | x;
21     mutex( sin ) sin | y | z;      // unnecessary, but ok because of recursive lock
22     sin | w;
23 } // implicitly release sin lock

```

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

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

```

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

```

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

```

32 int data = rtn( mon );
33 mutex( sout ) sout | "data: " | data;

```

## 22.8 Locale

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

```

36 12,345.123      // comma separator, period decimal-point
37 12.345,123      // period separator, comma decimal-point
38 12_345,123.     // space separator, comma decimal-point, period terminator

```

A locale is selected with function `setlocale`, and the corresponding locale package *must* be installed on the underlying system; `setlocale` returns 0p if the requested locale is unavailable. Furthermore, a locale covers the syntax for many cultural items, *e.g.*, address, measurement, money, etc. This discussion applies to item `LC_NUMERIC` for formatting non-monetary integral and floating-point values. Figure 8 shows selecting different cultural syntax, which may be associated with one or more countries.

## 23 String Stream

The stream types `ostrstream` and `istrstream` provide all the stream formatting capabilities to/from a C string rather than a stream file. Figure 9, p. 61 shows writing (output) to and reading (input) from a C string. The only string stream operations different from a file stream are:

```

#include <fstream.h>
#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 8: Stream Locale

- 1 • constructors to create a stream that writes to a write buffer (ostrstream) of size, or reads from a read buffer
- 2 (istrstream) containing a C string terminated with '\0'.
- 3     **void** ?{ }( ostrstream &, **char** buf[], size\_t size );
- 4     **void** ?{ }( istrstream & is, **char** buf[] );
- 5 • write (ostrstream only) writes all the buffered characters to the specified stream (stdout default).
- 6     ostrstream & write( ostrstream & os, FILE \* stream = stdout );
- 7     There is no read for istrstream.

## 8 24 Structures

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

```

#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_1.234568e+07_987.654n_abc
3_0x5_1.234568e+07_987.654n_abc
3_0x5_1.234568e+07_987.654n_abc
12_14_15_3.5_70000_abc

```

Figure 9: String Stream Processing

```

1  struct Point {
2      double x;
3      double y;
4  };
5
6  Point p = {0.0, 0.0};
7
8  CV does not support inheritance among types, but instead uses composition to enable reuse of structure fields.
9  Composition is achieved by embedding one type into another. When type A is embedded in type B, an object with
10 type B may be used as an object of type A, and the fields of type A are directly accessible. Embedding types is
11 achieved using anonymous members. For example, using Point from above:
12
13 void foo(Point p);
14
15 struct ColoredPoint {
16     Point; // anonymous member (no identifier)
17     int Color;
18 };
19 ...
20 ColoredPoint cp = ...;
21 cp.x = 10.3; // x from Point is accessed directly
22 cp.color = 0x33aaff; // color is accessed normally
23 foo(cp); // cp can be used directly as a Point

```

## 25 Constructors and Destructors

CV supports C initialization of structures, but it also adds constructors for more advanced initialization. Additionally, CV adds destructors that are called when a variable is deallocated (variable goes out of scope or object is deleted). These functions take a reference to the structure as a parameter (see Section 12, p. 19 for more information).

```

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 10: Constructors and Destructors

## 1 26 Overloading

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

10 The compiler follows some overload resolution rules to determine the best interpretation of all of these overloads.  
11 The best valid interpretations are the valid interpretations that use the fewest unsafe conversions. Of these, the best  
12 are those where the functions and objects involved are the least polymorphic. Of these, the best have the lowest total  
13 conversion cost, including all implicit conversions in the argument expressions. Of these, the best have the highest  
14 total conversion cost for the implicit conversions (if any) applied to the argument expressions. If there is no single best



valid interpretation, or if the best valid interpretation is ambiguous, then the resulting interpretation is ambiguous. For details about type inference and overload resolution, please see the C $\forall$  Language Specification.

```

1  int foo(int a, int b) {
2      float sum = 0.0;
3      float special = 1.0;
4      {
5          int sum = 0;
6          // both the float and int versions of sum are available
7          float special = 4.0;
8          // this inner special hides the outer version
9          ...
10     }
11     ...
12 }
13 ...
14 }
```

## 26.1 Constant

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

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

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

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

```

30 struct Complex {
31     double real;
32     double imaginary;
33 }
34
35 const Complex 0 = {0, 0};
36 const Complex 1 = {1, 0};
37 ...
38
39 Complex a = 0;
40 ...
41
42 a++;
43 ...
44 if (a) { // same as if (a == 0)
45     ...
46 }
```

## 26.2 Variable

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

```

52 int pi = 3;
53 float pi = 3.14;
```

```
1 char pi = .p.;
```

## 2 26.3 Function Overloading

3 Overloaded functions in C $\forall$  are resolved based on the number and type of arguments, type of return value, and the  
4 level of specialization required (specialized functions are preferred over generic).

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

```
6 // Choose the one with less conversions
7 int doSomething(int value) {...} // option 1
8 int doSomething(short value) {...} // option 2
9
10 int a, b = 4;
11 short c = 2;
12
13 a = doSomething(b); // chooses option 1
14 a = doSomething(c); // chooses option 2
15
16 // Choose the specialized version over the generic
17
18 generic(type T)
19 T bar(T rhs, T lhs) {...} // option 3
20 float bar(float rhs, float lhs){...} // option 4
21 float a, b, c;
22 double d, e, f;
23 c = bar(a, b); // chooses option 4
24
25 // specialization is preferred over unsafe conversions
26
27 f = bar(d, e); // chooses option 5
```

## 28 26.4 Operator

29 C $\forall$  also allows operators to be overloaded, to simplify the use of user-defined types. Overloading the operators allows  
30 the users to use the same syntax for their custom types that they use for built-in types, increasing readability and  
31 improving productivity. C $\forall$  uses the following special identifiers to name overloaded operators:

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

Table 2: Operator Identifiers

32 These identifiers are defined such that the question marks in the name identify the location of the operands. These  
33 operands represent the parameters to the functions, and define how the operands are mapped to the function call. For  
34 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.

```

1  type Complex = struct { // define a Complex type
2      double real;
3      double imag;
4  }
5
6  // Constructor with default values
7
8  void ?{ }(Complex &c, double real = 0.0, double imag = 0.0) {
9      c.real = real;
10     c.imag = imag;
11 }
12
13 Complex ?+?(Complex lhs, Complex rhs) {
14     Complex sum;
15     sum.real = lhs.real + rhs.real;
16     sum.imag = lhs.imag + rhs.imag;
17     return sum;
18 }
19
20 String ()?(const Complex c) {
21     // use the string conversions for the structure members
22     return (String)c.real + . + . + (String)c.imag + .i.;
23 }
24 ...
25
26 Complex a, b, c = {1.0}; // constructor for c w/ default imag
27 ...
28 c = a + b;
29 print(.sum = . + c);

```

## 27 Auto Type-Inferencing

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

C++

gcc

```

34 auto j = 3.0 * 4;      #define expr 3.0 * i
    int i;               typeof(expr) j = expr; // use type of initialization expression
    auto k = i;           int i;
                        typeof(i) k = i;      // use type of primary variable

```

The two important capabilities are:

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

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

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

and the need to write a routine to compute using j

```

void rtn( ... parm );
    rtn( j );

```

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

```

#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 11: Fibonacci Coroutine

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

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

## 28 Concurrency

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

### 28.1 Coroutine

Coroutines are the precursor to threads. Figure 11 shows a coroutine that computes the Fibonacci numbers.

### 28.2 Monitors

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

```

#include <fstream.h>
#include <thread.h>

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 12: Atomic-Counter Monitor

```

1 monitor is shown here:
2     type Account = monitor {
3         const unsigned long number; // account number
4         float balance; // account balance
5     };

```

## 28.3 Threads

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

## 29 Language Comparisons

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

### 29.1 C++

C++ is a general-purpose programming language. It is an imperative, object-oriented and generic programming language, while also providing facilities for low-level memory manipulation. The primary focus of C++ was adding object-oriented programming to C, and this is the primary difference between C++ and CV. C++ uses classes to encapsulate data and the functions that operate on that data, and to hide the internal representation of the data. CV uses

```

#include <fstream>
#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 13: Simple Threads

1 modules instead to perform these same tasks. Classes in C++ also enable inheritance among types. Instead of inheri-  
 2 tance, C# embraces composition and interfaces to achieve the same goals with more flexibility. There are many studies  
 3 and articles comparing inheritance and composition (or is-a versus has-a relationships), so we will not go into more  
 4 detail here (Venners, 1998) (Pike, Go at Google: Language Design in the Service of Software Engineering , 2012).

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

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

## 29.2 Go

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)

Go and C<sub>V</sub> differ significantly in syntax and implementation, but the underlying core concepts of the two languages are aligned. Both Go and C<sub>V</sub> 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 concurrency features that attempt to simplify the effort and thought process required for writing parallel programs while maintaining high performance.

Go has a significant runtime which handles the scheduling of its light weight threads, and performs garbage collection, among other tasks. C<sub>V</sub> uses a cooperative scheduling algorithm for its tasks, and uses automatic reference counting to enable advanced memory management without garbage collection. This results in Go requiring significant overhead to interface with C libraries while C<sub>V</sub> has no overhead.

## 29.3 Rust

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

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

Aside from those key differences, Rust and C<sub>V</sub> also have several similarities. Both languages support no overhead interoperability with C and have minimal runtimes. Both languages support inheritance and polymorphism through the use of interfaces (traits).

## 29.4 D

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

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

## A Syntax Ambiguities

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

In C<sub>V</sub>, there are ambiguous cases with dereference and operator identifiers, *e.g.*, `int *?*?()`, where the string `*?*?` can be interpreted as:

```
*?_ *?           // dereference operator, dereference operator
*_? *?           // dereference, multiplication operator
```

By default, the first interpretation is selected, which does not yield a meaningful parse. Therefore, C<sub>V</sub> does a lexical look-ahead for the second case, and backtracks to return the leading unary operator and reparses the trailing operator



identifier. Otherwise a space is needed between the unary operator and operator identifier to disambiguate this common case.

A similar issue occurs with the dereference, `*?(...)`, and routine-call, `?()(...)` identifiers. The ambiguity occurs when the dereference operator has no parameters:

```
*?()_... ;
*?()_...(...);
```

requiring arbitrary whitespace look-ahead for the routine-call parameter-list to disambiguate. However, the dereference operator *must* have a parameter/argument to dereference `*?(...)`. Hence, always interpreting the string `*?()` as `*_?()` does not preclude any meaningful program.

The remaining cases are with the increment/decrement operators and conditional expression, *e.g.*:

```
i++?_...(...);
i?+_...(...);
```

requiring arbitrary whitespace look-ahead for the operator parameter-list, even though that interpretation is an incorrect expression (juxtaposed identifiers). Therefore, it is necessary to disambiguate these cases with a space:

```
i++_? i : 0;
i?_++i : 0;
```

## B C Incompatibles

The following incompatibles exist between C<sub>V</sub> and C, and are similar to Annex C for C++ [22].

### 1. **Change:** add new keywords

New keywords are added to C<sub>V</sub> (see Section C, p. 73).

**Rationale:** keywords added to implement new semantics of C<sub>V</sub>.

**Effect on original feature:** change to semantics of well-defined feature.

Any C11 programs using these keywords as identifiers are invalid C<sub>V</sub> programs.

**Difficulty of converting:** keyword clashes are accommodated by syntactic transformations using the C<sub>V</sub> back-quote escape-mechanism (see Section 6, p. 5).

**How widely used:** clashes among new C<sub>V</sub> keywords and existing identifiers are rare.

### 2. **Change:** drop K&R C declarations

K&R declarations allow an implicit base-type of **int**, if no type is specified, plus an alternate syntax for declaring parameters. *e.g.*:

```
x;                // int x
*y;               // int *y
f( p1, p2 );      // int f( int p1, int p2 );
g( p1, p2 ) int p1, p2; // int g( int p1, int p2 );
```

C<sub>V</sub> continues to support K&R routine definitions:

```
f( a, b, c )      // default int return
  int a, b; char c; // K&R parameter declarations
{
  ...
}
```

**Rationale:** dropped from C11 standard.<sup>17</sup>

**Effect on original feature:** original feature is deprecated.

Any old C programs using these K&R declarations are invalid C<sub>V</sub> programs.

**Difficulty of converting:** trivial to convert to C<sub>V</sub>.

**How widely used:** existing usages are rare.

### 3. **Change:** type of character literal **int** to **char** to allow more intuitive overloading:

```
int rtn( int i );
int rtn( char c );
```

<sup>17</sup>At least one type specifier shall be given in the declaration specifiers in each declaration, and in the specifier-qualifier list in each structure declaration and type name [21, § 6.7.2(2)]

```
1         rtn( 'x' );                // programmer expects 2nd rtn to be called
```

2 **Rationale:** it is more intuitive for the call to `rtn` to match the second version of definition of `rtn` rather than the  
3 first. In particular, output of **char** variable now print a character rather than the decimal ASCII value of the  
4 character.

```
5         sout | 'x' | " " | (int)'x';  
6         x 120
```

7 Having to cast `'x'` to **char** is non-intuitive.

8 **Effect on original feature:** change to semantics of well-defined feature that depend on:

```
9         sizeof( 'x' ) == sizeof( int )
```

10 no long work the same in C∀ programs.

11 **Difficulty of converting:** simple

12 **How widely used:** programs that depend upon `sizeof( 'x' )` are rare and can be changed to `sizeof(char)`.

13 4. **Change:** make string literals **const**:

```
14         char * p = "abc";           // valid in C, deprecated in C∀  
15         char * q = expr ? "abc" : "de"; // valid in C, invalid in C∀
```

16 The type of a string literal is changed from `[] char` to **const** `[] char`. Similarly, the type of a wide string literal  
17 is changed from `[] wchar_t` to **const** `[] wchar_t`.

18 **Rationale:** This change is a safety issue:

```
19         char * p = "abc";  
20         p[0] = 'w';                // segment fault or change constant literal
```

21 The same problem occurs when passing a string literal to a routine that changes its argument.

22 **Effect on original feature:** change to semantics of well-defined feature.

23 **Difficulty of converting:** simple syntactic transformation, because string literals can be converted to **char \***.

24 **How widely used:** programs that have a legitimate reason to treat string literals as pointers to potentially modi-  
25 fiable memory are rare.

26 5. **Change:** remove *tentative definitions*, which only occurs at file scope:

```
27         int i;                      // forward definition  
28         int *j = &i;                // forward reference, valid in C, invalid in C∀  
29         int i = 0;                  // definition
```

30 is valid in C, and invalid in C∀ because duplicate overloaded object definitions at the same scope level  
31 are disallowed. This change makes it impossible to define mutually referential file-local static objects, if  
32 initializers are restricted to the syntactic forms of C. For example,

```
33         struct X { int i; struct X *next; };  
34         static struct X a;           // forward definition  
35         static struct X b = { 0, &a }; // forward reference, valid in C, invalid in C∀  
36         static struct X a = { 1, &b }; // definition
```

37 **Rationale:** avoids having different initialization rules for builtin types and user-defined types.

38 **Effect on original feature:** change to semantics of well-defined feature.

39 **Difficulty of converting:** the initializer for one of a set of mutually-referential file-local static objects must  
40 invoke a routine call to achieve the initialization.

41 **How widely used:** seldom

42 6. **Change:** have **struct** introduce a scope for nested types:

```
43         enum Colour { R, G, B, Y, C, M };  
44         struct Person {  
45             enum Colour { R, G, B }; // nested type  
46             struct Face {           // nested type  
47                 Colour Eyes, Hair; // type defined outside (1 level)  
48             };  
49             .Colour shirt;          // type defined outside (top level)  
50             Colour pants;           // type defined same level
```

```

1      Face looks[10];           // type defined same level
2  };
3  Colour c = R;                 // type/enum defined same level
4  Person.Colour pc = Person.R; // type/enum defined inside
5  Person.Face pretty;           // type defined inside

```

In C, the name of the nested types belongs to the same scope as the name of the outermost enclosing structure, *i.e.*, the nested types are hoisted to the scope of the outer-most type, which is not useful and confusing. C $\forall$  is C *incompatible* on this issue, and provides semantics similar to C++. Nested types are not hoisted and can be referenced using the field selection operator “.”, unlike the C++ scope-resolution operator “::”.

**Rationale:** **struct** scope is crucial to C $\forall$  as an information structuring and hiding mechanism.

**Effect on original feature:** change to semantics of well-defined feature.

**Difficulty of converting:** Semantic transformation.

**How widely used:** C programs rarely have nest types because they are equivalent to the hoisted version.

7. **Change:** In C++, the name of a nested class is local to its enclosing class.

**Rationale:** C++ classes have member functions which require that classes establish scopes.

**Difficulty of converting:** Semantic transformation. To make the struct type name visible in the scope of the enclosing struct, the struct tag could be declared in the scope of the enclosing struct, before the enclosing struct is defined. Example:

```

19      struct Y;                 // struct Y and struct X are at the same scope
20      struct X {
21          struct Y { /* ... */ } y;
22      };

```

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:

```

28      void foo() {
29          int * b = malloc( sizeof(int) ); // implicitly convert void * to int *
30          char * c = b;                   // implicitly convert int * to void *, and then void * to char *
31      }

```

**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):

```

35      int * b = (int *)malloc( sizeof(int) );
36      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,

```

40      p = (void*)(struct x {int i;} *)0;
41      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.

## 1 C CV Keywords

2 CV introduces the following new keywords, which cannot be used as identifiers.

```
3     basetypeof, choose, coroutine, disable, enable, exception, fallthrough, fallthru, finally, fixup, forall, generator,
4     int128, monitor, mutex, one_t, report, suspend, throw, throwResume, trait, try, virtual, waitfor, when, with,
5     zero_t
```

6 CV introduces the following new quasi-keywords, which can be used as identifiers.

```
7     catch, catchResume, finally, fixup, or, timeout
```

## 8 D Standard Headers

9 C11 prescribes the following standard header-files [21, § 7.1.2]:

```
10     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,
11     stdalign.h, stdarg.h, stdatomic.h, stdbool.h, stddef.h, stdint.h, stdio.h, stdlib.h, stdnoreturn.h, string.h, tgmath.h,
12     threads.h, time.h, uchar.h, wchar.h, wctype.h
```

13 and CV adds to this list:

```
14     gmp.h, malloc.h, unistd.h
```

15 For the prescribed head-files, CV uses header interposition to wraps these includes in an **extern "C"**; hence, names  
16 in these include files are not mangled (see Section 4, p. 2). All other C header files must be explicitly wrapped in  
17 **extern "C"** to prevent name mangling. This approach is different from C++ where the name-mangling issue is handled  
18 internally in C header-files through checks for preprocessor variable `__cplusplus`, which adds appropriate **extern "C"**  
19 qualifiers.

## 20 E Standard Library

21 The CV standard-library wraps explicitly-polymorphic C routines into implicitly-polymorphic versions.

### 22 E.1 Storage Management

23 The storage-management routines extend their C equivalents by overloading, alternate names, providing shallow type-  
24 safety, and removing the need to specify the allocation size for non-array types.

25 C storage management provides the following capabilities:

26 **filled** after allocation with a specified character or value.

27 **resize** an existing allocation to decreased or increased its size. In either case, new storage may or may not be allocated  
28 and, if there is a new allocation, as much data from the existing allocation is copied into the new allocation. For an  
29 increase in storage size, new storage after the copied data may be filled.

30 **align** an allocation on a specified memory boundary, *e.g.*, an address multiple of 64 or 128 for cache-line purposes.

31 **array** the allocation size is scaled to the specified number of array elements. An array may be filled, resized, or  
32 aligned.

33 Table 3 shows allocation routines supporting different combinations of storage-management capabilities.

34 CV memory management extends the type safety of all allocations by using the type of the left-hand-side type to  
35 determine the allocation size and return a matching type for the new storage. Type-safe allocation is provided for all  
36 C allocation routines and new CV allocation routines, *e.g.*, in

```
37     int * ip = (int *)malloc( sizeof(int) ); // C
38     int * ip = malloc(); // CV type-safe version of C malloc
39     int * ip = alloc(); // CV type-safe uniform alloc
```

40 the latter two allocations determine the allocation size from the type of `p` (**int**) and cast the pointer to the allocated  
41 storage to **int \***.

42 CV memory management extends allocation safety by implicitly honouring all alignment requirements, *e.g.*, in

```
43     struct S { int i; } __attribute__(( aligned( 128 ) )); // cache-line alignment
44     S * sp = malloc(); // honour type alignment
```

		fill	resize	alignment	array
C	malloc	no	no	no	no
	calloc	yes (0 only)	no	no	yes
	realloc	copy	yes	no	no
	memalign	no	no	yes	no
	aligned_alloc <sup>a</sup>	no	no	yes	no
	posix_memalign	no	no	yes	no
	valloc	no	no	yes (page size)	no
	pvalloc <sup>b</sup>	no	no	yes (page size)	no
Cv	cmemalign	yes (0 only)	no	yes	yes
	realloc	copy	yes	yes	no
	alloc	no	yes	no	yes
	alloc_set	yes	yes	no	yes
	alloc_align	no	yes	yes	yes
	alloc_align_set	yes	yes	yes	yes

<sup>a</sup>Same as memalign but size is an integral multiple of alignment, which is universally ignored.

<sup>b</sup>Same as valloc but rounds size to multiple of page size.

Table 3: Allocation Routines versus Storage-Management Capabilities

the storage allocation is implicitly aligned to 128 rather than the default 16. The alignment check is performed at compile time so there is no runtime cost.

Cv memory management extends the resize capability with the notion of *sticky properties*. Hence, initial allocation capabilities are remembered and maintained when resize requires copying. For example, an initial alignment and fill capability are preserved during a resize copy so the copy has the same alignment and extended storage is filled. Without sticky properties it is dangerous to use realloc, resulting in an idiom of manually performing the reallocation to maintain correctness.

Cv memory management extends allocation to support constructors for initialization of allocated storage, *e.g.*, in

```

struct S { int i; }; // cache-line alignment
void ?{ S & s, int i } { s.i = i; }
// assume ?? operator for printing an S

S & sp = *new( 3 ); // call constructor after allocation
sout | sp.i;
delete( &sp );

S * spa = new( 10, 5 ); // allocate array and initialize each array element
for ( i; 10 ) sout | spa[i] | nonl;
sout | nl;
adelete( 10, spa );

```

Allocation routines new/new allocate a variable/array and initialize storage using the allocated type's constructor. Note, the matching deallocation routines delete/adelete.

```

extern "C" {
    // C unsafe allocation
    void * malloc( size_t size );
    void * calloc( size_t dim, size_t size );
    void * realloc( void * ptr, size_t size );
    void * memalign( size_t align, size_t size );
    void * aligned_alloc( size_t align, size_t size );
    int posix_memalign( void ** ptr, size_t align, size_t size );
    void * cmemalign( size_t alignment, size_t noOfElems, size_t elemSize ); // CFA

```

```

1      // C unsafe initialization/copy
2      void * memset( void * dest, int c, size_t size );
3      void * memcpy( void * dest, const void * src, size_t size );
4  }
5
6  void * realloc( void * oaddr, size_t nalign, size_t size ); // CFA heap
7
8  forall( dtype T | sized(T) ) {
9      // CV safe equivalents, i.e., implicit size specification
10     T * malloc( void );
11     T * calloc( size_t dim );
12     T * realloc( T * ptr, size_t size );
13     T * memalign( size_t align );
14     T * cmemalign( size_t align, size_t dim );
15     T * aligned_alloc( size_t align );
16     int posix_memalign( T ** ptr, size_t align );
17
18     // CV safe general allocation, fill, resize, alignment, array
19     T * alloc( void ); // variable, T size
20     T * alloc( size_t dim ); // array[dim], T size elements
21     T * alloc( T ptr[], size_t dim ); // realloc array[dim], T size elements
22
23     T * alloc_set( char fill ); // variable, T size, fill bytes with value
24     T * alloc_set( T fill ); // variable, T size, fill with value
25     T * alloc_set( size_t dim, char fill ); // array[dim], T size elements, fill bytes with value
26     T * alloc_set( size_t dim, T fill ); // array[dim], T size elements, fill elements with value
27     T * alloc_set( size_t dim, const T fill[] ); // array[dim], T size elements, fill elements with array
28     T * alloc_set( T ptr[], size_t dim, char fill ); // realloc array[dim], T size elements, fill bytes with value
29
30     T * alloc_align( size_t align ); // aligned variable, T size
31     T * alloc_align( size_t align, size_t dim ); // aligned array[dim], T size elements
32     T * alloc_align( T ptr[], size_t align ); // realloc new aligned array
33     T * alloc_align( T ptr[], size_t align, size_t dim ); // realloc new aligned array[dim]
34
35     T * alloc_align_set( size_t align, char fill ); // aligned variable, T size, fill bytes with value
36     T * alloc_align_set( size_t align, T fill ); // aligned variable, T size, fill with value
37     T * alloc_align_set( size_t align, size_t dim, char fill ); // aligned array[dim], T size elements, fill bytes with value
38     T * alloc_align_set( size_t align, size_t dim, T fill ); // aligned array[dim], T size elements, fill elements with value
39     T * alloc_align_set( size_t align, size_t dim, const T fill[] ); // aligned array[dim], T size elements, fill elements with array
40     T * alloc_align_set( T ptr[], size_t align, size_t dim, char fill ); // realloc new aligned array[dim], fill new bytes with value
41
42     // CV safe initialization/copy, i.e., implicit size specification
43     T * memset( T * dest, char fill );
44     T * memcpy( T * dest, const T * src );
45
46     // CV safe initialization/copy, i.e., implicit size specification, array types
47     T * amemset( T dest[], char fill, size_t dim );
48     T * amemcpy( T dest[], const T src[], size_t dim );
49 }
50
51 // CV allocation/deallocation and constructor/destructor, non-array types
52 forall( dtype T | sized(T), ttype Params | { void ?{}( T &, Params ); } ) T * new( Params p );
53 forall( dtype T | sized(T) | { void ^{}( T & ); } ) void delete( T * ptr );
54 forall( dtype T, ttype Params | sized(T) | { void ^{}( T & ); void delete( Params ); } )
55     void delete( T * ptr, Params rest );
56
57 // CV allocation/deallocation and constructor/destructor, array types
58 forall( dtype T | sized(T), ttype Params | { void ?{}( T &, Params ); } ) T * anew( size_t dim, Params p );
59 forall( dtype T | sized(T) | { void ^{}( T & ); } ) void adelete( size_t dim, T arr[] );
60 forall( dtype T | sized(T) | { void ^{}( T & ); }, ttype Params | { void adelete( Params ); } )

```

```
1 void adelete( size_t dim, T arr[], Params rest );
```

## 2 E.2 String to Value Conversion

```
3 int ato( const char * ptr );
4 unsigned int ato( const char * ptr );
5 long int ato( const char * ptr );
6 unsigned long int ato( const char * ptr );
7 long long int ato( const char * ptr );
8 unsigned long long int ato( const char * ptr );
9 float ato( const char * ptr );
10 double ato( const char * ptr );
11 long double ato( const char * ptr );
12 float _Complex ato( const char * ptr );
13 double _Complex ato( const char * ptr );
14 long double _Complex ato( const char * ptr );
15
16 int strtol( const char * sptr, char ** eptr, int base );
17 unsigned int strtol( const char * sptr, char ** eptr, int base );
18 long int strtol( const char * sptr, char ** eptr, int base );
19 unsigned long int strtol( const char * sptr, char ** eptr, int base );
20 long long int strtol( const char * sptr, char ** eptr, int base );
21 unsigned long long int strtol( const char * sptr, char ** eptr, int base );
22 float strtol( const char * sptr, char ** eptr );
23 double strtol( const char * sptr, char ** eptr );
24 long double strtol( const char * sptr, char ** eptr );
25 float _Complex strtol( const char * sptr, char ** eptr );
26 double _Complex strtol( const char * sptr, char ** eptr );
27 long double _Complex strtol( const char * sptr, char ** eptr );
```

## 28 E.3 Search / Sort

```
29 forall( T | { int ?<?( T, T ); } ) // location
30 T * bsearch( T key, const T * arr, size_t dim );
31
32 forall( T | { int ?<?( T, T ); } ) // position
33 unsigned int bsearch( T key, const T * arr, size_t dim );
34
35 forall( T | { int ?<?( T, T ); } )
36 void qsort( const T * arr, size_t dim );
37
38 forall( E | { int ?<?( E, E ); } ) {
39     E * bsearch( E key, const E * vals, size_t dim ); // location
40     size_t bsearch( E key, const E * vals, size_t dim ); // position
41     E * bsearchl( E key, const E * vals, size_t dim );
42     size_t bsearchl( E key, const E * vals, size_t dim );
43     E * bsearchu( E key, const E * vals, size_t dim );
44     size_t bsearchu( E key, const E * vals, size_t dim );
45 }
46
47 forall( K, E | { int ?<?( K, K ); K getKey( const E & ); } ) {
48     E * bsearch( K key, const E * vals, size_t dim );
49     size_t bsearch( K key, const E * vals, size_t dim );
50     E * bsearchl( K key, const E * vals, size_t dim );
51     size_t bsearchl( K key, const E * vals, size_t dim );
52     E * bsearchu( K key, const E * vals, size_t dim );
53     size_t bsearchu( K key, const E * vals, size_t dim );
54 }
55
56 forall( E | { int ?<?( E, E ); } ) {
```



```

1      void qsort( E * vals, size_t dim );
2  }
```

### 3 E.4 Absolute Value

```

4      unsigned char abs( signed char );
5      int abs( int );
6      unsigned long int abs( long int );
7      unsigned long long int abs( long long int );
8      float abs( float );
9      double abs( double );
10     long double abs( long double );
11     float abs( float _Complex );
12     double abs( double _Complex );
13     long double abs( long double _Complex );
14     forall( T | { void ?( T *, zero_t ); int ?<?( T, T ); T -?( T ); } )
15     T abs( T );
```

### 16 E.5 C Random Numbers

```

17     void srand( unsigned int seed );
18     char random( void );
19     char random( char u ); // [0,u]
20     char random( char l, char u ); // [l,u]
21     int random( void );
22     int random( int u ); // [0,u]
23     int random( int l, int u ); // [l,u]
24     unsigned int random( void );
25     unsigned int random( unsigned int u ); // [0,u]
26     unsigned int random( unsigned int l, unsigned int u ); // [l,u]
27     long int random( void );
28     long int random( long int u ); // [0,u]
29     long int random( long int l, long int u ); // [l,u]
30     unsigned long int random( void );
31     unsigned long int random( unsigned long int u ); // [0,u]
32     unsigned long int random( unsigned long int l, unsigned long int u ); // [l,u]
33     float random( void ); // [0.0, 1.0]
34     double random( void ); // [0.0, 1.0]
35     float _Complex random( void ); // [0.0, 1.0)+[0.0, 1.0)i
36     double _Complex random( void ); // [0.0, 1.0)+[0.0, 1.0)i
37     long double _Complex random( void ); // [0.0, 1.0)+[0.0, 1.0)i
```

### 38 E.6 Algorithms

```

39     forall( T | { int ?<?( T, T ); } ) T min( T t1, T t2 );
40     forall( T | { int ?>?( T, T ); } ) T max( T t1, T t2 );
41     forall( T | { T min( T, T ); T max( T, T ); } ) T clamp( T value, T min_val, T max_val );
42     forall( T ) void swap( T * t1, T * t2 );
```

## 43 F Math Library

44 The CV math-library wraps explicitly-polymorphic C math-routines into implicitly-polymorphic versions.

### 45 F.1 General

```

46     float ?%( float, float );
47     float fmod( float, float );
48     double ?%( double, double );
49     double fmod( double, double );
50     long double ?%( long double, long double );
51     long double fmod( long double, long double );
```

```

1
2  float remainder( float, float );
3  double remainder( double, double );
4  long double remainder( long double, long double );
5
6  float remquo( float, float, int * );
7  double remquo( double, double, int * );
8  long double remquo( long double, long double, int * );
9  [ int, float ] remquo( float, float );
10 [ int, double ] remquo( double, double );
11 [ int, long double ] remquo( long double, long double );
12
13 [ int, float ] div( float, float );
14 [ int, double ] div( double, double );
15 [ int, long double ] div( long double, long double );
16
17 float fma( float, float, float );
18 double fma( double, double, double );
19 long double fma( long double, long double, long double );
20
21 float fdim( float, float );
22 double fdim( double, double );
23 long double fdim( long double, long double );
24
25 float nan( const char * );
26 double nan( const char * );
27 long double nan( const char * );

```

## 28 F.2 Exponential

```

29 float exp( float );
30 double exp( double );
31 long double exp( long double );
32 float _Complex exp( float _Complex );
33 double _Complex exp( double _Complex );
34 long double _Complex exp( long double _Complex );
35
36 float exp2( float );
37 double exp2( double );
38 long double exp2( long double );
39 // float _Complex exp2( float _Complex );
40 // double _Complex exp2( double _Complex );
41 // long double _Complex exp2( long double _Complex );
42
43 float expm1( float );
44 double expm1( double );
45 long double expm1( long double );
46
47 float pow( float, float );
48 double pow( double, double );
49 long double pow( long double, long double );
50 float _Complex pow( float _Complex, float _Complex );
51 double _Complex pow( double _Complex, double _Complex );
52 long double _Complex pow( long double _Complex, long double _Complex );

```

## 53 F.3 Logarithm

```

54 float log( float );
55 double log( double );
56 long double log( long double );

```

```

1  float _Complex log( float _Complex );
2  double _Complex log( double _Complex );
3  long double _Complex log( long double _Complex );
4
5  int log2( unsigned int );
6  long int log2( unsigned long int );
7  long long int log2( unsigned long long int )
8  float log2( float );
9  double log2( double );
10 long double log2( long double );
11 // float _Complex log2( float _Complex );
12 // double _Complex log2( double _Complex );
13 // long double _Complex log2( long double _Complex );
14
15 float log10( float );
16 double log10( double );
17 long double log10( long double );
18 // float _Complex log10( float _Complex );
19 // double _Complex log10( double _Complex );
20 // long double _Complex log10( long double _Complex );
21
22 float log1p( float );
23 double log1p( double );
24 long double log1p( long double );
25
26 int ilogb( float );
27 int ilogb( double );
28 int ilogb( long double );
29
30 float logb( float );
31 double logb( double );
32 long double logb( long double );
33
34 float sqrt( float );
35 double sqrt( double );
36 long double sqrt( long double );
37 float _Complex sqrt( float _Complex );
38 double _Complex sqrt( double _Complex );
39 long double _Complex sqrt( long double _Complex );
40
41 float cbrt( float );
42 double cbrt( double );
43 long double cbrt( long double );
44
45 float hypot( float, float );
46 double hypot( double, double );
47 long double hypot( long double, long double );

```

#### 48 F.4 Trigonometric

```

49 float sin( float );
50 double sin( double );
51 long double sin( long double );
52 float _Complex sin( float _Complex );
53 double _Complex sin( double _Complex );
54 long double _Complex sin( long double _Complex );
55
56 float cos( float );
57 double cos( double );
58 long double cos( long double );

```

```

1  float _Complex cos( float _Complex );
2  double _Complex cos( double _Complex );
3  long double _Complex cos( long double _Complex );
4
5  float tan( float );
6  double tan( double );
7  long double tan( long double );
8  float _Complex tan( float _Complex );
9  double _Complex tan( double _Complex );
10 long double _Complex tan( long double _Complex );
11
12 float asin( float );
13 double asin( double );
14 long double asin( long double );
15 float _Complex asin( float _Complex );
16 double _Complex asin( double _Complex );
17 long double _Complex asin( long double _Complex );
18
19 float acos( float );
20 double acos( double );
21 long double acos( long double );
22 float _Complex acos( float _Complex );
23 double _Complex acos( double _Complex );
24 long double _Complex acos( long double _Complex );
25
26 float atan( float );
27 double atan( double );
28 long double atan( long double );
29 float _Complex atan( float _Complex );
30 double _Complex atan( double _Complex );
31 long double _Complex atan( long double _Complex );
32
33 float atan2( float, float );
34 double atan2( double, double );
35 long double atan2( long double, long double );
36
37 float atan( float, float );           // alternative name for atan2
38 double atan( double, double );
39 long double atan( long double, long double );

```

## 40 F.5 Hyperbolic

```

41 float sinh( float );
42 double sinh( double );
43 long double sinh( long double );
44 float _Complex sinh( float _Complex );
45 double _Complex sinh( double _Complex );
46 long double _Complex sinh( long double _Complex );
47
48 float cosh( float );
49 double cosh( double );
50 long double cosh( long double );
51 float _Complex cosh( float _Complex );
52 double _Complex cosh( double _Complex );
53 long double _Complex cosh( long double _Complex );
54
55 float tanh( float );
56 double tanh( double );
57 long double tanh( long double );
58 float _Complex tanh( float _Complex );

```

```

1  double _Complex tanh( double _Complex );
2  long double _Complex tanh( long double _Complex );
3
4  float asinh( float );
5  double asinh( double );
6  long double asinh( long double );
7  float _Complex asinh( float _Complex );
8  double _Complex asinh( double _Complex );
9  long double _Complex asinh( long double _Complex );
10
11 float acosh( float );
12 double acosh( double );
13 long double acosh( long double );
14 float _Complex acosh( float _Complex );
15 double _Complex acosh( double _Complex );
16 long double _Complex acosh( long double _Complex );
17
18 float atanh( float );
19 double atanh( double );
20 long double atanh( long double );
21 float _Complex atanh( float _Complex );
22 double _Complex atanh( double _Complex );
23 long double _Complex atanh( long double _Complex );

```

## 24 F.6 Error / Gamma

```

25 float erf( float );
26 double erf( double );
27 long double erf( long double );
28 float _Complex erf( float _Complex );
29 double _Complex erf( double _Complex );
30 long double _Complex erf( long double _Complex );
31
32 float erfc( float );
33 double erfc( double );
34 long double erfc( long double );
35 float _Complex erfc( float _Complex );
36 double _Complex erfc( double _Complex );
37 long double _Complex erfc( long double _Complex );
38
39 float lgamma( float );
40 double lgamma( double );
41 long double lgamma( long double );
42 float lgamma( float, int * );
43 double lgamma( double, int * );
44 long double lgamma( long double, int * );
45
46 float tgamma( float );
47 double tgamma( double );
48 long double tgamma( long double );

```

## 49 F.7 Nearest Integer

```

50 // n / align * align
51 signed char floor( signed char n, signed char align );
52 unsigned char floor( unsigned char n, unsigned char align );
53 short int floor( short int n, short int align );
54 unsigned short int floor( unsigned short int n, unsigned short int align );
55 int floor( int n, int align );
56 unsigned int floor( unsigned int n, unsigned int align );

```

```

1  long int floor( long int n, long int align );
2  unsigned long int floor( unsigned long int n, unsigned long int align );
3  long long int floor( long long int n, long long int align );
4  unsigned long long int floor( unsigned long long int n, unsigned long long int align );
5
6  // (n + (align - 1)) / align
7  signed char ceiling_div( signed char n, char align );
8  unsigned char ceiling_div( unsigned char n, unsigned char align );
9  short int ceiling_div( short int n, short int align );
10 unsigned short int ceiling_div( unsigned short int n, unsigned short int align );
11 int ceiling_div( int n, int align );
12 unsigned int ceiling_div( unsigned int n, unsigned int align );
13 long int ceiling_div( long int n, long int align );
14 unsigned long int ceiling_div( unsigned long int n, unsigned long int align );
15 long long int ceiling_div( long long int n, long long int align );
16 unsigned long long int ceiling_div( unsigned long long int n, unsigned long long int align );
17
18 // floor( n + (n % align != 0 ? align - 1 : 0), align )
19 signed char ceiling( signed char n, signed char align );
20 unsigned char ceiling( unsigned char n, unsigned char align );
21 short int ceiling( short int n, short int align );
22 unsigned short int ceiling( unsigned short int n, unsigned short int align );
23 int ceiling( int n, int align );
24 unsigned int ceiling( unsigned int n, unsigned int align );
25 long int ceiling( long int n, long int align );
26 unsigned long int ceiling( unsigned long int n, unsigned long int align );
27 long long int ceiling( long long int n, long long int align );
28 unsigned long long int ceiling( unsigned long long int n, unsigned long long int align );
29
30 float floor( float );
31 double floor( double );
32 long double floor( long double );
33
34 float ceil( float );
35 double ceil( double );
36 long double ceil( long double );
37
38 float trunc( float );
39 double trunc( double );
40 long double trunc( long double );
41
42 float rint( float );
43 long double rint( long double );
44 long int rint( float );
45 long int rint( double );
46 long int rint( long double );
47 long long int rint( float );
48 long long int rint( double );
49 long long int rint( long double );
50
51 long int lrint( float );
52 long int lrint( double );
53 long int lrint( long double );
54 long long int llrint( float );
55 long long int llrint( double );
56 long long int llrint( long double );
57
58 float nearbyint( float );
59 double nearbyint( double );
60 long double nearbyint( long double );

```

```

1
2  float round( float );
3  long double round( long double );
4  long int round( float );
5  long int round( double );
6  long int round( long double );
7  long long int round( float );
8  long long int round( double );
9  long long int round( long double );
10
11 long int lround( float );
12 long int lround( double );
13 long int lround( long double );
14 long long int llround( float );
15 long long int llround( double );
16 long long int llround( long double );

```

## 17 F.8 Manipulation

```

18 float copysign( float, float );
19 double copysign( double, double );
20 long double copysign( long double, long double );
21
22 float frexp( float, int * );
23 double frexp( double, int * );
24 long double frexp( long double, int * );
25
26 float ldexp( float, int );
27 double ldexp( double, int );
28 long double ldexp( long double, int );
29
30 [ float, float ] modf( float );
31 float modf( float, float * );
32 [ double, double ] modf( double );
33 double modf( double, double * );
34 [ long double, long double ] modf( long double );
35 long double modf( long double, long double * );
36
37 float nextafter( float, float );
38 double nextafter( double, double );
39 long double nextafter( long double, long double );
40
41 float nexttoward( float, long double );
42 double nexttoward( double, long double );
43 long double nexttoward( long double, long double );
44
45 float scalbn( float, int );
46 double scalbn( double, int );
47 long double scalbn( long double, int );
48
49 float scalbln( float, long int );
50 double scalbln( double, long int );
51 long double scalbln( long double, long int );

```

## 52 G Time Keeping

### 53 G.1 Duration

```

54 struct Duration {
55     int64_t tn;
56     // nanoseconds

```



```

1   };
2
3   void ?{}( Duration & dur );
4   void ?{}( Duration & dur, zero_t );
5   void ?{}( Duration & dur, timeval t );
6   void ?{}( Duration & dur, timespec t );
7
8   Duration ?=( Duration & dur, zero_t );
9   Duration ?=( Duration & dur, timeval t );
10  Duration ?=( Duration & dur, timespec t );
11
12  Duration +?( Duration rhs );
13  Duration ?+?( Duration & lhs, Duration rhs );
14  Duration ?+=?( Duration & lhs, Duration rhs );
15
16  Duration -( Duration rhs );
17  Duration ?-( Duration & lhs, Duration rhs );
18  Duration ?-=?( Duration & lhs, Duration rhs );
19
20  Duration ?*( Duration lhs, int64_t rhs );
21  Duration ?*( int64_t lhs, Duration rhs );
22  Duration ?*=( Duration & lhs, int64_t rhs );
23
24  int64_t ?/( Duration lhs, Duration rhs );
25  Duration ?/( Duration lhs, int64_t rhs );
26  Duration ?/=( Duration & lhs, int64_t rhs );
27  double div( Duration lhs, Duration rhs );
28
29  Duration ?%( Duration lhs, Duration rhs );
30  Duration ?%=( Duration & lhs, Duration rhs );
31
32  bool ?==( Duration lhs, zero_t );
33  bool ?!==( Duration lhs, zero_t );
34  bool ?<? ( Duration lhs, zero_t );
35  bool ?<=( Duration lhs, zero_t );
36  bool ?>? ( Duration lhs, zero_t );
37  bool ?>=( Duration lhs, zero_t );
38
39  bool ?==( Duration lhs, Duration rhs );
40  bool ?!==( Duration lhs, Duration rhs );
41  bool ?<? ( Duration lhs, Duration rhs );
42  bool ?<=( Duration lhs, Duration rhs );
43  bool ?>? ( Duration lhs, Duration rhs );
44  bool ?>=( Duration lhs, Duration rhs );
45
46  Duration abs( Duration rhs );
47
48  Duration ?`ns( int64_t nsec );
49  Duration ?`us( int64_t usec );
50  Duration ?`ms( int64_t msec );
51  Duration ?`s( int64_t sec );
52  Duration ?`s( double sec );
53  Duration ?`m( int64_t min );
54  Duration ?`m( double min );
55  Duration ?`h( int64_t hours );
56  Duration ?`h( double hours );
57  Duration ?`d( int64_t days );
58  Duration ?`d( double days );
59  Duration ?`w( int64_t weeks );
60  Duration ?`w( double weeks );

```

```

1
2  int64_t ?`ns( Duration dur );
3  int64_t ?`us( Duration dur );
4  int64_t ?`ms( Duration dur );
5  int64_t ?`s( Duration dur );
6  int64_t ?`m( Duration dur );
7  int64_t ?`h( Duration dur );
8  int64_t ?`d( Duration dur );
9  int64_t ?`w( Duration dur );
10
11  double ?`dns( Duration dur );
12  double ?`dus( Duration dur );
13  double ?`dms( Duration dur );
14  double ?`ds( Duration dur );
15  double ?`dm( Duration dur );
16  double ?`dh( Duration dur );
17  double ?`dd( Duration dur );
18  double ?`dw( Duration dur );
19
20  Duration max( Duration lhs, Duration rhs );
21  Duration min( Duration lhs, Duration rhs );
22
23  forall( ostream & | ostream( ostream ) ) ostream & ?|?( ostream & os, Duration dur );

```

## 24 G.2 timeval

```

25  void ?{}( timeval & t );
26  void ?{}( timeval & t, zero_t );
27  void ?{}( timeval & t, time_t sec, suseconds_t usec );
28  void ?{}( timeval & t, time_t sec );
29  void ?{}( timeval & t, Time time );
30
31  timeval ?=( timeval & t, zero_t );
32  timeval ?+?( timeval & lhs, timeval rhs );
33  timeval ?-?( timeval & lhs, timeval rhs );
34  bool ?==?( timeval lhs, timeval rhs );
35  bool ?!==( timeval lhs, timeval rhs );

```

## 36 G.3 timespec

```

37  void ?{}( timespec & t );
38  void ?{}( timespec & t, zero_t );
39  void ?{}( timespec & t, time_t sec, __syscall_slong_t nsec );
40  void ?{}( timespec & t, time_t sec );
41  void ?{}( timespec & t, Time time );
42
43  timespec ?=( timespec & t, zero_t );
44  timespec ?+?( timespec & lhs, timespec rhs );
45  timespec ?-?( timespec & lhs, timespec rhs );
46  bool ?==?( timespec lhs, timespec rhs );
47  bool ?!==( timespec lhs, timespec rhs );

```

## 48 G.4 itimerval

```

49  void ?{}( itimerval & itv, Duration alarm );
50  void ?{}( itimerval & itv, Duration alarm, Duration interval );

```

## 51 G.5 Time

```

52  struct Time {
53      uint64_t tn;

```

*// nanoseconds since UNIX epoch*

```

1   };
2
3   void ?{}( Time & time );
4   void ?{}( Time & time, zero_t );
5   void ?{}( Time & time, timeval t );
6   void ?{}( Time & time, timespec t );
7
8   Time ?=( Time & time, zero_t );
9   Time ?=( Time & time, timeval t );
10  Time ?=( Time & time, timespec t );
11
12  Time ?+?( Time & lhs, Duration rhs );
13  Time ?+?( Duration lhs, Time rhs );
14  Time ?+=( Time & lhs, Duration rhs );
15
16  Duration ?-( Time lhs, Time rhs );
17  Time ?-( Time lhs, Duration rhs );
18  Time ?-=( Time & lhs, Duration rhs );
19  bool ?==( Time lhs, Time rhs );
20  bool ?!==( Time lhs, Time rhs );
21  bool ?<?( Time lhs, Time rhs );
22  bool ?<=( Time lhs, Time rhs );
23  bool ?>?( Time lhs, Time rhs );
24  bool ?>=( Time lhs, Time rhs );
25
26  int64_t ?`ns( Time t );
27
28  char * yy_mm_dd( Time time, char * buf );
29  char * ?`ymd( Time time, char * buf ); // short form
30
31  char * mm_dd_yy( Time time, char * buf );
32  char * ?`mdy( Time time, char * buf ); // short form
33
34  char * dd_mm_yy( Time time, char * buf );
35  char * ?`dmy( Time time, char * buf ); // short form
36
37  size_t strftime( char * buf, size_t size, const char * fmt, Time time );
38
39  forall( ostype & | ostream( ostype ) ) ostype & ?|?( ostype & os, Time time );

```

## 40 H Clock

### 41 H.1 C time

```

42  char * ctime( time_t tp );
43  char * ctime_r( time_t tp, char * buf );
44  tm * gmtime( time_t tp );
45  tm * gmtime_r( time_t tp, tm * result );
46  tm * localtime( time_t tp );
47  tm * localtime_r( time_t tp, tm * result );

```

### 48 H.2 Clock

```

49  struct Clock {                                     // virtual clock
50      Duration offset;                             // offset from computer real-time
51  };
52
53  void ?{}( Clock & clk );                          // create no offset
54  void ?{}( Clock & clk, Duration adj );             // create with offset
55  void reset( Clock & clk, Duration adj );          // change offset

```

```

1
2   Duration resolutionHi();           // clock resolution in nanoseconds (fine)
3   Duration resolution();             // clock resolution without nanoseconds (coarse)
4
5   Time timeHiRes();                  // real time with nanoseconds
6   Time time();                       // real time without nanoseconds
7   Time time( Clock & clk );          // real time for given clock
8   Time ?()( Clock & clk );           // alternative syntax
9   timeval time( Clock & clk );       // convert to C time format
10  tm time( Clock & clk );
11  Duration processor();               // non-monotonic duration of kernel thread
12  Duration program();                // non-monotonic duration of program CPU
13  Duration boot();                   // monotonic duration since computer boot

```

## 14 I Pseudo Random Number Generator

Random numbers are values generated independently, i.e., new values do not depend on previous values (independent trials), *e.g.*, lottery numbers, shuffled cards, dice roll, coin flip. While a primary goal of programming is computing values that are *not* random, random values are useful in simulation, cryptography, games, etc. A random-number generator is an algorithm that computes independent values. If the algorithm uses deterministic computation (a predictable sequence of values), it generates *pseudo* random numbers versus *true* random numbers.

All *pseudo random-number generators (PRNG)* involve some technique to scramble bits of a value, *e.g.*, multiplicative recurrence:

```

22   rand = 33967 * (rand + 1063); // scramble bits

```

Multiplication of large values adds new least-significant bits and drops most-significant bits.

	bits 63–32 (most)	bits 31–0 (least)
	0x0	0x3e8e36
	0x5f	0x718c25e1
	0xad3e	0x7b5f1dbe
	0xbc3b	0xac69ff19
	0x1070f	0x2d258dc6

By dropping bits 63–32, bits 31–0 become scrambled after each multiply. The least-significant bits *appear* random but the same bits are always generated given a fixed starting value, called the *seed* (value 0x3e8e36 above). Hence, if a program uses the same seed, the same sequence of pseudo-random values is generated from the PRNG. Often the seed is set to another random value like a program's process identifier (getpid) or time when the program is run; hence, one random value bootstraps another. Finally, a PRNG usually generates a range of large values, *e.g.*, [0, UINT\_MAX], which are scaled using the modulus operator, *e.g.*, prng() % 5 produces random values in the range 0–4.

Cv provides 32/64-bit sequential PRNG type only accessible by a single thread (not thread-safe) and a set of global routines and companion thread PRNG functions accessible by multiple threads without contention. To use the PRNG interface **requires including `stdlib.h`**.

• The PRNG types for sequential programs, including coroutines, are:

```

35   struct PRNG32 {};                  // opaque type, no copy or assignment
36   void ?{}( PRNG32 & prng, uint32_t seed ); // fixed seed
37   void ?{}( PRNG32 & prng );          // random seed
38   void set_seed( PRNG32 & prng, uint32_t seed ); // set seed
39   uint32_t get_seed( PRNG32 & prng ); // get seed
40   uint32_t prng( PRNG32 & prng );      // [0,UINT_MAX]
41   uint32_t prng( PRNG32 & prng, uint32_t u ); // [0,u]
42   uint32_t prng( PRNG32 & prng, uint32_t l, uint32_t u ); // [l,u]
43   uint32_t calls( PRNG32 & prng );     // number of calls
44   void copy( PRNG32 & dst, PRNG32 & src ); // checkpoint PRNG state

45   struct PRNG64 {};                  // opaque type, no copy or assignment
46   void ?{}( PRNG64 & prng, uint64_t seed ); // fixed seed
47   void ?{}( PRNG64 & prng );          // random seed

```

```

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 14: Sequential PRNG

```

1 void set_seed( PRNG64 & prng, uint64_t seed );           // set seed
2 uint64_t get_seed( PRNG64 & prng );                     // get seed
3 uint64_t prng( PRNG64 & prng );                         // [0,UINT_MAX]
4 uint64_t prng( PRNG64 & prng, uint64_t u );              // [0,u]
5 uint64_t prng( PRNG64 & prng, uint64_t l, uint64_t u ); // [l,u]
6 uint64_t calls( PRNG64 & prng );                        // number of calls
7 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 14 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 )`.

```

18 void set_seed( size_t seed );                           // set global seed
19 size_t get_seed();                                     // get global seed
20 // SLOWER, global routines
21 size_t prng( void );                                   // [0,UINT_MAX]
22 size_t prng( size_t u );                               // [0,u]
23 size_t prng( size_t l, size_t u );                     // [l,u]
24 // FASTER, thread members
25 size_t prng( thread$ & th );                           // [0,UINT_MAX]
26 size_t prng( thread$ & th, size_t u );                 // [0,u]
27 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 sequen-

```

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 15: Concurrent PRNG

1 tial PRNG in each thread. The slower prng global functions, *i.e.*, *without* a thread argument, call `active_thread`  
2 internally to indirectly access the current thread's PRNG state, while the faster prng functions, *i.e.*, *with* a thread  
3 argument, directly access the thread through the thread parameter. If a thread pointer is available, *e.g.*, in thread  
4 main, eliminating the call to `active_thread` significantly reduces the cost of accessing the thread's PRNG state.  
5 Figure 15 shows an example using the slower/faster concurrent PRNG in the program main and a thread.

## 6 J Multi-precision Integers

7 C $\forall$  has an interface to the GMP multi-precision signed-integers [17], similar to the C++ interface provided by GMP. The  
8 C $\forall$  interface wraps GMP routines into operator routines to make programming with multi-precision integers identical  
9 to using fixed-sized integers. The C $\forall$  type name for multi-precision signed-integers is `Int` and the header file is `gmp`.

```

10 void ?{}( Int * this ); // constructor/destructor
11 void ?{}( Int * this, Int init );
12 void ?{}( Int * this, zero_t );
13 void ?{}( Int * this, one_t );
14 void ?{}( Int * this, signed long int init );
15 void ?{}( Int * this, unsigned long int init );
16 void ?{}( Int * this, const char * val );
17 void ^?{}( Int * this );

18
19 Int ?=( Int * lhs, Int rhs ); // assignment
20 Int ?=( Int * lhs, long int rhs );
21 Int ?=( Int * lhs, unsigned long int rhs );
22 Int ?=( Int * lhs, const char * rhs );
23

```

```

1  char ?=( char * lhs, Int rhs );
2  short int ?=( short int * lhs, Int rhs );
3  int ?=( int * lhs, Int rhs );
4  long int ?=( long int * lhs, Int rhs );
5  unsigned char ?=( unsigned char * lhs, Int rhs );
6  unsigned short int ?=( unsigned short int * lhs, Int rhs );
7  unsigned int ?=( unsigned int * lhs, Int rhs );
8  unsigned long int ?=( unsigned long int * lhs, Int rhs );
9
10 long int narrow( Int val );
11 unsigned long int narrow( Int val );
12
13 int ?==( Int oper1, Int oper2 );           // comparison
14 int ?==( Int oper1, long int oper2 );
15 int ?==( long int oper2, Int oper1 );
16 int ?==( Int oper1, unsigned long int oper2 );
17 int ?==( unsigned long int oper2, Int oper1 );
18
19 int ?!=( Int oper1, Int oper2 );
20 int ?!=( Int oper1, long int oper2 );
21 int ?!=( long int oper1, Int oper2 );
22 int ?!=( Int oper1, unsigned long int oper2 );
23 int ?!=( unsigned long int oper1, Int oper2 );
24
25 int ?<( Int oper1, Int oper2 );
26 int ?<( Int oper1, long int oper2 );
27 int ?<( long int oper2, Int oper1 );
28 int ?<( Int oper1, unsigned long int oper2 );
29 int ?<( unsigned long int oper2, Int oper1 );
30
31 int ?<=( Int oper1, Int oper2 );
32 int ?<=( Int oper1, long int oper2 );
33 int ?<=( long int oper2, Int oper1 );
34 int ?<=( Int oper1, unsigned long int oper2 );
35 int ?<=( unsigned long int oper2, Int oper1 );
36
37 int ?>( Int oper1, Int oper2 );
38 int ?>( Int oper1, long int oper2 );
39 int ?>( long int oper1, Int oper2 );
40 int ?>( Int oper1, unsigned long int oper2 );
41 int ?>( unsigned long int oper1, Int oper2 );
42
43 int ?>=( Int oper1, Int oper2 );
44 int ?>=( Int oper1, long int oper2 );
45 int ?>=( long int oper1, Int oper2 );
46 int ?>=( Int oper1, unsigned long int oper2 );
47 int ?>=( unsigned long int oper1, Int oper2 );
48
49 Int +( Int oper );           // arithmetic
50 Int -( Int oper );
51 Int ~( Int oper );
52
53 Int &&( Int oper1, Int oper2 );
54 Int &&( Int oper1, long int oper2 );
55 Int &&( long int oper1, Int oper2 );
56 Int &&( Int oper1, unsigned long int oper2 );
57 Int &&( unsigned long int oper1, Int oper2 );
58 Int &&=( Int * lhs, Int rhs );
59
60 Int ?|( Int oper1, Int oper2 );

```



```

1      Int ?|( Int oper1, long int oper2 );
2      Int ?|( long int oper1, Int oper2 );
3      Int ?|( Int oper1, unsigned long int oper2 );
4      Int ?|( unsigned long int oper1, Int oper2 );
5      Int ?|=( Int * lhs, Int rhs );
6
7      Int ?^( Int oper1, Int oper2 );
8      Int ?^( Int oper1, long int oper2 );
9      Int ?^( long int oper1, Int oper2 );
10     Int ?^( Int oper1, unsigned long int oper2 );
11     Int ?^( unsigned long int oper1, Int oper2 );
12     Int ?^=( Int * lhs, Int rhs );
13
14     Int ?+?( Int addend1, Int addend2 );
15     Int ?+?( Int addend1, long int addend2 );
16     Int ?+?( long int addend2, Int addend1 );
17     Int ?+?( Int addend1, unsigned long int addend2 );
18     Int ?+?( unsigned long int addend2, Int addend1 );
19     Int ?+=?( Int * lhs, Int rhs );
20     Int ?+=?( Int * lhs, long int rhs );
21     Int ?+=?( Int * lhs, unsigned long int rhs );
22     Int ++?( Int * lhs );
23     Int ?++( Int * lhs );
24
25     Int ?-( Int minuend, Int subtrahend );
26     Int ?-( Int minuend, long int subtrahend );
27     Int ?-( long int minuend, Int subtrahend );
28     Int ?-( Int minuend, unsigned long int subtrahend );
29     Int ?-( unsigned long int minuend, Int subtrahend );
30     Int ?-=( Int * lhs, Int rhs );
31     Int ?-=( Int * lhs, long int rhs );
32     Int ?-=( Int * lhs, unsigned long int rhs );
33     Int --?( Int * lhs );
34     Int ?--( Int * lhs );
35
36     Int ?*( Int multiplicator, Int multiplicand );
37     Int ?*( Int multiplicator, long int multiplicand );
38     Int ?*( long int multiplicand, Int multiplicator );
39     Int ?*( Int multiplicator, unsigned long int multiplicand );
40     Int ?*( unsigned long int multiplicand, Int multiplicator );
41     Int ?*=( Int * lhs, Int rhs );
42     Int ?*=( Int * lhs, long int rhs );
43     Int ?*=( Int * lhs, unsigned long int rhs );
44
45     Int ?/( Int dividend, Int divisor );
46     Int ?/( Int dividend, unsigned long int divisor );
47     Int ?/( unsigned long int dividend, Int divisor );
48     Int ?/( Int dividend, long int divisor );
49     Int ?/( long int dividend, Int divisor );
50     Int ?/=( Int * lhs, Int rhs );
51     Int ?/=( Int * lhs, long int rhs );
52     Int ?/=( Int * lhs, unsigned long int rhs );
53
54     [ Int, Int ] div( Int dividend, Int divisor );
55     [ Int, Int ] div( Int dividend, unsigned long int divisor );
56
57     Int ?%?( Int dividend, Int divisor );
58     Int ?%?( Int dividend, unsigned long int divisor );
59     Int ?%?( unsigned long int dividend, Int divisor );
60     Int ?%?( Int dividend, long int divisor );

```

```

1   Int ?%( long int dividend, Int divisor );
2   Int ?%=( Int * lhs, Int rhs );
3   Int ?%=( Int * lhs, long int rhs );
4   Int ?%=( Int * lhs, unsigned long int rhs );
5
6   Int ?<<?( Int shiften, mp_bitcnt_t shift );
7   Int ?<=<?( Int * lhs, mp_bitcnt_t shift );
8   Int ?>>?( Int shiften, mp_bitcnt_t shift );
9   Int ?>=>?( Int * lhs, mp_bitcnt_t shift );
10
11  Int abs( Int oper );                               // number functions
12  Int fact( unsigned long int N );
13  Int gcd( Int oper1, Int oper2 );
14  Int pow( Int base, unsigned long int exponent );
15  Int pow( unsigned long int base, unsigned long int exponent );
16  void srand( gmp_randstate_t state );
17  Int random( gmp_randstate_t state, mp_bitcnt_t n );
18  Int random( gmp_randstate_t state, Int n );
19  Int random( gmp_randstate_t state, mp_size_t max_size );
20  int sgn( Int oper );
21  Int sqrt( Int oper );
22
23  forall( dtype istype | istream( istype ) ) istype * ?!( istype * is, Int * mp ); // I/O
24  forall( dtype ostype | ostream( ostype ) ) ostype * ?!( ostype * os, Int mp );

```

Figure 16 shows C++ and C factorial programs using the GMP interfaces. (Compile with flag `-lgmp` to link with the GMP library.)

## 27 K Rational Numbers

Rational numbers are numbers written as a ratio, *i.e.*, as a fraction, where the numerator (top number) and the denominator (bottom number) are whole numbers. When creating and computing with rational numbers, results are constantly reduced to keep the numerator and denominator as small as possible.

```

31  // implementation
32  struct Rational {
33      long int numerator, denominator;           // invariant: denominator > 0
34  }; // Rational
35
36  Rational rational();                           // constructors
37  Rational rational( long int n );
38  Rational rational( long int n, long int d );
39  void ?!?( Rational * r, zero_t );
40  void ?!?( Rational * r, one_t );
41
42  long int numerator( Rational r );               // numerator/denominator getter/setter
43  long int numerator( Rational r, long int n );
44  long int denominator( Rational r );
45  long int denominator( Rational r, long int d );
46
47  int ?==( Rational l, Rational r );              // comparison
48  int ?!==( Rational l, Rational r );
49  int ?<?( Rational l, Rational r );
50  int ?<=( Rational l, Rational r );
51  int ?>?( Rational l, Rational r );
52  int ?>=( Rational l, Rational r );
53
54  Rational -( Rational r );                       // arithmetic
55  Rational ?+?( Rational l, Rational r );
56  Rational ?-?( Rational l, Rational r );
57  Rational ?*?( Rational l, Rational r );

```

C	Cv
<pre> #include &lt;gmp.h&gt; 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 &lt;= 40; i += 1 ) {         mpz_mul_ui( fact, fact, i );         gmp_printf( "%d %Zd\n", i, fact );     } } </pre>	<pre> #include &lt;gmp.h&gt; int main( void ) {     sout   "Factorial Numbers";     Int fact = 1;      sout   0   fact;     for ( i; 40 ) {         fact *= i;         sout   i   fact;     } } </pre>
<p>Factorial Numbers</p> <pre> 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 16: Multi-precision Factorials

```

1   Rational ?/?( Rational l, Rational r );
2
3   double widen( Rational r );           // conversion
4   Rational narrow( double f, long int md );
5
6   forall( dtype istype | istream( istype ) ) istype * ?/?( istype *, Rational * ); // I/O
7   forall( dtype ostype | ostream( ostype ) ) ostype * ?/?( ostype *, Rational );

```

## 8 References

- [1] Ada16. *Ada Reference Manual ISO/IEC 8652:2012(E) with COR.1:2016*. AXE Consultants, Madison WI, USA, 3rd with technical corrigendum 1 for ada 2012 edition, 2016. <https://docs.adacore.com/live/wave/arm12/pdf/arm12/arm-12.pdf>. 2
- [2] Richard C. Bilson. Implementing overloading and polymorphism in C $\forall$ . Master's thesis, School of Computer Science, University of Waterloo, 2003. <http://plg.uwaterloo.ca/theses/BilsonThesis.pdf>. 2, 45
- [3] Walter Bright and Andrei Alexandrescu. *D Programming Language*. Digital Mars, Vienna Virginia, U.S.A., 2016. <http://dlang.org/spec/spec.html>. 2
- [4] P. A. Buhr. A case for teaching multi-exit loops to beginning programmers. *SIGPLAN Not.*, 20(11):14–22, November 1985. 13
- [5] P. A. Buhr, David Till, and C. R. Zarnke. Assignment as the sole means of updating objects. *Softw. Pract. Exper.*, 24(9):835–870, September 1994. 2, 40
- [6] Cobol14. *Programming Languages – Cobol ISO/IEC 1989:2014*. International Standard Organization, Geneva, Switzerland, 2nd edition, 2014. <https://www.iso.org/standard/51416.html>. 2
- [7] G. V. Cormack and A. K. Wright. Polymorphism in the compiled language ForceOne. In *Proceedings of the 20th Hawaii International Conference on Systems Sciences*, pages 284–292, January 1987. 2
- [8] G. V. Cormack and A. K. Wright. Type-dependent parameter inference. *SIGPLAN Not.*, 25(6):127–136, June 1990. Proceedings of the ACM Sigplan'90 Conference on Programming Language Design and Implementation June 20–22, 1990, White Plains, New York, U.S.A. 2, 35
- [9] Glen Jeffrey Ditchfield. *Contextual Polymorphism*. PhD thesis, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 1992. <http://plg.uwaterloo.ca/theses/DitchfieldThesis.pdf>. 2
- [10] Tom Duff. Duff's device. <http://www.lysator.liu.se/c/duffs-device.html>, November 1983. 8, 9
- [11] Dominic Duggan, Gordon V. Cormack, and John Ophel. Kindred type inference for parametric overloading. *Acta Infomatica*, 33(1):21–68, 1996. 2
- [12] Margaret A. Ellis and Bjarne Stroustrup. *The Annotated C++ Reference Manual*. Addison-Wesley, Boston, 1st edition, 1990. 2
- [13] Rodolfo Gabriel Esteves. C $\forall$ , a study in evolutionary design in programming languages. Master's thesis, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 2004. <http://plg.uwaterloo.ca/theses/EstevesThesis.pdf>. 2
- [14] Fortran08. *Programming Languages – Fortran Part 1:Base Language ISO/IEC 1539-1:2010*. International Standard Organization, Geneva, Switzerland, 3rd edition, 2010. <https://www.iso.org/standard/50459.html>. 2
- [15] Nissim Francez. Another advantage of keyword notation for parameter communication with subprograms. *Communications of the ACM*, 20(8):604–605, August 1977. 34
- [16] John Galletly. *OCCAM 2: Including OCCAM 2.1*. UCL (University College London) Press, London, 2nd edition, 1996. 32, 48
- [17] *GNU Multiple Precision Arithmetic Library*. GNU, 2016. <https://gmplib.org>. 89
- [18] James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. *Java Language Specification*, Java SE 8 edition, 2015. 2

- 1 [19] Robert Griesemer, Rob Pike, and Ken Thompson. *Go Programming Language*. Google, Mountain View, CA,  
2 USA, 2009. <http://golang.org/ref/spec>. 2, 69
- 3 [20] W. T. Hardgrave. Positional versus keyword parameter communication in programming languages. *SIGPLAN*  
4 *Not.*, 11(5):52–58, May 1976. 34
- 5 [21] International Standard Organization, Geneva, Switzerland. *C Programming Language ISO/IEC 9889:2011-12*,  
6 3rd edition, 2012. <https://www.iso.org/standard/57853.html>. 2, 19, 21, 44, 70, 73
- 7 [22] International Standard Organization, Geneva, Switzerland. *C++ Programming Language ISO/IEC 14882:2014*,  
8 4th edition, 2014. <https://www.iso.org/standard/64029.html>. 2, 70
- 9 [23] Kathleen Jensen and Niklaus Wirth. *Pascal User Manual and Report, ISO Pascal Standard*. Springer-Verlag,  
10 4th edition, 1991. Revised by Andrew B. Mickel and James F. Miner. 15
- 11 [24] C. H. Lindsey and S. G. van der Meulen. *Informal Introduction to ALGOL 68*. North-Holland, London, 1977.  
12 20
- 13 [25] Barbara Liskov, Russell Atkinson, Toby Bloom, Eliot Moss, J. Craig Schaffert, Robert Scheifler, and Alan  
14 Snyder. *CLU Reference Manual*, volume 114 of *Lecture Notes in Computer Science*. Springer, New York,  
15 1981. 32, 47, 48
- 16 [26] Python. *Python Language Reference, Release 3.7.2*. Python Software Foundation, [https://docs.python.org/3/](https://docs.python.org/3/reference/index.html)  
17 [reference/index.html](https://docs.python.org/3/reference/index.html), 2018. 1
- 18 [27] Dennis M. Ritchie. The development of the C language. *SIGPLAN Not.*, 28(3):201–208, March 1993. 18
- 19 [28] *Rust Programming Language*, 2015. <https://doc.rust-lang.org/reference.html>. 2
- 20 [29] *Scala Language Specification, Version 2.11*. École Polytechnique Fédérale de Lausanne, 2016. [http://www.](http://www.scala-lang.org/files/archive/spec/2.11)  
21 [scala-lang.org/files/archive/spec/2.11](http://www.scala-lang.org/files/archive/spec/2.11). 1
- 22 [30] Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley Series in Computer Science. Addison-  
23 Wesley, Boston, 1st edition, 1986. 1
- 24 [31] David W. Till. Tuples in imperative programming languages. Master’s thesis, Department of Computer Science,  
25 University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1, 1989. 2, 40, 41, 43
- 26 [32] TIOBE Index. [http://www.tiobe.com/tiobe\\_index](http://www.tiobe.com/tiobe_index). 1
- 27 [33] Ben Werther and Damian Conway. A modest proposal: C++ resyntaxed. *SIGPLAN Not.*, 31(11):74–82,  
28 November 1996. 2

## Index

- 1 Italic page numbers give the location of the main entry for  
2 the referenced term. Plain page numbers denote uses of  
3 the indexed term. Entries for grammar non-terminals are  
4 italicized. A typewriter font is used for grammar terminals  
5 and program identifiers.
- 6 **!?**, [64](#)  
7 **&&**, [7](#)  
8 **\***, [26](#)  
9 **\*=**, [26](#)  
10 **\*?**, [64](#)  
11 **+**, [25](#)  
12 **++?**, [64](#)  
13 **+=**, [11](#), [25](#)  
14 **+?**, [64](#)  
15 **--?**, [64](#)  
16 **--colors**, [4](#)  
17 **--deterministic-out**, [4](#)  
18 **--gdb**, [4](#)  
19 **--help**, [4](#)  
20 **--invariant**, [4](#)  
21 **--libcfa**, [4](#)  
22 **--linemarks**, [4](#)  
23 **--no-linemarks**, [4](#)  
24 **--no-main**, [4](#)  
25 **--no-prelude**, [4](#)  
26 **--prelude-dir**, [5](#)  
27 **--print**, [4](#)  
28 **--prototypes**, [4](#)  
29 **--statistics**, [5](#)  
30 **--tree**, [5](#)  
31 **--**, [11](#)  
32 **-?**, [64](#)  
33 **-CFA**, [3](#)  
34 **-L**, [4](#)  
35 **-N**, [4](#)  
36 **-P**, [4](#)  
37 **-S**, [5](#)  
38 **-XCFA**, [3](#), [4](#)  
39 **-\~**, [11](#)  
40 **-\~=**, [11](#)  
41 **-c**, [4](#)  
42 **-d**, [4](#)  
43 **-debug**, [3](#)  
44 **-fgnu89-inline**, [3](#)  
45 **-g**, [4](#)  
46 **-h**, [4](#)  
47 **-help**, [4](#)  
48 **-i**, [4](#)  
49 **-l**, [4](#)  
50 **-lgmp**, [92](#)  
51 **-m**, [4](#)
- n**, [4](#)  
**-nodebug**, [3](#)  
**-nohelp**, [4](#)  
**-noquiet**, [4](#)  
**-p**, [4](#)  
**-quiet**, [4](#)  
**-std=gnu11**, [3](#)  
**-t**, [5](#)  
**-~**, [11](#)  
**-~=**, [11](#)  
**?!=?**, [64](#)  
**?%=?**, [64](#)  
**?%??**, [64](#)  
**?&=?**, [64](#)  
**?&??**, [64](#)  
**?()**, [64](#)  
**?\*=?**, [64](#)  
**?\*??**, [64](#)  
**?++**, [64](#)  
**?+=?**, [64](#)  
**?+?**, [64](#)  
**?--**, [64](#)  
**?-=?**, [64](#)  
**?-?**, [64](#)  
**?/=?**, [64](#)  
**?/?**, [64](#)  
**?<=?**, [64](#)  
**?<<?**, [64](#)  
**?<=?**, [64](#)  
**?<?**, [64](#)  
**?==?**, [64](#)  
**?=?**, [64](#)  
**?>=?**, [64](#)  
**?>>=?**, [64](#)  
**?>?**, [64](#)  
**?[?]**, [64](#)  
**?|=?**, [64](#)  
**?\?**, [64](#)  
**?^=?**, [64](#)  
**?^?**, [64](#)  
**?|=?**, [64](#)  
**?|?**, [64](#)  
**\~**, [11](#)  
**\~=**, [11](#)  
**\_\_CFA\_MAJOR\_\_**, [4](#)  
**\_\_CFA\_MINOR\_\_**, [4](#)  
**\_\_CFA\_PATCH\_\_**, [4](#)  
**\_\_CFA\_\_**, [4](#)  
**\_\_CFORALL\_\_**, [4](#)  
**\_\_cforall**, [4](#)

- 1 ~, 11
- 2 ~=, 11
- 3 ~?, 64
- 4 0, 11
- 5 1, 11
- 6 abs, 77
- 7 acos, 80
- 8 acosh, 81
- 9 Ada, 2, 6, 36
- 10 address, 19
  - 11 duality, 20
- 12 address duality, 22
- 13 addressing errors, 21
- 14 adelete, 75
- 15 aggregate, 14
- 16 aggregation, 38, 39
- 17 Algol68, 20
- 18 aliasing, 38
- 19 aligned\_alloc, 74
- 20 alloc, 75
- 21 alloc\_set, 75
- 22 anew, 75
- 23 arguments
  - 24 default, 34
  - 25 named, 34
- 26 array, 22
  - 27 reference, 22
- 28 ASCII, 52
  - 29 extended, 52
- 30 ascodegen, 4
- 31 asin, 80
- 32 asinh, 81
- 33 assert.h, 73
- 34 assignment, 22
  - 35 cascade, 48
  - 36 pointer, 20
- 37 ast, 4
- 38 asterr, 4
- 39 atan, 80
- 40 atan2, 80
- 41 atanh, 81
- 42 ato, 76
- 43 backward-compatible, 3
- 44 basetypeof, 73
- 45 bbox, 5
- 46 bcodegen, 5
- 47 Bilson, Richard, 2
- 48 bin, 54
- 49 bit field, 18
- 50 brand, 66
- 51 break, 7, 13
- 52 break
  - 53 labelled, 13
- 54 bresolver, 5
- 55 bsearch, 76
- 56 bsearchl, 76
- 57 bsearchu, 76
- 58 C linkage, 3
- 59 C++, 1–3, 22, 36, 67, 72, 73
- 60 C-for-all, 1
- 61 C11, 2
- 62 calling convention, 2
- 63 calloc, 74
- 64 cancellation
  - 65 pointer/reference, 21
- 66 cascade assignment, 48
- 67 case, 7, 10
- 68 catch, 73
- 69 catchResume, 73
- 70 cbrt, 79
- 71 ceil, 82
- 72 cfa, 3, 4
- 73 CV, 1
- 74 char \*, 24
- 75 choose, 9, 10, 73
- 76 clamp, 77
- 77 clear, 50, 51
- 78 close, 51
- 79 closing coercion, 47
- 80 cmemalign, 74
- 81 Cobol, 2
- 82 coercion, 21
- 83 comma expression, 38
- 84 compilation
  - 85 cfa, 3
- 86 compilation option
  - 87 –CFA, 3
  - 88 –XCFA, 3
  - 89 –debug, 3
  - 90 –fgnu89–inline, 3
  - 91 –help, 4
  - 92 –nodebug, 3
  - 93 –nohelp, 4
  - 94 –noquiet, 4
  - 95 –quiet, 4
  - 96 –std=gnu11, 3
- 97 compiler, 3
- 98 complex, 6
- 99 complex constant, 6
- 100 complex.h, 73
- 101 const hell, 23
- 102 constant
  - 103 underscore, 5
- 104 continue, 13
- 105 continue



- 1       labelled, 13
- 2       copysign, 83
- 3       **coroutine**, 73
- 4       coroutines, 66
- 5       cos, 79
- 6       cosh, 80
- 7       ctordtor, 5
- 8       ctype.h, 73
- 9       D, 2, 67, 69
- 10      dangling else, 14
- 11      Dave Till, 2
- 12      deadlock, 59
- 13      declstats, 4
- 14      **default**, 10
- 15      default arguments, 34
- 16      delete, 75
- 17      describe not prescribe, 1
- 18      **disable**, 73
- 19      Ditchfield, Glen, 2
- 20      div, 38
- 21      **do**, 6, 11, 13, 14
- 22      Doug Lea, 32
- 23      down-to exclusive range, 11
- 24      down-to inclusive range, 11
- 25      duality, 20
- 26      Duff's device, 8–10
- 27      **enable**, 73
- 28      eng, 55
- 29      enum, 30
- 30      enumeration, 30
- 31      enumeration constant, 30
- 32      eof, 51
- 33      equality operator, 25
- 34      erf, 81
- 35      erfc, 81
- 36      errno.h, 73
- 37      Esteves, Rodolfo, 2
- 38      **exception**, 73
- 39      excl, 58
- 40      excpdecl, 4
- 41      exp, 78
- 42      exp2, 78
- 43      expm1, 78
- 44      exponentiation
  - 45       floating, 6
  - 46       logarithm, 6
  - 47       operator, 6
  - 48       unsigned integral, 6
- 49      expralt, 5
- 50      expranly, 5
- 51      extended **else**, 14
- 52      eye candy, 13
- 53      fail, 50, 51
- 54      **fallthrough**, 9, 10, 73
- 55      **fallthru**, 9, 73
- 56      fdim, 78
- 57      fenv.h, 73
- 58      Fibonacci, 66
- 59      **finally**, 73
- 60      **fixup**, 73
- 61      flattening coercion, 47
- 62      float.h, 73
- 63      floor, 82
- 64      flush, 51
- 65      fma, 78
- 66      fmod, 77
- 67      **for**, 6, 11, 13, 14
- 68      **forall**, 73
- 69      Fortran, 2
- 70      frexp, 83
- 71      fstream, 1, 4
- 72      fstream.hfa, 49
- 73      garbage-collection, 1
  - 74       regional, 1
- 75      gcc, 3, 4, 65
- 76      **generator**, 73
- 77      getline, 57
- 78      getpid, 87
- 79      Glen Ditchfield, 2
- 80      gmp, 89, 93
- 81      gmp.h, 73, 93
- 82      gnu, 3
- 83      Go, 2, 66–69
- 84      **goto**, 13
- 85      **goto**
  - 86       restricted, 13
- 87      hex, 55
- 88      hypot, 79
- 89      I/O
  - 90       common case, 49
  - 91       separator, 51
- 92      **if**, 6, 13, 14
- 93      ifstream
  - 94       ?{}, 51
  - 95       ^?{}, 51
  - 96       clear, 51
  - 97       close, 51
  - 98       eof, 51
  - 99       fail, 51
  - 100       open, 51
  - 101       read, 51
  - 102       unget, 51
- 103     ignore, 58
- 104     ilogb, 79

- 1 immutable, 19
- 2 implicit output separation, 53
- 3 implicit referencing, 22
- 4 implicit separator, 51
- 5 incl, 58
- 6 inheritance
  - 7 enumeration, 31
- 8 initialization, 22
- 9 input, 49
- 10 instgen, 5
- 11 instruction decoding, 20
- 12 Int, 89
- 13 **int128**, 73
- 14 interoperability, 2
- 15 interposition, 5
- 16 inttypes.h, 73
- 17 iostream, 1
- 18 isacquire, 58
- 19 iso646.h, 73
- 20 istream
  - 21 isacquire, 58
- 22 Java, 2, 6, 22, 37, 69
- 23 K-W C, 2
- 24 keyword, 73
- 25 labelled
  - 26 **break**, 13
  - 27 **continue**, 13
- 28 lambda, 38
- 29 ldexp, 83
- 30 Lea, Doug, 32
- 31 left, 56
- 32 level number, 4
- 33 lexical list, 45
- 34 lgamma, 81
- 35 library
  - 36 stream, 49
- 37 library support, 3
- 38 limits.h, 73
- 39 locale, 59
  - 40 stream, 59
- 41 locale.h, 73
- 42 log, 78
- 43 log10, 79
- 44 log1p, 79
- 45 log2, 79
- 46 logarithm, 6
- 47 logb, 79
- 48 lrint, 82
- 49 lround, 83
- 50 lvalue, 21, 23, 47, 48
- 51 malloc, 74
- 52 malloc.h, 73
- 53 mangling
  - 54 name, 3, 73
- 55 manipulator, 53, 54, 57
  - 56 bin, 54
  - 57 eng, 55
  - 58 excl, 58
  - 59 getline, 57
  - 60 hex, 55
  - 61 ignore, 58
  - 62 incl, 58
  - 63 left, 56
  - 64 nl, 54
  - 65 nlOff, 54
  - 66 nlOn, 54
  - 67 nobase, 55
  - 68 nodp, 55
  - 69 nosepl, 53
  - 70 oct, 54
  - 71 pad0, 56
  - 72 quoted, 57, 58
  - 73 sci, 55
  - 74 sep, 53
  - 75 sepGet, 53
  - 76 sepGetTuple, 53
  - 77 sepOff, 53
  - 78 sepOn, 53
  - 79 sepSet, 53
  - 80 sepSetTuple, 53
  - 81 sepTupleVal, 53
  - 82 sepVal, 53
  - 83 sign, 55
  - 84 skip, 57
  - 85 unit, 55
  - 86 upcase, 55
  - 87 wd, 55
  - 88 wdi, 57
  - 89 ws, 56
- 90 manipulators, 54
- 91 mass, 41
- 92 math.h, 73
- 93 max, 77
- 94 memalign, 74
- 95 memcpy, 75
- 96 memory fault, 19
- 97 memset, 75
- 98 Michael Tiemann, 32
- 99 min, 77
- 100 modf, 38, 83
- 101 **monitor**, 73
- 102 multi-level exit, 13
- 103 multi-precision, 89
- 104 multiple assignment, 47
- 105 multiple, 41

- 1 multiple derivation, 18
- 2 **mutex**, 73
- 3 name clash, 3
- 4 name hiding, 36
- 5 name mangling, 3
- 6 named arguments, 34
- 7 named return values, 32
- 8 naming, 38
- 9 nan, 78
- 10 nearbyint, 82
- 11 nested control-structure, 13
- 12 nested locking, 59
- 13 new, 75
- 14 newline separation, 54
- 15 nextafter, 83
- 16 nexttoward, 83
- 17 nl, 54
- 18 nlOff, 54
- 19 nlOn, 54
- 20 nobase, 55
- 21 nodp, 55
- 22 nose, 53
- 23 null-pointer constant, 19
- 24 object, 19
- 25 object-oriented, 1
- 26 oct, 54
- 27 ofstream
  - 28 `{}`, 51
  - 29 `^{}` , 51
  - 30 clear, 51
  - 31 close, 51
  - 32 fail, 51
  - 33 flush, 51
  - 34 open, 51
  - 35 write, 51
- 36 **one\_t**, 73
- 37 onion, 18
- 38 open, 51
- 39 opening coercion, 47
- 40 operator
  - 41 exponentiation, 6
- 42 **or**, 73
- 43 ostream
  - 44 **mutex**, 58
  - 45 osacquire, 58
- 46 out parameter, 38
- 47 output, 49
- 48 overload, 2, 3
- 49 pad0, 56
- 50 parametric-polymorphic, 2
- 51 parse, 4
- 52 particle and wave, 23
- 53 plan-9, 31
- 54 pointer, 18–20
  - 55 assignment, 20
  - 56 cancellation, 21
- 57 pointer type, 19
- 58 posix\_memalign, 74
- 59 pow, 78
- 60 pow(x,y), 6
- 61 power of a name, 3
- 62 prelude, 3
- 63 preprocessor, 30
- 64 preprocessor variables
  - 65 `__CFA_MINOR__`, 4
  - 66 `__CFA_PATCH__`, 4
  - 67 `__CFA__`, 4
  - 68 `__CFA__`, 4
  - 69 `__CFORALL__`, 4
  - 70 `__cforall`, 4
- 71 prescribing, 18
- 72 pretty, 4
- 73 prng, 87
- 74 productivity, 1, 18
- 75 pseudo random-number generators, 87
- 76 Python, 1, 49, 69
- 77 qsort, 76, 77
- 78 quasi-keyword, 73
- 79 quoted, 57, 58
- 80 raii, 58
- 81 random, 77
- 82 Rational, 92
- 83 read, 51
- 84 realloc, 74
- 85 recursive lock, 58
- 86 reference, 20
- 87 reference type, 19, 20
- 88 referenced type, 19
- 89 regional garbage-collection, 1
- 90 relational operator, 25
- 91 remainder, 78
- 92 remquo, 78
- 93 **report**, 73
- 94 return list, 32
- 95 Richard Bilson, 2
- 96 rint, 82
- 97 Rodolfo Esteves, 2
- 98 round, 83
- 99 routine object, 23
- 100 rproto, 4
- 101 rsteps, 4
- 102 Rust, 2, 67, 69
- 103 rvalue, 21, 23
- 104 safety, 1, 18

- 1 scala, 1
- 2 scalbln, 83
- 3 scalbn, 83
- 4 scanset, 56
- 5 sci, 55
- 6 seed, 87
- 7 sentinel value, 19
- 8 sep, 53
- 9 sepGet, 53
- 10 sepGetTuple, 53
- 11 sepOff, 53
- 12 sepOn, 53
- 13 sepSet, 53
- 14 sepSetTuple, 53
- 15 sepTupleVal, 53
- 16 sepVal, 53
- 17 setjmp.h, 73
- 18 shell pipe-operator, 49
- 19 sign, 55
- 20 signal.h, 73
- 21 sin, 79
- 22 sinh, 80
- 23 skip, 57
- 24 sound, 19
- 25 sout, 1
- 26 sqrt, 79
- 27 srandom, 77
- 28 static multi-level exit, 13
- 29 stdalign.h, 73
- 30 stdarg.h, 73
- 31 stdatomic.h, 73
- 32 stdbool.h, 73
- 33 stddef.h, 73
- 34 stdint.h, 73
- 35 stdio.h, 1, 4, 73
- 36 stdlib.h, 73
- 37 stdlib.hfa, 87
- 38 stdnoreturn.h, 73
- 39 sticky properties, 74
- 40 stream
  - 41 locale, 59
- 42 stream library, 49
- 43 string, 24
- 44 string.h, 73
- 45 string.hfa, 25
- 46 structuring coercion, 47
- 47 **suspend**, 73
- 48 swap, 77
- 49 **switch**, 7, 10, 13
- 50 symevt, 4
- 51 tan, 80
- 52 tanh, 80
- 53 tentative definitions, 71
- 54 tgamma, 81
- 55 tgmth.h, 73
- 56 threads.h, 73
- 57 **throw**, 73
- 58 **throwResume**, 73
- 59 Tiemann, Michael, 32
- 60 Till, Dave, 2
- 61 time.h, 73
- 62 **timeout**, 73
- 63 **trait**, 73
- 64 transpiler, 3, 4
- 65 transpiler option
  - 66 --colors, 4
  - 67 --deterministic-out, 4
  - 68 --gdb, 4
  - 69 --help, 4
  - 70 --invariant, 4
  - 71 --libcfa, 4
  - 72 --linemarks, 4
  - 73 --no-linemarks, 4
  - 74 --no-main, 4
  - 75 --no-prelude, 4
  - 76 --print, 4
  - 77 -print
    - 78 ascodegen, 4
    - 79 asterr, 4
    - 80 ast, 4
    - 81 bbox, 5
    - 82 bcodegen, 5
    - 83 bresolver, 5
    - 84 ctordtor, 5
    - 85 declstats, 4
    - 86 excpdecl, 4
    - 87 expralt, 5
    - 88 expranly, 5
    - 89 instgen, 5
    - 90 parse, 4
    - 91 pretty, 4
    - 92 rproto, 4
    - 93 rsteps, 4
    - 94 symevt, 4
    - 95 tuple, 5
    - 96 valdecl, 5
    - 97 --prototypes, 4
    - 98 --statistics
      - 99 counters,heap,time,all,none, 5
    - 100 --tree, 5
    - 101 -L, 4
    - 102 -N, 4
    - 103 -P, 4
      - 104 ascodegen, 4
      - 105 asterr, 4
      - 106 ast, 4
      - 107 bbox, 5

- 1      bcodegen, 5
- 2      bresolver, 5
- 3      ctordtor, 5
- 4      declstats, 4
- 5      excpdecl, 4
- 6      expralt, 5
- 7      expranly, 5
- 8      instgen, 5
- 9      parse, 4
- 10     pretty, 4
- 11     rproto, 4
- 12     rsteps, 4
- 13     symevt, 4
- 14     tuple, 5
- 15     valdecl, 5
- 16     -S
- 17        counters,heap,time,all,none, 5
- 18     -XCFA, 4
- 19     -c, 4
- 20     -d, 4
- 21     -g, 4
- 22     -h, 4
- 23     -i, 4
- 24     -l, 4
- 25     -m, 4
- 26     -n, 4
- 27     -p, 4
- 28     -t, 5
- 29     trunc, 82
- 30     **try**, 73
- 31     tuple, 38, 49
- 32     tuple, 5
- 33     tuple assignment, 41
- 34     tuple expression, 39
- 35     tuple type, 39, 46
- 36     tuple variable, 39
- 37     tuple-returning functions, 39
- 38     type hoisting, 36
- 39     type nesting, 36
- 40     uchar.h, 73
- 41     undefined, 19, 33
- 42     underscore, 5
- 43     unget, 51
- 44     ungetc, 51
- 45     unistd.h, 73
- 46     unit, 55
- 47     unsound, 19
- 48     up-to exclusive range, 11
- 49     up-to inclusive range, 11
- 50     upcase, 55
- 51     valdecl, 5
- 52     version number, 4
- 53     **virtual**, 73
- 54     **waitfor**, 14, 73
- 55     warning, 24
- 56     wchar.h, 73
- 57     wctype.h, 73
- 58     wd, 55
- 59     wdi, 57
- 60     What goes around, comes around., 2
- 61     **when**, 73
- 62     **while**, 6, 11, 13, 14
- 63     **with**, 73
- 64     write, 51, 60
- 65     ws, 56
- 66     **zero\_t**, 73