

# Beej's Guide to C Programming

Library Reference

Brian "Beej Jorgensen" Hall

v0.9.11, Copyright © April 16, 2025

# Contents

|          |                                                  |          |
|----------|--------------------------------------------------|----------|
| <b>1</b> | <b>Foreword</b>                                  | <b>1</b> |
| 1.1      | Audience . . . . .                               | 1        |
| 1.2      | How to Read This Book . . . . .                  | 1        |
| 1.3      | Platform and Compiler . . . . .                  | 2        |
| 1.4      | Official Homepage . . . . .                      | 2        |
| 1.5      | Email Policy . . . . .                           | 2        |
| 1.6      | Mirroring . . . . .                              | 2        |
| 1.7      | Note for Translators . . . . .                   | 3        |
| 1.8      | Copyright and Distribution . . . . .             | 3        |
| 1.9      | Dedication . . . . .                             | 3        |
| <br>     |                                                  |          |
| <b>2</b> | <b>The C Language</b>                            | <b>4</b> |
| 2.1      | Background . . . . .                             | 4        |
| 2.1.1    | Comments . . . . .                               | 4        |
| 2.1.2    | Separators . . . . .                             | 4        |
| 2.1.3    | Expressions . . . . .                            | 4        |
| 2.1.4    | Statements . . . . .                             | 4        |
| 2.1.5    | Booleans . . . . .                               | 4        |
| 2.1.6    | Blocks . . . . .                                 | 4        |
| 2.1.7    | Code Examples . . . . .                          | 5        |
| 2.2      | Operators . . . . .                              | 5        |
| 2.2.1    | Arithmetic Operators . . . . .                   | 5        |
| 2.2.2    | Pre- and Post-Increment and -Decrement . . . . . | 5        |
| 2.2.3    | Comparison Operators . . . . .                   | 5        |
| 2.2.4    | Pointer Operators . . . . .                      | 5        |
| 2.2.5    | Structure and Union Operators . . . . .          | 6        |
| 2.2.6    | Array Operators . . . . .                        | 6        |
| 2.2.7    | Bitwise Operators . . . . .                      | 6        |
| 2.2.8    | Assignment Operators . . . . .                   | 6        |
| 2.2.9    | The <code>sizeof</code> Operator . . . . .       | 6        |
| 2.2.10   | Type Casts . . . . .                             | 7        |
| 2.2.11   | <code>_Alignof</code> Operator . . . . .         | 7        |
| 2.2.12   | Comma Operator . . . . .                         | 7        |
| 2.3      | Type Specifiers . . . . .                        | 8        |
| 2.4      | Constant Types . . . . .                         | 8        |
| 2.5      | Composite Types . . . . .                        | 9        |
| 2.5.1    | <code>struct</code> Types . . . . .              | 9        |
| 2.5.2    | <code>union</code> Types . . . . .               | 9        |
| 2.5.3    | <code>enum</code> Types . . . . .                | 10       |
| 2.6      | Initializers . . . . .                           | 10       |
| 2.7      | Compound Literals . . . . .                      | 11       |
| 2.8      | Type Aliases . . . . .                           | 12       |
| 2.9      | Additional Type-Related Specifiers . . . . .     | 13       |

|          |                                                                |           |
|----------|----------------------------------------------------------------|-----------|
| 2.9.1    | Storage Class Specifiers . . . . .                             | 13        |
| 2.9.2    | Type Qualifiers . . . . .                                      | 13        |
| 2.9.3    | Function Specifiers . . . . .                                  | 14        |
| 2.9.4    | Alignment Specifier . . . . .                                  | 15        |
| 2.10     | if Statement . . . . .                                         | 15        |
| 2.11     | for Statement . . . . .                                        | 15        |
| 2.12     | while Statement . . . . .                                      | 16        |
| 2.13     | do-while Statement . . . . .                                   | 16        |
| 2.14     | switch Statement . . . . .                                     | 16        |
| 2.15     | break Statement . . . . .                                      | 17        |
| 2.16     | continue Statement . . . . .                                   | 17        |
| 2.17     | goto Statement . . . . .                                       | 18        |
| 2.18     | return Statement . . . . .                                     | 18        |
| 2.19     | _Static_assert Statement . . . . .                             | 19        |
| 2.20     | Functions . . . . .                                            | 19        |
| 2.20.1   | main() Function . . . . .                                      | 19        |
| 2.20.2   | Variadic Functions . . . . .                                   | 20        |
| <b>3</b> | <b>&lt;assert.h&gt; Runtime and Compile-time Diagnostics</b>   | <b>21</b> |
| 3.1      | Macros . . . . .                                               | 21        |
| 3.2      | assert() . . . . .                                             | 21        |
| 3.3      | static_assert() . . . . .                                      | 23        |
| <b>4</b> | <b>&lt;complex.h&gt; Complex Number Functionality</b>          | <b>25</b> |
| 4.1      | cacos(), cacosf(), cacosl() . . . . .                          | 27        |
| 4.2      | casin(), casinl(), casinf() . . . . .                          | 28        |
| 4.3      | catan(), catanf(), catanl() . . . . .                          | 29        |
| 4.4      | ccos(), ccosf(), ccosl() . . . . .                             | 30        |
| 4.5      | csin(), csinf(), csinl() . . . . .                             | 31        |
| 4.6      | ctan(), ctanf(), ctanl() . . . . .                             | 32        |
| 4.7      | cacosh(), cacoshf(), cacoshl() . . . . .                       | 33        |
| 4.8      | casinh(), casinhf(), casinhl() . . . . .                       | 34        |
| 4.9      | catanh(), catanhf(), catanhl() . . . . .                       | 35        |
| 4.10     | ccosh(), ccoshf(), ccoshl() . . . . .                          | 36        |
| 4.11     | csinh(), csinhf(), csinhl() . . . . .                          | 37        |
| 4.12     | ctanh(), ctanhf(), ctanhl() . . . . .                          | 38        |
| 4.13     | cexp(), cexpf(), cexpl() . . . . .                             | 39        |
| 4.14     | clog(), clogf(), clogl() . . . . .                             | 40        |
| 4.15     | cabs(), cabsf(), cabsl() . . . . .                             | 40        |
| 4.16     | cpow(), cpowf(), cpowl() . . . . .                             | 41        |
| 4.17     | csqrt(), csqrtf(), csqrtl() . . . . .                          | 42        |
| 4.18     | carg(), cargf(), cargl() . . . . .                             | 43        |
| 4.19     | cimag(), cimagf(), cimagl() . . . . .                          | 44        |
| 4.20     | CMPLX(), CMPLXF(), CMPLXL() . . . . .                          | 45        |
| 4.21     | conj(), conjf(), conjl() . . . . .                             | 47        |
| 4.22     | cproj(), cprojf(), cprojl() . . . . .                          | 48        |
| 4.23     | creal(), crealf(), creall() . . . . .                          | 49        |
| <b>5</b> | <b>&lt;ctype.h&gt; Character Classification and Conversion</b> | <b>51</b> |
| 5.1      | isalnum() . . . . .                                            | 52        |
| 5.2      | isalpha() . . . . .                                            | 52        |
| 5.3      | isblank() . . . . .                                            | 54        |
| 5.4      | iscntrl() . . . . .                                            | 54        |
| 5.5      | isdigit() . . . . .                                            | 55        |

|           |                                                                              |            |
|-----------|------------------------------------------------------------------------------|------------|
| 5.6       | <code>isgraph()</code>                                                       | 56         |
| 5.7       | <code>islower()</code>                                                       | 57         |
| 5.8       | <code>isprint()</code>                                                       | 58         |
| 5.9       | <code>ispunct()</code>                                                       | 59         |
| 5.10      | <code>isspace()</code>                                                       | 60         |
| 5.11      | <code>isupper()</code>                                                       | 61         |
| 5.12      | <code>isxdigit()</code>                                                      | 62         |
| 5.13      | <code>tolower()</code>                                                       | 62         |
| 5.14      | <code>toupper()</code>                                                       | 63         |
| <b>6</b>  | <b><code>&lt;errno.h&gt;</code> Error Information</b>                        | <b>65</b>  |
| 6.1       | <code>errno</code>                                                           | 65         |
| <b>7</b>  | <b><code>&lt;fenv.h&gt;</code> Floating Point Exceptions and Environment</b> | <b>69</b>  |
| 7.1       | Types and Macros                                                             | 69         |
| 7.2       | Pragmas                                                                      | 70         |
| 7.3       | <code>feclearexcept()</code>                                                 | 70         |
| 7.4       | <code>fegetexceptflag()</code> <code>fesetexceptflag()</code>                | 71         |
| 7.5       | <code>feraiseexcept()</code>                                                 | 72         |
| 7.6       | <code>fetestexcept()</code>                                                  | 73         |
| 7.7       | <code>fegetround()</code> <code>fesetround()</code>                          | 74         |
| 7.8       | <code>fegetenv()</code> <code>fesetenv()</code>                              | 76         |
| 7.9       | <code>feholdexcept()</code>                                                  | 77         |
| 7.10      | <code>feupdateenv()</code>                                                   | 78         |
| <b>8</b>  | <b><code>&lt;float.h&gt;</code> Floating Point Limits</b>                    | <b>81</b>  |
| 8.1       | Background                                                                   | 82         |
| 8.2       | <code>FLT_ROUNDS</code> Details                                              | 83         |
| 8.3       | <code>FLT_EVAL_METHOD</code> Details                                         | 83         |
| 8.4       | Subnormal Numbers                                                            | 83         |
| 8.5       | How Many Decimal Places Can I Use?                                           | 84         |
| 8.6       | Comprehensive Example                                                        | 85         |
| <b>9</b>  | <b><code>&lt;inttypes.h&gt;</code> More Integer Conversions</b>              | <b>88</b>  |
| 9.1       | Macros                                                                       | 88         |
| 9.2       | <code>imaxabs()</code>                                                       | 89         |
| 9.3       | <code>imaxdiv()</code>                                                       | 90         |
| 9.4       | <code>strtoimax()</code> <code>strtoumax()</code>                            | 91         |
| 9.5       | <code>wcstoimax()</code> <code>wcstoumax()</code>                            | 92         |
| <b>10</b> | <b><code>&lt;iso646.h&gt;</code> Alternative Operator Spellings</b>          | <b>95</b>  |
| <b>11</b> | <b><code>&lt;limits.h&gt;</code> Numeric Limits</b>                          | <b>97</b>  |
| 11.1      | <code>CHAR_MIN</code> and <code>CHAR_MAX</code>                              | 97         |
| 11.2      | Choosing the Correct Type                                                    | 98         |
| 11.3      | Whither Two's Complement?                                                    | 98         |
| 11.4      | Demo Program                                                                 | 98         |
| <b>12</b> | <b><code>&lt;locale.h&gt;</code> locale handling</b>                         | <b>100</b> |
| 12.1      | <code>setlocale()</code>                                                     | 100        |
| 12.2      | <code>localeconv()</code>                                                    | 102        |
| <b>13</b> | <b><code>&lt;math.h&gt;</code> Mathematics</b>                               | <b>106</b> |
| 13.1      | Math Function Idioms                                                         | 107        |
| 13.2      | Math Types                                                                   | 108        |

|       |                                                                                                                                                              |     |
|-------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 13.3  | Math Macros . . . . .                                                                                                                                        | 108 |
| 13.4  | Math Errors . . . . .                                                                                                                                        | 108 |
| 13.5  | Math Pragmas . . . . .                                                                                                                                       | 109 |
| 13.6  | <code>fpclassify()</code> . . . . .                                                                                                                          | 110 |
| 13.7  | <code>isfinite()</code> , <code>isinf()</code> , <code>isnan()</code> , <code>isnormal()</code> . . . . .                                                    | 111 |
| 13.8  | <code>signbit()</code> . . . . .                                                                                                                             | 112 |
| 13.9  | <code>acos()</code> , <code>acosf()</code> , <code>acosl()</code> . . . . .                                                                                  | 113 |
| 13.10 | <code>asin()</code> , <code>asinf()</code> , <code>asinl()</code> . . . . .                                                                                  | 114 |
| 13.11 | <code>atan()</code> , <code>atanf()</code> , <code>atanl()</code> , <code>atan2()</code> , <code>atan2f()</code> , <code>atan2l()</code> . . . . .           | 115 |
| 13.12 | <code>cos()</code> , <code>cosf()</code> , <code>cosl()</code> . . . . .                                                                                     | 116 |
| 13.13 | <code>sin()</code> , <code>sinf()</code> , <code>sinl()</code> . . . . .                                                                                     | 117 |
| 13.14 | <code>tan()</code> , <code>tanf()</code> , <code>tanl()</code> . . . . .                                                                                     | 118 |
| 13.15 | <code>acosh()</code> , <code>acoshf()</code> , <code>acoshl()</code> . . . . .                                                                               | 119 |
| 13.16 | <code>asinh()</code> , <code>asinhf()</code> , <code>asinhf()</code> . . . . .                                                                               | 120 |
| 13.17 | <code>atanh()</code> , <code>atanhf()</code> , <code>atanhl()</code> . . . . .                                                                               | 121 |
| 13.18 | <code>cosh()</code> , <code>coshf()</code> , <code>coshl()</code> . . . . .                                                                                  | 121 |
| 13.19 | <code>sinh()</code> , <code>sinhf()</code> , <code>sinhl()</code> . . . . .                                                                                  | 122 |
| 13.20 | <code>tanh()</code> , <code>tanhf()</code> , <code>tanhf()</code> . . . . .                                                                                  | 123 |
| 13.21 | <code>exp()</code> , <code>expf()</code> , <code>expl()</code> . . . . .                                                                                     | 124 |
| 13.22 | <code>exp2()</code> , <code>exp2f()</code> , <code>exp2l()</code> . . . . .                                                                                  | 125 |
| 13.23 | <code>expm1()</code> , <code>expm1f()</code> , <code>expm1l()</code> . . . . .                                                                               | 125 |
| 13.24 | <code>frexp()</code> , <code>frexpf()</code> , <code>frexpl()</code> . . . . .                                                                               | 126 |
| 13.25 | <code>ilogb()</code> , <code>ilogbf()</code> , <code>ilogbl()</code> . . . . .                                                                               | 128 |
| 13.26 | <code>ldexp()</code> , <code>ldexpf()</code> , <code>ldexpl()</code> . . . . .                                                                               | 129 |
| 13.27 | <code>log()</code> , <code>logf()</code> , <code>logl()</code> . . . . .                                                                                     | 129 |
| 13.28 | <code>log10()</code> , <code>log10f()</code> , <code>log10l()</code> . . . . .                                                                               | 130 |
| 13.29 | <code>log1p()</code> , <code>log1pf()</code> , <code>log1pl()</code> . . . . .                                                                               | 131 |
| 13.30 | <code>log2()</code> , <code>log2f()</code> , <code>log2l()</code> . . . . .                                                                                  | 132 |
| 13.31 | <code>logb()</code> , <code>logbf()</code> , <code>logbl()</code> . . . . .                                                                                  | 133 |
| 13.32 | <code>modf()</code> , <code>modff()</code> , <code>modfl()</code> . . . . .                                                                                  | 134 |
| 13.33 | <code>scalbn()</code> , <code>scalbnf()</code> , <code>scalbnl()</code> <code>scalbln()</code> , <code>scalblnf()</code> , <code>scalblnl()</code> . . . . . | 135 |
| 13.34 | <code>cbrt()</code> , <code>cbrtf()</code> , <code>cbrtl()</code> . . . . .                                                                                  | 137 |
| 13.35 | <code>fabs()</code> , <code>fabsf()</code> , <code>fabsl()</code> . . . . .                                                                                  | 138 |
| 13.36 | <code>hypot()</code> , <code>hypotf()</code> , <code>hypotl()</code> . . . . .                                                                               | 139 |
| 13.37 | <code>pow()</code> , <code>powf()</code> , <code>powl()</code> . . . . .                                                                                     | 139 |
| 13.38 | <code>sqrt()</code> . . . . .                                                                                                                                | 140 |
| 13.39 | <code>erf()</code> , <code>erff()</code> , <code>erfl()</code> . . . . .                                                                                     | 141 |
| 13.40 | <code>erfc()</code> , <code>erfcf()</code> , <code>erfcl()</code> . . . . .                                                                                  | 142 |
| 13.41 | <code>lgamma()</code> , <code>lgammaf()</code> , <code>lgammal()</code> . . . . .                                                                            | 143 |
| 13.42 | <code>tgamma()</code> , <code>tgammaf()</code> , <code>tgammal()</code> . . . . .                                                                            | 144 |
| 13.43 | <code>ceil()</code> , <code>ceilf()</code> , <code>ceilf()</code> . . . . .                                                                                  | 145 |
| 13.44 | <code>floor()</code> , <code>floorf()</code> , <code>floorf()</code> . . . . .                                                                               | 146 |
| 13.45 | <code>nearbyint()</code> , <code>nearbyintf()</code> , <code>nearbyintl()</code> . . . . .                                                                   | 147 |
| 13.46 | <code>rint()</code> , <code>rintf()</code> , <code>rintl()</code> . . . . .                                                                                  | 148 |
| 13.47 | <code>lrint()</code> , <code>lrintf()</code> , <code>lrintl()</code> , <code>llrint()</code> , <code>llrintf()</code> , <code>llrintl()</code> . . . . .     | 149 |
| 13.48 | <code>round()</code> , <code>roundf()</code> , <code>roundf()</code> . . . . .                                                                               | 150 |
| 13.49 | <code>lround()</code> , <code>lroundf()</code> , <code>lroundf()</code> <code>llround()</code> , <code>llroundf()</code> , <code>llroundf()</code> . . . . . | 151 |
| 13.50 | <code>trunc()</code> , <code>truncf()</code> , <code>truncf()</code> . . . . .                                                                               | 152 |
| 13.51 | <code>fmod()</code> , <code>fmodf()</code> , <code>fmodf()</code> . . . . .                                                                                  | 153 |
| 13.52 | <code>remainder()</code> , <code>remainderf()</code> , <code>remainderf()</code> . . . . .                                                                   | 154 |
| 13.53 | <code>remquo()</code> , <code>remquoof()</code> , <code>remquoof()</code> . . . . .                                                                          | 155 |
| 13.54 | <code>copysign()</code> , <code>copysignf()</code> , <code>copysignf()</code> . . . . .                                                                      | 156 |
| 13.55 | <code>nan()</code> , <code>nanf()</code> , <code>nanf()</code> . . . . .                                                                                     | 157 |
| 13.56 | <code>nextafter()</code> , <code>nextafterf()</code> , <code>nextafterf()</code> . . . . .                                                                   | 159 |

|           |                                                                                                                                       |            |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------|------------|
| 13.57     | <code>nexttoward()</code> , <code>nexttowardf()</code> , <code>nexttowardl()</code>                                                   | 160        |
| 13.58     | <code>fdim()</code> , <code>fdimf()</code> , <code>fdiml()</code>                                                                     | 161        |
| 13.59     | <code>fmax()</code> , <code>fmaxf()</code> , <code>fmaxl()</code> , <code>fmin()</code> , <code>fminf()</code> , <code>fminl()</code> | 162        |
| 13.60     | <code>fma()</code> , <code>fmaf()</code> , <code>fmal()</code>                                                                        | 162        |
| 13.61     | <code>isgreater()</code> , <code>isgreaterequal()</code> , <code>isless()</code> , <code>islessequal()</code>                         | 163        |
| 13.62     | <code>islessgreater()</code>                                                                                                          | 164        |
| 13.63     | <code>isunordered()</code>                                                                                                            | 165        |
| <b>14</b> | <b>&lt;setjmp.h&gt; Non-local Goto</b>                                                                                                | <b>167</b> |
| 14.1      | <code>setjmp()</code>                                                                                                                 | 167        |
| 14.2      | <code>longjmp()</code>                                                                                                                | 169        |
| <b>15</b> | <b>&lt;signal.h&gt; signal handling</b>                                                                                               | <b>172</b> |
| 15.1      | <code>signal()</code>                                                                                                                 | 172        |
| 15.2      | <code>raise()</code>                                                                                                                  | 176        |
| <b>16</b> | <b>&lt;stdalign.h&gt; Macros for Alignment</b>                                                                                        | <b>178</b> |
| 16.1      | <code>alignas()</code> <code>_Alignas()</code>                                                                                        | 178        |
| 16.2      | <code>alignof()</code> <code>_Alignof()</code>                                                                                        | 180        |
| <b>17</b> | <b>&lt;stdarg.h&gt; Variable Arguments</b>                                                                                            | <b>182</b> |
| 17.1      | <code>va_arg()</code>                                                                                                                 | 182        |
| 17.2      | <code>va_copy()</code>                                                                                                                | 183        |
| 17.3      | <code>va_end()</code>                                                                                                                 | 185        |
| 17.4      | <code>va_start()</code>                                                                                                               | 187        |
| <b>18</b> | <b>&lt;stdatomic.h&gt; Atomic-Related Functions</b>                                                                                   | <b>189</b> |
| 18.1      | Atomic Types                                                                                                                          | 190        |
| 18.2      | Lock-free Macros                                                                                                                      | 191        |
| 18.3      | Atomic Flag                                                                                                                           | 191        |
| 18.4      | Memory Order                                                                                                                          | 191        |
| 18.5      | <code>ATOMIC_VAR_INIT()</code>                                                                                                        | 192        |
| 18.6      | <code>atomic_init()</code>                                                                                                            | 193        |
| 18.7      | <code>kill_dependency()</code>                                                                                                        | 194        |
| 18.8      | <code>atomic_thread_fence()</code>                                                                                                    | 194        |
| 18.9      | <code>atomic_signal_fence()</code>                                                                                                    | 196        |
| 18.10     | <code>atomic_is_lock_free()</code>                                                                                                    | 197        |
| 18.11     | <code>atomic_store()</code>                                                                                                           | 199        |
| 18.12     | <code>atomic_load()</code>                                                                                                            | 200        |
| 18.13     | <code>atomic_exchange()</code>                                                                                                        | 201        |
| 18.14     | <code>atomic_compare_exchange_*()</code>                                                                                              | 202        |
| 18.15     | <code>atomic_fetch_*()</code>                                                                                                         | 204        |
| 18.16     | <code>atomic_flag_test_and_set()</code>                                                                                               | 206        |
| 18.17     | <code>atomic_flag_clear()</code>                                                                                                      | 208        |
| <b>19</b> | <b>&lt;stdbit.h&gt; Bit-Related Functions</b>                                                                                         | <b>210</b> |
| 19.1      | Existence Macro                                                                                                                       | 210        |
| 19.2      | Endianness Macros                                                                                                                     | 210        |
| 19.3      | General Structure of These Functions                                                                                                  | 211        |
| 19.4      | <code>stdc_leading_zeros()</code>                                                                                                     | 212        |
| 19.5      | <code>stdc_leading_ones()</code>                                                                                                      | 213        |
| 19.6      | <code>stdc_trailing_zeros()</code>                                                                                                    | 214        |
| 19.7      | <code>stdc_trailing_ones()</code>                                                                                                     | 215        |
| 19.8      | <code>stdc_first_leading_zero()</code>                                                                                                | 216        |
| 19.9      | <code>stdc_first_leading_one()</code>                                                                                                 | 217        |

|           |                                                                                                       |            |
|-----------|-------------------------------------------------------------------------------------------------------|------------|
| 19.10     | <code>stdc_first_trailing_zero()</code>                                                               | 218        |
| 19.11     | <code>stdc_first_trailing_one()</code>                                                                | 219        |
| 19.12     | <code>stdc_count_zeros()</code>                                                                       | 220        |
| 19.13     | <code>stdc_count_ones()</code>                                                                        | 221        |
| 19.14     | <code>stdc_has_single_bit()</code>                                                                    | 222        |
| 19.15     | <code>stdc_bit_width()</code>                                                                         | 223        |
| 19.16     | <code>stdc_bit_floor()</code>                                                                         | 224        |
| 19.17     | <code>stdc_bit_ceil()</code>                                                                          | 225        |
| <b>20</b> | <b><code>&lt;stdbool.h&gt;</code> Boolean Types</b>                                                   | <b>227</b> |
| 20.1      | Example                                                                                               | 227        |
| 20.2      | <code>_Bool?</code>                                                                                   | 228        |
| <b>21</b> | <b><code>&lt;stddef.h&gt;</code> A Few Standard Definitions</b>                                       | <b>229</b> |
| 21.1      | <code>ptrdiff_t</code>                                                                                | 229        |
| 21.2      | <code>size_t</code>                                                                                   | 230        |
| 21.3      | <code>max_align_t</code>                                                                              | 230        |
| 21.4      | <code>wchar_t</code>                                                                                  | 231        |
| 21.5      | <code>offsetof</code>                                                                                 | 231        |
| <b>22</b> | <b><code>&lt;stdint.h&gt;</code> More Integer Types</b>                                               | <b>232</b> |
| 22.1      | Specific-Width Integers                                                                               | 232        |
| 22.2      | Other Integer Types                                                                                   | 233        |
| 22.3      | Macros                                                                                                | 233        |
| 22.4      | Other Limits                                                                                          | 234        |
| 22.5      | Macros for Declaring Constants                                                                        | 234        |
| <b>23</b> | <b><code>&lt;stdio.h&gt;</code> Standard I/O Library</b>                                              | <b>236</b> |
| 23.1      | <code>remove()</code>                                                                                 | 238        |
| 23.2      | <code>rename()</code>                                                                                 | 239        |
| 23.3      | <code>tmpfile()</code>                                                                                | 240        |
| 23.4      | <code>tmpnam()</code>                                                                                 | 241        |
| 23.5      | <code>fclose()</code>                                                                                 | 242        |
| 23.6      | <code>fflush()</code>                                                                                 | 243        |
| 23.7      | <code>fopen()</code>                                                                                  | 245        |
| 23.8      | <code>freopen()</code>                                                                                | 246        |
| 23.9      | <code>setbuf()</code> , <code>setvbuf()</code>                                                        | 248        |
| 23.10     | <code>printf()</code> , <code>fprintf()</code> , <code>sprintf()</code> , <code>snprintf()</code>     | 249        |
| 23.11     | <code>scanf()</code> , <code>fscanf()</code> , <code>sscanf()</code>                                  | 256        |
| 23.12     | <code>vprintf()</code> , <code>vfprintf()</code> , <code>vsprintf()</code> , <code>vsnprintf()</code> | 262        |
| 23.13     | <code>vscanf()</code> , <code>vfscanf()</code> , <code>vsscanf()</code>                               | 264        |
| 23.14     | <code>getc()</code> , <code>fgetc()</code> , <code>getchar()</code>                                   | 265        |
| 23.15     | <code>gets()</code> , <code>fgets()</code>                                                            | 266        |
| 23.16     | <code>putc()</code> , <code>fputc()</code> , <code>putchar()</code>                                   | 268        |
| 23.17     | <code>puts()</code> , <code>fputs()</code>                                                            | 269        |
| 23.18     | <code>ungetc()</code>                                                                                 | 270        |
| 23.19     | <code>fread()</code>                                                                                  | 272        |
| 23.20     | <code>fwrite()</code>                                                                                 | 273        |
| 23.21     | <code>fgetpos()</code> , <code>fsetpos()</code>                                                       | 274        |
| 23.22     | <code>fseek()</code> , <code>rewind()</code>                                                          | 275        |
| 23.23     | <code>ftell()</code>                                                                                  | 277        |
| 23.24     | <code>feof()</code> , <code>ferror()</code> , <code>clearerr()</code>                                 | 278        |
| 23.25     | <code>perror()</code>                                                                                 | 279        |
| <b>24</b> | <b><code>&lt;stdlib.h&gt;</code> Standard Library Functions</b>                                       | <b>282</b> |

|           |                                                                                                   |            |
|-----------|---------------------------------------------------------------------------------------------------|------------|
| 24.1      | <code>&lt;stdlib.h&gt;</code> Types and Macros                                                    | 283        |
| 24.2      | <code>atof()</code>                                                                               | 283        |
| 24.3      | <code>atoi()</code> , <code>atol()</code> , <code>atoll()</code>                                  | 284        |
| 24.4      | <code>strtod()</code> , <code>strtof()</code> , <code>strtold()</code>                            | 286        |
| 24.5      | <code>strtol()</code> , <code>strtoll()</code> , <code>strtoul()</code> , <code>strtoull()</code> | 288        |
| 24.6      | <code>rand()</code>                                                                               | 290        |
| 24.7      | <code>srand()</code>                                                                              | 291        |
| 24.8      | <code>aligned_alloc()</code>                                                                      | 293        |
| 24.9      | <code>calloc()</code> , <code>malloc()</code>                                                     | 294        |
| 24.10     | <code>free()</code>                                                                               | 296        |
| 24.11     | <code>realloc()</code>                                                                            | 297        |
| 24.12     | <code>abort()</code>                                                                              | 298        |
| 24.13     | <code>atexit()</code> , <code>at_quick_exit()</code>                                              | 299        |
| 24.14     | <code>exit()</code> , <code>quick_exit()</code> , <code>_Exit()</code>                            | 301        |
| 24.15     | <code>getenv()</code>                                                                             | 302        |
| 24.16     | <code>system()</code>                                                                             | 303        |
| 24.17     | <code>bsearch()</code>                                                                            | 304        |
| 24.18     | <code>qsort()</code>                                                                              | 306        |
| 24.19     | <code>abs()</code> , <code>labs()</code> , <code>llabs()</code>                                   | 308        |
| 24.20     | <code>div()</code> , <code>ldiv()</code> , <code>lldiv()</code>                                   | 309        |
| 24.21     | <code>mblen()</code>                                                                              | 310        |
| 24.22     | <code>mbtowc()</code>                                                                             | 312        |
| 24.23     | <code>wctomb()</code>                                                                             | 313        |
| 24.24     | <code>mbstowcs()</code>                                                                           | 314        |
| 24.25     | <code>wcstombs()</code>                                                                           | 316        |
| 24.26     | <code>memalignment()</code>                                                                       | 317        |
| <b>25</b> | <b><code>&lt;stdnoreturn.h&gt;</code> Macros for Non-Returning Functions</b>                      | <b>319</b> |
| <b>26</b> | <b><code>&lt;string.h&gt;</code> String Manipulation</b>                                          | <b>320</b> |
| 26.1      | <code>memcpy()</code> , <code>memccpy()</code> , <code>memmove()</code>                           | 321        |
| 26.2      | <code>strcpy()</code> , <code>strncpy()</code>                                                    | 322        |
| 26.3      | <code>strdup()</code> , <code>strndup()</code>                                                    | 323        |
| 26.4      | <code>strcat()</code> , <code>strncat()</code>                                                    | 324        |
| 26.5      | <code>strcmp()</code> , <code>strncmp()</code> , <code>memcmp()</code>                            | 326        |
| 26.6      | <code>strcoll()</code>                                                                            | 327        |
| 26.7      | <code>strxfrm()</code>                                                                            | 328        |
| 26.8      | <code>strchr()</code> , <code>strrchr()</code> , <code>memchr()</code>                            | 330        |
| 26.9      | <code>strspn()</code> , <code>strcspn()</code>                                                    | 332        |
| 26.10     | <code>strpbrk()</code>                                                                            | 333        |
| 26.11     | <code>strstr()</code>                                                                             | 334        |
| 26.12     | <code>strtok()</code>                                                                             | 335        |
| 26.13     | <code>memset()</code> , <code>'memset_explicit</code>                                             | 336        |
| 26.14     | <code>strerror()</code>                                                                           | 337        |
| 26.15     | <code>strlen()</code>                                                                             | 339        |
| <b>27</b> | <b><code>&lt;tgmath.h&gt;</code> Type-Generic Math Functions</b>                                  | <b>340</b> |
| 27.1      | Example                                                                                           | 341        |
| <b>28</b> | <b><code>&lt;threads.h&gt;</code> Multithreading Functions</b>                                    | <b>343</b> |
| 28.1      | <code>call_once()</code>                                                                          | 344        |
| 28.2      | <code>cnd_broadcast()</code>                                                                      | 345        |
| 28.3      | <code>cnd_destroy()</code>                                                                        | 348        |
| 28.4      | <code>cnd_init()</code>                                                                           | 349        |
| 28.5      | <code>cnd_signal()</code>                                                                         | 351        |



|           |                                                                            |            |
|-----------|----------------------------------------------------------------------------|------------|
| 28.6      | <code>cnd_timedwait()</code>                                               | 353        |
| 28.7      | <code>cnd_wait()</code>                                                    | 355        |
| 28.8      | <code>mtx_destroy()</code>                                                 | 356        |
| 28.9      | <code>mtx_init()</code>                                                    | 358        |
| 28.10     | <code>mtx_lock()</code>                                                    | 360        |
| 28.11     | <code>mtx_timedlock()</code>                                               | 361        |
| 28.12     | <code>mtx_trylock()</code>                                                 | 363        |
| 28.13     | <code>mtx_unlock()</code>                                                  | 365        |
| 28.14     | <code>thrd_create()</code>                                                 | 367        |
| 28.15     | <code>thrd_current()</code>                                                | 368        |
| 28.16     | <code>thrd_detach()</code>                                                 | 370        |
| 28.17     | <code>thrd_equal()</code>                                                  | 371        |
| 28.18     | <code>thrd_exit()</code>                                                   | 373        |
| 28.19     | <code>thrd_join()</code>                                                   | 374        |
| 28.20     | <code>thrd_sleep()</code>                                                  | 376        |
| 28.21     | <code>thrd_yield()</code>                                                  | 377        |
| 28.22     | <code>tss_create()</code>                                                  | 378        |
| 28.23     | <code>tss_delete()</code>                                                  | 381        |
| 28.24     | <code>tss_get()</code>                                                     | 383        |
| 28.25     | <code>tss_set()</code>                                                     | 385        |
| <b>29</b> | <b>&lt;time.h&gt; Date and Time Functions</b>                              | <b>388</b> |
| 29.1      | Thread Safety Warning                                                      | 389        |
| 29.2      | <code>clock()</code>                                                       | 389        |
| 29.3      | <code>difftime()</code>                                                    | 390        |
| 29.4      | <code>mktime()</code>                                                      | 391        |
| 29.5      | <code>timegm()</code>                                                      | 393        |
| 29.6      | <code>time()</code>                                                        | 394        |
| 29.7      | <code>timespec_get()</code>                                                | 395        |
| 29.8      | <code>asctime()</code>                                                     | 397        |
| 29.9      | <code>ctime()</code>                                                       | 398        |
| 29.10     | <code>gmtime()</code>                                                      | 399        |
| 29.11     | <code>localtime()</code>                                                   | 400        |
| 29.12     | <code>strftime()</code>                                                    | 401        |
| <b>30</b> | <b>&lt;uchar.h&gt; Unicode utility functions</b>                           | <b>406</b> |
| 30.1      | Types                                                                      | 406        |
| 30.2      | OS X issue                                                                 | 407        |
| 30.3      | <code>mbrtoc16()</code> <code>mbrtoc32()</code>                            | 407        |
| 30.4      | <code>c16rtomb()</code> <code>c32rtomb()</code>                            | 410        |
| <b>31</b> | <b>&lt;wchar.h&gt; Wide Character Handling</b>                             | <b>413</b> |
| 31.1      | Restartable Functions                                                      | 414        |
| 31.2      | <code>wprintf()</code> , <code>fwprintf()</code> , <code>swprintf()</code> | 415        |
| 31.3      | <code>wscanf()</code> <code>fwscanf()</code> <code>swscanf()</code>        | 416        |
| 31.4      | <code>vwprintf()</code> <code>vfwprintf()</code> <code>vswprintf()</code>  | 417        |
| 31.5      | <code>wscanf()</code> , <code>vfwscanf()</code> , <code>vswscanf()</code>  | 418        |
| 31.6      | <code>getwc()</code> <code>fgetwc()</code> <code>getwchar()</code>         | 420        |
| 31.7      | <code>fgetws()</code>                                                      | 421        |
| 31.8      | <code>putwchar()</code> <code>putwc()</code> <code>fputwc()</code>         | 422        |
| 31.9      | <code>fputws()</code>                                                      | 423        |
| 31.10     | <code>fwide()</code>                                                       | 424        |
| 31.11     | <code>ungetwc()</code>                                                     | 426        |
| 31.12     | <code>wcstod()</code> <code>wcstof()</code> <code>wcstold()</code>         | 427        |

|           |                                                                                                       |            |
|-----------|-------------------------------------------------------------------------------------------------------|------------|
| 31.13     | <code>wcstol()</code> <code>wcstoll()</code> <code>wcstoul()</code> <code>wcstoull()</code> . . . . . | 428        |
| 31.14     | <code>wscpy()</code> <code>wcsncpy()</code> . . . . .                                                 | 430        |
| 31.15     | <code>wmemcpy()</code> <code>wmemmove()</code> . . . . .                                              | 431        |
| 31.16     | <code>wscat()</code> <code>wcsncat()</code> . . . . .                                                 | 432        |
| 31.17     | <code>wscmp()</code> , <code>wcsncmp()</code> , <code>wmemcmp()</code> . . . . .                      | 433        |
| 31.18     | <code>wscoll()</code> . . . . .                                                                       | 434        |
| 31.19     | <code>wcsxfrm()</code> . . . . .                                                                      | 435        |
| 31.20     | <code>wcschr()</code> <code>wcsrchr()</code> . . . . .                                                | 437        |
| 31.21     | <code>wcsspn()</code> <code>wcscspn()</code> . . . . .                                                | 438        |
| 31.22     | <code>wcspbrk()</code> . . . . .                                                                      | 440        |
| 31.23     | <code>wcsstr()</code> . . . . .                                                                       | 440        |
| 31.24     | <code>wcstok()</code> . . . . .                                                                       | 441        |
| 31.25     | <code>wcslen()</code> . . . . .                                                                       | 443        |
| 31.26     | <code>wcsftime()</code> . . . . .                                                                     | 443        |
| 31.27     | <code>btowc()</code> <code>wctob()</code> . . . . .                                                   | 445        |
| 31.28     | <code>mbsinit()</code> . . . . .                                                                      | 446        |
| 31.29     | <code>mbrlen()</code> . . . . .                                                                       | 447        |
| 31.30     | <code>mbrtowc()</code> . . . . .                                                                      | 448        |
| 31.31     | <code>wcrtomb()</code> . . . . .                                                                      | 450        |
| 31.32     | <code>mbsrtowcs()</code> . . . . .                                                                    | 451        |
| 31.33     | <code>wcsrtombs()</code> . . . . .                                                                    | 453        |
| <b>32</b> | <b>&lt;wctype.h&gt; Wide Character Classification and Transformation</b>                              | <b>457</b> |
| 32.1      | <code>iswalnum()</code> . . . . .                                                                     | 457        |
| 32.2      | <code>iswalpha()</code> . . . . .                                                                     | 458        |
| 32.3      | <code>iswblank()</code> . . . . .                                                                     | 459        |
| 32.4      | <code>iswcntrl()</code> . . . . .                                                                     | 460        |
| 32.5      | <code>iswdigit()</code> . . . . .                                                                     | 461        |
| 32.6      | <code>iswgraph()</code> . . . . .                                                                     | 462        |
| 32.7      | <code>iswlower()</code> . . . . .                                                                     | 463        |
| 32.8      | <code>iswprint()</code> . . . . .                                                                     | 464        |
| 32.9      | <code>iswpunct()</code> . . . . .                                                                     | 464        |
| 32.10     | <code>iswspace()</code> . . . . .                                                                     | 465        |
| 32.11     | <code>iswupper()</code> . . . . .                                                                     | 466        |
| 32.12     | <code>iswxdigit()</code> . . . . .                                                                    | 467        |
| 32.13     | <code>iswctype()</code> . . . . .                                                                     | 468        |
| 32.14     | <code>wctype()</code> . . . . .                                                                       | 470        |
| 32.15     | <code>towlower()</code> . . . . .                                                                     | 471        |
| 32.16     | <code>towupper()</code> . . . . .                                                                     | 472        |
| 32.17     | <code>towctrans()</code> . . . . .                                                                    | 473        |
| 32.18     | <code>wctrans()</code> . . . . .                                                                      | 475        |

# Chapter 1

## Foreword

The door slowly creaks open revealing a long hall with dusty stacks of books of lore...

I admit, maybe not that.

But you have found the Library Reference portion of Beej's Guide to C!

This isn't a tutorial, but rather is a comprehensive set of manual pages (or *man pages* as Unix hackers like to say) that define *every* function in the C Standard Library, complete with examples.

“*This book, sir, contains every word in our beloved language.*”

“*Every single one, sir?*”

“*Every single one, sir!*”

“*Ah, well in that case, sir, I hope you will not object if I also offer the doctor my most enthusiastic contrafibularities.*”

—Blackadder toying with Dr. Samuel Johnson

There are, in fact, a number of functions left out of this guide, most notably all the optional “safe” functions (with a `_s` suffix).

But everything you're likely to want is definitely covered in here. With examples.

Probably.

### 1.1 Audience

This guide is for people who are at least modestly proficient in C.

If you are not one of those people and wish to become one of those people, I can wholeheartedly recommend with zero bias the book *Beej's Guide to C Programming*<sup>1</sup>, freely available wherever the Internet is sold.

### 1.2 How to Read This Book

Use the contents or index to find the function or category you're after.

Then grab a bowl of your favorite cereal and devour the delicious, delicious verbiage.

---

<sup>1</sup><https://beej.us/guide/bgc/>

## 1.3 Platform and Compiler

I'll try to stick to Plain Ol'-Fashioned ISO-standard C<sup>2</sup>. Well, for the most part. Here and there I might go crazy and start talking about POSIX<sup>3</sup> or something, but we'll see.

**Unix** users (e.g. Linux, BSD, etc.) try running `cc` or `gcc` from the command line—you might already have a compiler installed. If you don't, search your distribution for installing `gcc` or `clang`.

**Windows** users should check out Visual Studio Community<sup>4</sup>. Or, if you're looking for a more Unix-like experience (recommended!), install WSL<sup>5</sup> and `gcc`.

**Mac** users will want to install XCode<sup>6</sup>, and in particular the command line tools.

There are a lot of compilers out there, and virtually all of them will work for this book. And a C++ compiler will compile a lot of (but not all!) C code. Best use a proper C compiler if you can.

## 1.4 Official Homepage

This official location of this document is <https://beej.us/guide/bgclr/><sup>7</sup>. There used to be a note here about migrating off Chico State's computers (my alma mater), but that's something that happened roughly a zillion years ago and the wording remained here only because it was copied over from the Network Guide, [*breath*] which I apparently haven't read in its entirety for quite some time.

The End.

## 1.5 Email Policy

I'm generally available to help out with email questions so feel free to write in, but I can't guarantee a response. I lead a pretty busy life and there are times when I just can't answer a question you have. When that's the case, I usually just delete the message. It's nothing personal; I just won't ever have the time to give the detailed answer you require.

As a rule, the more complex the question, the less likely I am to respond. If you can narrow down your question before mailing it and be sure to include any pertinent information (like platform, compiler, error messages you're getting, and anything else you think might help me troubleshoot), you're much more likely to get a response.

If you don't get a response, hack on it some more, try to find the answer, and if it's still elusive, then write me again with the information you've found and hopefully it will be enough for me to help out.

Now that I've badgered you about how to write and not write me, I'd just like to let you know that I *fully* appreciate all the praise the guide has received over the years. It's a real morale boost, and it gladdens me to hear that it is being used for good! :- ) Thank you!

## 1.6 Mirroring

You are more than welcome to mirror this site, whether publicly or privately. If you publicly mirror the site and want me to link to it from the main page, drop me a line at [beej@beej.us](mailto:beej@beej.us).

---

<sup>2</sup>[https://en.wikipedia.org/wiki/ANSI\\_C](https://en.wikipedia.org/wiki/ANSI_C)

<sup>3</sup><https://en.wikipedia.org/wiki/POSIX>

<sup>4</sup><https://visualstudio.microsoft.com/vs/community/>

<sup>5</sup><https://docs.microsoft.com/en-us/windows/wsl/install-win10>

<sup>6</sup><https://developer.apple.com/xcode/>

<sup>7</sup><https://beej.us/guide/bgclr/>

## 1.7 Note for Translators

If you want to translate the guide into another language, write me at [beej@beej.us](mailto:beej@beej.us) and I'll link to your translation from the main page. Feel free to add your name and contact info to the translation.

Please note the license restrictions in the Copyright and Distribution section, below.

## 1.8 Copyright and Distribution

Beej's Guide to C Programming—Library Reference is Copyright © 2021 Brian “Beej Jorgensen” Hall.

With specific exceptions for source code and translations, below, this work is licensed under the Creative Commons Attribution-Noncommercial-No Derivative Works 3.0 License. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

One specific exception to the “No Derivative Works” portion of the license is as follows: this guide may be freely translated into any language, provided the translation is accurate, and the guide is reprinted in its entirety. The same license restrictions apply to the translation as to the original guide. The translation may also include the name and contact information for the translator.

The C source code presented in this document is hereby granted to the public domain, and is completely free of any license restriction.

Educators are freely encouraged to recommend or supply copies of this guide to their students.

Contact [beej@beej.us](mailto:beej@beej.us) for more information.

## 1.9 Dedication

The hardest things about writing these guides are:

- Learning the material in enough detail to be able to explain it
- Figuring out the best way to explain it clearly, a seemingly-endless iterative process
- Putting myself out there as a so-called *authority*, when really I'm just a regular human trying to make sense of it all, just like everyone else
- Keeping at it when so many other things draw my attention

A lot of people have helped me through this process, and I want to acknowledge those who have made this book possible.

- Everyone on the Internet who decided to help share their knowledge in one form or another. The free sharing of instructive information is what makes the Internet the great place that it is.
- The volunteers at [cppreference.com](http://cppreference.com)<sup>8</sup> who provide the bridge that leads from the spec to the real world.
- The helpful and knowledgeable folks on [comp.lang.c](http://comp.lang.c)<sup>9</sup> and [r/C\\_Programming](http://r/C_Programming)<sup>10</sup> who got me through the tougher parts of the language.
- Everyone who submitted corrections and pull-requests on everything from misleading instructions to typos.

Thank you! ♥

---

<sup>8</sup><https://en.cppreference.com/>

<sup>9</sup><https://groups.google.com/g/comp.lang.c>

<sup>10</sup>[https://www.reddit.com/r/C\\_Programming/](https://www.reddit.com/r/C_Programming/)

## Chapter 2

# The C Language

This is just a quick overview of the fashionable and fun highlights of the syntax, keywords, and other animals in the C menagerie.

### 2.1 Background

Some things you'll need to make sense of the examples, below.

#### 2.1.1 Comments

Comments in C start with `//` and go to the end of a line.

Multiline comments begin with `/*` and continue until a closing `*/`.

#### 2.1.2 Separators

Expressions in C are separated by semicolons (`;`). These tend to appear at the ends of lines.

#### 2.1.3 Expressions

If it's not a keyword or other reserved punctuation, it tends to be an expression. Think “math including function calls”.

#### 2.1.4 Statements

Think `if`, `while`, etc. Executable keywords.

#### 2.1.5 Booleans

Ignoring the `bool` type, zero is false and non-zero is true.

#### 2.1.6 Blocks

Multiple expressions and flow control keywords can be wrapped up in a block, made up of `{` followed by one or more expressions or statements, followed by `}`.

### 2.1.7 Code Examples

They are meant to give an idea of how to use various statements, but not be comprehensive in terms of examples.

In the examples, below, if either an expression or statement can be used, the word `code` is inserted.

## 2.2 Operators

### 2.2.1 Arithmetic Operators

The arithmetic operators: `+`, `-`, `*`, `/`, `%` (remainder).

Division is integer if all arguments are integers. Otherwise it's a floating result.

You can also negate an expression by putting `-` in front of it. (You can also put a `+` in front of it—this doesn't do anything mathematically, but it causes the Usual Arithmetic Conversions to be performed on the expression.)

### 2.2.2 Pre- and Post-Increment and -Decrement

The post-increment (`++`) and post-decrement (`--`) operators (after the variable) do their work *after* the rest of the expression has been evaluated.

```
int x = 10;
int y = 20;
int z = 30;

int w = (x++) + (y--) + (z++);

print("%d %d %d %d\n", x, y, z, w); // 11 19 31 60
```

The pre-increment (`++`) and pre-decrement (`--`) operators (before the variable) do their work *before* the rest of the expression has been evaluated.

```
int x = 10;
int y = 20;
int z = 30;

int w = (++x) + (--y) + (++z);

print("%d %d %d %d\n", x, y, z, w); // 11 19 31 61
```

### 2.2.3 Comparison Operators

All of these return a Boolean true-y or false-y value.

Less than, greater than, and equal to are: `<`, `>`, `==`, respectively.

Less than or equal to and greater than or equal to are `<=` and `>=`.

Not equal to is `!=`.

### 2.2.4 Pointer Operators

`*` in front of a pointer variable dereferences that variable.

`&` in front of a variable gives the address of that variable.

`+` and `-` arithmetic operators work on pointers for pointer arithmetic.

## 2.2.5 Structure and Union Operators

The dot operator (`.`) can get a field value out of a `struct` or `union`.

The arrow operator (`->`) can get a field value out of a pointer to a `struct` or `union`. These two are equivalent, assuming `p` is just such a pointer:

```
(*p).bar;
p->bar;
```

## 2.2.6 Array Operators

The square bracket operators can reference a value in an array:

```
a[10] = 99;
```

This is syntactic sugar over pointer arithmetic and referencing. The above line is equivalent to:

```
*(a + 10) = 99;
```

## 2.2.7 Bitwise Operators

Bit shift right: `>>`, bit shift left: `<<`.

```
int i = x << 3; // left shift 3 bits
```

Whether or not a right shift on a signed value is sign-extended is implementation-defined.

Bitwise AND, OR, NOT, and XOR are `&`, `|`, `~`, and `^`, respectively.

## 2.2.8 Assignment Operators

A standalone `=` is your basic assignment.

But there are also compound assignments that are like a shorthand version. For example, these two are basically equivalent:

```
x = x + 1;
x += 1;
```

There are compound assignment operators for many of the other operators.

Arithmetic: `+=`, `-=`, `*=`, `/=`, and `%=`.

Bitwise: `|=`, `&=`, `~=`, and `^=`.

## 2.2.9 The `sizeof` Operator

This is a compile-time operator that gives you the size in bytes of the type of the argument. The type of the expression is used; the expression is not evaluated. `sizeof` works with any type, even user-defined composite types.



The return type is the integer type `size_t`.

```
float f;  
size_t x = sizeof f;  
  
printf("f is %zu bytes\n", x);
```

You can also specify a raw type name in there by wrapping it in parentheses:

```
size_t x = sizeof(int);  
  
printf("int is %zu bytes\n", x);
```

### 2.2.10 Type Casts

You can force an expression to be another type (within reason) by *casting* to that type.

You give the new type name in parentheses.

Here we are forcing the subexpression `x` to be type `float` just before the division<sup>1</sup>. This causes the division, which would otherwise be an integer division, to be a floating point division.

```
int x = 17;  
int y = 2;  
  
float f = (float)x / y;
```

### 2.2.11 `_Alignof` Operator

You can get the byte alignment of any type with the `_Alignof` compile-time operator. If you include `<stdalign.h>`, you can use `alignof` instead.

Any type can be the argument to the operator, which must be in parenthesis. Unlike `sizeof`, the argument cannot be an expression.

```
printf("Alignment of int is %zu\n", alignof(int));
```

### 2.2.12 Comma Operator

You can separate subexpressions with commas, and each will be evaluated from left to right, and the value of the entire expression will be the value of the subexpression after the last comma.

```
int x = (1, 2, 3); // Silly way to assign `x = 3`
```

Usually this is used in the various clauses in loops. For example, we can do multiple assignments in a `for` loop, and have multiple post expressions like this:

```
for (i = 2, j = 10; i < 100; i++, j += 4) { ... }
```

<sup>1</sup>This doesn't change the type of `x` in other contexts—it's just in this one usage in this expression.

## 2.3 Type Specifiers

Integer types from smallest to largest capacity: `char`, `short`, `int`, `long`, `long long`.

Any integer type may be prefaced with `signed` (the default except for `char`) or `unsigned`.

Whether or not `char` is signed is implementation defined.

Floating types from least accuracy to most: `float`, `double`, `long double`.

`void` is a type representing lack of type.

`_Bool` is a Boolean type. This becomes `bool` in C23. Earlier versions of C must include `<stdbool.h>` to get `bool`.

`_Complex` indicates a complex floating type type, when paired with such a type. Include `<complex.h>` to use `complex` instead.

```
complex float x = 1.2 + 2.3*I;
complex double y = 1.2 + 2.3*I;
```

`_Imaginary` is an optional keyword used to specify an imaginary type (the imaginary part of a complex number) when paired with a floating type. Include `<complex.h>` to use `imaginary` instead. Neither GCC nor clang support this.

```
imaginary float f = 2.3*I;
```

`_Generic` is a type “switcher” that allows you to emit different code at compile time depending on the type of the data.

## 2.4 Constant Types

You can declare constants to be of specific types (though it might be a larger type). In the following example unqualified types, case doesn’t matter, and the `U` can come before or after the `L` or `LL`.

|                             |                                          |
|-----------------------------|------------------------------------------|
| <code>123</code>            | <code>int</code> or larger               |
| <code>123L</code>           | <code>long int</code> or larger          |
| <code>123LL</code>          | <code>long long int</code>               |
| <code>123U</code>           | <code>unsigned int</code> or larger      |
| <code>123UL</code>          | <code>unsigned long int</code> or larger |
| <code>123ULL</code>         | <code>unsigned long long int</code>      |
| <code>123.4F</code>         | <code>float</code>                       |
| <code>123.4</code>          | <code>double</code>                      |
| <code>123.4L</code>         | <code>long double</code>                 |
| <code>'a'</code>            | <code>char</code>                        |
| <code>"hello, world"</code> | <code>char*</code> (string)              |

You can specify the constant in other bases as well:

|                    |             |
|--------------------|-------------|
| <code>123</code>   | decimal     |
| <code>0x123</code> | hexadecimal |
| <code>0123</code>  | octal       |

You can also specify floating constants in base-10 exponential notation:

```
1.2e3          1.2 x 10^3
```

And you can specify floats in hex! Except in this case the exponent is still in decimal, and the base is 2 instead of 10:

```
0x1.2p3        0x1.2 x 2^3
```

## 2.5 Composite Types

### 2.5.1 struct Types

You can build a composite type made out of other types with `struct` and then declare variables to be of that type.

```
struct animal {  
    char *name;  
    int leg_count;  
};  
  
struct animal a;  
struct animal b = {"goat", 4};  
struct animal c = {.name="goat", .leg_count=4};
```

Accessing is done with the dot operator (`.`) or, if the variable is a pointer to a `struct`, the arrow operator (`->`).

```
struct animal *p = &b; // b from above  
  
printf("%d\n", b.leg_count);  
printf("%d\n", p->leg_count);
```

### 2.5.2 union Types

These are like `struct` types in usage, except that you can only use one field at a time. (The fields all use the same region of memory.)

```
union dt {  
    float distance;  
    int time;  
};  
  
union dt a;  
union dt b = {6};           // Initializes "distance", the first field  
union dt c = {.distance=6}; // Initializes "distance"  
union dt d = {.time=6};     // Initializes "time"
```

Accessing is done with the dot operator (`.`) or, if the variable is a pointer to a `union`, the arrow operator (`->`).

```
union dt *p = &b;

printf("%d\n", b.time);
printf("%d\n", p->time);
```

### 2.5.3 enum Types

Gives you a typed way to have named constant integer values. These can be used with `switch()`, or as an array size, or any other place constant values are needed.

Names are conventionally capitalized.

```
enum animal {
    ANTELOPE,
    BADGER,
    CAT,
    DOG,
    ELEPHANT,
    FISH
};

enum animal a = CAT;

if (a == CAT)
    printf("The animal is a cat.\n");
```

The names have numeric values starting with zero and counting up. (In the example above, `DOG` would be 3.)

The numeric value can be overridden by specifying an integer exactly. Subsequent values increment from the specified one.

```
enum animal {
    ANTELOPE = 4,
    BADGER,           // Will be 5
    CAT,              // Will be 6
    DOG = 3,
    ELEPHANT,         // Will be 4
    FISH              // Will be 5
};
```

As above, duplicate values are not illegal, but might be of marginal usefulness.

## 2.6 Initializers

You can do this when the variable is defined, but not elsewhere.

Initializing basic types:

```
int x = 12;
float y = 1.2;
char c = 'a';
```

```
char *s = "Hello, world!";
```

Initializing array types:

```
int a[3] = {1,2,3};
int a[] = {1,2,3};    // Same as a[3]

int a[3] = {1, 2};    // Same as {1, 2, 0}
int a[3] = {1};       // Same as {1, 0, 0}
int a[3] = {0};       // Same as {0, 0, 0}
```

Initializing pointer types:

```
int q;
int *p = &q;
```

Initializing structs:

```
struct s {
    int a;
    float b;
};

struct s x0 = {1, 2.2}; // Initialize fields in order

struct s x0 = {.a=1, .b=2.2}; // Initialize fields by name
struct s x0 = {.b=2.2, .a=1}; // Same thing

struct s x0 = {.b=2.2}; // All other fields initialized to 0
struct s x0 = {.b=2.2, .a=0}; // Same thing
```

Initializing unions:

```
union u {
    int a;
    float b;
};

union u x0 = {1}; // Initialize the first field (a)

union u x0 = {.a=1}; // Initialize fields by name
union u x0 = {.b=2.2};

//union u x0 = {1, 2}; // ILLEGAL
//union u x0 = {.a=1, .b=2}; // ILLEGAL
```

## 2.7 Compound Literals

You can declare “unnamed” objects in C. This is often useful for passing a `struct` to a function that otherwise doesn’t need a name.

You use the type name in parens followed by an initializer to make the object.

Here's an example of passing a compound literal to a function. Note that there's no `struct s` variable in `main()`:

```
1 #include <stdio.h>
2
3 struct s {
4     int a, b;
5 };
6
7 int add(struct s x)
8 {
9     return x.a + x.b;
10 }
11
12 int main(void)
13 {
14     int t = add((struct s){.a=2, .b=4}); // <-- Here
15
16     printf("%d\n", t);
17 }
```

Compound literals have the lifetime of their scope.

You can also pass a pointer to a compound literal by taking its address:

```
foo(&(struct s){1, 2});
```

## 2.8 Type Aliases

You can set up a type alias for convenience or abstraction.

Here we'll make a new type called `time_counter` that is just an `int`. It can only be used exactly like an `int`. It's just an alias for an `int`.

```
typedef int time_counter;

time_counter t = 3490;
```

Also works with `structs` or `unions`:

```
struct foo {
    int bar;
    float baz;
};

typedef struct foo funtype;

funtype f = {1, 2}; // "funtype" is an alias for "struct foo";
```

It also works inline, and with named or unnamed `structs` or `unions`:

```
typedef struct {
    int bar;
    float baz;
} funtype;

funtype f = {1, 2}; // "funtype" is an alias for the unnamed struct
```

## 2.9 Additional Type-Related Specifiers

You can give the compiler more hints about what qualities a type should have using these specifiers and qualifiers.

### 2.9.1 Storage Class Specifiers

These can be placed before a type to provide more guidance about how the type is used.

```
auto int a
register int a
static int a
extern int a
thread_local int a
```

**auto** is the default, so it's basically never used. Indicates automatic storage duration (things like local variables get freed automatically when they fall out of scope). In C23 this keyword changes to indicate type inference like C++.

**register** indicates that accessing this variable should be as quick as possible. Restricts some usage of the variable giving the compiler a chance to optimize. Rare in daily use.

**static** at function scope indicates that this variable's value should persist from call to call. At file scope indicates that this variable should not be visible outside of this source file.

**extern** indicates that this variable refers to one declared in another source file.

**\_Thread\_local** means that every thread gets its own copy of this variable. You can use **thread\_local** if you include `<threads.h>`.

### 2.9.2 Type Qualifiers

These can be placed before a type to provide more guidance about how the type is used.

```
const int a
const int *p
int * const p
const int * const p
int * restrict p
volatile int a
atomic int a
```

**const** means the value can't be modified. You can use it with pointers, as well:

```
const int a = 10;           // Can't modify "a"
```

```
const int *p = &b      // Can't modify the thing "p" points to ("b")
int *const p = &b      // Can't modify "p"
const int *const p = &b // Can't modify "p" or the thing it points to
```

`restrict` on a pointer means that there will only be one pointer to the item in question, freeing the compiler to make some optimizations.

`volatile` indicates that the value in a variable might change at any time and should be loaded from memory instead of being kept in a register. Usually used with memory-mapped hardware.

`_Atomic` (or `atomic` if you include `<stdatomic.h>`) tells the compiler that reads or writes to this type should happen atomically. (This might be accomplished with a lock depending on the platform and type.)

### 2.9.2.1 C23 Qualified Pseudo-Types: `QVoid*`, `QChar*`, etc.

There are some generic functions in C23 that will return a `const`-qualified type if one of the parameters is `const`, but not otherwise.

The spec makes up a fake type for this, with a `Q` at the front (for “qualified”). This is not a real type and will not compile—it’s just for documentation purposes.

These pseudotypes are:

- `QVoid *`
- `QChar *`
- `QWchar_t *`

For example, the `strchr()` function, which searches a string for a character, has this prototype in the spec:

```
QChar *strchr(QChar *s, int c);
```

What is it? It basically means that if `s` is type `const char *`, then the return type of the function will also be `const char *`.

If `s` is merely `char *`, the return type of the function will merely be `char *`.

In other words, the `const`-ness of `s` is preserved in the return value.

Another way to look at it is that this:

```
QChar *strchr(QChar *s, int c);
```

is the same as:

```
char *strchr(char *s, int c);
const char *strchr(const char *s, int c);
```

The TLDR is when you see this, drop the leading `Q` and change the next letter to lowercase and you’re there.

### 2.9.3 Function Specifiers

These are used on functions to provide additional guidance for the compiler.

`_Noreturn` indicates that a function will never return. It can only run forever or exit the program entirely. If you include `<stdnoreturn.h>`, you can use `noreturn` instead.

`inline` indicates that calls to this function should be as fast as possible. The intention here is that the code of the function be moved *inline* to remove the overhead of the call and return. The compiler regards `inline` as a suggestion, not a requirement.



### 2.9.4 Alignment Specifier

You can force the alignment of a variable with memory with `_Alignas`. If you include `<stdalign.h>` you can use `alignas` instead.

`alignas(0)` has no effect.

```
alignas(16) int a = 12;    // 16-byte alignment
alignas(long) int b = 34;  // Same alignment as "long"
```

### 2.10 if Statement

```
if (boolean_expression) code;

if (boolean_expression) {
    code;
    code;
    code;
}

if (boolean_expression) {
    code;
    code;
} else
    code;

if (boolean_expression) {
    code;
    code;
} else if {
    code;
    code;
    code;
} else {
    code;
}
```

### 2.11 for Statement

Classic `for`-loop.

The bit in parens comes in three parts separated by semicolons:

- Initialization, executed once.
- Block entry condition, evaluated every time before entering the loop body.
- Post expression, evaluated every time after the loop body.

For example, initialize `i` to `0`, enter the loop body while `i < 10`, and then increment `i` after each loop iteration:

```
for (i = 0; i < 10; i++) {
    code;
    code;
}
```

```
    code;
}
```

You can declare loop-local variables by specifying their type:

```
for (int i = 0; i < 10; i++) {
    code;
    code;
}
```

You can separate parts of the expressions with the comma operator:

```
for (i = 0, j = 5; i < 10; i++, j *= 3) {
    code;
    code;
}
```

## 2.12 while Statement

This loop won't enter if the Boolean expression is false. The continuation test happens before the loop.

```
while (boolean_expression) code;

while (boolean_expression) {
    code;
    code;
}
```

## 2.13 do-while Statement

This loop will run at least once even if the Boolean expression is false. The continuation test doesn't happen until after the loop.

```
do code while (boolean_expression);

do {
    code;
    code;
} while (boolean_expression);
```

## 2.14 switch Statement

Performs actions based on the value of an expression. The cases that it is compared against must be constant values.

If the optional `default` is present, that code is executed if none of the cases match. Braces are not required around the cases.

```
switch (expression) {  
    case constant:  
        code;  
        code;  
        break;  
  
    case constant:  
        code;  
        code;  
        break;  
  
    default:  
        code;  
        break;  
}
```

The final `break` in the `switch` is unnecessary if there are no cases after it.

If the `break` isn't present, the `case` falls through to the next one. It's nice to put a comment to that effect so other devs don't hate you.

```
switch (expression) {  
    case constant:  
        code;  
        code;  
        // fall through!  
  
    case constant:  
        code;  
        break;  
}
```

## 2.15 `break` Statement

This breaks out of a `switch` case, but it also can break out of any loop.

```
while (boolean_expression) {  
    code;  
  
    if (boolean_expression)  
        break;  
  
    code;  
}
```

## 2.16 `continue` Statement

This can be used to short-circuit a loop and go to the next continuation condition test without completing the body of the loop.

```
while (boolean_expression) {
    code;
    code;

    if (boolean_expression_2)
        continue;

    // If boolean_expression_2, code down here will be skipped:

    code;
    code;
}
```

## 2.17 goto Statement

You can just jump anywhere within a function with `goto`. (You can't `goto` between functions, only within the same function as the `goto`.)

The destination of the `goto` is a *label*, which is an identifier followed by a colon (:). Labels are typically left-justified all the way to the margin to make them visually stand out.

```
{
    // Abusive demo code that should be a while loop

    int i = 0;

loop:

    printf("%d\n", i++);

    if (i < 10)
        goto loop;
}
```

## 2.18 return Statement

This is how you get back from a function. You can `return` multiple times or just once.

If a function with `void` return type falls off the end, the `return` is implicit.

If the return type is not `void`, the `return` statement must specify a return value of the same type.

Parentheses around the return value are not necessary (as it's a statement, not a function).

```
int increment(int a)
{
    return a + 1;
}
```

## 2.19 `_Static_assert` Statement

This is a way to prevent *compilation* of a program if a certain constant condition is not met.

```
_Static_assert(__STDC_VERSION__ >= 201112L, "You need at least C11!")
```

## 2.20 Functions

You need to specify the return type and parameter types for the function, and the body goes in a block afterward.

Variables in the function are local to that function.

```
// Function that adds two numbers

int add(int x, int y)
{
    int sum = x + y;

    return sum;
}
```

Functions that return nothing should be return type `void`. Functions that accept no parameters should have `void` as the parameter list.

```
// All side effects, all the time!

void foo(void)
{
    some_global = 12;
    printf("Here we go!\n");
}
```

### 2.20.1 `main()` Function

This is the function that runs when you first start the program. It will be one of these forms:

```
int main(void)
int main(int argc, char *argv[])
```

The first form ignores all command line parameters.

The second form stores the count of the command line parameters in `argc`, and stores the parameters themselves as an array of strings in `argv`. The first of these, `argv[0]`, is typically the name of the executable. The last `argv` pointer has the value `NULL`.

The return values usually show up as exit status codes in the OS. If there is no `return`, falling off the end of `main()` is an implied `return 0`<sup>2</sup>.

---

<sup>2</sup>Note that this implication only for `main()`, and not for any other functions.

### 2.20.2 Variadic Functions

Some functions can take a variable number of arguments. Every function must have at least one argument. The remaining arguments are specified by `...` and can be read with the `va_start()`, `va_arg()`, and `va_end()` macros.

Here's an example that adds up a variable number of integer values.

```
int add(int count, ...)
{
    int total = 0;
    va_list va;

    va_start(va, count); // Start with arguments after "count"

    for (int i = 0; i < count; i++) {
        int n = va_arg(va, int); // Get the next int

        total += n;
    }

    va_end(va); // All done

    return total;
}
```

## Chapter 3

# <assert.h> Runtime and Compile-time Diagnostics

| Macro                        | Description            |
|------------------------------|------------------------|
| <code>assert()</code>        | Runtime assertion      |
| <code>static_assert()</code> | Compile-time assertion |

This functionality has to do with things that Should Never Happen™. If you have something that should never be true and you want your program to bomb out because it happened, this is the header file for you.

There are two types of assertions: compile-time assertions (called “static assertions”) and runtime assertions. If the assertion *fails* (i.e. the thing that you need to be true is not true) then the program will bomb out either at compile-time or runtime.

### 3.1 Macros

If you define the macro `NDEBUG` **before** you include `<assert.h>`, then the `assert()` macro will have no effect. You can define `NDEBUG` to be anything, but `1` seems like a good value.

Since `assert()` causes your program to bomb out at runtime, you might not desire this behavior when you go into production. Defining `NDEBUG` causes `assert()` to be ignored.

`NDEBUG` has no effect on `static_assert()`.

### 3.2 `assert()`

Bomb out at runtime if a condition fails

#### Synopsis

```
#include <assert.h>

void assert(scalar expression);
```

## Description

You pass in an expression to this macro. If it evaluates to false, the program will crash with an assertion failure (by calling the `abort()` function).

Basically, you're saying, "Hey, I'm assuming this condition is true, and if it's not, I don't want to continue running."

This is used while debugging to make sure no unexpected conditions arise. And if you find during development that the condition does arise, maybe you should modify the code to handle it before going to production.

If you've defined the macro `NDEBUG` to any value before `<assert.h>` was included, the `assert()` macro is ignored. This is a good idea before production.

Unlike `static_assert()`, this macro doesn't allow you to print an arbitrary message. If you want to do this, you can roll your own assert as a preprocessor macro:

```
#define ASSERT(c, m) \
do { \
    if (!(c)) { \
        fprintf(stderr, __FILE__ ":%d: assertion %s failed: %s\n", \
            __LINE__, #c, m); \
        exit(1); \
    } \
} while(0)
```

## Return Value

This macro doesn't return (since it calls `abort()` which never returns).

If `NDEBUG` is set, the macro evaluates to `((void)0)`, which does nothing.

## Example

Here's a function that divides the size of our goat herd. But we're assuming we'll never get a `0` passed to us.

So we assert that `amount != 0`... and if it is, the program aborts/

```
1  // #define NDEBUG 1 // uncomment this to disable the assert
2
3  #include <stdio.h>
4  #include <assert.h>
5
6  int goat_count = 10;
7
8  void divide_goat_herd_by(int amount)
9  {
10     assert(amount != 0);
11
12     goat_count /= amount;
13 }
14
15 int main(void)
16 {
17     divide_goat_herd_by(2); // OK
18 }
```



```
19     divide_goat_herd_by(0); // Causes the assert to fire
20 }
```

When I run this and pass 0 to the function, I get the following on my system (the exact output may vary):

```
assert: assert.c:10: divide_goat_herd_by: Assertion `amount != 0' failed.
```

## See Also

`static_assert()`, `abort()`

---

## 3.3 static\_assert()

Bomb out at compile-time if a condition fails

### Synopsis

```
#include <assert.h>

static_assert(constant-expression, string-literal);
```

### Description

This macro prevents your program from even compiling if a condition isn't true.

And it prints the string literal you give it.

Basically if `constant-expression` is false, then compilation will cease and the `string-literal` will be printed.

The constant expression must be truly constant—just values, no variables. And the same is true for the string literal: no variables, just a literal string in double quotes. (It has to be this way since the program's not running at this point.)

### Return Value

Not applicable, as this is a compile-time feature.

### Example

Here's a partial example with an algorithm that presumably has poor performance or memory issues if the size of the local array is too large. We prevent that eventuality at compile-time by catching it with the `static_assert()`.

```
1  #include <stdio.h>
2  #include <assert.h>
3
4  #define ARRAY_SIZE 16
5
6  int main(void)
```

```
7 {  
8     static_assert(ARRAY_SIZE > 32, "ARRAY_SIZE too small");  
9  
10    int a[ARRAY_SIZE];  
11  
12    a[32] = 10;  
13  
14    printf("%d\n", a[32]);  
15 }
```

On my system, when I try to compile it, this prints (your output may vary):

```
In file included from static_assert.c:2:  
static_assert.c: In function 'main':  
static_assert.c:8:5: error: static assertion failed: "ARRAY_SIZE too small"  
    8 |     static_assert(ARRAY_SIZE > 32, "ARRAY_SIZE too small");  
      |     ^~~~~~
```

## See Also

`assert()`

## Chapter 4

# <complex.h> Complex Number Functionality

The complex functions in this reference section come in three flavors each: `double complex`, `float complex`, and `long double complex`.

The `float` variants end with `f` and the `long double` variants end with `l`, e.g. for complex cosine:

```
ccos()    double complex
ccosf()   float complex
ccosl()   long double complex
```

The table below only lists the `double complex` version for brevity.

| Function              | Description                                         |
|-----------------------|-----------------------------------------------------|
| <code>cabs()</code>   | Compute the complex absolute value                  |
| <code>cacos()</code>  | Compute the complex arc-cosine                      |
| <code>cacosh()</code> | Compute the complex arc hyperbolic cosine           |
| <code>carg()</code>   | Compute the complex argument                        |
| <code>casin()</code>  | Compute the complex arc-sine                        |
| <code>casinh()</code> | Compute the complex arc hyperbolic sine             |
| <code>catan()</code>  | Compute the complex arc-tangent                     |
| <code>catanh()</code> | Compute the complex arc hyperbolic tangent          |
| <code>ccos()</code>   | Compute the complex cosine                          |
| <code>ccosh()</code>  | Compute the complex hyperbolic cosine               |
| <code>cexp()</code>   | Compute the complex base- <i>e</i> exponential      |
| <code>cimag()</code>  | Returns the imaginary part of a complex number      |
| <code>clog()</code>   | Compute the complex logarithm                       |
| <code>CMPLX()</code>  | Build a complex value from real and imaginary types |
| <code>conj()</code>   | Compute the conjugate of a complex number           |
| <code>cproj()</code>  | Compute the projection of a complex number          |
| <code>creal()</code>  | Returns the real part of a complex number           |
| <code>csin()</code>   | Compute the complex sine                            |
| <code>csinh()</code>  | Compute the complex hyperbolic sine                 |
| <code>csqrt()</code>  | Compute the complex square root                     |
| <code>ctan()</code>   | Compute the complex tangent                         |
| <code>ctanh()</code>  | Compute the complex hyperbolic tangent              |

You can test for complex number support by looking at the `_STDC_NO_COMPLEX` macro. If it's defined, complex numbers aren't available.

There are possibly two types of numbers defined: *complex* and *imaginary*. No system I'm currently aware of implements imaginary types.

The complex types, which are a real value plus a multiple of *i*, are:

```
float complex
double complex
long double complex
```

The imaginary types, which hold a multiple of *i*, are:

```
float imaginary
double imaginary
long double imaginary
```

The mathematical value  $i = \sqrt{-1}$  is represented by the symbol `_Complex_I` or `_Imaginary_I`, if it exists.

The macro `I` will be preferentially set to `_Imaginary_I` (if it exists), or to `_Complex_I` otherwise.

You can write imaginary literals (if supported) using this notation:

```
double imaginary x = 3.4 * I;
```

You can write complex literals using regular complex notation:

```
double complex x = 1.2 + 3.4 * I;
```

or build them with the `CMPLX()` macro:

```
double complex x = CMPLX(1.2, 3.4); // Like 1.2 + 3.4 * I
```

The latter has the advantage of handling special cases of complex numbers correctly (like those involving infinity or signed zeroes) as if `_Imaginary_I` were present, even if it's not.

All angular values are in radians.

Some functions have discontinuities called *branch cuts*. Now, I'm no mathematician so I can't really talk sensibly about this, but if you're here, I like to think you know what you're doing when it comes to this side of things.

If your system has signed zeroes, you can tell which side of the cut you're on by the sign. And you can't if you don't. The spec elaborates:

Implementations that do not support a signed zero [...] cannot distinguish the sides of branch cuts. These implementations shall map a cut so the function is continuous as the cut is approached coming around the finite endpoint of the cut in a counter clockwise direction. (Branch cuts for the functions specified here have just one finite endpoint.) For example, for the square root function, coming counter clockwise around the finite endpoint of the cut along the negative real axis approaches the cut from above, so the cut maps to the positive imaginary axis.

Finally, there's a pragma called `CX_LIMITED_RANGE` that can be turned on and off (default is off). You can turn it on with:

```
#pragma STDC CX_LIMITED_RANGE ON
```

It allows for certain intermediate operations to underflow, overflow, or deal badly with infinity, presumably for a tradeoff in speed. If you're sure these types of errors won't occur with the numbers you're using AND you're trying to get as much speed out as you can, you could turn this macro on.

The spec also elaborates here:

The purpose of the pragma is to allow the implementation to use the formulas:

$$(x + iy) \times (u + iv) = (xu - yv) + i(yu + xv)$$

$$(x + iy)/(u + iv) = [(xu + yv) + i(yu - xv)]/(u^2 + v^2)$$

$$|x + iy| = \sqrt{x^2 + y^2}$$

where the programmer can determine they are safe.

## 4.1 `cacos()`, `cacosf()`, `cacosl()`

Compute the complex arc-cosine

### Synopsis

```
#include <complex.h>

double complex cacos(double complex z);

float complex cacosf(float complex z);

long double complex cacosl(long double complex z);
```

### Description

Computes the complex arc-cosine of a complex number.

The complex number `z` will have an imaginary component in the range  $[0, \pi]$ , and the real component is unbounded.

There are branch cuts outside the interval  $[-1, +1]$  on the real axis.

### Return Value

Returns the complex arc-cosine of `z`.

### Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
```

```
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = cacos(x);  
9  
10    printf("Result: %f + %fi\n", creal(y), cimag(y));  
11 }
```

Output:

```
Result: 0.195321 + -2.788006i
```

## See Also

`ccos()`, `casin()`, `catan()`

---

## 4.2 `casin()`, `casinf()`, `casinl()`

Compute the complex arc-sine

### Synopsis

```
#include <complex.h>  
  
double complex casin(double complex z);  
  
float complex casinf(float complex z);  
  
long double complex casinl(long double complex z);
```

### Description

Computes the complex arc-sine of a complex number.

The complex number `z` will have an imaginary component in the range  $[-\pi/2, +\pi/2]$ , and the real component is unbounded.

There are branch cuts outside the interval  $[-1, +1]$  on the real axis.

### Return Value

Returns the complex arc-sine of `z`.

### Example

```
1  #include <stdio.h>  
2  #include <complex.h>  
3  
4  int main(void)  
5  {
```

```
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = casin(x);  
9  
10    printf("Result: %f + %fi\n", creal(y), cimag(y));  
11 }
```

Output:

```
Result: 1.375476 + 2.788006i
```

## See Also

`csin()`, `cacos()`, `catan()`

---

## 4.3 `catan()`, `catanf()`, `catanl()`

Compute the complex arc-tangent

### Synopsis

```
#include <complex.h>  
  
double complex catan(double complex z);  
  
float complex catanf(float complex z);  
  
long double complex catanl(long double complex z);
```

### Description

Computes the complex arc-tangent of a complex number.

The complex number `z` will have an real component in the range  $[-\pi/2, +\pi/2]$ , and the imaginary component is unbounded.

There are branch cuts outside the interval  $[-i, +i]$  on the imaginary axis.

### Return Value

Returns the complex arc-tangent of `z`.

### Example

```
1 #include <stdio.h>  
2 #include <complex.h>  
3  
4 int main(void)  
5 {
```

```
6     double wheat = 8;
7     double sheep = 1.5708;
8
9     double complex x = wheat + sheep * I;
10
11    double complex y = catan(x);
12
13    printf("Result: %f + %fi\n", creal(y), cimag(y));
14 }
```

Output:

```
Result: 1.450947 + 0.023299i
```

### See Also

`ctan()`, `cacos()`, `casin()`

---

## 4.4 `ccos()`, `ccosf()`, `ccosl()`

Compute the complex cosine

### Synopsis

```
#include <complex.h>

double complex ccos(double complex z);

float complex ccosf(float complex z);

long double complex ccosl(long double complex z);
```

### Description

Computes the complex cosine of a complex number.

### Return Value

Returns the complex cosine of `z`.

### Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 8 + 1.5708 * I;
```



```
7
8     double complex y = ccos(x);
9
10    printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: -0.365087 + -2.276818i
```

## See Also

`csin()`, `ctan()`, `cacos()`

---

## 4.5 `csin()`, `csinf()`, `csinl()`

Compute the complex sine

### Synopsis

```
#include <complex.h>

double complex csin(double complex z);

float complex csinf(float complex z);

long double complex csinl(long double complex z);
```

### Description

Computes the complex sine of a complex number.

### Return Value

Returns the complex sine of `z`.

### Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = 8 + 1.5708 * I;
7
8     double complex y = csin(x);
9
10    printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: 2.482485 + -0.334840i
```

## See Also

`ccos()`, `ctan()`, `casin()`

---

## 4.6 `ctan()`, `ctanf()`, `ctanl()`

Compute the complex tangent

### Synopsis

```
#include <complex.h>

double complex ctan(double complex z);

float complex ctanf(float complex z);

long double complex ctanl(long double complex z);
```

### Description

Computes the complex tangent of a complex number.

### Return Value

Returns the complex tangent of `z`.

### Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
```

```
5 {  
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = ctan(x);  
9  
10    printf("Result: %f + %fi\n", creal(y), cimag(y));  
11 }
```

Output:

```
Result: -0.027073 + 1.085990i
```

## See Also

`ccos()`, `csin()`, `catan()`

---

## 4.7 `cacosh()`, `cacoshf()`, `cacoshl()`

Compute the complex arc hyperbolic cosine

### Synopsis

```
#include <complex.h>  
  
double complex cacosh(double complex z);  
  
float complex cacoshf(float complex z);  
  
long double complex cacoshl(long double complex z);
```

### Description

Computes the complex arc hyperbolic cosine of a complex number.

There is a branch cut at values less than 1 on the real axis.

The return value will be non-negative on the real number axis, and in the range  $[-i\pi, +i\pi]$  on the imaginary axis.

### Return Value

Returns the complex arc hyperbolic cosine of `z`.

### Example

```
1 #include <stdio.h>  
2 #include <complex.h>  
3  
4 int main(void)
```

```
5 {  
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = cacosh(x);  
9  
10    printf("Result: %f + %fi\n", creal(y), cimag(y));  
11 }
```

Output:

```
Result: 2.788006 + 0.195321i
```

## See Also

`casinh()`, `catanh()`, `acosh()`

---

## 4.8 `casinh()`, `casinhf()`, `casinhl()`

Compute the complex arc hyperbolic sine

### Synopsis

```
#include <complex.h>  
  
double complex casinh(double complex z);  
  
float complex casinhf(float complex z);  
  
long double complex casinhl(long double complex z);
```

### Description

Computes the complex arc hyperbolic sine of a complex number.

There are branch cuts outside  $[-i, +i]$  on the imaginary axis.

The return value will be unbounded on the real number axis, and in the range  $[-i\pi/2, +i\pi/2]$  on the imaginary axis.

### Return Value

Returns the complex arc hyperbolic sine of `z`.

### Example

```
1 #include <stdio.h>  
2 #include <complex.h>  
3  
4 int main(void)
```

```
5 {  
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = casinh(x);  
9  
10    printf("Result: %f + %fi\n", creal(y), cimag(y));  
11 }
```

Output:

```
Result: 2.794970 + 0.192476i
```

## See Also

`cacosh()`, `catanh()`, `asinh()`

---

## 4.9 `catanh()`, `catanhf()`, `catanhl()`

Compute the complex arc hyperbolic tangent

### Synopsis

```
#include <complex.h>  
  
double complex catanh(double complex z);  
  
float complex catanhf(float complex z);  
  
long double complex catanhl(long double complex z);
```

### Description

Computes the complex arc hyperbolic tangent of a complex number.

There are branch cuts outside  $[-1, +1]$  on the real axis.

The return value will be unbounded on the real number axis, and in the range  $[-i\pi/2, +i\pi/2]$  on the imaginary axis.

### Return Value

Returns the complex arc hyperbolic tangent of `z`.

### Example

```
1 #include <stdio.h>  
2 #include <complex.h>  
3  
4 int main(void)
```

```
5 {  
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = catanh(x);  
9  
10    printf("Result: %f + %fi\n", creal(y), cimag(y));  
11 }
```

Output:

```
Result: 0.120877 + 1.546821i
```

## See Also

`cacosh()`, `casinh()`, `atanh()`

---

## 4.10 `ccosh()`, `ccoshf()`, `ccoshl()`

Compute the complex hyperbolic cosine

### Synopsis

```
#include <complex.h>  
  
double complex ccosh(double complex z);  
  
float complex ccoshf(float complex z);  
  
long double complex ccoshl(long double complex z);
```

### Description

Computes the complex hyperbolic cosine of a complex number.

### Return Value

Returns the complex hyperbolic cosine of `z`.

### Example

```
1 #include <stdio.h>  
2 #include <complex.h>  
3  
4 int main(void)  
5 {  
6     double complex x = 8 + 1.5708 * I;  
7  
8     double complex y = ccosh(x);
```

```
9
10     printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: -0.005475 + 1490.478826i
```

## See Also

`csinh()`, `ctanh()`, `ccos()`

---

## 4.11 `csinh()`, `csinhf()`, `csinhl()`

Compute the complex hyperbolic sine

### Synopsis

```
#include <complex.h>

double complex csinh(double complex z);

float complex csinhf(float complex z);

long double complex csinhl(long double complex z);
```

### Description

Computes the complex hyperbolic sine of a complex number.

### Return Value

Returns the complex hyperbolic sine of `z`.

### Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 8 + 1.5708 * I;
7
8      double complex y = csinh(x);
9
10     printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: -0.005475 + 1490.479161i
```

## See Also

`ccosh()`, `ctanh()`, `csin()`

---

## 4.12 `ctanh()`, `ctanhf()`, `ctanhl()`

Compute the complex hyperbolic tangent

### Synopsis

```
#include <complex.h>

double complex ctanh(double complex z);

float complex ctanhf(float complex z);

long double complex ctanhl(long double complex z);
```

### Description

Computes the complex hyperbolic tangent of a complex number.

### Return Value

Returns the complex hyperbolic tangent of `z`.

### Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 8 + 1.5708 * I;
7
8      double complex y = ctanh(x);
9
10     printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: 1.000000 + -0.000000i
```



## See Also

`ccosh()`, `csinh()`, `ctan()`

---

## 4.13 `cexp()`, `cexpf()`, `cexpl()`

Compute the complex base- $e$  exponential

### Synopsis

```
#include <complex.h>

double complex cexp(double complex z);

float complex cexpf(float complex z);

long double complex cexpl(long double complex z);
```

### Description

Computes the complex base- $e$  exponential of `z`.

### Return Value

Returns the complex base- $e$  exponential of `z`.

### Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = 1 + 2 * I;
7
8     double complex y = cexp(x);
9
10    printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: -1.131204 + 2.471727i
```

## See Also

`cpow()`, `clog()`, `exp()`

---

## 4.14 `clog()`, `clogf()`, `clogl()`

Compute the complex logarithm

### Synopsis

```
#include <complex.h>

double complex clog(double complex z);

float complex clogf(float complex z);

long double complex clogl(long double complex z);
```

### Description

Compute the base- $e$  complex logarithm of  $z$ . There is a branch cut on the negative real axis.

The returns value is unbounded on the real axis and in the range  $[-i\pi, +i\pi]$  on the imaginary axis.

### Return Value

Returns the base- $e$  complex logarithm of  $z$ .

### Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 1 + 2 * I;
7
8      double complex y = clog(x);
9
10     printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: 0.804719 + 1.107149i
```

### See Also

`cexp()`, `log()`

---

## 4.15 `cabs()`, `cabsf()`, `cabsl()`

Compute the complex absolute value

## Synopsis

```
#include <complex.h>

double cabs(double complex z);

float cabsf(float complex z);

long double cabsl(long double complex z);
```

## Description

Computes the complex absolute value of `z`.

## Return Value

Returns the complex absolute value of `z`.

## Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = 1 + 2 * I;
7
8     double complex y = cabs(x);
9
10    printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: 2.236068 + 0.000000i
```

## See Also

`fabs()`, `abs()`

---

## 4.16 `cpow()`, `cpowf()`, `cpowl()`

Compute complex power

## Synopsis

```
#include <complex.h>

double complex cpow(double complex x, double complex y);

float complex cpowf(float complex x, float complex y);

long double complex cpowl(long double complex x,
                           long double complex y);
```

## Description

Computes the complex  $x^y$ .

There is a branch cut for  $x$  along the negative real axis.

## Return Value

Returns the complex  $x^y$ .

## Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = 1 + 2 * I;
7     double complex y = 3 + 4 * I;
8
9     double r = cpow(x, y);
10
11     printf("Result: %f + %fi\n", creal(r), cimag(r));
12 }
```

Result:

```
Result: 0.129010 + 0.000000i
```

## See Also

`csqrt()`, `cexp()`

---

## 4.17 `csqrt()`, `csqrtf()`, `csqrtl()`

Compute the complex square root

## Synopsis

```
#include <complex.h>

double complex csqrt(double complex z);

float complex csqrtf(float complex z);

long double complex csqrtl(long double complex z);
```

## Description

Computes the complex square root of `z`.

There is a branch cut along the negative real axis.

The return value is in the right half of the complex plane and includes the imaginary axis.

## Return Value

Returns the complex square root of `z`.

## Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 1 + 2 * I;
7
8      double complex y = csqrt(x);
9
10     printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: 1.272020 + 0.786151i
```

## See Also

`cpow()`, `sqrt()`

---

## 4.18 `carg()`, `cargf()`, `cargl()`

Compute the complex argument

## Synopsis

```
#include <complex.h>

double carg(double complex z);

float cargf(float complex z);

long double cargl(long double complex z);
```

## Description

Computes the complex argument (AKA phase angle) of `z`.

There is a branch cut along the negative real axis.

Returns a value in the range  $[-\pi, +\pi]$ .

## Return Value

Returns the complex argument of `z`.

## Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 1 + 2 * I;
7
8      double y = carg(x);
9
10     printf("Result: %f\n", y);
11 }
```

Output:

```
Result: 1.107149
```

---

## 4.19 `cimag()`, `cimagf()`, `cimagl()`

Returns the imaginary part of a complex number

## Synopsis

```
#include <complex.h>

double cimag(double complex z);
```

```
float cimagf(float complex z);

long double cimagl(long double complex z);
```

## Description

Returns the imaginary part of `z`.

As a footnote, the spec points out that any complex number `x` is part of the following equivalency:

```
x == creal(x) + cimag(x) * I;
```

## Return Value

Returns the imaginary part of `z`.

## Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = 1 + 2 * I;
7
8     double y = cimag(x);
9
10    printf("Result: %f\n", y);
11 }
```

Output—just the imaginary part:

```
Result: 2.000000
```

## See Also

`creal()`

## 4.20 Cmplx(), Cmplxf(), Cmplxl()

Build a complex value from real and imaginary types

## Synopsis

```
#include <complex.h>

double complex Cmplx(double x, double y);
```

```
float complex CMPLXF(float x, float y);

long double complex CMPLXL(long double x, long double y);
```

## Description

These macros build a complex value from real and imaginary types.

Now I know what you’re thinking. “But I can already build a complex value from real and imaginary types using the `I` macro, like in the example you’re about to give us.”

```
double complex x = 1 + 2 * I;
```

And that’s true.

But the reality of the matter is weird and complex.

Maybe `I` got undefined, or maybe you redefined it.

Or maybe `I` was defined as `_Complex_I` which doesn’t necessarily preserve the sign of a zero value.

As the spec points out, these macros build complex numbers as if `_Imaginary_I` were defined (thus preserving your zero sign) even if it’s not. That is, they are defined equivalently to:

```
#define CMPLX(x, y) ((double complex)((double)(x) + \
    _Imaginary_I * (double)(y)))

#define CMPLXF(x, y) ((float complex)((float)(x) + \
    _Imaginary_I * (float)(y)))

#define CMPLXL(x, y) ((long double complex)((long double)(x) + \
    _Imaginary_I * (long double)(y)))
```

## Return Value

Returns the complex number for the given real `x` and imaginary `y` components.

## Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = CMPLX(1, 2); // Like 1 + 2 * I
7
8     printf("Result: %f + %fi\n", creal(x), cimag(x));
9 }
```

Output:

```
Result: 1.000000 + 2.000000i
```



## See Also

`creal()`, `cimag()`

---

## 4.21 `conj()`, `conjf()`, `conjl()`

Compute the conjugate of a complex number

### Synopsis

```
#include <complex.h>

double complex conj(double complex z);

float complex conjf(float complex z);

long double complex conjl(long double complex z);
```

### Description

This function computes the complex conjugate<sup>1</sup> of `z`. Apparently it does this by reversing the sign of the imaginary part, but dammit, I'm a programmer not a mathematician, Jim!

### Return Value

Returns the complex conjugate of `z`

### Example

```
1 #include <stdio.h>
2 #include <complex.h>
3
4 int main(void)
5 {
6     double complex x = 1 + 2 * I;
7
8     double complex y = conj(x);
9
10    printf("Result: %f + %fi\n", creal(y), cimag(y));
11 }
```

Output:

```
Result: 1.000000 + -2.000000i
```

---

<sup>1</sup>[https://en.wikipedia.org/wiki/Complex\\_conjugate](https://en.wikipedia.org/wiki/Complex_conjugate)

## 4.22 `cproj()`, `cprojf()`, `cprojl()`

Compute the projection of a complex number

### Synopsis

```
#include <complex.h>

double complex cproj(double complex z);

float complex cprojf(float complex z);

long double complex cprojl(long double complex z);
```

### Description

Computes the projection of `z` onto a Riemann sphere<sup>2</sup>.

Now we're *really* outside my expertise. The spec has this to say, which I'm quoting verbatim because I'm not knowledgeable enough to rewrite it sensibly. Hopefully it makes sense to anyone who would need to use this function.

`z` projects to `z` except that all complex infinities (even those with one infinite part and one NaN part) project to positive infinity on the real axis. If `z` has an infinite part, then `cproj(z)` is equivalent to

```
INFINITY + I * copysign(0.0, cimag(z))
```

So there you have it.

### Return Value

Returns the projection of `z` onto a Riemann sphere.

### Example

Fingers crossed this is a remotely sane example...

```
1  #include <stdio.h>
2  #include <complex.h>
3  #include <math.h>
4
5  int main(void)
6  {
7      double complex x = 1 + 2 * I;
8
9      double complex y = cproj(x);
10
11     printf("Result: %f + %fi\n", creal(y), cimag(y));
12
13     x = INFINITY + 2 * I;
14     y = cproj(x);
```

<sup>2</sup>[https://en.wikipedia.org/wiki/Riemann\\_sphere](https://en.wikipedia.org/wiki/Riemann_sphere)

```
15
16     printf("Result: %f + %fi\n", creal(y), cimag(y));
17 }
```

Output:

```
Result: 1.000000 + 2.000000i
Result: inf + 0.000000i
```

## 4.23 creal(), crealf(), creall()

Returns the real part of a complex number

### Synopsis

```
#include <complex.h>

double creal(double complex z);

float crealf(float complex z);

long double creall(long double complex z);
```

### Description

Returns the real part of `z`.

As a footnote, the spec points out that any complex number `x` is part of the following equivalency:

```
x == creal(x) + cimag(x) * I;
```

### Return Value

Returns the real part of `z`.

### Example

```
1  #include <stdio.h>
2  #include <complex.h>
3
4  int main(void)
5  {
6      double complex x = 1 + 2 * I;
7
8      double y = creal(x);
9
10     printf("Result: %f\n", y);
11 }
```

Output—just the real part:

```
Result: 1.000000
```

## See Also

`cimag()`

## Chapter 5

# <ctype.h> Character Classification and Conversion

| Function                | Description                                         |
|-------------------------|-----------------------------------------------------|
| <code>isalnum()</code>  | Tests if a character is alphabetic or is a digit    |
| <code>isalpha()</code>  | Returns true if a character is alphabetic           |
| <code>isblank()</code>  | Tests if a character is word-separating whitespace  |
| <code>iscntrl()</code>  | Test if a character is a control character          |
| <code>isdigit()</code>  | Tests if a character is a digit                     |
| <code>isgraph()</code>  | Tests if the character is printable and not a space |
| <code>islower()</code>  | Tests if a character is lowercase                   |
| <code>isprint()</code>  | Tests if a character is printable                   |
| <code>ispunct()</code>  | Test if a character is punctuation                  |
| <code>isspace()</code>  | Test if a character is whitespace                   |
| <code>isupper()</code>  | Tests if a character is uppercase                   |
| <code>isxdigit()</code> | Tests if a character is a hexadecimal digit         |
| <code>tolower()</code>  | Convert a letter to lowercase                       |
| <code>toupper()</code>  | Convert a letter to uppercase                       |

This collection of macros is good for testing characters to see if they're of a certain class, such as alphabetic, numeric, control characters, etc.

Surprisingly, they take `int` arguments instead of some kind of `char`. This is so you can feed `EOF` in for convenience if you have an integer representation of that. If not `EOF`, the value passed in has to be representable in an `unsigned char`. Otherwise it's (dun dun DUUNNNN) undefined behavior. So you can forget about passing in your UTF-8 multibyte characters.

You can portably avoid this undefined behavior by casting the arguments to these functions to `(unsigned char)`. This is irksome and ugly, admittedly. The values in the basic character set are all safe to use since they're positive values that fit into an `unsigned char`.

Also, the behavior of these functions varies based on locale.

In many of the pages in this section, I give some examples. These are from the "C" locale, and might vary if you've set a different locale.

Note that wide characters have their own set of classification functions, so don't try to use these on `wchar_ts`. Or else!

## 5.1 isalnum()

Tests if a character is alphabetic or is a digit

### Synopsis

```
#include <ctype.h>

int isalnum(int c);
```

### Description

Tests if a character is alphabetic (A-Z or a-z) or a digit (0-9).

Is equivalent to:

```
isalpha(c) || isdigit(c)
```

### Return Value

Returns true if a character is alphabetic (A-Z or a-z) or a digit (0-9).

### Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isalnum('a')? "yes": "no"); // yes
9      printf("%s\n", isalnum('B')? "yes": "no"); // yes
10     printf("%s\n", isalnum('5')? "yes": "no"); // yes
11     printf("%s\n", isalnum('?')? "yes": "no"); // no
12 }
```

### See Also

isalpha(), isdigit()

---

## 5.2 isalpha()

Returns true if a character is alphabetic

## Synopsis

```
#include <ctype.h>

int isalpha(int c);
```

## Description

Returns true for alphabetic characters (A-Z or a-z).

Technically (and in the “C” locale) equivalent to:

```
isupper(c) || islower(c)
```

Extra super technically, because I know you’re dying for this to be extra unnecessarily complex, it can also include some locale-specific characters for which this is true:

```
!isctrl(c) && !isdigit(c) && !ispunct(c) && !isspace(c)
```

and this is true:

```
isupper(c) || islower(c)
```

## Return Value

Returns true for alphabetic characters (A-Z or a-z).

Or for any of the other crazy stuff in the description, above.

## Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isalpha('a')? "yes": "no"); // yes
9      printf("%s\n", isalpha('B')? "yes": "no"); // yes
10     printf("%s\n", isalpha('5')? "yes": "no"); // no
11     printf("%s\n", isalpha('?')? "yes": "no"); // no
12 }
```

## See Also

`isalnum()`

---

## 5.3 isblank()

Tests if a character is word-separating whitespace

### Synopsis

```
#include <ctype.h>

int isblank(int c);
```

### Description

True if the character is a whitespace character used to separate words in a single line.

For example, space (' ') or horizontal tab ('\t'). Other locales might define other blank characters.

### Return Value

Returns true if the character is a whitespace character used to separate words in a single line.

### Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isblank(' ')? "yes": "no"); // yes
9      printf("%s\n", isblank('\t')? "yes": "no"); // yes
10     printf("%s\n", isblank('\n')? "yes": "no"); // no
11     printf("%s\n", isblank('a')? "yes": "no"); // no
12     printf("%s\n", isblank('?')? "yes": "no"); // no
13 }
```

### See Also

isspace()

---

## 5.4 iscntrl()

Test if a character is a control character



## Synopsis

```
#include <ctype.h>

int iscntrl(int c);
```

## Description

A *control character* is a locale-specific non-printing character.

For the “C” locale, this means control characters are in the range 0x00 to 0x1F (the character right before SPACE) and 0x7F (the DEL character).

Basically if it’s not an ASCII (or Unicode less than 128) printable character, it’s a control character in the “C” locale.

## Return Value

Returns true if `c` is a control character.

## Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", iscntrl('\t')? "yes": "no"); // yes (tab)
9      printf("%s\n", iscntrl('\n')? "yes": "no"); // yes (newline)
10     printf("%s\n", iscntrl('\r')? "yes": "no"); // yes (return)
11     printf("%s\n", iscntrl('\a')? "yes": "no"); // yes (bell)
12     printf("%s\n", iscntrl(' ')? "yes": "no"); // no
13     printf("%s\n", iscntrl('a')? "yes": "no"); // no
14     printf("%s\n", iscntrl('?')? "yes": "no"); // no
15 }
```

## See Also

`isgraph()`, `isprint()`

---

## 5.5 `isdigit()`

Tests if a character is a digit

## Synopsis

```
#include <ctype.h>

int isdigit(int c);
```

## Description

Tests if `c` is a digit in the range 0-9.

## Return Value

Returns true if the character is a digit, unsurprisingly.

## Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isdigit('0')? "yes": "no");    // yes
9      printf("%s\n", isdigit('5')? "yes": "no");    // yes
10     printf("%s\n", isdigit('a')? "yes": "no");    // no
11     printf("%s\n", isdigit('B')? "yes": "no");    // no
12     printf("%s\n", isdigit('?')? "yes": "no");    // no
13 }
```

## See Also

`isalnum()`, `isxdigit()`

---

## 5.6 `isgraph()`

Tests if the character is printable and not a space

## Synopsis

```
#include <ctype.h>

int isgraph(int c);
```

## Description

Tests if `c` is any printable character that isn't a space (' ').

## Return Value

Returns true if `c` is any printable character that isn't a space (' ').

## Example

```

1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isgraph('0')? "yes": "no"); // yes
9      printf("%s\n", isgraph('a')? "yes": "no"); // yes
10     printf("%s\n", isgraph('B')? "yes": "no"); // yes
11     printf("%s\n", isgraph('?')? "yes": "no"); // yes
12     printf("%s\n", isgraph(' ')? "yes": "no"); // no
13     printf("%s\n", isgraph('\n')? "yes": "no"); // no
14 }
```

## See Also

`isctrnl()`, `isprint()`

---

## 5.7 `islower()`

Tests if a character is lowercase

### Synopsis

```

#include <ctype.h>

int islower(int c);
```

### Description

Tests if a character is lowercase, in the range `a-z`.

In other locales, there could be other lowercase characters. In all cases, to be lowercase, the following must be true:

```
!isctrnl(c) && !isdigit(c) && !ispunct(c) && !isspace(c)
```

## Return Value

Returns true if the character is lowercase.

**Example**

```

1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", islower('c')? "yes": "no"); // yes
9      printf("%s\n", islower('0')? "yes": "no"); // no
10     printf("%s\n", islower('B')? "yes": "no"); // no
11     printf("%s\n", islower('?')? "yes": "no"); // no
12     printf("%s\n", islower(' ')? "yes": "no"); // no
13 }

```

**See Also**

isupper(), isalpha(), toupper(), tolower()

---

**5.8 isprint()**

Tests if a character is printable

**Synopsis**

```

#include <ctype.h>

int isprint(int c);

```

**Description**

Tests if a character is printable, including space (' '). So like `isgraph()`, except space isn't left out in the cold.

**Return Value**

Returns true if the character is printable, including space (' ').

**Example**

```

1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isprint('c')? "yes": "no"); // yes

```

```

9     printf("%s\n", isprint('0')? "yes": "no");    // yes
10    printf("%s\n", isprint(' ')? "yes": "no");    // yes
11    printf("%s\n", isprint('\r')? "yes": "no");    // no
12 }
```

**See Also**

`isgraph()`, `iscntrl()`

---

**5.9 ispunct()**

Test if a character is punctuation

**Synopsis**

```

#include <ctype.h>

int ispunct(int c);
```

**Description**

Tests if a character is punctuation.

In the “C” locale, this means:

```
!isspace(c) && !isalnum(c)
```

In other locales, there could be other punctuation characters (but they also can’t be space or alphanumeric).

**Return Value**

True if the character is punctuation.

**Example**

```

1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", ispunct(',')? "yes": "no");    // yes
9      printf("%s\n", ispunct('!')? "yes": "no");    // yes
10     printf("%s\n", ispunct('c')? "yes": "no");    // no
11     printf("%s\n", ispunct('0')? "yes": "no");    // no
12     printf("%s\n", ispunct(' ')? "yes": "no");    // no
13     printf("%s\n", ispunct('\n')? "yes": "no");    // no
14 }
```

## See Also

`isspace()`, `isalnum()`

---

## 5.10 `isspace()`

Test if a character is whitespace

### Synopsis

```
#include <ctype.h>

int isspace(int c);
```

### Description

Tests if `c` is a whitespace character. These are:

- Space (' ')
- Formfeed ('\f')
- Newline ('\n')
- Carriage Return ('\r')
- Horizontal Tab ('\t')
- Vertical Tab ('\v')

Other locales might specify other whitespace characters. `isalnum()` is false for all whitespace characters.

### Return Value

True if the character is whitespace.

### Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isspace(' ')? "yes": "no"); // yes
9      printf("%s\n", isspace('\n'? "yes": "no"); // yes
10     printf("%s\n", isspace('\t'? "yes": "no"); // yes
11     printf("%s\n", isspace(',')? "yes": "no"); // no
12     printf("%s\n", isspace('!'? "yes": "no"); // no
13     printf("%s\n", isspace('c'? "yes": "no"); // no
14 }
```

## See Also

`isblank()`

---

## 5.11 `isupper()`

Tests if a character is uppercase

### Synopsis

```
#include <ctype.h>

int isupper(int c);
```

### Description

Tests if a character is uppercase, in the range A-Z.

In other locales, there could be other uppercase characters. In all cases, to be uppercase, the following must be true:

```
!isctrl(c) && !isdigit(c) && !ispunct(c) && !isspace(c)
```

### Return Value

Returns true if the character is uppercase.

### Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isupper('B')? "yes": "no"); // yes
9      printf("%s\n", isupper('c')? "yes": "no"); // no
10     printf("%s\n", isupper('0')? "yes": "no"); // no
11     printf("%s\n", isupper('?')? "yes": "no"); // no
12     printf("%s\n", isupper(' ')? "yes": "no"); // no
13 }
```

## See Also

`islower()`, `isalpha()`, `toupper()`, `tolower()`

---

## 5.12 isxdigit()

Tests if a character is a hexadecimal digit

### Synopsis

```
#include <ctype.h>

int isxdigit(int c);
```

### Description

Returns true if the character is a hexadecimal digit. Namely if it's 0-9, a-f, or A-F.

### Return Value

True if the character is a hexadecimal digit.

### Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      printf("%s\n", isxdigit('B')? "yes": "no");    // yes
9      printf("%s\n", isxdigit('c')? "yes": "no");    // yes
10     printf("%s\n", isxdigit('2')? "yes": "no");    // yes
11     printf("%s\n", isxdigit('G')? "yes": "no");    // no
12     printf("%s\n", isxdigit('?')? "yes": "no");    // no
13 }
```

### See Also

isdigit()

---

## 5.13 tolower()

Convert a letter to lowercase

### Synopsis

```
#include <ctype.h>

int tolower(int c);
```



## Description

If the character is uppercase (i.e. `isupper(c)` is true), this function returns the corresponding lowercase letter.

Different locales might have different upper- and lowercase letters.

## Return Value

Returns the lowercase value for an uppercase letter. If the letter isn't uppercase, returns it unchanged.

## Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          changing this char
7      //          v
8      printf("%c\n", tolower('B')); // b (made lowercase!)
9      printf("%c\n", tolower('e')); // e (unchanged)
10     printf("%c\n", tolower('!')); // ! (unchanged)
11 }
```

## See Also

`toupper()`, `islower()`, `isupper()`

---

## 5.14 `toupper()`

Convert a letter to uppercase

## Synopsis

```
#include <ctype.h>

int toupper(int c);
```

## Description

If the character is lower (i.e. `islower(c)` is true), this function returns the corresponding uppercase letter.

Different locales might have different upper- and lowercase letters.

## Return Value

Returns the uppercase value for a lowercase letter. If the letter isn't lowercase, returns it unchanged.

## Example

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int main(void)
5  {
6      //          changing this char
7      //          v
8      printf("%c\n", toupper('B')); // B (unchanged)
9      printf("%c\n", toupper('e')); // E (made uppercase!)
10     printf("%c\n", toupper('!')); // ! (unchanged)
11 }
```

## See Also

`tolower()`, `islower()`, `isupper()`

## Chapter 6

# <errno.h> Error Information

| Variable           | Description                             |
|--------------------|-----------------------------------------|
| <code>errno</code> | Holds the error status of the last call |

This header defines a single variable<sup>1</sup>, `errno`, that can be checked to see if an error has occurred.

`errno` is set to 0 on startup, but no library function sets it to 0. If you're going to use solely it to check for errors, set it to 0 before the call and then check it after. Not only that, but if there's no error, all library functions will leave the value of `errno` unchanged.

Often, though, you'll get some error indication from the function you're calling then check `errno` to see what went wrong.

This is commonly used in conjunction with `perror()` to get a human-readable error message that corresponds to the specific error.

**Important Safety Tip:** You should never make your own variable called `errno`—that's undefined behavior.

Note that the C Spec defines less than a handful of values `errno` can take on. Unix defines a bunch more<sup>2</sup>, as does Windows<sup>3</sup>.

### 6.1 `errno`

Holds the error status of the last call

#### Synopsis

```
errno    // Type is undefined, but it's assignable
```

<sup>1</sup>Really it's just required to be a modifiable lvalue, so not necessarily a variable. But you can treat it as such.

<sup>2</sup><https://man.archlinux.org/man/errno.3.en>

<sup>3</sup><https://docs.microsoft.com/en-us/cpp/c-runtime-library/errno-constants?view=msvc-160>

## Description

Indicates the error status of the last call (note that not all calls will set this value).

| Value  | Description                                |
|--------|--------------------------------------------|
| 0      | No error                                   |
| EDOM   | Domain error (from math)                   |
| EILSEQ | Encoding error (from character conversion) |
| ERANGE | Range error (from math)                    |

If you're doing a number of math functions, you might come across `EDOM` or `ERANGE`.

With multibyte/wide character conversion functions, you might see `EILSEQ`.

And your system might define any other number of values that `errno` could be set to, all of which will begin with the letter `E`.

Fun Fact: you can use `EDOM`, `EILSEQ`, and `ERANGE` with preprocessor directives such as `#ifdef`. But, frankly, I'm not sure why you'd do that other than to test their existence.

## Example

The following prints an error message, since passing `2.0` to `acos()` is outside the function's domain.

```

1  #include <stdio.h>
2  #include <math.h>
3  #include <errno.h>
4
5  int main(void)
6  {
7      double x;
8
9      errno = 0;          // Make sure this is clear before the call
10
11     x = acos(2.0);      // Invalid argument to acos()
12
13     if (errno == EDOM)
14         perror("acos");
15     else
16         printf("Answer is %f\n", x);
17
18     return 0;
19 }
```

Output:

```
acos: Numerical argument out of domain
```

The following prints an error message (on my system), since passing `1e+30` to `exp()` produces a result that's outside the range of a double.

```

1  #include <stdio.h>
2  #include <math.h>
3  #include <errno.h>
```

```

4
5 int main(void)
6 {
7     double x;
8
9     errno = 0;          // Make sure this is clear before the call
10
11     x = exp(1e+30); // Pass in some too-huge number
12
13     if (errno == ERANGE)
14         perror("exp");
15     else
16         printf("Answer is %f\n", x);
17
18     return 0;
19 }

```

Output:

```
exp: Numerical result out of range
```

This example tries to convert an invalid character into a wide character, failing. This sets `errno` to `EILSEQ`. We then use `perror()` to print an error message.

```

1 #include <stdio.h>
2 #include <string.h>
3 #include <wchar.h>
4 #include <errno.h>
5 #include <locale.h>
6
7 int main(void)
8 {
9     setlocale(LC_ALL, "");
10
11     char *bad_str = "\xff"; // Probably invalid char in this locale
12     wchar_t wc;
13     size_t result;
14     mbstate_t ps;
15
16     memset(&ps, 0, sizeof ps);
17
18     result = mbrtowc(&wc, bad_str, 1, &ps);
19
20     if (result == (size_t)(-1))
21         perror("mbrtowc"); // mbrtowc: Illegal byte sequence
22     else
23         printf("Converted to L'%lc'\n", wc);
24
25     return 0;
26 }

```

Output:

`mbrtowc`: Invalid or incomplete multibyte or wide character

**See Also**

`perror()`, `mbrtoc16()`, `c16rtomb()`, `mbrtoc32()`, `c32rtomb()`, `fgetwc()`, `fputwc()`, `mbrtowc()`, `wcrtomb()`, `mbsrtowcs()`, `wcsrtombs()`, `<math.h>`,

## Chapter 7

# <fenv.h> Floating Point Exceptions and Environment

| Function                       | Description                                                    |
|--------------------------------|----------------------------------------------------------------|
| <code>feclearexcept()</code>   | Clear floating point exceptions                                |
| <code>fegetexceptflag()</code> | Save the floating point exception flags                        |
| <code>fesetexceptflag()</code> | Restore the floating point exception flags                     |
| <code>feraiseexcept()</code>   | Raise a floating point exception through software              |
| <code>fetestexcept()</code>    | Test to see if an exception has occurred                       |
| <code>fegetround()</code>      | Get the rounding direction                                     |
| <code>fesetround()</code>      | Set the rounding direction                                     |
| <code>fegetenv()</code>        | Save the entire floating point environment                     |
| <code>fesetenv()</code>        | Restore the entire floating point environment                  |
| <code>feholdexcept()</code>    | Save floating point state and install non-stop mode            |
| <code>feupdateenv()</code>     | Restore floating point environment and apply recent exceptions |

### 7.1 Types and Macros

There are two types defined in this header:

| Type                   | Description                           |
|------------------------|---------------------------------------|
| <code>fenv_t</code>    | The entire floating point environment |
| <code>fexcept_t</code> | A set of floating point exceptions    |

The “environment” can be thought of as the status at this moment of the floating point processing system: this includes the exceptions, rounding, etc. It’s an opaque type, so you won’t be able to access it directly, and it must be done through the proper functions.

If the functions in question exist on your system (they might not be!), then you’ll also have these macros defined to represent different exceptions:

| Macro                     | Description                       |
|---------------------------|-----------------------------------|
| <code>FE_DIVBYZERO</code> | Division by zero                  |
| <code>FE_INEXACT</code>   | Result was not exact, was rounded |

| Macro         | Description               |
|---------------|---------------------------|
| FE_INVALID    | Domain error              |
| FE_OVERFLOW   | Numeric overflow          |
| FE_UNDERFLOW  | Numeric underflow         |
| FE_ALL_EXCEPT | All of the above combined |

The idea is that you can bitwise-OR these together to represent multiple exceptions, e.g. `FE_INVALID | FE_OVERFLOW`.

The functions, below, that have an `excepts` parameter will take these values.

See <math.h> for which functions raise which exceptions and when.

## 7.2 Pragmas

Normally C is free to optimize all kinds of stuff that might cause the flags to not look like you might expect. So if you're going to use this stuff, be sure to set this pragma:

```
#pragma STDC FENV_ACCESS ON
```

If you do this at global scope, it remains in effect until you turn it off:

```
#pragma STDC FENV_ACCESS OFF
```

If you do it in block scope, it has to come before any statements or declarations. In this case, it has effect until the block ends (or until it is explicitly turned off.)

**A caveat:** this program isn't supported on either of the compilers I have (gcc and clang) as of this writing, so though I have built the code, below, it's not particularly well-tested.

## 7.3 `feclearexcept()`

Clear floating point exceptions

### Synopsis

```
#include <fenv.h>

int feclearexcept(int excepts);
```

### Description

If a floating point exception has occurred, this function can clear it.

Set `excepts` to a bitwise-OR list of exceptions to clear.

Passing `0` has no effect.

### Return Value

Returns `0` on success and non-zero on failure.



## Example

```
1 #include <stdio.h>
2 #include <math.h>
3 #include <fenv.h>
4
5 int main(void)
6 {
7     #pragma STDC FENV_ACCESS ON
8
9     double f = sqrt(-1);
10
11     int r = fectestexcept(FE_INVALID);
12
13     printf("%d %f\n", r, f);
14 }
```

## See Also

`feraiseexcept()`, `fetestexcept()`

---

## 7.4 fegetexceptflag() fesetexceptflag()

Save or restore the floating point exception flags

### Synopsis

```
#include <fenv.h>

int fegetexceptflag(fexcept_t *flagp, int excepts);

int fesetexceptflag(fexcept_t *flagp, int excepts);
```

### Description

Use these functions to save or restore the current floating point environment in a variable.

Set `excepts` to the set of exceptions you want to save or restore the state of. Setting it to `FE_ALL_EXCEPT` will save or restore the entire state.

Note that `fexcept_t` is an opaque type—you don't know what's in it.

`excepts` can be set to zero for no effect.

### Return Value

Returns 0 on success or if `excepts` is zero.

Returns non-zero on failure.

### Example

This program saves the state (before any error has happened), then deliberately causes a domain error by trying to take  $\sqrt{-1}$ .

After that, it restores the floating point state to before the error had occurred, thereby clearing it.

```

1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  int main(void)
6  {
7      #pragma STDC FENV_ACCESS ON
8
9      fexcept_t flag;
10
11     fegetexceptflag(&flag, FE_ALL_EXCEPT); // Save state
12
13     double f = sqrt(-1);                      // I imagine this won't work
14     printf("%f\n", f);                        // "nan"
15
16     if (fetetestexcept(FE_INVALID))
17         printf("1: Domain error\n");          // This prints!
18     else
19         printf("1: No domain error\n");
20
21     fesetexceptflag(&flag, FE_ALL_EXCEPT); // Restore to before error
22
23     if (fetetestexcept(FE_INVALID))
24         printf("2: Domain error\n");
25     else
26         printf("2: No domain error\n");      // This prints!
27 }
```

## 7.5 `feraiseexcept()`

Raise a floating point exception through software

### Synopsis

```

#include <fenv.h>

int feraiseexcept(int excepts);
```

### Description

This attempts to raise a floating point exception as if it had happened.

You can specify multiple exceptions to raise.

If either `FE_UNDERFLOW` or `FE_OVERFLOW` is raised, C *might* also raise `FE_INEXACT`.

If either `FE_UNDERFLOW` or `FE_OVERFLOW` is raised at the same time as `FE_INEXACT`, then `FE_UNDERFLOW` or `FE_OVERFLOW` will be raised *before* `FE_INEXACT` behind the scenes.

The order the other exceptions are raised is undefined.

## Return Value

Returns `0` if all the exceptions were raised or if `excepts` is `0`.

Returns non-zero otherwise.

## Example

This code deliberately raises a division-by-zero exception and then detects it.

```
1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  int main(void)
6  {
7      #pragma STDC FENV_ACCESS ON
8
9      feraisexcept(FE_DIVBYZERO);
10
11     if (fetestexcept(FE_DIVBYZERO) == FE_DIVBYZERO)
12         printf("Detected division by zero\n"); // This prints!!
13     else
14         printf("This is fine.\n");
15 }
```

## See Also

`feclearexcept()`, `fetestexcept()`

---

## 7.6 `fetestexcept()`

Test to see if an exception has occurred

### Synopsis

```
#include <fenv.h>

int fetestexcept(int excepts);
```

### Description

Put the exceptions you want to test in `excepts`, bitwise-ORing them together.

## Return Value

Returns the bitwise-OR of the exceptions that have been raised.

## Example

This code deliberately raises a division-by-zero exception and then detects it.

```

1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  int main(void)
6  {
7      #pragma STDC FENV_ACCESS ON
8
9      feraisexcept(FE_DIVBYZERO);
10
11     if (fetetestexcept(FE_DIVBYZERO) == FE_DIVBYZERO)
12         printf("Detected division by zero\n"); // This prints!!
13     else
14         printf("This is fine.\n");
15 }
```

## See Also

`feclearexcept()`, `feraisexcept()`

## 7.7 fegetround() fesetround()

Get or set the rounding direction

### Synopsis

```

#include <fenv.h>

int fegetround(void);

int fesetround(int round);
```

## Description

Use these to get or set the rounding direction used by a variety of math functions.

Basically when a function “rounds” a number, it wants to know how to do it. By default, it does it how we tend to expect: if the fractional part is less than 0.5, it rounds down closer to zero, otherwise up farther from zero.

| Macro                     | Description                                    |
|---------------------------|------------------------------------------------|
| <code>FE_TONEAREST</code> | Round to the nearest whole number, the default |

| Macro         | Description                                |
|---------------|--------------------------------------------|
| FE_TOWARDZERO | Round toward zero always                   |
| FE_DOWNWARD   | Round toward the next lesser whole number  |
| FE_UPWARD     | Round toward the next greater whole number |

Some implementations don't support rounding. If it does, the above macros will be defined.

Note that the `round()` function is always “to-nearest” and doesn't pay attention to the rounding mode.

## Return Value

`fegetround()` returns the current rounding direction, or a negative value on error.

`fesetround()` returns zero on success, or non-zero on failure.

## Example

This rounds some numbers

```

1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  // Helper function to print the rounding mode
6  const char *rounding_mode_str(int mode)
7  {
8      switch (mode) {
9          case FE_TONEAREST: return "FE_TONEAREST";
10         case FE_TOWARDZERO: return "FE_TOWARDZERO";
11         case FE_DOWNWARD:   return "FE_DOWNWARD";
12         case FE_UPWARD:     return "FE_UPWARD";
13     }
14
15     return "Unknown";
16 }
17
18 int main(void)
19 {
20     #pragma STDC FENV_ACCESS ON
21
22     int rm;
23
24     rm = fegetround();
25
26     printf("%s\n", rounding_mode_str(rm)); // Print current mode
27     printf("%f %f\n", rint(2.1), rint(2.7)); // Try rounding
28
29     fesetround(FE_TOWARDZERO); // Set the mode
30
31     rm = fegetround();
32
33     printf("%s\n", rounding_mode_str(rm)); // Print it
34     printf("%f %f\n", rint(2.1), rint(2.7)); // Try it now!

```

35

}

Output:

```
FE_TONEAREST
2.000000 3.000000
FE_TOWARDZERO
2.000000 2.000000
```

**See Also**

nearbyint(), nearbyintf(), nearbyintl(), rint(), rintf(), rintl(), lrint(), lrintf(),  
lrintl(), llrint(), llrintf(), llrintl()

**7.8 fegetenv() fesetenv()**

Save or restore the entire floating point environment

**Synopsis**

```
#include <fenv.h>

int fegetenv(fenv_t *envp);
int fesetenv(const fenv_t *envp);
```

**Description**

You can save the environment (exceptions, rounding direction, etc.) by calling `fegetenv()` and restore it with `fesetenv()`.

Use this if you want to restore the state after a function call, i.e. hide from the caller that some floating point exceptions or changes occurred.

**Return Value**

`fegetenv()` and `fesetenv()` return 0 on success, and non-zero otherwise.

**Example**

This example saves the environment, messes with the rounding and exceptions, then restores it. After the environment is restored, we see that the rounding is back to default and the exception is cleared.

```
1 #include <stdio.h>
2 #include <math.h>
3 #include <fenv.h>
4
5 void show_status(void)
6 {
7     printf("Rounding is FE_TOWARDZERO: %d\n",
```

```

8         fegetround() == FE_TOWARDZERO);
9
10        printf("FE_DIVBYZERO is set: %d\n",
11               fetestexcept(FE_DIVBYZERO) != 0);
12    }
13
14    int main(void)
15    {
16        #pragma STDC FENV_ACCESS ON
17
18        fenv_t env;
19
20        fegetenv(&env); // Save the environment
21
22        fesetround(FE_TOWARDZERO); // Change rounding
23        feraiseexcept(FE_DIVBYZERO); // Raise an exception
24
25        show_status();
26
27        fesetenv(&env); // Restore the environment
28
29        show_status();
30    }

```

Output:

```

Rounding is FE_TOWARDZERO: 1
FE_DIVBYZERO is set: 1
Rounding is FE_TOWARDZERO: 0
FE_DIVBYZERO is set: 0

```

## See Also

`feholdexcept()`, `feupdateenv()`

## 7.9 `feholdexcept()`

Save floating point state and install non-stop mode

### Synopsis

```

#include <fenv.h>

int feholdexcept(fenv_t *envp);

```

### Description

This is just like `fegetenv()` except that it updates the current environment to be in *non-stop* mode, namely it won't halt on any exceptions.

It remains in this state until you restore the state with `fesetenv()` or `feupdateenv()`.

## Return Value

### Example

This example saves the environment and goes into non-stop mode, messes with the rounding and exceptions, then restores it. After the environment is restored, we see that the rounding is back to default and the exception is cleared. We'll also be out of non-stop mode.

```
1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  void show_status(void)
6  {
7      printf("Rounding is FE_TOWARDZERO: %d\n",
8             fegetround() == FE_TOWARDZERO);
9
10     printf("FE_DIVBYZERO is set: %d\n",
11            fetestexcept(FE_DIVBYZERO) != 0);
12 }
13
14 int main(void)
15 {
16     #pragma STDC FENV_ACCESS ON
17
18     fenv_t env;
19
20     // Save the environment and don't stop on exceptions
21     feholdexcept(&env);
22
23     fesetround(FE_TOWARDZERO); // Change rounding
24     feraiseexcept(FE_DIVBYZERO); // Raise an exception
25
26     show_status();
27
28     fesetenv(&env); // Restore the environment
29
30     show_status();
31 }
```

## See Also

`fegetenv()`, `fesetenv()`, `feupdateenv()`

---

## 7.10 feupdateenv()

Restore floating point environment and apply recent exceptions



## Synopsis

```
#include <fenv.h>

int feupdateenv(const fenv_t *envp);
```

## Description

This is like `fesetenv()` except that it modifies the passed-in environment so that it is updated with exceptions that have happened in the meantime.

So let's say you had a function that might raise exceptions, but you wanted to hide those in the caller. One option might be to:

1. Save the environment with `fegetenv()` or `feholdexcept()`.
2. Do whatever you do that might raise exceptions.
3. Restore the environment with `fesetenv()`, thereby hiding the exceptions that happened in step 2.

But that hides *all* exceptions. What if you just wanted to hide some of them? You could use `feupdateenv()` like this:

1. Save the environment with `fegetenv()` or `feholdexcept()`.
2. Do whatever you do that might raise exceptions.
3. Call `feclearexcept()` to clear the exceptions you want to hide from the caller.
4. Call `feupdateenv()` to restore the previous environment and update it with the other exceptions that have occurred.

So it's like a more capable way of restoring the environment than simply `fegetenv()/fesetenv()`.

## Return Value

Returns 0 on success, non-zero otherwise.

## Example

This program saves state, raises some exceptions, then clears one of the exceptions, then restores and updates the state.

```
1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  void show_status(void)
6  {
7      printf("FE_DIVBYZERO: %d\n", fetestexcept(FE_DIVBYZERO) != 0);
8      printf("FE_INVALID : %d\n", fetestexcept(FE_INVALID) != 0);
9      printf("FE_OVERFLOW : %d\n\n", fetestexcept(FE_OVERFLOW) != 0);
10 }
11
12 int main(void)
13 {
14     #pragma STDC FENV_ACCESS ON
15
16     fenv_t env;
17
```

```
18     feholdexcept(&env); // Save the environment
19
20     // Pretend some bad math happened here:
21     feraiseexcept(FE_DIVBYZERO); // Raise an exception
22     feraiseexcept(FE_INVALID);   // Raise an exception
23     feraiseexcept(FE_OVERFLOW);  // Raise an exception
24
25     show_status();
26
27     feclearexcept(FE_INVALID);
28
29     feupdateenv(&env); // Restore the environment
30
31     show_status();
32 }
```

In the output, at first we have no exceptions. Then we have the three we raised. Then after we restore/update the environment, we see the one we cleared (`FE_INVALID`) hasn't been applied:

```
FE_DIVBYZERO: 0
FE_INVALID   : 0
FE_OVERFLOW  : 0

FE_DIVBYZERO: 1
FE_INVALID   : 1
FE_OVERFLOW  : 1

FE_DIVBYZERO: 1
FE_INVALID   : 0
FE_OVERFLOW  : 1
```

## See Also

`fegetenv()`, `fesetenv()`, `feholdexcept()`, `feclearexcept()`

## Chapter 8

# <float.h> Floating Point Limits

| Macro            | Minimum Magnitude | Description                                                                                |
|------------------|-------------------|--------------------------------------------------------------------------------------------|
| FLT_ROUNDS       |                   | Current rounding mode                                                                      |
| FLT_EVAL_METHOD  |                   | Types used for evaluation                                                                  |
| FLT_HAS_SUBNORM  |                   | Subnormal support for float                                                                |
| DBL_HAS_SUBNORM  |                   | Subnormal support for double                                                               |
| LDBL_HAS_SUBNORM |                   | Subnormal support for long double                                                          |
| FLT_RADIX        | 2                 | Floating point radix (base)                                                                |
| FLT_MANT_DIG     |                   | Number of base FLT_RADIX digits in a float                                                 |
| DBL_MANT_DIG     |                   | Number of base FLT_RADIX digits in a double                                                |
| LDBL_MANT_DIG    |                   | Number of base FLT_RADIX digits in a long double                                           |
| FLT_DECIMAL_DIG  | 6                 | Number of decimal digits required to encode a float                                        |
| DBL_DECIMAL_DIG  | 10                | Number of decimal digits required to encode a double                                       |
| LDBL_DECIMAL_DIG | 10                | Number of decimal digits required to encode a long double                                  |
| DECIMAL_DIG      | 10                | Number of decimal digits required to encode the the widest floating point number supported |
| FLT_DIG          | 6                 | Number of decimal digits that can be safely stored in a float                              |
| DBL_DIG          | 10                | Number of decimal digits that can be safely stored in a double                             |
| LDBL_DIG         | 10                | Number of decimal digits that can be safely stored in a long double                        |
| FLT_MIN_EXP      |                   | FLT_RADIX to the FLT_MIN_EXP-1 power is the smallest normalized float                      |
| DBL_MIN_EXP      |                   | FLT_RADIX to the DBL_MIN_EXP-1 power is the smallest normalized double                     |
| LDBL_MIN_EXP     |                   | FLT_RADIX to the LDBL_MIN_EXP-1 power is the smallest normalized long double               |
| FLT_MIN_10_EXP   | -37               | Minimum exponent such that 10 to this number is a normalized float                         |
| DBL_MIN_10_EXP   | -37               | Minimum exponent such that 10 to this number is a normalized double                        |
| LDBL_MIN_10_EXP  | -37               | Minimum exponent such that 10 to this number is a normalized long double                   |
| FLT_MAX_EXP      |                   | FLT_RADIX to the FLT_MAX_EXP-1 power is the largest finite float                           |

| Macro           | Minimum Magnitude | Description                                                             |
|-----------------|-------------------|-------------------------------------------------------------------------|
| DBL_MAX_EXP     |                   | FLT_RADIX to the DBL_MAX_EXP-1 power is the largest finite double       |
| LDBL_MAX_EXP    |                   | FLT_RADIX to the LDBL_MAX_EXP-1 power is the largest finite long double |
| FLT_MAX_10_EXP  | -37               | Minimum exponent such that 10 to this number is a finite float          |
| DBL_MAX_10_EXP  | -37               | Minimum exponent such that 10 to this number is a finite double         |
| LDBL_MAX_10_EXP | -37               | Minimum exponent such that 10 to this number is a finite long double    |
| FLT_MAX         | 1E+37             | Largest finite float                                                    |
| DBL_MAX         | 1E+37             | Largest finite double                                                   |
| LDBL_MAX        | 1E+37             | Largest finite long double                                              |

| Macro         | Maximum Value | Description                                                         |
|---------------|---------------|---------------------------------------------------------------------|
| FLT_EPSILON   | 1E-5          | Difference between 1 and the next biggest representable float       |
| DBL_EPSILON   | 1E-9          | Difference between 1 and the next biggest representable double      |
| LDBL_EPSILON  | 1E-9          | Difference between 1 and the next biggest representable long double |
| FLT_MIN       | 1E-37         | Minimum positive normalized float                                   |
| DBL_MIN       | 1E-37         | Minimum positive normalized double                                  |
| LDBL_MIN      | 1E-37         | Minimum positive normalized long double                             |
| FLT_TRUE_MIN  | 1E-37         | Minimum positive float                                              |
| DBL_TRUE_MIN  | 1E-37         | Minimum positive double                                             |
| LDBL_TRUE_MIN | 1E-37         | Minimum positive long double                                        |

The minimum and maximum values here are from the spec—they should what you can at least expect across all platforms. Your super dooper machine might do better, still!

## 8.1 Background

The spec allows a lot of leeway when it comes to how C represents floating point numbers. This header file spells out the limits on those numbers.

It gives a model that can describe any floating point number that I *know* you're going to absolutely love. It looks like this:

$$x = sb^e \sum_{k=1}^p f_k b^{-k}, e_{\min} \leq e \leq e_{\max}$$

where:

| Variable | Meaning                                            |
|----------|----------------------------------------------------|
| $s$      | Sign, $-1$ or $1$                                  |
| $b$      | Base (radix), probably 2 on your system            |
| $e$      | Exponent                                           |
| $p$      | Precision: how many base- $b$ digits in the number |

| Variable | Meaning                                              |
|----------|------------------------------------------------------|
| $f_k$    | The individual digits of the number, the significand |

But let's blissfully ignore all that for a second.

Let's assume your computer uses base 2 for its floating point (it probably does). And that in the example below the 1s-and-0s numbers are in binary, and the rest are in decimal.

The short of it is you could have floating point numbers like shown in this example:

$$-0.10100101 \times 2^5 = -10100.101 = -20.625$$

That's your fractional part multiplied by the base to the exponent's power. The exponent controls where the decimal point is. It "floats" around!

## 8.2 FLT\_ROUNDS Details

This tells you the rounding mode. It can be changed with a call to `fesetround()`.

| Mode | Description                             |
|------|-----------------------------------------|
| -1   | Indeterminable                          |
| 0    | Toward zero                             |
| 1    | To nearest                              |
| 2    | Toward positive infinity                |
| 3    | Toward negative infinity... and beyond! |

Unlike every other macro in this here header, `FLT_ROUNDS` might not be a constant expression.

## 8.3 FLT\_EVAL\_METHOD Details

This basically tells you how floating point values are promoted to different types in expressions.

| Method | Description                                                                                                                                        |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| -1     | Indeterminable                                                                                                                                     |
| 0      | Evaluate all operations and constants to the precision of their respective types                                                                   |
| 1      | Evaluate <code>float</code> and <code>double</code> operations as <code>double</code> and <code>long double</code> ops as <code>long double</code> |
| 2      | Evaluate all operations and constants as <code>long double</code>                                                                                  |

## 8.4 Subnormal Numbers

The macros `FLT_HAS_SUBNORM`, `DBL_HAS_SUBNORM`, and `LDBL_HAS_SUBNORM` all let you know if those types support subnormal numbers<sup>1</sup>.

| Value | Description                            |
|-------|----------------------------------------|
| -1    | Indeterminable                         |
| 0     | Subnormals not supported for this type |

<sup>1</sup>[https://en.wikipedia.org/wiki/Subnormal\\_number](https://en.wikipedia.org/wiki/Subnormal_number)

| Value | Description                        |
|-------|------------------------------------|
| 1     | Subnormals supported for this type |

## 8.5 How Many Decimal Places Can I Use?

It depends on what you want to do.

The safe thing is if you never use more than `FLT_DIG` base-10 digits in your `float`, you're good. (Same for `DBL_DIG` and `LDBL_DIG` for their types.)

And by “use” I mean print out, have in code, read from the keyboard, etc.

You can print out that many decimal places with `printf()` and the `%g` format specifier:

```

1  #include <stdio.h>
2  #include <float.h>
3
4  int main(void)
5  {
6      float pi = 3.1415926535897932384626433832795028841971;
7
8      // With %g or %G, the precision refers to the number of significant
9      // digits:
10
11     printf("%.*g\n", FLT_DIG, pi); // For me: 3.14159
12
13     // But %f prints too many, since the precision is the number of
14     // digits to the right of the decimal--it doesn't count the digits
15     // to the left of it:
16
17     printf("%.*f\n", FLT_DIG, pi); // For me: 3.14159... 3 ???
18 }
```

That's the end, but stay tuned for the exciting conclusion of “How Many Decimal Places Can I Use?”

Because base 10 and base 2 (your typical `FLT_RADIX`) don't mix very well, you can actually have more than `FLT_DIG` in your `float`; the bits of storage go out a little farther. But these might round in a way you don't expect.

But if you want to convert a floating point number to base 10 and then be able to convert it back again to the exact same floating point number, you'll need `FLT_DECIMAL_DIG` digits from your `float` to make sure you get those extra bits of storage represented. (And `DBL_DECIMAL_DIG` and `LDBL_DECIMAL_DIG` for those corresponding types.)

Here's some example output that shows how the value stored might have some extra decimal places at the end.

```

1  #include <stdio.h>
2  #include <math.h>
3  #include <assert.h>
4  #include <float.h>
5
6  int main(void)
7  {
```

```

8     printf("FLT_DIG = %d\n", FLT_DIG);
9     printf("FLT_DECIMAL_DIG = %d\n\n", FLT_DECIMAL_DIG);
10
11     assert(FLT_DIG == 6); // Code below assumes this
12
13     for (float x = 0.123456; x < 0.12346; x += 0.000001) {
14         printf("As written: %.*g\n", FLT_DIG, x);
15         printf("As stored: %.*g\n\n", FLT_DECIMAL_DIG, x);
16     }
17 }

```

And the output on my machine, starting at `0.123456` and incrementing by `0.000001` each time:

```

FLT_DIG = 6
FLT_DECIMAL_DIG = 9

As written: 0.123456
As stored: 0.123456001

As written: 0.123457
As stored: 0.123457

As written: 0.123458
As stored: 0.123457998

As written: 0.123459
As stored: 0.123458996

As written: 0.12346
As stored: 0.123459995

```

You can see that the value stored isn't always the value we're expecting since base-2 can't represent all base-10 fractions exactly. The best it can do is store more places and then round.

Also notice that even though we tried to stop the `for` loop *before* `0.123460`, it actually ran including that value since the stored version of that number was `0.123459995`, which is still less than `0.123460`.

Aren't floating point numbers fun?

## 8.6 Comprehensive Example

Here's a program that prints out the details for a particular machine:

```

1  #include <stdio.h>
2  #include <float.h>
3
4  int main(void)
5  {
6      printf("FLT_RADIX: %d\n", FLT_RADIX);
7      printf("FLT_ROUNDS: %d\n", FLT_ROUNDS);
8      printf("FLT_EVAL_METHOD: %d\n", FLT_EVAL_METHOD);
9      printf("DECIMAL_DIG: %d\n\n", DECIMAL_DIG);
10 }

```

```

11     printf("FLT_HAS_SUBNORM: %d\n", FLT_HAS_SUBNORM);
12     printf("FLT_MANT_DIG: %d\n", FLT_MANT_DIG);
13     printf("FLT_DECIMAL_DIG: %d\n", FLT_DECIMAL_DIG);
14     printf("FLT_DIG: %d\n", FLT_DIG);
15     printf("FLT_MIN_EXP: %d\n", FLT_MIN_EXP);
16     printf("FLT_MIN_10_EXP: %d\n", FLT_MIN_10_EXP);
17     printf("FLT_MAX_EXP: %d\n", FLT_MAX_EXP);
18     printf("FLT_MAX_10_EXP: %d\n", FLT_MAX_10_EXP);
19     printf("FLT_MIN: %. *e\n", FLT_DECIMAL_DIG, FLT_MIN);
20     printf("FLT_MAX: %. *e\n", FLT_DECIMAL_DIG, FLT_MAX);
21     printf("FLT_EPSILON: %. *e\n", FLT_DECIMAL_DIG, FLT_EPSILON);
22     printf("FLT_TRUE_MIN: %. *e\n\n", FLT_DECIMAL_DIG, FLT_TRUE_MIN);
23
24     printf("DBL_HAS_SUBNORM: %d\n", DBL_HAS_SUBNORM);
25     printf("DBL_MANT_DIG: %d\n", DBL_MANT_DIG);
26     printf("DBL_DECIMAL_DIG: %d\n", DBL_DECIMAL_DIG);
27     printf("DBL_DIG: %d\n", DBL_DIG);
28     printf("DBL_MIN_EXP: %d\n", DBL_MIN_EXP);
29     printf("DBL_MIN_10_EXP: %d\n", DBL_MIN_10_EXP);
30     printf("DBL_MAX_EXP: %d\n", DBL_MAX_EXP);
31     printf("DBL_MAX_10_EXP: %d\n", DBL_MAX_10_EXP);
32     printf("DBL_MIN: %. *e\n", DBL_DECIMAL_DIG, DBL_MIN);
33     printf("DBL_MAX: %. *e\n", DBL_DECIMAL_DIG, DBL_MAX);
34     printf("DBL_EPSILON: %. *e\n", DBL_DECIMAL_DIG, DBL_EPSILON);
35     printf("DBL_TRUE_MIN: %. *e\n\n", DBL_DECIMAL_DIG, DBL_TRUE_MIN);
36
37     printf("LDBL_HAS_SUBNORM: %d\n", LDBL_HAS_SUBNORM);
38     printf("LDBL_MANT_DIG: %d\n", LDBL_MANT_DIG);
39     printf("LDBL_DECIMAL_DIG: %d\n", LDBL_DECIMAL_DIG);
40     printf("LDBL_DIG: %d\n", LDBL_DIG);
41     printf("LDBL_MIN_EXP: %d\n", LDBL_MIN_EXP);
42     printf("LDBL_MIN_10_EXP: %d\n", LDBL_MIN_10_EXP);
43     printf("LDBL_MAX_EXP: %d\n", LDBL_MAX_EXP);
44     printf("LDBL_MAX_10_EXP: %d\n", LDBL_MAX_10_EXP);
45     printf("LDBL_MIN: %. *Le\n", LDBL_DECIMAL_DIG, LDBL_MIN);
46     printf("LDBL_MAX: %. *Le\n", LDBL_DECIMAL_DIG, LDBL_MAX);
47     printf("LDBL_EPSILON: %. *Le\n", LDBL_DECIMAL_DIG, LDBL_EPSILON);
48     printf("LDBL_TRUE_MIN: %. *Le\n\n", LDBL_DECIMAL_DIG, LDBL_TRUE_MIN);
49
50     printf("sizeof(float): %zu\n", sizeof(float));
51     printf("sizeof(double): %zu\n", sizeof(double));
52     printf("sizeof(long double): %zu\n", sizeof(long double));
53 }

```

And here's the output on my machine:

```

FLT_RADIX: 2
FLT_ROUNDS: 1
FLT_EVAL_METHOD: 0
DECIMAL_DIG: 21

FLT_HAS_SUBNORM: 1

```



```
FLT_MANT_DIG: 24
FLT_DECIMAL_DIG: 9
FLT_DIG: 6
FLT_MIN_EXP: -125
FLT_MIN_10_EXP: -37
FLT_MAX_EXP: 128
FLT_MAX_10_EXP: 38
FLT_MIN: 1.175494351e-38
FLT_MAX: 3.402823466e+38
FLT_EPSILON: 1.192092896e-07
FLT_TRUE_MIN: 1.401298464e-45

DBL_HAS_SUBNORM: 1
DBL_MANT_DIG: 53
DBL_DECIMAL_DIG: 17
DBL_DIG: 15
DBL_MIN_EXP: -1021
DBL_MIN_10_EXP: -307
DBL_MAX_EXP: 1024
DBL_MAX_10_EXP: 308
DBL_MIN: 2.22507385850720138e-308
DBL_MAX: 1.79769313486231571e+308
DBL_EPSILON: 2.22044604925031308e-16
DBL_TRUE_MIN: 4.94065645841246544e-324

LDBL_HAS_SUBNORM: 1
LDBL_MANT_DIG: 64
LDBL_DECIMAL_DIG: 21
LDBL_DIG: 18
LDBL_MIN_EXP: -16381
LDBL_MIN_10_EXP: -4931
LDBL_MAX_EXP: 16384
LDBL_MAX_10_EXP: 4932
LDBL_MIN: 3.362103143112093506263e-4932
LDBL_MAX: 1.189731495357231765021e+4932
LDBL_EPSILON: 1.084202172485504434007e-19
LDBL_TRUE_MIN: 3.645199531882474602528e-4951

sizeof(float): 4
sizeof(double): 8
sizeof(long double): 16
```

## Chapter 9

# <inttypes.h> More Integer Conversions

| Function                 | Description                                                  |
|--------------------------|--------------------------------------------------------------|
| <code>imaxabs()</code>   | Compute the absolute value of an <code>intmax_t</code>       |
| <code>imaxdiv()</code>   | Compute the quotient and remainder of <code>intmax_ts</code> |
| <code>strtoimax()</code> | Convert strings to type <code>intmax_t</code>                |
| <code>strtoumax()</code> | Convert strings to type <code>uintmax_t</code>               |
| <code>wcstoimax()</code> | Convert wide strings to type <code>intmax_t</code>           |
| <code>wcstoumax()</code> | Convert wide strings to type <code>uintmax_t</code>          |

This header does conversions to maximum sized integers, division with maximum sized integers, and also provides format specifiers for `printf()` and `scanf()` for a variety of types defined in `<stdint.h>`.

The header `<stdint.h>` is included by this one.

### 9.1 Macros

These are to help with `printf()` and `scanf()` when you use a type such as `int_least16_t`... what format specifiers do you use?

Let's start with `printf()`—all these macros start with `PRI` and then are followed by the format specifier you'd typically use for that type. Lastly, the number of bits is added on.

For example, the format specifier for a 64-bit integer is `PRId64`—the `d` is because you usually print integers with `"%d"`.

An unsigned 16-bit integer could be printed with `PRId16`.

These macros expand to string literals. We can take advantage of the fact that C automatically concatenates neighboring string literals and use these specifiers like this:

```
1 #include <stdio.h>      // for printf()
2 #include <inttypes.h>
3
4 int main(void)
5 {
6     int16_t x = 32;
```

```

7
8     printf("The value is %" PRId16 "!\\n", x);
9 }

```

There's nothing magical happening on line 8, above. Indeed, if I print out the value of the macro:

```
printf("%s\\n", PRId16);
```

we get this on my system:

```
hd
```

which is a `printf()` format specifier meaning “short signed integer”.

So back to line 8, after string literal concatenation, it's just as if I'd typed:

```
8     printf("The value is %hd!\\n", x);
```

Here's a table of all the macros you can use for `printf()` format specifiers... substitute the number of bits for *N*, usually 8, 16, 32, or 64.

|                |                     |                    |          |          |
|----------------|---------------------|--------------------|----------|----------|
| PRId <i>N</i>  | PRIdLEAST <i>N</i>  | PRIdFAST <i>N</i>  | PRIdMAX  | PRIdPTR  |
| PRIdi <i>N</i> | PRIdiLEAST <i>N</i> | PRIdiFAST <i>N</i> | PRIdiMAX | PRIdiPTR |
| PRIdo <i>N</i> | PRIdoLEAST <i>N</i> | PRIdoFAST <i>N</i> | PRIdoMAX | PRIdoPTR |
| PRIdu <i>N</i> | PRIduLEAST <i>N</i> | PRIduFAST <i>N</i> | PRIduMAX | PRIduPTR |
| PRIdx <i>N</i> | PRIdxLEAST <i>N</i> | PRIdxFAST <i>N</i> | PRIdxMAX | PRIdxPTR |
| PRIdX <i>N</i> | PRIdXLEAST <i>N</i> | PRIdXFAST <i>N</i> | PRIdXMAX | PRIdXPTR |

Note again how the lowercase center letter represents the usual format specifiers you'd pass to `printf()`: *d*, *i*, *o*, *u*, *x*, and *X*.

And we have a similar set of macros for `scanf()` for reading in these various types:

|               |                    |                   |         |         |
|---------------|--------------------|-------------------|---------|---------|
| SCNd <i>N</i> | SCNdLEAST <i>N</i> | SCNdFAST <i>N</i> | SCNdMAX | SCNdPTR |
| SCNi <i>N</i> | SCNiLEAST <i>N</i> | SCNiFAST <i>N</i> | SCNiMAX | SCNiPTR |
| SCNo <i>N</i> | SCNoLEAST <i>N</i> | SCNoFAST <i>N</i> | SCNoMAX | SCNoPTR |
| SCNu <i>N</i> | SCNuLEAST <i>N</i> | SCNuFAST <i>N</i> | SCNuMAX | SCNuPTR |
| SCNx <i>N</i> | SCNxLEAST <i>N</i> | SCNxFAST <i>N</i> | SCNxMAX | SCNxPTR |

The rule is that for each type defined in <stdint.h> there will be corresponding `printf()` and `scanf()` macros defined here.

## 9.2 imaxabs()

Compute the absolute value of an `intmax_t`

## Synopsis

```
#include <inttypes.h>

intmax_t imaxabs(intmax_t j);
```

## Description

When you need the absolute value of the biggest integer type on the system, this is the function for you.

The spec notes that if the absolute value of the number cannot be represented, the behavior is undefined. This would happen if you tried to take the absolute value of the smallest possible negative number in a two's-complement system.

## Return Value

Returns the absolute value of the input,  $|j|$ .

## Example

```
1 #include <stdio.h>
2 #include <inttypes.h>
3
4 int main(void)
5 {
6     intmax_t j = -3490;
7
8     printf("%jd\n", imaxabs(j));    // 3490
9 }
```

## See Also

`fabs()`

---

## 9.3 `imaxdiv()`

Compute the quotient and remainder of `intmax_t`s

## Synopsis

```
#include <inttypes.h>

imaxdiv_t imaxdiv(intmax_t numer, intmax_t denom);
```

## Description

When you want to do integer division and remainder in a single operation, this function will do it for you.

It computes `numer/denom` and `numer%denom` and returns the result in a structure of type `imaxdiv_t`.

This structure has two `imaxdiv_t` fields, `quot` and `rem`, that you use to retrieve the sought-after values.

## Return Value

Returns an `imaxdiv_t` containing the quotient and remainder of the operation.

## Example

```
1 #include <stdio.h>
2 #include <inttypes.h>
3
4 int main(void)
5 {
6     intmax_t numer = INTMAX_C(3490);
7     intmax_t denom = INTMAX_C(17);
8
9     imaxdiv_t r = imaxdiv(numer, denom);
10
11     printf("Quotient: %jd, remainder: %jd\n", r.quot, r.rem);
12 }
```

Output:

```
Quotient: 205, remainder: 5
```

## See Also

`remquo()`

---

## 9.4 `strtoimax()` `strtoumax()`

Convert strings to types `intmax_t` and `uintmax_t`

### Synopsis

```
#include <inttypes.h>

intmax_t strtoimax(const char * restrict nptr, char ** restrict endptr,
                  int base);

uintmax_t strtoumax(const char * restrict nptr, char ** restrict endptr,
                   int base);
```

### Description

These work just like the `strtol()` family of functions, except they return an `intmax_t` or `uintmax_t`.

See the `strtol()` reference page for details.

## Return Value

Returns the converted string as an `intmax_t` or `uintmax_t`.

If the result is out of range, the returned value will be `INTMAX_MAX`, `INTMAX_MIN`, or `UINTMAX_MAX`, as appropriate. And the `errno` variable will be set to `ERANGE`.

## Example

The following example converts a base-10 number to an `intmax_t`. Then it attempts to convert an invalid base-2 number, catching the error.

```

1  #include <stdio.h>
2  #include <inttypes.h>
3
4  int main(void)
5  {
6      intmax_t r;
7      char *endptr;
8
9      // Valid base-10 number
10     r = strtoumax("123456789012345", &endptr, 10);
11
12     if (*endptr != '\0')
13         printf("Invalid digit: %c\n", *endptr);
14     else
15         printf("Value is %jd\n", r);
16
17     // The following binary number contains an invalid digit
18     r = strtoumax("0100102010101101", &endptr, 2);
19
20     if (*endptr != '\0')
21         printf("Invalid digit: %c\n", *endptr);
22     else
23         printf("Value is %jd\n", r);
24 }
```

Output:

```
Value is 123456789012345
Invalid digit: 2
```

## See Also

`strtol()`, `errno`

## 9.5 `wcstoimax()` `wcstoumax()`

Convert wide strings to types `intmax_t` and `uintmax_t`

## Synopsis

```
#include <stddef.h> // for wchar_t
#include <inttypes.h>

intmax_t wcstoimax(const wchar_t * restrict nptr,
                  wchar_t ** restrict endptr, int base);

uintmax_t wcstoumax(const wchar_t * restrict nptr,
                   wchar_t ** restrict endptr, int base);
```

## Description

These work just like the `wcstol()` family of functions, except they return an `intmax_t` or `uintmax_t`.

See the `wcstol()` reference page for details.

## Return Value

Returns the converted wide string as an `intmax_t` or `uintmax_t`.

If the result is out of range, the returned value will be `INTMAX_MAX`, `INTMAX_MIN`, or `UINTMAX_MAX`, as appropriate. And the `errno` variable will be set to `ERANGE`.

## Example

The following example converts a base-10 number to an `intmax_t`. Then it attempts to convert an invalid base-2 number, catching the error.

```
1  #include <wchar.h>
2  #include <inttypes.h>
3
4  int main(void)
5  {
6      intmax_t r;
7      wchar_t *endptr;
8
9      // Valid base-10 number
10     r = wcstoimax(L"123456789012345", &endptr, 10);
11
12     if (*endptr != '\0')
13         wprintf(L"Invalid digit: %lc\n", *endptr);
14     else
15         wprintf(L"Value is %jd\n", r);
16
17     // The following binary number contains an invalid digit
18     r = wcstoimax(L"0100102010101101", &endptr, 2);
19
20     if (*endptr != '\0')
21         wprintf(L"Invalid digit: %lc\n", *endptr);
22     else
23         wprintf(L"Value is %jd\n", r);
24 }
```

```
Value is 123456789012345  
Invalid digit: 2
```

### See Also

`wcstol()`, `errno`



## Chapter 10

# <iso646.h> Alternative Operator Spellings

ISO-646 is a character encoding standard that's very similar to ASCII. But it's missing a few notable characters, like `|`, `^`, and `~`.

Since these are operators or parts of operators in C, this header file defines a number of macros you can use in case those characters aren't found on your keyboard. (And also C++ can use these same alternate spellings.)

| Operator                | <iso646.h> equivalent |
|-------------------------|-----------------------|
| <code>&amp;&amp;</code> | <code>and</code>      |
| <code>&amp;=</code>     | <code>and_eq</code>   |
| <code>&amp;</code>      | <code>bitand</code>   |
| <code> </code>          | <code>bitor</code>    |
| <code>~</code>          | <code>compl</code>    |
| <code>!</code>          | <code>not</code>      |
| <code>!=</code>         | <code>not_eq</code>   |
| <code>  </code>         | <code>or</code>       |
| <code> =</code>         | <code>or_eq</code>    |
| <code>^</code>          | <code>xor</code>      |
| <code>^=</code>         | <code>xor_eq</code>   |

Interestingly, there is no `eq` for `==`, and `&` and `!` are included despite being in ISO-646.

Example usage:

```
1  #include <stdio.h>
2  #include <iso646.h>
3
4  int main(void)
5  {
6      int x = 12;
7      int y = 30;
8
9      if (x == 12 and y not_eq 40)
10         printf("Now we know.\n");
11 }
```

I've personally never seen this file included, but I'm sure it gets used from time to time.

# Chapter 11

## <limits.h> Numeric Limits

Important note: the “minimum magnitude” in the table below is the minimum allowed by the spec. It’s very likely that the values on your bad-ass system exceed those, below.

| Macro      | Minimum Magnitude      | Description                                                    |
|------------|------------------------|----------------------------------------------------------------|
| CHAR_BIT   | 8                      | Number of bits in a byte                                       |
| SCHAR_MIN  | -127                   | Minimum value of a signed char                                 |
| SCHAR_MAX  | 127                    | Maximum value of a signed char                                 |
| UCHAR_MAX  | 255                    | Maximum value of an unsigned char <sup>1</sup>                 |
| CHAR_MIN   | 0 or SCHAR_MIN         | More detail below                                              |
| CHAR_MAX   | SCHAR_MAX or UCHAR_MAX | More detail below                                              |
| MB_LEN_MAX | 1                      | Maximum number of bytes in a multibyte character on any locale |
| SHRT_MIN   | -32767                 | Minimum value of a short                                       |
| SHRT_MAX   | 32767                  | Maximum value of a short                                       |
| USHRT_MAX  | 65535                  | Maximum value of an unsigned short                             |
| INT_MIN    | -32767                 | Minimum value of an int                                        |
| INT_MAX    | 32767                  | Maximum value of an int                                        |
| UINT_MAX   | 65535                  | Maximum value of an unsigned int                               |
| LONG_MIN   | -2147483647            | Minimum value of a long                                        |
| LONG_MAX   | 2147483647             | Maximum value of a long                                        |
| ULONG_MAX  | 4294967295             | Maximum value of an unsigned long                              |
| LLONG_MIN  | -9223372036854775807   | Minimum value of a long long                                   |
| LLONG_MAX  | 9223372036854775807    | Maximum value of a long long                                   |
| ULLONG_MAX | 18446744073709551615   | Maximum value of an unsigned long long                         |

### 11.1 CHAR\_MIN and CHAR\_MAX

When it comes to the CHAR\_MIN and CHAR\_MAX macros, it all depends on if your char type is signed or unsigned by default. Remember that C leaves that up to the implementation? No? Well, it does.

So if it’s signed, the values of CHAR\_MIN and CHAR\_MAX are the same as SCHAR\_MIN and SCHAR\_MAX.

And if it’s unsigned, the values of CHAR\_MIN and CHAR\_MAX are the same as 0 and UCHAR\_MAX.

<sup>1</sup>The minimum value of an unsigned char is 0. Same for an unsigned short and unsigned long. Or any unsigned type, for that matter.

Side benefit: you can tell at runtime if the system has signed or unsigned chars by checking to see if `CHAR_MIN` is 0.

```

1 #include <stdio.h>
2 #include <limits.h>
3
4 int main(void)
5 {
6     printf("chars are %signed\n", CHAR_MIN == 0? "un": "");
7 }
```

On my system, chars are signed.

## 11.2 Choosing the Correct Type

If you want to be super portable, choose a type you know will be at least as big as you need by the table, above.

That said, a lot of code, for better or (likely) worse, assumes ints are 32-bits, when in actuality it's only guaranteed to be 16.

If you need a guaranteed bit size, check out the `int_leastN_t` types in `<stdint.h>`.

## 11.3 Whither Two's Complement?

If you were looking closely and have *a priori* knowledge of the matter, you might have thought I erred in the minimum values of the macros, above.

“short goes from 32767 to -32767? Shouldn't it go to -32768?”

No, I have it right. The spec lists the minimum magnitudes for those macros, and some old-timey systems might have used a different encoding for their signed values that could only go that far.

Virtually every modern system uses Two's Complement<sup>2</sup> for signed numbers, and those would go from 32767 to -32768 for a short. Your system probably does, too.

## 11.4 Demo Program

Here's a program to print out the values of the macros:

```

1 #include <stdio.h>
2 #include <limits.h>
3
4 int main(void)
5 {
6     printf("CHAR_BIT = %d\n", CHAR_BIT);
7     printf("SCHAR_MIN = %d\n", SCHAR_MIN);
8     printf("SCHAR_MAX = %d\n", SCHAR_MAX);
9     printf("UCHAR_MAX = %d\n", UCHAR_MAX);
10    printf("CHAR_MIN = %d\n", CHAR_MIN);
11    printf("CHAR_MAX = %d\n", CHAR_MAX);
12    printf("MB_LEN_MAX = %d\n", MB_LEN_MAX);
13 }
```

<sup>2</sup>[https://en.wikipedia.org/wiki/Two%27s\\_complement](https://en.wikipedia.org/wiki/Two%27s_complement)

```

13     printf("SHRT_MIN = %d\n", SHRT_MIN);
14     printf("SHRT_MAX = %d\n", SHRT_MAX);
15     printf("USHRT_MAX = %u\n", USHRT_MAX);
16     printf("INT_MIN = %d\n", INT_MIN);
17     printf("INT_MAX = %d\n", INT_MAX);
18     printf("UINT_MAX = %u\n", UINT_MAX);
19     printf("LONG_MIN = %ld\n", LONG_MIN);
20     printf("LONG_MAX = %ld\n", LONG_MAX);
21     printf("ULONG_MAX = %lu\n", ULONG_MAX);
22     printf("LLONG_MIN = %lld\n", LLONG_MIN);
23     printf("LLONG_MAX = %lld\n", LLONG_MAX);
24     printf("ULLONG_MAX = %llu\n", ULLONG_MAX);
25 }

```

On my 64-bit Intel system with clang, this outputs:

```

CHAR_BIT = 8
SCHAR_MIN = -128
SCHAR_MAX = 127
UCHAR_MAX = 255
CHAR_MIN = -128
CHAR_MAX = 127
MB_LEN_MAX = 6
SHRT_MIN = -32768
SHRT_MAX = 32767
USHRT_MAX = 65535
INT_MIN = -2147483648
INT_MAX = 2147483647
UINT_MAX = 4294967295
LONG_MIN = -9223372036854775808
LONG_MAX = 9223372036854775807
ULONG_MAX = 18446744073709551615
LLONG_MIN = -9223372036854775808
LLONG_MAX = 9223372036854775807
ULLONG_MAX = 18446744073709551615

```

Looks like my system probably uses two's-complement encoding for signed numbers, my `char`s are signed, and my `ints` are 32-bit.

# Chapter 12

## <locale.h> locale handling

| Function                  | Description                              |
|---------------------------|------------------------------------------|
| <code>setlocale()</code>  | Set the locale                           |
| <code>localeconv()</code> | Get information about the current locale |

The “locale” is the details of how the program should run given its physical location on the planet.

For example, in one locale, a unit of money might be printed as \$123, and in another €123.

Or one locale might use ASCII encoding and another UTF-8 encoding.

By default, the program runs in the “C” locale. It has a basic set of characters with a single-byte encoding. If you try to print UTF-8 characters in the C locale, nothing will print. You have to switch to a proper locale.

### 12.1 `setlocale()`

Set the locale

#### Synopsis

```
#include <locale.h>

char *setlocale(int category, const char *locale);
```

#### Description

Sets the `locale` for the given `category`.

Category is one of the following:

| Category                | Description                                                             |
|-------------------------|-------------------------------------------------------------------------|
| <code>LC_ALL</code>     | All of the following categories                                         |
| <code>LC_COLLATE</code> | Affects the <code>strcoll()</code> and <code>strxfrm()</code> functions |
| <code>LC_CTYPE</code>   | Affects the functions in <code>&lt;ctype.h&gt;</code>                   |

| Category    | Description                                                                                                                                      |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| LC_MONETARY | Affects the monetary information returned from <code>localeconv()</code>                                                                         |
| LC_NUMERIC  | Affects the decimal point for formatted I/O and formatted string functions, and the monetary information returned from <code>localeconv()</code> |
| LC_TIME     | Affects the <code>strftime()</code> and <code>wcsftime()</code> functions                                                                        |

And there are three portable things you can pass in for `locale`; any other string passed in is implementation-defined and non-portable.

| Locale | Description                                                        |
|--------|--------------------------------------------------------------------|
| "C"    | Set the program to the C locale                                    |
| ""     | (Empty string) Set the program to the native locale of this system |
| NULL   | Change nothing; just return the current locale                     |
| Other  | Set the program to an implementation-defined locale                |

The most common call, I'd wager, is this:

```
// Set all locale settings to the local, native locale

setlocale(LC_ALL, "");
```

Handily, `setlocale()` returns the locale that was just set, so you could see what the actual locale is on your system.

## Return Value

On success, returns a pointer to the string representing the current locale. You may not modify this string, and it might be changed by subsequent calls to `setlocale()`.

On failure, returns `NULL`.

## Example

Here we get the current locale. Then we set it to the native locale, and print out what that is.

```
1  #include <stdio.h>
2  #include <locale.h>
3
4  int main(void)
5  {
6      char *loc;
7
8      // Get the current locale
9      loc = setlocale(LC_ALL, NULL);
10
11     printf("Starting locale: %s\n", loc);
12
```

```

13 // Set (and get) the locale to native locale
14 loc = setlocale(LC_ALL, "");
15
16 printf("Native locale: %s\n", loc);
17 }

```

Output on my system:

```

Starting locale: C
Native locale: en_US.UTF-8

```

Note that my native locale (on a Linux box) might be different from what you see.

Nevertheless, I can explicitly set it on my system without a problem, or to any other locale I have installed:

```

13 loc = setlocale(LC_ALL, "en_US.UTF-8"); // Non-portable

```

But again, your system might have different locales defined.

## See Also

localeconv(), strcoll(), strxfrm(), strftime(), wcsftime(), printf(), scanf(), <ctype.h>

## 12.2 localeconv()

Get information about the current locale

### Synopsis

```

#include <locale.h>

struct lconv *localeconv(void);

```

### Description

This function just returns a pointer to a `struct lconv`, but is still a bit of a powerhouse.

The returned structure contains *tons* of information about the locale. Here are the fields of `struct lconv` and their meanings.

First, some conventions. In the field names, below, a `_p_` means “positive”, and `_n_` means “negative”, and `int_` means “international”. Though a lot of these are type `char` or `char*`, most (or the strings they point to) are actually treated as integers<sup>1</sup>.

Before we go further, know that `CHAR_MAX` (from `<limits.h>`) is the maximum value that can be held in a `char`. And that many of the following `char` values use that to indicate the value isn’t available in the given locale.

<sup>1</sup>Remember that `char` is just a byte-sized integer.



| Field                                | Description                                                                                                        |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| <code>char *mon_decimal_point</code> | Decimal pointer character for money, e.g. ".".                                                                     |
| <code>char *mon_thousands_sep</code> | Thousands separator character for money, e.g. ",".                                                                 |
| <code>char *mon_grouping</code>      | Grouping description for money (see below).                                                                        |
| <code>char *positive_sign</code>     | Positive sign for money, e.g. "+" or "".                                                                           |
| <code>char *negative_sign</code>     | Negative sign for money, e.g. "-".                                                                                 |
| <code>char *currency_symbol</code>   | Currency symbol, e.g. "\$".                                                                                        |
| <code>char frac_digits</code>        | When printing monetary amounts, how many digits to print past the decimal point, e.g. 2.                           |
| <code>char p_cs_precedes</code>      | 1 if the <code>currency_symbol</code> comes before the value for a non-negative monetary amount, 0 if after.       |
| <code>char n_cs_precedes</code>      | 1 if the <code>currency_symbol</code> comes before the value for a negative monetary amount, 0 if after.           |
| <code>char p_sep_by_space</code>     | Determines the separation of the <code>currency symbol</code> from the value for non-negative amounts (see below). |
| <code>char n_sep_by_space</code>     | Determines the separation of the <code>currency symbol</code> from the value for negative amounts (see below).     |
| <code>char p_sign_posn</code>        | Determines the <code>positive_sign</code> position for non-negative values.                                        |
| <code>char p_sign_posn</code>        | Determines the <code>positive_sign</code> position for negative values.                                            |
| <code>char *int_curr_symbol</code>   | International currency symbol, e.g. "USD ".                                                                        |
| <code>char int_frac_digits</code>    | International value for <code>frac_digits</code> .                                                                 |
| <code>char int_p_cs_precedes</code>  | International value for <code>p_cs_precedes</code> .                                                               |
| <code>char int_n_cs_precedes</code>  | International value for <code>n_cs_precedes</code> .                                                               |
| <code>char int_p_sep_by_space</code> | International value for <code>p_sep_by_space</code> .                                                              |
| <code>char int_n_sep_by_space</code> | International value for <code>n_sep_by_space</code> .                                                              |
| <code>char int_p_sign_posn</code>    | International value for <code>p_sign_posn</code> .                                                                 |
| <code>char int_n_sign_posn</code>    | International value for <code>n_sign_posn</code> .                                                                 |

Even though many of these have `char` type, the value stored within is meant to be accessed as an integer.

All the `sep_by_space` variants deal with spacing around the currency sign. Valid values are:

| Value | Description                                                                                                                                   |
|-------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| 0     | No space between currency symbol and value.                                                                                                   |
| 1     | Separate the currency symbol (and sign, if any) from the value with a space.                                                                  |
| 2     | Separate the sign symbol from the currency symbol (if adjacent) with a space, otherwise separate the sign symbol from the value with a space. |

The `sign_posn` variants are determined by the following values:

| Value | Description                                                    |
|-------|----------------------------------------------------------------|
| 0     | Put parens around the value and the currency symbol.           |
| 1     | Put the sign string in front of the currency symbol and value. |
| 2     | Put the sign string after the currency symbol and value.       |
| 3     | Put the sign string directly in front of the currency symbol.  |
| 4     | Put the sign string directly behind the currency symbol.       |

## Return Value

Returns a pointer to the structure containing the locale information.

The program may not modify this structure.

Subsequent calls to `localeconv()` may overwrite this structure, as might calls to `setlocale()` with `LC_ALL`, `LC_MONETARY`, or `LC_NUMERIC`.

## Example

Here's a program to print the locale information for the native locale.

```

1  #include <stdio.h>
2  #include <locale.h>
3  #include <limits.h> // for CHAR_MAX
4
5  void print_grouping(char *mg)
6  {
7      int done = 0;
8
9      while (!done) {
10         if (*mg == CHAR_MAX)
11             printf("CHAR_MAX ");
12         else
13             printf("%c ", *mg + '0');
14         done = *mg == CHAR_MAX || *mg == 0;
15         mg++;
16     }
17 }
18
19 int main(void)
20 {
21     setlocale(LC_ALL, "");
22
23     struct lconv *lc = localeconv();
24
25     printf("mon_decimal_point : %s\n", lc->mon_decimal_point);
26     printf("mon_thousands_sep : %s\n", lc->mon_thousands_sep);
27     printf("mon_grouping       : ");
28     print_grouping(lc->mon_grouping);
29     printf("\n");
30     printf("positive_sign      : %s\n", lc->positive_sign);
31     printf("negative_sign     : %s\n", lc->negative_sign);
32     printf("currency_symbol    : %s\n", lc->currency_symbol);
33     printf("frac_digits        : %c\n", lc->frac_digits);
34     printf("p_cs_precedes      : %c\n", lc->p_cs_precedes);
35     printf("n_cs_precedes      : %c\n", lc->n_cs_precedes);
36     printf("p_sep_by_space     : %c\n", lc->p_sep_by_space);
37     printf("n_sep_by_space     : %c\n", lc->n_sep_by_space);
38     printf("p_sign_posn        : %c\n", lc->p_sign_posn);
39     printf("n_sign_posn        : %c\n", lc->n_sign_posn);
40     printf("int_curr_symbol     : %s\n", lc->int_curr_symbol);
41     printf("int_frac_digits     : %c\n", lc->int_frac_digits);
42     printf("int_p_cs_precedes   : %c\n", lc->int_p_cs_precedes);
43     printf("int_n_cs_precedes   : %c\n", lc->int_n_cs_precedes);
44     printf("int_p_sep_by_space : %c\n", lc->int_p_sep_by_space);

```

```
45     printf("int_n_sep_by_space: %c\n", lc->int_n_sep_by_space);
46     printf("int_p_sign_posn    : %c\n", lc->int_p_sign_posn);
47     printf("int_n_sign_posn    : %c\n", lc->int_n_sign_posn);
48 }
```

Output on my system:

```
mon_decimal_point : .
mon_thousands_sep : ,
mon_grouping       : 3 3 0
positive_sign      : 
negative_sign       : -
currency_symbol    : $
frac_digits        : 2
p_cs_precedes      : 1
n_cs_precedes      : 1
p_sep_by_space     : 0
n_sep_by_space     : 0
p_sign_posn        : 1
p_sign_posn        : 1
int_curr_symbol     : USD
int_frac_digits     : 2
int_p_cs_precedes   : 1
int_n_cs_precedes   : 1
int_p_sep_by_space : 1
int_n_sep_by_space : 1
int_p_sign_posn     : 1
int_n_sign_posn     : 1
```

## See Also

`setlocale()`

## Chapter 13

# <math.h> Mathematics

Many of the following functions have `float` and `long double` variants as described below (e.g. `pow()`, `powf()`, `powl()`). The `float` and `long double` variants are omitted from the following table to keep your eyeballs from melting out.

| Function                                   | Description                                                             |
|--------------------------------------------|-------------------------------------------------------------------------|
| <code>acos()</code>                        | Calculate the arc cosine of a number.                                   |
| <code>acosh()</code>                       | Compute arc hyperbolic cosine.                                          |
| <code>asin()</code>                        | Calculate the arc sine of a number.                                     |
| <code>asinh()</code>                       | Compute arc hyperbolic sine.                                            |
| <code>atan()</code> , <code>atan2()</code> | Calculate the arc tangent of a number.                                  |
| <code>atanh()</code>                       | Compute the arc hyperbolic tangent.                                     |
| <code>cbrt()</code>                        | Compute the cube root.                                                  |
| <code>ceil()</code>                        | Ceiling—return the next whole number not smaller than the given number. |
| <code>copysign()</code>                    | Copy the sign of one value into another.                                |
| <code>cos()</code>                         | Calculate the cosine of a number.                                       |
| <code>cosh()</code>                        | Compute the hyperbolic cosine.                                          |
| <code>erf()</code>                         | Compute the error function of the given value.                          |
| <code>erfc()</code>                        | Compute the complementary error function of a value.                    |
| <code>exp()</code>                         | Compute $e$ raised to a power.                                          |
| <code>exp2()</code>                        | Compute 2 to a power.                                                   |
| <code>expm1()</code>                       | Compute $e^x - 1$ .                                                     |
| <code>fabs()</code>                        | Compute the absolute value.                                             |
| <code>fdim()</code>                        | Return the positive difference between two numbers clamped at 0.        |
| <code>floor()</code>                       | Compute the largest whole number not larger than the given value.       |
| <code>fma()</code>                         | Floating (AKA “Fast”) multiply and add.                                 |
| <code>fmax()</code> , <code>fmin()</code>  | Return the maximum or minimum of two numbers.                           |
| <code>fmod()</code>                        | Compute the floating point remainder.                                   |
| <code>fpclassify()</code>                  | Return the classification of a given floating point number.             |
| <code>frexp()</code>                       | Break a number into its fraction part and exponent (as a power of 2).   |
| <code>hypot()</code>                       | Compute the length of the hypotenuse of a triangle.                     |
| <code>ilogb()</code>                       | Return the exponent of a floating point number.                         |
| <code>isfinite()</code>                    | True if the number is not infinite or NaN.                              |
| <code>isgreater()</code>                   | True if one argument is greater than another.                           |
| <code>isgreaterequal()</code>              | True if one argument is greater than or equal to another.               |
| <code>isinf()</code>                       | True if the number is infinite.                                         |
| <code>isless()</code>                      | True if one argument is less than another.                              |
| <code>islessequal()</code>                 | True if one argument is less than or equal to another.                  |

| Function                         | Description                                                                     |
|----------------------------------|---------------------------------------------------------------------------------|
| <code>islessgreater()</code>     | Test if a floating point number is less than or greater than another.           |
| <code>isnan()</code>             | True if the number is Not-a-Number.                                             |
| <code>isnormal()</code>          | True if the number is normal.                                                   |
| <code>isunordered()</code>       | Macro returns true if either floating point argument is NaN.                    |
| <code>ldexp()</code>             | Multiply a number by an integral power of 2.                                    |
| <code>lgamma()</code>            | Compute the natural logarithm of the absolute value of $\Gamma(x)$ .            |
| <code>log()</code>               | Compute the natural logarithm.                                                  |
| <code>log10()</code>             | Compute the log-base-10 of a number.                                            |
| <code>log2()</code>              | Compute the base-2 logarithm of a number.                                       |
| <code>logb()</code>              | Extract the exponent of a number given <code>FLT_RADIX</code> .                 |
| <code>log1p()</code>             | Compute the natural logarithm of a number plus 1.                               |
| <code>lrint()</code>             | Returns <code>x</code> rounded in the current rounding direction as an integer. |
| <code>lround(), llround()</code> | Round a number in the good old-fashioned way, returning an integer.             |
| <code>modf()</code>              | Extract the integral and fractional parts of a number.                          |
| <code>nan()</code>               | Return <code>NAN</code> .                                                       |
| <code>nearbyint()</code>         | Rounds a value in the current rounding direction.                               |
| <code>nextafter()</code>         | Get the next (or previous) representable floating point value.                  |
| <code>nexttoward()</code>        | Get the next (or previous) representable floating point value.                  |
| <code>pow()</code>               | Compute a value raised to a power.                                              |
| <code>remainder()</code>         | Compute the remainder IEC 60559-style.                                          |
| <code>remquo()</code>            | Compute the remainder and (some of the) quotient.                               |
| <code>rint()</code>              | Rounds a value in the current rounding direction.                               |
| <code>round()</code>             | Round a number in the good old-fashioned way.                                   |
| <code>scalbn(), scalbln()</code> | Efficiently compute $x \times r^n$ , where $r$ is <code>FLT_RADIX</code> .      |
| <code>signbit()</code>           | Return the sign of a number.                                                    |
| <code>sin()</code>               | Calculate the sine of a number.                                                 |
| <code>sqrt()</code>              | Calculate the square root of a number.                                          |
| <code>tan()</code>               | Calculate the tangent of a number.                                              |
| <code>tanh()</code>              | Compute the hyperbolic tangent.                                                 |
| <code>tgamma()</code>            | Compute the gamma function, $\Gamma(x)$ .                                       |
| <code>trunc()</code>             | Truncate the fractional part off a floating point value.                        |

It's your favorite subject: Mathematics! Hello, I'm Doctor Math, and I'll be making math FUN and EASY!

[vomiting sounds]

Ok, I know math isn't the grandest thing for some of you out there, but these are merely functions that quickly and easily do math you either know, want, or just don't care about. That pretty much covers it.

## 13.1 Math Function Idioms

Many of these math functions exist in three forms, each corresponding to the argument and/or return types the function uses, `float`, `double`, or `long double`.

The alternate form for `float` is made by appending `f` to the end of the function name.

The alternate form for `long double` is made by appending `l` to the end of the function name.

For example, the `pow()` function, which computes  $x^y$ , exists in these forms:

```
double    pow(double x, double y);           // double
float     powf(float x, float y);           // float
```

```
long double powl(long double x, long double y); // long double
```

Remember that parameters are given values as if you assigned into them. So if you pass a `double` to `powf()`, it'll choose the closest `float` it can to hold the double. If the `double` doesn't fit, undefined behavior happens.

## 13.2 Math Types

We have two exciting new types in <math.h>:

- `float_t`
- `double_t`

The `float_t` type is at least as accurate as a `float`, and the `double_t` type is at least as accurate as a `double`.

The idea with these types is they can represent the most efficient way of storing numbers for maximum speed.

Their actual types vary by implementation, but can be determined by the value of the `FLT_EVAL_METHOD` macro.

| FLT_EVAL_METHOD | float_t type           | double_t type          |
|-----------------|------------------------|------------------------|
| 0               | float                  | double                 |
| 1               | double                 | double                 |
| 2               | long double            | long double            |
| Other           | Implementation-defined | Implementation-defined |

For all defined values of `FLT_EVAL_METHOD`, `float_t` is the least-precise type used for all floating calculations.

## 13.3 Math Macros

There are actually a number of these defined, but we'll cover most of them in their relevant reference sections, below.

But here are a couple:

`NAN` represents Not-A-Number.

Defined in <float.h> is `FLT_RADIX`: the number base used by floating point numbers. This is commonly 2, but could be anything.

## 13.4 Math Errors

As we know, nothing can ever go wrong with math... except *everything*!

So there are just a couple errors that might occur when using some of these functions.

- **Range errors** mean that some result is beyond what can be stored in the result type.
- **Domain errors** mean that you've passed in an argument that doesn't have a defined result for this function.
- **Pole errors** mean that the limit of the function as  $x$  approaches the given argument is infinite.
- **Overflow errors** are when the result is really large, but can't be stored without incurring large roundoff error.

- **Underflow errors** are like overflow errors, except with very small numbers.

Now, the C math library can do a couple things when these errors occur:

- Set `errno` to some value, or...
- Raise a floating point exception.

Your system might vary on what happens. You can check it by looking at the value of the variable `math_errhandling`. It will be equivalent to one of the following<sup>1</sup>:

| <code>math_errhandling</code>            | Description                                         |
|------------------------------------------|-----------------------------------------------------|
| <code>MATH_ERRNO</code>                  | The system uses <code>errno</code> for math errors. |
| <code>MATH_ERREXCEPT</code>              | The system uses exceptions for math errors.         |
| <code>MATH_ERRNO   MATH_ERREXCEPT</code> | The system does both! (That's a bitwise-OR!)        |

You are not allowed to change `math_errhandling`.

For a fuller description on how exceptions work and their meanings, see the `<fenv.h>` section.

## 13.5 Math Pragmas

In a nutshell, pragmas offer various ways to control the compiler's behavior. In this case, we're talking about controlling how C's math library works.

In specific, we have a pragma `FP_CONTRACT` that can be turned off and on.

What does it mean?

First of all, keep in mind that any operation in an expression can cause rounding error. So each step of the expression can introduce more rounding error.

But what if the compiler knows a *double secret* way of taking the expression you wrote and converting it to a single instruction that reduced the number of steps such that the intermediate rounding error didn't occur?

Could it use it? I mean, the results would be different than if you let the rounding error settle each step of the way...

Because the results would be different, you can tell the compiler if you want to allow it to do this or not.

If you want to allow it:

```
#pragma STDC FP_CONTRACT ON
```

and to disallow it:

```
#pragma STDC FP_CONTRACT OFF
```

If you do this at global scope, it stays at whatever state you set it to until you change it.

If you do it at block scope, it reverts to the value outside the block when the block ends.

The initial value of the `FP_CONTRACT` pragma varies from system to system.

<sup>1</sup>Though the system defines `MATH_ERRNO` as 1 and `MATH_ERREXCEPT` as 2, it's best to always use their symbolic names. Just in case.

## 13.6 fpclassify()

Return the classification of a given floating point number.

### Synopsis

```
#include <math.h>

int fpclassify(any_floating_type x);
```

### Description

What kind of entity does this floating point number represent? What are the options?

We're used to floating point numbers being regular old things like 3.14 or 3490.0001.

But floating point numbers can also represent things like infinity. Or Not-A-Number (NaN). This function will let you know which type of floating point number the argument is.

This is a macro, so you can use it with `float`, `double`, `long double` or anything similar.

### Return Value

Returns one of these macros depending on the argument's classification:

| Classification | Description                    |
|----------------|--------------------------------|
| FP_INFINITE    | Number is infinite.            |
| FP_NAN         | Number is Not-A-Number (NaN).  |
| FP_NORMAL      | Just a regular number.         |
| FP_SUBNORMAL   | Number is a sub-normal number. |
| FP_ZERO        | Number is zero.                |

A discussion of subnormal numbers is beyond the scope of the guide, and is something that most devs go their whole lives without dealing with. In a nutshell, it's a way to represent really small numbers that might normally round down to zero. If you want to know more, see the Wikipedia page on denormal numbers<sup>2</sup>.

### Example

Print various number classifications.

```
1 #include <stdio.h>
2 #include <math.h>
3
4 const char *get_classification(double n)
5 {
6     switch (fpclassify(n)) {
7         case FP_INFINITE: return "infinity";
8         case FP_NAN: return "not a number";
9         case FP_NORMAL: return "normal";
10        case FP_SUBNORMAL: return "subnormal";
11        case FP_ZERO: return "zero";
```

<sup>2</sup>[https://en.wikipedia.org/wiki/Denormal\\_number](https://en.wikipedia.org/wiki/Denormal_number)



```

12     }
13
14     return "unknown";
15 }
16
17 int main(void)
18 {
19     printf("    1.23: %s\n", get_classification(1.23));
20     printf("    0.0: %s\n", get_classification(0.0));
21     printf("sqrt(-1): %s\n", get_classification(sqrt(-1)));
22     printf("1/tan(0): %s\n", get_classification(1/tan(0)));
23     printf("    1e-310: %s\n", get_classification(1e-310)); // very small!
24 }

```

Output<sup>3</sup>:

```

    1.23: normal
    0.0: zero
sqrt(-1): not a number
1/tan(0): infinity
    1e-310: subnormal

```

## See Also

`isfinite()`, `isinf()`, `isnan()`, `isnormal()`, `signbit()`

## 13.7 `isfinite()`, `isinf()`, `isnan()`, `isnormal()`

Return true if a number matches a classification.

### Synopsis

```

#include <math.h>

int isfinite(any_floating_type x);

int isinf(any_floating_type x);

int isnan(any_floating_type x);

int isnormal(any_floating_type x);

```

### Description

These are helper macros to `fpclassify()`. Being macros, they work on any floating point type.

<sup>3</sup>This is on my system. Some systems will have different points at which numbers become subnormal, or they might not support subnormal values at all.

| Macro                   | Description                                |
|-------------------------|--------------------------------------------|
| <code>isfinite()</code> | True if the number is not infinite or NaN. |
| <code>isinf()</code>    | True if the number is infinite.            |
| <code>isnan()</code>    | True if the number is Not-a-Number.        |
| <code>isnormal()</code> | True if the number is normal.              |

For more superficial discussion on normal and subnormal numbers, see `fpclassify()`.

## Return Value

Returns non-zero for true, and zero for false.

## Example

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf(" isfinite(1.23): %d\n", isfinite(1.23));    // 1
7      printf(" isinf(1/tan(0)): %d\n", isinf(1/tan(0))); // 1
8      printf(" isnan(sqrt(-1)): %d\n", isnan(sqrt(-1))); // 1
9      printf(" isnormal(1e-310): %d\n", isnormal(1e-310)); // 0
10 }
```

## See Also

`fpclassify()`, `signbit()`,

## 13.8 `signbit()`

Return the sign of a number.

## Synopsis

```

#include <math.h>

int signbit(any_floating_type x);
```

## Description

This macro takes any floating point number and returns a value indicating the sign of the number, positive or negative.

## Return Value

Returns `1` if the sign is negative, otherwise `0`.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%d\n", signbit(3490.0)); // 0
7     printf("%d\n", signbit(-37.0)); // 1
8 }
```

## See Also

fpclassify(), isfinite(), isinf(), isnan(), isnormal(), copysign()

---

## 13.9 acos(), acosf(), acosl()

Calculate the arc cosine of a number.

### Synopsis

```
#include <math.h>

double acos(double x);
float acosf(float x);
long double acosl(long double x);
```

### Description

Calculates the arc cosine of a number in radians. (That is, the value whose cosine is `x`.) The number must be in the range -1.0 to 1.0.

For those of you who don't remember, radians are another way of measuring an angle, just like degrees. To convert from degrees to radians or the other way around, use the following code:

```
pi = 3.14159265358979;
degrees = radians * 180 / pi;
radians = degrees * pi / 180;
```

### Return Value

Returns the arc cosine of `x`, unless `x` is out of range. In that case, `errno` will be set to EDOM and the return value will be NaN. The variants return different types.

**Example**

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      double acosx;
7      long double ldacosx;
8
9      acosx = acos(0.2);
10     ldacosx = acosl(0.3L);
11
12     printf("%f\n", acosx);
13     printf("%Lf\n", ldacosx);
14 }

```

**See Also**

asin(), atan(), atan2(), cos()

---

**13.10 asin(), asinf(), asinl()**

Calculate the arc sine of a number.

**Synopsis**

```

#include <math.h>

double asin(double x);
float asinf(float x);
long double asinl(long double x);

```

**Description**

Calculates the arc sine of a number in radians. (That is, the value whose sine is `x`.) The number must be in the range -1.0 to 1.0.

For those of you who don't remember, radians are another way of measuring an angle, just like degrees. To convert from degrees to radians or the other way around, use the following code:

```

pi = 3.14159265358979;
degrees = radians * 180 / pi;
radians = degrees * pi / 180;

```

**Return Value**

Returns the arc sine of `x`, unless `x` is out of range. In that case, `errno` will be set to `EDOM` and the return value will be NaN. The variants return different types.

**Example**

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      double asinx;
7      long double ldasinx;
8
9      asinx = asin(0.2);
10     ldasinx = asinl(0.3L);
11
12     printf("%f\n", asinx);
13     printf("%Lf\n", ldasinx);
14 }

```

**See Also**

acos(), atan(), atan2(), sin()

---

**13.11 atan(), atanf(), atanl(), atan2(), atan2f(), atan2l()**

Calculate the arc tangent of a number.

**Synopsis**

```

#include <math.h>

double atan(double x);
float atanf(float x);
long double atanl(long double x);

double atan2(double y, double x);
float atan2f(float y, float x);
long double atan2l(long double y, long double x);

```

**Description**

Calculates the arc tangent of a number in radians. (That is, the value whose tangent is `x`.)

The `atan2()` variants are pretty much the same as using `atan()` with `y/x` as the argument...except that `atan2()` will use those values to determine the correct quadrant of the result.

For those of you who don't remember, radians are another way of measuring an angle, just like degrees. To convert from degrees to radians or the other way around, use the following code:

```

pi = 3.14159265358979;
degrees = radians * 180 / pi;
radians = degrees * pi / 180;

```

## Return Value

The `atan()` functions return the arc tangent of `x`, which will be between  $\text{PI}/2$  and  $-\text{PI}/2$ . The `atan2()` functions return an angle between  $\text{PI}$  and  $-\text{PI}$ .

## Example

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      double atanx;
7      long double ldatanx;
8
9      atanx = atan(0.7);
10     ldatanx = atanl(0.3L);
11
12     printf("%f\n", atanx);
13     printf("%Lf\n", ldatanx);
14
15     atanx = atan2(7, 10);
16     ldatanx = atan2l(3L, 10L);
17
18     printf("%f\n", atanx);
19     printf("%Lf\n", ldatanx);
20 }
```

## See Also

`tan()`, `asin()`, `atan()`

---

## 13.12 `cos()`, `cosf()`, `cosl()`

Calculate the cosine of a number.

### Synopsis

```

#include <math.h>

double cos(double x)
float cosf(float x)
long double cosl(long double x)
```

### Description

Calculates the cosine of the value `x`, where `x` is in radians.

For those of you who don't remember, radians are another way of measuring an angle, just like degrees. To convert from degrees to radians or the other way around, use the following code:

```
pi = 3.14159265358979;  
degrees = radians * 180 / pi;  
radians = degrees * pi / 180;
```

## Return Value

Returns the cosine of `x`. The variants return different types.

## Example

```
1  #include <stdio.h>  
2  #include <math.h>  
3  
4  int main(void)  
5  {  
6      double cosx;  
7      long double ldcosx;  
8  
9      cosx = cos(3490.0); // round and round we go!  
10     ldcosx = cosl(3.490L);  
11  
12     printf("%f\n", cosx);  
13     printf("%Lf\n", ldcosx);  
14 }
```

## See Also

`sin()`, `tan()`, `acos()`

---

## 13.13 `sin()`, `sinf()`, `sinl()`

Calculate the sine of a number.

## Synopsis

```
#include <math.h>  
  
double sin(double x);  
float sinf(float x);  
long double sinl(long double x);
```

## Description

Calculates the sine of the value `x`, where `x` is in radians.

For those of you who don't remember, radians are another way of measuring an angle, just like degrees. To convert from degrees to radians or the other way around, use the following code:

```
pi = 3.14159265358979;  
degrees = radians * 180 / pi;  
radians = degrees * pi / 180;
```

## Return Value

Returns the sine of `x`. The variants return different types.

## Example

```
1  #include <stdio.h>  
2  #include <math.h>  
3  
4  int main(void)  
5  {  
6      double sinx;  
7      long double ldsinx;  
8  
9      sinx = sin(3490.0); // round and round we go!  
10     ldsinx = sinl(3.490L);  
11  
12     printf("%f\n", sinx);  
13     printf("%Lf\n", ldsinx);  
14 }
```

## See Also

`cos()`, `tan()`, `asin()`

---

## 13.14 `tan()`, `tanf()`, `tanl()`

Calculate the tangent of a number.

## Synopsis

```
#include <math.h>  
  
double tan(double x)  
float tanf(float x)  
long double tanl(long double x)
```

## Description

Calculates the tangent of the value `x`, where `x` is in radians.

For those of you who don't remember, radians are another way of measuring an angle, just like degrees. To convert from degrees to radians or the other way around, use the following code:



```
pi = 3.14159265358979;
degrees = radians * 180 / pi;
radians = degrees * pi / 180;
```

## Return Value

Returns the tangent of `x`. The variants return different types.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     double tanx;
7     long double ldtanx;
8
9     tanx = tan(3490.0); // round and round we go!
10    ldtanx = tanl(3.490L);
11
12    printf("%f\n", tanx);
13    printf("%Lf\n", ldtanx);
14 }
```

## See Also

`sin()`, `cos()`, `atan()`, `atan2()`

## 13.15 `acosh()`, `acoshf()`, `acoshl()`

Compute arc hyperbolic cosine.

## Synopsis

```
#include <math.h>

double acosh(double x);

float acoshf(float x);

long double acoshl(long double x);
```

## Description

Trig lovers can rejoice! C has arc hyperbolic cosine!

These functions return the nonnegative `acosh` of `x`, which must be greater than or equal to `1`.

## Return Value

Returns the arc hyperbolic cosince in the range  $[0, +\infty]$ .

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("acosh 1.8 = %f\n", acosh(1.8)); // 1.192911
7 }
```

## See Also

`asinh()`

---

## 13.16 `asinh()`, `asinhf()`, `asinhf_l()`

Compute arc hyperbolic sine.

## Synopsis

```
#include <math.h>

double asinh(double x);

float asinhf(float x);

long double asinhf_l(long double x);
```

## Description

Trig lovers can rejoice! C has arc hyperbolic sine!

These functions return the asinh of `x`.

## Return Value

Returns the arc hyperbolic sine.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
```

```
6     printf("asinh 1.8 = %f\n", asinh(1.8)); // 1.350441
7 }
```

## See Also

`acosh()`

---

## 13.17 `atanh()`, `atanhf()`, `atanhl()`

Compute the arc hyperbolic tangent.

### Synopsis

```
#include <math.h>

double atanh(double x);

float atanhf(float x);

long double atanhL(long double x);
```

### Description

These functions compute the arc hyperbolic tangent of `x`, which must be in the range  $[-1, +1]$ . Passing exactly  $-1$  or  $+1$  might result in a pole error.

### Return Value

Returns the arc hyperbolic tangent of `x`.

### Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("atanh 0.5 = %f\n", atanh(0.5)); // 0.549306
7 }
```

## See Also

`acosh()`, `asinh()`

---

## 13.18 `cosh()`, `coshf()`, `coshl()`

Compute the hyperbolic cosine.

## Synopsis

```
#include <math.h>

double cosh(double x);

float coshf(float x);

long double coshl(long double x);
```

## Description

These functions predictably compute the hyperbolic cosine of `x`. A range error might occur if `x` is too large.

## Return Value

Returns the hyperbolic cosine of `x`.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("cosh 0.5 = %f\n", cosh(0.5)); // 1.127626
7 }
```

## See Also

`sinh()`, `tanh()`

---

## 13.19 `sinh()`, `sinhf()`, `sinhl()`

Compute the hyperbolic sine.

## Synopsis

```
#include <math.h>

double sinh(double x);

float sinhf(float x);

long double sinhl(long double x);
```

## Description

These functions predictably compute the hyperbolic sine of `x`. A range error might occur if `x` is too large.

## Return Value

Returns the hyperbolic sine of `x`.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("sinh 0.5 = %f\n", sinh(0.5)); // 0.521095
7 }
```

## See Also

`sinh()`, `tanh()`

---

## 13.20 `tanh()`, `tanhf()`, `tanhL()`

Compute the hyperbolic tangent.

## Synopsis

```
#include <math.h>

double tanh(double x);

float tanhf(float x);

long double tanhL(long double x);
```

## Description

These functions predictably compute the hyperbolic tangent of `x`.

Mercifully, this is the last trig-related man page I'm going to write.

## Return Value

Returns the hyperbolic tangent of `x`.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
```

```

6     printf("tanh 0.5 = %f\n", tanh(0.5)); // 0.462117
7 }

```

**See Also**

cosh(), sinh()

---

**13.21 exp(), expf(), expl()**

Compute  $e$  raised to a power.

**Synopsis**

```

#include <math.h>

double exp(double x);

float expf(float x);

long double expl(long double x);

```

**Description**

Compute  $e^x$  where  $e$  is Euler's number<sup>4</sup>.

The number  $e$  is named after Leonard Euler, born April 15, 1707, who is responsible, among other things, for making this reference page longer than it needed to be.

**Return Value**

Returns  $e^x$ .

**Example**

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("exp(1) = %f\n", exp(1)); // 2.718282
7      printf("exp(2) = %f\n", exp(2)); // 7.389056
8  }

```

**See Also**

exp2(), expm1(), pow(), log()

---

<sup>4</sup>[https://en.wikipedia.org/wiki/E\\_\(mathematical\\_constant\)](https://en.wikipedia.org/wiki/E_(mathematical_constant))

## 13.22 `exp2()`, `exp2f()`, `exp2l()`

Compute 2 to a power.

### Synopsis

```
#include <math.h>

double exp2(double x);

float exp2f(float x);

long double exp2l(long double x);
```

### Description

These functions raise 2 to a power. Very exciting, since computers are all about twos-to-powers!

These are likely to be faster than using `pow()` to do the same thing.

They support fractional exponents, as well.

A range error occurs if `x` is too large.

### Return Value

`exp2()` returns  $2^x$ .

### Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("2^3 = %f\n", exp2(3));      // 2^3 = 8.000000
7     printf("2^8 = %f\n", exp2(8));      // 2^8 = 256.000000
8     printf("2^0.5 = %f\n", exp2(0.5));  // 2^0.5 = 1.414214
9 }
```

### See Also

`exp()`, `pow()`

---

## 13.23 `expm1()`, `expm1f()`, `expm1l()`

Compute  $e^x - 1$ .

## Synopsis

```
#include <math.h>

double expm1(double x);

float expm1f(float x);

long double expm1l(long double x);
```

## Description

This is just like `exp()` except—*plot twist!*—it computes that result minus one.

For more discussion about what  $e$  is, see the `exp()` man page.

If  $x$  is giant, a range error might occur.

For small values of  $x$  near zero, `expm1(x)` might be more accurate than computing `exp(x) - 1`.

## Return Value

Returns  $e^x - 1$ .

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%f\n", expm1(2.34)); // 9.381237
7 }
```

## See Also

`exp()`

---

## 13.24 frexp(), frexpf(), frexpl()

Break a number into its fraction part and exponent (as a power of 2).

## Synopsis

```
#include <math.h>

double frexp(double value, int *exp);

float frexpf(float value, int *exp);
```



```
long double frexpl(long double value, int *exp);
```

## Description

If you have a floating point number, you can break it into its fractional part and exponent part (as a power of 2).

For example, if you have the number 1234.56, this can be represented as a multiple of a power of 2 like so:

$$1234.56 = 0.6028125 \times 2^{11}$$

And you can use this function to get the 0.6028125 and 11 parts of that equation.

As for why, I have a simple answer: I don't know. I can't find a use. K&R2 and everyone else I can find just says *how* to use it, but not *why* you might want to.

The C99 Rationale document says:

The functions `frexp`, `ldexp`, and `modf` are primitives used by the remainder of the library. There was some sentiment for dropping them for the same reasons that `ecvt`, `fcvt`, and `gcvt` were dropped, but their adherents rescued them for general use. Their use is problematic: on non-binary architectures, `ldexp` may lose precision and `frexp` may be inefficient.

So there you have it. If you need it.

## Return Value

`frexp()` returns the fractional part of `value` in the range 0.5 (inclusive) to 1 (exclusive), or 0. And it stores the exponent power-of-2 in the variable pointed to by `exp`.

If you pass in zero, the return value and the variable `exp` points to are both zero.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     double frac;
7     int expt;
8
9     frac = frexp(1234.56, &expt);
10    printf("1234.56 = %.7f x 2^%d\n", frac, expt);
11 }
```

Output:

```
1234.56 = 0.6028125 x 2^11
```

## See Also

`ldexp()`, `ilogb()`, `modf()`

---

## 13.25 `ilogb()`, `ilogbf()`, `ilogbl()`

Return the exponent of a floating point number.

### Synopsis

```
#include <math.h>

int ilogb(double x);

int ilogbf(float x);

int ilogbl(long double x);
```

### Description

This gives you the exponent of the given number... it's a little weird, because the exponent depends on the value of `FLT_RADIX`. Now, this is very often 2—but no guarantees!

It actually returns  $\log_r |x|$  where  $r$  is `FLT_RADIX`.

Domain or range errors might occur for invalid values of `x`, or for return values that are outside the range of the return type.

### Return Value

The exponent of the absolute value of the given number, depending on `FLT_RADIX`.

Specifically  $\log_r |x|$  where  $r$  is `FLT_RADIX`.

If you pass in 0, it'll return `FP_ILOGB0`.

If you pass in infinity, it'll return `INT_MAX`.

If you pass in NaN, it'll return `FP_ILOGBNAN`.

The spec goes on to say that the value of `FP_ILOGB0` will be either `INT_MIN` or `-INT_MAX`. And the value of `FP_ILOGBNAN` shall be either `INT_MAX` or `INT_MIN`, if that's useful in any way.

### Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%d\n", ilogb(257)); // 8
7     printf("%d\n", ilogb(256)); // 8
8     printf("%d\n", ilogb(255)); // 7
9 }
```

### See Also

`frexp()`, `logb()`

---

## 13.26 ldexp(), ldexpf(), ldexpl()

Multiply a number by an integral power of 2.

### Synopsis

```
#include <math.h>

double ldexp(double x, int exp);

float ldexpf(float x, int exp);

long double ldexpl(long double x, int exp);
```

### Description

These functions multiply the given number `x` by 2 raised to the `exp` power.

### Return Value

Returns  $x \times 2^{exp}$ .

### Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("1 x 2^10 = %f\n", ldexp(1, 10));
7     printf("5.67 x 2^7 = %f\n", ldexp(5.67, 7));
8 }
```

Output:

```
1 x 2^10 = 1024.000000
5.67 x 2^7 = 725.760000
```

### See Also

`exp()`

---

## 13.27 log(), logf(), logl()

Compute the natural logarithm.

## Synopsis

```
#include <math.h>

double log(double x);

float logf(float x);

long double logl(long double x);
```

## Description

Natural logarithms! And there was much rejoicing.

These compute the base- $e$  logarithm of a number,  $\log_e x$ ,  $\ln x$ .

In other words, for a given  $x$ , solves  $x = e^y$  for  $y$ .

## Return Value

The base- $e$  logarithm of the given value,  $\log_e x$ ,  $\ln x$ .

## Example

```
1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      const double e = 2.718281828459045;
7
8      printf("%f\n", log(3490.2)); // 8.157714
9      printf("%f\n", log(e));     // 1.000000
10 }
```

## See Also

`exp()`, `log10()`, `log1p()`

---

## 13.28 `log10()`, `log10f()`, `log10l()`

Compute the log-base-10 of a number.

## Synopsis

```
#include <math.h>

double log10(double x);

float log10f(float x);
```

```
long double log10l(long double x);
```

## Description

Just when you thought you might have to use Laws of Logarithms to compute this, here's a function coming out of the blue to save you.

These compute the base-10 logarithm of a number,  $\log_{10} x$ .

In other words, for a given  $x$ , solves  $x = 10^y$  for  $y$ .

## Return Value

Returns the log base-10 of `x`,  $\log_{10} x$ .

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%f\n", log10(3490.2)); // 3.542850
7     printf("%f\n", log10(10));    // 1.000000
8 }
```

## See Also

`pow()`, `log()`

---

## 13.29 `log1p()`, `log1pf()`, `log1pl()`

Compute the natural logarithm of a number plus 1.

## Synopsis

```
#include <math.h>

double log1p(double x);

float log1pf(float x);

long double log1pl(long double x);
```

## Description

This computes  $\log_e(1 + x)$ ,  $\ln(1 + x)$ .

This works just like calling:

```
log(1 + x)
```

except it could be more accurate for small values of  $x$ .

So if your  $x$  is small magnitude, use this.

## Return Value

Returns  $\log_e(1 + x)$ ,  $\ln(1 + x)$ .

## Example

Compute some big and small logarithm values to see the difference between `log1p()` and `log()`:

```
1 #include <stdio.h>
2 #include <float.h> // for LDBL_DECIMAL_DIG
3 #include <math.h>
4
5 int main(void)
6 {
7     printf("Big log1p() : %.Lf\n", LDBL_DECIMAL_DIG-1, log1pl(9));
8     printf("Big log()   : %.Lf\n", LDBL_DECIMAL_DIG-1, logl(1 + 9));
9
10    printf("Small log1p(): %.Lf\n", LDBL_DECIMAL_DIG-1, log1pl(0.01));
11    printf("Small log()  : %.Lf\n", LDBL_DECIMAL_DIG-1, logl(1 + 0.01));
12 }
```

Output on my system:

```
Big log1p() : 2.30258509299404568403
Big log()   : 2.30258509299404568403
Small log1p(): 0.00995033085316808305
Small log()  : 0.00995033085316809164
```

## See Also

`log()`

## 13.30 `log2()`, `log2f()`, `log2l()`

Compute the base-2 logarithm of a number.

## Synopsis

```
#include <math.h>

double log2(double x);

float log2f(float x);
```

```
long double log2l(long double x);
```

## Description

Wow! Were you thinking we were done with the logarithm functions? We're only getting started!

This one computes  $\log_2 x$ . That is, computes  $y$  that satisfies  $x = 2^y$ .

Love me those powers of 2!

## Return Value

Returns the base-2 logarithm of the given value,  $\log_2 x$ .

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%f\n", log2(3490.2)); // 11.769094
7     printf("%f\n", log2(256));   // 8.000000
8 }
```

## See Also

`log()`

---

## 13.31 `logb()`, `logbf()`, `logbl()`

Extract the exponent of a number given `FLT_RADIX`.

## Synopsis

```
#include <math.h>

double logb(double x);

float logbf(float x);

long double logbl(long double x);
```

## Description

This function returns the whole number portion of the exponent of the number with radix `FLT_RADIX`, namely the whole number portion  $\log_r |x|$  where  $r$  is `FLT_RADIX`. Fractional numbers are truncated.

If the number is subnormal<sup>5</sup>, `logb()` treats it as if it were normalized.

<sup>5</sup>[https://en.wikipedia.org/wiki/Denormal\\_number](https://en.wikipedia.org/wiki/Denormal_number)

If  $x$  is 0, there could be a domain error or pole error.

## Return Value

This function returns the whole number portion of  $\log_r |x|$  where  $r$  is FLT\_RADIX.

## Example

```

1  #include <stdio.h>
2  #include <float.h>  // For FLT_RADIX
3  #include <math.h>
4
5  int main(void)
6  {
7      printf("FLT_RADIX = %d\n", FLT_RADIX);
8      printf("%f\n", logb(3490.2));
9      printf("%f\n", logb(256));
10 }
```

Output:

```

FLT_RADIX = 2
11.000000
8.000000
```

## See Also

ilogb()

## 13.32 modf(), modff(), modfl()

Extract the integral and fractional parts of a number.

## Synopsis

```

#include <math.h>

double modf(double value, double *iptr);

float modff(float value, float *iptr);

long double modfl(long double value, long double *iptr);
```

## Description

If you have a floating point number, like 123.456, this function will extract the integral part (123.0) and the fractional part (0.456). It's total coincidence that this is exactly the plot for the latest Jason Statham action spectacular.

Both the integral part and fractional parts keep the sign of the passed in value.



The integral part is stored in the address pointed to by `iptr`.

See the note in `frexp()` regarding why this is in the library.

## Return Value

These functions return the fractional part of the number. The integral part is stored in the address pointed to by `iptr`. Both the integral and fractional parts preserve the sign of the passed-in `value`.

## Example

```

1  #include <stdio.h>
2  #include <math.h>
3
4  void print_parts(double x)
5  {
6      double i, f;
7
8      f = modf(x, &i);
9
10     printf("Entire number   : %f\n", x);
11     printf("Integral part   : %f\n", i);
12     printf("Fractional part: %f\n", f);
13 }
14
15 int main(void)
16 {
17     print_parts(123.456);
18     print_parts(-123.456);
19 }
```

Output:

```

Entire number   : 123.456000
Integral part   : 123.000000
Fractional part: 0.456000

Entire number   : -123.456000
Integral part   : -123.000000
Fractional part: -0.456000
```

## See Also

`frexp()`

---

## 13.33 `scalbn()`, `scalbnf()`, `scalbnl()` `scalbln()`, `scalblnf()`, `scalblnl()`

Efficiently compute  $x \times r^n$ , where  $r$  is `FLT_RADIX`.

## Synopsis

```
#include <math.h>

double scalbn(double x, int n);

float scalbnf(float x, int n);

long double scalbnl(long double x, int n);

double scalbln(double x, long int n);

float scalblnf(float x, long int n);

long double scalblnl(long double x, long int n);
```

## Description

These functions efficiently compute  $x \times r^n$ , where  $r$  is `FLT_RADIX`.

If `FLT_RADIX` happens to be 2 (no guarantees!), then this works like `exp2()`.

The name of this function should have an obvious meaning to you. Clearly they all start with the prefix “scalb” which means...

...OK, I confess! I have no idea what it means. My searches are futile!

But let’s look at the suffixes:

| Suffix | Meaning                                     |
|--------|---------------------------------------------|
| n      | scalbn()—exponent n is an int               |
| nf     | scalbnf()—float version of scalbn()         |
| nl     | scalbnl()—long double version of scalbn()   |
| ln     | scalbln()—exponent n is a long int          |
| lnf    | scalblnf()—float version of scalbln()       |
| lnl    | scalblnl()—long double version of scalbln() |

So while I’m still in the dark about “scalb”, at least I have that part down.

A range error might occur for large values.

## Return Value

Returns  $x \times r^n$ , where  $r$  is `FLT_RADIX`.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3 #include <float.h>
4
5 int main(void)
6 {
```

```

7   printf("FLT_RADIX = %d\n\n", FLT_RADIX);
8   printf("scalbn(3, 8)      = %f\n", scalbn(2, 8));
9   printf("scalbnf(10.2, 20) = %f\n", scalbnf(10.2, 20));
10  }

```

Output on my system:

```

FLT_RADIX = 2

scalbn(3, 8)      = 512.000000
scalbn(10.2, 20.7) = 10695475.200000

```

## See Also

`exp2()`, `pow()`

## 13.34 `cbrt()`, `cbrtf()`, `cbrtl()`

Compute the cube root.

### Synopsis

```

#include <math.h>

double cbrt(double x);

float cbrtf(float x);

long double cbrtl(long double x);

```

### Description

Computes the cube root of `x`,  $x^{1/3}$ ,  $\sqrt[3]{x}$ .

### Return Value

Returns the cube root of `x`,  $x^{1/3}$ ,  $\sqrt[3]{x}$ .

### Example

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("cbrt(1729.03) = %f\n", cbrt(1729.03));
7  }

```

Output:

```
cbirt(1729.03) = 12.002384
```

## See Also

`sqrt()`, `pow()`

---

## 13.35 `fabs()`, `fabsf()`, `fabsl()`

Compute the absolute value.

### Synopsis

```
#include <math.h>

double fabs(double x);

float fabsf(float x);

long double fabsl(long double x);
```

### Description

These functions straightforwardly return the absolute value of `x`, that is  $|x|$ .

If you're rusty on your absolute values, all it means is that the result will be positive, even if `x` is negative. It's just strips negative signs off.

### Return Value

Returns the absolute value of `x`,  $|x|$ .

### Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("fabs(3490.0) = %f\n", fabs(3490.0)); // 3490.000000
7     printf("fabs(-3490.0) = %f\n", fabs(3490.0)); // 3490.000000
8 }
```

## See Also

`abs()`, `copysign()`, `imaxabs()`

---

### 13.36 `hypot()`, `hypotf()`, `hypotl()`

Compute the length of the hypotenuse of a triangle.

#### Synopsis

```
#include <math.h>

double hypot(double x, double y);

float hypotf(float x, float y);

long double hypotl(long double x, long double y);
```

#### Description

Pythagorean Theorem<sup>6</sup> fans rejoice! This is the function you've been waiting for!

If you know the lengths of the two sides of a right triangle, `x` and `y`, you can compute the length of the hypotenuse (the longest, diagonal side) with this function.

In particular, it computes the square root of the sum of the squares of the sides:  $\sqrt{x^2 + y^2}$ .

#### Return Value

Returns the length of the hypotenuse of a right triangle with side lengths `x` and `y`:  $\sqrt{x^2 + y^2}$ .

#### Example

```
1 printf("%f\n", hypot(3, 4)); // 5.000000
```

#### See Also

`sqrt()`

---

### 13.37 `pow()`, `powf()`, `powl()`

Compute a value raised to a power.

#### Synopsis

```
#include <math.h>

double pow(double x, double y);

float powf(float x, float y);

long double powl(long double x, long double y);
```

<sup>6</sup>[https://en.wikipedia.org/wiki/Pythagorean\\_theorem](https://en.wikipedia.org/wiki/Pythagorean_theorem)

**Description**

Computes  $x$  raised to the  $y$ th power:  $x^y$ .

These arguments can be fractional.

**Return Value**

Returns  $x$  raised to the  $y$ th power:  $x^y$ .

A domain error can occur if:

- $x$  is a finite negative number and  $y$  is a finite non-integer
- $x$  is zero and  $y$  is zero.

A domain error or pole error can occur if  $x$  is zero and  $y$  is negative.

A range error can occur for large values.

**Example**

```
1 printf("%f\n", pow(3, 4));    // 3^4    = 81.000000
2 printf("%f\n", pow(2, 0.5)); // sqrt 2 = 1.414214
```

**See Also**

`exp()`, `exp2()`, `sqrt()`, `cbrt()`

**13.38 sqrt()**

Calculate the square root of a number.

**Synopsis**

```
#include <math.h>

double sqrt(double x);

float sqrtf(float x);

long double sqrtl(long double x);
```

**Description**

Computes the square root of a number:  $\sqrt{x}$ . To those of you who don't know what a square root is, I'm not going to explain. Suffice it to say, the square root of a number delivers a value that when squared (multiplied by itself) results in the original number.

Ok, fine—I did explain it after all, but only because I wanted to show off. It's not like I'm giving you examples or anything, such as the square root of nine is three, because when you multiply three by three you get nine, or anything like that. No examples. I hate examples!

And I suppose you wanted some actual practical information here as well. You can see the usual trio of functions here—they all compute square root, but they take different types as arguments. Pretty straightforward, really.

A domain error occurs if `x` is negative.

## Return Value

Returns (and I know this must be something of a surprise to you) the square root of `x`:  $\sqrt{x}$ .

## Example

```

1 // example usage of sqrt()
2
3 float something = 10;
4
5 double x1 = 8.2, y1 = -5.4;
6 double x2 = 3.8, y2 = 34.9;
7 double dx, dy;
8
9 printf("square root of 10 is %.2f\n", sqrtf(something));
10
11 dx = x2 - x1;
12 dy = y2 - y1;
13 printf("distance between points (x1, y1) and (x2, y2): %.2f\n",
14        sqrt(dx*dx + dy*dy));

```

And the output is:

```

square root of 10 is 3.16
distance between points (x1, y1) and (x2, y2): 40.54

```

## See Also

`hypot()`, `pow()`

## 13.39 `erf()`, `erff()`, `erfl()`

Compute the error function of the given value.

## Synopsis

```

#include <math.h>

double erfc(double x);

float erfcf(float x);

long double erfcl(long double x);

```

## Description

These functions compute the error function<sup>7</sup> of a value.

## Return Value

Returns the error function of `x`:

$$\frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

## Example

```
1 for (float i = -2; i <= 2; i += 0.5)
2     printf("%.1f: %f\n", i, erf(i));
```

Output:

```
-2.0: -0.995322
-1.5: -0.966105
-1.0: -0.842701
-0.5: -0.520500
 0.0: 0.000000
 0.5: 0.520500
 1.0: 0.842701
 1.5: 0.966105
 2.0: 0.995322
```

## See Also

`erfc()`

---

## 13.40 `erfc()`, `erfcf()`, `erfcfcl()`

Compute the complementary error function of a value.

## Synopsis

```
#include <math.h>

double erfc(double x);

float erfcf(float x);

long double erfcfcl(long double x);
```

<sup>7</sup>[https://en.wikipedia.org/wiki/Error\\_function](https://en.wikipedia.org/wiki/Error_function)



## Description

These functions compute the complementary error function<sup>8</sup> of a value.

This is the same as:

```
1 - erf(x)
```

A range error can occur if `x` is too large.

## Return Value

Returns `1 - erf(x)`, namely:

$$\frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} dt$$

## Example

```
1 for (float i = -2; i <= 2; i += 0.5)
2     printf("%.1f: %f\n", i, erfc(i));
```

Output:

```
-2.0: 1.995322
-1.5: 1.966105
-1.0: 1.842701
-0.5: 1.520500
 0.0: 1.000000
 0.5: 0.479500
 1.0: 0.157299
 1.5: 0.033895
 2.0: 0.004678
```

## See Also

`erf()`

## 13.41 `lgamma()`, `lgammaf()`, `lgammal()`

Compute the natural logarithm of the absolute value of  $\Gamma(x)$ .

## Synopsis

```
#include <math.h>

double lgamma(double x);

float lgammaf(float x);
```

<sup>8</sup>[https://en.wikipedia.org/wiki/Error\\_function](https://en.wikipedia.org/wiki/Error_function)

```
long double lgammal(long double x);
```

### Description

Compute the natural log of the absolute value of gamma<sup>9</sup>  $x$ ,  $\log_e |\Gamma(x)|$ .

A range error can occur if  $x$  is too large.

A pole error can occur if  $x$  is non-positive.

### Return Value

Returns  $\log_e |\Gamma(x)|$ .

### Example

```
1 for (float i = 0.5; i <= 4; i += 0.5)
2     printf("%.1f: %f\n", i, lgamma(i));
```

Output:

```
0.5: 0.572365
1.0: 0.000000
1.5: -0.120782
2.0: 0.000000
2.5: 0.284683
3.0: 0.693147
3.5: 1.200974
4.0: 1.791759
```

### See Also

`tgamma()`

## 13.42 `tgamma()`, `tgammaf()`, `tgammal()`

Compute the gamma function,  $\Gamma(x)$ .

### Synopsis

```
#include <math.h>

double tgamma(double x);

float tgammaf(float x);

long double tgammal(long double x);
```

<sup>9</sup>[https://en.wikipedia.org/wiki/Gamma\\_function](https://en.wikipedia.org/wiki/Gamma_function)

## Description

Computes the gamma function<sup>10</sup> of `x`,  $\Gamma(x)$ .

A domain or pole error might occur if `x` is non-positive.

A range error might occur if `x` is too large or too small.

## Return Value

Returns the gamma function of `x`,  $\Gamma(x)$ .

## Example

```
1 for (float i = 0.5; i <= 4; i += 0.5)
2     printf("%.1f: %f\n", i, tgamma(i));
```

Output:

```
0.5: 1.772454
1.0: 1.000000
1.5: 0.886227
2.0: 1.000000
2.5: 1.329340
3.0: 2.000000
3.5: 3.323351
4.0: 6.000000
```

## See Also

`lgamma()`

---

## 13.43 `ceil()`, `ceilf()`, `ceil()`

Ceiling—return the next whole number not smaller than the given number.

## Synopsis

```
#include <math.h>

double ceil(double x);

float ceilf(float x);

long double ceil(long double x);
```

---

<sup>10</sup>[https://en.wikipedia.org/wiki/Gamma\\_function](https://en.wikipedia.org/wiki/Gamma_function)

## Description

Returns the ceiling of the `x`:  $\lceil x \rceil$ .

This is the next whole number not smaller than `x`.

Beware this minor dragon: it's not just "rounding up". Well, it is for positive numbers, but negative numbers effectively round toward zero. (Because the ceiling function is headed for the next largest whole number and  $-4$  is larger than  $-5$ .)

## Return Value

Returns the next largest whole number larger than `x`.

## Example

Notice for the negative numbers it heads toward zero, i.e. toward the next largest whole number—just like the positives head toward the next largest whole number.

```
1 printf("%f\n", ceil(4.0)); // 4.000000
2 printf("%f\n", ceil(4.1)); // 5.000000
3 printf("%f\n", ceil(-2.0)); // -2.000000
4 printf("%f\n", ceil(-2.1)); // -2.000000
5 printf("%f\n", ceil(-3.1)); // -3.000000
```

## See Also

`floor()`, `round()`

## 13.44 `floor()`, `floorf()`, `floorl()`

Compute the largest whole number not larger than the given value.

## Synopsis

```
#include <math.h>
double floor(double x);
float floorf(float x);
long double floorl(long double x);
```

## Description

Returns the floor of the value:  $\lfloor x \rfloor$ . This is the opposite of `ceil()`.

This is the largest whole number that is not greater than `x`.

For positive numbers, this is like rounding down: `4.5` becomes `4.0`.

For negative numbers, it's like rounding up: `-3.6` becomes `-4.0`.

In both cases, those results are the largest whole number not bigger than the given number.

## Return Value

Returns the largest whole number not greater than  $x$ :  $\lfloor x \rfloor$ .

## Example

Note how the negative numbers effectively round away from zero, unlike the positives.

```
1 printf("%f\n", floor(4.0)); // 4.000000
2 printf("%f\n", floor(4.1)); // 4.000000
3 printf("%f\n", floor(-2.0)); // -2.000000
4 printf("%f\n", floor(-2.1)); // -3.000000
5 printf("%f\n", floor(-3.1)); // -4.000000
```

## See Also

`ceil()`, `round()`

## 13.45 `nearbyint()`, `nearbyintf()`, `nearbyintl()`

Rounds a value in the current rounding direction.

## Synopsis

```
#include <math.h>

double nearbyint(double x);

float nearbyintf(float x);

long double nearbyintl(long double x);
```

## Description

This function rounds  $x$  to the nearest integer in the current rounding direction.

The rounding direction can be set with `fesetround()` in `<fenv.h>`.

`nearbyint()` won't raise the "inexact" floating point exception.

## Return Value

Returns  $x$  rounded in the current rounding direction.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3 #include <fenv.h>
4
```

```

5 int main(void)
6 {
7     #pragma STDC FENV_ACCESS ON           // If supported
8
9     fesetround(FE_TONEAREST);             // round to nearest
10
11     printf("%f\n", nearbyint(3.14));      // 3.000000
12     printf("%f\n", nearbyint(3.74));      // 4.000000
13
14     fesetround(FE_TOWARDZERO);            // round toward zero
15
16     printf("%f\n", nearbyint(1.99));      // 1.000000
17     printf("%f\n", nearbyint(-1.99));     // -1.000000
18 }

```

## See Also

`rint()`, `lrint()`, `round()`, `fesetround()`, `fegetround()`

---

## 13.46 `rint()`, `rintf()`, `rintl()`

Rounds a value in the current rounding direction.

## Synopsis

```

#include <math.h>

double rint(double x);

float rintf(float x);

long double rintl(long double x);

```

## Description

This works just like `nearbyint()` except that it can raise the “inexact” floating point exception.

## Return Value

Returns `x` rounded in the current rounding direction.

## Example

```

1 #include <stdio.h>
2 #include <math.h>
3 #include <fenv.h>
4
5 int main(void)

```

```

6 {
7     #pragma STDC FENV_ACCESS ON
8
9     fesetround(FE_TONEAREST);
10
11     printf("%f\n", rint(3.14)); // 3.000000
12     printf("%f\n", rint(3.74)); // 4.000000
13
14     fesetround(FE_TOWARDZERO);
15
16     printf("%f\n", rint(1.99)); // 1.000000
17     printf("%f\n", rint(-1.99)); // -1.000000
18 }

```

### See Also

`nearbyint()`, `lrint()`, `round()`, `fesetround()`, `fegetround()`

---

## 13.47 `lrint()`, `lrintf()`, `lrintl()`, `llrint()`, `llrintf()`, `llrintl()`

Returns `x` rounded in the current rounding direction as an integer.

### Synopsis

```

#include <math.h>

long int lrint(double x);
long int lrintf(float x);
long int lrintl(long double x);

long long int llrint(double x);
long long int llrintf(float x);
long long int llrintl(long double x);

```

### Description

Round a floating point number in the current rounding direction, but this time return an integer instead of a float. You know, just to mix it up.

These come in two variants:

- `lrint()`—returns `long int`
- `llrint()`—returns `long long int`

If the result doesn't fit in the return type, a domain or range error might occur.

### Return Value

The value of `x` rounded to an integer in the current rounding direction.

## Example

```
1  #include <stdio.h>
2  #include <math.h>
3  #include <fenv.h>
4
5  int main(void)
6  {
7      #pragma STDC FENV_ACCESS ON
8
9      fesetround(FE_TONEAREST);
10
11     printf("%ld\n", lrint(3.14));    // 3
12     printf("%ld\n", lrint(3.74));    // 4
13
14     fesetround(FE_TOWARDZERO);
15
16     printf("%ld\n", lrint(1.99));    // 1
17     printf("%ld\n", lrint(-1.99));   // -1
18 }
```

## See Also

`nearbyint()`, `rint()`, `round()`, `fesetround()`, `fegetround()`

---

## 13.48 `round()`, `roundf()`, `roundl()`

Round a number in the good old-fashioned way.

## Synopsis

```
#include <math.h>

double round(double x);

float roundf(float x);

long double roundl(long double x);
```

## Description

Rounds a number to the nearest whole value.

In case of halvesies, rounds away from zero (i.e. “round up” in magnitude).

The current rounding direction’s Jedi mind tricks don’t work on this function.

## Return Value

The rounded value of `x`.



**Example**

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("%f\n", round(3.14)); // 3.000000
7      printf("%f\n", round(3.5));  // 4.000000
8
9      printf("%f\n", round(-1.5)); // -2.000000
10     printf("%f\n", round(-1.14)); // -1.000000
11 }

```

**See Also**

lround(), nearbyint(), rint(), lrint(), trunc()

---

### 13.49 lround(), lroundf(), lroundl() llround(), llroundf(), llroundl()

Round a number in the good old-fashioned way, returning an integer.

**Synopsis**

```

#include <math.h>

long int lround(double x);
long int lroundf(float x);
long int lroundl(long double x);

long long int llround(double x);
long long int llroundf(float x);
long long int llroundl(long double x);

```

**Description**

These are just like `round()` except they return integers.

Halfway values round away from zero, e.g. 1.5 rounds to 2 and −1.5 rounds to −2.

The functions are grouped by return type:

- `lround()`—returns a `long int`
- `llround()`—returns a `long long int`

If the rounded value can't fit in the return type, a domain or range error can occur.

**Return Value**

Returns the rounded value of `x` as an integer.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%ld\n", lround(3.14));    // 3
7     printf("%ld\n", lround(3.5));    // 4
8
9     printf("%ld\n", lround(-1.5));    // -2
10    printf("%ld\n", lround(-1.14));   // -1
11 }
```

## See Also

`round()`, `nearbyint()`, `rint()`, `lrint()`, `trunc()`

---

## 13.50 `trunc()`, `truncf()`, `trunc1()`

Truncate the fractional part off a floating point value.

## Synopsis

```
#include <math.h>

double trunc(double x);

float truncf(float x);

long double trunc1(long double x);
```

## Description

These functions just drop the fractional part of a floating point number. Boom.

In other words, they always round toward zero.

## Return Value

Returns the truncated floating point number.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
```

```
6     printf("%f\n", trunc(3.14));    // 3.000000
7     printf("%f\n", trunc(3.8));    // 3.000000
8
9     printf("%f\n", trunc(-1.5));    // -1.000000
10    printf("%f\n", trunc(-1.14));   // -1.000000
11 }
```

## See Also

`round()`, `lround()`, `nearbyint()`, `rint()`, `lrint()`

---

## 13.51 `fmod()`, `fmodf()`, `fmodl()`

Compute the floating point remainder.

### Synopsis

```
#include <math.h>

double fmod(double x, double y);

float fmodf(float x, float y);

long double fmodl(long double x, long double y);
```

### Description

Returns the remainder of  $\frac{x}{y}$ . The result will have the same sign as `x`.

Under the hood, the computation performed is:

```
x - trunc(x / y) * y
```

But it might be easier just to think of the remainder.

### Return Value

Returns the remainder of  $\frac{x}{y}$  with the same sign as `x`.

### Example

```

1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%f\n", fmod(-9.2, 5.1)); // -4.100000
7     printf("%f\n", fmod(9.2, 5.1));  //  4.100000
8 }

```

## See Also

`remainder()`

## 13.52 `remainder()`, `remainderf()`, `remainderl()`

Compute the remainder IEC 60559-style.

## Synopsis

```

#include <math.h>

double remainder(double x, double y);

float remainderf(float x, float y);

long double remainderl(long double x, long double y);

```

## Description

This is similar to `fmod()`, but not quite the same. `fmod()` is probably what you're after if you're expecting remainders to wrap around like an odometer.

The C spec quotes IEC 60559 on how this works:

When  $y \neq 0$ , the remainder  $r = x \text{ REM } y$  is defined regardless of the rounding mode by the mathematical relation  $r = x - ny$ , where  $n$  is the integer nearest the exact value of  $x/y$ ; whenever  $|n - x/y| = 1/2$ , then  $n$  is even. If  $r = 0$ , its sign shall be that of  $x$ .

Hope that clears it up!

OK, maybe not. Here's the upshot:

You know how if you `fmod()` something by, say `2.0` you get a result that is somewhere between `0.0` and `2.0`? And how if you just increase the number that you're modding by `2.0`, you can see the result climb up to `2.0` and then wrap around to `0.0` like your car's odometer?

`remainder()` works just like that, except if `y` is `2.0`, it wraps from `-1.0` to `1.0` instead of from `0.0` to `2.0`.

In other words, the range of the function runs from `-y/2` to `y/2`. Contrasted to `fmod()` that runs from `0.0` to `y`, `remainder()`'s output is just shifted down half a `y`.

And zero-remainder-anything is `0`.

Except if `y` is zero, the function might return zero or a domain error might occur.

## Return Value

The IEC 60559 result of `x-remainder-y`.

## Example

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("%f\n", remainder(3.7, 4)); // -0.300000
7      printf("%f\n", remainder(4.3, 4)); //  0.300000
8  }
```

## See Also

`fmod()`, `remquo()`

---

## 13.53 `remquo()`, `remquof()`, `remquol()`

Compute the remainder and (some of the) quotient.

## Synopsis

```

#include <math.h>

double remquo(double x, double y, int *quo);

float remquof(float x, float y, int *quo);

long double remquol(long double x, long double y, int *quo);
```

## Description

This is a funky little thing.

First of all, the return value is the remainder, the same as the `remainder()` function, so check that out.

And the quotient comes back in the `quo` pointer.

Or at least *some of it* does. You'll get at least 3 bits worth of the quotient.

But *why*?

So a couple things.

One is that the quotient of some very large floating point numbers can easily be far too gigantic to fit in even a `long long unsigned int`. So some of it might very well need to be lopped off, anyway.

But at 3 bits? How's that even useful? That only gets you from 0 to 7!

The C99 Rationale document states:

The `remquo` functions are intended for implementing argument reductions which can exploit a few low-order bits of the quotient. Note that  $x$  may be so large in magnitude relative to  $y$  that an exact representation of the quotient is not practical.

So... implementing argument reductions... which can exploit a few low-order bits... Ooookay.

CPPReference has this to say<sup>11</sup> on the matter, which is spoken so well, I will quote wholesale:

This function is useful when implementing periodic functions with the period exactly representable as a floating-point value: when calculating  $\sin(\pi x)$  for a very large  $x$ , calling `sin` directly may result in a large error, but if the function argument is first reduced with `remquo`, the low-order bits of the quotient may be used to determine the sign and the octant of the result within the period, while the remainder may be used to calculate the value with high precision.

And there you have it. If you have another example that works for you... congratulations! :)

## Return Value

Returns the same as `remainder`: The IEC 60559 result of  $x$ -remainder- $y$ .

In addition, at least the lowest 3 bits of the quotient will be stored in `quo` with the same sign as  $x/y$ .

## Example

There's a great `cos()` example at CPPReference<sup>12</sup> that covers a genuine use case.

But instead of stealing it, I'll just post a simple example here and you can visit their site for a real one.

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      int quo;
7      double rem;
8
9      rem = remquo(12.75, 2.25, &quo);
10
11     printf("%d remainder %f\n", quo, rem); // 6 remainder -0.750000
12 }
```

## See Also

`remainder()`, `imaxdiv()`

## 13.54 `copysign()`, `copysignf()`, `copysignl()`

Copy the sign of one value into another.

<sup>11</sup><https://en.cppreference.com/w/c/numeric/math/remquo>

<sup>12</sup><https://en.cppreference.com/w/c/numeric/math/remquo>

## Synopsis

```
#include <math.h>

double copysign(double x, double y);

float copysignf(float x, float y);

long double copysignl(long double x, long double y);
```

## Description

These functions return a number that has the magnitude of `x` and the sign of `y`. You can use them to coerce the sign to that of another value.

Neither `x` nor `y` are modified, of course. The return value holds the result.

## Return Value

Returns a value with the magnitude of `x` and the sign of `y`.

## Example

```
1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      double x = 34.9;
7      double y = -999.9;
8      double z = 123.4;
9
10     printf("%f\n", copysign(x, y)); // -34.900000
11     printf("%f\n", copysign(x, z)); // 34.900000
12 }
```

## See Also

`signbit()`

---

## 13.55 `nan()`, `nanf()`, `nanl()`

Return `NAN`.

## Synopsis

```
#include <math.h>

double nan(const char *tagp);

float nanf(const char *tagp);

long double nanl(const char *tagp);
```

## Description

These functions return a quiet NaN<sup>13</sup>. It is produced as if calling `strtod()` with `"NaN"` (or a variant thereof) as an argument.

`tagp` points to a string which could be several things, including empty. The contents of the string determine which variant of NaN might get returned depending on the implementation.

Which *version* of NaN? Did you even know it was possible to get this far into the weeds with something that wasn't a number?

Case 1 in which you pass in an empty string, in which case these are the same:

```
nan("");

strtod("NaN()", NULL);
```

Case 2 in which the string contains only digits 0-9, letters a-z, letters A-Z, and/or underscore:

```
nan("goats");

strtod("NaN(goats)", NULL);
```

And Case 3, in which the string contains anything else and is ignored:

```
nan("!");

strtod("NaN", NULL);
```

As for what `strtod()` does with those values in parens, see the `[strtod()]` reference page. Spoiler: it's implementation-defined.

## Return Value

Returns the requested quiet NaN, or 0 if such things aren't supported by your system.

---

<sup>13</sup>A *quiet NaN* is one that doesn't raise any exceptions.



## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%f\n", nan(""));           // nan
7     printf("%f\n", nan("goats"));      // nan
8     printf("%f\n", nan("!"));          // nan
9 }
```

## See Also

strtod()

---

## 13.56 nextafter(), nextafterf(), nextafterl()

Get the next (or previous) representable floating point value.

## Synopsis

```
#include <math.h>

double nextafter(double x, double y);

float nextafterf(float x, float y);

long double nextafterl(long double x, long double y);
```

## Description

As you probably know, floating point numbers can't represent *every* possible real number. There are limits.

And, as such, there exists a “next” and “previous” number after or before any floating point number.

These functions return the next (or previous) representable number. That is, no floating point numbers exist between the given number and the next one.

The way it figures it out is it works from `x` in the direction of `y`, answering the question of “what is the next representable number from `x` as we head toward `y`.”

## Return Value

Returns the next representable floating point value from `x` in the direction of `y`.

If `x` equals `y`, returns `y`. And also `x`, I suppose.

**Example**

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("%. *f\n", DBL_DECIMAL_DIG, nextafter(0.5, 1.0));
7      printf("%. *f\n", DBL_DECIMAL_DIG, nextafter(0.349, 0.0));
8  }

```

Output on my system:

```

0.500000000000000011
0.34899999999999992

```

**See Also**

`nexttoward()`

**13.57 nexttoward(), nexttowardf(), nexttowardl()**

Get the next (or previous) representable floating point value.

**Synopsis**

```

include <math.h>

double nexttoward(double x, long double y);

float nexttowardf(float x, long double y);

long double nexttowardl(long double x, long double y);

```

**Description**

These functions are the same as `nextafter()` except the second parameter is always `long double`.

**Return Value**

Returns the same as `nextafter()` except if `x` equals `y`, returns `y` cast to the function's return type.

**Example**

```

1  #include <stdio.h>
2  #include <float.h>
3  #include <math.h>
4
5  int main(void)

```

```

6 {
7     printf("%. *f\n", DBL_DECIMAL_DIG, nexttoward(0.5, 1.0));
8     printf("%. *f\n", DBL_DECIMAL_DIG, nexttoward(0.349, 0.0));
9 }

```

Output on my system:

```

0.500000000000000011
0.34899999999999992

```

## See Also

`nextafter()`

## 13.58 `fdim()`, `fdimf()`, `fdiml()`

Return the positive difference between two numbers clamped at 0.

### Synopsis

```

#include <math.h>

double fdim(double x, double y);

float fdimf(float x, float y);

long double fdiml(long double x, long double y);

```

### Description

The positive difference between `x` and `y` is the difference... except if the difference is less than 0, it's clamped to 0.

These functions might throw a range error.

### Return Value

Returns the difference of `x-y` if the difference is greater than 0. Otherwise it returns 0.

### Example

```

1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("%f\n", fdim(10.0, 3.0));    // 7.000000
7      printf("%f\n", fdim(3.0, 10.0));    // 0.000000, clamped
8  }

```

---

### 13.59 `fmax()`, `fmaxf()`, `fmaxl()`, `fmin()`, `fminf()`, `fminl()`

Return the maximum or minimum of two numbers.

#### Synopsis

```
#include <math.h>

double fmax(double x, double y);

float fmaxf(float x, float y);

long double fmaxl(long double x, long double y);

double fmin(double x, double y);

float fminf(float x, float y);

long double fminl(long double x, long double y);
```

#### Description

Straightforwardly, these functions return the minimum or maximum of two given numbers.

If one of the numbers is NaN, the functions return the non-NaN number. If both arguments are NaN, the functions return NaN.

#### Return Value

Returns the minimum or maximum values, with NaN handled as mentioned above.

#### Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%f\n", fmin(10.0, 3.0)); // 3.000000
7     printf("%f\n", fmax(3.0, 10.0)); // 10.000000
8 }
```

---

### 13.60 `fma()`, `fmaf()`, `fmal()`

Floating (AKA “Fast”) multiply and add.

## Synopsis

```
#include <math.h>

double fma(double x, double y, double z);

float fmaf(float x, float y, float z);

long double fmal(long double x, long double y, long double z);
```

## Description

This performs the operation  $(x \times y) + z$ , but does so in a nifty way. It does the computation as if it had infinite precision, and then rounds the final result to the final data type according to the current rounding mode.

Contrast to if you'd do the math yourself, where it would have rounded each step of the way, potentially.

Also some architectures have a CPU instruction to do exactly this calculation, so it can do it super quick. (If it doesn't, it's considerably slower.)

You can tell if your CPU supports the fast version by checking that the macro `FP_FAST_FMA` is set to `1`. (The `float` and `long` variants of `fma()` can be tested with `FP_FAST_FMAF` and `FP_FAST_FMAL`, respectively.)

These functions might cause a range error to occur.

## Return Value

Returns  $(x * y) + z$ .

## Example

```
1 printf("%f\n", fma(1.0, 2.0, 3.0)); // 5.000000
```

## 13.61 isgreater(), isgreaterequal(), isless(), islessequal()

Floating point comparison macros.

## Synopsis

```
#include <math.h>

int isgreater(any_floating_type x, any_floating_type y);

int isgreaterequal(any_floating_type x, any_floating_type y);

int isless(any_floating_type x, any_floating_type y);

int islessequal(any_floating_type x, any_floating_type y);
```

## Description

These macros compare floating point numbers. Being macros, we can pass in any floating point type.

You might think you can already do that with just regular comparison operators—and you’d be right!

One exception: the comparison operators raise the “invalid” floating exception if one or more of the operands is NaN. These macros do not.

Note that you must only pass floating point types into these functions. Passing an integer or any other type is undefined behavior.

## Return Value

`isgreater()` returns the result of `x > y`.

`isgreaterequal()` returns the result of `x >= y`.

`isless()` returns the result of `x < y`.

`islessequal()` returns the result of `x <= y`.

## Example

```
1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("%d\n", isgreater(10.0, 3.0));      // 1
7      printf("%d\n", isgreaterequal(10.0, 10.0)); // 1
8      printf("%d\n", isless(10.0, 3.0));         // 0
9      printf("%d\n", islessequal(10.0, 3.0));    // 0
10 }
```

## See Also

`islessgreater()`, `isunordered()`

---

## 13.62 `islessgreater()`

Test if a floating point number is less than or greater than another.

## Synopsis

```
#include <math.h>

int islessgreater(any_floating_type x, any_floating_type y);
```

## Description

This macro is similar to `isgreater()` and all those, except it made the section name too long if I included it up there. So it gets its own spot.

This returns true if  $x < y$  or  $x > y$ .

Even though it's a macro, we can rest assured that `x` and `y` are only evaluated once.

And even if `x` or `y` are NaN, this will not throw an “invalid” exception, unlike the normal comparison operators.

If you pass in a non-floating type, the behavior is undefined.

## Return Value

Returns `(x < y) || (x > y)`.

## Example

```
1 #include <stdio.h>
2 #include <math.h>
3
4 int main(void)
5 {
6     printf("%d\n", islessgreater(10.0, 3.0)); // 1
7     printf("%d\n", islessgreater(10.0, 30.0)); // 1
8     printf("%d\n", islessgreater(10.0, 10.0)); // 0
9 }
```

## See Also

`isgreater()`, `isgreaterequal()`, `isless()`, `islessequal()`, `isunordered()`

## 13.63 `isunordered()`

Macro returns true if either floating point argument is NaN.

## Synopsis

```
#include <math.h>

int isunordered(any_floating_type x, any_floating_type y);
```

## Description

The spec writes:

| The `isunordered` macro determines whether its arguments are unordered. |

See? Told you C was easy!

It does also elaborate that the arguments are unordered if one or both of them are NaN.

## Return Value

This macro returns true if one or both of the arguments are NaN.

## Example

```
1  #include <stdio.h>
2  #include <math.h>
3
4  int main(void)
5  {
6      printf("%d\n", isunordered(1.0, 2.0));    // 0
7      printf("%d\n", isunordered(1.0, sqrt(-1))); // 1
8      printf("%d\n", isunordered(NAN, 30.0));   // 1
9      printf("%d\n", isunordered(NAN, NAN));    // 1
10 }
```

## See Also

`isgreater()`, `isgreaterequal()`, `isless()`, `islessequal()`, `islessgreater()`



## Chapter 14

# <setjmp.h> Non-local Goto

These functions enable you to rewind the call stack to an earlier point, with a bunch of gotchas. It is rarely used.

| Function               | Description                              |
|------------------------|------------------------------------------|
| <code>longjmp()</code> | Return to the previously-placed bookmark |
| <code>setjmp()</code>  | Bookmark this place to return to later   |

There's also a new opaque type, `jmp_buf`, that holds all the information needed to pull off this magic trick.

If you want your automatic local variables to be correct after a call to `longjmp()`, declare them as `volatile` where you called `setjmp()`.

### 14.1 `setjmp()`

Save this location as one to return to later

#### Synopsis

```
#include <setjmp.h>

int setjmp(jmp_buf env);
```

#### Description

This is how you save your position so you can `longjmp()` back it, later. Think of it as setting up a warp destination for later use.

Basically, you call this, giving it an `env` it can fill in with all the information it needs to come back here later. This `env` is one you'll pass to `longjmp()` later when you want to teleport back here.

And the really funky part is this can return two different ways:

1. It can return 0 from the call where you set up the jump destination.
2. It can return non-zero when you actually warp back here as the result of a call to `longjmp()`.

What you can do is check the return value to see which case has occurred.

You're only allowed to call `setjmp()` in a limited number of circumstances.

1. As a standalone expression:

```
setjmp(env);
```

You can also cast it to `(void)` if you really wanted to do such a thing.

2. As the complete controlling expression in an `if` or `switch`.

```
if (setjmp(env)) { ... }

switch (setjmp(env)) { ... }
```

But not this as it's not the complete controlling expression in this case:

```
if (x == 2 && setjmp()) { ... } // Undefined behavior
```

3. The same as (2), above, except with a comparison to an integer constant:

```
if (setjmp(env) == 0) { ... }

if (setjmp(env) > 2) { ... }
```

4. As the operand to the not (!) operator:

```
if (!setjmp(env)) { ... }
```

Anything else is (you guessed it) undefined behavior!

This can be a macro or a function, but you'll treat it the same way in any case.

## Return Value

This one is funky. It returns one of two things:

Returns 0 if this was the call to `setjmp()` to set it up.

Returns non-zero if being here was the result of a call to `longjmp()`. (Namely, it returns the value passed into the `longjmp()` function.)

## Example

Here's a function that calls `setjmp()` to set things up (where it returns 0), then calls a couple levels deep into functions, and finally short-circuits the return path by `longjmp()`ing back to the place where `setjmp()` was called, earlier. This time, it passes 3490 as a value, which `setjmp()` returns.

```
1 #include <stdio.h>
2 #include <setjmp.h>
3
4 jmp_buf env;
5
6 void depth2(void)
```

```

7  {
8      printf("Entering depth 2\n");
9      longjmp(env, 3490); // Jump back to setjmp()!!
10     printf("Leaving depth 2\n"); // This won't happen
11 }
12
13 void depth1(void)
14 {
15     printf("Entering depth 1\n");
16     depth2();
17     printf("Leaving depth 1\n"); // This won't happen
18 }
19
20 int main(void)
21 {
22     switch (setjmp(env)) {
23         case 0:
24             printf("Calling into functions, setjmp() returned 0\n");
25             depth1();
26             printf("Returned from functions\n"); // This won't happen
27             break;
28
29         case 3490:
30             printf("Bailed back to main, setjmp() returned 3490\n");
31             break;
32     }
33 }

```

When run, this outputs:

```

Calling into functions, setjmp() returned 0
Entering depth 1
Entering depth 2
Bailed back to main, setjmp() returned 3490

```

Notice that the second `printf()` in `case 0` didn't run; it got jumped over by `longjmp()`!

## See Also

`longjmp()`

## 14.2 `longjmp()`

Return to the previous `setjmp()` location

## Synopsis

```
#include <setjmp.h>

_Noreturn void longjmp(jmp_buf env, int val);
```

## Description

This returns to a previous call to `setjmp()` back in the call history. `setjmp()` will return the `val` passed into `longjmp()`.

The `env` passed to `setjmp()` should be the same one you pass into `longjmp()`.

There are a bunch of potential issues with doing this, so you'll want to be careful that you avoid undefined behavior by not doing the following:

1. Don't call `longjmp()` if the corresponding `setjmp()` was in a different thread.
2. Don't call `longjmp()` if you didn't call `setjmp()` first.
3. Don't call `longjmp()` if the function that called `setjmp()` has completed.
4. Don't call `longjmp()` if the call to `setjmp()` had a variable length array (VLA) in scope and the scope has ended.
5. Don't call `longjmp()` if there are any VLAs in any active scopes between the `setjmp()` and the `longjmp()`. A good rule of thumb here is to not mix VLAs and `longjmp()`.

Though `longjmp()` attempts to restore the machine to the state at the `setjmp()`, including local variables, there are some things that aren't brought back to life:

- Non-volatile local variables that might have changed
- Floating point status flags
- Open files
- Any other component of the abstract machine

## Return Value

This one is also funky in that it is one of the few functions in C that never returns!

## Example

Here's a function that calls `setjmp()` to set things up (where it returns 0), then calls a couple levels deep into functions, and finally short-circuits the return path by `longjmp()`ing back to the place where `setjmp()` was called, earlier. This time, it passes 3490 as a value, which `setjmp()` returns.

```
1  #include <stdio.h>
2  #include <setjmp.h>
3
4  jmp_buf env;
5
6  void depth2(void)
7  {
8      printf("Entering depth 2\n");
9      longjmp(env, 3490); // Jump back to setjmp()!!
10     printf("Leaving depth 2\n"); // This won't happen
11 }
```

```
12
13 void depth1(void)
14 {
15     printf("Entering depth 1\n");
16     depth2();
17     printf("Leaving depth 1\n"); // This won't happen
18 }
19
20 int main(void)
21 {
22     switch (setjmp(env)) {
23         case 0:
24             printf("Calling into functions, setjmp() returned 0\n");
25             depth1();
26             printf("Returned from functions\n"); // This won't happen
27             break;
28
29         case 3490:
30             printf("Bailed back to main, setjmp() returned 3490\n");
31             break;
32     }
33 }
```

When run, this outputs:

```
Calling into functions, setjmp() returned 0
Entering depth 1
Entering depth 2
Bailed back to main, setjmp() returned 3490
```

Notice that the second `printf()` in `case 0` didn't run; it got jumped over by `longjmp()`!

## See Also

`setjmp()`

## Chapter 15

# <signal.h> signal handling

| Function              | Description                             |
|-----------------------|-----------------------------------------|
| <code>signal()</code> | Set a signal handler for a given signal |
| <code>raise()</code>  | Cause a signal to be raised             |

Handle signals in a portable way, kind of!

These signals get raised for a variety of reasons such as CTRL-C being hit, requests to terminate for external programs, memory access violations, and so on.

Your OS likely defines a plethora of other signals, as well.

This system is pretty limited, as seen below. If you're on Unix, it's almost certain your OS has far superior signal handling capabilities than the C standard library. Check out `sigaction`<sup>1</sup>.

### 15.1 `signal()`

Set a signal handler for a given signal

#### Synopsis

```
#include <signal.h>

void (*signal(int sig, void (*func)(int)))(int);
```

#### Description

How's *that* for a function declaration?

Let's ignore it for a moment and just talk about what this function does.

When a signal is raised, *something* is going to happen. This function lets you decide to do one of these things when the signal is raised:

---

<sup>1</sup><https://man.archlinux.org/man/sigaction.2.en>

- Ignore the signal
- Perform the default action
- Have a specific function called

The `signal()` function takes two arguments. The first, `sig`, is the name of the signal to handle.

| Signal               | Description                                                   |
|----------------------|---------------------------------------------------------------|
| <code>SIGABRT</code> | Raised when <code>abort()</code> is called                    |
| <code>SIGFPE</code>  | Floating-point arithmetic exception                           |
| <code>SIGILL</code>  | CPU tried to execute an illegal instruction                   |
| <code>SIGINT</code>  | Interrupt signal, as if <code>CTRL-C</code> were pressed      |
| <code>SIGSEGV</code> | Segmentation Violation: attempted to access restricted memory |
| <code>SIGTERM</code> | Termination request <sup>2</sup>                              |

So that's the first bit when you call `signal()`—tell it the signal in question:

```
signal(SIGINT, ...
```

But what's that `func` parameter?

For spoilers, it's a pointer to a function that takes an `int` argument and returns `void`. We can use this to call an arbitrary function when the signal occurs.

Before we do that, though, let's look at the easy ones: telling the system to ignore the signal or perform the default action (which it does by default if you never call `signal()`).

You can set `func` to one of two special values to make this happen:

| func                 | Description                               |
|----------------------|-------------------------------------------|
| <code>SIG_DFL</code> | Perform the default action on this signal |
| <code>SIG_IGN</code> | Ignore this signal                        |

For example:

```
signal(SIGTERM, SIG_DFL); // Default action on SIGTERM
signal(SIGINT, SIG_IGN);  // Ignore SIGINT
```

But what if you want to have your own handler do something instead of the default or ignoring it? You can pass in your own function to be called. That's what the crazy function signature is partially about. It's saying that the argument can be a pointer to a function that takes an `int` argument and returns `void`.

So if you wanted to call your handler, you could have code like this:

```
int handler(int sig)
{
    // Handle the signal
}

int main(void)
{
    signal(SIGINT, handler);
```

<sup>2</sup>As if might be sent from Unix's `kill` command.]

What can you do in the signal handler? Not much.

If the signal is due to `abort()` or `raise()`, the handler can't call `raise()`.

If the signal is **not** due to `abort()` or `raise()`, you're only allowed to call these functions from the standard library (though the spec doesn't prohibit calling other non-library functions):

- `abort()`
- `_Exit()`
- `quick_exit()`
- Functions in `<stdatomic.h>` when the atomic arguments are lock-free
- `signal()` with a first argument equivalent to the argument that was passed into the handler

In addition, if the signal was **not** due to `abort()` or `raise()`, the handler can't access any object with static or thread-storage duration unless it's lock-free.

An exception is that you can assign to (but not read from!) a variable of type `volatile sig_atomic_t`.

It's up to the implementation, but the signal handler might be reset to `SIG_DFL` just before the handler is called.

It's undefined behavior to call `signal()` in a multithreaded program.

It's undefined behavior to return from the handler for `SIGFPE`, `SIGILL`, `SIGSEGV`, or any implementation-defined value. You must exit.

The implementation might or might not prevent other signals from arising while in the signal handler.

## Return Value

On success, `signal()` returns a pointer to the previous signal handler set by a call to `signal()` for that particular signal number. If you haven't called it set, returns `SIG_DFL`.

On failure, `SIG_ERR` is returned and `errno` is set to a positive value.

## Example

Here's a program that causes `SIGINT` to be ignored. Commonly you trigger this signal by hitting `CTRL-C`.

```

1  #include <stdio.h>
2  #include <signal.h>
3
4  int main(void)
5  {
6      signal(SIGINT, SIG_IGN);
7
8      printf("You can't hit CTRL-C to exit this program. Try it!\n\n");
9      printf("Press return to exit, instead.");
10     fflush(stdout);
11     getchar();
12 }
```

Output:

```

You can't hit CTRL-C to exit this program. Try it!

Press return to exit, instead.^^^C^^^C^^^C^^^C^^^C^^^C
```

This program sets the signal handler, then raises the signal. The signal handler fires.



```
1 #include <stdio.h>
2 #include <signal.h>
3
4 void handler(int sig)
5 {
6     // Undefined behavior to call printf() if this handler was not
7     // as the result of a raise(), i.e. if you hit CTRL-C.
8
9     printf("Got signal %d!\n", sig);
10
11     // Common to reset the handler just in case the implementation set
12     // it to SIG_DFL when the signal occurred.
13
14     signal(sig, handler);
15 }
16
17 int main(void)
18 {
19     signal(SIGINT, handler);
20
21     raise(SIGINT);
22     raise(SIGINT);
23     raise(SIGINT);
24 }
```

Output:

```
Got signal 2!
Got signal 2!
Got signal 2!
```

This example catches `SIGINT` but then sets a flag to 1. Then the main loop sees the flag and exits.

```
1 #include <stdio.h>
2 #include <signal.h>
3
4 volatile sig_atomic_t x;
5
6 void handler(int sig)
7 {
8     x = 1;
9 }
10
11 int main(void)
12 {
13     signal(SIGINT, handler);
14
15     printf("Hit CTRL-C to exit\n");
16     while (x != 1);
17 }
```

## See Also

`raise()`, `abort()`

---

## 15.2 `raise()`

Cause a signal to be raised

### Synopsis

```
#include <signal.h>

int raise(int sig);
```

### Description

Causes the signal handler for the signal `sig` to be called. If the handler is `SIG_DFL` or `SIG_IGN`, then the default action or no action happens.

`raise()` returns after the signal handler has finished running.

Interestingly, if you cause a signal to happen with `raise()`, you can call library functions from within the signal handler without causing undefined behavior. I'm not sure how this fact is practically useful, though.

### Return Value

Returns `0` on success. Nonzero otherwise.

### Example

This program sets the signal handler, then raises the signal. The signal handler fires.

```
1  #include <stdio.h>
2  #include <signal.h>
3
4  void handler(int sig)
5  {
6      // Undefined behavior to call printf() if this handler was not
7      // as the result of a raise(), i.e. if you hit CTRL-C.
8
9      printf("Got signal %d!\n", sig);
10
11     // Common to reset the handler just in case the implementation set
12     // it to SIG_DFL when the signal occurred.
13
14     signal(sig, handler);
15 }
16
17 int main(void)
18 {
19     signal(SIGINT, handler);
```

```
20
21     raise(SIGINT);
22     raise(SIGINT);
23     raise(SIGINT);
24 }
```

Output:

```
Got signal 2!
Got signal 2!
Got signal 2!
```

## See Also

`signal()`

## Chapter 16

# <stdalign.h> Macros for Alignment

If you're coding up something low-level like a memory allocator that interfaces with your OS, you might need this header file. But most C devs go their careers without using it.

*Alignment*<sup>1</sup> is all about multiples of addresses on which objects can be stored. Can you store this at any address? Or must it be a starting address that's divisible by 2? Or 8? Or 16?

| Name                   | Description                                         |
|------------------------|-----------------------------------------------------|
| <code>alignas()</code> | Specify alignment, expands to <code>_Alignas</code> |
| <code>alignof()</code> | Get alignment, expands to <code>_Alignof</code>     |

These two additional macros are defined to be 1:

```
__alignas_is_defined
__alignof_is_defined
```

Quick note: alignments greater than that of `max_align_t` are known as *overalignments* and are implementation-defined.

### 16.1 `alignas()` `_Alignas()`

Force a variable to have a certain alignment

#### Synopsis

```
#include <stdalign.h>

alignas(type-name)
alignas(constant-expression)
```

```
_Alignas(type-name)
_Alignas(constant-expression)
```

<sup>1</sup>[https://en.wikipedia.org/wiki/Data\\_structure\\_alignment](https://en.wikipedia.org/wiki/Data_structure_alignment)

## Description

Use this *alignment specifier* to force the alignment of particular variables. For instance, we can declare `c` to be `char`, but aligned as if it were an `int`:

```
char alignas(int) c;
```

You can put a constant integer expression in there, as well. The compiler will probably impose limits on what these values can be. Small powers of 2 (1, 2, 4, 8, and 16) are generally safe bets.

```
char alignas(8) c;    // align on 8-byte boundaries
```

For convenience, you can also specify `0` if you want the default alignment (as if you hadn't said `alignas()` at all):

```
char alignas(0) c;    // use default alignment for this type
```

## Example

```
1 #include <stdalign.h>
2 #include <stdio.h>      // for printf()
3 #include <stddef.h>     // for max_align_t
4
5 int main(void)
6 {
7     int i, j;
8     char alignas(max_align_t) a, b;
9     char alignas(int) c, d;
10    char e, f;
11
12    printf("i: %p\n", (void *)&i);
13    printf("j: %p\n", (void *)&j);
14    printf("a: %p\n", (void *)&a);
15    printf("b: %p\n", (void *)&b);
16    printf("c: %p\n", (void *)&c);
17    printf("d: %p\n", (void *)&d);
18    printf("e: %p\n", (void *)&e);
19    printf("f: %p\n", (void *)&f);
20 }
```

Output on my system follows. Notice the difference between the pairs of values.

- `i` and `j`, both `ints`, are aligned on 4-byte boundaries.
- `a` and `b` have been forced to the boundary of the type `max_align_t`, which is every 16 bytes on my system.
- `c` and `d` have been forced to the same alignment as `int`, which is 4 bytes, just like with `i` and `j`.
- `e` and `f` do not have an alignment specified, so they were stored with their default alignment of 1 byte.

```
i: 0x7ffee7dfb4cc    <-- difference of 4 bytes
j: 0x7ffee7dfb4c8

a: 0x7ffee7dfb4c0    <-- difference of 16 bytes
```

```

b: 0x7ffee7dfb4b0

c: 0x7ffee7dfb4ac    <-- difference of 4 bytes
d: 0x7ffee7dfb4a8

e: 0x7ffee7dfb4a7    <-- difference of 1 byte
f: 0x7ffee7dfb4a6

```

## See Also

`alignof`, `max_align_t`, `memalignment()`

---

## 16.2 `alignof()` `_Alignof()`

Get the alignment of a type

### Synopsis

```
#include <stdalign.h>
```

```
alignof(type-name)
```

```
_Alignof(type-name)
```

### Description

This evaluates to a value of type `size_t` that gives the alignment of a particular type on your system.

### Return Value

Returns the alignment value, i.e. the address of the beginning of the given type of object must begin on an address boundary divisible by this number.

### Example

Print out the alignments of a variety of different types.

```

1  #include <stdalign.h>
2  #include <stdio.h>      // for printf()
3  #include <stddef.h>     // for max_align_t
4
5  struct t {
6      int a;
7      char b;
8      float c;
9  };
10
11 int main(void)

```

```
12 {  
13     printf("char      : %zu\n", alignof(char));  
14     printf("short     : %zu\n", alignof(short));  
15     printf("int       : %zu\n", alignof(int));  
16     printf("long      : %zu\n", alignof(long));  
17     printf("long long  : %zu\n", alignof(long long));  
18     printf("double     : %zu\n", alignof(double));  
19     printf("long double: %zu\n", alignof(long double));  
20     printf("struct t   : %zu\n", alignof(struct t));  
21     printf("max_align_t: %zu\n", alignof(max_align_t));  
22 }
```

Output on my system:

```
char      : 1  
short     : 2  
int       : 4  
long      : 8  
long long : 8  
double    : 8  
long double: 16  
struct t   : 16  
max_align_t: 16
```

## See Also

`alignas`, `max_align_t`, `memalignment()`

## Chapter 17

# <stdarg.h> Variable Arguments

| Macro                   | Description                                                             |
|-------------------------|-------------------------------------------------------------------------|
| <code>va_arg()</code>   | Get the next variable argument                                          |
| <code>va_copy()</code>  | Copy a <code>va_list</code> and the work done so far                    |
| <code>va_end()</code>   | Signify we're done processing variable arguments                        |
| <code>va_start()</code> | Initialize a <code>va_list</code> to start variable argument processing |

This header file is what allows you to write functions that take a variable number of arguments.

In addition to the macros, you get a new type that helps C keep track of where it is in the variable-number-of-arguments-processing: `va_list`. This type is opaque, and you'll be passing it around to the various macros to help get at the arguments.

Note that every variadic function requires at least one non-variable parameter. You need this to kick off processing with `va_start()`.

### 17.1 `va_arg()`

Get the next variable argument

#### Synopsis

```
#include <stdarg.h>

type va_arg(va_list ap, type);
```

#### Description

If you have a variable argument list you've initialized with `va_start()`, pass it to this one along with the type of argument you're trying to get, e.g.

```
int x = va_arg(args, int);
float y = va_arg(args, float);
```



## Return Value

Evaluates to the value and type of the next variable argument.

## Example

Here's a demo that adds together an arbitrary number of integers. The first argument is the number of integers to add together. We'll make use of that to figure out how many times we have to call `va_arg()`.

```

1  #include <stdio.h>
2  #include <stdarg.h>
3
4  int add(int count, ...)
5  {
6      int total = 0;
7      va_list va;
8
9      va_start(va, count); // Start with arguments after "count"
10
11     for (int i = 0; i < count; i++) {
12         int n = va_arg(va, int); // Get the next int
13
14         total += n;
15     }
16
17     va_end(va); // All done
18
19     return total;
20 }
21
22 int main(void)
23 {
24     printf("%d\n", add(4, 6, 2, -4, 17)); // 6 + 2 - 4 + 17 = 21
25     printf("%d\n", add(2, 22, 44));      // 22 + 44 = 66
26 }
```

## See Also

`va_start()`, `va_end()`

---

## 17.2 `va_copy()`

Copy a `va_list` and the work done so far

### Synopsis

```

#include <stdarg.h>

void va_copy(va_list dest, va_list src);
```

## Description

The main intended use of this is to save your state partway through processing variable arguments so you can scan ahead and then rewind back to the save point.

You pass in a `src va_list` and it copies it to `dest`.

If you've already called this once for a particular `dest`, you can't call it (or `va_start()`) again with the same `dest` unless you call `va_end()` on that `dest` first.

```
va_copy(dest, src);
va_copy(dest, src2); // BAD!

va_copy(dest, src);
va_start(dest, var); // BAD!

va_copy(dest, src);
va_end(dest);
va_copy(dest, src2); // OK!

va_copy(dest, src);
va_end(dest);
va_start(dest, var); // OK!
```

## Return Value

Returns nothing.

## Example

Here's an example where we're adding together all the variable arguments, but then we want to go back and add on all the numbers past the first two, for example if the arguments are:

```
10 20 30 40
```

First we add them all for `100`, and then we add on everything from the third number on, so add on `30+40` for a total of `170`.

We'll do this by saving our place in the variable argument processing with `va_copy` and then using that later to reprocess the trailing arguments.

(And yes, I know there's a mathematical way to do this without all the rewinding, but I'm having an heck of a time coming up with a good example!)

```
1  #include <stdio.h>
2  #include <stdarg.h>
3
4  // Add all the numbers together, but then add on all the numbers
5  // past the second one again.
6  int contrived_adder(int count, ...)
7  {
8      if (count < 3) return 0; // OK, I'm being lazy. You got me.
9
10     int total = 0;
```

```

12     va_list args, mid_args;
13
14     va_start(args, count);
15
16     for (int i = 0; i < count; i++) {
17
18         // If we're at the second number, save our place in
19         // mid_args:
20
21         if (i == 2)
22             va_copy(mid_args, args);
23
24         total += va_arg(args, int);
25     }
26
27     va_end(args); // Done with this
28
29     // But now let's start with mid_args and add all those on:
30     for (int i = 0; i < count - 2; i++)
31         total += va_arg(mid_args, int);
32
33     va_end(mid_args); // Done with this, too
34
35     return total;
36 }
37
38 int main(void)
39 {
40     // 10+20+30 + 30 == 90
41     printf("%d\n", contrived_adder(3, 10, 20, 30));
42
43     // 10+20+30+40+50 + 30+40+50 == 270
44     printf("%d\n", contrived_adder(5, 10, 20, 30, 40, 50));
45 }

```

## See Also

`va_start()`, `va_arg()`, `va_end()`

---

## 17.3 `va_end()`

Signify we're done processing variable arguments

### Synopsis

```

#include <stdarg.h>

void va_end(va_list ap);

```

## Description

After you've `va_start()`ed or `va_copy()`d a new `va_list`, you **must** call `va_end()` with it before it goes out of scope.

You also have to do this if you're going to call `va_start()` or `va_copy()` *again* on a variable you've already done that to.

There's the rules if you want to avoid undefined behavior.

But just think of it as cleanup. You called `va_start()`, so you'll call `va_end()` when you're done.

## Return Value

Returns nothing.

## Example

Here's a demo that adds together an arbitrary number of integers. The first argument is the number of integers to add together. We'll make use of that to figure out how many times we have to call `va_arg()`.

```
1  #include <stdio.h>
2  #include <stdarg.h>
3
4  int add(int count, ...)
5  {
6      int total = 0;
7      va_list va;
8
9      va_start(va, count); // Start with arguments after "count"
10
11     for (int i = 0; i < count; i++) {
12         int n = va_arg(va, int); // Get the next int
13
14         total += n;
15     }
16
17     va_end(va); // All done
18
19     return total;
20 }
21
22 int main(void)
23 {
24     printf("%d\n", add(4, 6, 2, -4, 17)); // 6 + 2 - 4 + 17 = 21
25     printf("%d\n", add(2, 22, 44));      // 22 + 44 = 66
26 }
```

## See Also

`va_start()`, `va_copy()`

---

## 17.4 va\_start()

Initialize a `va_list` to start variable argument processing

### Synopsis

```
#include <stdarg.h>

void va_start(va_list ap, parmN);
```

### Description

You've declared a variable of type `va_list` to keep track of the variable argument processing... now how to initialize it so you can start calling `va_arg()` to get those arguments?

`va_start()` to the rescue!

What you do is pass in your `va_list`, here shown as parameter `ap`. Just pass the list, not a pointer to it.

Then for the second argument to `va_start()`, you give the name of the parameter that you want to start processing arguments *after*. This must be the parameter right before the `...` in the argument list.

If you've already called `va_start()` on a particular `va_list` and you want to call `va_start()` on it again, you **must** call `va_end()` first!

### Return Value

Returns nothing!

### Example

Here's a demo that adds together an arbitrary number of integers. The first argument is the number of integers to add together. We'll make use of that to figure out how many times we have to call `va_arg()`.

```
1  #include <stdio.h>
2  #include <stdarg.h>
3
4  int add(int count, ...)
5  {
6      int total = 0;
7      va_list va;
8
9      va_start(va, count); // Start with arguments after "count"
10
11     for (int i = 0; i < count; i++) {
12         int n = va_arg(va, int); // Get the next int
13
14         total += n;
15     }
16
17     va_end(va); // All done
18
19     return total;
20 }
```

```
21
22 int main(void)
23 {
24     printf("%d\n", add(4, 6, 2, -4, 17)); // 6 + 2 - 4 + 17 = 21
25     printf("%d\n", add(2, 22, 44));      // 22 + 44 = 66
26 }
```

## See Also

`va_arg()`, `va_end()`

## Chapter 18

# <stdatomic.h> Atomic-Related Functions

| Function                                               | Description                                          |
|--------------------------------------------------------|------------------------------------------------------|
| <code>atomic_compare_exchange_strong_explicit()</code> | Atomic compare and exchange, strong, explicit        |
| <code>atomic_compare_exchange_strong()</code>          | Atomic compare and exchange, strong                  |
| <code>atomic_compare_exchange_weak_explicit()</code>   | Atomic compare and exchange, weak, explicit          |
| <code>atomic_compare_exchange_weak()</code>            | Atomic compare and exchange, weak                    |
| <code>atomic_exchange_explicit()</code>                | Replace a value in an atomic object, explicit        |
| <code>atomic_exchange()</code>                         | Replace a value in an atomic object                  |
| <code>atomic_fetch_add_explicit()</code>               | Atomically add to an atomic integer, explicit        |
| <code>atomic_fetch_add()</code>                        | Atomically add to an atomic integer                  |
| <code>atomic_fetch_and_explicit()</code>               | Atomically bitwise-AND an atomic integer, explicit   |
| <code>atomic_fetch_and()</code>                        | Atomically bitwise-AND an atomic integer             |
| <code>atomic_fetch_or_explicit()</code>                | Atomically bitwise-OR an atomic integer, explicit    |
| <code>atomic_fetch_or()</code>                         | Atomically bitwise-OR an atomic integer              |
| <code>atomic_fetch_sub_explicit()</code>               | Atomically subtract from an atomic integer, explicit |
| <code>atomic_fetch_sub()</code>                        | Atomically subtract from an atomic integer           |
| <code>atomic_fetch_xor_explicit()</code>               | Atomically bitwise-XOR an atomic integer, explicit   |
| <code>atomic_fetch_xor()</code>                        | Atomically bitwise-XOR an atomic integer             |
| <code>atomic_flag_clear_explicit()</code>              | Clear an atomic flag, explicit                       |
| <code>atomic_flag_clear()</code>                       | Clear an atomic flag                                 |
| <code>atomic_flag_test_and_set_explicit()</code>       | Test and set an atomic flag, explicit                |
| <code>atomic_flag_test_and_set()</code>                | Test and set an atomic flag                          |
| <code>atomic_init()</code>                             | Initialize an atomic variable                        |
| <code>atomic_is_lock_free()</code>                     | Determine if an atomic type is lock free             |
| <code>atomic_load_explicit()</code>                    | Return a value from an atomic variable, explicit     |
| <code>atomic_load()</code>                             | Return a value from an atomic variable               |
| <code>atomic_signal_fence()</code>                     | Fence for intra-thread signal handlers               |
| <code>atomic_store_explicit()</code>                   | Store a value in an atomic variable, explicit        |
| <code>atomic_store()</code>                            | Store a value in an atomic variable                  |
| <code>atomic_thread_fence()</code>                     | Set up a fence                                       |
| <code>ATOMIC_VAR_INIT()</code>                         | Create an initializer for an atomic variable         |
| <code>kill_dependency()</code>                         | End a dependency chain                               |

You might need to add `-latomic` to your compilation command line on Unix-like operating systems.

## 18.1 Atomic Types

A bunch of types are predefined by this header:

| Atomic type                        | Longhand equivalent                     |
|------------------------------------|-----------------------------------------|
| <code>atomic_bool</code>           | <code>_Atomic_Bool</code>               |
| <code>atomic_char</code>           | <code>_Atomic char</code>               |
| <code>atomic_schar</code>          | <code>_Atomic signed char</code>        |
| <code>atomic_uchar</code>          | <code>_Atomic unsigned char</code>      |
| <code>atomic_short</code>          | <code>_Atomic short</code>              |
| <code>atomic_ushort</code>         | <code>_Atomic unsigned short</code>     |
| <code>atomic_int</code>            | <code>_Atomic int</code>                |
| <code>atomic_uint</code>           | <code>_Atomic unsigned int</code>       |
| <code>atomic_long</code>           | <code>_Atomic long</code>               |
| <code>atomic_ulong</code>          | <code>_Atomic unsigned long</code>      |
| <code>atomic_llong</code>          | <code>_Atomic long long</code>          |
| <code>atomic_ullong</code>         | <code>_Atomic unsigned long long</code> |
| <code>atomic_char16_t</code>       | <code>_Atomic char16_t</code>           |
| <code>atomic_char32_t</code>       | <code>_Atomic char32_t</code>           |
| <code>atomic_wchar_t</code>        | <code>_Atomic wchar_t</code>            |
| <code>atomic_int_least8_t</code>   | <code>_Atomic int_least8_t</code>       |
| <code>atomic_uint_least8_t</code>  | <code>_Atomic uint_least8_t</code>      |
| <code>atomic_int_least16_t</code>  | <code>_Atomic int_least16_t</code>      |
| <code>atomic_uint_least16_t</code> | <code>_Atomic uint_least16_t</code>     |
| <code>atomic_int_least32_t</code>  | <code>_Atomic int_least32_t</code>      |
| <code>atomic_uint_least32_t</code> | <code>_Atomic uint_least32_t</code>     |
| <code>atomic_int_least64_t</code>  | <code>_Atomic int_least64_t</code>      |
| <code>atomic_uint_least64_t</code> | <code>_Atomic uint_least64_t</code>     |
| <code>atomic_int_fast8_t</code>    | <code>_Atomic int_fast8_t</code>        |
| <code>atomic_uint_fast8_t</code>   | <code>_Atomic uint_fast8_t</code>       |
| <code>atomic_int_fast16_t</code>   | <code>_Atomic int_fast16_t</code>       |
| <code>atomic_uint_fast16_t</code>  | <code>_Atomic uint_fast16_t</code>      |
| <code>atomic_int_fast32_t</code>   | <code>_Atomic int_fast32_t</code>       |
| <code>atomic_uint_fast32_t</code>  | <code>_Atomic uint_fast32_t</code>      |
| <code>atomic_int_fast64_t</code>   | <code>_Atomic int_fast64_t</code>       |
| <code>atomic_uint_fast64_t</code>  | <code>_Atomic uint_fast64_t</code>      |
| <code>atomic_intptr_t</code>       | <code>_Atomic intptr_t</code>           |
| <code>atomic_uintptr_t</code>      | <code>_Atomic uintptr_t</code>          |
| <code>atomic_size_t</code>         | <code>_Atomic size_t</code>             |
| <code>atomic_ptrdiff_t</code>      | <code>_Atomic ptrdiff_t</code>          |
| <code>atomic_intmax_t</code>       | <code>_Atomic intmax_t</code>           |
| <code>atomic_uintmax_t</code>      | <code>_Atomic uintmax_t</code>          |

You can make your own additional types with the `_Atomic` type qualifier:

```
_Atomic double x;
```

or the `_Atomic()` type specifier:

```
_Atomic(double) x;
```



## 18.2 Lock-free Macros

These macros let you know if a type is lock-free or not. Maybe.

They can be used at compile time with `#if`. They apply to both signed and unsigned types.

| Atomic Type                  | Lock Free Macro                        |
|------------------------------|----------------------------------------|
| <code>atomic_bool</code>     | <code>ATOMIC_BOOL_LOCK_FREE</code>     |
| <code>atomic_char</code>     | <code>ATOMIC_CHAR_LOCK_FREE</code>     |
| <code>atomic_char16_t</code> | <code>ATOMIC_CHAR16_T_LOCK_FREE</code> |
| <code>atomic_char32_t</code> | <code>ATOMIC_CHAR32_T_LOCK_FREE</code> |
| <code>atomic_wchar_t</code>  | <code>ATOMIC_WCHAR_T_LOCK_FREE</code>  |
| <code>atomic_short</code>    | <code>ATOMIC_SHORT_LOCK_FREE</code>    |
| <code>atomic_int</code>      | <code>ATOMIC_INT_LOCK_FREE</code>      |
| <code>atomic_long</code>     | <code>ATOMIC_LONG_LOCK_FREE</code>     |
| <code>atomic_llong</code>    | <code>ATOMIC_LLONG_LOCK_FREE</code>    |
| <code>atomic_intptr_t</code> | <code>ATOMIC_POINTER_LOCK_FREE</code>  |

These macros can interestingly have *three* different values:

| Value | Meaning                                   |
|-------|-------------------------------------------|
| 0     | Never lock-free.                          |
| 1     | <i>Sometimes</i> lock-free <sup>1</sup> . |
| 2     | Always lock-free.                         |

## 18.3 Atomic Flag

The `atomic_flag` opaque type is the only type guaranteed to be lock-free. Though your PC implementation probably does a lot more.

It is accessed through the `atomic_flag_test_and_set()` and `atomic_flag_clear()` functions.

Before use, it can be initialized to a clear state with:

```
atomic_flag f = ATOMIC_FLAG_INIT;
```

## 18.4 Memory Order

This header introduces a new `enum` type called `memory_order`. This is used by a bunch of the functions to specify memory orders other than sequential consistency.

| <code>memory_order</code>         | Description            |
|-----------------------------------|------------------------|
| <code>memory_order_seq_cst</code> | Sequential Consistency |
| <code>memory_order_acq_rel</code> | Acquire/Release        |
| <code>memory_order_release</code> | Release                |
| <code>memory_order_acquire</code> | Acquire                |
| <code>memory_order_consume</code> | Consume                |

<sup>1</sup>Maybe it depends on the run-time environment and can't be known at compile-time.

| memory_order         | Description |
|----------------------|-------------|
| memory_order_relaxed | Relaxed     |

You can feed these into atomic functions with the `_explicit` suffix.

The non-`_explicit` versions of the functions are the same as if you'd called the `_explicit` counterpart with `memory_order_seq_cst`.

## 18.5 ATOMIC\_VAR\_INIT()

Create an initializer for an atomic variable

### Synopsis

```
#include <stdatomic.h>

#define ATOMIC_VAR_INIT(C value)    // Deprecated
```

### Description

This macro expands to an initializer, so you can use it when a variable is defined.

The type of the `value` should be the base type of the atomic variable.

The initialization itself is *not* an atomic operation, ironically.

CPPReference says this is deprecated<sup>2</sup> and likely to be removed. Standards document p1138r0<sup>3</sup> elaborates that the macro is limited in that it can't properly initialize atomic `structs`, and its original *raison d'être* turned out to not be useful.

Just initialize the variable straight-up, instead.

### Return Value

Expands to an initializer suitable for this atomic variable.

### Example

```
1 #include <stdio.h>
2 #include <stdatomic.h>
3
4 int main(void)
5 {
6     atomic_int x = ATOMIC_VAR_INIT(3490); // Deprecated
7     printf("%d\n", x);
8 }
```

<sup>2</sup>[https://en.cppreference.com/w/cpp/atomic/ATOMIC\\_VAR\\_INIT](https://en.cppreference.com/w/cpp/atomic/ATOMIC_VAR_INIT)

<sup>3</sup><http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2018/p1138r0.pdf>

## See Also

`atomic_init()`

---

## 18.6 `atomic_init()`

Initialize an atomic variable

### Synopsis

```
#include <stdatomic.h>

void atomic_init(volatile A *obj, C value);
```

### Description

You can use this to initialize an atomic variable.

The type of the `value` should be the base type of the atomic variable.

The initialization itself is *not* an atomic operation, ironically.

As far as I can tell, there's no difference between this and assigning directly to the atomic variable. The spec says it's there to allow the compiler to inject any additional initialization that needs doing, but everything seems fine without it. If anyone has more info, send it my way.

### Return Value

Returns nothing!

### Example

```
1  #include <stdio.h>
2  #include <stdatomic.h>
3
4  int main(void)
5  {
6      atomic_int x;
7
8      atomic_init(&x, 3490);
9
10     printf("%d\n", x);
11 }
```

## See Also

`ATOMIC_VAR_INIT()`, `atomic_store()`, `atomic_store_explicit()`

---

## 18.7 kill\_dependency()

End a dependency chain

### Synopsis

```
#include <stdatomic.h>

type kill_dependency(type y);
```

### Description

This is potentially useful for optimizing if you're using `memory_order_consume` anywhere.

And if you know what you're doing. If unsure, learn more before trying to use this.

### Return Value

Returns the value passed in.

### Example

In this example, `i` carries a dependency into `x`. And would do into `y`, except for the call to `kill_dependency()`.

```
1  #include <stdio.h>
2  #include <stdatomic.h>
3
4  int main(void)
5  {
6      atomic_int a;
7      int i = 10, x, y;
8
9      atomic_store_explicit(&a, 3490, memory_order_release);
10
11     i = atomic_load_explicit(&a, memory_order_consume);
12     x = i;
13     y = kill_dependency(i);
14
15     printf("%d %d\n", x, y); // 3490 and either 3490 or 10
16 }
```

---

## 18.8 atomic\_thread\_fence()

Set up a fence

## Synopsis

```
#include <stdatomic.h>

void atomic_thread_fence(memory_order order);
```

## Description

This sets up a memory fence with the specified `order`.

| order                             | Description                                    |
|-----------------------------------|------------------------------------------------|
| <code>memory_order_seq_cst</code> | Sequentially consistency acquire/release fence |
| <code>memory_order_acq_rel</code> | Acquire/release fence                          |
| <code>memory_order_release</code> | Release fence                                  |
| <code>memory_order_acquire</code> | Acquire fence                                  |
| <code>memory_order_consume</code> | Acquire fence (again)                          |
| <code>memory_order_relaxed</code> | No fence at all—no point in calling with this  |

You might try to avoid using these and just stick with the different modes with `atomic_store_explicit()` and `atomic_load_explicit()`. Or not.

## Return Value

Returns nothing!

## Example

```
1  #include <stdio.h>
2  #include <threads.h>
3  #include <stdatomic.h>
4
5  atomic_int shared_1 = 1;
6  atomic_int shared_2 = 2;
7
8  int thread_1(void *arg)
9  {
10     (void)arg;
11
12     atomic_store_explicit(&shared_1, 10, memory_order_relaxed);
13
14     atomic_thread_fence(memory_order_release);
15
16     atomic_store_explicit(&shared_2, 20, memory_order_relaxed);
17
18     return 0;
19 }
20
21 int thread_2(void *arg)
22 {
23     (void)arg;
24
```

```

25 // If this fence runs after the release fence, we're
26 // guaranteed to see thread_1's changes to the shared
27 // variables.
28
29 atomic_thread_fence(memory_order_acquire);
30
31 if (shared_2 == 20) {
32     printf("Shared_1 better be 10 and it's %d\n", shared_1);
33 } else {
34     printf("Anything's possible: %d %d\n", shared_1, shared_2);
35 }
36
37 return 0;
38 }
39
40 int main(void)
41 {
42     thrd_t t1, t2;
43
44     thrd_create(&t2, thread_2, NULL);
45     thrd_create(&t1, thread_1, NULL);
46
47     thrd_join(t1, NULL);
48     thrd_join(t2, NULL);
49 }

```

## See Also

`atomic_store_explicit()`, `atomic_load_explicit()`, `atomic_signal_fence()`

---

## 18.9 `atomic_signal_fence()`

Fence for intra-thread signal handlers

### Synopsis

```

#include <stdatomic.h>

void atomic_signal_fence(memory_order order);

```

### Description

This works like `atomic_thread_fence()` except its purpose is within in a single thread; notably for use in a signal handler in that thread.

Since signals can happen at any time, we might need a way to be certain that any writes by the thread that happened before the signal handler be visible within that signal handler.

## Return Value

Returns nothing!

## Example

Partial demo. (Note that it's technically undefined behavior to call `printf()` in a signal handler.)

```
1  #include <stdio.h>
2  #include <signal.h>
3  #include <stdatomic.h>
4
5  int global;
6
7  void handler(int sig)
8  {
9      (void)sig;
10
11     // If this runs before the release, the handler will
12     // potentially see global == 0.
13     //
14     // Otherwise, it will definitely see global == 10.
15
16     atomic_signal_fence(memory_order_acquire);
17
18     printf("%d\n", global);
19 }
20
21 int main(void)
22 {
23     signal(SIGINT, handler);
24
25     global = 10;
26
27     atomic_signal_fence(memory_order_release);
28
29     // If the signal handler runs after the release
30     // it will definitely see the value 10 in global.
31 }
```

## See Also

`atomic_thread_fence()`, `signal()`

---

## 18.10 `atomic_is_lock_free()`

Determine if an atomic type is lock free

## Synopsis

```
#include <stdatomic.h>

_Bool atomic_is_lock_free(const volatile A *obj);
```

## Description

Determines if the variable `obj` of type `A` is lock-free. Can be used with any type.

Unlike the lock-free macros which can be used at compile-time, this is strictly a run-time function. So in places where the macros say “maybe”, this function will definitely tell you one way or another if the atomic variable is lock-free.

This is useful when you’re defining your own atomic variables and want to know their lock-free status.

## Return Value

True if the variable is lock-free, false otherwise.

## Example

Test if a couple `structs` and an atomic `double` are lock-free. On my system, the larger `struct` is too big to be lock-free, but the other two are OK.

```
1  #include <stdio.h>
2  #include <stdatomic.h>
3
4  int main(void)
5  {
6      struct foo {
7          int x, y;
8      };
9
10     struct bar {
11         int x, y, z;
12     };
13
14     _Atomic(double) a;
15     struct foo b;
16     struct bar c;
17
18     printf("a is lock-free: %d\n", atomic_is_lock_free(&a));
19     printf("b is lock-free: %d\n", atomic_is_lock_free(&b));
20     printf("c is lock-free: %d\n", atomic_is_lock_free(&c));
21 }
```

Output on my system (YMMV):

```
a is lock-free: 1
b is lock-free: 1
c is lock-free: 0
```



## See Also

Lock-free Macros

---

### 18.11 atomic\_store()

Store a value in an atomic variable

#### Synopsis

```
#include <stdatomic.h>

void atomic_store(volatile A *object, C desired);

void atomic_store_explicit(volatile A *object,
                           C desired, memory_order order);
```

#### Description

Store a value in an atomic variable, possibly synchronized.

This is like a plain assignment, but with more flexibility.

These have the same storage effect for an `atomic_int x`:

```
x = 10;
atomic_store(&x, 10);
atomic_store_explicit(&x, 10, memory_order_seq_cst);
```

But the last function, `atomic_store_explicit()`, lets you specify the memory order.

Since this is a “release-y” operation, none of the “acquire-y” memory orders are legal. `order` can be only be `memory_order_seq_cst`, `memory_order_release`, or `memory_order_relaxed`.

`order` cannot be `memory_order_acq_rel`, `memory_order_acquire`, or `memory_order_consume`.

#### Return Value

Returns nothing!

#### Example

```
1  #include <stdio.h>
2  #include <stdatomic.h>
3
4  int main(void)
5  {
6      atomic_int x = 0;
7      atomic_int y = 0;
8
9      atomic_store(&x, 10);
10
```

```

11     atomic_store_explicit(&y, 20, memory_order_relaxed);
12
13     // Will print either "10 20" or "10 0":
14     printf("%d %d\n", x, y);
15 }

```

## See Also

atomic\_init(), atomic\_load(), atomic\_load\_explicit(), atomic\_exchange(),  
 atomic\_exchange\_explicit(), atomic\_compare\_exchange\_strong(),  
 atomic\_compare\_exchange\_strong\_explicit(), atomic\_compare\_exchange\_weak(),  
 atomic\_compare\_exchange\_weak\_explicit(), atomic\_fetch\_\*

## 18.12 atomic\_load()

Return a value from an atomic variable

### Synopsis

```

#include <stdatomic.h>

C atomic_load(const volatile A *object);

C atomic_load_explicit(const volatile A *object, memory_order order);

```

### Description

For a pointer to an `object` of type `A`, atomically returns its value `C`. This is a generic function that can be used with any type.

The function `atomic_load_explicit()` lets you specify the memory order.

Since this is an “acquire-y” operation, none of the “release-y” memory orders are legal. `order` can be only be `memory_order_seq_cst`, `memory_order_acquire`, `memory_order_consume`, or `memory_order_relaxed`.

`order` cannot be `memory_order_acq_rel` or `memory_order_release`.

### Return Value

Returns the value stored in `object`.

### Example

```

1  #include <stdio.h>
2  #include <stdatomic.h>
3
4  int main(void)
5  {
6      atomic_int x = 10;

```

```
7
8     int v = atomic_load(&x);
9
10    printf("%d\n", v); // 10
11 }
```

## See Also

`atomic_store()`, `atomic_store_explicit()`

---

## 18.13 `atomic_exchange()`

Replace a value in an atomic object

### Synopsis

```
#include <stdatomic.h>

C atomic_exchange(volatile A *object, C desired);

C atomic_exchange_explicit(volatile A *object, C desired,
                           memory_order order);
```

### Description

Sets the value in `object` to `desired`.

`object` is type `A`, some atomic type.

`desired` is type `C`, the respective non-atomic type to `A`.

This is very similar to `atomic_store()`, except the previous value is atomically returned.

### Return Value

Returns the previous value of `object`.

### Example

```
1  #include <stdio.h>
2  #include <stdatomic.h>
3
4  int main(void)
5  {
6      atomic_int x = 10;
7
8      int previous = atomic_exchange(&x, 20);
9
10     printf("x is %d\n", x);
11     printf("x was %d\n", previous);
```

12

}

Output:

```
x is 20
x was 10
```

**See Also**

```
atomic_init(), atomic_load(), atomic_load_explicit(), atomic_store(),
atomic_store_explicit() atomic_compare_exchange_strong(),
atomic_compare_exchange_strong_explicit(), atomic_compare_exchange_weak(),
atomic_compare_exchange_weak_explicit()
```

**18.14 atomic\_compare\_exchange\_\*()**

Atomic compare and exchange

**Synopsis**

```
#include <stdatomic.h>

_Bool atomic_compare_exchange_strong(volatile A *object,
                                     C *expected, C desired);

_Bool atomic_compare_exchange_strong_explicit(volatile A *object,
                                              C *expected, C desired,
                                              memory_order success,
                                              memory_order failure);

_Bool atomic_compare_exchange_weak(volatile A *object,
                                   C *expected, C desired);

_Bool atomic_compare_exchange_weak_explicit(volatile A *object,
                                            C *expected, C desired,
                                            memory_order success,
                                            memory_order failure);
```

**Description**

The venerable basis for some many things lock-free: compare and exchange.

In the above prototypes, `A` is the type of the atomic object, and `C` is the equivalent base type.

Ignoring the `_explicit` versions for a moment, what these do is:

- If the value pointed to by `object` is equal to the value pointed to by `expected`, then the value pointed to by `object` is set to `desired`. And the function returns `true` indicating the exchange did take place.
- Else the value pointed to by `expected` (yes, `expected`) is set to `desired` and the function returns `false` indicating the exchange did not take place.

Pseudocode for the exchange would look like this<sup>4</sup>:

```
bool compare_exchange(atomic_A *object, C *expected, C desired)
{
    if (*object is the same as *expected) {
        *object = desired
        return true
    }

    *expected = desired
    return false
}
```

The `_weak` variants might spontaneously fail, so even if `*object == *desired`, it might not change the value and will return `false`. So you'll want that in a loop if you use it<sup>5</sup>.

The `_explicit` variants have two memory orders: `success` if `*object` is set to `desired`, and `failure` if it is not.

These are test-and-set functions, so you can use `memory_order_acq_rel` with the `_explicit` variants.

## Return Value

Returns `true` if `*object` was `*expected`. Otherwise, `false`.

## Example

A contrived example where multiple threads add 2 to a shared value in a lock-free way.

(It would be better to use `+= 2` to get this done in real life unless you were using some `_explicit` wizardry.)

```
1  #include <stdio.h>
2  #include <threads.h>
3  #include <stdatomic.h>
4
5  #define LOOP_COUNT 10000
6
7  atomic_int value;
8
9  int run(void *arg)
10 {
11     (void)arg;
12
13     for(int i = 0; i < LOOP_COUNT; i++) {
14
15         int cur = value;
16         int next;
17
18         do {
19             next = cur + 2;
```

<sup>4</sup>This effectively does the same thing, but it's clearly not atomic.

<sup>5</sup>The spec says, "This spurious failure enables implementation of compare-and-exchange on a broader class of machines, e.g. load-locked store-conditional machines." And adds, "When a compare-and-exchange is in a loop, the weak version will yield better performance on some platforms. When a weak compare-and-exchange would require a loop and a strong one would not, the strong one is preferable."

```

20     } while (!atomic_compare_exchange_strong(&value, &cur, next));
21 }
22
23 return 0;
24 }
25
26 int main(void)
27 {
28     thrd_t t1, t2;
29
30     thrd_create(&t1, run, NULL);
31     thrd_create(&t2, run, NULL);
32
33     thrd_join(t1, NULL);
34     thrd_join(t2, NULL);
35
36     printf("%d should equal %d\n", value, LOOP_COUNT * 4);
37 }

```

Just replacing this with `value = value + 2` causes data trampling.

## See Also

`atomic_load()`, `atomic_load_explicit()`, `atomic_store()`, `atomic_store_explicit()`, `atomic_exchange()`, `atomic_exchange_explicit()`, `atomic_fetch_*`()

## 18.15 `atomic_fetch_*`()

Atomically modify atomic variables

### Synopsis

```

#include <stdatomic.h>

C atomic_fetch_KEY(volatile A *object, M operand);

C atomic_fetch_KEY_explicit(volatile A *object, M operand,
                           memory_order order);

```

### Description

These are actually a group of 10 functions. You substitute one of the following for `KEY` to perform that operation:

- `add`
- `sub`
- `or`
- `xor`
- `and`

So these functions can add or subtract values to or from an atomic variable, or can perform bitwise-OR, XOR, or AND on them.

Use it with integer or pointer types. Though the spec is a little vague on the matter, other types make C unhappy. It goes out of its way to avoid undefined behavior with signed integers, as well:

C18 §7.17.7.5 ¶3:

For signed integer types, arithmetic is defined to use two's complement representation with silent wrap-around on overflow; there are no undefined results.

In the synopsis, above, **A** is an atomic type, and **M** is the corresponding non-atomic type for **A** (or `ptrdiff_t` for atomic pointers), and **C** is the corresponding non-atomic type for **A**.

For example, here are some operations on an `atomic_int`.

```
atomic_fetch_add(&x, 20);
atomic_fetch_sub(&x, 37);
atomic_fetch_xor(&x, 3490);
```

They are the same as `+=`, `-=`, `|=`, `^=` and `&=`, except the return value is the *previous* value of the atomic object. (With the assignment operators, the value of the expression is that *after* its evaluation.)

```
atomic_int x = 10;
int prev = atomic_fetch_add(&x, 20);
printf("%d %d\n", prev, x); // 10 30
```

versus:

```
atomic_int x = 10;
int prev = (x += 20);
printf("%d %d\n", prev, x); // 30 30
```

And, of course, the `_explicit` version allows you to specify a memory order and all the assignment operators are `memory_order_seq_cst`.

## Return Value

Returns the previous value of the atomic object before the modification.

## Example

```
1 #include <stdio.h>
2 #include <stdatomic.h>
3
4 int main(void)
5 {
6     atomic_int x = 0;
7     int prev;
8
9     atomic_fetch_add(&x, 3490);
10    atomic_fetch_sub(&x, 12);
11    atomic_fetch_xor(&x, 444);
12    atomic_fetch_or(&x, 12);
```

```

13     prev = atomic_fetch_and(&x, 42);
14
15     printf("%d %d\n", prev, x);    // 3118 42
16 }

```

## See Also

atomic\_exchange(), atomic\_exchange\_explicit(), atomic\_compare\_exchange\_strong(),  
atomic\_compare\_exchange\_strong\_explicit(), atomic\_compare\_exchange\_weak(),  
atomic\_compare\_exchange\_weak\_explicit()

## 18.16 atomic\_flag\_test\_and\_set()

Test and set an atomic flag

### Synopsis

```

#include <stdatomic.h>

_Bool atomic_flag_test_and_set(volatile atomic_flag *object);

_Bool atomic_flag_test_and_set_explicit(volatile atomic_flag *object,
                                       memory_order order);

```

### Description

One of the venerable old functions of lock-free programming, this function sets the given atomic flag in `object`, and returns the previous value of the flag.

As usual, the `_explicit` allows you to specify an alternate memory order.

### Return Value

Returns `true` if the flag was set previously, and `false` if it wasn't.

### Example

Using test-and-set to implement a spin lock<sup>6</sup>:

```

1  #include <stdio.h>
2  #include <threads.h>
3  #include <stdatomic.h>
4
5  // Shared non-atomic struct
6  struct {
7      int x, y, z;
8  } s = {1, 2, 3};

```

<sup>6</sup>Don't use this unless you know what you're doing—use the thread mutex functionality instead. It'll let your blocked thread sleep and stop chewing up CPU.



```

9
10 atomic_flag f = ATOMIC_FLAG_INIT;
11
12 int run(void *arg)
13 {
14     int tid = *(int*)arg;
15
16     printf("Thread %d: waiting for lock...\n", tid);
17
18     while (atomic_flag_test_and_set(&f));
19
20     printf("Thread %d: got lock, s is {%d, %d, %d}\n", tid,
21                                                  s.x, s.y, s.z);
22     s.x = (tid + 1) * 5 + 0;
23     s.y = (tid + 1) * 5 + 1;
24     s.z = (tid + 1) * 5 + 2;
25     printf("Thread %d: set s to {%d, %d, %d}\n", tid, s.x, s.y, s.z);
26
27     printf("Thread %d: releasing lock...\n", tid);
28     atomic_flag_clear(&f);
29
30     return 0;
31 }
32
33 int main(void)
34 {
35     thrd_t t1, t2;
36     int tid[] = {0, 1};
37
38     thrd_create(&t1, run, tid+0);
39     thrd_create(&t2, run, tid+1);
40
41     thrd_join(t1, NULL);
42     thrd_join(t2, NULL);
43 }

```

Example output (varies run to run):

```

Thread 0: waiting for lock...
Thread 0: got lock, s is {1, 2, 3}
Thread 1: waiting for lock...
Thread 0: set s to {5, 6, 7}
Thread 0: releasing lock...
Thread 1: got lock, s is {5, 6, 7}
Thread 1: set s to {10, 11, 12}
Thread 1: releasing lock...

```

## See Also

`atomic_flag_clear()`

---

## 18.17 atomic\_flag\_clear()

Clear an atomic flag

### Synopsis

```
#include <stdatomic.h>

void atomic_flag_clear(volatile atomic_flag *object);

void atomic_flag_clear_explicit(volatile atomic_flag *object,
                               memory_order order);
```

### Description

Clears an atomic flag.

As usual, the `_explicit` allows you to specify an alternate memory order.

### Return Value

Returns nothing!

### Example

Using test-and-set to implement a spin lock<sup>7</sup>:

```
1  #include <stdio.h>
2  #include <threads.h>
3  #include <stdatomic.h>
4
5  // Shared non-atomic struct
6  struct {
7      int x, y, z;
8  } s = {1, 2, 3};
9
10 atomic_flag f = ATOMIC_FLAG_INIT;
11
12 int run(void *arg)
13 {
14     int tid = *(int*)arg;
15
16     printf("Thread %d: waiting for lock...\n", tid);
17
18     while (atomic_flag_test_and_set(&f));
19
20     printf("Thread %d: got lock, s is {%d, %d, %d}\n", tid,
21                                                s.x, s.y, s.z);
22     s.x = (tid + 1) * 5 + 0;
23     s.y = (tid + 1) * 5 + 1;
```

<sup>7</sup>Don't use this unless you know what you're doing—use the thread mutex functionality instead. It'll let your blocked thread sleep and stop chewing up CPU.

```
24     s.z = (tid + 1) * 5 + 2;
25     printf("Thread %d: set s to {%d, %d, %d}\n", tid, s.x, s.y, s.z);
26
27     printf("Thread %d: releasing lock...\n", tid);
28     atomic_flag_clear(&f);
29
30     return 0;
31 }
32
33 int main(void)
34 {
35     thrd_t t1, t2;
36     int tid[] = {0, 1};
37
38     thrd_create(&t1, run, tid+0);
39     thrd_create(&t2, run, tid+1);
40
41     thrd_join(t1, NULL);
42     thrd_join(t2, NULL);
43 }
```

Example output (varies run to run):

```
Thread 0: waiting for lock...
Thread 0: got lock, s is {1, 2, 3}
Thread 1: waiting for lock...
Thread 0: set s to {5, 6, 7}
Thread 0: releasing lock...
Thread 1: got lock, s is {5, 6, 7}
Thread 1: set s to {10, 11, 12}
Thread 1: releasing lock...
```

## See Also

`atomic_flag_test_and_set()`

## Chapter 19

# <stdbit.h> Bit-Related Functions

This header file and all the included functions are new in C23!

Caveat: none of my compilers support this, so all the code is untested!

| Function                                | Description                                                  |
|-----------------------------------------|--------------------------------------------------------------|
| <code>stdc_bit_ceil()</code>            | Return the smallest power of 2 not less than a number        |
| <code>stdc_bit_floor()</code>           | Return the largest power of 2 not greater than a number      |
| <code>stdc_bit_width()</code>           | Return the number of bits needed to store a value            |
| <code>stdc_count_ones()</code>          | Count the ones in an unsigned number                         |
| <code>stdc_count_zeros()</code>         | Count the zeros in an unsigned number                        |
| <code>stdc_first_leading_one()</code>   | Find the first leading one in an unsigned number             |
| <code>stdc_first_leading_zero()</code>  | Find the first leading zero in an unsigned number            |
| <code>stdc_first_trailing_one()</code>  | Find the first trailing one in an unsigned number            |
| <code>stdc_first_trailing_zero()</code> | Find the first trailing zero in an unsigned number           |
| <code>stdc_has_single_bit()</code>      | Test to see if an unsigned integer has only a single bit set |
| <code>stdc_leading_ones()</code>        | Count the leading ones in an unsigned number                 |
| <code>stdc_leading_zeros()</code>       | Count the leading zeros in an unsigned number                |
| <code>stdc_trailing_ones()</code>       | Count the trailing ones in an unsigned number                |
| <code>stdc_trailing_zeros()</code>      | Count the trailing zeros in an unsigned number               |

### 19.1 Existence Macro

Since this is a new feature, you can test for its existence by making sure the `__STDC_VERSION_STDBIT_H__` macro exists and is `202311L` or greater.

### 19.2 Endianness Macros

This header file defines a number of macros that can be used to determine the *endianness* of the system. (That is, in a multibyte value, do the first byte represent the most significant part of the value, or least? Or neither?)

There are two different values for the defined endiannesses `__STDC_ENDIAN_BIG__` and `__STDC_ENDIAN_LITTLE__`.

Additionally, there's a macro `__STDC_ENDIAN_NATIVE__` that gives you the endianness of *this* system. You can compare that against the other two macros to see what you've got. If it doesn't compare equal to either of them, you must be on some other mixed-endian system.

Let's see what this system is:

```

#include <stdio.h>
#include <stdbit.h>

int main(void)
{
    switch(__STDC_ENDIAN_NATIVE__) {
        case __STDC_ENDIAN_BIG__:
            puts("Big-endian!");
            break;

        case __STDC_ENDIAN_LITTLE__:
            puts("Little-endian!");
            break;

        default:
            puts("Other-endian!");
    }
}

```

### 19.3 General Structure of These Functions

All the functions in this header follow a standard pattern, so we might as well go over it once here up-front to save extraneous and excessive verbosity later.

Each function comes in six amazing forms, depending on the types involved.

First of all, they're just going to operate on unsigned integer types, so we'll get that out of the way.

And they're *mostly* going to return `unsigned int`.

And they come in two subforms: the type-specific and type-generic forms.

Let's look at the type specific forms first. We'll use the function that counts the leading zero bits in a value as an example. Here it is:

```

unsigned int stdc_leading_zeros_uc(unsigned char value);
unsigned int stdc_leading_zeros_us(unsigned short value);
unsigned int stdc_leading_zeros_ui(unsigned int value);
unsigned int stdc_leading_zeros_ul(unsigned long value);
unsigned int stdc_leading_zeros_ull(unsigned long long value);

```

Yikes. OK, what do we have?

Well, in order to prevent namespace pollution, they all start with `stdc_`.

And what about this `uc`, `us`, `ui`, etc. stuff? Those correspond to the type of the argument, you'll see with a little inspection.

| Suffix            | Parameter Type     |
|-------------------|--------------------|
| <code>_uc</code>  | unsigned char      |
| <code>_us</code>  | unsigned short     |
| <code>_ui</code>  | unsigned int       |
| <code>_ul</code>  | unsigned long      |
| <code>_ull</code> | unsigned long long |

But wait! There's more!

Each of these function families in the header also has a *generic* variant.

So, again with `count_leading_zeros` as an example, there's a version that doesn't need types to be specified, and it's just the function without any suffix.

```
generic_return_type stdc_leading_zeros(generic_value_type value);
```

In that example, the words `generic_return_type` and `generic_value_type` are **not** C keywords. They're just placeholders to let you know you need to put the right thing there.

```
unsigned short x = 3490;

unsigned int a = stdc_leading_zeros(x);
auto b = stdc_leading_zeros(1234);
```

The generic functions work with all unsigned types, not counting `bool`.

---

## 19.4 `stdc_leading_zeros()`

Count the leading zeros in an unsigned number

### Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_leading_zeros_uc(unsigned char value);
unsigned int stdc_leading_zeros_us(unsigned short value);
unsigned int stdc_leading_zeros_ui(unsigned int value);
unsigned int stdc_leading_zeros_ul(unsigned long value);
unsigned int stdc_leading_zeros_ull(unsigned long long value);

generic_return_type stdc_leading_zeros(generic_value_type value);
```

### Description

This function returns the number of leading zero bits on a particular value, starting from the most-significant bit. The number will be influenced by the size of the argument.

### Return Value

Returns the number of leading zero bits.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 3490;
7      unsigned int count = stdc_leading_zeros_ui(value);
8
9      printf("%u\n", count);
10
11     unsigned long long value2 = 3490;
12     auto count2 = stdc_leading_zeros(value2);
13
14     printf("%u\n", count2);
15 }
```

## See Also

`stdc_leading_ones()`, `stdc_trailing_zeros()`

---

## 19.5 `stdc_leading_ones()`

Count the leading ones in an unsigned number

### Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_leading_ones_uc(unsigned char value);
unsigned int stdc_leading_ones_us(unsigned short value);
unsigned int stdc_leading_ones_ui(unsigned int value);
unsigned int stdc_leading_ones_ul(unsigned long value);
unsigned int stdc_leading_ones_ull(unsigned long long value);

generic_return_type stdc_leading_ones(generic_value_type value);
```

### Description

This function returns the number of leading one bits on a particular value starting from the most-significant bit. The number will be influenced by the size of the argument.

### Return Value

Returns the number of leading one bits.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 3490;
7      unsigned int count = stdc_leading_ones_ui(value);
8
9      printf("%u\n", count);
10
11     unsigned long long value2 = 3490;
12     auto count2 = stdc_leading_ones(value2);
13
14     printf("%u\n", count2);
15 }
```

## See Also

`stdc_leading_zeros()`, `stdc_trailing_ones()`

---

## 19.6 `stdc_trailing_zeros()`

Count the trailing zeros in an unsigned number

### Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_trailing_zeros_uc(unsigned char value);
unsigned int stdc_trailing_zeros_us(unsigned short value);
unsigned int stdc_trailing_zeros_ui(unsigned int value);
unsigned int stdc_trailing_zeros_ul(unsigned long value);
unsigned int stdc_trailing_zeros_ull(unsigned long long value);

generic_return_type stdc_trailing_zeros(generic_value_type value);
```

### Description

This function returns the number of trailing zero bits on a particular value, starting from the least-significant bit. The number will be influenced by the size of the argument.

### Return Value

Returns the number of trailing zero bits.



## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 3490;
7      unsigned int count = stdc_trailing_zeros_ui(value);
8
9      printf("%u\n", count);
10
11     unsigned long long value2 = 3490;
12     auto count2 = stdc_trailing_zeros(value2);
13
14     printf("%u\n", count2);
15 }
```

## See Also

`stdc_trailing_ones()`, `stdc_leading_zeros()`

---

## 19.7 `stdc_trailing_ones()`

Count the trailing ones in an unsigned number

### Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_trailing_ones_uc(unsigned char value);
unsigned int stdc_trailing_ones_us(unsigned short value);
unsigned int stdc_trailing_ones_ui(unsigned int value);
unsigned int stdc_trailing_ones_ul(unsigned long value);
unsigned int stdc_trailing_ones_ull(unsigned long long value);

generic_return_type stdc_trailing_ones(generic_value_type value);
```

### Description

This function returns the number of trailing one bits on a particular value, starting from the least-significant bit. The number will be influenced by the size of the argument.

### Return Value

Returns the number of trailing one bits.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 3490;
7      unsigned int count = stdc_trailing_ones_ui(value);
8
9      printf("%u\n", count);
10
11     unsigned long long value2 = 3490;
12     auto count2 = stdc_trailing_ones(value2);
13
14     printf("%u\n", count2);
15 }
```

## See Also

`stdc_trailing_zeros()`, `stdc_leading_ones()`

---

## 19.8 `stdc_first_leading_zero()`

Find the first leading zero in an unsigned number

### Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_first_leading_zero_uc(unsigned char value);
unsigned int stdc_first_leading_zero_us(unsigned short value);
unsigned int stdc_first_leading_zero_ui(unsigned int value);
unsigned int stdc_first_leading_zero_ul(unsigned long value);
unsigned int stdc_first_leading_zero_ull(unsigned long long value);

generic_return_type stdc_first_leading_zero(generic_value_type value);
```

### Description

This finds the index of the first leading zero in a number. Indexes are numbered starting from **1** being the most significant (“leftmost”) bit position. (This might be in contrast to how you might be used to numbering bit indexes.)

It’s one-based so that you can quickly use the return value as a Boolean expression for if it found a zero bit or not.

## Return Value

Returns the 1-based index from the most significant bit of the first zero bit in the `value`, or `0` if there are no zero bits.

## Example

```

1  #include <stdio.h>
2  #include <limits.h>
3  #include <stdbit.h>
4
5  int main(void)
6  {
7      unsigned int value = UINT_MAX;
8      unsigned int index = stdc_first_leading_zero_ui(value);
9
10     printf("%u\n", index);
11
12     unsigned long long value2 = UINT_MAX >> 2;
13     auto index2 = stdc_first_leading_zero(value2);
14
15     printf("%u\n", index2);
16 }
```

## See Also

`stdc_first_leading_one()`, `stdc_first_trailing_zero()`

## 19.9 `stdc_first_leading_one()`

Find the first leading one in an unsigned number

### Synopsis

New in C23!

```

#include <stdbit.h>

unsigned int stdc_first_leading_one_uc(unsigned char value);
unsigned int stdc_first_leading_one_us(unsigned short value);
unsigned int stdc_first_leading_one_ui(unsigned int value);
unsigned int stdc_first_leading_one_ul(unsigned long value);
unsigned int stdc_first_leading_one_ull(unsigned long long value);

generic_return_type stdc_first_leading_one(generic_value_type value);
```

### Description

This finds the index of the first leading one in a number. Indexes are numbered starting from `1` being the most significant (“leftmost”) bit position. (This might be in contrast to how you might be used to numbering bit indexes.)

It's one-based so that you can quickly use the return value as a Boolean expression for if it found a one bit or not.

## Return Value

Returns the 1-based index from the most significant bit of the first one bit in the `value`, or 0 if there are no one bits.

## Example

```

1  #include <stdio.h>
2  #include <limits.h>
3  #include <stdbit.h>
4
5  int main(void)
6  {
7      unsigned int value = UINT_MAX;
8      unsigned int index = stdc_first_leading_one_ui(value);
9
10     printf("%u\n", index);
11
12     unsigned long long value2 = UINT_MAX >> 2;
13     auto index2 = stdc_first_leading_one(value2);
14
15     printf("%u\n", index2);
16 }
```

## See Also

`stdc_first_leading_zero()`, `stdc_first_trailing_one()`

## 19.10 `stdc_first_trailing_zero()`

Find the first trailing zero in an unsigned number

## Synopsis

New in C23!

```

#include <stdbit.h>

unsigned int stdc_first_trailing_zero_uc(unsigned char value);
unsigned int stdc_first_trailing_zero_us(unsigned short value);
unsigned int stdc_first_trailing_zero_ui(unsigned int value);
unsigned int stdc_first_trailing_zero_ul(unsigned long value);
unsigned int stdc_first_trailing_zero_ull(unsigned long long value);

generic_return_type stdc_first_trailing_zero(generic_value_type value);
```

## Description

This finds the index of the first trailing zero in a number. Indexes are numbered starting from 1 being the most significant (“leftmost”) bit position. (This might be in contrast to how you might be used to numbering bit indexes.)

It’s one-based so that you can quickly use the return value as a Boolean expression for if it found a zero bit or not.

## Return Value

Returns the 1-based index from the most significant bit of the first zero bit in the `value`, or 0 if there are no zero bits.

## Example

```

1  #include <stdio.h>
2  #include <limits.h>
3  #include <stdbit.h>
4
5  int main(void)
6  {
7      unsigned int value = UINT_MAX;
8      unsigned int index = stdc_first_trailing_zero_ui(value);
9
10     printf("%u\n", index);
11
12     unsigned long long value2 = UINT_MAX >> 2;
13     auto index2 = stdc_first_trailing_zero(value2);
14
15     printf("%u\n", index2);
16 }
```

## See Also

`stdc_first_leading_zero()`, `stdc_first_trailing_one()`

## 19.11 `stdc_first_trailing_one()`

Find the first trailing one in an unsigned number

## Synopsis

New in C23!

```

#include <stdbit.h>

unsigned int stdc_first_trailing_one_uc(unsigned char value);
unsigned int stdc_first_trailing_one_us(unsigned short value);
unsigned int stdc_first_trailing_one_ui(unsigned int value);
unsigned int stdc_first_trailing_one_ul(unsigned long value);
```

```
unsigned int stdc_first_trailing_one_ull(unsigned long long value);

generic_return_type stdc_first_trailing_one(generic_value_type value);
```

## Description

This finds the index of the first trailing one in a number. Indexes are numbered starting from 1 being the most significant (“leftmost”) bit position. (This might be in contrast to how you might be used to numbering bit indexes.)

It’s one-based so that you can quickly use the return value as a Boolean expression for if it found a one bit or not.

## Return Value

Returns the 1-based index from the most significant bit of the first one bit in the `value`, or 0 if there are no one bits.

## Example

```
1  #include <stdio.h>
2  #include <limits.h>
3  #include <stdbit.h>
4
5  int main(void)
6  {
7      unsigned int value = UINT_MAX;
8      unsigned int index = stdc_first_trailing_one_ui(value);
9
10     printf("%u\n", index);
11
12     unsigned long long value2 = UINT_MAX >> 2;
13     auto index2 = stdc_first_trailing_one(value2);
14
15     printf("%u\n", index2);
16 }
```

## See Also

`stdc_first_leading_one()`, `stdc_first_trailing_zero()`

## 19.12 `stdc_count_zeros()`

Count the zeros in an unsigned number

## Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_count_zeros_uc(unsigned char value);
unsigned int stdc_count_zeros_us(unsigned short value);
unsigned int stdc_count_zeros_ui(unsigned int value);
unsigned int stdc_count_zeros_ul(unsigned long value);
unsigned int stdc_count_zeros_ull(unsigned long long value);

generic_return_type stdc_count_zeros(generic_value_type value);
```

## Description

This function returns the number of zero bits in a particular value. The number will be influenced by the size of the argument.

## Return Value

Returns the number of zero bits.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 3490;
7      unsigned int count = stdc_count_zeros_ui(value);
8
9      printf("%u\n", count);
10
11     unsigned long long value2 = 123456;
12     auto count2 = stdc_count_zeros(value2);
13
14     printf("%u\n", count2);
15 }
```

## See Also

`stdc_count_ones()`, `stdc_leading_zeros()`, `stdc_trailing_zeros()`

## 19.13 `stdc_count_ones()`

Count the ones in an unsigned number

## Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_count_ones_uc(unsigned char value);
unsigned int stdc_count_ones_us(unsigned short value);
unsigned int stdc_count_ones_ui(unsigned int value);
unsigned int stdc_count_ones_ul(unsigned long value);
unsigned int stdc_count_ones_ull(unsigned long long value);

generic_return_type stdc_count_ones(generic_value_type value);
```

## Description

This function returns the number of one bits in a particular value. The number will be influenced by the size of the argument.

## Return Value

Returns the number of one bits.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 3490;
7      unsigned int count = stdc_count_ones_ui(value);
8
9      printf("%u\n", count);
10
11     unsigned long long value2 = 123456;
12     auto count2 = stdc_count_ones(value2);
13
14     printf("%u\n", count2);
15 }
```

## See Also

`stdc_count_zeros()`, `stdc_leading_ones()`, `stdc_trailing_ones()`

---

## 19.14 `stdc_has_single_bit()`

Test to see if an unsigned integer has only a single bit set

## Synopsis

New in C23!



```
#include <stdbit.h>

bool stdc_has_single_bit_uc(unsigned char value);
bool stdc_has_single_bit_us(unsigned short value);
bool stdc_has_single_bit_ui(unsigned int value);
bool stdc_has_single_bit_ul(unsigned long value);
bool stdc_has_single_bit_ull(unsigned long long value);

bool stdc_has_single_bit(generic_value_type value);
```

## Description

These functions return true if exactly one bit is **1** in an unsigned **value**.

If true, it also means the **value** is a power of 2.

## Return Value

Return true if **value** has exactly a single bit set to **1**.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 0b0001000;
7      unsigned int result = stdc_has_single_bit_ui(value);
8
9      printf("%d\n", result);    // 1
10
11     unsigned long long value2 = 0b000011000
12     auto result2 = stdc_has_single_bit(value2);
13
14     printf("%d\n", result2);   // 0
15 }
```

## 19.15 stdc\_bit\_width()

Return the number of bits needed to store a value

## Synopsis

New in C23!

```
#include <stdbit.h>

unsigned int stdc_bit_width_uc(unsigned char value);
```

```

unsigned int stdc_bit_width_us(unsigned short value);
unsigned int stdc_bit_width_ui(unsigned int value);
unsigned int stdc_bit_width_ul(unsigned long value);
unsigned int stdc_bit_width_ull(unsigned long long value);

generic_return_type stdc_bit_width(generic_value_type value);

```

## Description

Given an unsigned integer value, what's the smallest number of bits needed to store it? That's the question this function answers.

## Return Value

Returns the number of bits needed to store a positive `value`. Returns `0` if the `value` passed in is `0`.

## Example

```

1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 0b0001000;
7      unsigned int result = stdc_bit_width_ui(value);
8
9      printf("%d\n", result);    // 4
10
11     unsigned long long value2 = 0b000011010
12     auto result2 = stdc_bit_width(value2);
13
14     printf("%d\n", result2);    // 5
15 }

```

## 19.16 stdc\_bit\_floor()

Return the largest power of 2 not greater than a number

## Synopsis

New in C23!

```

#include <stdbit.h>

unsigned char stdc_bit_floor_uc(unsigned char value);
unsigned short stdc_bit_floor_us(unsigned short value);
unsigned int stdc_bit_floor_ui(unsigned int value);
unsigned long stdc_bit_floor_ul(unsigned long value);
unsigned long long stdc_bit_floor_ull(unsigned long long value);

```

```
generic_value_type stdc_bit_floor(generic_value_type value);
```

## Description

This function returns the largest power of 2 that is not greater than the `value`.

In other words “is less than or equal to” the `value`.

In other words, floor down to the nearest power of two.

In other words, return the value of the highest set bit in a number

For example: `value` Return | - | `0b1010b100` `0b10000b1000` `0b1001010b100000`

## Return Value

Returns the largest power of 2 not greater than `value`. Returns 0 if the `value` is 0.

## Example

```
1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 0b0001000;
7      unsigned int result = stdc_bit_floor_ui(value);
8
9      printf("%b\n", result);    // 1000
10
11     unsigned long long value2 = 0b000011010
12     auto result2 = stdc_bit_floor(value2);
13
14     printf("%b\n", result2);    // 10000
15 }
```

## See Also

`stdc_bit_ceil()`

## 19.17 `stdc_bit_ceil()`

Return the smallest power of 2 not less than a number

## Synopsis

New in C23!

```
#include <stdbit.h>
```

```

unsigned char stdc_bit_ceil_uc(unsigned char value);
unsigned short stdc_bit_ceil_us(unsigned short value);
unsigned int stdc_bit_ceil_ui(unsigned int value);
unsigned long stdc_bit_ceil_ul(unsigned long value);
unsigned long long stdc_bit_ceil_ull(unsigned long long value);

generic_value_type stdc_bit_ceil(generic_value_type value);

```

## Description

This function returns the smallest power of 2 that is not less than the `value`.

In other words “is greater than or equal to” the `value`.

In other words, ceil up to the nearest power of two.

For example: `value` | Return | `0b1010b1000` | `0b10000b1000` | `0b1001010b1000000`

If the result doesn’t fit in the return type, the behavior is undefined.

## Return Value

Returns the smallest power of 2 not less than `value`.

## Example

```

1  #include <stdio.h>
2  #include <stdbit.h>
3
4  int main(void)
5  {
6      unsigned int value = 0b0001000;
7      unsigned int result = stdc_bit_ceil_ui(value);
8
9      printf("%b\n", result);    // 1000
10
11     unsigned long long value2 = 0b000011010
12     auto result2 = stdc_bit_ceil(value2);
13
14     printf("%b\n", result2);    // 100000
15 }

```

## See Also

`stdc_bit_floor()`

## Chapter 20

# <stdbool.h> Boolean Types

This is a small header file that defines a number of convenient Boolean macros. If you really need that kind of thing.

| Macro              | Description                                     |
|--------------------|-------------------------------------------------|
| <code>bool</code>  | Type for Boolean, expands to <code>_Bool</code> |
| <code>true</code>  | True value, expands to <code>1</code>           |
| <code>false</code> | False value, expands to <code>0</code>          |

There's one more macro that I'm not putting in the table because it's such a long name it'll blow up the table alignment:

```
__bool_true_false_are_defined
```

which expands to `1`.

### 20.1 Example

Here's a lame example that shows off these macros.

```
1  #include <stdio.h>
2  #include <stdbool.h>
3
4  int main(void)
5  {
6      bool x;
7
8      x = (3 > 2);
9
10     if (x == true)
11         printf("The universe still makes sense.\n");
12
13     x = false;
14
15     printf("x is now %d\n", x); // 0
```

16

}

Output:

```
The universe still makes sense.  
x is now 0
```

## 20.2 `_Bool`?

What's the deal with `_Bool`? Why didn't they just make it `bool`?

Well, there was a lot of C code out there where people had defined their own `bool` type and adding an official `bool` would have broken those `typedefs`.

But C has already reserved all identifiers that start with an underscore followed by a capital letter, so it was clear to make up a new `_Bool` type and go with that.

And, if you know your code can handle it, you can include this header to get all this juicy syntax.

One more note on conversions: unlike converting to `int`, the *only* thing that converts to `false` in a `_Bool` is a scalar zero value. Anything at all that's not zero, like `-3490`, `0.12`, or `NaN`, converts to `true`.

# Chapter 21

## <stddef.h> A Few Standard Definitions

Despite its name, I've haven't seen this frequently included.

It includes several types and macros.

| Name                     | Description                                                              |
|--------------------------|--------------------------------------------------------------------------|
| <code>ptrdiff_t</code>   | Signed integer difference between two pointers                           |
| <code>size_t</code>      | Unsigned integer type returned by <code>sizeof</code>                    |
| <code>max_align_t</code> | Declare a type with the biggest possible alignment                       |
| <code>wchar_t</code>     | Wide character type                                                      |
| <code>NULL</code>        | NULL pointer, as defined a number of places                              |
| <code>offsetof</code>    | Get the byte offsets of <code>struct</code> or <code>union</code> fields |

### 21.1 `ptrdiff_t`

This holds the different between two pointers. You could store this in another type, but the result of a pointer subtraction is an implementation-defined type; you can be maximally portable by using `ptrdiff_t`.

```
1  #include <stdio.h>
2  #include <stddef.h>
3
4  int main(void)
5  {
6      int cats[100];
7
8      int *f = cats + 20;
9      int *g = cats + 60;
10
11     ptrdiff_t d = g - f; // difference is 40
```

And you can print it by prefixing the integer format specifier with `t`:

```
13     printf("%td\n", d); // Print decimal: 40
14     printf("%tX\n", d); // Print hex:      28
15 }
```

## 21.2 `size_t`

This is the type returned by `sizeof` and used in a few other places. It's an unsigned integer.

You can print it using the `z` prefix in `printf()`:

```
1  #include <stdio.h>
2  #include <uchar.h>
3  #include <string.h>
4  #include <stddef.h>
5
6  int main(void)
7  {
8      size_t x;
9
10     x = sizeof(int);
11
12     printf("%zu\n", x);
```

Some functions return negative numbers cast to `size_t` as error values (such as `mbrtoc16()`). If you want to print these as negative values, you can do it with `%zd`:

```
14  char16_t a;
15  mbstate_t mbs;
16  memset(&mbs, 0, sizeof mbs);
17
18  x = mbrtoc16(&a, "b", 8, &mbs);
19
20  printf("%zd\n", x);
21 }
```

## 21.3 `max_align_t`

As far as I can tell, this exists to allow the runtime computation of the maximum fundamental alignment<sup>1</sup> on the current platform. Someone please mail me if there's another use.

Maybe you need this if you're writing your own memory allocator or somesuch.

```
1  #include <stddef.h>
2  #include <stdio.h>      // For printf()
3  #include <stdalign.h>   // For alignof
4
5  int main(void)
6  {
7      int max = alignof(max_align_t);
8
9      printf("Maximum fundamental alignment: %d\n", max);
10 }
```

On my system, this prints:

<sup>1</sup>[https://en.wikipedia.org/wiki/Data\\_structure\\_alignment](https://en.wikipedia.org/wiki/Data_structure_alignment)



```
Maximum fundamental alignment: 16
```

See also `alignas`, `alignof`.

## 21.4 `wchar_t`

This is analogous to `char`, except it's for wide characters.

It's an integer type that has enough range to hold unique values for all characters in all supported locales.

The value `0` is the wide NUL character.

Finally, the values of character constants from the basic character set will be the same as their corresponding `wchar_t` values... unless `__STDC_MB_MIGHT_NEQ_WC__` is defined.

## 21.5 `offsetof`

If you have a `struct` or `union`, you can use this to get the byte offset of fields within that type.

Usage is:

```
offsetof(type, fieldname);
```

The resulting value has type `size_t`.

Here's an example that prints the field offsets of a `struct`:

```
1  #include <stdio.h>
2  #include <stddef.h>
3
4  struct foo {
5      int a;
6      char b;
7      char c;
8      float d;
9  };
10
11 int main(void)
12 {
13     printf("a: %zu\n", offsetof(struct foo, a));
14     printf("b: %zu\n", offsetof(struct foo, b));
15     printf("c: %zu\n", offsetof(struct foo, c));
16     printf("d: %zu\n", offsetof(struct foo, d));
17 }
```

On my system, this outputs:

```
a: 0
b: 4
c: 5
d: 8
```

And you can't use `offsetof` on a bitfield, so don't get your hopes up.

## Chapter 22

# <stdint.h> More Integer Types

This header gives us access to (potentially) types of a fixed number of bits, or, at the very least, types that are at least that many bits.

It also gives us handy macros to use.

### 22.1 Specific-Width Integers

There are three main classes of types defined here, signed and unsigned:

- Integers of exactly a certain size (`intN_t`, `uintN_t`)
- Integers that are at least a certain size (`int_leastN_t`, `uint_leastN_t`)
- Integers that are at least a certain size and are as fast as possible (`int_fastN_t`, `uint_fastN_t`)

Where the *N* occurs, you substitute the number of bits, commonly multiples of 8, e.g. `uint16_t`.

The following types are guaranteed to be defined:

```
int_least8_t      uint_least8_t
int_least16_t     uint_least16_t
int_least32_t     uint_least32_t
int_least64_t     uint_least64_t

int_fast8_t       uint_fast8_t
int_fast16_t      uint_fast16_t
int_fast32_t      uint_fast32_t
int_fast64_t      uint_fast64_t
```

Everything else is optional, but you'll probably also have the following, which are required when a system has integers of these sizes with no padding and two's-complement representation... which is the case for Macs and PCs and a lot of other systems. In short, you very likely have these:

```
int8_t           uint8_t
int16_t          uint16_t
int32_t          uint32_t
int64_t          uint64_t
```

Other numbers of bits can also be supported by an implementation if it wants to go all crazy with it.

Examples:

```
#include <stdint.h>

int main(void)
{
    int16_t x = 32;
    int_fast32_t y = 3490;

    // ...
}
```

## 22.2 Other Integer Types

There are a couple optional types that are integers capable of holding pointer types.

```
intptr_t
uintptr_t
```

You can convert a `void*` to one of these types, and back again. And the `void*`s will compare equal.

The use case is any place you need an integer that represents a pointer for some reason.

Also, there are a couple types that are just there to be the biggest possible integers your system supports:

```
intmax_t
uintmax_t
```

Fun fact: you can print these types with the `"%jd"` and `"%ju"` `printf()` format specifiers.

There are also a bunch of macros in `<inttypes.h>`(#inttypes) that you can use to print any of the types mentioned, above.

## 22.3 Macros

The following macros define the minimum and maximum values for these types:

|                 |                 |                  |
|-----------------|-----------------|------------------|
| INT8_MAX        | INT8_MIN        | UINT8_MAX        |
| INT16_MAX       | INT16_MIN       | UINT16_MAX       |
| INT32_MAX       | INT32_MIN       | UINT32_MAX       |
| INT64_MAX       | INT64_MIN       | UINT64_MAX       |
| INT_LEAST8_MAX  | INT_LEAST8_MIN  | UINT_LEAST8_MAX  |
| INT_LEAST16_MAX | INT_LEAST16_MIN | UINT_LEAST16_MAX |
| INT_LEAST32_MAX | INT_LEAST32_MIN | UINT_LEAST32_MAX |
| INT_LEAST64_MAX | INT_LEAST64_MIN | UINT_LEAST64_MAX |
| INT_FAST8_MAX   | INT_FAST8_MIN   | UINT_FAST8_MAX   |
| INT_FAST16_MAX  | INT_FAST16_MIN  | UINT_FAST16_MAX  |
| INT_FAST32_MAX  | INT_FAST32_MIN  | UINT_FAST32_MAX  |
| INT_FAST64_MAX  | INT_FAST64_MIN  | UINT_FAST64_MAX  |
| INTMAX_MAX      | INTMAX_MIN      | UINTMAX_MAX      |
| INTPTR_MAX      | INTPTR_MIN      | UINTPTR_MAX      |

For the exact-bit-size signed types, the minimum is exactly  $-(2^{N-1})$  and the maximum is exactly  $2^{N-1} - 1$ . And for the exact-bit-size unsigned types, the max is exactly  $2^N - 1$ .

For the signed “least” and “fast” variants, the magnitude and sign of the minimum is at least  $-(2^{N-1} - 1)$  and the maximum is at least  $2^{N-1} - 1$ . And for unsigned it’s at least  $2^N - 1$ .

`INTMAX_MAX` is at least  $2^{63} - 1$ , `INTMAX_MIN` is at least  $-(2^{63} - 1)$  in sign and magnitude. And `UINTMAX_MAX` is at least  $2^{64} - 1$ .

Finally, `INTPTR_MAX` is at least  $2^{15} - 1$ , `INTPTR_MIN` is at least  $-(2^{15} - 1)$  in sign and magnitude. And `UINTPTR_MAX` is at least  $2^{16} - 1$ .

## 22.4 Other Limits

There are a bunch of types in <inttypes.h> (#inttypes) that have their limits defined here. (<inttypes.h> includes <stdint.h>.)

| Macro                       | Description                             |
|-----------------------------|-----------------------------------------|
| <code>PTRDIFF_MIN</code>    | Minimum <code>ptrdiff_t</code> value    |
| <code>PTRDIFF_MAX</code>    | Maximum <code>ptrdiff_t</code> value    |
| <code>SIG_ATOMIC_MIN</code> | Minimum <code>sig_atomic_t</code> value |
| <code>SIG_ATOMIC_MAX</code> | Maximum <code>sig_atomic_t</code> value |
| <code>SIZE_MAX</code>       | Maximum <code>size_t</code> value       |
| <code>WCHAR_MIN</code>      | Minimum <code>wchar_t</code> value      |
| <code>WCHAR_MAX</code>      | Maximum <code>wchar_t</code> value      |
| <code>WINT_MIN</code>       | Minimum <code>wint_t</code> value       |
| <code>WINT_MAX</code>       | Maximum <code>wint_t</code> value       |

The spec says that `PTRDIFF_MIN` will be at least -65535 in magnitude. And `PTRDIFF_MAX` and `SIZE_MAX` will be at least 65535.

`SIG_ATOMIC_MIN` and `MAX` will be either -127 and 127 (if it’s signed) or 0 and 255 (if it’s unsigned).

Same for `WCHAR_MIN` and `MAX`.

`WINT_MIN` and `MAX` will be either -32767 and 32767 (if it’s signed) or 0 and 65535 (if it’s unsigned).

## 22.5 Macros for Declaring Constants

If you recall, you can specify a type for integer constants:

```
int x = 12;
long int y = 12L;
unsigned long long int z = 12ULL;
```

You can use the macros `INTN_C()` and `UINTN()` where *N* is 8, 16, 32 or 64.

```
uint_least16_t x = INT16_C(3490);
uint_least64_t y = INT64_C(1122334455);
```

A variant on these is `INTMAX_C()` and `UINTMAX_C()`. They will make a constant suitable for storing in an `intmax_t` or `uintmax_t`.

```
intmax_t x = INTMAX_C(3490);  
uintmax_t x = UINTMAX_C(1122334455);
```

## Chapter 23

# <stdio.h> Standard I/O Library

| Function                | Description                                                      |
|-------------------------|------------------------------------------------------------------|
| <code>clearerr()</code> | Clear the <code>feof</code> and <code>ferror</code> status flags |
| <code>fclose()</code>   | Close an open file                                               |
| <code>feof()</code>     | Return the file end-of-file status                               |
| <code>ferror()</code>   | Return the file error status                                     |
| <code>fflush()</code>   | Flush all buffered output to a file                              |
| <code>fgetc()</code>    | Read a character in a file                                       |
| <code>fgetpos()</code>  | Get the file I/O position                                        |
| <code>fgets()</code>    | Read a line from a file                                          |
| <code>fopen()</code>    | Open a file                                                      |
| <code>fprintf()</code>  | Print formatted output to a file                                 |
| <code>fputc()</code>    | Print a character to a file                                      |
| <code>fputs()</code>    | Print a string to a file                                         |
| <code>fread()</code>    | Read binary data from a file                                     |
| <code>freopen()</code>  | Change file associated with a stream                             |
| <code>fscanf()</code>   | Read formatted input from a file                                 |
| <code>fseek()</code>    | Set the file I/O position                                        |
| <code>fsetpos()</code>  | Set the file I/O position                                        |
| <code>ftell()</code>    | Get the file I/O position                                        |
| <code>fwrite()</code>   | Write binary data to a file                                      |
| <code>getc()</code>     | Get a character from <code>stdin</code>                          |
| <code>getchar()</code>  | Get a character from <code>stdin</code>                          |
| <code>gets()</code>     | Get a string from <code>stdin</code> (removed in C11)            |
| <code>perror()</code>   | Print a human-formatted error message                            |
| <code>printf()</code>   | Print formatted output to <code>stdout</code>                    |
| <code>putc()</code>     | Print a character to <code>stdout</code>                         |
| <code>putchar()</code>  | Print a character to <code>stdout</code>                         |
| <code>puts()</code>     | Print a string to <code>stdout</code>                            |
| <code>remove()</code>   | Delete a file from disk                                          |
| <code>rename()</code>   | Rename or move a file on disk                                    |
| <code>rewind()</code>   | Set the I/O position to the beginning of a file                  |
| <code>scanf()</code>    | Read formatted input from <code>stdin</code>                     |
| <code>setbuf()</code>   | Configure buffering for I/O operations                           |
| <code>setvbuf()</code>  | Configure buffering for I/O operations                           |
| <code>snprintf()</code> | Print length-limited formatted output to a string                |
| <code>sprintf()</code>  | Print formatted output to a string                               |

| Function                | Description                                                |
|-------------------------|------------------------------------------------------------|
| <code>sscanf()</code>   | Read formatted input from a string                         |
| <code>tmpfile()</code>  | Create a temporary file                                    |
| <code>tmpnam()</code>   | Generate a unique name for a temporary file                |
| <code>ungetc()</code>   | Push a character back on the input stream                  |
| <code>vfprintf()</code> | Variadic print formatted output to a file                  |
| <code>vfscanf()</code>  | Variadic read formatted input from a file                  |
| <code>vprintf()</code>  | Variadic print formatted output to <code>stdout</code>     |
| <code>vscanf()</code>   | Variadic read formatted input from <code>stdin</code>      |
| <code>vsprintf()</code> | Variadic length-limited print formatted output to a string |
| <code>vsprintf()</code> | Variadic print formatted output to a string                |
| <code>vsscanf()</code>  | Variadic read formatted input to a string                  |

The most basic of all libraries in the whole of the standard C library is the standard I/O library. It's used for reading from and writing to files. I can see you're very excited about this.

So I'll continue. It's also used for reading and writing to the console, as we've already often seen with the `printf()` function.

(A little secret here—many many things in various operating systems are secretly files deep down, and the console is no exception. “*Everything in Unix is a file!*” :-))

You'll probably want some prototypes of the functions you can use, right? To get your grubby little mittens on those, you'll want to include `stdio.h`.

Anyway, so we can do all kinds of cool stuff in terms of file I/O. LIE DETECTED. Ok, ok. We can do all kinds of stuff in terms of file I/O. Basically, the strategy is this:

1. Use `fopen()` to get a pointer to a file structure of type `FILE*`. This pointer is what you'll be passing to many of the other file I/O calls.
2. Use some of the other file calls, like `fscanf()`, `fgets()`, `fprintf()`, or etc. using the `FILE*` returned from `fopen()`.
3. When done, call `fclose()` with the `FILE*`. This let's the operating system know that you're truly done with the file, no take-backs.

What's in the `FILE*`? Well, as you might guess, it points to a `struct` that contains all kinds of information about the current read and write position in the file, how the file was opened, and other stuff like that. But, honestly, who cares. No one, that's who. The `FILE` structure is *opaque* to you as a programmer; that is, you don't need to know what's in it, and you don't even *want* to know what's in it. You just pass it to the other standard I/O functions and they know what to do.

This is actually pretty important: try to not muck around in the `FILE` structure. It's not even the same from system to system, and you'll end up writing some really non-portable code.

One more thing to mention about the standard I/O library: a lot of the functions that operate on files use an “f” prefix on the function name. The same function that is operating on the console will leave the “f” off. For instance, if you want to print to the console, you use `printf()`, but if you want to print to a file, use `fprintf()`, see?

Wait a moment! If writing to the console is, deep down, just like writing to a file, since everything in Unix is a file, why are there two functions? Answer: it's more convenient. But, more importantly, is there a `FILE*` associated with the console that you can use? Answer: YES!

There are, in fact, *three* (count 'em!) special `FILE*`s you have at your disposal merely for just including `stdio.h`. There is one for input, and two for output.

That hardly seems fair—why does output get two files, and input only get one?

That’s jumping the gun a bit—let’s just look at them:

| Stream              | Description                                     |
|---------------------|-------------------------------------------------|
| <code>stdin</code>  | Input from the console.                         |
| <code>stdout</code> | Output to the console.                          |
| <code>stderr</code> | Output to the console on the error file stream. |

So standard input (`stdin`) is by default just what you type at the keyboard. You can use that in `fscanf()` if you want, just like this:

```
/* this line: */
scanf("%d", &x);

/* is just like this line: */
fscanf(stdin, "%d", &x);
```

And `stdout` works the same way:

```
printf("Hello, world!\n");
fprintf(stdout, "Hello, world!\n"); /* same as previous line! */
```

So what is this `stderr` thing? What happens when you output to that? Well, generally it goes to the console just like `stdout`, but people use it for error messages, specifically. Why? On many systems you can redirect the output from the program into a file from the command line...and sometimes you’re interested in getting just the error output. So if the program is good and writes all its errors to `stderr`, a user can redirect just `stderr` into a file, and just see that. It’s just a nice thing you, as a programmer, can do.

Finally, a lot of these functions return `int` where you might expect `char`. This is because the function can return a character or end-of-file (EOF), and EOF is potentially an integer. If you don’t get EOF as a return value, you can safely store the result in a `char`.

## 23.1 `remove()`

Delete a file

### Synopsis

```
#include <stdio.h>

int remove(const char *filename);
```

### Description

Removes the specified file from the filesystem. It just deletes it. Nothing magical. Simply call this function and sacrifice a small chicken and the requested file will be deleted.



## Return Value

Returns zero on success, and `-1` on error, setting `errno`.

## Example

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     char *filename = "evidence.txt";
6
7     remove(filename);
8 }
```

## See Also

`rename()`

---

## 23.2 `rename()`

Renames a file and optionally moves it to a new location

## Synopsis

```
#include <stdio.h>

int rename(const char *old, const char *new);
```

## Description

Renames the file `old` to name `new`. Use this function if you're tired of the old name of the file, and you are ready for a change. Sometimes simply renaming your files makes them feel new again, and could save you money over just getting all new files!

One other cool thing you can do with this function is actually move a file from one directory to another by specifying a different path for the new name.

## Return Value

Returns zero on success, and `-1` on error, setting `errno`.

## Example

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     // Rename a file
```

```
6     rename("foo", "bar");
7
8     // Rename and move to another directory:
9     rename("/home/beej/evidence.txt", "/tmp/nothing.txt");
10 }
```

## See Also

`remove()`

---

## 23.3 `tmpfile()`

Create a temporary file

### Synopsis

```
#include <stdio.h>

FILE *tmpfile(void);
```

### Description

This is a nifty little function that will create and open a temporary file for you, and will return a `FILE*` to it that you can use. The file is opened with mode “`r+b`”, so it’s suitable for reading, writing, and binary data.

By using a little magic, the temp file is automatically deleted when it is `close()`’d or when your program exits. (Specifically, in Unix terms, `tmpfile()` *unlinks*<sup>1</sup> the file right after it opens it. This means that it’s primed to be deleted from disk, but still exists because your process still has it open. As soon as your process exits, all open files are closed, and the temp file vanishes into the ether.)

### Return Value

This function returns an open `FILE*` on success, or `NULL` on failure.

### Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *temp;
6      char s[128];
7
8      temp = tmpfile();
9
10     fprintf(temp, "What is the frequency, Alexander?\n");
11 }
```

---

<sup>1</sup><https://man.archlinux.org/man/unlinkat.2.en#DESCRIPTION>

```
12     rewind(temp); // back to the beginning
13
14     fscanf(temp, "%s", s); // read it back out
15
16     fclose(temp); // close (and magically delete)
17 }
```

## See Also

`fopen()`, `fclose()`, `tmpnam()`

---

## 23.4 `tmpnam()`

Generate a unique name for a temporary file

### Synopsis

```
#include <stdio.h>

char *tmpnam(char *s);
```

### Description

This function takes a good hard look at the existing files on your system, and comes up with a unique name for a new file that is suitable for temporary file usage.

Let's say you have a program that needs to store off some data for a short time so you create a temporary file for the data, to be deleted when the program is done running. Now imagine that you called this file `foo.txt`. This is all well and good, except what if a user already has a file called `foo.txt` in the directory that you ran your program from? You'd overwrite their file, and they'd be unhappy and stalk you forever. And you wouldn't want that, now would you?

Ok, so you get wise, and you decide to put the file in `/tmp` so that it won't overwrite any important content. But wait! What if some other user is running your program at the same time and they both want to use that filename? Or what if some other program has already created that file?

See, all of these scary problems can be completely avoided if you just use `tmpnam()` to get a safe-ready-to-use filename.

So how do you use it? There are two amazing ways. One, you can declare an array (or `malloc()` it—whatever) that is big enough to hold the temporary file name. How big is that? Fortunately there has been a macro defined for you, `L_tmpnam`, which is how big the array must be.

And the second way: just pass `NULL` for the filename. `tmpnam()` will store the temporary name in a static array and return a pointer to that. Subsequent calls with a `NULL` argument will overwrite the static array, so be sure you're done using it before you call `tmpnam()` again.

Again, this function just makes a file name for you. It's up to you to later `fopen()` the file and use it.

One more note: some compilers warn against using `tmpnam()` since some systems have better functions (like the Unix function `mkstemp()`.) You might want to check your local documentation to see if there's a better option. Linux documentation goes so far as to say, "Never use this function. Use `mkstemp()` instead."

I, however, am going to be a jerk and not talk about `mkstemp()`<sup>2</sup> because it's not in the standard I'm writing about. Nyaah.

The macro `TMP_MAX` holds the number of unique filenames that can be generated by `tmpnam()`. Ironically, it is the *minimum* number of such filenames.

## Return Value

Returns a pointer to the temporary file name. This is either a pointer to the string you passed in, or a pointer to internal static storage if you passed in `NULL`. On error (like it can't find any temporary name that is unique), `tmpnam()` returns `NULL`.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      char filename[L_tmpnam];
6      char *another_filename;
7
8      if (tmpnam(filename) != NULL)
9          printf("We got a temp file name: \"%s\\n\"", filename);
10     else
11         printf("Something went wrong, and we got nothing!\\n");
12
13     another_filename = tmpnam(NULL);
14
15     printf("We got another temp file name: \"%s\\n\"", another_filename);
16     printf("And we didn't error check it because we're too lazy!\\n");
17 }
```

On my Linux system, this generates the following output:

```
We got a temp file name: "/tmp/filew9PMuZ"
We got another temp file name: "/tmp/file0wrgP0"
And we didn't error check it because we're too lazy!
```

## See Also

`fopen()`, `tmpfile()`

---

## 23.5 `fclose()`

The opposite of `fopen()`—closes a file when you're done with it so that it frees system resources

---

<sup>2</sup><https://man.archlinux.org/man/mkstemp.3.en>

## Synopsis

```
#include <stdio.h>

int fclose(FILE *stream);
```

## Description

When you open a file, the system sets aside some resources to maintain information about that open file. Usually it can only open so many files at once. In any case, the Right Thing to do is to close your files when you're done using them so that the system resources are freed.

Also, you might not find that all the information that you've written to the file has actually been written to disk until the file is closed. (You can force this with a call to `fflush()`.)

When your program exits normally, it closes all open files for you. Lots of times, though, you'll have a long-running program, and it'd be better to close the files before then. In any case, not closing a file you've opened makes you look bad. So, remember to `fclose()` your file when you're done with it!

## Return Value

On success, `0` is returned. Typically no one checks for this. On error `EOF` is returned. Typically no one checks for this, either.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *fp;
6
7      fp = fopen("spoon.txt", "r");
8
9      if (fp == NULL) {
10         printf("Error opening file\n");
11     } else {
12         printf("Opened file just fine!\n");
13         fclose(fp); // All done!
14     }
15 }
```

## See Also

`fopen()`

---

## 23.6 `fflush()`

Process all buffered I/O for a stream right now

## Synopsis

```
#include <stdio.h>

int fflush(FILE *stream);
```

## Description

When you do standard I/O, as mentioned in the section on the `setvbuf()` function, it is usually stored in a buffer until a line has been entered or the buffer is full or the file is closed. Sometimes, though, you really want the output to happen *right this second*, and not wait around in the buffer. You can force this to happen by calling `fflush()`.

The advantage to buffering is that the OS doesn't need to hit the disk every time you call `fprintf()`. The disadvantage is that if you look at the file on the disk after the `fprintf()` call, it might not have actually been written to yet. ("I called `fputs()`, but the file is still zero bytes long! Why?!" ) In virtually all circumstances, the advantages of buffering outweigh the disadvantages; for those other circumstances, however, use `fflush()`.

Note that `fflush()` is only designed to work on output streams according to the spec. What will happen if you try it on an input stream? Use your spooky voice: *who knooooows!*

## Return Value

On success, `fflush()` returns zero. If there's an error, it returns `EOF` and sets the error condition for the stream (see `ferror()`.)

## Example

In this example, we're going to use the carriage return, which is `'\r'`. This is like newline (`'\n'`), except that it doesn't move to the next line. It just returns to the front of the current line.

What we're going to do is a little text-based status bar like so many command line programs implement. It'll do a countdown from 10 to 0 printing over itself on the same line.

What is the catch and what does this have to do with `fflush()`? The catch is that the terminal is most likely "line buffered" (see the section on `setvbuf()` for more info), meaning that it won't actually display anything until it prints a newline. But we're not printing newlines; we're just printing carriage returns, so we need a way to force the output to occur even though we're on the same line. Yes, it's `fflush()`!

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  void sleep_seconds(int s)
5  {
6      thrd_sleep(&(struct timespec){.tv_sec=s}, NULL);
7  }
8
9  int main(void)
10 {
11     int count;
12
13     for(count = 10; count >= 0; count--) {
14         printf("\rSeconds until launch: "); // lead with a CR
```

```

15     if (count > 0)
16         printf("%2d", count);
17     else
18         printf("blastoff!\n");
19
20     // force output now!!
21     fflush(stdout);
22
23     sleep_seconds(1);
24 }
25 }

```

**See Also**

setbuf(), setvbuf()

**23.7 fopen()**

Opens a file for reading or writing

**Synopsis**

```

#include <stdio.h>

FILE *fopen(const char *path, const char *mode);

```

**Description**

The `fopen()` opens a file for reading or writing.

Parameter `path` can be a relative or fully-qualified path and file name to the file in question.

Parameter `mode` tells `fopen()` how to open the file (reading, writing, or both), and whether or not it's a binary file. Possible modes are:

| Mode            | Description                                                                                                                                                                                                      |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>r</code>  | Open the file for reading (read-only).                                                                                                                                                                           |
| <code>w</code>  | Open the file for writing (write-only). The file is created if it doesn't exist.                                                                                                                                 |
| <code>r+</code> | Open the file for reading and writing. The file has to already exist.                                                                                                                                            |
| <code>w+</code> | Open the file for writing and reading. The file is created if it doesn't already exist.                                                                                                                          |
| <code>a</code>  | Open the file for append. This is just like opening a file for writing, but it positions the file pointer at the end of the file, so the next write appends to the end. The file is created if it doesn't exist. |
| <code>a+</code> | Open the file for reading and appending. The file is created if it doesn't exist.                                                                                                                                |

Any of the modes can have the letter “b” appended to the end, as is “wb” (“write binary”), to signify that the file in question is a *binary* file. (“Binary” in this case generally means that the file contains non-alphanumeric characters that look like garbage to human eyes.) Many systems (like Unix) don’t differentiate between binary and non-binary files, so the “b” is extraneous. But if your data is binary, it doesn’t hurt to throw the “b” in there, and it might help someone who is trying to port your code to another system.

The macro `FOPEN_MAX` tells you how many streams (at least) you can have open at once.

The macro `FILENAME_MAX` tells you what the longest valid filename can be. Don’t go crazy, now.

## Return Value

`fopen()` returns a `FILE*` that can be used in subsequent file-related calls.

If something goes wrong (e.g. you tried to open a file for read that didn’t exist), `fopen()` will return `NULL`.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *fp;
6
7      fp = fopen("spoon.txt", "r");
8
9      if (fp == NULL) {
10         printf("Error opening file\n");
11     } else {
12         printf("Opened file just fine!\n");
13         fclose(fp); // All done!
14     }
15 }
```

## See Also

`fclose()`, `freopen()`

---

## 23.8 `freopen()`

Reopen an existing `FILE*`, associating it with a new path

### Synopsis

```
#include <stdio.h>

FILE *freopen(const char *filename, const char *mode, FILE *stream);
```



## Description

Let's say you have an existing `FILE*` stream that's already open, but you want it to suddenly use a different file than the one it's using. You can use `freopen()` to "re-open" the stream with a new file.

Why on Earth would you ever want to do that? Well, the most common reason would be if you had a program that normally would read from `stdin`, but instead you wanted it to read from a file. Instead of changing all your `scanf()`s to `fscanf()`s, you could simply reopen `stdin` on the file you wanted to read from.

Another usage that is allowed on some systems is that you can pass `NULL` for `filename`, and specify a new mode for `stream`. So you could change a file from "r+" (read and write) to just "r" (read), for instance. It's implementation dependent which modes can be changed.

When you call `freopen()`, the old `stream` is closed. Otherwise, the function behaves just like the standard `fopen()`.

## Return Value

`freopen()` returns `stream` if all goes well.

If something goes wrong (e.g. you tried to open a file for read that didn't exist), `freopen()` will return `NULL`.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      int i, i2;
6
7      scanf("%d", &i); // read i from stdin
8
9      // now change stdin to refer to a file instead of the keyboard
10     freopen("someints.txt", "r", stdin);
11
12     scanf("%d", &i2); // now this reads from the file "someints.txt"
13
14     printf("Hello, world!\n"); // print to the screen
15
16     // change stdout to go to a file instead of the terminal:
17     freopen("output.txt", "w", stdout);
18
19     printf("This goes to the file \"output.txt\"\n");
20
21     // this is allowed on some systems--you can change the mode of a file:
22     freopen(NULL, "wb", stdout); // change to "wb" instead of "w"
23 }
```

## See Also

`fclose()`, `fopen()`

---

## 23.9 `setbuf()`, `setvbuf()`

Configure buffering for standard I/O operations

### Synopsis

```
#include <stdio.h>

void setbuf(FILE *stream, char *buf);

int setvbuf(FILE *stream, char *buf, int mode, size_t size);
```

### Description

Now brace yourself because this might come as a bit of a surprise to you: when you `printf()` or `fprintf()` or use any I/O functions like that, *it does not normally work immediately*. For the sake of efficiency, and to irritate you, the I/O on a `FILE*` stream is buffered away safely until certain conditions are met, and only then is the actual I/O performed. The functions `setbuf()` and `setvbuf()` allow you to change those conditions and the buffering behavior.

So what are the different buffering behaviors? The biggest is called “full buffering”, wherein all I/O is stored in a big buffer until it is full, and then it is dumped out to disk (or whatever the file is). The next biggest is called “line buffering”; with line buffering, I/O is stored up a line at a time (until a newline (`'\n'`) character is encountered) and then that line is processed. Finally, we have “unbuffered”, which means I/O is processed immediately with every standard I/O call.

You might have seen and wondered why you could call `putchar()` time and time again and not see any output until you called `putchar('\n')`; that’s right—`stdout` is line-buffered!

Since `setbuf()` is just a simplified version of `setvbuf()`, we’ll talk about `setvbuf()` first.

The `stream` is the `FILE*` you wish to modify. The standard says you *must* make your call to `setvbuf()` *before* any I/O operation is performed on the stream, or else by then it might be too late.

The next argument, `buf` allows you to make your own buffer space (using `malloc()` or just a `char` array) to use for buffering. If you don’t care to do this, just set `buf` to `NULL`.

Now we get to the real meat of the function: `mode` allows you to choose what kind of buffering you want to use on this `stream`. Set it to one of the following:

| Mode                | Description                                 |
|---------------------|---------------------------------------------|
| <code>_IOFBF</code> | <code>stream</code> will be fully buffered. |
| <code>_IOLBF</code> | <code>stream</code> will be line buffered.  |
| <code>_IONBF</code> | <code>stream</code> will be unbuffered.     |

Finally, the `size` argument is the size of the array you passed in for `buf`...unless you passed `NULL` for `buf`, in which case it will resize the existing buffer to the size you specify.

Now what about this lesser function `setbuf()`? It’s just like calling `setvbuf()` with some specific parameters, except `setbuf()` doesn’t return a value. The following example shows the equivalency:

```
// these are the same:
setbuf(stream, buf);
setvbuf(stream, buf, _IOFBF, BUFSIZ); // fully buffered
```

```
// and these are the same:
setbuf(stream, NULL);
setvbuf(stream, NULL, _IONBF, BUFSIZ); // unbuffered
```

## Return Value

`setvbuf()` returns zero on success, and nonzero on failure. `setbuf()` has no return value.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *fp;
6      char lineBuf[1024];
7
8      fp = fopen("somefile.txt", "w");
9      setvbuf(fp, lineBuf, _IOLBF, 1024); // set to line buffering
10     fprintf(fp, "You won't see this in the file yet. ");
11     fprintf(fp, "But now you will because of this newline.\n");
12     fclose(fp);
13
14     fp = fopen("anotherfile.txt", "w");
15     setbuf(fp, NULL); // set to unbuffered
16     fprintf(fp, "You will see this in the file now.");
17     fclose(fp);
18 }
```

## See Also

`fflush()`

## 23.10 `printf()`, `fprintf()`, `sprintf()`, `snprintf()`

Print a formatted string to the console or to a file

### Synopsis

```
#include <stdio.h>

int printf(const char *format, ...);

int fprintf(FILE *stream, const char *format, ...);

int sprintf(char * restrict s, const char * restrict format, ...);

int snprintf(char * restrict s, size_t n, const char * restrict format, ...);
```

## Description

These functions print formatted output to a variety of destinations.

| Function                | Output Destination                               |
|-------------------------|--------------------------------------------------|
| <code>printf()</code>   | Print to console (screen by default, typically). |
| <code>fprintf()</code>  | Print to a file.                                 |
| <code>sprintf()</code>  | Print to a string.                               |
| <code>snprintf()</code> | Print to a string (safely).                      |

The only differences between these is are the leading parameters that you pass to them before the `format` string.

| Function                | What you pass before <code>format</code>                             |
|-------------------------|----------------------------------------------------------------------|
| <code>printf()</code>   | Nothing comes before <code>format</code> .                           |
| <code>fprintf()</code>  | Pass a <code>FILE*</code> .                                          |
| <code>sprintf()</code>  | Pass a <code>char*</code> to a buffer to print into.                 |
| <code>snprintf()</code> | Pass a <code>char*</code> to the buffer and a maximum buffer length. |

The `printf()` function is legendary as being one of the most flexible outputting systems ever devised. It can also get a bit freaky here or there, most notably in the `format` string. We'll take it a step at a time here.

The easiest way to look at the format string is that it will print everything in the string as-is, *unless* a character has a percent sign (%) in front of it. That's when the magic happens: the next argument in the `printf()` argument list is printed in the way described by the percent code. These percent codes are called *format specifiers*.

Here are the most common format specifiers.

| Specifier       | Description                                                                                                                                                                                         |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>%d</code> | Print the next argument as a signed decimal number, like 3490. The argument printed this way should be an <code>int</code> , or something that gets promoted to <code>int</code> .                  |
| <code>%f</code> | Print the next argument as a signed floating point number, like 3.14159. The argument printed this way should be a <code>double</code> , or something that gets promoted to a <code>double</code> . |
| <code>%c</code> | Print the next argument as a character, like 'B'. The argument printed this way should be a <code>char</code> variant.                                                                              |
| <code>%s</code> | Print the next argument as a string, like "Did you remember your mittens?". The argument printed this way should be a <code>char*</code> or <code>char[]</code> .                                   |
| <code>%%</code> | No arguments are converted, and a plain old run-of-the-mill percent sign is printed. This is how you print a '%' using <code>printf()</code> .                                                      |

So those are the basics. I'll give you some more of the format specifiers in a bit, but let's get some more breadth before then. There's actually a lot more that you can specify in there after the percent sign.

For one thing, you can put a field width in there—this is a number that tells `printf()` how many spaces to put on one side or the other of the value you're printing. That helps you line things up in nice columns. If the number is negative, the result becomes left-justified instead of right-justified. Example:

```
printf("%10d", x); /* prints X on the right side of the 10-space field */
printf("%-10d", x); /* prints X on the left side of the 10-space field */
```

If you don't know the field width in advance, you can use a little kung-foo to get it from the argument list just before the argument itself. Do this by placing your seat and tray tables in the fully upright position. The seatbelt is fastened by placing the—cough. I seem to have been doing way too much flying lately. Ignoring that useless fact completely, you can specify a dynamic field width by putting a `*` in for the width. If you are not willing or able to perform this task, please notify a flight attendant and we will reseat you.

```
int width = 12;
int value = 3490;

printf("%*d\n", width, value);
```

You can also put a "0" in front of the number if you want it to be padded with zeros:

```
int x = 17;
printf("%05d", x); /* "00017" */
```

When it comes to floating point, you can also specify how many decimal places to print by making a field width of the form "`x.y`" where `x` is the field width (you can leave this off if you want it to be just wide enough) and `y` is the number of digits past the decimal point to print:

```
float f = 3.1415926535;

printf("%.2f", f); /* "3.14" */
printf("%7.3f", f); /* " 3.141" <-- 7 spaces across */
```

Ok, those above are definitely the most common uses of `printf()`, but let's get *total coverage*.

### 23.10.0.1 Format Specifier Layout

Technically, the layout of the format specifier is these things in this order:

1. `%`, followed by...
2. Optional: zero or more flags, left justify, leading zeros, etc.
3. Optional: Field width, how wide the output field should be.
4. Optional: Precision, or how many decimal places to print.
5. Optional: Length modifier, for printing things bigger than `int` or `double`.
6. Conversion specifier, like `d`, `f`, etc.

In short, the whole format specifier is laid out like this:

```
%[flags][fieldwidth][.precision][lengthmodifier]conversionspecifier
```

What could be easier?

### 23.10.0.2 Conversion Specifiers

Let's talk conversion specifiers first. Each of the following specifies what type it can print, but it can also print anything that gets promoted to that type. For example, `%d` can print `int`, `short`, and `char`.

The binary specifier is new in C23!

| Conversion Specifier | Description                                             |
|----------------------|---------------------------------------------------------|
| <code>d</code>       | Print an <code>int</code> argument as a decimal number. |

| Conversion Specifier | Description                                                                                                                                                                                                         |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>i</code>       | Identical to <code>d</code> .                                                                                                                                                                                       |
| <code>b</code>       | Print an unsigned int in binary (base 2).                                                                                                                                                                           |
| <code>B</code>       | Identical to <code>b</code> , except for the alternate form, below.                                                                                                                                                 |
| <code>o</code>       | Print an unsigned int in octal (base 8).                                                                                                                                                                            |
| <code>u</code>       | Print an unsigned int in decimal.                                                                                                                                                                                   |
| <code>x</code>       | Print an unsigned int in hexadecimal with lowercase letters.                                                                                                                                                        |
| <code>X</code>       | Print an unsigned int in hexadecimal with uppercase letters; also note the alternate form, below.                                                                                                                   |
| <code>f</code>       | Print a double in decimal notation. Infinity is printed as <code>infinity</code> or <code>inf</code> , and NaN is printed as <code>nan</code> , any of which could have a leading minus sign.                       |
| <code>F</code>       | Same as <code>f</code> , except it prints out <code>INFINITY</code> , <code>INF</code> , or <code>NAN</code> in all caps.                                                                                           |
| <code>e</code>       | Print a number in scientific notation, e.g. <code>1.234e56</code> . Does infinity and NaN like <code>f</code> .                                                                                                     |
| <code>E</code>       | Just like <code>e</code> , except prints the exponent <code>E</code> (and infinity and NaN) in uppercase.                                                                                                           |
| <code>g</code>       | Print small numbers like <code>f</code> and large numbers like <code>e</code> . See note below.                                                                                                                     |
| <code>G</code>       | Print small numbers like <code>F</code> and large numbers like <code>E</code> . See note below.                                                                                                                     |
| <code>a</code>       | Print a double in hexadecimal form <code>0xh.hhhhp</code> where <code>h</code> is a lowercase hex digit and <code>d</code> is a decimal exponent of 2. Infinity and NaN in the form of <code>f</code> . More below. |
| <code>A</code>       | Like <code>a</code> except everything's uppercase.                                                                                                                                                                  |
| <code>c</code>       | Convert int argument to unsigned char and print as a character.                                                                                                                                                     |
| <code>s</code>       | Print a string starting at the given char*.                                                                                                                                                                         |
| <code>p</code>       | Print a void* out as a number, probably the numeric address, possibly in hex.                                                                                                                                       |
| <code>n</code>       | Store the number of characters written so far in the given int*. Doesn't print anything. See below.                                                                                                                 |
| <code>%</code>       | Print a literal percent sign.                                                                                                                                                                                       |

**23.10.0.2.1 Note on %a and %A** When printing floating point numbers in hex form, there is one number before the decimal point, and the rest of are out to the precision.

```
double pi = 3.14159265358979;

printf("%.3a\n", pi); // 0x1.922p+1
```

C can choose the leading number in such a way to ensure subsequent digits align to 4-bit boundaries.

If the precision is left out and the macro `FLT_RADIX` is a power of 2, enough precision is used to represent the number exactly. If `FLT_RADIX` is not a power of two, enough precision is used to be able to tell any two floating values apart.

If the precision is 0 and the `#` flag isn't specified, the decimal point is omitted.

**23.10.0.2.2 Note on %g and %G** The gist of this is to use scientific notation when the number gets too "extreme", and regular decimal notation otherwise.

The exact behavior for whether these print as `%f` or `%e` depends on a number of factors:

If the number's exponent is greater than or equal to -4 and the precision is greater than the exponent, we use `%f`. In this case, the precision is converted according to  $p = p - (x + 1)$ , where  $p$  is the specified precision and  $x$  is the exponent.

Otherwise we use `%e`, and the precision becomes  $p - 1$ .

Trailing zeros in the decimal portion are removed. And if there are none left, the decimal point is removed, too. All this unless the `#` flag is specified.

**23.10.0.2.3 Note on %n** This specifier is cool and different, and rarely needed. It doesn't actually print anything, but stores the number of characters printed so far in the next pointer argument in the list.

```
int numChars;
float a = 3.14159;
int b = 3490;

printf("%f %d%n\n", a, b, &numChars);
printf("The above line contains %d characters.\n", numChars);
```

The above example will print out the values of `a` and `b`, and then store the number of characters printed so far into the variable `numChars`. The next call to `printf()` prints out that result.

```
3.141590 3490
The above line contains 13 characters
```

### 23.10.0.3 Length Modifiers

You can stick a *length* modifier in front of each of the conversion specifiers, if you want. most of those format specifiers work on `int` or `double` types, but what if you want larger or smaller types? That's what these are good for.

For example, you could print out a long long int with the `ll` modifier:

```
long long int x = 3490;

printf("%lld\n", x); // 3490
```

| Length Modifier | Conversion Specifier                | Description                                                                                     |
|-----------------|-------------------------------------|-------------------------------------------------------------------------------------------------|
| <code>hh</code> | <code>b, d, i, o, u, x, X</code>    | Convert argument to <code>char</code> (signed or unsigned as appropriate) before printing.      |
| <code>h</code>  | <code>b, d, i, o, u, x, X</code>    | Convert argument to <code>short int</code> (signed or unsigned as appropriate) before printing. |
| <code>l</code>  | <code>b, d, i, o, u, x, X</code>    | Argument is a <code>long int</code> (signed or unsigned as appropriate).                        |
| <code>ll</code> | <code>b, d, i, o, u, x, X</code>    | Argument is a <code>long long int</code> (signed or unsigned as appropriate).                   |
| <code>j</code>  | <code>b, d, i, o, u, x, X</code>    | Argument is a <code>intmax_t</code> or <code>uintmax_t</code> (as appropriate).                 |
| <code>z</code>  | <code>b, d, i, o, u, x, X</code>    | Argument is a <code>size_t</code> .                                                             |
| <code>t</code>  | <code>b, d, i, o, u, x, X</code>    | Argument is a <code>ptrdiff_t</code> .                                                          |
| <code>L</code>  | <code>a, A, e, E, f, F, g, G</code> | Argument is a <code>long double</code> .                                                        |
| <code>l</code>  | <code>c</code>                      | Argument is in a <code>wint_t</code> , a wide character.                                        |
| <code>l</code>  | <code>s</code>                      | Argument is in a <code>wchar_t*</code> , a wide character string.                               |
| <code>hh</code> | <code>n</code>                      | Store result in <code>signed char*</code> argument.                                             |
| <code>h</code>  | <code>n</code>                      | Store result in <code>short int*</code> argument.                                               |
| <code>l</code>  | <code>n</code>                      | Store result in <code>long int*</code> argument.                                                |
| <code>ll</code> | <code>n</code>                      | Store result in <code>long long int*</code> argument.                                           |
| <code>j</code>  | <code>n</code>                      | Store result in <code>intmax_t*</code> argument.                                                |
| <code>z</code>  | <code>n</code>                      | Store result in <code>size_t*</code> argument.                                                  |
| <code>t</code>  | <code>n</code>                      | Store result in <code>ptrdiff_t*</code> argument.                                               |

### 23.10.0.4 Precision

In front of the length modifier, you can put a precision, which generally means how many decimal places you want on your floating point numbers.

To do this, you put a decimal point (.) and the decimal places afterward.

For example, we could print  $\pi$  rounded to two decimal places like this:

```
double pi = 3.14159265358979;

printf("%.2f\n", pi); // 3.14
```

| Conversion Specifier     | Precision Value Meaning                                                                                                                                       |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| b, d, i, o, u, x, X      | For integer types, minimum number of digits (will pad with leading zeros)                                                                                     |
| a, e, f, A, E, F<br>g, G | For floating types, the precision is the number of digits past the decimal.<br>For floating types, the precision is the number of significant digits printed. |
| s                        | The maximum number of bytes (not multibyte characters!) to be written.                                                                                        |

If no number is specified in the precision after the decimal point, the precision is zero.

If an `*` is specified after the decimal, something amazing happens! It means the `int` argument to `printf()` before the number to be printed holds the precision. You can use this if you don't know the precision at compile time.

```
int precision;
double pi = 3.14159265358979;

printf("Enter precision: "); fflush(stdout);
scanf("%d", &precision);

printf("%.*f\n", precision, pi);
```

Which gives:

```
Enter precision: 4
3.1416
```

### 23.10.0.5 Field Width

In front of the optional precision, you can indicate a field width. This is a decimal number that indicates how wide the region should be in which the argument is printed. The region is padding with leading (or trailing) spaces to make sure it's wide enough.

If the field width specified is too small to hold the output, it is ignored.

As a preview, you can give a negative field width to justify the item the other direction.

So let's print a number in a field of width 10. We'll put some angle brackets around it so we can see the padding spaces in the output.



```
printf("<<%10d>>\n", 3490); // right justified
printf("<<%-10d>>\n", 3490); // left justified
```

```
<<      3490>>
<<3490      >>
```

Like with the precision, you can use an asterisk (\*) as the field width

```
int field_width;
int val = 3490;

printf("Enter field_width: "); fflush(stdout);
scanf("%d", &field_width);

printf("<<.*d>>\n", field_width, val);
```

### 23.10.0.6 Flags

Before the field width, you can put some optional flags that further control the output of the subsequent fields. We just saw that the - flag can be used to left- or right-justify fields. But there are plenty more!

| Flag    | Description                                                                 |
|---------|-----------------------------------------------------------------------------|
| -       | For a field width, left justify in the field (right is default).            |
| +       | If the number is signed, always prefix a + or - on the front.               |
| [SPACE] | If the number is signed, prefix a space for positive, or a - for negative.  |
| 0       | Pad the right-justified field with leading zeros instead of leading spaces. |
| #       | Print using an alternate form. See below.                                   |

For example, we could pad a hexadecimal number with leading zeros to a field width of 8 with:

```
printf("%08x\n", 0x1234); // 00001234
```

The # “alternate form” result depends on the conversion specifier.

| Conversion Specifier | Alternate Form (#) Meaning                                                                    |
|----------------------|-----------------------------------------------------------------------------------------------|
| o                    | Increase precision of a non-zero number just enough to get one leading 0 on the octal number. |
| b                    | Prefix a non-zero number with 0b.                                                             |
| B                    | Same as b, except capital 0B.                                                                 |
| x                    | Prefix a non-zero number with 0x.                                                             |
| X                    | Same as x, except capital 0X.                                                                 |
| a, e, f              | Always print a decimal point, even if nothing follows it.                                     |
| A, E, F              | Identical to a, e, f.                                                                         |
| g, G                 | Always print a decimal point, even if nothing follows it, and keep trailing zeros.            |

### 23.10.0.7 sprintf() and snprintf() Details

Both sprintf() and snprintf() have the quality that if you pass in NULL as the buffer, nothing is written—but you can still check the return value to see how many characters *would* have been written.

`snprintf()` **always** terminates the string with a `NUL` character. So if you try to write out more than the maximum specified characters, the universe ends.

Just kidding. If you do, `snprintf()` will write  $n - 1$  characters so that it has enough room to write the terminator at the end.

## Return Value

Returns the number of characters outputted, or a negative number on error.

## Example

```

1  #include <stdio.h>
2
3  int main(void)
4  {
5      int a = 100;
6      float b = 2.717;
7      char *c = "beej!";
8      char d = 'X';
9      int e = 5;
10
11     printf("%d\n", a); /* "100" */
12     printf("%f\n", b); /* "2.717000" */
13     printf("%s\n", c); /* "beej!" */
14     printf("%c\n", d); /* "X" */
15     printf("110%\n"); /* "110%" */
16
17     printf("%10d\n", a); /* "      100" */
18     printf("%-10d\n", a); /* "100      " */
19     printf("%*d\n", e, a); /* " 100" */
20     printf("%.2f\n", b); /* "2.72" */
21
22     printf("%hhhd\n", d); /* "88" <-- ASCII code for 'X' */
23
24     printf("%5d %5.2f %c\n", a, b, d); /* " 100 2.72 X" */
25 }
```

## See Also

`sprintf()`, `vprintf()`

## 23.11 `scanf()`, `fscanf()`, `sscanf()`

Read formatted string, character, or numeric data from the console or from a file

## Synopsis

```
#include <stdio.h>

int scanf(const char *format, ...);

int fscanf(FILE *stream, const char *format, ...);

int sscanf(const char * restrict s, const char * restrict format, ...);
```

## Description

These functions read formatted output from a variety of sources.

| Function              | Input Source                                            |
|-----------------------|---------------------------------------------------------|
| <code>scanf()</code>  | Read from the console (keyboard by default, typically). |
| <code>fscanf()</code> | Read from a file.                                       |
| <code>sscanf()</code> | Read from a string.                                     |

The only differences between these is are the leading parameters that you pass to them before the `format` string.

| Function              | What you pass before <code>format</code>            |
|-----------------------|-----------------------------------------------------|
| <code>scanf()</code>  | Nothing comes before <code>format</code> .          |
| <code>fscanf()</code> | Pass a <code>FILE*</code> .                         |
| <code>sscanf()</code> | Pass a <code>char*</code> to a buffer to read from. |

The `scanf()` family of functions reads data from the console or from a `FILE` stream, parses it, and stores the results away in variables you provide in the argument list.

The format string is very similar to that in `printf()` in that you can tell it to read a `"%d"`, for instance for an `int`. But it also has additional capabilities, most notably that it can eat up other characters in the input that you specify in the format string.

But let's start simple, and look at the most basic usage first before plunging into the depths of the function. We'll start by reading an `int` from the keyboard:

```
int a;

scanf("%d", &a);
```

`scanf()` obviously needs a pointer to the variable if it is going to change the variable itself, so we use the address-of operator to get the pointer.

In this case, `scanf()` walks down the format string, finds a `"%d"`, and then knows it needs to read an integer and store it in the next variable in the argument list, `a`.

Here are some of the other format specifiers you can put in the format string:

| Format Specifier | Description                                                                        |
|------------------|------------------------------------------------------------------------------------|
| <code>%d</code>  | Reads an integer to be stored in an <code>int</code> . This integer can be signed. |
| <code>%u</code>  | Reads an integer to be stored in an <code>unsigned int</code> .                    |
| <code>%f</code>  | Reads a floating point number, to be stored in a <code>float</code> .              |
| <code>%s</code>  | Reads a string up to the first whitespace character.                               |
| <code>%c</code>  | Reads a <code>char</code> .                                                        |

And that's the end of the story!

Ha! Just kidding. If you've just arrived from the `printf()` page, you know there's a near-infinite amount of additional material.

### 23.11.0.1 Consuming Other Characters

`scanf()` will move along the format string matching any characters you include.

For example, you could read a hyphenated date like so:

```
scanf("%u-%u-%u", &yyyy, &mm, &dd);
```

In that case, `scanf()` will attempt to consume an unsigned decimal number, then a hyphen, then another unsigned number, then another hyphen, then another unsigned number.

If it fails to match at any point (e.g. the user entered "foo"), `scanf()` will bail without consuming the offending characters.

And it will return the number of variables successfully converted. In the example above, if the user entered a valid string, `scanf()` would return 3, one for each variable successfully read.

### 23.11.0.2 Problems with `scanf()`

I (and the C FAQ and a lot of people) recommend *against* using `scanf()` to read directly from the keyboard. It's too easy for it to stop consuming characters when the user enters some bad data.

If you have data in a file and you're confident it's in good shape, `fscanf()` can be really useful.

But in the case of the keyboard or file, you can always use `fgets()` to read a complete line into a buffer, and then use `sscanf()` to scan things out of the buffer. This gives you the best of both worlds.

### 23.11.0.3 Problems with `sscanf()`

A while back, a third-party programmer rose to fame for figuring out how to cut *GTA Online* load times by 70%<sup>3</sup>.

What they'd discovered was that the implementation of `sscanf()` first effectively calls `strlen()`... so even if you're just using `sscanf()` to peel the first few characters off the string, it still runs all the way out to the end of the string first.

On small strings, no big deal, but on large strings with repeated calls (which is what was happening in *GTA*) it got *slooooooooooooo*www...

So if you're just converting a string to a number, consider `atoi()`, `atof()`, or the `strtol()` and `strtod()` families of functions, instead.

(The programmer collected a \$10,000 bug bounty for the effort.)

<sup>3</sup><https://nee.lv/2021/02/28/How-I-cut-GTA-Online-loading-times-by-70/>

### 23.11.0.4 The Deep Details

Let's check out what a `scanf()`

And here are some more codes, except these don't tend to be used as often. You, of course, may use them as often as you wish!

First, the format string. Like we mentioned, it can hold ordinary characters as well as `%` format specifiers. And whitespace characters.

Whitespace characters have a special role: a whitespace character will cause `scanf()` to consume as many whitespace characters as it can up to the next non-whitespace character. You can use this to ignore all leading or trailing whitespace.

Also, all format specifiers except for `s`, `c`, and `[]` automatically consume leading whitespace.

But I know what you're thinking: the meat of this function is in the format specifiers. What do those look like?

These consist of the following, in sequence:

1. A `%` sign
2. Optional: an `*` to suppress assignment—more later
3. Optional: a field width—max characters to read
4. Optional: length modifier, for specifying longer or shorter types
5. A conversion specifier, like `d` or `f` indicating the type to read

### 23.11.0.5 The Conversion Specifier

Let's start with the best and last: the *conversion specifier*.

This is the part of the format specifier that tells us what type of variable `scanf()` should be reading into, like `%d` or `%f`.

The binary conversion is new in C23!

| Conversion Specifier | Description                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>d</code>       | Matches a decimal <code>int</code> . Can have a leading sign.                                                                                               |
| <code>b</code>       | Matches a binary (base 2) <code>unsigned int</code> . Can have a leading sign.                                                                              |
| <code>i</code>       | Like <code>d</code> , except will handle it if you put a leading <code>0x</code> (hex) or <code>0</code> (octal) or <code>0b</code> (binary) on the number. |
| <code>o</code>       | Matches an octal (base 8) <code>unsigned int</code> . Leading zeros are ignored.                                                                            |
| <code>u</code>       | Matches a decimal <code>unsigned int</code> .                                                                                                               |
| <code>x</code>       | Matches a hex (base 16) <code>unsigned int</code> .                                                                                                         |
| <code>f</code>       | Match a floating point number (or scientific notation, or anything <code>strtod()</code> can handle).                                                       |
| <code>c</code>       | Match a <code>char</code> , or multiple <code>chars</code> if a field width is given.                                                                       |
| <code>s</code>       | Match a sequence of non-whitespace <code>chars</code> .                                                                                                     |
| <code>[]</code>      | Match a sequence of characters from a set. The set ends with <code>]</code> . More below.                                                                   |
| <code>p</code>       | Match a pointer, the opposite of <code>%p</code> for <code>printf()</code> .                                                                                |
| <code>n</code>       | Store the number of characters written so far in the given <code>int*</code> . Doesn't consume anything.                                                    |
| <code>%</code>       | Match a literal percent sign.                                                                                                                               |

All of the following are equivalent to the `f` specifier: `a`, `e`, `g`, `A`, `E`, `F`, `G`.

And capital `X` is equivalent to lowercase `x`.

**23.11.0.5.1 The Scanset %[ ] Conversion Specifier** This is about the weirdest format specifier there is. It allows you to specify a set of characters (the *scanset*) to be stored away (likely in an array of `chars`). Conversion stops when a character that is not in the set is matched.

For example, `%[0-9]` means “match all numbers zero through nine.” And `%[AD-G34]` means “match A, D through G, 3, or 4”.

Now, to convolute matters, you can tell `scanf()` to match characters that are *not* in the set by putting a caret (^) directly after the `%[` and following it with the set, like this: `%[^A-C]`, which means “match all characters that are *not* A through C.”

To match a close square bracket, make it the first character in the set, like this: `%[ ]A-C]` or `%[^ ]A-C]`. (I added the “A-C” just so it was clear that the “]” was first in the set.)

To match a hyphen, make it the last character in the set, e.g. to match A-through-C or hyphen: `%[A-C-]`.

So if we wanted to match all letters *except* “%”, “^”, “]”, “B”, “C”, “D”, “E”, and “-”, we could use this format string: `%[ ^]%^B-E-]`.

Got it? Now we can go onto the next func—no wait! There’s more! Yes, still more to know about `scanf()`. Does it never end? Try to imagine how I feel writing about it!

### 23.11.0.6 The Length Modifier

So you know that “%d” stores into an `int`. But how do you store into a `long`, `short`, or `double`?

Well, like in `printf()`, you can add a modifier before the type specifier to tell `scanf()` that you have a longer or shorter type. The following is a table of the possible modifiers:

| Length Modifier | Conversion Specifier                | Description                                                                                  |
|-----------------|-------------------------------------|----------------------------------------------------------------------------------------------|
| <code>hh</code> | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>char</code> (signed or unsigned as appropriate) before printing.      |
| <code>h</code>  | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>short int</code> (signed or unsigned as appropriate) before printing. |
| <code>l</code>  | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>long int</code> (signed or unsigned as appropriate).                  |
| <code>ll</code> | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>long long int</code> (signed or unsigned as appropriate).             |
| <code>j</code>  | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>intmax_t</code> or <code>uintmax_t</code> (as appropriate).           |
| <code>z</code>  | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>size_t</code> .                                                       |
| <code>t</code>  | <code>b, d, i, o, u, x, X</code>    | Convert input to <code>ptrdiff_t</code> .                                                    |
| <code>L</code>  | <code>a, A, e, E, f, F, g, G</code> | Convert input to <code>long double</code> .                                                  |
| <code>l</code>  | <code>c, s, [</code>                | Convert input to <code>wchar_t</code> , a wide character.                                    |
| <code>l</code>  | <code>s</code>                      | Argument is in a <code>wchar_t*</code> , a wide character string.                            |
| <code>hh</code> | <code>n</code>                      | Store result in <code>signed char*</code> argument.                                          |
| <code>h</code>  | <code>n</code>                      | Store result in <code>short int*</code> argument.                                            |
| <code>l</code>  | <code>n</code>                      | Store result in <code>long int*</code> argument.                                             |
| <code>ll</code> | <code>n</code>                      | Store result in <code>long long int*</code> argument.                                        |
| <code>j</code>  | <code>n</code>                      | Store result in <code>intmax_t*</code> argument.                                             |
| <code>z</code>  | <code>n</code>                      | Store result in <code>size_t*</code> argument.                                               |
| <code>t</code>  | <code>n</code>                      | Store result in <code>ptrdiff_t*</code> argument.                                            |

### 23.11.0.7 Field Widths

The field width generally allows you to specify a maximum number of characters to consume. If the thing you’re trying to match is shorter than the field width, that input will stop being processed before the field width is reached.

So a string will stop being consumed when whitespace is found, even if fewer than the field width characters are matched.

And a float will stop being consumed at the end of the number, even if fewer characters than the field width are matched.

But `%c` is an interesting one—it doesn't stop consuming characters on anything. So it'll go exactly to the field width. (Or 1 character if no field width is given.)

### 23.11.0.8 Skip Input with `*`

If you put an `*` in the format specifier, it tells `scanf()` to do the conversion specified, but not store it anywhere. It simply discards the data as it reads it. This is what you use if you want `scanf()` to eat some data but you don't want to store it anywhere; you don't give `scanf()` an argument for this conversion.

```
// Read 3 ints, but discard the middle one
scanf("%d %*d %d", &int1, &int3);
```

## Return Value

`scanf()` returns the number of items assigned into variables. Since assignment into variables stops when given invalid input for a certain format specifier, this can tell you if you've input all your data correctly.

Also, `scanf()` returns EOF on end-of-file.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      int a;
6      long int b;
7      unsigned int c;
8      float d;
9      double e;
10     long double f;
11     char s[100];
12
13     scanf("%d", &a); // store an int
14     scanf(" %d", &a); // eat any whitespace, then store an int
15     scanf("%s", s); // store a string
16     scanf("%Lf", &f); // store a long double
17
18     // store an unsigned, read all whitespace, then store a long int:
19     scanf("%u %ld", &c, &b);
20
21     // store an int, read whitespace, read "blendo", read whitespace,
22     // and store a float:
23     scanf("%d blendo %f", &a, &d);
24
25     // read all whitespace, then store all characters up to a newline
26     scanf(" %[^n]", s);
27
```

```

28 // store a float, read (and ignore) an int, then store a double:
29 scanf("%f %*d %lf", &d, &e);
30
31 // store 10 characters:
32 scanf("%10c", s);
33 }

```

## See Also

sscanf(), vscanf(), vsscanf(), vfscanf()

## 23.12 vprintf(), vfprintf(), vsprintf(), vsnprintf()

printf() variants using variable argument lists (va\_list)

### Synopsis

```

#include <stdio.h>
#include <stdarg.h>

int vprintf(const char * restrict format, va_list arg);

int vfprintf(FILE * restrict stream, const char * restrict format,
             va_list arg);

int vsprintf(char * restrict s, const char * restrict format, va_list arg);

int vsnprintf(char * restrict s, size_t n, const char * restrict format,
              va_list arg);

```

### Description

These are just like the printf() variants except instead of taking an actual variable number of arguments, they take a fixed number—the last of which is a va\_list that refers to the variable arguments.

Like with printf(), the different variants send output different places.

| Function    | Output Destination                               |
|-------------|--------------------------------------------------|
| vprintf()   | Print to console (screen by default, typically). |
| vfprintf()  | Print to a file.                                 |
| vsprintf()  | Print to a string.                               |
| vsnprintf() | Print to a string (safely).                      |

Both vsprintf() and vsnprintf() have the quality that if you pass in NULL as the buffer, nothing is written—but you can still check the return value to see how many characters *would* have been written.

If you try to write out more than the maximum number of characters, vsnprintf() will graciously write only  $n - 1$  characters so that it has enough room to write the terminator at the end.



As for why in the heck would you ever want to do this, the most common reason is to create your own specialized versions of `printf()`-type functions, piggybacking on all that `printf()` functionality goodness.

See the example for an example, predictably.

## Return Value

`vprintf()` and `vfprintf()` return the number of characters printed, or a negative value on error.

`vsprintf()` returns the number of characters printed to the buffer, not counting the NUL terminator, or a negative value if an error occurred.

`vnsprintf()` returns the number of characters printed to the buffer. Or the number that *would* have been printed if the buffer had been large enough.

## Example

In this example, we make our own version of `printf()` called `logger()` that timestamps output. Notice how the calls to `logger()` have all the bells and whistles of `printf()`.

```

1  #include <stdio.h>
2  #include <stdarg.h>
3  #include <time.h>
4
5  int logger(char *format, ...)
6  {
7      va_list va;
8      time_t now_secs = time(NULL);
9      struct tm *now = gmtime(&now_secs);
10
11     // Output timestamp in format "YYYY-MM-DD hh:mm:ss : "
12     printf("%04d-%02d-%02d %02d:%02d:%02d : ",
13           now->tm_year + 1900, now->tm_mon + 1, now->tm_mday,
14           now->tm_hour, now->tm_min, now->tm_sec);
15
16     va_start(va, format);
17     int result = vprintf(format, va);
18     va_end(va);
19
20     printf("\n");
21
22     return result;
23 }
24
25 int main(void)
26 {
27     int x = 12;
28     float y = 3.2;
29
30     logger("Hello!");
31     logger("x = %d and y = %.2f", x, y);
32 }
```

Output:

```
2021-03-30 04:25:49 : Hello!
2021-03-30 04:25:49 : x = 12 and y = 3.20
```

## See Also

`printf()`

## 23.13 `vscanf()`, `vfscanf()`, `vsscanf()`

`scanf()` variants using variable argument lists (`va_list`)

### Synopsis

```
#include <stdio.h>
#include <stdarg.h>

int vscanf(const char * restrict format, va_list arg);

int vfscanf(FILE * restrict stream, const char * restrict format,
            va_list arg);

int vsscanf(const char * restrict s, const char * restrict format,
            va_list arg);
```

### Description

These are just like the `scanf()` variants except instead of taking an actual variable number of arguments, they take a fixed number—the last of which is a `va_list` that refers to the variable arguments.

| Function               | Input Source                                            |
|------------------------|---------------------------------------------------------|
| <code>vscanf()</code>  | Read from the console (keyboard by default, typically). |
| <code>vfscanf()</code> | Read from a file.                                       |
| <code>vsscanf()</code> | Read from a string.                                     |

Like with the `vprintf()` functions, this would be a good way to add additional functionality that took advantage of the power `scanf()` has to offer.

### Return Value

Returns the number of items successfully scanned, or `EOF` on end-of-file or error.

### Example

I have to admit I was wracking my brain to think of when you'd ever want to use this. The best example I could find was one on Stack Overflow<sup>4</sup> that error-checks the return value from `scanf()` against the expected. A variant of that is shown below.

<sup>4</sup><https://stackoverflow.com/questions/17017331/c99-vscanf-for-dummies/17018046#17018046>

```

1  #include <stdio.h>
2  #include <stdarg.h>
3  #include <assert.h>
4
5  int error_check_scanf(int expected_count, char *format, ...)
6  {
7      va_list va;
8
9      va_start(va, format);
10     int count = vsscanf(format, va);
11     va_end(va);
12
13     // This line will crash the program if the condition is false:
14     assert(count == expected_count);
15
16     return count;
17 }
18
19 int main(void)
20 {
21     int a, b;
22     float c;
23
24     error_check_scanf(3, "%d, %d/%f", &a, &b, &c);
25     error_check_scanf(2, "%d", &a);
26 }

```

## See Also

scanf()

## 23.14 getc(), fgetc(), getchar()

Get a single character from the console or from a file

### Synopsis

```

#include <stdio.h>

int getc(FILE *stream);

int fgetc(FILE *stream);

int getchar(void);

```

### Description

All of these functions in one way or another, read a single character from the console or from a **FILE**. The differences are fairly minor, and here are the descriptions:

`getc()` returns a character from the specified `FILE`. From a usage standpoint, it's equivalent to the same `fgetc()` call, and `fgetc()` is a little more common to see. Only the implementation of the two functions differs.

`fgetc()` returns a character from the specified `FILE`. From a usage standpoint, it's equivalent to the same `getc()` call, except that `fgetc()` is a little more common to see. Only the implementation of the two functions differs.

Yes, I cheated and used cut-n-paste to do that last paragraph.

`getchar()` returns a character from `stdin`. In fact, it's the same as calling `getc(stdin)`.

## Return Value

All three functions return the `unsigned char` that they read, except it's cast to an `int`.

If end-of-file or an error is encountered, all three functions return `EOF`.

## Example

This example reads all the characters from a file, outputting only the letter 'b's it finds..

```

1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *fp;
6      int c;
7
8      fp = fopen("spoon.txt", "r"); // error check this!
9
10     // this while-statement assigns into c, and then checks against EOF:
11
12     while((c = fgetc(fp)) != EOF) {
13         if (c == 'b') {
14             putchar(c);
15         }
16     }
17
18     putchar('\n');
19
20     fclose(fp);
21 }
```

## See Also

---

### 23.15 `gets()`, `fgets()`

Read a string from console or file

## Synopsis

```
#include <stdio.h>

char *fgets(char *s, int size, FILE *stream);
char *gets(char *s);
```

## Description

These are functions that will retrieve a newline-terminated string from the console or a file. In other normal words, it reads a line of text. The behavior is slightly different, and, as such, so is the usage. For instance, here is the usage of `gets()`:

Don't use `gets()`. In fact, as of C11, it ceases to exist! This is one of the rare cases of a function being *removed* from the standard.

Admittedly, rationale would be useful, yes? For one thing, `gets()` doesn't allow you to specify the length of the buffer to store the string in. This would allow people to keep entering data past the end of your buffer, and believe me, this would be Bad News.

And that's what the `size` parameter in `fgets()` is for. `fgets()` will read at most `size-1` characters and then stick a `NUL` terminator on after that.

I was going to add another reason, but that's basically the primary and only reason not to use `gets()`. As you might suspect, `fgets()` allows you to specify a maximum string length.

One difference here between the two functions: `gets()` will devour and throw away the newline at the end of the line, while `fgets()` will store it at the end of your string (space permitting).

Here's an example of using `fgets()` from the console, making it behave more like `gets()` (with the exception of the newline inclusion):

```
char s[100];
gets(s); // don't use this--read a line (from stdin)
fgets(s, sizeof(s), stdin); // read a line from stdin
```

In this case, the `sizeof()` operator gives us the total size of the array in bytes, and since a `char` is a byte, it conveniently gives us the total size of the array.

Of course, like I keep saying, the string returned from `fgets()` probably has a newline at the end that you might not want. You can write a short function to chop the newline off—in fact, let's just roll that into our own version of `gets()`

```
#include <stdio.h>
#include <string.h>

char *ngets(char *s, int size)
{
    char *rv = fgets(s, size, stdin);

    if (rv == NULL)
        return NULL;

    char *p = strchr(s, '\n'); // Find a newline
```

```

    if (p != NULL) // if there's a newline
        *p = '\0'; // truncate the string there

    return s;
}

```

So, in summary, use `fgets()` to read a line of text from the keyboard or a file, and don't use `gets()`.

## Return Value

Both `gets()` and `fgets()` return a pointer to the string passed.

On error or end-of-file, the functions return `NULL`.

## Example

```

1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *fp;
6      char s[100];
7
8      gets(s); // read from standard input (don't use this--use fgets())
9
10     fgets(s, sizeof s, stdin); // read 100 bytes from standard input
11
12     fp = fopen("spoon.txt", "r"); // (you should error-check this)
13     fgets(s, 100, fp); // read 100 bytes from the file datafile.dat
14     fclose(fp);
15
16     fgets(s, 20, stdin); // read a maximum of 20 bytes from stdin
17 }

```

## See Also

`getc()`, `fgetc()`, `getchar()`, `puts()`, `fputs()`, `ungetc()`

## 23.16 `putc()`, `fputc()`, `putchar()`

Write a single character to the console or to a file

### Synopsis

```

#include <stdio.h>

int putc(int c, FILE *stream);

int fputc(int c, FILE *stream);

```

```
int putchar(int c);
```

## Description

All three functions output a single character, either to the console or to a `FILE`.

`putc()` takes a character argument, and outputs it to the specified `FILE`. `fputc()` does exactly the same thing, and differs from `putc()` in implementation only. Most people use `fputc()`.

`putchar()` writes the character to the console, and is the same as calling `putc(c, stdout)`.

## Return Value

All three functions return the character written on success, or `EOF` on error.

## Example

Print the alphabet:

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      char i;
6
7      for(i = 'A'; i <= 'Z'; i++)
8          putchar(i);
9
10     putchar('\n'); // put a newline at the end to make it pretty
11 }
```

## See Also

---

### 23.17 puts(), fputs()

Write a string to the console or to a file

## Synopsis

```
#include <stdio.h>

int puts(const char *s);

int fputs(const char *s, FILE *stream);
```

## Description

Both these functions output a NUL-terminated string. `puts()` outputs to the console, while `fputs()` allows you to specify the file for output.

## Return Value

Both functions return non-negative on success, or `EOF` on error.

## Example

Read strings from the console and save them in a file:

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     FILE *fp;
6     char s[100];
7
8     fp = fopen("somefile.txt", "w"); // error check this!
9
10    while(fgets(s, sizeof(s), stdin) != NULL) { // read a string
11        fputs(s, fp); // write it to the file we opened
12    }
13
14    fclose(fp);
15 }
```

## See Also

---

### 23.18 `ungetc()`

Pushes a character back into the input stream

## Synopsis

```
#include <stdio.h>

int ungetc(int c, FILE *stream);
```

## Description

You know how `getc()` reads the next character from a file stream? Well, this is the opposite of that—it pushes a character back into the file stream so that it will show up again on the very next read from the stream, as if you'd never gotten it from `getc()` in the first place.

Why, in the name of all that is holy would you want to do that? Perhaps you have a stream of data that you're reading a character at a time, and you won't know to stop reading until you get a certain character, but you want to be able to read that character again later. You can read the character, see that it's what you're supposed to stop on, and then `ungetc()` it so it'll show up on the next read.

Yeah, that doesn't happen very often, but there we are.

Here's the catch: the standard only guarantees that you'll be able to push back *one character*. Some implementations might allow you to push back more, but there's really no way to tell and still be portable.



## Return Value

On success, `ungetc()` returns the character you passed to it. On failure, it returns `EOF`.

## Example

This example reads a piece of punctuation, then everything after it up to the next piece of punctuation. It returns the leading punctuation, and stores the rest in a string.

```
1  #include <stdio.h>
2  #include <ctype.h>
3
4  int read_punctstring(FILE *fp, char *s)
5  {
6      int origpunct, c;
7
8      origpunct = fgetc(fp);
9
10     if (origpunct == EOF) // return EOF on end-of-file
11         return EOF;
12
13     while (c = fgetc(fp), !ispunct(c) && c != EOF)
14         *s++ = c; // save it in the string
15
16     *s = '\0'; // nul-terminate the string
17
18     // if we read punctuation last, ungetc it so we can fgetc it next
19     // time:
20     if (ispunct(c))
21         ungetc(c, fp);
22
23     return origpunct;
24 }
25
26 int main(void)
27 {
28     char s[128];
29     char c;
30
31     while((c = read_punctstring(stdin, s)) != EOF) {
32         printf("%c: %s\n", c, s);
33     }
34 }
```

Sample Input:

```
!foo#bar*baz
```

Sample output:

```
!: foo
#: bar
*: baz
```

## See Also

`fgetc()`

---

## 23.19 `fread()`

Read binary data from a file

### Synopsis

```
#include <stdio.h>

size_t fread(void *p, size_t size, size_t nmemb, FILE *stream);
```

### Description

You might remember that you can call `fopen()` with the “b” flag in the open mode string to open the file in “binary” mode. Files open in not-binary (ASCII or text mode) can be read using standard character-oriented calls like `fgetc()` or `fgets()`. Files open in binary mode are typically read using the `fread()` function.

All this function does is says, “Hey, read this many things where each thing is a certain number of bytes, and store the whole mess of them in memory starting at this pointer.”

This can be very useful, believe me, when you want to do something like store 20 `ints` in a file.

But wait—can’t you use `fprintf()` with the “%d” format specifier to save the `ints` to a text file and store them that way? Yes, sure. That has the advantage that a human can open the file and read the numbers. It has the disadvantage that it’s slower to convert the numbers from `ints` to text and that the numbers are likely to take more space in the file. (Remember, an `int` is likely 4 bytes, but the string “12345678” is 8 bytes.)

So storing the binary data can certainly be more compact and faster to read.

### Return Value

This function returns the number of items successfully read. If all requested items are read, the return value will be equal to that of the parameter `nmemb`. If EOF occurs, the return value will be zero.

To make you confused, it will also return zero if there’s an error. You can use the functions `feof()` or `ferror()` to tell which one really happened.

### Example

Read 10 numbers from a file and store them in an array:

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      int i;
6      int n[10]
7      FILE *fp;
8
9      fp = fopen("numbers.dat", "rb");
```

```
10     fread(n, sizeof(int), 10, fp); // read 10 ints
11     fclose(fp);
12
13     // print them out:
14     for(i = 0; i < 10; i++)
15         printf("n[%d] == %d\n", i, n[i]);
16 }
```

## See Also

fopen(), fwrite(), feof(), ferror()

---

## 23.20 fwrite()

Write binary data to a file

## Synopsis

```
#include <stdio.h>

size_t fwrite(const void *p, size_t size, size_t nmemb, FILE *stream);
```

## Description

This is the counterpart to the `fread()` function. It writes blocks of binary data to disk. For a description of what this means, see the entry for `fread()`.

## Return Value

`fwrite()` returns the number of items successfully written, which should hopefully be `nmemb` that you passed in. It'll return zero on error.

## Example

Save 10 random numbers to a file:

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      int i;
7      int n[10];
8      FILE *fp;
9
10     // populate the array with random numbers:
11     for(i = 0; i < 10; i++) {
12         n[i] = rand();
13         printf("n[%d] = %d\n", i, n[i]);
```

```
14     }
15
16     // save the random numbers (10 ints) to the file
17     fp = fopen("numbers.dat", "wb");
18     fwrite(n, sizeof(int), 10, fp); // write 10 ints
19     fclose(fp);
20 }
```

## See Also

fopen(), fread()

---

## 23.21 fgetpos(), fsetpos()

Get the current position in a file, or set the current position in a file. Just like `ftell()` and `fseek()` for most systems

## Synopsis

```
#include <stdio.h>

int fgetpos(FILE *stream, fpos_t *pos);

int fsetpos(FILE *stream, fpos_t *pos);
```

## Description

These functions are just like `ftell()` and `fseek()`, except instead of counting in bytes, they use an *opaque* data structure to hold positional information about the file. (Opaque, in this case, means you're not supposed to know what the data type is made up of.)

On virtually every system (and certainly every system that I know of), people don't use these functions, using `ftell()` and `fseek()` instead. These functions exist just in case your system can't remember file positions as a simple byte offset.

Since the `pos` variable is opaque, you have to assign to it using the `fgetpos()` call itself. Then you save the value for later and use it to reset the position using `fsetpos()`.

## Return Value

Both functions return zero on success, and `-1` on error.

## Example

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     char s[100];
```

```
6     fpos_t pos;
7     FILE *fp;
8
9     fp = fopen("spoon.txt", "r");
10
11     fgets(s, sizeof(s), fp); // read a line from the file
12     printf("%s", s);
13
14     fgetpos(fp, &pos); // save the position after the read
15
16     fgets(s, sizeof(s), fp); // read another line from the file
17     printf("%s", s);
18
19     fsetpos(fp, &pos); // now restore the position to where we saved
20
21     fgets(s, sizeof(s), fp); // read the earlier line again
22     printf("%s", s);
23
24     fclose(fp);
25 }
```

## See Also

`fseek()`, `ftell()`, `rewind()`

---

## 23.22 `fseek()`, `rewind()`

Position the file pointer in anticipation of the next read or write

### Synopsis

```
#include <stdio.h>

int fseek(FILE *stream, long offset, int whence);

void rewind(FILE *stream);
```

### Description

When doing reads and writes to a file, the OS keeps track of where you are in the file using a counter generically known as the file pointer. You can reposition the file pointer to a different point in the file using the `fseek()` call. Think of it as a way to randomly access you file.

The first argument is the file in question, obviously. `offset` argument is the position that you want to seek to, and `whence` is what that offset is relative to.

Of course, you probably like to think of the offset as being from the beginning of the file. I mean, “Seek to position 3490, that should be 3490 bytes from the beginning of the file.” Well, it *can* be, but it doesn’t have to be. Imagine the power you’re wielding here. Try to command your enthusiasm.

You can set the value of `whence` to one of three things:

| whence   | Description                                                                                                                                                                                                                                  |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SEEK_SET | offset is relative to the beginning of the file. This is probably what you had in mind anyway, and is the most commonly used value for whence.                                                                                               |
| SEEK_CUR | offset is relative to the current file pointer position. So, in effect, you can say, “Move to my current position plus 30 bytes,” or, “move to my current position minus 20 bytes.”                                                          |
| SEEK_END | offset is relative to the end of the file. Just like SEEK_SET except from the other end of the file. Be sure to use negative values for offset if you want to back up from the end of the file, instead of going past the end into oblivion. |

Speaking of seeking off the end of the file, can you do it? Sure thing. In fact, you can seek way off the end and then write a character; the file will be expanded to a size big enough to hold a bunch of zeros way out to that character.

Now that the complicated function is out of the way, what’s this `rewind()` that I briefly mentioned? It repositions the file pointer at the beginning of the file:

```
fseek(fp, 0, SEEK_SET); // same as rewind()
rewind(fp);             // same as fseek(fp, 0, SEEK_SET)
```

## Return Value

For `fseek()`, on success zero is returned; `-1` is returned on failure.

The call to `rewind()` never fails.

## Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      FILE *fp;
6
7      fp = fopen("spoon.txt", "r");
8
9      fseek(fp, 100, SEEK_SET); // seek to the 100th byte of the file
10     printf("100: %c\n", fgetc(fp));
11
12     fseek(fp, -31, SEEK_CUR); // seek backward 30 bytes from the current pos
13     printf("31 back: %c\n", fgetc(fp));
14
15     fseek(fp, -12, SEEK_END); // seek to the 10th byte before the end of file
16     printf("12 from end: %c\n", fgetc(fp));
17
18     fseek(fp, 0, SEEK_SET);    // seek to the beginning of the file
19     rewind(fp);               // seek to the beginning of the file, too
20     printf("Beginning: %c\n", fgetc(fp));
21
22     fclose(fp);
23 }
```

## See Also

`ftell()`, `fgetpos()`, `fsetpos()`

---

## 23.23 `ftell()`

Tells you where a particular file is about to read from or write to

### Synopsis

```
#include <stdio.h>

long ftell(FILE *stream);
```

### Description

This function is the opposite of `fseek()`. It tells you where in the file the next file operation will occur relative to the beginning of the file.

It's useful if you want to remember where you are in the file, `fseek()` somewhere else, and then come back later. You can take the return value from `ftell()` and feed it back into `fseek()` (with `whence` parameter set to `SEEK_SET`) when you want to return to your previous position.

### Return Value

Returns the current offset in the file, or `-1` on error.

### Example

```
1  #include <stdio.h>
2
3  int main(void)
4  {
5      char c[6];
6      FILE *fp;
7
8      fp = fopen("spoon.txt", "r");
9
10     long pos;
11
12     // seek ahead 10 bytes:
13     fseek(fp, 10, SEEK_SET);
14
15     // store the current position in variable "pos":
16     pos = ftell(fp);
17
18     // Read some bytes
19     fread(c, sizeof c - 1, 1, fp);
20     c[5] = '\0';
21     printf("Read: \"%s\"\n", c);
```

```
22 // and return to the starting position, stored in "pos":
23 fseek(fp, pos, SEEK_SET);
24
25 // Read the same bytes again
26 fread(c, sizeof c - 1, 1, fp);
27 c[5] = '\0';
28 printf("Read: \"%s\\n\"", c);
29
30 fclose(fp);
31 }
32
```

## See Also

`fseek()`, `rewind()`, `fgetpos()`, `fsetpos()`

---

## 23.24 `feof()`, `ferror()`, `clearerr()`

Determine if a file has reached end-of-file or if an error has occurred

### Synopsis

```
#include <stdio.h>

int feof(FILE *stream);

int ferror(FILE *stream);

void clearerr(FILE *stream);
```

### Description

Each `FILE*` that you use to read and write data from and to a file contains flags that the system sets when certain events occur. If you get an error, it sets the error flag; if you reach the end of the file during a read, it sets the EOF flag. Pretty simple really.

The functions `feof()` and `ferror()` give you a simple way to test these flags: they'll return non-zero (true) if they're set.

Once the flags are set for a particular stream, they stay that way until you call `clearerr()` to clear them.

### Return Value

`feof()` and `ferror()` return non-zero (true) if the file has reached EOF or there has been an error, respectively.

### Example

Read binary data, checking for EOF or error:



```

1  #include <stdio.h>
2
3  int main(void)
4  {
5      int a;
6      FILE *fp;
7
8      fp = fopen("numbers.dat", "r");
9
10     // read single ints at a time, stopping on EOF or error:
11
12     while(fread(&a, sizeof(int), 1, fp), !feof(fp) && !ferror(fp)) {
13         printf("Read %d\n", a);
14     }
15
16     if (feof(fp))
17         printf("End of file was reached.\n");
18
19     if (ferror(fp))
20         printf("An error occurred.\n");
21
22     fclose(fp);
23 }

```

**See Also**

fopen(), fread()

---

**23.25 perror()**

Print the last error message to `stderr`

**Synopsis**

```

#include <stdio.h>
#include <errno.h> // only if you want to directly use the "errno" var

void perror(const char *s);

```

**Description**

Many functions, when they encounter an error condition for whatever reason, will set a global variable called `errno` (in `<errno.h>`) for you. `errno` is just an integer representing a unique error.

But to you, the user, some number isn't generally very useful. For this reason, you can call `perror()` after an error occurs to print what error has actually happened in a nice human-readable string.

And to help you along, you can pass a parameter, `s`, that will be prepended to the error string for you.

One more clever trick you can do is check the value of the `errno` (you have to include `errno.h` to see it) for specific errors and have your code do different things. Perhaps you want to ignore certain errors but not

others, for instance.

The standard only defines three values for `errno`, but your system undoubtedly defines more. The three that are defined are:

| <code>errno</code> | Description                                               |
|--------------------|-----------------------------------------------------------|
| EDOM               | Math operation outside domain.                            |
| EILSEQ             | Invalid sequence in multibyte to wide character encoding. |
| ERANGE             | Result of operation doesn't fit in specified type.        |

The catch is that different systems define different values for `errno`, so it's not very portable beyond the above 3. The good news is that at least the values are *largely* portable between Unix-like systems, at least.

## Return Value

Returns nothing at all! Sorry!

## Example

`fseek()` returns `-1` on error, and sets `errno`, so let's use it. Seeking on `stdin` makes no sense, so it should generate an error:

```

1  #include <stdio.h>
2  #include <errno.h> // must include this to see "errno" in this example
3
4  int main(void)
5  {
6      if (fseek(stdin, 10L, SEEK_SET) < 0)
7          perror("fseek");
8
9      fclose(stdin); // stop using this stream
10
11     if (fseek(stdin, 20L, SEEK_CUR) < 0) {
12
13         // specifically check errno to see what kind of
14         // error happened...this works on Linux, but your
15         // mileage may vary on other systems!
16
17         if (errno == EBADF) {
18             perror("fseek again, EBADF");
19         } else {
20             perror("fseek again");
21         }
22     }
23 }
```

And the output is:

```
fseek: Illegal seek
fseek again, EBADF: Bad file descriptor
```

## See Also

`feof()`, `ferror()`, `strerror()`

## Chapter 24

# <stdlib.h> Standard Library Functions

Some of the following functions have variants that handle different types: `atoi()`, `strtod()`, `strtol()`, `abs()`, and `div()`. Only a single one is listed here for brevity.

| Function                     | Description                                            |
|------------------------------|--------------------------------------------------------|
| <code>_Exit()</code>         | Exit the currently-running program and don't look back |
| <code>abort()</code>         | Abruptly end program execution                         |
| <code>abs()</code>           | Compute the absolute value of an integer               |
| <code>aligned_alloc()</code> | Allocate specifically-aligned memory                   |
| <code>at_quick_exit()</code> | Set up handlers to run when the program quickly exits  |
| <code>atexit()</code>        | Set up handlers to run when the program exits          |
| <code>atof()</code>          | Convert a string to a floating point value             |
| <code>atoi()</code>          | Convert an integer in a string into a integer type     |
| <code>bsearch()</code>       | Binary Search (maybe) an array of objects              |
| <code>calloc()</code>        | Allocate and clear memory for arbitrary use            |
| <code>div()</code>           | Compute the quotient and remainder of two numbers      |
| <code>exit()</code>          | Exit the currently-running program                     |
| <code>free()</code>          | Free a memory region                                   |
| <code>getenv()</code>        | Get the value of an environment variable               |
| <code>malloc()</code>        | Allocate memory for arbitrary use                      |
| <code>mblen()</code>         | Return the number of bytes in a multibyte character    |
| <code>mbstowcs()</code>      | Convert a multibyte string to a wide character string  |
| <code>mbtowc()</code>        | Convert a multibyte character to a wide character      |
| <code>qsort()</code>         | Quicksort (maybe) some data                            |
| <code>quick_exit()</code>    | Exit the currently-running program quickly             |
| <code>rand()</code>          | Return a pseudorandom number                           |
| <code>realloc()</code>       | Resize a previously allocated stretch of memory        |
| <code>srand()</code>         | Seed the built-in pseudorandom number generator        |
| <code>strtod()</code>        | Convert a string to a floating point number            |
| <code>strtol()</code>        | Convert a string to an integer                         |
| <code>system()</code>        | Run an external program                                |
| <code>wcstombs()</code>      | Convert a wide character string to a multibyte string  |
| <code>wctomb()</code>        | Convert a wide character to a multibyte character      |

The `<stdlib.h>` header has all kinds of—dare I say—miscellaneous functions bundled into it. This func-

tionality includes:

- Conversions from numbers to strings
- Conversions from strings to numbers
- Pseudorandom number generation
- Dynamic memory allocation
- Various ways to exit the program
- Ability to run external programs
- Binary search (or some fast search)
- Quicksort (or some fast sort)
- Integer arithmetic functions
- Multibyte and wide character and string conversions

So, you know... a little of everything.

## 24.1 <stdlib.h> Types and Macros

A couple new types and macros are introduced, though some of these might also be defined elsewhere:

| Type                 | Description                                          |
|----------------------|------------------------------------------------------|
| <code>size_t</code>  | Returned from <code>sizeof</code> and used elsewhere |
| <code>wchar_t</code> | For wide character operations                        |
| <code>div_t</code>   | For the <code>div()</code> function                  |
| <code>ldiv_t</code>  | For the <code>ldiv()</code> function                 |
| <code>lldiv_t</code> | for the <code>lldiv()</code> function                |

And some macros:

| Type                      | Description                                                                |
|---------------------------|----------------------------------------------------------------------------|
| <code>NULL</code>         | Our good pointer friend                                                    |
| <code>EXIT_SUCCESS</code> | Good exit status when things go well                                       |
| <code>EXIT_FAILURE</code> | Good exit status when things go poorly                                     |
| <code>RAND_MAX</code>     | The maximum value that can be returned by the <code>rand()</code> function |
| <code>MB_CUR_MAX</code>   | Maximum number of bytes in a multibyte character in the current locale     |

And there you have it. Just a lot of fun, useful functions in here. Let's check 'em out!

## 24.2 `atof()`

Convert a string to a floating point value

## Synopsis

```
#include <stdlib.h>

double atof(const char *nptr);
```

## Description

This stood for “ASCII-To-Floating” back in the day<sup>1</sup>, but no one would dare to use such coarse language now.

But the gist is the same: we’re going to convert a string with numbers and (optionally) a decimal point into a floating point value. Leading whitespace is ignored, and translation stops at the first invalid character.

If the result doesn’t fit in a `double`, behavior is undefined.

It generally works as if you’d called `strtod()`:

```
strtod(nptr, NULL)
```

So check out that reference page for more info.

In fact, `strtod()` is just better and you should probably use that.

## Return Value

Returns the string converted to a `double`.

## Example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main(void)
5 {
6     double x = atof("3.141593");
7
8     printf("%f\n", x); // 3.141593
9 }
```

## See Also

`atoi()`, `strtod()`

---

## 24.3 `atoi()`, `atol()`, `atoll()`

Convert an integer in a string into a integer type

---

<sup>1</sup><http://man.cat-v.org/unix-1st3/atof>

## Synopsis

```
#include <stdlib.h>

int atoi(const char *nptr);

long int atol(const char *nptr);

long long int atoll(const char *nptr);
```

## Description

Back in the day, `atoi()` stood for “ASCII-To\_Integer”<sup>2</sup> but now the spec makes no mention of that.

These functions take a string with a number in them and convert it to an integer of the specified return type. Leading whitespace is ignored. Translation stops at the first invalid character.

If the result doesn’t fit in the return type, behavior is undefined.

It generally works as if you’d called `strtol()` family of functions:

```
atoi(nptr)           // is basically the same as...
(int)strtol(nptr, NULL, 10)

atol(nptr)            // is basically the same as...
strtol(nptr, NULL, 10)

atoll(nptr)           // is basically the same as...
strtoll(nptr, NULL, 10)
```

Again, the `strtol()` functions are generally better, so I recommend them instead of these.

## Return Value

Returns an integer result corresponding to the return type.

## Example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main(void)
5 {
6     int x = atoi("3490");
7
8     printf("%d\n", x); // 3490
9 }
```

## See Also

`atof()`, `strtol()`

---

<sup>2</sup><http://man.cat-v.org/unix-1st/3/atoi>

## 24.4 strtod(), strtof(), strtold()

Convert a string to a floating point number

### Synopsis

```
#include <stdlib.h>

double strtod(const char * restrict nptr, char ** restrict endptr);

float strtof(const char * restrict nptr, char ** restrict endptr);

long double strtold(const char * restrict nptr, char ** restrict endptr);
```

### Description

These are some neat functions that convert strings to floating point numbers (or even NaN or Infinity) and provide some error checking, besides.

Firstly, leading whitespace is skipped.

Then the functions attempt to convert characters into the floating point result. Finally, when an invalid character (or NUL character) is reached, they set `endptr` to point to the invalid character.

Set `endptr` to `NULL` if you don't care about where the first invalid character is.

If you didn't set `endptr` to `NULL`, it will point to a NUL character if the translation didn't find any bad characters. That is:

```
if (*endptr == '\0') {
    printf("What a perfectly-formed number!\n");
} else {
    printf("I found badness in your number: \"%s\"\n", endptr);
}
```

But guess what! You can also translate strings into special values, like NaN and Infinity!

If `nptr` points to a string containing `INF` or `INFINITY` (upper or lowercase), the value for Infinity will be returned.

If `nptr` points to a string containing `NAN`, then (a quiet, non-signalling) NaN will be returned. You can tag the `NAN` with a sequence of characters from the set `0-9`, `a-z`, `A-Z`, and `_` by enclosing them in parens:

```
NAN(foobar_3490)
```

What your compiler does with this is implementation-defined, but it can be used to specify different kinds of NaN.

You can also specify a number in hexadecimal with a power-of-two exponent ( $2^x$ ) if you lead with `0x` (or `0X`). For the exponent, use a `p` followed by a base 10 exponent. (You can't use `e` because that's a valid hex digit!)

Example:

```
0xabc.123p15
```



Which computes to  $0xabc.123 \times 2^{15}$ .

You can put in `FLT_DECIMAL_DIG`, `DBL_DECIMAL_DIG`, or `LDBL_DECIMAL_DIG` digits and get a correctly-rounded result for the type.

## Return Value

Returns the converted number. If there was no number, returns 0. `endptr` is set to point to the first invalid character, or the NUL terminator if all characters were consumed.

If there's an overflow, `HUGE_VAL`, `HUGE_VALF`, or `HUGE_VALL` is returned, signed like the input, and `errno` is set to `ERANGE`.

If there's an underflow, it returns the smallest number closest to zero with the input sign. `errno` may be set to `ERANGE`.

## Example

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      char *inp = " 123.4567beej";
7      char *badchar;
8
9      double val = strtod(inp, &badchar);
10
11     printf("Converted string to %f\n", val);
12     printf("Encountered bad characters: %s\n", badchar);
13
14     val = strtod("987.654321beej", NULL);
15     printf("Ignoring bad chars for result: %f\n", val);
16
17     val = strtod("11.2233", &badchar);
18
19     if (*badchar == '\0')
20         printf("No bad chars: %f\n", val);
21     else
22         printf("Found bad chars: %f, %s\n", val, badchar);
23 }
```

Output:

```

Converted string to 123.456700
Encountered bad characters: beej
Ignoring bad chars: 987.654321
No bad chars: 11.223300
```

## See Also

`atof()`, `strtoul()`

## 24.5 `strtol()`, `strtoll()`, `strtoul()`, `strtoull()`

Convert a string to an integer

### Synopsis

```
#include <stdlib.h>

long int strtol(const char * restrict nptr,
               char ** restrict endptr, int base);

long long int strtoll(const char * restrict nptr,
                    char ** restrict endptr, int base);

unsigned long int strtoul(const char * restrict nptr,
                        char ** restrict endptr, int base);

unsigned long long int strtoull(const char * restrict nptr,
                              char ** restrict endptr, int base);
```

### Description

These convert a string to an integer like `atoi()`, but they have a few more bells and whistles.

Most notable, they can tell you where conversion started going wrong, i.e. where invalid characters, if any, appear. Leading spaces are ignored. A `+` or `-` sign may precede the number.

The basic idea is that if things go well, these functions will return the integer values contained in the strings. And if you pass in the `char**` typed `endptr`, it'll set it to point at the NUL at the end of the string.

If things don't go well, they'll set `endptr` to point at the first character where things have gone awry. That is, if you're converting a value `103z2!` in base 10, they'll send `endptr` to point at the `z` because that's the first non-numeric character.

You can pass in `NULL` for `endptr` if you don't care to do any of that kind of error checking.

Wait—did I just say we could set the number base for the conversion? Yes! Yes, I did. Now number bases<sup>3</sup> are out of scope for this document, but certainly some of the more well-known are binary (base 2), octal (base 8), decimal (base 10), and hexadecimal (base 16).

You can specify the number base for the conversion as the third parameter. Bases from 2 to 36 are supported, with case-insensitive digits running from `0` to `Z`.

If you specify a base of `0`, the function will make an effort to determine it. It'll default to base 10 except for a couple cases:

The binary conversion is new in C23!

- If the number has a leading `0b` or `0B`, it will be binary (base 2)
- If the number has a leading `0`, it will be octal (base 8)
- If the number has a leading `0x` or `0X`, it will be hex (base 16)

The locale might affect the behavior of these functions.

---

<sup>3</sup><https://en.wikipedia.org/wiki/Radix>

## Return Value

Returns the converted value.

`endptr`, if not `NULL` is set to the first invalid character, or to the beginning of the string if no conversion was performed, or to the string terminal `NUL` if all characters were valid.

If there's overflow, one of these values will be returned: `LONG_MIN`, `LONG_MAX`, `LLONG_MIN`, `LLONG_MAX`, `ULONG_MAX`, `ULLONG_MAX`. And `errno` is set to `ERANGE`.

## Example

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      // All output in decimal (base 10)
7
8      printf("%ld\n", strtol("123", NULL, 0));    // 123
9      printf("%ld\n", strtol("123", NULL, 10));   // 123
10     printf("%ld\n", strtol("101010", NULL, 2));  // binary, 42
11     printf("%ld\n", strtol("123", NULL, 8));     // octal, 83
12     printf("%ld\n", strtol("123", NULL, 16));    // hex, 291
13
14     printf("%ld\n", strtol("0123", NULL, 0));    // octal, 83
15     printf("%ld\n", strtol("0x123", NULL, 0));   // hex, 291
16
17     char *badchar;
18     long int x = strtol(" 1234beej", &badchar, 0);
19
20     printf("Value is %ld\n", x);                  // Value is 1234
21     printf("Bad chars at \"%s\"\n", badchar);     // Bad chars at "beej"
22 }
```

Output:

```

123
123
42
83
291
83
291
Value is 1234
Bad chars at "beej"
```

## See Also

`atoi()`, `strtod()`, `setlocale()`, `strtoimax()`, `strtoumax()`

---

## 24.6 rand()

Return a pseudorandom number

### Synopsis

```
#include <stdlib.h>

int rand(void);
```

### Description

This gives us back a pseudorandom number in the range 0 to `RAND_MAX`, inclusive. (`RAND_MAX` will be at least 32767.)

If you want to force this to a certain range, the classic way to do this is to force it with the modulo operator `%`, although this introduces biases<sup>4</sup> if `RAND_MAX+1` is not a multiple of the number you're modding by. Dealing with this is out of scope for this guide.

If you want to make a floating point number between 0 and 1 inclusive, you can divide the result by `RAND_MAX`. Or `RAND_MAX+1` if you don't want to include 1. But of course, there are out-of-scope problems with this, as well<sup>5</sup>.

In short, `rand()` is a great way to get potentially poor random numbers with ease. Probably good enough for the game you're writing.

The spec elaborates:

There are no guarantees as to the quality of the random sequence produced and some implementations are known to produce sequences with distressingly non-random low-order bits. Applications with particular requirements should use a generator that is known to be sufficient for their needs.

Your system probably has a good random number generator on it if you need a stronger source. Linux users have `getrandom()`, for example, and Windows has `CryptGenRandom()`.

For other more demanding random number work, you might find a library like the GNU Scientific Library<sup>6</sup> of use.

With most implementations, the numbers produced by `rand()` will be the same from run to run. To get around this, you need to start it off in a different place by passing a *seed* into the random number generator. You can do this with `srand()`.

### Return Value

Returns a random number in the range 0 to `RAND_MAX`, inclusive.

### Example

Note that all of these examples don't produce perfectly uniform distributions. But good enough for the untrained eye, and really common in general use when mediocre random number quality is acceptable.

<sup>4</sup><https://stackoverflow.com/questions/10984974/why-do-people-say-there-is-modulo-bias-when-using-a-random-number-generator>

<sup>5</sup><https://mumble.net/~campbell/2014/04/28/uniform-random-float>

<sup>6</sup><https://www.gnu.org/software/gsl/doc/html/rng.html>

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      printf("RAND_MAX = %d\n", RAND_MAX);
7
8      printf("0 to 9: %d\n", rand() % 10);
9
10     printf("10 to 44: %d\n", rand() % 35 + 10);
11     printf("0 to 0.99999: %f\n", rand() / ((float)RAND_MAX + 1));
12     printf("10.5 to 15.7: %f\n", 10.5 + 5.2 * rand() / (float)RAND_MAX);
13 }

```

Output on my system:

```

RAND_MAX = 2147483647
0 to 9: 3
10 to 44: 21
0 to 0.99999: 0.783099
10.5 to 15.7: 14.651888

```

Example of seeding the RNG with the time:

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <time.h>
4
5  int main(void)
6  {
7      // time(NULL) very likely returns the number of seconds since
8      // January 1, 1970:
9
10     srand(time(NULL));
11
12     for (int i = 0; i < 5; i++)
13         printf("%d\n", rand());
14 }

```

## See Also

srand()

## 24.7 srand()

Seed the built-in pseudorandom number generator

## Synopsis

```
#include <stdlib.h>

void srand(unsigned int seed);
```

## Description

The dirty little secret of pseudorandom number generation is that they’re completely deterministic. There’s nothing random about them. They just look random.

If you use `rand()` and run your program several times, you might notice something *fishy*: they produce the same random numbers over and over again.

To mix it up, we need to give the pseudorandom number generator a new “starting point”, if you will. We call that the *seed*. It’s just a number, but it is used as the basis for subsequent number generation. Give a different seed, and you’ll get a different sequence of random numbers. Give the same seed, and you’ll get the same sequence of random numbers corresponding to it<sup>7</sup>.

So if you call `srand(3490)` before you start generating numbers with `rand()`, you’ll get the same sequence every time. `srand(37)` would also give you the same sequence every time, but it would be a different sequence than the one you got with `srand(3490)`.

But if you can’t hardcode the seed (because that would give you the same sequence every time), how are you supposed to do this?

It’s really common to use the number of seconds since January 1, 1970 (this date is known as the *Unix epoch*<sup>8</sup>) to seed the generator. This sounds pretty arbitrary except for the fact that it’s exactly the value most implementations return from the library call `time(NULL)`<sup>9</sup>.

We’ll do that in the example.

If you don’t call `srand()`, it’s as if you called `srand(1)`.

## Return Value

Returns nothing!

## Example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <time.h>    // for the time() call
4
5 int main(void)
6 {
7     srand(time(NULL));
8
9     for (int i = 0; i < 5; i++)
10         printf("%d\n", rand() % 32);
```

<sup>7</sup>Minecraft enthusiasts might recall that when generating a new world, they were given the option to enter a random number seed. That single value is used to generate that entire random world. And if your friend starts a world with the same seed you did, they’ll get the same world you did.

<sup>8</sup>[https://en.wikipedia.org/wiki/Unix\\_time](https://en.wikipedia.org/wiki/Unix_time)

<sup>9</sup>The C spec doesn’t say exactly what `time(NULL)` will return, but the POSIX spec does! And virtually everyone returns exactly that: the number of seconds since epoch.

11

}

Output:

```
4
20
22
14
9
```

Output from a subsequent run:

```
19
0
31
31
24
```

**See Also**`rand()`, `time()`**24.8 aligned\_alloc()**

Allocate specifically-aligned memory

**Synopsis**

```
#include <stdlib.h>

void *aligned_alloc(size_t alignment, size_t size);
```

**Description**

Maybe you wanted `malloc()` or `calloc()` instead of this. But if you're sure you don't, read on!

Normally you don't have to think about this, since `malloc()` and `realloc()` both provide memory regions that are suitably aligned<sup>10</sup> for use with any data type.

But if you need a more specific alignment, you can specify it with this function.

When you're done using the memory region, be sure to free it with a call to `free()`.

Don't pass in 0 for the size. It probably won't do anything you want.

In case you're wondering, all dynamically-allocated memory is automatically freed by the system when the program ends. That said, it's considered to be *Good Form* to explicitly `free()` everything you allocate. This way other programmers don't think you were being sloppy.

<sup>10</sup>[https://en.wikipedia.org/wiki/Data\\_structure\\_alignment](https://en.wikipedia.org/wiki/Data_structure_alignment)

## Return Value

Returns a pointer to the newly-allocated memory, aligned as specified. Returns `NULL` if something goes wrong.

## Example

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <stdint.h>
4
5  int main(void)
6  {
7      int *p = aligned_alloc(256, 10 * sizeof(int));
8
9      // Just for fun, let's convert to intptr_t and mod with 256
10     // to make sure we're actually aligned on a 256-byte boundary.
11     //
12     // This is probably some kind of implementation-defined
13     // behavior, but I'll bet it works.
14
15     intptr_t ip = (intptr_t)p;
16
17     printf("%ld\n", ip % 256);    // 0!
18
19     // Free it up
20     free(p);
21 }
```

## See Also

`malloc()`, `calloc()`, `free()`

---

## 24.9 `calloc()`, `malloc()`

Allocate memory for arbitrary use

### Synopsis

```
#include <stdlib.h>

void *calloc(size_t nmem, size_t size);

void *malloc(size_t size);
```

### Description

Both of these functions allocate memory for general-purpose use. It will be aligned such that it's useable for storing any data type.



`malloc()` allocates exactly the specified number of bytes of memory in a contiguous block. The memory might be full of garbage data. (You can clear it with `memset()`, if you wish.)

`calloc()` is different in that it allocates space for `nmemb` objects of `size` bytes each. (You can do the same with `malloc()`, but you have to do the multiplication yourself.)

`calloc()` has an additional feature: it clears all the memory to 0.

So if you're planning to zero the memory anyway, `calloc()` is probably the way to go. If you're not, you can avoid that overhead by calling `malloc()`.

When you're done using the memory region, free it with a call to `free()`.

Don't pass in 0 for the size. It probably won't do anything you want.

In case you're wondering, all dynamically-allocated memory is automatically freed by the system when the program ends. That said, it's considered to be *Good Form* to explicitly `free()` everything you allocate. This way other programmers don't think you were being sloppy.

## Return Value

Both functions return a pointer to the shiny, newly-allocated memory. Or `NULL` if something's gone awry.

## Example

Comparison of `malloc()` and `calloc()` for allocating 5 ints:

```
1  #include <stdlib.h>
2
3  int main(void)
4  {
5      // Allocate space for 5 ints
6      int *p = malloc(5 * sizeof(int));
7
8      p[0] = 12;
9      p[1] = 30;
10
11     // Allocate space for 5 ints
12     // (Also clear that memory to 0)
13     int *q = calloc(5, sizeof(int));
14
15     q[0] = 12;
16     q[1] = 30;
17
18     // All done
19     free(p);
20     free(q);
21 }
```

## See Also

`aligned_alloc()`, `free()`

---

## 24.10 free()

Free a memory region

### Synopsis

```
#include <stdlib.h>

void free(void *ptr);
```

### Description

You know that pointer you got back from `malloc()`, `calloc()`, or `aligned_alloc()`? You pass that pointer to `free()` to free the memory associated with it.

If you don't do this, the memory will stay allocated FOREVER AND EVER! (Well, until your program exits, anyway.)

Fun fact: `free(NULL)` does nothing. You can safely call that. Sometimes it's convenient.

Don't `free()` a pointer that's already been `free()`d. Don't `free()` a pointer that you didn't get back from one of the allocation functions. It would be *Bad*<sup>11</sup>.

### Return Value

Returns nothing!

### Example

```
1  #include <stdlib.h>
2
3  int main(void)
4  {
5      // Allocate space for 5 ints
6      int *p = malloc(5 * sizeof(int));
7
8      p[0] = 12;
9      p[1] = 30;
10
11     // Free that space
12     free(p);
13 }
```

### See Also

`malloc()`, `calloc()`, `aligned_alloc()`

---

<sup>11</sup>“Try to imagine all life as you know it stopping instantaneously and every molecule in your body exploding at the speed of light.”  
—Egon Spengler

## 24.11 realloc()

Resize a previously allocated stretch of memory

### Synopsis

```
#include <stdlib.h>

void *realloc(void *ptr, size_t size);
```

### Description

This takes a pointer to some memory previously allocated with `malloc()` or `calloc()` and resizes it to the new size.

If the new size is smaller than the old size, any data larger than the new size is discarded.

If the new size is larger than the old size, the new larger part is uninitialized. (You can clear it with `memset()`.)

Important note: the memory might move! If you resize, the system might need to relocate the memory to a larger contiguous chunk. If this happens, `realloc()` will copy the old data to the new location for you.

Because of this, it's important to save the returned value to your pointer to update it to the new location if things move. (Also, be sure to error-check so that you don't overwrite your old pointer with `NULL`, leaking the memory.)

You can also `realloc()` memory allocated with `aligned_alloc()`, but it will potentially lose its alignment if the block is moved.

### Return Value

Returns a pointer to the resized memory region. This might be equivalent to the `ptr` passed in, or it might be some other location.

### Example

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      // Allocate space for 5 ints
7      int *p = malloc(5 * sizeof(int));
8
9      p[0] = 12;
10     p[1] = 30;
11
12     // Reallocate for 10 bytes
13     int *new_p = realloc(p, 10 * sizeof(int));
14
15     if (new_p == NULL) {
16         printf("Error reallocating\n");
17     } else {
18         p = new_p; // It's good; let's keep it
```

```
19     p[7] = 99;
20 }
21
22 // All done
23 free(p);
24 }
```

## See Also

`malloc()`, `calloc()`

---

## 24.12 `abort()`

Abruptly end program execution

### Synopsis

```
#include <stdlib.h>

_Noreturn void abort(void);
```

### Description

This ends program execution *abnormally* and immediately. Use this in rare, unexpected circumstances.

Open streams might not be flushed. Temporary files created might not be removed. Exit handlers are not called.

A non-zero exit status is returned to the environment.

On some systems, `abort()` might dump core<sup>12</sup>, but this is outside the scope of the spec.

You can cause the equivalent of an `abort()` by calling `raise(SIGABRT)`, but I don't know why you'd do that.

The only portable way to stop an `abort()` call midway is to use `signal()` to catch `SIGABRT` and then `exit()` in the signal handler.

### Return Value

This function never returns.

### Example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main(void)
5 {
```

---

<sup>12</sup>[https://en.wikipedia.org/wiki/Core\\_dump](https://en.wikipedia.org/wiki/Core_dump)

```

6      int bad_thing = 1;
7
8      if (bad_thing) {
9          printf("This should never have happened!\n");
10         fflush(stdout); // Make sure the message goes out
11         abort();
12     }
13 }

```

On my system, this outputs:

```

This should never have happened!
zsh: abort (core dumped) ./foo

```

## See Also

`signal()`

## 24.13 `atexit()`, `at_quick_exit()`

Set up handlers to run when the program exits

### Synopsis

```

#include <stdlib.h>

int atexit(void (*func)(void));

int at_quick_exit(void (*func)(void));

```

### Description

When the program does a normal exit with `exit()` or returns from `main()`, it looks for previously-registered handlers to call on the way out. These handlers are registered with the `atexit()` call.

Think of it like, “Hey, when you’re about to exit, do these extra things.”

For the `quick_exit()` call, you can use the `at_quick_exit()` function to register handlers for that<sup>13</sup>. There’s no crossover in handlers from `exit()` to `quick_exit()`, i.e. for a call to one, none of the other’s handlers will fire.

You can register multiple handlers to fire—at least 32 handlers are supported by both `exit()` and `quick_exit()`.

The argument `func` to the functions looks a little weird—it’s a pointer to a function to call. Basically just put the function name to call in there (without parentheses after). See the example, below.

If you call `atexit()` from inside your `atexit()` handler (or equivalent in your `at_quick_exit()` handler), it’s unspecified if it will get called. So get them all registered before you exit.

When exiting, the functions will be called in the reverse order they were registered.

<sup>13</sup>`quick_exit()` differs from `exit()` in that open files might not be flushed and temporary files might not be removed.

## Return Value

These functions return 0 on success, or nonzero on failure.

## Example

`atexit()`:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 void exit_handler_1(void)
5 {
6     printf("Exit handler 1 called!\n");
7 }
8
9 void exit_handler_2(void)
10 {
11     printf("Exit handler 2 called!\n");
12 }
13
14 int main(void)
15 {
16     atexit(exit_handler_1);
17     atexit(exit_handler_2);
18
19     exit(0);
20 }
```

For the output:

```
Exit handler 2 called!
Exit handler 1 called!
```

And a similar example with `quick_exit()`:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 void exit_handler_1(void)
5 {
6     printf("Exit handler 1 called!\n");
7 }
8
9 void exit_handler_2(void)
10 {
11     printf("Exit handler 2 called!\n");
12 }
13
14 int main(void)
15 {
16     at_quick_exit(exit_handler_1);
17     at_quick_exit(exit_handler_2);
18 }
```

```

18
19     quick_exit(0);
20 }

```

## See Also

`exit()`, `quick_exit()`

## 24.14 `exit()`, `quick_exit()`, `_Exit()`

Exit the currently-running program

### Synopsis

```

#include <stdlib.h>

_Noreturn void exit(int status);

_Noreturn void quick_exit(int status);

_Noreturn void _Exit(int status);

```

### Description

All these functions cause the program to exit, with various levels of cleanup performed.

`exit()` does the most cleanup and is the most normal exit.

`quick_exit()` is the second most.

`_Exit()` unceremoniously drops everything and ragequits on the spot.

Calling either of `exit()` or `quick_exit()` causes their respective `atexit()` or `at_quick_exit()` handlers to be called in the reverse order in which they were registered.

`exit()` will flush all streams and delete all temporary files.

`quick_exit()` or `_Exit()` might not perform that nicety.

`_Exit()` doesn't call any of the at-exit handlers, either.

For all functions, the exit `status` is returned to the environment.

Defined exit statuses are:

| Status                    | Description                                      |
|---------------------------|--------------------------------------------------|
| <code>EXIT_SUCCESS</code> | Typically returned when good things happen       |
| <code>0</code>            | Same as <code>EXIT_SUCCESS</code>                |
| <code>EXIT_FAILURE</code> | Oh noes! Definitely failure!                     |
| Any positive value        | Generally indicates another failure of some kind |

OS X note: `quick_exit()` is not supported.

## Return Value

None of these functions ever return.

## Example

```
1  #include <stdlib.h>
2
3  int main(void)
4  {
5      int contrived_exit_type = 1;
6
7      switch(contrived_exit_type) {
8          case 1:
9              exit(EXIT_SUCCESS);
10
11         case 2:
12             // Not supported in OS X
13             quick_exit(EXIT_SUCCESS);
14
15         case 3:
16             _Exit(2);
17     }
18 }
```

## See Also

`atexit()`, `at_quick_exit()`

---

## 24.15 `getenv()`

Get the value of an environment variable

### Synopsis

```
#include <stdlib.h>

char *getenv(const char *name);
```

### Description

The environment often provides variables that are set before the program run that you can access at runtime.

Of course the exact details are system dependent, but these variables are key/value pairs, and you can get the value by passing the key to `getenv()` as the `name` parameter.

You're not allowed to overwrite the string that's returned.

This is pretty limited in the standard, but your OS often provides better functionality.



## Return Value

Returns a pointer to the environment variable value, or `NULL` if the variable doesn't exist.

## Example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main(void)
5 {
6     printf("PATH is %s\n", getenv("PATH"));
7 }
```

Output (truncated in my case):

```
PATH is /usr/bin:/usr/local/bin:/usr/sbin:/home/beej/.cargo/bin [...]
```

## 24.16 system()

Run an external program

### Synopsis

```
#include <stdlib.h>

int system(const char *string);
```

### Description

This will run an external program and then return to the caller.

The manner in which it runs the program is system-defined, but typically you can pass something to it just like you'd run on the command line, searching the `PATH`, etc.

Not all systems have this capability, but you can test for it by passing `NULL` to `system()` and seeing if it returns 0 (no command processor is available) or non-zero (a command processor is available! Yay!)

If you're getting user input and passing it to the `system()` call, be extremely careful to escape all special shell characters (everything that's not alphanumeric) with a backslash to keep a villain from running something you don't want them to.

### Return Value

If `NULL` is passed, returns nonzero if a command processor is available (i.e. `system()` will work at all).

Otherwise returns an implementation-defined value.

**Example**

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      printf("Here's a directory listing:\n\n");
7
8      system("ls -l");    // Run this command and return
9
10     printf("\nAll done!\n");
11 }

```

Output:

```

Here's a directory listing:

total 92
drwxr-xr-x 3 beej beej  4096 Oct 14 21:38 bin
drwxr-xr-x 2 beej beej  4096 Dec 20 20:07 examples
-rwxr-xr-x 1 beej beej 16656 Feb 23 21:49 foo
-rw-rw-rw- 1 beej beej   155 Feb 23 21:49 foo.c
-rw-r--r-- 1 beej beej  1350 Jan 27 22:11 Makefile
-rw-r--r-- 1 beej beej  4644 Jan 18 09:12 README.md
drwxr-xr-x 3 beej beej  4096 Feb 23 20:21 src
drwxr-xr-x 6 beej beej  4096 Feb 21 20:24 stage
drwxr-xr-x 2 beej beej  4096 Sep 27 20:54 translations
drwxr-xr-x 2 beej beej  4096 Sep 27 20:54 website

All done!

```

**24.17 bsearch()**

Binary Search (maybe) an array of objects

**Synopsis**

```

#include <stdlib.h>

// Pre-C23:

void *bsearch(const void *key, const void *base,
              size_t nmemb, size_t size,
              int (*compar)(const void *, const void *));

// C23:

QVoid *bsearch(const void *key, QVoid *base,
               size_t nmemb, size_t size,

```

```
int (*compar)(const void *, const void *));
```

## Description

This crazy-looking function searches an array for a value.

It probably is a binary search or some fast, efficient search. But the spec doesn't really say.

However, the array must be sorted! So binary search seems likely.

- `key` is a pointer to the value to find.
- `base` is a pointer to the start of the array—the array must be sorted!
- `nmemb` is the number of elements in the array.
- `size` is the `sizeof` each element in the array.
- `compar` is a pointer to a function that will compare the key against other values.

The comparison function takes the key as the first argument and the value to compare against as the second. It should return a negative number if the key is less than the value, 0 if the key equals the value, and a positive number if the key is greater than the value.

This is commonly computed by taking the difference between the key and the value to be compared. If subtraction is supported.

The return value from the `strcmp()` function can be used for comparing strings.

Again, the array must be sorted according to the order of the comparison function before running `bsearch()`. Luckily for you, you can just call `qsort()` with the same comparison function to get this done.

It's a general-purpose function—it'll search any type of array for anything. The catch is you have to write the comparison function.

And that's not as scary as it looks. Jump down to the example

## Return Value

The function returns a pointer to the found value, or `NULL` if it can't be found.

New in C23: If `base` is `const`, the return type of the `bsearch()` function will also be `const`.

## Example

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int compar(const void *key, const void *value)
5  {
6      const int *k = key, *v = value; // Need ints, not voids
7
8      return *k - *v;
9  }
10
11 int main(void)
12 {
13     int a[9] = {2, 6, 9, 12, 13, 18, 20, 32, 47};
14
15     int *r, key;
16
```

```

17     key = 12; // 12 is in there
18     r = bsearch(&key, a, 9, sizeof(int), compar);
19     printf("Found %d\n", *r);
20
21     key = 30; // Won't find a 30
22     r = bsearch(&key, a, 9, sizeof(int), compar);
23     if (r == NULL)
24         printf("Didn't find 30\n");
25
26     // Searching with an unnamed key, pointer to 32
27     r = bsearch(&(int){32}, a, 9, sizeof(int), compar);
28     printf("Found %d\n", *r); // Found it
29 }

```

Output:

```

Found 12
Didn't find 30
Found 32

```

## See Also

`strcmp()`, `qsort()`

## 24.18 `qsort()`

Quicksort (maybe) some data

### Synopsis

```

#include <stdlib.h>

void qsort(void *base, size_t nmemb, size_t size,
           int (*compar)(const void *, const void *));

```

### Description

This function will quicksort (or some other sort, probably speedy) an array of data in-place<sup>14</sup>.

Like `bsearch()`, it's data-agnostic. Any data for which you can define a relative ordering can be sorted, whether ints, structs, or anything else.

Also like `bsearch()`, you have to give a comparison function to do the actual compare.

- `base` is a pointer to the start of the array to be sorted.
- `nmemb` is the number of elements in the array.
- `size` is the `sizeof` each element.
- `compar` is a pointer to the comparison function.

<sup>14</sup>"In-place" meaning that the original array will hold the results; no new array is allocated.

The comparison function takes pointers to two elements of the array as arguments and compares them. It should return a negative number if the first argument is less than the second, 0 if they are equal, and a positive number if the first argument is greater than the second.

This is commonly computed by taking the difference between the first argument and the second. If subtraction is supported.

The return value from the `strcmp()` function can provide sort order for strings.

If you have to sort a `struct`, just subtract the specific field you want to sort by.

This comparison function can be used by `bsearch()` to do searches after the list is sorted.

To reverse the sort, subtract the second argument from the first, i.e. negate the return value from `compar()`.

## Return Value

Returns nothing!

## Example

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int compar(const void *elem0, const void *elem1)
5  {
6      const int *x = elem0, *y = elem1; // Need ints, not voids
7
8      if (*x > *y) return 1;
9      if (*x < *y) return -1;
10     return 0;
11 }
12
13 int main(void)
14 {
15     int a[9] = {14, 2, 3, 17, 10, 8, 6, 1, 13};
16
17     // Sort the list
18
19     qsort(a, 9, sizeof(int), compar);
20
21     // Print sorted list
22
23     for (int i = 0; i < 9; i++)
24         printf("%d ", a[i]);
25
26     putchar('\n');
27
28     // Use the same compar() function to binary search
29     // for 17 (passed in as an unnamed object)
30
31     int *r = bsearch(&(int){17}, a, 9, sizeof(int), compar);
32     printf("Found %d!\n", *r);
33 }
```

Output:

```
1 2 3 6 8 10 13 14 17
Found 17!
```

## See Also

`strcmp()`, `bsearch()`

---

## 24.19 `abs()`, `labs()`, `llabs()`

Compute the absolute value of an integer

### Synopsis

```
#include <stdlib.h>

int abs(int j);

long int labs(long int j);

long long int llabs(long long int j);
```

### Description

Compute the absolute value of `j`. If you don't remember, that's how far from zero `j` is.

In other words, if `j` is negative, return it as a positive. If it's positive, return it as a positive. Always be positive. Enjoy life.

If the result cannot be represented, the behavior is undefined. Be especially aware of the upper half of unsigned numbers.

### Return Value

Returns the absolute value of `j`,  $|j|$ .

### Example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 int main(void)
5 {
6     printf("|-2| = %d\n", abs(-2));
7     printf("|4| = %d\n", abs(4));
8 }
```

Output:

```
| -2 | = 2  
| 4 | = 4
```

## See Also

`fabs()`

---

## 24.20 `div()`, `ldiv()`, `lldiv()`

Compute the quotient and remainder of two numbers

### Synopsis

```
#include <stdlib.h>  
  
div_t div(int numer, int denom);  
  
ldiv_t ldiv(long int numer, long int denom);  
  
lldiv_t lldiv(long long int numer, long long int denom);
```

### Description

These functions get you the quotient and remainder of a pair of numbers in one go.

They return a structure that has two fields, `quot`, and `rem`, the types of which match types of `numer` and `denom`. Note how each function returns a different variant of `div_t`.

These `div_t` variants are equivalent to the following:

```
typedef struct {  
    int quot, rem;  
} div_t;  
  
typedef struct {  
    long int quot, rem;  
} ldiv_t;  
  
typedef struct {  
    long long int quot, rem;  
} lldiv_t;
```

Why use these instead of the division operator?

The C99 Rationale says:

Because C89 had implementation-defined semantics for division of signed integers when negative operands were involved, `div` and `ldiv`, and `lldiv` in C99, were invented to provide well-specified semantics for signed integer division and remainder operations. The semantics were adopted to be the same as in Fortran. Since these functions return both the quotient and the remainder, they also serve as a convenient way of efficiently modeling underlying hardware that computes both results as

| part of the same operation. Table 7.2 summarizes the semantics of these functions.

Indeed, K&R2 (C89) says:

| The direction of truncation for `/` and the sign of the result for `%` are machine-dependent for negative operands [...]

The Rationale then goes on to spell out what the signs of the quotient and remainder will be given the signs of a numerator and denominator when using the `div()` functions:

| numer | denom | quot | rem |
|-------|-------|------|-----|
| +     | +     | +    | +   |
| −     | +     | −    | −   |
| +     | −     | −    | +   |
| −     | −     | +    | −   |

## Return Value

A `div_t`, `ldiv_t`, or `lldiv_t` structure with the `quot` and `rem` fields loaded with the quotient and remainder of the operation of `number/denom`.

## Example

```

1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main(void)
5  {
6      div_t d = div(64, -7);
7
8      printf("64 / -7 = %d\n", d.quot);
9      printf("64 %% -7 = %d\n", d.rem);
10 }
```

Output:

```

64 / -7 = -9
64 % -7 = 1
```

## See Also

`fmod()`, `remainder()`

## 24.21 `mblen()`

Return the number of bytes in a multibyte character



## Synopsis

```
#include <stdlib.h>

int mblen(const char *s, size_t n);
```

## Description

If you have a multibyte character in a string, this will tell you how many bytes long it is.

`n` is the maximum number of bytes `mblen()` will scan before giving up.

If `s` is a `NULL` pointer, tests if this encoding has state dependency, as noted in the return value, below. It also resets the state, if there is one.

The behavior of this function is influenced by the locale.

## Return Value

Returns the number of bytes used to encode this character, or `-1` if there is no valid multibyte character in the next `n` bytes.

Or, if `s` is `NULL`, returns true if this encoding has state dependency.

## Example

For the example, I used my extended character set to put Unicode characters in the source. If this doesn't work for you, use the `\uXXXX` escape.

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <locale.h>
4
5  int main(void)
6  {
7      setlocale(LC_ALL, "");
8
9      printf("State dependency: %d\n", mblen(NULL, 0));
10     printf("Bytes for €: %d\n", mblen("€", 5));
11     printf("Bytes for \u00e9: %d\n", mblen("\u00e9", 5)); // \u00e9 == é
12     printf("Bytes for &: %d\n", mblen("&", 5));
13 }
```

Output (in my case, the encoding is UTF-8, but your mileage may vary):

```
State dependency: 0
Bytes for €: 3
Bytes for é: 2
Bytes for &: 1
```

## See Also

`mbtowc()`, `mbstowcs()`, `setlocale()`

## 24.22 mbtowc()

Convert a multibyte character to a wide character

### Synopsis

```
#include <stdlib.h>

int mbtowc(wchar_t * restrict pwc, const char * restrict s, size_t n);
```

### Description

If you have a multibyte character, this function will convert it to a wide character and stored at the address pointed to by `pwc`. Up to `n` bytes of the multibyte character will be analyzed.

If `pwc` is `NULL`, the resulting character will not be stored. (Useful for just getting the return value.)

If `s` is a `NULL` pointer, tests if this encoding has state dependency, as noted in the return value, below. It also resets the state, if there is one.

The behavior of this function is influenced by the locale.

### Return Value

Returns the number of bytes used in the encoded wide character, or `-1` if there is no valid multibyte character in the next `n` bytes.

Returns `0` if `s` points to the NUL character.

Or, if `s` is `NULL`, returns true if this encoding has state dependency.

### Example

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <locale.h>
4  #include <wchar.h>
5
6  int main(void)
7  {
8      setlocale(LC_ALL, "");
9
10     printf("State dependency: %d\n", mbtowc(NULL, NULL, 0));
11
12     wchar_t wc;
13     int bytes;
14
15     bytes = mbtowc(&wc, "€", 5);
16
17     printf("L'%lc' takes %d bytes as multibyte char '€'\n", wc, bytes);
18 }
```

Output on my system:

```
State dependency: 0
L'€' takes 3 bytes as multibyte char '€'
```

## See Also

`mblen()`, `mbstowcs()`, `wcstombs()`, `setlocale()`

---

## 24.23 `wctomb()`

Convert a wide character to a multibyte character

### Synopsis

```
#include <stdlib.h>

int wctomb(char *s, wchar_t wc);
```

### Description

If you have your hands on a wide character, you can use this to make it multibyte.

The wide character `wc` is stored as a multibyte character in the string pointed to by `s`. The buffer `s` points to should be at least `MB_CUR_MAX` characters long. Note that `MB_CUR_MAX` changes with locale.

If `wc` is a NUL wide character, a NUL is stored in `s` after the bytes needed to reset the shift state (if any).

If `s` is a `NULL` pointer, tests if this encoding has state dependency, as noted in the return value, below. It also resets the state, if there is one.

The behavior of this function is influenced by the locale.

### Return Value

Returns the number of bytes used in the encoded multibyte character, or `-1` if `wc` does not correspond to any valid multibyte character.

Or, if `s` is `NULL`, returns true if this encoding has state dependency.

### Example

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <locale.h>
4  #include <wchar.h>
5
6  int main(void)
7  {
8      setlocale(LC_ALL, "");
9
10     printf("State dependency: %d\n", mbtowc(NULL, NULL, 0));
11
```

```

12     int bytes;
13     char mb[MB_CUR_MAX + 1];
14
15     bytes = wctomb(mb, L'€');
16     mb[bytes] = '\0';
17
18     printf("L'€' takes %d bytes as multibyte char '%s'\n", bytes, mb);
19 }

```

Output on my system:

```

State dependency: 0
L'€' takes 3 bytes as multibyte char '€'

```

## See Also

`mbtowl()`, `mbstowcs()`, `wcstombs()`, `setlocale()`

## 24.24 mbstowcs()

Convert a multibyte string to a wide character string

### Synopsis

```

#include <stdlib.h>

size_t mbstowcs(wchar_t * restrict pwcs, const char * restrict s, size_t n);

```

### Description

If you have a multibyte string (AKA a regular string), you can convert it wto a wide character string with this function.

At most `n` wide characters are written to the destination `pwcs` from the source `s`.

A NUL character is stored as a wide NUL character.

Non-portable POSIX extension: if you're using a POSIX-complaint library, this function allows `pwcs` to be `NULL` if you're only interested in the return value. Most notably, this will give you the number of characters in a multibyte string (as opposed to `strlen()` which counts the bytes.)

### Return Value

Returns the number of wide characters written to the destination `pwcs`.

If an invalid multibyte character was found, returns `(size_t)(-1)`.

If the return value is `n`, it means the result was *not* NUL-terminated.

## Example

This source uses an extended character set. If your compiler doesn't support it, you'll have to replace them with `\u` escapes.

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <locale.h>
4  #include <string.h>
5
6  int main(void)
7  {
8      setlocale(LC_ALL, "");
9
10     wchar_t wcs[128];
11     char *s = "€200 for this spoon?"; // 20 characters
12
13     size_t char_count, byte_count;
14
15     char_count = mbstowcs(wcs, s, 128);
16     byte_count = strlen(s);
17
18     printf("Wide string: L\"%ls\"\n", wcs);
19     printf("Char count : %zu\n", char_count); // 20
20     printf("Byte count : %zu\n\n", byte_count); // 22 on my system
21
22     // POSIX Extension that allows you to pass NULL for
23     // the destination so you can just use the return
24     // value (which is the character count of the string,
25     // if no errors have occurred)
26
27     s = "§¶°±π€•"; // 7 characters
28
29     char_count = mbstowcs(NULL, s, 0); // POSIX-only, nonportable
30     byte_count = strlen(s);
31
32     printf("Multibyte str: \"%s\"\n", s);
33     printf("Char count : %zu\n", char_count); // 7
34     printf("Byte count : %zu\n", byte_count); // 16 on my system
35 }

```

Output on my system (byte count will depend on your encoding):

```

Wide string: L"€200 for this spoon?"
Char count : 20
Byte count : 22

Multibyte str: "§¶°±π€•"
Char count : 7
Byte count : 16

```

## See Also

`mblen()`, `mbtowc()`, `wcstombs()`, `setlocale()`

## 24.25 wcstombs()

Convert a wide character string to a multibyte string

### Synopsis

```
#include <stdlib.h>

size_t wcstombs(char * restrict s, const wchar_t * restrict pwcs, size_t n);
```

### Description

If you have a wide character string and you want it as multibyte string, this is the function for you!

It'll take the wide characters pointed to by `pwcs` and convert them to multibyte characters stored in `s`. No more than `n` bytes will be written to `s`.

Non-portable POSIX extension: if you're using a POSIX-complaint library, this function allows `s` to be `NULL` if you're only interested in the return value. Most notably, this will give you the number of bytes needed to encode the wide characters in a multibyte string.

### Return Value

Returns the number of bytes written to `s`, or `(size_t)(-1)` if one of the characters can't be encoded into a multibyte string.

If the return value is `n`, it means the result was *not* NUL-terminated.

### Example

This source uses an extended character set. If your compiler doesn't support it, you'll have to replace them with `\u` escapes.

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <locale.h>
4  #include <string.h>
5
6  int main(void)
7  {
8      setlocale(LC_ALL, "");
9
10     char mbs[128];
11     wchar_t *wcs = L"€200 for this spoon?"; // 20 characters
12
13     size_t byte_count;
14
15     byte_count = wcstombs(mbs, wcs, 128);
16
17     printf("Wide string: L\"%ls\"\n", wcs);
18     printf("Multibyte  : \"%s\"\n", mbs);
```

```

19     printf("Byte count : %zu\n\n", byte_count); // 22 on my system
20
21     // POSIX Extension that allows you to pass NULL for
22     // the destination so you can just use the return
23     // value (which is the character count of the string,
24     // if no errors have occurred)
25
26     wcs = L"§¶°±π€•"; // 7 characters
27
28     byte_count = wcstombs(NULL, wcs, 0); // POSIX-only, nonportable
29
30     printf("Wide string: L\"%ls\"\n", wcs);
31     printf("Byte count : %zu\n", byte_count); // 16 on my system
32 }

```

Output on my system (byte count will depend on your encoding):

```

Wide string: L"€200 for this spoon?"
Multibyte   : "€200 for this spoon?"
Byte count  : 22

Wide string: L"§¶°±π€•"
Byte count  : 16

```

## See Also

`mblen()`, `wctomb()`, `mbstowcs()`, `setlocale()`

## 24.26 memalignment()

### Synopsis

New in C23!

```

#include <stdlib.h>

size_t memalignment(const void * p);

```

### Description

This function will give you the memory alignment (in bytes) of the object pointed to.

The idea is that we're going to get the same information as we get with `alignof`, except we can do it with an untyped `void*`.

And you might want to do this because misaligned data could be slower or unusable on various platforms.

### Return Value

Returns the integer alignment of `p`, which will be a power of 2. Returns `0` if `NULL` is passed in.

`0` means the pointer can't be used to access an object of any type.

## Example

(Caveat: none of my compilers support this function yet, so the code is largely untested.)

The proposal for this feature suggests that a macro might be useful for determining if a pointer is well-aligned for a particular type, so we'll just steal that as an example:

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

#define isaligned(P, T) (memalignment (P) >= alignof(T))

void check(void *p)
{
    printf("%lu %lu\n", memalignment(p), alignof(int));
    if (isaligned(p, int)) {
        puts("The pointer p is well-aligned for an int! :)");

        // So I'm happy doing this:

        int *i = p;
        *i = 3490;

    } else
        puts("The pointer p is poorly-aligned for an int! :(");
}

int main(void)
{
    char *p = malloc(10);

    check(p);
    check(p + 1);
}
```

## See Also

`alignof()`, `alignas()`, `max_align_t`



## Chapter 25

# <stdnoreturn.h> Macros for Non-Returning Functions

This header provides a macro `noreturn` that is a handy alias for `_Noreturn`.

Use this macro to indicate to the compiler that a function will never return to the caller. It's undefined behavior if the so-marked function does return.

Here's a usage example:

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <stdnoreturn.h>
4
5  noreturn void foo(void) // This function should never return!
6  {
7      printf("Happy days\n");
8
9      exit(1);           // And it doesn't return--it exits here!
10 }
11
12 int main(void)
13 {
14     foo();
15 }
```

That's all there is to it.

## Chapter 26

# <string.h> String Manipulation

| Function                       | Description                                                                  |
|--------------------------------|------------------------------------------------------------------------------|
| <code>memcpy()</code>          | Copy a region of memory to another, stopping at a specified character.       |
| <code>memchr()</code>          | Find the first occurrence of a character in memory.                          |
| <code>memcmp()</code>          | Compare two regions of memory.                                               |
| <code>memcpy()</code>          | Copy a region of memory to another.                                          |
| <code>memmove()</code>         | Move a (potentially overlapping) region of memory.                           |
| <code>memset()</code>          | Set a region of memory to a value.                                           |
| <code>memset_explicit()</code> | Set a region of memory to a value in a manner that cannot be optimized away. |
| <code>strcat()</code>          | Concatenate (join) two strings together.                                     |
| <code>strchr()</code>          | Find the first occurrence of a character in a string.                        |
| <code>strcmp()</code>          | Compare two strings.                                                         |
| <code>strcoll()</code>         | Compare two strings accounting for locale.                                   |
| <code>strcpy()</code>          | Copy a string.                                                               |
| <code>strcspn()</code>         | Find length of a string not consisting of a set of characters.               |
| <code>strdup()</code>          | Duplicate a string on the heap.                                              |
| <code>strerror()</code>        | Return a human-readable error message for a given code.                      |
| <code>strlen()</code>          | Return the length of a string.                                               |
| <code>strncat()</code>         | Concatenate (join) two strings, length-limited.                              |
| <code>strncmp()</code>         | Compare two strings, length-limited.                                         |
| <code>strncpy()</code>         | Copy two strings, length-limited.                                            |
| <code>strndup()</code>         | Duplicate a string on the heap, length-limited.                              |
| <code>strpbrk()</code>         | Search a string for one of a set of character.                               |
| <code>strrchr()</code>         | Find the last occurrence of a character in a string.                         |
| <code>strspn()</code>          | Find length of a string consisting of a set of characters.                   |
| <code>strstr()</code>          | Find a substring in a string.                                                |
| <code>strtok()</code>          | Tokenize a string.                                                           |
| <code>strxfrm()</code>         | Prepare a string for comparison as if by <code>strcoll()</code> .            |

As has been mentioned earlier in the guide, a string in C is a sequence of bytes in memory, terminated by a NUL character (`'\0'`). The NUL at the end is important, since it lets all these string functions (and `printf()` and `puts()` and everything else that deals with a string) know where the end of the string actually is.

Fortunately, when you operate on a string using one of these many functions available to you, they add the NUL terminator on for you, so you actually rarely have to keep track of it yourself. (Sometimes you do, especially if you're building a string from scratch a character at a time or something.)

In this section you'll find functions for pulling substrings out of strings, concatenating strings together, getting the length of a string, and so forth and so on.

---

## 26.1 `memcpy()`, `memccpy()`, `memmove()`

Copy bytes of memory from one location to another

### Synopsis

`memccpy()` is new in C23!

```
#include <string.h>

void *memcpy(void * restrict s1, const void * restrict s2, size_t n);

void *memccpy(void * restrict s1, const void * restrict s2, int c,
              size_t n);

void *memmove(void *s1, const void *s2, size_t n);
```

### Description

These functions copy memory—as many bytes as you want! From source to destination for as many bytes as you specify!

`memccpy()` gives you the additional option of stopping after a certain character is encountered in the source.

The main difference between the `memcpy()` variants and `memmove()` is that the former cannot safely copy overlapping memory regions, whereas the latter can.

I'm not sure why you'd want to ever use `memcpy()` instead of `memmove()`, but I'll bet it's possibly more performant.

The parameters are in a particular order: destination first, then source. I remember this order because it behaves like an “=” assignment: the destination is on the left.

### Return Value

Both functions return whatever you passed in for parameter `s1` for your convenience.

### Example

```
1 #include <string.h>
2
3 int main(void)
4 {
5     char s[100] = "Goats";
6     char t[100];
7
```

```

8     memcpy(t, s, 6);           // Copy non-overlapping memory
9
10    memmove(s + 2, s, 6);      // Copy overlapping memory
11 }

```

## See Also

strcpy(), strncpy()

---

## 26.2 strcpy(), strncpy()

Copy a string

### Synopsis

```

#include <string.h>

char *strcpy(char *dest, char *src);

char *strncpy(char *dest, char *src, size_t n);

```

### Description

These functions copy a string from one address to another, stopping at the NUL terminator on the `src` string. `strncpy()` is just like `strcpy()`, except only the first `n` characters are actually copied. Beware that if you hit the limit, `n` before you get a NUL terminator on the `src` string, your `dest` string won't be NUL-terminated. Beware! BEWARE!

(If the `src` string has fewer than `n` characters, it works just like `strcpy()`.)

You can terminate the string yourself by sticking the `'\0'` in there yourself:

```

char s[10];
char foo = "My hovercraft is full of eels."; // more than 10 chars

strncpy(s, foo, 9); // only copy 9 chars into positions 0-8
s[9] = '\0';        // position 9 gets the terminator

```

### Return Value

Both functions return `dest` for your convenience, at no extra charge.

### Example

```

1  #include <string.h>
2
3  int main(void)
4  {

```

```

5     char *src = "hockey hockey hockey hockey hockey hockey hockey hockey";
6     char dest[20];
7
8     int len;
9
10    strcpy(dest, "I like "); // dest is now "I like "
11
12    len = strlen(dest);
13
14    // tricky, but let's use some pointer arithmetic and math to append
15    // as much of src as possible onto the end of dest, -1 on the length to
16    // leave room for the terminator:
17    strncpy(dest+len, src, sizeof(dest)-len-1);
18
19    // remember that sizeof() returns the size of the array in bytes
20    // and a char is a byte:
21    dest[sizeof(dest)-1] = '\0'; // terminate
22
23    // dest is now:          v null terminator
24    // I like hockey hocke
25    // 01234567890123456789012345
26 }

```

## See Also

`memcpy()`, `strcat()`, `strncat()`

---

## 26.3 `strdup()`, `strndup()`

Duplicate a string on the heap

### Synopsis

New in C23!

```

#include <string.h>

char *strdup(const char *s);

char *strndup(const char *s, size_t size);

```

### Description

This is like `strcpy()`, except it allocates space for the new string, just as if it had been done with `malloc()`. The pointer to the duplicate string it returns should be freed up with a call to `free()` once you're done with it.

`strdup()` works the same way, except you can limit the number of bytes duplicated. `strndup()` always produces a NUL-terminated string.

## Return Value

Both functions return a pointer to the newly created string, or `NULL` if the memory for the duplicate cannot be allocated.

## Example

Here's an example that duplicates a string and then capitalizes the first letter of the duplicate.

```
#include <string.h>
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char *dup_cap(const char *s)
{
    char *d = strdup(s);

    if (d != NULL)
        d[0] = toupper(d[0]);

    return d;
}

int main(void)
{
    char *s = dup_cap("hello, world!");

    if (s != NULL)
        puts(s); // Hello, world!
    else
        puts("Error duplicating string");

    free(s);
}
```

## See Also

`strcpy()`, `strncpy()`, `malloc()`, `free()`

---

## 26.4 `strcat()`, `strncat()`

Concatenate two strings into a single string

### Synopsis

```
#include <string.h>

int strcat(const char *dest, const char *src);
```

```
int strncat(const char *dest, const char *src, size_t n);
```

## Description

“Concatenate”, for those not in the know, means to “stick together”. These functions take two strings, and stick them together, storing the result in the first string.

These functions don’t take the size of the first string into account when it does the concatenation. What this means in practical terms is that you can try to stick a 2 megabyte string into a 10 byte space. This will lead to unintended consequences, unless you intended to lead to unintended consequences, in which case it will lead to intended unintended consequences.

Technical banter aside, your boss and/or professor will be irate.

If you want to make sure you don’t overrun the first string, be sure to check the lengths of the strings first and use some highly technical subtraction to make sure things fit.

You can actually only concatenate the first `n` characters of the second string by using `strncat()` and specifying the maximum number of characters to copy.

## Return Value

Both functions return a pointer to the destination string, like most of the string-oriented functions.

## Example

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      char dest[30] = "Hello";
7      char *src = ", World!";
8      char numbers[] = "12345678";
9
10     printf("dest before strcat: \"%s\"\n", dest); // "Hello"
11
12     strcat(dest, src);
13     printf("dest after strcat:  \"%s\"\n", dest); // "Hello, world!"
14
15     strncat(dest, numbers, 3); // strcat first 3 chars of numbers
16     printf("dest after strncat: \"%s\"\n", dest); // "Hello, world!123"
17 }
```

Notice I mixed and matched pointer and array notation there with `src` and `numbers`; this is just fine with string functions.

## See Also

`strlen()`

---

## 26.5 strcmp(), strncmp(), memcmp()

Compare two strings or memory regions and return a difference

### Synopsis

```
#include <string.h>

int strcmp(const char *s1, const char *s2);

int strncmp(const char *s1, const char *s2, size_t n);

int memcmp(const void *s1, const void *s2, size_t n);
```

### Description

All these functions compare chunks of bytes in memory.

`strcmp()` and `strncmp()` operate on NUL-terminated strings, whereas `memcmp()` will compare the number of bytes you specify, brazenly ignoring any NUL characters it finds along the way.

`strcmp()` compares the entire string down to the end, while `strncmp()` only compares the first `n` characters of the strings.

It's a little funky what they return. Basically it's a difference of the strings, so if the strings are the same, it'll return zero (since the difference is zero). It'll return non-zero if the strings differ; basically it will find the first mismatched character and return less-than zero if that character in `s1` is less than the corresponding character in `s2`. It'll return greater-than zero if that character in `s1` is greater than that in `s2`.

So if they return `0`, the comparison was equal (i.e. the difference was `0`.)

These functions can be used as comparison functions for `qsort()` if you have an array of `char*`s you want to sort.

### Return Value

Returns zero if the strings or memory are the same, less-than zero if the first different character in `s1` is less than that in `s2`, or greater-than zero if the first difference character in `s1` is greater than than in `s2`.

### Example

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      char *s1 = "Muffin";
7      char *s2 = "Muffin Sandwich";
8      char *s3 = "Muffin";
9
10     int r1 = strcmp("Biscuits", "Kittens");
11     printf("%d\n", r1); // prints < 0 since 'B' < 'K'
12
13     int r2 = strcmp("Kittens", "Biscuits");
```



```

14     printf("%d\n", r2); // prints > 0 since 'K' > 'B'
15
16     if (strcmp(s1, s2) == 0)
17         printf("This won't get printed because the strings differ\n");
18
19     if (strcmp(s1, s3) == 0)
20         printf("This will print because s1 and s3 are the same\n");
21
22     // this is a little weird...but if the strings are the same, it'll
23     // return zero, which can also be thought of as "false". Not-false
24     // is "true", so (!strcmp()) will be true if the strings are the
25     // same. yes, it's odd, but you see this all the time in the wild
26     // so you might as well get used to it:
27
28     if (!strcmp(s1, s3))
29         printf("The strings are the same!\n");
30
31     if (!strncmp(s1, s2, 6))
32         printf("The first 6 characters of s1 and s2 are the same\n");
33 }

```

**See Also**

`memcmp()`, `qsort()`

**26.6 strcoll()**

Compare two strings accounting for locale

**Synopsis**

```

#include <string.h>

int strcoll(const char *s1, const char *s2);

```

**Description**

This is basically `strcmp()`, except that it handles accented characters better depending on the locale.

For example, my `strcmp()` reports that the character “é” (with accent) is greater than “f”. But that’s hardly useful for alphabetizing.

By setting the `LC_COLLATE` locale value (either by name or via `LC_ALL`), you can have `strcoll()` sort in a way that’s more meaningful by the current locale. For example, by having “é” appear sanely *before* “f”.

It’s also a lot slower than `strcmp()` so use it only if you have to. See `strxfrm()` for a potential speedup.

**Return Value**

Like the other string comparison functions, `strcoll()` returns a negative value if `s1` is less than `s2`, or a positive value if `s1` is greater than `s2`. Or 0 if they are equal.

### Example

```

1  #include <stdio.h>
2  #include <string.h>
3  #include <locale.h>
4
5  int main(void)
6  {
7      setlocale(LC_ALL, "");
8
9      // If your source character set doesn't support "é" in a string
10     // you can replace it with `        `, the Unicode code point
11     // for "  ".
12
13     printf("%d  n", strcmp("  ", "f"));    // Reports    > f, yuck.
14     printf("%d  n", strcoll("  ", "f"));    // Reports    < f, yay!
15 }

```

## See Also

strcmp()

## 26.7 strxfrm()

## Transform a string for comparing based on locale

## Synopsis

```
#include <string.h>

size_t strxfrm(char * restrict s1, const char * restrict s2, size_t n);
```

## Description

This is a strange little function, so bear with me.

Firstly, if you haven't done so, get familiar with `strcoll()` because this is closely related to that.

OK! Now that you're back, you can think of `strxfrm()` as the first part of the `strcoll()` internals. Basically, `strcoll()` has to transform a string into a form that can be compared with `strcmp()`. And it does this with `strxfrm()` for both strings every time you call it.

`strxform()` takes string `s2` and transforms it (readies it for `strcmp()`) storing the result in `s1`. It writes no more than `n` bytes, protecting us from terrible buffer overflows.

But hang on—there’s another mode! If you pass `NULL` for `s1` and `0` for `n`, it will return the number of bytes that the transformed string *would have used*<sup>1</sup>. This is useful if you need to allocate some space to hold the transformed string before you `strcmp()` it against another.

<sup>1</sup>It always returns the number of bytes the transformed string took, but in this case because `s1` was `NULL`, it doesn't actually write a transformed string.

What I'm getting at, not to be too blunt, is that `strcoll()` is slow compared to `strcmp()`. It does a lot of extra work running `strxfrm()` on all its strings.

In fact, we can see how it works by writing our own like this:

```

1 int my_strcoll(char *s1, char *s2)
2 {
3     // Use n = 0 to just get the lengths of the transformed strings
4     int len1 = strxfrm(NULL, s1, 0) + 1;
5     int len2 = strxfrm(NULL, s2, 0) + 1;
6
7     // Allocate enough room for each
8     char *d1 = malloc(len1);
9     char *d2 = malloc(len2);
10
11    // Transform the strings for comparison
12    strxfrm(d1, s1, len1);
13    strxfrm(d2, s2, len2);
14
15    // Compare the transformed strings
16    int result = strcmp(d1, d2);
17
18    // Free up the transformed strings
19    free(d2);
20    free(d1);
21
22    return result;
23 }
```

You see on lines 12, 13, and 16, above how we transform the two input strings and then call `strcmp()` on the result.

So why do we have this function? Can't we just call `strcoll()` and be done with it?

The idea is that if you have one string that you're going to be comparing against a whole lot of other ones, maybe you just want to transform that string one time, then use the faster `strcmp()` saving yourself a bunch of the work we had to do in the function, above.

We'll do that in the example.

## Return Value

Returns the number of bytes in the transformed sequence. If the value is greater than `n`, the results in `s1` are meaningless.

## Example

```

1 #include <stdio.h>
2 #include <string.h>
3 #include <locale.h>
4 #include <stdlib.h>
5
6 // Transform a string for comparison, returning a malloc'd
7 // result
```

```

8 char *get_xfrm_str(char *s)
9 {
10     int len = strxfrm(NULL, s, 0) + 1;
11     char *d = malloc(len);
12
13     strxfrm(d, s, len);
14
15     return d;
16 }
17
18 // Does half the work of a regular strcoll() because the second
19 // string arrives already transformed.
20 int half_strcoll(char *s1, char *s2_transformed)
21 {
22     char *s1_transformed = get_xfrm_str(s1);
23
24     int result = strcmp(s1_transformed, s2_transformed);
25
26     free(s1_transformed);
27
28     return result;
29 }
30
31 int main(void)
32 {
33     setlocale(LC_ALL, "");
34
35     // Pre-transform the string to compare against
36     char *s = get_xfrm_str("éfg");
37
38     // Repeatedly compare against "éfg"
39     printf("%d\n", half_strcoll("fgh", s)); // "fgh" > "éfg"
40     printf("%d\n", half_strcoll("àbc", s)); // "àbc" < "éfg"
41     printf("%d\n", half_strcoll("ñij", s)); // "ñij" > "éfg"
42
43     free(s);
44 }

```

**See Also**

strcoll()

**26.8 strchr(), strrchr(), memchr()**

Find a character in a string

## Synopsis

```
// Pre-C23:

#include <string.h>

char *strchr(char *str, int c);

char *strrchr(char *str, int c);

void *memchr(const void *s, int c, size_t n);

// C23:

QChar *strchr(QChar *s, int c);

QChar *strrchr(QChar *s, int c);

QVoid *memchr(QVoid *s, int c, size_t n);
```

## Description

The functions `strchr()` and `strrchr()` find the first or last occurrence of a letter in a string, respectively. (The extra “r” in `strrchr()` stands for “reverse”—it looks starting at the end of the string and working backward.) Each function returns a pointer to the char in question, or `NULL` if the letter isn’t found in the string.

`memchr()` is similar, except that instead of stopping on the first NUL character, it continues searching for however many bytes you specify.

Quite straightforward.

One thing you can do if you want to find the next occurrence of the letter after finding the first, is call the function again with the previous return value plus one. (Remember pointer arithmetic?) Or minus one if you’re looking in reverse. Don’t accidentally go off the end of the string!

## Return Value

Returns a pointer to the occurrence of the letter in the string, or `NULL` if the letter is not found.

## Example

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      // "Hello, world!"
7      //      ^  ^  ^
8      //      A  B  C
9
10     char *str = "Hello, world!";
11     char *p;
```

```

12
13     p = strchr(str, ',');           // p now points at position A
14     p = strchr(str, 'o');           // p now points at position B
15
16     p = memchr(str, '!', 13);       // p now points at position C
17
18     // repeatedly find all occurrences of the letter 'B'
19     str = "A BIG BROWN BAT BIT BEEJ";
20
21     for(p = strchr(str, 'B'); p != NULL; p = strchr(p + 1, 'B')) {
22         printf("Found a 'B' here: %s\n", p);
23     }
24 }

```

Output:

```

Found a 'B' here: BIG BROWN BAT BIT BEEJ
Found a 'B' here: BROWN BAT BIT BEEJ
Found a 'B' here: BAT BIT BEEJ
Found a 'B' here: BIT BEEJ
Found a 'B' here: BEEJ

```

## 26.9 strspn(), strcspn()

Return the length of a string consisting entirely of a set of characters, or of not a set of characters

### Synopsis

```

#include <string.h>

size_t strspn(char *str, const char *accept);

size_t strcspn(char *str, const char *reject);

```

### Description

`strspn()` will tell you the length of a string consisting entirely of the set of characters in `accept`. That is, it starts walking down `str` until it finds a character that is *not* in the set (that is, a character that is not to be accepted), and returns the length of the string so far.

`strcspn()` works much the same way, except that it walks down `str` until it finds a character in the `reject` set (that is, a character that is to be rejected.) It then returns the length of the string so far.

### Return Value

The length of the string consisting of all characters in `accept` (for `strspn()`), or the length of the string consisting of all characters except `reject` (for `strcspn()`).

**Example**

```

1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      char str1[] = "a banana";
7      char str2[] = "the bolivian navy on maenuvers in the south pacific";
8      int n;
9
10     // how many letters in str1 until we reach something that's not a vowel?
11     n = strspn(str1, "aeiou");
12     printf("%d\n", n); // n == 1, just "a"
13
14     // how many letters in str1 until we reach something that's not a, b,
15     // or space?
16     n = strspn(str1, "ab ");
17     printf("%d\n", n); // n == 4, "a ba"
18
19     // how many letters in str2 before we get a "y"?
20     n = strcspn(str2, "y");
21     printf("%d\n", n); // n = 16, "the bolivian nav"
22 }

```

**See Also**

strchr(), strchr()

---

**26.10 strpbrk()**

Search a string for one of a set of characters

**Synopsis**

```

#include <string.h>

// Pre-C23:
char *strpbrk(const char *s1, const char *s2);

// C23:
QChar *strpbrk(QChar *s1, const char *s2);

```

**Description**

This function searches string `s1` for any of the characters that are found in string `s2`.

It's just like how `strchr()` searches for a specific character in a string, except it will match *any* of the characters found in `s2`.

Think of the power!

## Return Value

Returns a pointer to the first character matched in `s1`, or NULL if the string isn't found.

## Example

```

1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      // p points here after strpbrk
7      //           v
8      char *s1 = "Hello, world!";
9      char *s2 = "dow!"; // Match any of these chars
10
11     char *p = strpbrk(s1, s2); // p points to the o
12
13     printf("%s\n", p); // "o, world!"
14 }
```

## See Also

`strchr()`, `memchr()`

---

## 26.11 `strstr()`

Find a string in another string

## Synopsis

```

#include <string.h>

// Pre-C23:
char *strstr(const char *str, const char *substr);

// C23:
QChar *strstr(QChar *s1, const char *s2);
```

## Description

Let's say you have a big long string, and you want to find a word, or whatever substring strikes your fancy, inside the first string. Then `strstr()` is for you! It'll return a pointer to the `substr` within the `str`!



## Return Value

You get back a pointer to the occurrence of the `substr` inside the `str`, or `NULL` if the substring can't be found.

## Example

```

1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      char *str = "The quick brown fox jumped over the lazy dogs.";
7      char *p;
8
9      p = strstr(str, "lazy");
10     printf("%s\n", p == NULL? "null": p); // "lazy dogs."
11
12     // p is NULL after this, since the string "wombat" isn't in str:
13     p = strstr(str, "wombat");
14     printf("%s\n", p == NULL? "null": p); // "null"
15 }
```

## See Also

`strchr()`, `strrchr()`, `strspn()`, `strcspn()`

## 26.12 strtok()

Tokenize a string

### Synopsis

```

#include <string.h>

char *strtok(char *str, const char *delim);
```

### Description

If you have a string that has a bunch of separators in it, and you want to break that string up into individual pieces, this function can do it for you.

The usage is a little bit weird, but at least whenever you see the function in the wild, it's consistently weird.

Basically, the first time you call it, you pass the string, `str` that you want to break up in as the first argument. For each subsequent call to get more tokens out of the string, you pass `NULL`. This is a little weird, but `strtok()` remembers the string you originally passed in, and continues to strip tokens off for you.

Note that it does this by actually putting a NUL terminator after the token, and then returning a pointer to the start of the token. So the original string you pass in is destroyed, as it were. If you need to preserve the string, be sure to pass a copy of it to `strtok()` so the original isn't destroyed.

## Return Value

A pointer to the next token. If you're out of tokens, `NULL` is returned.

## Example

```

1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      // break up the string into a series of space or
7      // punctuation-separated words
8      char str[] = "Where is my bacon, dude?";
9      char *token;
10
11     // Note that the following if-do-while construct is very very
12     // very very very common to see when using strtok().
13
14     // grab the first token (making sure there is a first token!)
15     if ((token = strtok(str, ".,?! ")) != NULL) {
16         do {
17             printf("Word: \"%s\"\n", token);
18
19             // now, the while continuation condition grabs the
20             // next token (by passing NULL as the first param)
21             // and continues if the token's not NULL:
22         } while ((token = strtok(NULL, ".,?! ")) != NULL);
23     }
24 }
```

Output:

```

Word: "Where"
Word: "is"
Word: "my"
Word: "bacon"
Word: "dude"
```

## See Also

`strchr()`, `strrchr()`, `strspn()`, `strcspn()`

## 26.13 `memset()`, `memset_explicit`

Set a region of memory to a certain value

### Synopsis

`memset_explicit()` is new in C23!

```
#include <string.h>

void *memset(void *s, int c, size_t n);

void *memset_explicit(void *s, int c, size_t n);
```

## Description

This function is what you use to set a region of memory to a particular value, namely `c` converted into `unsigned char`.

The most common usage is to zero out an array or `struct`.

`memset_explicit()` only differs in that it will never be optimized away (like `memset()` might be). The idea is that you could use it to most-definitely remove sensitive information (like passwords) from memory before any nasty hackers get their hands on it.

## Return Value

`memset()` and `memset_explicit()` return whatever you passed in as `s` for happy convenience.

## Example

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      struct banana {
7          float ripeness;
8          char *peel_color;
9          int grams;
10     };
11
12     struct banana b;
13
14     memset(&b, 0, sizeof b);
15
16     printf("%d\n", b.ripeness == 0.0);    // True
17     printf("%d\n", b.peel_color == NULL); // True
18     printf("%d\n", b.grams == 0);        // True
19 }
```

## See Also

`memcpy()`, `memmove()`

---

## 26.14 `strerror()`

Get a string version of an error number

## Synopsis

```
#include <string.h>

char *strerror(int errnum);
```

## Description

This function ties closely into `perror()` (which prints a human-readable error message corresponding to `errno`). But instead of printing, `strerror()` returns a pointer to the locale-specific error message string.

So if you ever need that string back for some reason (e.g. you're going to `fprintf()` it to a file or something), this function will give it to you. All you need to do is pass in `errno` as an argument. (Recall that `errno` gets set as an error status by a variety of functions.)

You can actually pass in any integer for `errnum` you want. The function will return *some* message, even if the number doesn't correspond to any known value for `errno`.

The values of `errno` and the strings returned by `strerror()` are system-dependent.

## Return Value

A string error message corresponding to the given error number.

You are not allowed to modify the returned string.

## Example

```
1  #include <stdio.h>
2  #include <string.h>
3  #include <errno.h>
4
5  int main(void)
6  {
7      FILE *fp = fopen("NONEXISTENT_FILE.TXT", "r");
8
9      if (fp == NULL) {
10         char *errmsg = strerror(errno);
11         printf("Error %d opening file: %s\n", errno, errmsg);
12     }
13 }
```

Output:

```
Error 2 opening file: No such file or directory
```

## See Also

`perror()`

---

## 26.15 `strlen()`

Returns the length of a string

### Synopsis

```
#include <string.h>

size_t strlen(const char *s);
```

### Description

This function returns the length of the passed null-terminated string (not counting the NUL character at the end). It does this by walking down the string and counting the bytes until the NUL character, so it's a little time consuming. If you have to get the length of the same string repeatedly, save it off in a variable somewhere.

### Return Value

Returns the number of bytes in the string. Note that this might be different than the number of characters in a multibyte string.

### Example

```
1  #include <stdio.h>
2  #include <string.h>
3
4  int main(void)
5  {
6      char *s = "Hello, world!"; // 13 characters
7
8      // prints "The string is 13 characters long.":
9
10     printf("The string is %zu characters long.\n", strlen(s));
11 }
```

### See Also

## Chapter 27

# <tgmath.h> Type-Generic Math Functions

These are type-generic macros that are wrappers around the math functions in `<math.h>` and `<complex.h>`. This header includes both of those.

But on the surface, you can think of them as being able to use, say, the `sqrt()` function with any type without needed to think about if it's `double` or `long double` or even `complex`.

These are the defined macros—some of them don't have a counterpart in the real or complex space. Type suffixes are omitted in the table on the Real and Complex columns. None of the generic macros have type suffixes.

| Real Function         | Complex Function    | Generic Macro         |
|-----------------------|---------------------|-----------------------|
| <code>acos</code>     | <code>cacos</code>  | <code>acos</code>     |
| <code>asin</code>     | <code>casin</code>  | <code>asin</code>     |
| <code>atan</code>     | <code>catan</code>  | <code>atan</code>     |
| <code>acosh</code>    | <code>cacosh</code> | <code>acosh</code>    |
| <code>asinh</code>    | <code>casinh</code> | <code>asinh</code>    |
| <code>atanh</code>    | <code>catanh</code> | <code>atanh</code>    |
| <code>cos</code>      | <code>ccos</code>   | <code>cos</code>      |
| <code>sin</code>      | <code>csin</code>   | <code>sin</code>      |
| <code>tan</code>      | <code>ctan</code>   | <code>tan</code>      |
| <code>cosh</code>     | <code>ccosh</code>  | <code>cosh</code>     |
| <code>sinh</code>     | <code>csinh</code>  | <code>sinh</code>     |
| <code>tanh</code>     | <code>ctanh</code>  | <code>tanh</code>     |
| <code>exp</code>      | <code>cexp</code>   | <code>exp</code>      |
| <code>log</code>      | <code>clog</code>   | <code>log</code>      |
| <code>pow</code>      | <code>cpow</code>   | <code>pow</code>      |
| <code>sqrt</code>     | <code>csqrt</code>  | <code>sqrt</code>     |
| <code>fabs</code>     | <code>cabs</code>   | <code>fabs</code>     |
| <code>atan2</code>    | —                   | <code>atan2</code>    |
| <code>fdim</code>     | —                   | <code>fdim</code>     |
| <code>cbrt</code>     | —                   | <code>cbrt</code>     |
| <code>floor</code>    | —                   | <code>floor</code>    |
| <code>ceil</code>     | —                   | <code>ceil</code>     |
| <code>fma</code>      | —                   | <code>fma</code>      |
| <code>copysign</code> | —                   | <code>copysign</code> |
| <code>fmax</code>     | —                   | <code>fmax</code>     |

| Real Function | Complex Function | Generic Macro |
|---------------|------------------|---------------|
| erf           | —                | erf           |
| fmin          | —                | fmin          |
| erfc          | —                | erfc          |
| fmod          | —                | fmod          |
| exp2          | —                | exp2          |
| frexp         | —                | frexp         |
| expm1         | —                | expm1         |
| hypot         | —                | hypot         |
| ilogb         | —                | ilogb         |
| ldexp         | —                | ldexp         |
| lgamma        | —                | lgamma        |
| llrint        | —                | llrint        |
| llround       | —                | llround       |
| log10         | —                | log10         |
| log1p         | —                | log1p         |
| log2          | —                | log2          |
| logb          | —                | logb          |
| lrint         | —                | lrint         |
| lround        | —                | lround        |
| nearbyint     | —                | nearbyint     |
| nextafter     | —                | nextafter     |
| nexttoward    | —                | nexttoward    |
| remainder     | —                | remainder     |
| remquo        | —                | remquo        |
| rint          | —                | rint          |
| round         | —                | round         |
| scalbn        | —                | scalbn        |
| scalbln       | —                | scalbln       |
| tgamma        | —                | tgamma        |
| trunc         | —                | trunc         |
| —             | carg             | carg          |
| —             | cimag            | cimag         |
| —             | conj             | conj          |
| —             | cproj            | cproj         |
| —             | creal            | creal         |

## 27.1 Example

Here's an example where we call the type-generic `sqrt()` function on a variety of types.

```

1  #include <stdio.h>
2  #include <tgmath.h>
3
4  int main(void)
5  {
6      double x = 12.8;
7      long double y = 34.9;
8      double complex z = 1 + 2 * I;
9
10     double x_result;
```

```
11     long double y_result;  
12     double complex z_result;  
13  
14     // We call the same sqrt() function--it's type-generic!  
15     x_result = sqrt(x);  
16     y_result = sqrt(y);  
17     z_result = sqrt(z);  
18  
19     printf("x_result: %f\n", x_result);  
20     printf("y_result: %Lf\n", y_result);  
21     printf("z_result: %f + %fi\n", creal(z_result), cimag(z_result));  
22 }
```

Output:

```
x_result: 3.577709  
y_result: 5.907622  
z_result: 1.272020 + 0.786151i
```



## Chapter 28

# <threads.h> Multithreading Functions

| Function                     | Description                                              |
|------------------------------|----------------------------------------------------------|
| <code>call_once()</code>     | Call a function one time no matter how many threads try  |
| <code>cnd_broadcast()</code> | Wake up all threads waiting on a condition variable      |
| <code>cnd_destroy()</code>   | Free up resources from a condition variable              |
| <code>cnd_init()</code>      | Initialize a condition variable to make it ready for use |
| <code>cnd_signal()</code>    | Wake up a thread waiting on a condition variable         |
| <code>cnd_timedwait()</code> | Wait on a condition variable with a timeout              |
| <code>cnd_wait()</code>      | Wait for a signal on a condition variable                |
| <code>mtx_destroy()</code>   | Cleanup a mutex when done with it                        |
| <code>mtx_init()</code>      | Initialize a mutex for use                               |
| <code>mtx_lock()</code>      | Acquire a lock on a mutex                                |
| <code>mtx_timedlock()</code> | Lock a mutex allowing for timeout                        |
| <code>mtx_trylock()</code>   | Try to lock a mutex, returning if not possible           |
| <code>mtx_unlock()</code>    | Free a mutex when you're done with the critical section  |
| <code>thrd_create()</code>   | Create a new thread of execution                         |
| <code>thrd_current()</code>  | Get the ID of the calling thread                         |
| <code>thrd_detach()</code>   | Automatically clean up threads when they exit            |
| <code>thrd_equal()</code>    | Compare two thread descriptors for equality              |
| <code>thrd_exit()</code>     | Stop and exit this thread                                |
| <code>thrd_join()</code>     | Wait for a thread to exit                                |
| <code>thrd_yield()</code>    | Stop running that other threads might run                |
| <code>tss_create()</code>    | Create new thread-specific storage                       |
| <code>tss_delete()</code>    | Clean up a thread-specific storage variable              |
| <code>tss_get()</code>       | Get thread-specific data                                 |
| <code>tss_set()</code>       | Set thread-specific data                                 |

We have a bunch of good things at our disposal with this one:

- Threads
- Mutexes
- Condition Variables
- Thread-Specific Storage
- And, last but not least, the always-fun `call_once()` function!

Enjoy!

---

## 28.1 call\_once()

Call a function one time no matter how many threads try

### Synopsis

```
#include <threads.h>

void call_once(once_flag *flag, void (*func)(void));
```

### Description

If you have a bunch of threads running over the same piece of code that calls a function, but you only want that function to run one time, `call_once()` can help you out.

The catch is the function that is called doesn't return anything and takes no arguments.

If you need more than that, you'll have to set a threadsafe flag such as `atomic_flag`, or one that you protect with a mutex.

To use this, you need to pass it a pointer to a function to execute, `func`, and also a pointer to a flag of type `once_flag`.

`once_flag` is an opaque type, so all you need to know is that you initialize it to the value `ONCE_FLAG_INIT`.

### Return Value

Returns nothing.

### Example

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  once_flag of = ONCE_FLAG_INIT; // Initialize it like this
5
6  void run_once_function(void)
7  {
8      printf("I'll only run once!\n");
9  }
10
11 int run(void *arg)
12 {
13     (void)arg;
14
15     printf("Thread running!\n");
16
17     call_once(&of, run_once_function);
18
```

```

19     return 0;
20 }
21
22 #define THREAD_COUNT 5
23
24 int main(void)
25 {
26     thrd_t t[THREAD_COUNT];
27
28     for (int i = 0; i < THREAD_COUNT; i++)
29         thrd_create(t + i, run, NULL);
30
31     for (int i = 0; i < THREAD_COUNT; i++)
32         thrd_join(t[i], NULL);
33 }

```

Output (might vary per run):

```

Thread running!
Thread running!
I'll only run once!
Thread running!
Thread running!
Thread running!

```

## 28.2 cnd\_broadcast()

Wake up all threads waiting on a condition variable

### Synopsis

```

#include <threads.h>

int cnd_broadcast(cnd_t *cond);

```

### Description

This is just like `cnd_signal()` in that it wakes up threads that are waiting on a condition variable.... except instead of just rousing one thread, it wakes them all.

Of course, only one will get the mutex, and the rest will have to wait their turn. But instead of being asleep waiting for a signal, they'll be asleep waiting to reacquire the mutex. They're rearin' to go, in other words.

This can make a difference in a specific set of circumstances where `cnd_signal()` might leave you hanging.

If you're relying on subsequent threads to issue the next `cnd_signal()`, but you have the `cnd_wait()` in a `while` loop<sup>1</sup> that doesn't allow any threads to escape, you'll be stuck. No more threads will be woken up from the wait.

---

<sup>1</sup>Which you should because of spurious wakeups.

But if you `cnd_broadcast()`, all the threads will be woken, and presumably at least one of them will be allowed to escape the `while` loop, freeing it up to broadcast the next wakeup when its work is done.

## Return Value

Returns `thrd_success` or `thrd_error` depending on how well things went.

## Example

In the example below, we launch a bunch of threads, but they're only allowed to run if their ID matches the current ID. If it doesn't, they go back to waiting.

If you `cnd_signal()` to wake the next thread, it might not be the one with the proper ID to run. If it's not, it goes back to sleep and we hang (because no thread is awake to hit `cnd_signal()` again).

But if you `cnd_broadcast()` to wake them all, then they'll all try (one after another) to get out of the `while` loop. And one of them will make it.

Try switching the `cnd_broadcast()` to `cnd_signal()` to see likely deadlocks. It doesn't happen every time, but usually does.

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      int id = *(int*)arg;
10
11      static int current_id = 0;
12
13      mtx_lock(&mutex);
14
15      while (id != current_id) {
16          printf("THREAD %d: waiting\n", id);
17          cnd_wait(&condvar, &mutex);
18
19          if (id != current_id)
20              printf("THREAD %d: woke up, but it's not my turn!\n", id);
21          else
22              printf("THREAD %d: woke up, my turn! Let's go!\n", id);
23      }
24
25      current_id++;
26
27      printf("THREAD %d: signaling thread %d to run\n", id, current_id);
28
29      //cnd_signal(&condvar);
30      cnd_broadcast(&condvar);
31      mtx_unlock(&mutex);
32
33      return 0;
34  }
```

```

35
36 #define THREAD_COUNT 5
37
38 int main(void)
39 {
40     thrd_t t[THREAD_COUNT];
41     int id[] = {4, 3, 2, 1, 0};
42
43     mtx_init(&mutex, mtx_plain);
44     cnd_init(&condvar);
45
46     for (int i = 0; i < THREAD_COUNT; i++)
47         thrd_create(t + i, run, id + i);
48
49     for (int i = 0; i < THREAD_COUNT; i++)
50         thrd_join(t[i], NULL);
51
52     mtx_destroy(&mutex);
53     cnd_destroy(&condvar);
54 }

```

Example run with `cnd_broadcast()`:

```

THREAD 4: waiting
THREAD 1: waiting
THREAD 3: waiting
THREAD 2: waiting
THREAD 0: signaling thread 1 to run
THREAD 2: woke up, but it's not my turn!
THREAD 2: waiting
THREAD 4: woke up, but it's not my turn!
THREAD 4: waiting
THREAD 3: woke up, but it's not my turn!
THREAD 3: waiting
THREAD 1: woke up, my turn! Let's go!
THREAD 1: signaling thread 2 to run
THREAD 4: woke up, but it's not my turn!
THREAD 4: waiting
THREAD 3: woke up, but it's not my turn!
THREAD 3: waiting
THREAD 2: woke up, my turn! Let's go!
THREAD 2: signaling thread 3 to run
THREAD 4: woke up, but it's not my turn!
THREAD 4: waiting
THREAD 3: woke up, my turn! Let's go!
THREAD 3: signaling thread 4 to run
THREAD 4: woke up, my turn! Let's go!
THREAD 4: signaling thread 5 to run

```

Example run with `cnd_signal()`:

```

THREAD 4: waiting
THREAD 1: waiting
THREAD 3: waiting
THREAD 2: waiting
THREAD 0: signaling thread 1 to run
THREAD 4: woke up, but it's not my turn!
THREAD 4: waiting

[deadlock at this point]

```

See how `THREAD 0` signaled that it was `THREAD 1`'s turn? But—bad news—it was `THREAD 4` that got woken up. So no one continued the process. `cnd_broadcast()` would have woken them all, so eventually `THREAD 1` would have run, gotten out of the `while`, and broadcast for the next thread to run.

## See Also

`cnd_signal()`, `mtx_lock()`, `mtx_unlock()`

---

## 28.3 `cnd_destroy()`

Free up resources from a condition variable

### Synopsis

```

#include <threads.h>

void cnd_destroy(cnd_t *cond);

```

### Description

This is the opposite of `cnd_init()` and should be called when all threads are done using a condition variable.

### Return Value

Returns nothing!

### Example

General-purpose condition variable example here, but you can see the `cnd_destroy()` down at the end.

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10

```

```

11     mtx_lock(&mutex);
12
13     printf("Thread: waiting...\n");
14     cnd_wait(&condvar, &mutex);
15     printf("Thread: running again!\n");
16
17     mtx_unlock(&mutex);
18
19     return 0;
20 }
21
22 int main(void)
23 {
24     thrd_t t;
25
26     mtx_init(&mutex, mtx_plain);
27     cnd_init(&condvar);
28
29     printf("Main creating thread\n");
30     thrd_create(&t, run, NULL);
31
32     // Sleep 0.1s to allow the other thread to wait
33     thrd_sleep(&(struct timespec){.tv_nsec=100000000L}, NULL);
34
35     mtx_lock(&mutex);
36     printf("Main: signaling thread\n");
37     cnd_signal(&condvar);
38     mtx_unlock(&mutex);
39
40     thrd_join(t, NULL);
41
42     mtx_destroy(&mutex);
43     cnd_destroy(&condvar); // <-- DESTROY CONDITION VARIABLE
44 }

```

Output:

```

Main creating thread
Thread: waiting...
Main: signaling thread
Thread: running again!

```

## See Also

`cnd_init()`

## 28.4 `cnd_init()`

Initialize a condition variable to make it ready for use

## Synopsis

```
#include <threads.h>

int cnd_init(cnd_t *cond);
```

## Description

This is the opposite of `cnd_destroy()`. This prepares a condition variable for use, doing behind-the-scenes work on it.

Don't use a condition variable without calling this first!

## Return Value

If all goes well, returns `thrd_success`. If all doesn't go well, it could return `thrd_nomem` if the system is out of memory, or `thread_error` in the case of any other error.

## Example

General-purpose condition variable example here, but you can see the `cnd_init()` down at the start of `main()`.

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10
11     mtx_lock(&mutex);
12
13     printf("Thread: waiting...\n");
14     cnd_wait(&condvar, &mutex);
15     printf("Thread: running again!\n");
16
17     mtx_unlock(&mutex);
18
19     return 0;
20 }
21
22 int main(void)
23 {
24     thrd_t t;
25
26     mtx_init(&mutex, mtx_plain);
27     cnd_init(&condvar);      // <-- INITIALIZE CONDITION VARIABLE
28
29     printf("Main creating thread\n");
30     thrd_create(&t, run, NULL);
```



```
31
32 // Sleep 0.1s to allow the other thread to wait
33 thrd_sleep(&(struct timespec){.tv_nsec=100000000L}, NULL);
34
35 mtx_lock(&mutex);
36 printf("Main: signaling thread\n");
37 cnd_signal(&condvar);
38 mtx_unlock(&mutex);
39
40 thrd_join(t, NULL);
41
42 mtx_destroy(&mutex);
43 cnd_destroy(&condvar);
44 }
```

Output:

```
Main creating thread
Thread: waiting...
Main: signaling thread
Thread: running again!
```

## See Also

`cnd_destroy()`

---

## 28.5 `cnd_signal()`

Wake up a thread waiting on a condition variable

### Synopsis

```
#include <threads.h>

int cnd_signal(cnd_t *cond);
```

### Description

If you have a thread (or a bunch of threads) waiting on a condition variable, this function will wake one of them up to run.

Compare to `cnd_broadcast()` that wakes up all the threads. See the `cnd_broadcast()` page for more information on when you're want to use that versus this.

### Return Value

Returns `thrd_success` or `thrd_error` depending on how happy your program is.

## Example

General-purpose condition variable example here, but you can see the `cnd_signal()` in the middle of `main()`.

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10
11     mtx_lock(&mutex);
12
13     printf("Thread: waiting...\n");
14     cnd_wait(&condvar, &mutex);
15     printf("Thread: running again!\n");
16
17     mtx_unlock(&mutex);
18
19     return 0;
20 }
21
22 int main(void)
23 {
24     thrd_t t;
25
26     mtx_init(&mutex, mtx_plain);
27     cnd_init(&condvar);
28
29     printf("Main creating thread\n");
30     thrd_create(&t, run, NULL);
31
32     // Sleep 0.1s to allow the other thread to wait
33     thrd_sleep(&(struct timespec){.tv_nsec=100000000L}, NULL);
34
35     mtx_lock(&mutex);
36     printf("Main: signaling thread\n");
37     cnd_signal(&condvar);    // <-- SIGNAL CHILD THREAD HERE!
38     mtx_unlock(&mutex);
39
40     thrd_join(t, NULL);
41
42     mtx_destroy(&mutex);
43     cnd_destroy(&condvar);
44 }
```

Output:

```

Main creating thread
Thread: waiting...
Main: signaling thread
Thread: running again!

```

## See Also

`cnd_init()`, `cnd_destroy()`

---

## 28.6 `cnd_timedwait()`

Wait on a condition variable with a timeout

### Synopsis

```

#include <threads.h>

int cnd_timedwait(cnd_t *restrict cond, mtx_t *restrict mtx,
                  const struct timespec *restrict ts);

```

### Description

This is like `cnd_wait()` except we get to specify a timeout, as well.

Note that the thread still must reacquire the mutex to get more work done even after the timeout. The the main difference is that regular `cnd_wait()` will only try to get the mutex after a `cnd_signal()` or `cnd_broadcast()`, whereas `cnd_timedwait()` will do that, too, **and** try to get the mutex after the timeout.

The timeout is specified as an absolute UTC time since Epoch. You can get this with the `timespec_get()` function and then add values on to the result to timeout later than now, as shown in the example.

Beware that you can't have more than 999999999 nanoseconds in the `tv_nsec` field of the `struct timespec`. Mod those so they stay in range.

### Return Value

If the thread wakes up for a non-timeout reason (e.g. signal or broadcast), returns `thrd_success`. If woken up due to timeout, returns `thrd_timedout`. Otherwise returns `thrd_error`.

### Example

This example has a thread wait on a condition variable for a maximum of 1.75 seconds. And it always times out because no one ever sends a signal. Tragic.

```

1  #include <stdio.h>
2  #include <time.h>
3  #include <threads.h>
4
5  cnd_t condvar;
6  mtx_t mutex;
7

```

```

8  int run(void *arg)
9  {
10     (void)arg;
11
12     mtx_lock(&mutex);
13
14     struct timespec ts;
15
16     // Get the time now
17     timespec_get(&ts, TIME_UTC);
18
19     // Add on 1.75 seconds from now
20     ts.tv_sec += 1;
21     ts.tv_nsec += 750000000L;
22
23     // Handle nsec overflow
24     ts.tv_sec += ts.tv_nsec / 1000000000L;
25     ts.tv_nsec = ts.tv_nsec % 1000000000L;
26
27     printf("Thread: waiting...\n");
28     int r = cnd_timedwait(&condvar, &mutex, &ts);
29
30     switch (r) {
31         case thrd_success:
32             printf("Thread: signaled!\n");
33             break;
34
35         case thrd_timedout:
36             printf("Thread: timed out!\n");
37             return 1;
38
39         case thrd_error:
40             printf("Thread: Some kind of error\n");
41             return 2;
42     }
43
44     mtx_unlock(&mutex);
45
46     return 0;
47 }
48
49 int main(void)
50 {
51     thrd_t t;
52
53     mtx_init(&mutex, mtx_plain);
54     cnd_init(&condvar);
55
56     printf("Main creating thread\n");
57     thrd_create(&t, run, NULL);
58
59     // Sleep 3s to allow the other thread to timeout
60     thrd_sleep(&(struct timespec){.tv_sec=3}, NULL);

```

```

61
62     thrd_join(t, NULL);
63
64     mtx_destroy(&mutex);
65     cnd_destroy(&condvar);
66 }

```

Output:

```

Main creating thread
Thread: waiting...
Thread: timed out!

```

## See Also

`cnd_wait()`, `timespec_get()`

---

## 28.7 cnd\_wait()

Wait for a signal on a condition variable

### Synopsis

```

#include <threads.h>

int cnd_wait(cnd_t *cond, mtx_t *mtx);

```

### Description

This puts the calling thread to sleep until it is awakened by a call to `cnd_signal()` or `cnd_broadcast()`.

### Return Value

If everything's fantastic, returns `thrd_success`. Otherwise it returns `thrd_error` to report that something has gone fantastically, horribly awry.

### Example

General-purpose condition variable example here, but you can see the `cnd_wait()` in the `run()` function.

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;

```

```

10
11     mtx_lock(&mutex);
12
13     printf("Thread: waiting...\n");
14     cnd_wait(&condvar, &mutex);    // <-- WAIT HERE!
15     printf("Thread: running again!\n");
16
17     mtx_unlock(&mutex);
18
19     return 0;
20 }
21
22 int main(void)
23 {
24     thrd_t t;
25
26     mtx_init(&mutex, mtx_plain);
27     cnd_init(&condvar);
28
29     printf("Main creating thread\n");
30     thrd_create(&t, run, NULL);
31
32     // Sleep 0.1s to allow the other thread to wait
33     thrd_sleep(&(struct timespec){.tv_nsec=100000000L}, NULL);
34
35     mtx_lock(&mutex);
36     printf("Main: signaling thread\n");
37     cnd_signal(&condvar);    // <-- SIGNAL CHILD THREAD HERE!
38     mtx_unlock(&mutex);
39
40     thrd_join(t, NULL);
41
42     mtx_destroy(&mutex);
43     cnd_destroy(&condvar);
44 }

```

Output:

```

Main creating thread
Thread: waiting...
Main: signaling thread
Thread: running again!

```

## See Also

`cnd_timedwait()`

## 28.8 `mtx_destroy()`

Cleanup a mutex when done with it

## Synopsis

```
#include <threads.h>

void mtx_destroy(mtx_t *mtx);
```

## Description

The opposite of `mtx_init()`, this function frees up any resources associated with the given mutex. You should call this when all threads are done using the mutex.

## Return Value

Returns nothing, the selfish ingrate!

## Example

General-purpose mutex example here, but you can see the `mtx_destroy()` down at the end.

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10
11     static int count = 0;
12
13     mtx_lock(&mutex);
14
15     printf("Thread: I got %d!\n", count);
16     count++;
17
18     mtx_unlock(&mutex);
19
20     return 0;
21 }
22
23 #define THREAD_COUNT 5
24
25 int main(void)
26 {
27     thrd_t t[THREAD_COUNT];
28
29     mtx_init(&mutex, mtx_plain);
30
31     for (int i = 0; i < THREAD_COUNT; i++)
32         thrd_create(t + i, run, NULL);
33
```

```

34     for (int i = 0; i < THREAD_COUNT; i++)
35         thrd_join(t[i], NULL);
36
37     mtx_destroy(&mutex);    // <-- DESTROY THE MUTEX HERE
38 }

```

Output:

```

Thread: I got 0!
Thread: I got 1!
Thread: I got 2!
Thread: I got 3!
Thread: I got 4!

```

## See Also

`mtx_init()`

## 28.9 `mtx_init()`

Initialize a mutex for use

### Synopsis

```

#include <threads.h>

int mtx_init(mtx_t *mtx, int type);

```

### Description

Before you can use a mutex variable, you have to initialize it with this call to get it all prepped and ready to go.

But wait! It's not quite that simple. You have to tell it what `type` of mutex you want to create.

| Type                                 | Description                            |
|--------------------------------------|----------------------------------------|
| <code>mtx_plain</code>               | Regular ol' mutex                      |
| <code>mtx_timed</code>               | Mutex that supports timeouts           |
| <code>mtx_plain mtx_recursive</code> | Recursive mutex                        |
| <code>mtx_timed mtx_recursive</code> | Recursive mutex that supports timeouts |

As you can see, you can make a plain or timed mutex *recursive* by bitwise-ORing the value with `mtx_recursive`.

“Recursive” means that the holder of a lock can call `mtx_lock()` multiple times on the same lock. (They have to unlock it an equal number of times before anyone else can take the mutex.) This might ease coding from time to time, especially if you call a function that needs to lock the mutex when you already hold the mutex.



And the timeout gives a thread a chance to *try* to get the lock for a while, but then bail out if it can't get it in that timeframe. You use the `mtx_timedlock()` function with `mtx_timed` mutexes.

## Return Value

Returns `thrd_success` in a perfect world, and potentially `thrd_error` in an imperfect one.

## Example

General-purpose mutex example here, but you can see the `mtx_init()` down at the top of `main()`:

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10
11     static int count = 0;
12
13     mtx_lock(&mutex);
14
15     printf("Thread: I got %d!\n", count);
16     count++;
17
18     mtx_unlock(&mutex);
19
20     return 0;
21 }
22
23 #define THREAD_COUNT 5
24
25 int main(void)
26 {
27     thrd_t t[THREAD_COUNT];
28
29     mtx_init(&mutex, mtx_plain); // <-- CREATE THE MUTEX HERE
30
31     for (int i = 0; i < THREAD_COUNT; i++)
32         thrd_create(t + i, run, NULL);
33
34     for (int i = 0; i < THREAD_COUNT; i++)
35         thrd_join(t[i], NULL);
36
37     mtx_destroy(&mutex); // <-- DESTROY THE MUTEX HERE
38 }

```

Output:

```
Thread: I got 0!  
Thread: I got 1!  
Thread: I got 2!  
Thread: I got 3!  
Thread: I got 4!
```

## See Also

`mtx_destroy()`

---

## 28.10 `mtx_lock()`

Acquire a lock on a mutex

### Synopsis

```
#include <threads.h>  
  
int mtx_lock(mtx_t *mtx);
```

### Description

If you're a thread and want to enter a critical section, do I have the function for you!

A thread that calls this function will wait until it can acquire the mutex, then it will grab it, wake up, and run!

If the mutex is recursive and is already locked by this thread, it will be locked again and the lock count will increase. If the mutex is not recursive and the thread already holds it, this call will error out.

### Return Value

Returns `thrd_success` on goodness and `thrd_error` on badness.

### Example

General-purpose mutex example here, but you can see the `mtx_lock()` in the `run()` function:

```
1  #include <stdio.h>  
2  #include <threads.h>  
3  
4  cnd_t condvar;  
5  mtx_t mutex;  
6  
7  int run(void *arg)  
8  {  
9      (void)arg;  
10  
11     static int count = 0;  
12  
13     mtx_lock(&mutex); // <-- LOCK HERE
```

```

14
15     printf("Thread: I got %d!\n", count);
16     count++;
17
18     mtx_unlock(&mutex);
19
20     return 0;
21 }
22
23 #define THREAD_COUNT 5
24
25 int main(void)
26 {
27     thrd_t t[THREAD_COUNT];
28
29     mtx_init(&mutex, mtx_plain); // <-- CREATE THE MUTEX HERE
30
31     for (int i = 0; i < THREAD_COUNT; i++)
32         thrd_create(t + i, run, NULL);
33
34     for (int i = 0; i < THREAD_COUNT; i++)
35         thrd_join(t[i], NULL);
36
37     mtx_destroy(&mutex); // <-- DESTROY THE MUTEX HERE
38 }

```

Output:

```

Thread: I got 0!
Thread: I got 1!
Thread: I got 2!
Thread: I got 3!
Thread: I got 4!

```

## See Also

`mtx_unlock()`, `mtx_trylock()`, `mtx_timedlock()`

## 28.11 `mtx_timedlock()`

Lock a mutex allowing for timeout

### Synopsis

```

#include <threads.h>

int mtx_timedlock(mtx_t *restrict mtx, const struct timespec *restrict ts);

```

## Description

This is just like `mtx_lock()` except you can add a timeout if you don't want to wait forever.

The timeout is specified as an absolute UTC time since Epoch. You can get this with the `timespec_get()` function and then add values on to the result to timeout later than now, as shown in the example.

Beware that you can't have more than 999999999 nanoseconds in the `tv_nsec` field of the struct `timespec`. Mod those so they stay in range.

## Return Value

If everything works and the mutex is obtained, returns `thrd_success`. If a timeout happens first, returns `thrd_timedout`.

Otherwise, returns `thrd_error`. Because if nothing is right, everything is wrong.

## Example

This example has a thread wait on a mutex for a maximum of 1.75 seconds. And it always times out because no one ever sends a signal.

```

1  #include <stdio.h>
2  #include <time.h>
3  #include <threads.h>
4
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10
11     struct timespec ts;
12
13     // Get the time now
14     timespec_get(&ts, TIME_UTC);
15
16     // Add on 1.75 seconds from now
17     ts.tv_sec += 1;
18     ts.tv_nsec += 750000000L;
19
20     // Handle nsec overflow
21     ts.tv_sec += ts.tv_nsec / 1000000000L;
22     ts.tv_nsec = ts.tv_nsec % 1000000000L;
23
24     printf("Thread: waiting for lock...\n");
25     int r = mtx_timedlock(&mutex, &ts);
26
27     switch (r) {
28         case thrd_success:
29             printf("Thread: grabbed lock!\n");
30             break;
31
32         case thrd_timedout:
33             printf("Thread: timed out!\n");

```

```
34         break;
35
36     case thrd_error:
37         printf("Thread: Some kind of error\n");
38         break;
39     }
40
41     mtx_unlock(&mutex);
42
43     return 0;
44 }
45
46 int main(void)
47 {
48     thrd_t t;
49
50     mtx_init(&mutex, mtx_plain);
51
52     mtx_lock(&mutex);
53
54     printf("Main creating thread\n");
55     thrd_create(&t, run, NULL);
56
57     // Sleep 3s to allow the other thread to timeout
58     thrd_sleep(&(struct timespec){.tv_sec=3}, NULL);
59
60     mtx_unlock(&mutex);
61
62     thrd_join(t, NULL);
63
64     mtx_destroy(&mutex);
65 }
```

Output:

```
Main creating thread
Thread: waiting for lock...
Thread: timed out!
```

## See Also

`mtx_lock()`, `mtx_trylock()`, `timespec_get()`

---

## 28.12 `mtx_trylock()`

Try to lock a mutex, returning if not possible

## Synopsis

```
#include <threads.h>

int mtx_trylock(mtx_t *mtx);
```

## Description

This works just like `mtx_lock` except that it returns instantly if a lock can't be obtained.

The spec notes that there's a chance that `mtx_trylock()` might spuriously fail with `thrd_busy` even if there are no other threads holding the lock. I'm not sure why this is, but you should defensively code against it.

## Return Value

Returns `thrd_success` if all's well. Or `thrd_busy` if some other thread holds the lock. Or `thrd_error`, which means something went right. I mean "wrong".

## Example

```
1  #include <stdio.h>
2  #include <time.h>
3  #include <threads.h>
4
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      int id = *(int*)arg;
10
11     int r = mtx_trylock(&mutex);    // <-- TRY TO GRAB THE LOCK
12
13     switch (r) {
14         case thrd_success:
15             printf("Thread %d: grabbed lock!\n", id);
16             break;
17
18         case thrd_busy:
19             printf("Thread %d: lock already taken :(\n", id);
20             return 1;
21
22         case thrd_error:
23             printf("Thread %d: Some kind of error\n", id);
24             return 2;
25     }
26
27     mtx_unlock(&mutex);
28
29     return 0;
30 }
31
32 #define THREAD_COUNT 5
```

```

33
34 int main(void)
35 {
36     thrd_t t[THREAD_COUNT];
37     int id[THREAD_COUNT];
38
39     mtx_init(&mutex, mtx_plain);
40
41     for (int i = 0; i < THREAD_COUNT; i++) {
42         id[i] = i;
43         thrd_create(t + i, run, id + i);
44     }
45
46     for (int i = 0; i < THREAD_COUNT; i++)
47         thrd_join(t[i], NULL);
48
49     mtx_destroy(&mutex);
50 }

```

Output (varies by run):

```

Thread 0: grabbed lock!
Thread 1: lock already taken :(
Thread 4: lock already taken :(
Thread 3: grabbed lock!
Thread 2: lock already taken :(

```

## See Also

`mtx_lock()`, `mtx_timedlock()`, `mtx_unlock()`

## 28.13 `mtx_unlock()`

Free a mutex when you're done with the critical section

### Synopsis

```

#include <threads.h>

int mtx_unlock(mtx_t *mtx);

```

### Description

After you've done all the dangerous stuff you have to do, wherein the involved threads should not be stepping on each other's toes... you can free up your stranglehold on the mutex by calling `mtx_unlock()`.

### Return Value

Returns `thrd_success` on success. Or `thrd_error` on error. It's not very original in this regard.

## Example

General-purpose mutex example here, but you can see the `mtx_unlock()` in the `run()` function:

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  cnd_t condvar;
5  mtx_t mutex;
6
7  int run(void *arg)
8  {
9      (void)arg;
10
11     static int count = 0;
12
13     mtx_lock(&mutex);
14
15     printf("Thread: I got %d!\n", count);
16     count++;
17
18     mtx_unlock(&mutex); // <-- UNLOCK HERE
19
20     return 0;
21 }
22
23 #define THREAD_COUNT 5
24
25 int main(void)
26 {
27     thrd_t t[THREAD_COUNT];
28
29     mtx_init(&mutex, mtx_plain);
30
31     for (int i = 0; i < THREAD_COUNT; i++)
32         thrd_create(t + i, run, NULL);
33
34     for (int i = 0; i < THREAD_COUNT; i++)
35         thrd_join(t[i], NULL);
36
37     mtx_destroy(&mutex);
38 }

```

Output:

```

Thread: I got 0!
Thread: I got 1!
Thread: I got 2!
Thread: I got 3!
Thread: I got 4!

```

## See Also

`mtx_lock()`, `mtx_timedlock()`, `mtx_trylock()`



## 28.14 `thrd_create()`

Create a new thread of execution

### Synopsis

```
#include <threads.h>

int thrd_create(thrd_t *thr, thrd_start_t func, void *arg);
```

### Description

Now you have the POWER!

Right?

This is how you launch new threads to make your program do multiple things at once<sup>2</sup>!

In order to make this happen, you need to pass a pointer to a `thrd_t` that will be used to represent the thread you're spawning.

That thread will start running the function you pass a pointer to in `func`. This is a value of type `thrd_start_t`, which is a pointer to a function that returns an `int` and takes a single `void*` as a parameter, i.e.:

```
int thread_run_func(void *arg)
```

And, as you might have guessed, the pointer you pass to `thrd_create()` for the `arg` parameter is passed on to the `func` function. This is how you can give additional information to the thread when it starts up.

Of course, for `arg`, you have to be sure to pass a pointer to an object that is thread-safe or per-thread.

If the thread returns from the function, it exits just as if it had called `thrd_exit()`.

Finally, the value that the `func` function returns can be picked up by the parent thread with `thrd_join()`.

### Return Value

In the case of goodness, returns `thrd_success`. If you're out of memory, will return `thrd_nomem`. Otherwise, `thrd_error`.

### Example

```
1 #include <stdio.h>
2 #include <threads.h>
3
4 int run(void *arg)
5 {
6     int id = *(int*)arg;
7
8     printf("Thread %d: I'm alive!!\n", id);
```

<sup>2</sup>Well, as at least as many things as you have free cores. Your OS will schedule them as it can.

```

9
10     return id;
11 }
12
13 #define THREAD_COUNT 5
14
15 int main(void)
16 {
17     thrd_t t[THREAD_COUNT];
18     int id[THREAD_COUNT]; // One of these per thread
19
20     for (int i = 0; i < THREAD_COUNT; i++) {
21         id[i] = i; // Let's pass in the thread number as the ID
22         thrd_create(t + i, run, id + i);
23     }
24
25     for (int i = 0; i < THREAD_COUNT; i++) {
26         int res;
27
28         thrd_join(t[i], &res);
29
30         printf("Main: thread %d exited with code %d\n", i, res);
31     }
32 }

```

Output (might vary from run to run):

```

Thread 1: I'm alive!!
Thread 0: I'm alive!!
Thread 3: I'm alive!!
Thread 2: I'm alive!!
Main: thread 0 exited with code 0
Main: thread 1 exited with code 1
Main: thread 2 exited with code 2
Main: thread 3 exited with code 3
Thread 4: I'm alive!!
Main: thread 4 exited with code 4

```

## See Also

`thrd_exit()`, `thrd_join()`

## 28.15 `thrd_current()`

Get the ID of the calling thread

## Synopsis

```
#include <threads.h>

thrd_t thrd_current(void);
```

## Description

Each thread has an opaque ID of type `thrd_t`. This is the value we see get initialized when we call `thrd_create()`.

But what if you want to get the ID of the currently running thread?

No problem! Just call this function and it will be returned to you.

Why? Who knows!

Well, to be honest, I could see it being used a couple places.

1. You could use it to have a thread detach itself with `thrd_detach()`. I'm not sure why you'd want to do this, however.
2. You could use it to compare this thread's ID with another you have stored in a variable somewhere by using the `thrd_equal()` function. Seems like the most legit use.
3. ...
4. Profit!

If anyone has another use, please let me know.

## Return Value

Returns the calling thread's ID.

## Example

Here's a general example that shows getting the current thread ID and comparing it to a previously-recorded thread ID and taking exciting action based on the result! Starring Arnold Schwarzenegger!

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  thrd_t first_thread_id;
5
6  int run(void *arg)
7  {
8      (void)arg;
9
10     thrd_t my_id = thrd_current();    // <-- GET MY THREAD ID
11
12     if (thrd_equal(my_id, first_thread_id))
13         printf("I'm the first thread!\n");
14     else
15         printf("I'm not the first!\n");
16
17     return 0;
18 }
```

```

19
20 int main(void)
21 {
22     thrd_t t;
23
24     thrd_create(&first_thread_id, run, NULL);
25     thrd_create(&t, run, NULL);
26
27     thrd_join(first_thread_id, NULL);
28     thrd_join(t, NULL);
29 }

```

Output:

```
Come on, you got what you want, Cohaagen! Give deez people ay-ah!
```

No, wait, that's an Arnold Schwarzenegger quote from *Total Recall*, one of the best science fiction films of all time. Watch it now and then come back to finish this reference page.

Man—what an ending! And Johnny Cab? So excellent. Anyway!

Output:

```
I'm the first thread!
I'm not the first!
```

## See Also

`thrd_equal()`, `thrd_detach()`

---

## 28.16 `thrd_detach()`

Automatically clean up threads when they exit

### Synopsis

```

#include <threads.h>

int thrd_detach(thrd_t thr);

```

### Description

Normally you have to `thrd_join()` to get resources associated with a deceased thread cleaned up. (Most notably, its exit status is still floating around waiting to get picked up.)

But if you call `thrd_detach()` on the thread first, manual cleanup isn't necessary. They just exit and are cleaned up by the OS.

(Note that when the main thread dies, all the threads die in any case.)

## Return Value

`thrd_success` if the thread successfully detaches, `thrd_error` otherwise.

## Example

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  thrd_t first_thread_id;
5
6  int run(void *arg)
7  {
8      (void)arg;
9
10     printf("Thread running!\n");
11
12     return 0;
13 }
14
15 #define THREAD_COUNT 5
16
17 int main(void)
18 {
19     thrd_t t;
20
21     for (int i = 0; i < THREAD_COUNT; i++) {
22         thrd_create(&t, run, NULL);
23         thrd_detach(t);
24     }
25
26     // No need to thrd_join()!
27
28     // Sleep a quarter second to let them all finish
29     thrd_sleep(&(struct timespec){.tv_nsec=250000000}, NULL);
30 }
```

## See Also

`thrd_join()`, `thrd_exit()`

---

## 28.17 `thrd_equal()`

Compare two thread descriptors for equality

## Synopsis

```
#include <threads.h>

int thrd_equal(thrd_t thr0, thrd_t thr1);
```

## Description

If you have two thread descriptors in `thrd_t` variables, you can test them for equality with this function.

For example, maybe one of the threads has special powers the others don't, and the run function needs to be able to tell them apart, as in the example.

## Return Value

Returns non-zero if the threads are equal. Returns 0 if they're not.

## Example

Here's a general example that shows getting the current thread ID and comparing it to a previously-recorded thread ID and taking boring action based on the result.

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  thrd_t first_thread_id;
5
6  int run(void *arg)
7  {
8      (void)arg;
9
10     thrd_t my_id = thrd_current();
11
12     if (thrd_equal(my_id, first_thread_id)) // <-- COMPARE!
13         printf("I'm the first thread!\n");
14     else
15         printf("I'm not the first!\n");
16
17     return 0;
18 }
19
20 int main(void)
21 {
22     thrd_t t;
23
24     thrd_create(&first_thread_id, run, NULL);
25     thrd_create(&t, run, NULL);
26
27     thrd_join(first_thread_id, NULL);
28     thrd_join(t, NULL);
29 }
```

Output:

```
I'm the first thread!  
I'm not the first!
```

## See Also

`thrd_current()`

---

## 28.18 `thrd_exit()`

Stop and exit this thread

### Synopsis

```
#include <threads.h>  
  
_Noreturn void thrd_exit(int res);
```

### Description

A thread commonly exits by returning from its run function. But if it wants to exit early (perhaps from deeper in the call stack), this function will get that done.

The `res` code can be picked up by a thread calling `thrd_join()`, and is equivalent to returning a value from the run function.

Like with returning from the run function, this will also properly clean up all the thread-specific storage associated with this thread—all the destructors for the threads TSS variables will be called. If there are any remaining TSS variables with destructors after the first round of destruction<sup>3</sup>, the remaining destructors will be called. This happens repeatedly until there are no more, or the number of rounds of carnage reaches `TSS_DTOR_ITERATIONS`.

If the main thread calls this, it's as if you called `exit(EXIT_SUCCESS)`.

### Return Value

This function never returns because the thread calling it is killed in the process. Trippy!

### Example

Threads in this example exit early with result `22` if they get a `NULL` value for `arg`.

```
1 #include <stdio.h>  
2 #include <threads.h>  
3  
4 thrd_t first_thread_id;  
5  
6 int run(void *arg)  
7 {  
8     (void)arg;
```

---

<sup>3</sup>For example, if a destructor caused more variables to be set.

```

9
10     if (arg == NULL)
11         thrd_exit(22);
12
13     return 0;
14 }
15
16 #define THREAD_COUNT 5
17
18 int main(void)
19 {
20     thrd_t t[THREAD_COUNT];
21
22     for (int i = 0; i < THREAD_COUNT; i++)
23         thrd_create(t + i, run, i == 2? NULL: "spatula");
24
25
26     for (int i = 0; i < THREAD_COUNT; i++) {
27         int res;
28         thrd_join(t[i], &res);
29
30         printf("Thread %d exited with code %d\n", i, res);
31     }
32 }

```

Output:

```

Thread 0 exited with code 0
Thread 1 exited with code 0
Thread 2 exited with code 22
Thread 3 exited with code 0
Thread 4 exited with code 0

```

## See Also

`thrd_join()`

## 28.19 `thrd_join()`

Wait for a thread to exit

### Synopsis

```

#include <threads.h>

int thrd_join(thrd_t thr, int *res);

```

### Description

When a parent thread fires off some child threads, it can wait for them to complete with this call



## Return Value

### Example

Threads in this example exit early with result 22 if they get a `NULL` value for `arg`. The parent thread picks up this result code with `thrd_join()`.

```

1  #include <stdio.h>
2  #include <threads.h>
3
4  thrd_t first_thread_id;
5
6  int run(void *arg)
7  {
8      (void)arg;
9
10     if (arg == NULL)
11         thrd_exit(22);
12
13     return 0;
14 }
15
16 #define THREAD_COUNT 5
17
18 int main(void)
19 {
20     thrd_t t[THREAD_COUNT];
21
22     for (int i = 0; i < THREAD_COUNT; i++)
23         thrd_create(t + i, run, i == 2? NULL: "spatula");
24
25
26     for (int i = 0; i < THREAD_COUNT; i++) {
27         int res;
28         thrd_join(t[i], &res);
29
30         printf("Thread %d exited with code %d\n", i, res);
31     }
32 }

```

Output:

```

Thread 0 exited with code 0
Thread 1 exited with code 0
Thread 2 exited with code 22
Thread 3 exited with code 0
Thread 4 exited with code 0

```

### See Also

`thrd_exit()`

## 28.20 `thrd_sleep()`

Sleep for a specific number of seconds and nanoseconds

### Synopsis

```
#include <threads.h>

int thrd_sleep(const struct timespec *duration, struct timespec *remaining);
```

### Description

This function puts the current thread to sleep for a while<sup>4</sup> allowing other threads to run.

The calling thread will wake up after the time has elapsed, or if it gets interrupted by a signal or something.

If it doesn't get interrupted, it'll sleep at least as long as you asked. Maybe a tad longer. You know how hard it can be to get out of bed.

The structure looks like this:

```
struct timespec {
    time_t tv_sec;    // Seconds
    long   tv_nsec;   // Nanoseconds (billionths of a second)
};
```

Don't set `tv_nsec` greater than 999,999,999. I can't see what officially happens if you do, but on my system `thrd_sleep()` returns `-2` and fails.

### Return Value

Returns `0` on timeout, or `-1` if interrupted by a signal. Or any negative value on some other error. Weirdly, the spec allows this "other error negative value" to also be `-1`, so good luck with that.

### Example

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  int main(void)
5  {
6      // Sleep for 3.25 seconds
7      thrd_sleep(&(struct timespec){.tv_sec=3, .tv_nsec=250000000}, NULL);
8
9      return 0;
10 }
```

### See Also

`thrd_yield()`

---

<sup>4</sup>Unix-like systems have a `sleep()` syscall that sleeps for an integer number of seconds. But `thrd_sleep()` is likely more portable and gives subsecond resolution, besides!

## 28.21 `thrd_yield()`

Stop running that other threads might run

### Synopsis

```
#include <threads.h>

void thrd_yield(void);
```

### Description

If you have a thread that's hogging the CPU and you want to give your other threads time to run, you can call `thrd_yield()`. If the system sees fit, it will put the calling thread to sleep and one of the other threads will run instead.

It's a good way to be "polite" to the other threads in your program if you want to encourage them to run instead.

### Return Value

Returns nothing!

### Example

This example's kinda poor because the OS is probably going to reschedule threads on the output anyway, but it gets the point across.

The main thread is giving other threads a chance to run after every block of dumb work it does.

```
1  #include <stdio.h>
2  #include <threads.h>
3
4  int run(void *arg)
5  {
6      int main_thread = arg != NULL;
7
8      if (main_thread) {
9          long int total = 0;
10
11         for (int i = 0; i < 10; i++) {
12             for (long int j = 0; j < 1000L; j++)
13                 total++;
14
15             printf("Main thread yielding\n");
16             thrd_yield();           // <-- YIELD HERE
17         }
18     } else
19         printf("Other thread running!\n");
20
```

```

21     return 0;
22 }
23
24 #define THREAD_COUNT 10
25
26 int main(void)
27 {
28     thrd_t t[THREAD_COUNT];
29
30     for (int i = 0; i < THREAD_COUNT; i++)
31         thrd_create(t + i, run, i == 0? "main": NULL);
32
33     for (int i = 0; i < THREAD_COUNT; i++)
34         thrd_join(t[i], NULL);
35
36     return 0;
37 }

```

The output will vary from run to run. Notice that even after `thrd_yield()` other threads might not yet be ready to run and the main thread will continue.

```

Main thread yielding
Main thread yielding
Main thread yielding
Other thread running!
Other thread running!
Other thread running!
Other thread running!
Main thread yielding
Other thread running!
Other thread running!
Main thread yielding
Main thread yielding
Main thread yielding
Other thread running!
Main thread yielding
Main thread yielding
Main thread yielding
Other thread running!
Other thread running!

```

## See Also

`thrd_sleep()`

---

## 28.22 `tss_create()`

Create new thread-specific storage

## Synopsis

```
#include <threads.h>

int tss_create(tss_t *key, tss_dtor_t dtor);
```

## Description

This helps when you need per-thread storage of different values.

A common place this comes up is if you have a file scope variable that is shared between a bunch of functions and often returned. That's not threadsafe. One way to refactor is to replace it with thread-specific storage so that each thread gets their own code and doesn't step on other thread's toes.

To make this work, you pass in a pointer to a `tss_t` key—this is the variable you will use in subsequent `tss_set()` and `tss_get()` calls to set and get the value associated with the key.

The interesting part of this is the `dtor` destructor pointer of type `tss_dtor_t`. This is actually a pointer to a function that takes a `void*` argument and returns `void`, i.e.

```
void dtor(void *p) { ... }
```

This function will be called per thread when the thread exits with `thrd_exit()` (or returns from the run function).

It's unspecified behavior to call this function while other threads' destructors are running.

## Return Value

Returns nothing!

## Example

This is a general-purpose TSS example. Note the TSS variable is created near the top of `main()`.

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <threads.h>
4
5  tss_t str;
6
7  void some_function(void)
8  {
9      // Retrieve the per-thread value of this string
10     char *tss_string = tss_get(str);
11
12     // And print it
13     printf("TSS string: %s\n", tss_string);
14 }
15
16 int run(void *arg)
17 {
18     int serial = *(int*)arg; // Get this thread's serial number
19     free(arg);
```

```

20
21 // malloc() space to hold the data for this thread
22 char *s = malloc(64);
23 sprintf(s, "thread %d! :)", serial); // Happy little string
24
25 // Set this TSS variable to point at the string
26 tss_set(str, s);
27
28 // Call a function that will get the variable
29 some_function();
30
31 return 0; // Equivalent to thrd_exit(0); fires destructors
32 }
33
34 #define THREAD_COUNT 15
35
36 int main(void)
37 {
38     thrd_t t[THREAD_COUNT];
39
40     // Make a new TSS variable, the free() function is the destructor
41     tss_create(&str, free); // <-- CREATE TSS VAR!
42
43     for (int i = 0; i < THREAD_COUNT; i++) {
44         int *n = malloc(sizeof *n); // Holds a thread serial number
45         *n = i;
46         thrd_create(t + i, run, n);
47     }
48
49     for (int i = 0; i < THREAD_COUNT; i++) {
50         thrd_join(t[i], NULL);
51     }
52
53     // And all threads are done, so let's free this
54     tss_delete(str);
55 }

```

Output:

```

TSS string: thread 0! :)
TSS string: thread 2! :)
TSS string: thread 1! :)
TSS string: thread 5! :)
TSS string: thread 3! :)
TSS string: thread 6! :)
TSS string: thread 4! :)
TSS string: thread 7! :)
TSS string: thread 8! :)
TSS string: thread 9! :)
TSS string: thread 10! :)
TSS string: thread 13! :)
TSS string: thread 12! :)
TSS string: thread 11! :)

```

```
TSS string: thread 14! :)
```

## See Also

`tss_delete()`, `tss_set()`, `tss_get()`, `thr_exit()`

---

## 28.23 `tss_delete()`

Clean up a thread-specific storage variable

### Synopsis

```
#include <threads.h>

void tss_delete(tss_t key);
```

### Description

This is the opposite of `tss_create()`. You create (initialize) the TSS variable before using it, then, when all the threads are done that need it, you delete (deinitialize/free) it with this.

This doesn't call any destructors! Those are all called by `thr_exit()`!

### Return Value

Returns nothing!

### Example

This is a general-purpose TSS example. Note the TSS variable is deleted near the bottom of `main()`.

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <threads.h>
4
5  tss_t str;
6
7  void some_function(void)
8  {
9      // Retrieve the per-thread value of this string
10     char *tss_string = tss_get(str);
11
12     // And print it
13     printf("TSS string: %s\n", tss_string);
14 }
15
16 int run(void *arg)
17 {
18     int serial = *(int*)arg; // Get this thread's serial number
```

```

19     free(arg);
20
21     // malloc() space to hold the data for this thread
22     char *s = malloc(64);
23     sprintf(s, "thread %d! :)", serial); // Happy little string
24
25     // Set this TSS variable to point at the string
26     tss_set(str, s);
27
28     // Call a function that will get the variable
29     some_function();
30
31     return 0; // Equivalent to thrd_exit(0); fires destructors
32 }
33
34 #define THREAD_COUNT 15
35
36 int main(void)
37 {
38     thrd_t t[THREAD_COUNT];
39
40     // Make a new TSS variable, the free() function is the destructor
41     tss_create(&str, free);
42
43     for (int i = 0; i < THREAD_COUNT; i++) {
44         int *n = malloc(sizeof *n); // Holds a thread serial number
45         *n = i;
46         thrd_create(t + i, run, n);
47     }
48
49     for (int i = 0; i < THREAD_COUNT; i++) {
50         thrd_join(t[i], NULL);
51     }
52
53     // And all threads are done, so let's free this
54     tss_delete(str); // <-- DELETE TSS VARIABLE!
55 }

```

Output:

```

TSS string: thread 0! :)
TSS string: thread 2! :)
TSS string: thread 1! :)
TSS string: thread 5! :)
TSS string: thread 3! :)
TSS string: thread 6! :)
TSS string: thread 4! :)
TSS string: thread 7! :)
TSS string: thread 8! :)
TSS string: thread 9! :)
TSS string: thread 10! :)
TSS string: thread 13! :)

```



```
TSS string: thread 12! :)
TSS string: thread 11! :)
TSS string: thread 14! :)
```

## See Also

`tss_create()`, `tss_set()`, `tss_get()`, `thr_exit()`

---

## 28.24 `tss_get()`

Get thread-specific data

### Synopsis

```
#include <threads.h>

void *tss_get(tss_t key);
```

### Description

Once you've set a variable with `tss_set()`, you can retrieve the value with `tss_get()`—just pass in the key and you'll get a pointer to the value back.

Don't call this from a destructor.

### Return Value

Returns the value stored for the given `key`, or `NULL` if there's trouble.

### Example

This is a general-purpose TSS example. Note the TSS variable is retrieved in `some_function()`, below.

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <threads.h>
4
5  tss_t str;
6
7  void some_function(void)
8  {
9      // Retrieve the per-thread value of this string
10     char *tss_string = tss_get(str);    // <-- GET THE VALUE
11
12     // And print it
13     printf("TSS string: %s\n", tss_string);
14 }
15
16 int run(void *arg)
```

```

17 {
18     int serial = *(int*)arg; // Get this thread's serial number
19     free(arg);
20
21     // malloc() space to hold the data for this thread
22     char *s = malloc(64);
23     sprintf(s, "thread %d! :)", serial); // Happy little string
24
25     // Set this TSS variable to point at the string
26     tss_set(str, s);
27
28     // Call a function that will get the variable
29     some_function();
30
31     return 0; // Equivalent to thrd_exit(0); fires destructors
32 }
33
34 #define THREAD_COUNT 15
35
36 int main(void)
37 {
38     thrd_t t[THREAD_COUNT];
39
40     // Make a new TSS variable, the free() function is the destructor
41     tss_create(&str, free);
42
43     for (int i = 0; i < THREAD_COUNT; i++) {
44         int *n = malloc(sizeof *n); // Holds a thread serial number
45         *n = i;
46         thrd_create(t + i, run, n);
47     }
48
49     for (int i = 0; i < THREAD_COUNT; i++) {
50         thrd_join(t[i], NULL);
51     }
52
53     // And all threads are done, so let's free this
54     tss_delete(str);
55 }

```

Output:

```

TSS string: thread 0! :)
TSS string: thread 2! :)
TSS string: thread 1! :)
TSS string: thread 5! :)
TSS string: thread 3! :)
TSS string: thread 6! :)
TSS string: thread 4! :)
TSS string: thread 7! :)
TSS string: thread 8! :)
TSS string: thread 9! :)

```

```
TSS string: thread 10! :)
TSS string: thread 13! :)
TSS string: thread 12! :)
TSS string: thread 11! :)
TSS string: thread 14! :)
```

## See Also

`tss_set()`

---

## 28.25 `tss_set()`

Set thread-specific data

### Synopsis

```
#include <threads.h>

int tss_set(tss_t key, void *val);
```

### Description

Once you've set up your TSS variable with `tss_create()`, you can set it on a per thread basis with `tss_set()`.

`key` is the identifier for this data, and `val` is a pointer to it.

The destructor specified in `tss_create()` will be called for the value set when the thread exits.

Also, if there's a destructor *and* there is already a value for this key in place, the destructor will not be called for the already-existing value. In fact, this function will never cause a destructor to be called. So you're on your own, there—best clean up the old value before overwriting it with the new one.

### Return Value

Returns `thrd_success` when happy, and `thrd_error` when not.

### Example

This is a general-purpose TSS example. Note the TSS variable is set in `run()`, below.

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <threads.h>
4
5 tss_t str;
6
7 void some_function(void)
8 {
9     // Retrieve the per-thread value of this string
```

```

10     char *tss_string = tss_get(str);
11
12     // And print it
13     printf("TSS string: %s\n", tss_string);
14 }
15
16 int run(void *arg)
17 {
18     int serial = *(int*)arg; // Get this thread's serial number
19     free(arg);
20
21     // malloc() space to hold the data for this thread
22     char *s = malloc(64);
23     sprintf(s, "thread %d! :)", serial); // Happy little string
24
25     // Set this TSS variable to point at the string
26     tss_set(str, s); // <-- SET THE TSS VARIABLE
27
28     // Call a function that will get the variable
29     some_function();
30
31     return 0; // Equivalent to thrd_exit(0); fires destructors
32 }
33
34 #define THREAD_COUNT 15
35
36 int main(void)
37 {
38     thrd_t t[THREAD_COUNT];
39
40     // Make a new TSS variable, the free() function is the destructor
41     tss_create(&str, free);
42
43     for (int i = 0; i < THREAD_COUNT; i++) {
44         int *n = malloc(sizeof *n); // Holds a thread serial number
45         *n = i;
46         thrd_create(t + i, run, n);
47     }
48
49     for (int i = 0; i < THREAD_COUNT; i++) {
50         thrd_join(t[i], NULL);
51     }
52
53     // And all threads are done, so let's free this
54     tss_delete(str);
55 }

```

Output:

```

TSS string: thread 0! :)
TSS string: thread 2! :)
TSS string: thread 1! :)

```

```
TSS string: thread 5! :)
TSS string: thread 3! :)
TSS string: thread 6! :)
TSS string: thread 4! :)
TSS string: thread 7! :)
TSS string: thread 8! :)
TSS string: thread 9! :)
TSS string: thread 10! :)
TSS string: thread 13! :)
TSS string: thread 12! :)
TSS string: thread 11! :)
TSS string: thread 14! :)
```

## See Also

`tss_get()`

## Chapter 29

# <time.h> Date and Time Functions

| Function                    | Description                                                 |
|-----------------------------|-------------------------------------------------------------|
| <code>clock()</code>        | How much processor time has been used by this process       |
| <code>difftime()</code>     | Compute the difference between two times                    |
| <code>mktime()</code>       | Convert a <code>struct tm</code> into a <code>time_t</code> |
| <code>time()</code>         | Get the current calendar time                               |
| <code>timespec_get()</code> | Get a higher resolution time, probably now                  |
| <code>asctime()</code>      | Return a human-readable version of a <code>struct tm</code> |
| <code>ctime()</code>        | Return a human-readable version of a <code>time_t</code>    |
| <code>gmtime()</code>       | Convert a calendar time into a UTC broken-down time         |
| <code>localtime()</code>    | Convert a calendar time into a broken-down local time       |
| <code>strftime()</code>     | Formatted date and time output                              |

When it comes to time and C, there are two main types to look for:

- `time_t` holds a *calendar time*. This is an potentially opaque numeric type that represents an absolute time that can be converted to UTC<sup>1</sup> or local time.
- `struct tm` holds a *broken-down time*. This has things like the day of the week, the day of the month, the hour, the minute, the second, etc.

On POSIX systems and Windows, `time_t` is an integer and represents the number of seconds that have elapsed since January 1, 1970 at 00:00 UTC.

A `struct tm` contains the following fields:

```
struct tm {
    int tm_sec;    // seconds after the minute -- [0, 60]
    int tm_min;    // minutes after the hour -- [0, 59]
    int tm_hour;   // hours since midnight -- [0, 23]
    int tm_mday;   // day of the month -- [1, 31]
    int tm_mon;    // months since January -- [0, 11]
    int tm_year;   // years since 1900
    int tm_wday;   // days since Sunday -- [0, 6]
    int tm_yday;   // days since January 1 -- [0, 365]
    int tm_isdst;  // Daylight Saving Time flag
```

<sup>1</sup>When you say GMT, unless you're talking specifically about the time zone and not the time, you probably mean "UTC".

```
};
```

You can convert between the two with `mktime()`, `gmtime()`, and `localtime()`.

You can print time information to strings with `ctime()`, `asctime()`, and `strftime()`.

## 29.1 Thread Safety Warning

`asctime()`, `ctime()`: These two functions return a pointer to a `static` memory region. They both might return the same pointer. If you need thread safety, you'll need a mutex across them. If you need both results at once, `strncpy()` one of them out.

All these problems with `asctime()` and `ctime()` can be avoided by using the more flexible and thread-safe `strftime()` function instead.

`localtime()`, `gmtime()`: These other two functions also return a pointer to a `static` memory region. They both might return the same pointer. If you need thread safety, you'll need a mutex across them. If you need both results at once, copy the `struct` to another.

---

## 29.2 `clock()`

How much processor time has been used by this process

### Synopsis

```
#include <time.h>

clock_t clock(void);
```

### Description

Your processor is juggling a lot of things right now. Just because a process has been alive for 20 minutes doesn't mean that it used 20 minutes of "CPU time".

Most of the time your average process spends asleep, and that doesn't count toward the CPU time spent.

This function returns an opaque type representing the number of "clock ticks"<sup>2</sup> the process has spent in operation.

You can get the number of seconds out of that by dividing by the macro `CLOCKS_PER_SEC`. This is an integer, so you will have to cast part of the expression to a floating type to get a fractional time.

Note that this is not the "wall clock time" of the program. If you want to get that loosely use `time()` and `difftime()` (which might only offer 1-second resolution) or `timespec_get()` (which might only also offer low resolution, but at least it *might* go to nanosecond level).

### Return Value

Returns the amount of CPU time spent by this process. This comes back in a form that can be divided by `CLOCKS_PER_SEC` to determine the time in seconds.

---

<sup>2</sup>The spec doesn't actually say "clock ticks", but I... am.

## Example

```
1 #include <stdio.h>
2 #include <time.h>
3
4 // Deliberately naive Fibonacci
5 long long int fib(long long int n) {
6     if (n <= 1) return n;
7
8     return fib(n-1) + fib(n-2);
9 }
10
11 int main(void)
12 {
13     printf("The 42nd Fibonacci Number is %lld\n", fib(42));
14
15     printf("CPU time: %f\n", clock() / (double)CLOCKS_PER_SEC);
16 }
```

Output on my system:

```
The 42nd Fibonacci Number is 267914296
CPU time: 1.863078
```

## See Also

`time()`, `difftime()`, `timespec_get()`

---

## 29.3 `difftime()`

Compute the difference between two times

### Synopsis

```
#include <time.h>

double difftime(time_t time1, time_t time0);
```

### Description

Since the `time_t` type is technically opaque, you can't just straight-up subtract to get the difference between two of them<sup>3</sup>. Use this function to do it.

There is no guarantee as to the resolution of this difference, but it's probably to the second.

### Return Value

Returns the difference between two `time_ts` in seconds.

---

<sup>3</sup>Unless you're on a POSIX system where `time_t` is definitely an integer, in which case you can subtract. But you should still use `difftime()` for maximum portability.



## Example

```

1  #include <stdio.h>
2  #include <time.h>
3
4  int main(void)
5  {
6      // April 12, 1982 and change
7      struct tm time_a = { .tm_year=82, .tm_mon=3, .tm_mday=12,
8                          .tm_hour=4, .tm_min=00, .tm_sec=04, .tm_isdst=-1,
9                          };
10
11     // November 15, 2020 and change
12     struct tm time_b = { .tm_year=120, .tm_mon=10, .tm_mday=15,
13                        .tm_hour=16, .tm_min=27, .tm_sec=00, .tm_isdst=-1,
14                        };
15
16     time_t cal_a = mktime(&time_a);
17     time_t cal_b = mktime(&time_b);
18
19     double diff = difftime(cal_b, cal_a);
20
21     double years = diff / 60 / 60 / 24 / 365.2425; // close enough
22
23     printf("%f seconds (%f years) between events\n", diff, years);
24 }

```

Output:

```
1217996816.000000 seconds (38.596783 years) between events
```

## See Also

`time()`, `mktime()`

## 29.4 mktime()

Convert a `struct tm` with local time into a `time_t`

### Synopsis

```

#include <time.h>

time_t mktime(struct tm *timeptr);

```

### Description

If you have a local date and time and want it converted to a `time_t` (so that you can `difftime()` it or whatever), you can convert it with this function.

Basically you fill out the fields in your `struct tm` in local time and `mktime()` will convert those to the UTC `time_t` equivalent.

A couple notes:

- Don't bother filling out `tm_wday` or `tm_yday`. `mktime()` will fill these out for you.
- You can set `tm_isdst` to `0` to indicate your time isn't Daylight Saving Time (DST), `1` to indicate it is, and `-1` to have `mktime()` fill it in according to your locale's preference.

If you don't have a C23 compiler and you need input in UTC, see the non-standard functions `timegm()`<sup>4</sup> for Unix-likes and `_mkgmtime()`<sup>5</sup> for Windows.

## Return Value

Returns the local time in the `struct tm` as a `time_t` calendar time.

Returns `(time_t)(-1)` on error.

## Example

In the following example, we have `mktime()` tell us if that time was DST or not.

```

1  #include <stdio.h>
2  #include <time.h>
3
4  int main(void)
5  {
6      struct tm broken_down_time = {
7          .tm_year=82,    // years since 1900
8          .tm_mon=3,      // months since January -- [0, 11]
9          .tm_mday=12,    // day of the month -- [1, 31]
10         .tm_hour=4,      // hours since midnight -- [0, 23]
11         .tm_min=00,      // minutes after the hour -- [0, 59]
12         .tm_sec=04,      // seconds after the minute -- [0, 60]
13         .tm_isdst=-1,    // Daylight Saving Time flag
14     };
15
16     time_t calendar_time = mktime(&broken_down_time);
17
18     char *days[] = {"Sunday", "Monday", "Tuesday",
19                     "Wednesday", "Thursday", "Friday", "Saturday"};
20
21     // This will print what was in broken_down_time
22     printf("Local time : %s", asctime(localtime(&calendar_time)));
23     printf("Is DST      : %d\n", broken_down_time.tm_isdst);
24     printf("Day of week: %s\n\n", days[broken_down_time.tm_wday]);
25
26     // This will print UTC for the local time, above
27     printf("UTC         : %s", asctime(gmtime(&calendar_time)));
28 }

```

Output (for me in Pacific Time—UTC is 8 hours ahead):

<sup>4</sup><https://man.archlinux.org/man/timegm.3.en>

<sup>5</sup><https://docs.microsoft.com/en-us/cpp/c-runtime-library/reference/mkgmtime-mkgmtime32-mkgmtime64?view=msvc-160>

```
Local time : Mon Apr 12 04:00:04 1982
Is DST      : 0
Day of week: Monday

UTC         : Mon Apr 12 12:00:04 1982
```

**See Also**

`timegm()`, `localtime()`, `gmtime()`

---

**29.5 timegm()**

Convert a `struct tm` with UTC time into a `time_t`

**Synopsis**

New in C23!

```
#include <time.h>

time_t timegm(struct tm *timeptr);
```

**Description**

If you have a UTC date and time and want it converted to a `time_t` (so that you can `difftime()` it or whatever), you can convert it with this function.

Basically you fill out the fields in your `struct tm` in local time and `mktime()` will convert those to the UTC `time_t` equivalent.

A couple notes:

- Don't bother filling out `tm_wday` or `tm_yday`. `mktime()` will fill these out for you.
- The spec doesn't say anything about the `tm_isdst` Daylight Saving flag, but since UTC is immune to DST, I'm assuming it is ignored or set to 0 for us.

If you don't have a C23 compiler and you need input in UTC, see the non-standard functions `timegm()`<sup>6</sup> for Unix-likes and `_mkgmtime()`<sup>7</sup> for Windows.

**Return Value**

Returns the UTC time in the `struct tm` as a `time_t` calendar time.

Returns `(time_t)(-1)` on error.

---

<sup>6</sup><https://man.archlinux.org/man/timegm.3.en>

<sup>7</sup><https://docs.microsoft.com/en-us/cpp/c-runtime-library/reference/mkgmtime-mkgmtime32-mkgmtime64?view=msvc-160>

## Example

```

1  #include <stdio.h>
2  #include <time.h>
3
4  int main(void)
5  {
6      struct tm broken_down_time = {
7          .tm_year=82,    // years since 1900
8          .tm_mon=3,      // months since January -- [0, 11]
9          .tm_mday=12,    // day of the month -- [1, 31]
10         .tm_hour=4,      // hours since midnight -- [0, 23]
11         .tm_min=00,      // minutes after the hour -- [0, 59]
12         .tm_sec=04,      // seconds after the minute -- [0, 60]
13         .tm_isdst=-1,    // Daylight Saving Time flag
14     };
15
16     time_t calendar_time = timegm(&broken_down_time);
17
18     char *days[] = {"Sunday", "Monday", "Tuesday",
19                     "Wednesday", "Thursday", "Friday", "Saturday"};
20
21     // This will print what was in broken_down_time
22     printf("UTC          : %s", asctime(gmtime(&calendar_time)));
23     printf("Day of week: %s\n\n", days[broken_down_time.tm_wday]);
24
25     // This will print UTC for the local time, above
26     printf("Local time : %s", asctime(localtime(&calendar_time)));
27 }

```

## See Also

`mktime()`, `timegm()`, `localtime()`, `gmtime()`

## 29.6 time()

Get the current calendar time

### Synopsis

```

#include <time.h>

time_t time(time_t *timer);

```

### Description

Returns the current calendar time right now. I mean, now. No, now!

If `timer` is not `NULL`, it gets loaded with the current time, as well.

This can be converted into a `struct tm` with `localtime()` or `gmtime()`, or printed directly with `ctime()`.

## Return Value

Returns the current calendar time. Also loads `timer` with the current time if it's not `NULL`.

Or returns `(time_t)(-1)` if the time isn't available because you've fallen out of the space-time continuum and/or the system doesn't support times.

## Example

```
1 #include <stdio.h>
2 #include <time.h>
3
4 int main(void)
5 {
6     time_t now = time(NULL);
7
8     printf("The local time is %s", ctime(&now));
9 }
```

Example output:

```
The local time is Mon Mar  1 18:45:14 2021
```

## See Also

`localtime()`, `gmtime()`, `ctime()`

---

## 29.7 timespec\_get()

Get a higher resolution time, probably now

### Synopsis

```
#include <time.h>

int timespec_get(struct timespec *ts, int base);
```

### Description

This function loads the current time UTC (unless directed otherwise) into the given `struct timespec`, `ts`.

That structure has two fields:

```
struct timespec {
    time_t tv_sec;    // Whole seconds
    long   tv_nsec;   // Nanoseconds, 0-999999999
}
```

Nanoseconds are billionths of a second. You can divide by 1000000000.0 to convert to seconds.

The `base` parameter has only one defined value, by the spec: `TIME_UTC`. So portably make it that. This will load `ts` with the current time in seconds since a system-defined Epoch<sup>8</sup>, often January 1, 1970 at 00:00 UTC.

Your implementation might define other values for `base`.

## Return Value

When `base` is `TIME_UTC`, loads `ts` with the current UTC time.

On success, returns `base`, valid values for which will always be non-zero. On error, returns `0`.

## Example

```
1 struct timespec ts;
2
3 timespec_get(&ts, TIME_UTC);
4
5 printf("%ld s, %ld ns\n", ts.tv_sec, ts.tv_nsec);
6
7 double float_time = ts.tv_sec + ts.tv_nsec/1000000000.0;
8 printf("%f seconds since epoch\n", float_time);
```

Example output:

```
1614654187 s, 825540756 ns
1614654187.825541 seconds since epoch
```

Here's a helper function to add values to a `struct timespec` that handles negative values and nanosecond overflow.

```
#include <stdlib.h>

// Add delta seconds and delta nanoseconds to ts.
// Negative values are allowed. Each component is added individually.
//
// Subtract 1.5 seconds from the current value:
//
// timespec_add(&ts, -1, -500000000L);

struct timespec *timespec_add(struct timespec *ts, long dsec, long dnsec)
{
    long sec = (long)ts->tv_sec + dsec;
    long nsec = ts->tv_nsec + dnsec;

    ldiv_t qr = ldiv(nsec, 1000000000L);

    if (qr.rem < 0) {
        nsec = 1000000000L + qr.rem;
        sec += qr.quot - 1;
    } else {
        nsec = qr.rem;
        sec += qr.quot;
    }
}
```

<sup>8</sup>[https://en.wikipedia.org/wiki/Unix\\_time](https://en.wikipedia.org/wiki/Unix_time)

```

    }

    ts->tv_sec = sec;
    ts->tv_nsec = nsec;

    return ts;
}

```

And here are some functions to convert from `long double` to `struct timespec` and back, just in case you like thinking in decimals. This is more limited in significant figures than using the integer values.

```

#include <math.h>

// Convert a struct timespec into a long double
long double timespec_to_ld(struct timespec *ts)
{
    return ts->tv_sec + ts->tv_nsec / 1000000000.0;
}

// Convert a long double to a struct timespec
struct timespec ld_to_timespec(long double t)
{
    long double f;
    struct timespec ts;
    ts.tv_nsec = modfl(t, &f) * 1000000000L;
    ts.tv_sec = f;

    return ts;
}

```

## See Also

`time()`, `mtx_timedlock()`, `cnd_timedwait()`

---

## 29.8 asctime()

Return a human-readable version of a `struct tm`

### Synopsis

```

#include <time.h>

char *asctime(const struct tm *timeptr)

```

### Description

This takes a time in a `struct tm` and returns a string with that date in the form:

```
Sun Sep 16 01:03:52 1973
```

with a newline included at the end, rather unhelpfully. (`strftime()` will give you more flexibility.)

It's just like `ctime()`, except it takes a `struct tm` instead of a `time_t`.

**WARNING:** This function returns a pointer to a static `char*` region that isn't thread-safe and might be shared with the `ctime()` function. If you need thread safety, use `strftime()` or use a mutex that covers `ctime()` and `asctime()`.

Behavior is undefined for:

- Years less than 1000
- Years greater than 9999
- Any members of `timeptr` are out of range

## Return Value

Returns a pointer to the human-readable date string.

## Example

```
1 #include <stdio.h>
2 #include <time.h>
3
4 int main(void)
5 {
6     time_t now = time(NULL);
7
8     printf("Local: %s", asctime(localtime(&now)));
9     printf("UTC  : %s", asctime(gmtime(&now)));
10 }
```

Sample output:

```
Local: Mon Mar  1 21:17:34 2021
UTC   : Tue Mar  2 05:17:34 2021
```

## See Also

`ctime()`, `localtime()`, `gmtime()`

## 29.9 `ctime()`

Return a human-readable version of a `time_t`



## Synopsis

```
#include <time.h>

char *ctime(const time_t *timer);
```

## Description

This takes a time in a `time_t` and returns a string with the local time and date in the form:

```
Sun Sep 16 01:03:52 1973
```

with a newline included at the end, rather unhelpfully. (`strftime()` will give you more flexibility.)

It's just like `asctime()`, except it takes a `time_t` instead of a `struct tm`.

**WARNING:** This function returns a pointer to a `static char*` region that isn't thread-safe and might be shared with the `asctime()` function. If you need thread safety, use `strftime()` or use a mutex that covers `ctime()` and `asctime()`.

Behavior is undefined for:

- Years less than 1000
- Years greater than 9999
- Any members of `timeptr` are out of range

## Return Value

A pointer to the human-readable local time and data string.

## Example

```
1 time_t now = time(NULL);
2
3 printf("Local: %s", ctime(&now));
```

Sample output:

```
Local: Mon Mar  1 21:32:23 2021
```

## See Also

`asctime()`

---

## 29.10 `gmtime()`

Convert a calendar time into a UTC broken-down time

## Synopsis

```
#include <time.h>

struct tm *gmtime(const time_t *timer);
```

## Description

If you have a `time_t`, you can run it through this function to get a `struct tm` back full of the corresponding broken-down UTC time information.

This is just like `localtime()`, except it does UTC instead of local time.

Once you have that `struct tm`, you can feed it to `strftime()` to print it out.

**WARNING:** This function returns a pointer to a `static struct tm*` region that isn't thread-safe and might be shared with the `localtime()` function. If you need thread safety use a mutex that covers `gmtime()` and `localtime()`.

## Return Value

Returns a pointer to the broken-down UTC time, or `NULL` if it can't be obtained.

## Example

```
1 #include <stdio.h>
2 #include <time.h>
3
4 int main(void)
5 {
6     time_t now = time(NULL);
7
8     printf("UTC   : %s", asctime(gmtime(&now)));
9     printf("Local: %s", asctime(localtime(&now)));
10 }
```

Sample output:

```
UTC   : Tue Mar  2 05:40:05 2021
Local: Mon Mar  1 21:40:05 2021
```

## See Also

`localtime()`, `asctime()`, `strftime()`

---

## 29.11 `localtime()`

Convert a calendar time into a broken-down local time

## Synopsis

```
#include <time.h>

struct tm *localtime(const time_t *timer);
```

## Description

If you have a `time_t`, you can run it through this function to get a `struct tm` back full of the corresponding broken-down local time information.

This is just like `gmtime()`, except it does local time instead of UTC.

Once you have that `struct tm`, you can feed it to `strftime()` to print it out.

**WARNING:** This function returns a pointer to a `static struct tm*` region that isn't thread-safe and might be shared with the `gmtime()` function. If you need thread safety use a mutex that covers `gmtime()` and `localtime()`.

## Return Value

Returns a pointer to the broken-down local time, or `NULL` if it can't be obtained.

## Example

```
1 #include <stdio.h>
2 #include <time.h>
3
4 int main(void)
5 {
6     time_t now = time(NULL);
7
8     printf("Local: %s", asctime(localtime(&now)));
9     printf("UTC : %s", asctime(gmtime(&now)));
10 }
```

Sample output:

```
Local: Mon Mar  1 21:40:05 2021
UTC   : Tue Mar  2 05:40:05 2021
```

## See Also

`gmtime()`, `asctime()`, `strftime()`

---

## 29.12 strftime()

Formatted date and time output

## Synopsis

```
#include <time.h>

size_t strftime(char * restrict s, size_t maxsize,
                const char * restrict format,
                const struct tm * restrict timeptr);
```

## Description

This is the `sprintf()` of date and time functions. It'll take a `struct tm` and produce a string in just about whatever form you desire, for example:

```
2021-03-01
Monday, March 1 at 9:54 PM
It's Monday!
```

It's a super flexible version of `asctime()`. And thread-safe, besides, since it doesn't rely on a static buffer to hold the results.

Basically what you do is give it a destination, `s`, and its max size in bytes in `maxsize`. Also, provide a `format` string that's analogous to `printf()`'s format string, but with different format specifiers. And lastly, a `struct tm` with the broken-down time information to use for printing.

The `format` string works like this, for example:

```
"It's %A, %B %d!"
```

Which produces:

```
It's Monday, March 1!
```

The `%A` is the full day-of-week name, the `%B` is the full month name, and the `%d` is the day of the month. `strftime()` substitutes the right thing to produce the result. Brilliant!

So what are all the format specifiers? Glad you asked!

I'm going to be lazy and just drop this table in right from the spec.

| Specifier | Description                                                                                                                                          |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| %a        | Locale's abbreviated weekday name. [ <code>tm_wday</code> ]                                                                                          |
| %A        | Locale's full weekday name. [ <code>tm_wday</code> ]                                                                                                 |
| %b        | Locale's abbreviated month name. [ <code>tm_mon</code> ]                                                                                             |
| %B        | Locale's full month name. [ <code>tm_mon</code> ]                                                                                                    |
| %c        | Locale's appropriate date and time representation.                                                                                                   |
| %C        | Year divided by 100 and truncated to an integer, as a decimal number (00–99). [ <code>tm_year</code> ]                                               |
| %d        | Day of the month as a decimal number (01–31). [ <code>tm_mday</code> ]                                                                               |
| %D        | Equivalent to "%m/%d/%y". [ <code>tm_mon</code> , <code>tm_mday</code> , <code>tm_year</code> ]                                                      |
| %e        | Day of the month as a decimal number (1–31); a single digit is preceded by a space. [ <code>tm_mday</code> ]                                         |
| %F        | Equivalent to "%Y-%m-%d" (the ISO 8601 date format). [ <code>tm_year</code> , <code>tm_mon</code> , <code>tm_mday</code> ]                           |
| %g        | Last 2 digits of the week-based year (see below) as a decimal number (00–99). [ <code>tm_year</code> , <code>tm_wday</code> , <code>tm_yday</code> ] |

| Specifier | Description                                                                                                                                                                |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| %G        | Week-based year (see below) as a decimal number (e.g., 1997). [tm_year, tm_wday, tm_yday]                                                                                  |
| %h        | Equivalent to “%b”. [tm_mon]                                                                                                                                               |
| %H        | Hour (24-hour clock) as a decimal number (00–23). [tm_hour]                                                                                                                |
| %I        | Hour (12-hour clock) as a decimal number (01–12). [tm_hour]                                                                                                                |
| %j        | Day of the year as a decimal number (001–366). [tm_yday]                                                                                                                   |
| %m        | Month as a decimal number (01–12).                                                                                                                                         |
| %M        | Minute as a decimal number (00–59). [tm_min]                                                                                                                               |
| %n        | A new-line character.                                                                                                                                                      |
| %p        | Locale’s equivalent of the AM/PM designations associated with a 12-hour clock. [tm_hour]                                                                                   |
| %r        | Locale’s 12-hour clock time. [tm_hour, tm_min, tm_sec]                                                                                                                     |
| %R        | Equivalent to “%H:%M”. [tm_hour, tm_min]                                                                                                                                   |
| %S        | Second as a decimal number (00–60). [tm_sec]                                                                                                                               |
| %t        | A horizontal-tab character.                                                                                                                                                |
| %T        | Equivalent to “%H:%M:%S” (the ISO 8601 time format). [tm_hour, tm_min, tm_sec]                                                                                             |
| %u        | ISO 8601 weekday as a decimal number (1–7), where Monday is 1. [tm_wday]                                                                                                   |
| %U        | Week number of the year (the first Sunday as the first day of week 1) as a decimal number (00–53). [tm_year, tm_wday, tm_yday]                                             |
| %V        | ISO 8601 week number (see below) as a decimal number (01–53). [tm_year, tm_wday, tm_yday]                                                                                  |
| %w        | Weekday as a decimal number (0–6), where Sunday is 0.                                                                                                                      |
| %W        | Week number of the year (the first Monday as the first day of week 1) as a decimal number (00–53). [tm_year, tm_wday, tm_yday]                                             |
| %x        | Locale’s appropriate date representation.                                                                                                                                  |
| %X        | Locale’s appropriate time representation.                                                                                                                                  |
| %y        | Last 2 digits of the year as a decimal number (00–99). [tm_year]                                                                                                           |
| %Y        | Year as a decimal number (e.g., 1997). [tm_year]                                                                                                                           |
| %z        | Offset from UTC in the ISO 8601 format “-0430” (meaning 4 hours 30 minutes behind UTC, west of Greenwich), or by no characters if no time zone is determinable. [tm_isdst] |
| %Z        | Locale’s time zone name or abbreviation, or by no characters if no time zone is determinable. [tm_isdst]                                                                   |
| %%        | A plain ol’ %                                                                                                                                                              |

Phew. That’s love.

%G, %g, and %v are a little funky in that they use something called the ISO 8601 week-based year. I’d never heard of it. But, again stealing from the spec, these are the rules:

%g, %G, and %V give values according to the ISO 8601 week-based year. In this system, weeks begin on a Monday and week 1 of the year is the week that includes January 4th, which is also the week that includes the first Thursday of the year, and is also the first week that contains at least four days in the year. If the first Monday of January is the 2nd, 3rd, or 4th, the preceding days are part of the last week of the preceding year; thus, for Saturday 2nd January 1999, %G is replaced by 1998 and %V is replaced by 53. If December 29th, 30th, or 31st is a Monday, it and any following days are part of week 1 of the following year. Thus, for Tuesday 30th December 1997, %G is replaced by 1998 and %V is replaced by 01.

Learn something new every day! If you want to know more, Wikipedia has a page on it<sup>9</sup>.

If you’re in the “C” locale, the specifiers produce the following (again, stolen from the spec):

<sup>9</sup>[https://en.wikipedia.org/wiki/ISO\\_week\\_date](https://en.wikipedia.org/wiki/ISO_week_date)

| Specifier | Description                               |
|-----------|-------------------------------------------|
| %a        | The first three characters of %A.         |
| %A        | One of Sunday, Monday, ... , Saturday.    |
| %b        | The first three characters of %B.         |
| %B        | One of January, February, ... , December. |
| %c        | Equivalent to %a %b %e %T %Y.             |
| %p        | One of AM or PM.                          |
| %r        | Equivalent to %I:%M:%S %p.                |
| %x        | Equivalent to %m/%d/%y.                   |
| %X        | Equivalent to %T.                         |
| %Z        | Implementation-defined.                   |

There are additional variants of the format specifiers that indicate you want to use a locale’s alternative format. These don’t exist for all locales. It’s one of the format specifiers above, with either an **E** or **O** prefix:

```
%Ec %EC %Ex %EX %Ey %EY %Od %Oe %OH %OI
%Om %OM %OS %Ou %OU %OV %Ow %OW %Oy
```

The **E** and **O** prefixes are ignored in the “C” locale.

## Return Value

Returns the total number of bytes put into the result string, not including the NUL terminator.

If the result doesn’t fit in the string, zero is returned and the value in **s** is indeterminate.

## Example

```
1  #include <stdio.h>
2  #include <time.h>
3
4  int main(void)
5  {
6      char s[128];
7      time_t now = time(NULL);
8
9      // %c: print date as per current locale
10     strftime(s, sizeof s, "%c", localtime(&now));
11     puts(s);    // Sun Feb 28 22:29:00 2021
12
13     // %A: full weekday name
14     // %B: full month name
15     // %d: day of the month
16     strftime(s, sizeof s, "%A, %B %d", localtime(&now));
17     puts(s);    // Sunday, February 28
18
19     // %I: hour (12 hour clock)
20     // %M: minute
21     // %S: second
22     // %p: AM or PM
23     strftime(s, sizeof s, "It's %I:%M:%S %p", localtime(&now));
24     puts(s);    // It's 10:29:00 PM
```

```
25
26     // %F: ISO 8601 yyyy-mm-dd
27     // %T: ISO 8601 hh:mm:ss
28     // %z: ISO 8601 time zone offset
29     strftime(s, sizeof s, "ISO 8601: %FT%T%Z", localtime(&now));
30     puts(s);    // ISO 8601: 2021-02-28T22:29:00-0800
31 }
```

## See Also

`ctime()`, `asctime()`

## Chapter 30

# <uchar.h> Unicode utility functions

| Function   | Description                                              |
|------------|----------------------------------------------------------|
| c16rtomb() | Convert a <code>char16_t</code> to a multibyte character |
| c32rtomb() | Convert a <code>char32_t</code> to a multibyte character |
| mbrtoc16() | Convert a multibyte character to a <code>char16_t</code> |
| mbrtoc32() | Convert a multibyte character to a <code>char32_t</code> |

These functions are *restartable*, meaning multiple threads can safely call them at once. They handle this by having their own conversion state variable (of type `mbstate_t`) per call.

### 30.1 Types

This header file defines four types.

| Type                   | Description                                                                                          |
|------------------------|------------------------------------------------------------------------------------------------------|
| <code>char16_t</code>  | Type to hold 16-bit characters                                                                       |
| <code>char32_t</code>  | Type to hold 32-bit characters                                                                       |
| <code>mbstate_t</code> | Holds the conversion state for restartable functions (also defined in <code>&lt;wchar.h&gt;</code> ) |
| <code>size_t</code>    | To hold various counts (also defined in <code>&lt;stddef.h&gt;</code> )                              |

String literals for the character types are `u` for `char16_t` and `U` for `char32_t`.

```
char16_t *str1 = u"Hello, world!";
char32_t *str2 = U"Hello, world!";

char16_t *chr1 = u'A';
char32_t *chr2 = U'B';
```

Note that `char16_t` and `char32_t` *might* contain Unicode. Or not. If `__STDC_UTF_16__` or `__STDC_UTF_32__` is defined as `1`, then `char16_t` and `char32_t` use Unicode, respectively. Otherwise they don't and the actual value stored depend on the locale. And if you're not using Unicode, you have my commiserations.



## 30.2 OS X issue

This header file doesn't exist on OS X—bummer. If you just want the types, you can:

```
#include <stdint.h>

typedef int_least16_t char16_t;
typedef int_least32_t char32_t;
```

But if you also want the functions, that's all on you.

---

## 30.3 mbrtoc16() mbrtoc32()

Convert a multibyte character to a `char16_t` or `char32_t` restartably

### Synopsis

```
#include <uchar.h>

size_t mbrtoc16(char16_t * restrict pc16, const char * restrict s, size_t n,
                mbstate_t * restrict ps);

size_t mbrtoc32(char32_t * restrict pc32, const char * restrict s, size_t n,
                mbstate_t * restrict ps);
```

### Description

Given a source string `s` and a destination buffer `pc16` (or `pc32` for `mbrtoc32()`), convert the first character of the source to `char16_t`s (or `char32_t`s for `mbrtoc32()`).

Basically you have a regular character and you want it as `char16_t` or `char32_t`. Use these functions to do it. Note that only one character is converted no matter how many characters in `s`.

As the functions scan `s`, you don't want them to overrun the end. So you pass in `n` as the maximum number of bytes to inspect. The functions will quit after that many bytes or when they have a complete multibyte character, whichever comes first.

Since they're restartable, pass in a conversion state variable for the functions to do their work.

And the result will be placed in `pc16` (or `pc32` for `mbrtoc32()`).

### Return Value

When successful this function returns a number between `1` and `n` inclusive representing the number of bytes that made up the multibyte character.

Or, also in the success category, they can return `0` if the source character is the NUL character (value `0`).

When not entirely successful, they can return a variety of codes. These are all of type `size_t`, but negative values cast to that type.

| Return Value | Description                                                                                             |
|--------------|---------------------------------------------------------------------------------------------------------|
| (size_t)(-1) | Encoding error—this isn't a valid sequence of bytes. <code>errno</code> is set to <code>EILSEQ</code> . |
| (size_t)(-2) | <code>n</code> bytes were examined and were a <i>partial</i> valid character, but not a complete one.   |
| (size_t)(-3) | A subsequent value of a character that can't be represented as a single value. See below.               |

Case `(size_t)(-3)` is an odd one. Basically there are some characters that can't be represented with 16 bits and so can't be stored in a `char16_t`. These characters are store in something called (in the Unicode world) *surrogate pairs*. That is, there are *two* 16-bit values back to back that represent a larger Unicode value.

For example, if you want to read the Unicode character `\U0001fbc5` (which is a stick figure<sup>1</sup>—I'm just not putting it in the text because my font doesn't render it) that's more than 16 bits. But each call to `mbrtoc16()` only returns a single `char16_t`!

So subsequent calls to `mbrtoc16()` resolves the *next* value in the surrogate pair and returns `(size_t)(-3)` to let you know this has happened.

You can also pass `NULL` for `pc16` or `pc32`. This will cause no result to be stored, but you can use it if you're only interested in the return value from the functions.

Finally, if you pass `NULL` for `s`, the call is equivalent to:

```
mbrtoc16(NULL, "", 1, ps)
```

Since the character is a NUL in that case, this has the effect of setting the state in `ps` to the initial conversion state.

## Example

Normal use case example where we get the first two character values from the multibyte string `"€Zillion"`:

```
#include <uchar.h>
#include <stdio.h>    // for printf()
#include <locale.h>    // for setlocale()
#include <string.h>    // for memset()

int main(void)
{
    char *s = "\u20acZillion"; // 20ac is "€"
    char16_t pc16;
    size_t r;
    mbstate_t mbs;

    setlocale(LC_ALL, "");
    memset(&mbs, 0, sizeof mbs);

    // Examine the next 8 bytes to see if there's a character in there
    r = mbrtoc16(&pc16, s, 8, &mbs);

    printf("%zu\n", r);    // Prints a value >= 1 (3 in UTF-8 locale)
    printf("%#x\n", pc16); // Prints 0x20ac for "€"
```

<sup>1</sup>[https://en.wikipedia.org/wiki/Symbols\\_for\\_Legacy\\_Computing](https://en.wikipedia.org/wiki/Symbols_for_Legacy_Computing)

```

    s += r; // Move to next character

    // Examine the next 8 bytes to see if there's a character in there
    r = mbrtoc16(&pc16, s, 8, &mbs);

    printf("%zu\n", r); // Prints 1
    printf("%#x\n", pc16); // Prints 0x5a for "Z"
}

```

Example with a surrogate pair. In this case we read plenty to get the entire character, but the result must be stored in two `char16_ts`, requiring two calls to get them both.

```

1  #include <uchar.h>
2  #include <stdio.h> // for printf()
3  #include <string.h> // for memset()
4  #include <locale.h> // for setlocale()
5
6  int main(void)
7  {
8      char *s = "\U0001fbc5"; // Stick figure glyph, more than 16 bits
9      char16_t pc16;
10     mbstate_t mbs;
11     size_t r;
12
13     setlocale(LC_ALL, "");
14     memset(&mbs, 0, sizeof mbs);
15
16     r = mbrtoc16(&pc16, s, 8, &mbs);
17
18     printf("%zd\n", r); // r is 4 bytes in UTF-8 locale
19     printf("%#x\n", pc16); // First value of surrogate pair
20
21     s += r; // Move to next character
22
23     r = mbrtoc16(&pc16, s, 8, &mbs);
24
25     printf("%zd\n", r); // r is (size_t)(-3) here to indicate...
26     printf("%#x\n", pc16); // ...Second value of surrogate pair
27
28     // Since r is -3, it means we're still processing the same
29     // character, so DON'T move to the next character this time
30     //s += r; // Commented out
31
32     r = mbrtoc16(&pc16, s, 8, &mbs);
33
34     printf("%zd\n", r); // 1 byte for ""
35     printf("%#x\n", pc16); // 0x2a for ""
36 }

```

Output on my system, indicating the first character is represented by the pair (0xd83e, 0xdfc5) and the second character is represented by 0x2a:

```
4
0xd83e
-3
0xdfc5
1
0x2a
```

## See Also

`c16rtomb()`, `c32rtomb()`

---

## 30.4 `c16rtomb()` `c32rtomb()`

Convert a `char16_t` or `char32_t` to a multibyte character restartably

### Synopsis

```
#include <uchar.h>

size_t c16rtomb(char * restrict s, char16_t c16, mbstate_t * restrict ps);

size_t c32rtomb(char * restrict s, char32_t c32, mbstate_t * restrict ps);
```

### Description

If you have a character in a `char16_t` or `char32_t`, use these functions to convert them into a multibyte character.

These functions figure out how many bytes are needed for the multibyte character in the current locale and stores them in the buffer pointed to by `s`.

But how big to make that buffer? Luckily there is a macro to help: it needs be no larger than `MB_CUR_MAX`.

As a special case, if `s` is `NULL`, it's the same as calling

```
c16rtomb(buf, L'\0', ps); // or...
c32rtomb(buf, L'\0', ps);
```

where `buf` is a buffer maintained by the system that you don't have access to.

This has the effect of setting the `ps` state to the initial state.

Finally for surrogate pairs (where the character has been split into two `char16_ts`), you call this once with the first of the pair—at this point, the function will return `0`. Then you call it again with the second of the pair, and the function will return the number of bytes and store the result in the array `s`.

### Return Value

Returns the number of bytes stored in the array pointed to by `s`.

Returns `0` if processing is not yet complete for the current character, as in the case of surrogate pairs.

If there is an encoding error, the functions return `(size_t)(-1)` and `errno` is set to `EILSEQ`.

## Example

```

1  #include <uchar.h>
2  #include <stdlib.h> // for MB_CUR_MAX
3  #include <stdio.h>  // for printf()
4  #include <string.h> // for memset()
5  #include <locale.h> // for setlocale()
6
7  int main(void)
8  {
9      char16_t c16 = 0x20ac; // Unicode for Euro symbol
10     char dest[MB_CUR_MAX];
11     size_t r;
12     mbstate_t mbs;
13
14     setlocale(LC_ALL, "");
15     memset(&mbs, 0, sizeof mbs); // Reset conversion state
16
17     // Convert
18     r = c16rtomb(dest, c16, &mbs);
19
20     printf("r == %zd\n", r); // r == 3 on my system
21
22     // And this should print a Euro symbol
23     printf("dest == \"%s\"\n", dest);
24 }

```

Output on my system:

```

r == 3
dest == "€"

```

This is a more complex example that converts a large-valued character in a multibyte string into a surrogate pair (as in the `mbrtoc16()` example, above) and then converts it back again into a multibyte string to print.

```

1  #include <uchar.h>
2  #include <stdlib.h> // for MB_CUR_MAX
3  #include <stdio.h>  // for printf()
4  #include <string.h> // for memset()
5  #include <locale.h> // for setlocale()
6
7  int main(void)
8  {
9      char *src = "\U0001fbc5*"; // Stick figure glyph, more than 16 bits
10     char dest[MB_CUR_MAX];
11     char16_t surrogate0, surrogate1;
12     mbstate_t mbs;
13     size_t r;
14
15     setlocale(LC_ALL, "");
16     memset(&mbs, 0, sizeof mbs); // Reset conversion state
17
18     // Get first surrogate character

```

```

19     r = mbrtoc16(&surrogate0, src, 8, &mbs);
20
21     // Get next surrogate character
22     src += r; // Move to next character
23     r = mbrtoc16(&surrogate1, src, 8, &mbs);
24
25     printf("Surrogate pair: %#x, %#x\n", surrogate0, surrogate1);
26
27     // Now reverse it
28     memset(&mbs, 0, sizeof mbs); // Reset conversion state
29
30     // Process first surrogate character
31     r = c16rtomb(dest, surrogate0, &mbs);
32
33     // r should be 0 at this point, because the character hasn't been
34     // processed yet. And dest won't have anything useful... yet!
35     printf("r == %zd\n", r); // r == 0
36
37     // Process second surrogate character
38     r = c16rtomb(dest, surrogate1, &mbs);
39
40     // Now we should be in business. r should have the number of
41     // bytes, and dest should hold the character.
42     printf("r == %zd\n", r); // r == 4 on my system
43
44     // And this should print a stick figure, if your font supports it
45     printf("dest == \"%s\"\n", dest);
46 }

```

## See Also

`mbrtoc16()`, `mbrtoc32()`

## Chapter 31

# <wchar.h> Wide Character Handling

| Function                 | Description                                                          |
|--------------------------|----------------------------------------------------------------------|
| <code>btowc()</code>     | Convert a single byte character to a wide character                  |
| <code>fgetwc()</code>    | Get a wide character from a wide stream                              |
| <code>fgetws()</code>    | Read a wide string from a wide stream                                |
| <code>fputwc()</code>    | Write a wide character to a wide stream                              |
| <code>fputws()</code>    | Write a wide string to a wide stream                                 |
| <code>fwide()</code>     | Get or set the orientation of the stream                             |
| <code>fwprintf()</code>  | Formatted wide output to a wide stream                               |
| <code>fwscanf()</code>   | Formatted wide input from a wide stream                              |
| <code>getwchar()</code>  | Get a wide character from <code>stdin</code>                         |
| <code>getwc()</code>     | Get a wide character from <code>stdin</code>                         |
| <code>mbrlen()</code>    | Compute the number of bytes in a multibyte character restartably     |
| <code>mbrtowc()</code>   | Convert multibyte to wide characters restartably                     |
| <code>mbsinit()</code>   | Test if an <code>mbstate_t</code> is in the initial conversion state |
| <code>mbstowcs()</code>  | Convert a multibyte string to a wide character string restartably    |
| <code>putwchar()</code>  | Write a wide character to <code>stdout</code>                        |
| <code>putwc()</code>     | Write a wide character to <code>stdout</code>                        |
| <code>swprintf()</code>  | Formatted wide output to a wide string                               |
| <code>swscanf()</code>   | Formatted wide input from a wide string                              |
| <code>ungetwc()</code>   | Pushes a wide character back into the input stream                   |
| <code>vfwprintf()</code> | Variadic formatted wide output to a wide stream                      |
| <code>vfwscanf()</code>  | Variadic formatted wide input from a wide stream                     |
| <code>vswprintf()</code> | Variadic formatted wide output to a wide string                      |
| <code>vswscanf()</code>  | Variadic formatted wide input from a wide string                     |
| <code>vwprintf()</code>  | Variadic formatted wide output                                       |
| <code>vwscanf()</code>   | Variadic formatted wide input                                        |
| <code>wscat()</code>     | Concatenate wide strings dangerously                                 |
| <code>wcschr()</code>    | Find a wide character in a wide string                               |
| <code>wscmp()</code>     | Compare wide strings                                                 |
| <code>wscoll()</code>    | Compare two wide strings accounting for locale                       |
| <code>wscpy()</code>     | Copy a wide string dangerously                                       |
| <code>wscspn()</code>    | Count characters not from a start at the front of a wide string      |
| <code>wcsftime()</code>  | Formatted date and time output                                       |
| <code>wcslen()</code>    | Returns the length of a wide string                                  |
| <code>wcsncat()</code>   | Concatenate wide strings more safely                                 |
| <code>wcsncmp()</code>   | Compare wide strings, length limited                                 |

| Function    | Description                                                       |
|-------------|-------------------------------------------------------------------|
| wcsncpy()   | Copy a wide string more safely                                    |
| wcsprk()    | Search a wide string for one of a set of wide characters          |
| wcsrchr()   | Find a wide character in a wide string from the end               |
| wcsrtombs() | Convert a wide character string to a multibyte string restartably |
| wcsspn()    | Count characters from a set at the front of a wide string         |
| wcsstr()    | Find a wide string in another wide string                         |
| wcstod()    | Convert a wide string to a double                                 |
| wcstof()    | Convert a wide string to a float                                  |
| wcstok()    | Tokenize a wide string                                            |
| wcstold()   | Convert a wide string to a long double                            |
| wcstoll()   | Convert a wide string to a long long                              |
| wcstol()    | Convert a wide string to a long                                   |
| wcstoull()  | Convert a wide string to an unsigned long long                    |
| wcstoul()   | Convert a wide string to an unsigned long                         |
| wcsxfrm()   | Transform a wide string for comparing based on locale             |
| wctob()     | Convert a wide character to a single byte character               |
| wctombr()   | Convert wide to multibyte characters restartably                  |
| wmemcmp()   | Compare wide characters in memory                                 |
| wmemcpy()   | Copy wide character memory                                        |
| wmemmove()  | Copy wide character memory, potentially overlapping               |
| wprintf()   | Formatted wide output                                             |
| wscanf()    | Formatted wide input                                              |

These are the wide character variants of the functions found in <stdio.h>.

Remember that you can't mix-and-match multibyte output functions (like `printf()`) with wide character output functions (like `wprintf()`). The output stream has an *orientation* to either multibyte or wide that gets set on the first I/O call to that stream. (Or it can be set with `fwide()`.)

So choose one or the other and stick with it.

And you can specify wide character constants and string literals by prefixing `L` to the front of it:

```
wchar_t *s = L"Hello, world!";
wchar_t c = L'B';
```

This header also introduces a type `wint_t` that is used by the character I/O functions. It's a type that can hold any single wide character, but *also* the macro `WEOF` to indicate wide end-of-file.

## 31.1 Restartable Functions

Finally, a note on the “restartable” functions that are included here. When conversion is happening, some encodings require C to keep track of some *state* about the progress of the conversion so far.

For a lot of the functions, C uses an internal variable for the state that is shared between function calls. The problem is if you're writing multithreaded code, this state might get trampled by other threads.

To avoid this, each thread needs to maintain its own state in a variable of the opaque type `mbstate_t`. And the “restartable” functions allow you to pass in this state so that each thread can use their own.



## 31.2 wprintf(), fwprintf(), swprintf()

Formatted output with a wide string

### Synopsis

```
#include <stdio.h>    // For fwprintf()
#include <wchar.h>

int wprintf(const wchar_t * restrict format, ...);

int fwprintf(FILE * restrict stream, const wchar_t * restrict format, ...);

int swprintf(wchar_t * restrict s, size_t n,
             const wchar_t * restrict format, ...);
```

### Description

These are the wide versions of `printf()`, `fprintf()` (`#man-printf`), and `sprintf()`.

See those pages for exact substantial usage.

These are the same except the `format` string is a wide character string instead of a multibyte string.

And that `swprintf()` is analogous to `snprintf()` in that they both take the size of the destination array as an argument.

And one more thing: the precision specified for a `%s` specifier corresponds to the number of wide characters printed, not the number of bytes. If you know of other difference, let me know.

### Return Value

Returns the number of wide characters outputted, or `-1` if there's an error.

### Example

```
1  #include <stdio.h>
2  #include <wchar.h>
3
4  int main(void)
5  {
6      char *mbs = "multibyte";
7      wchar_t *ws = L"wide";
8
9      wprintf(L"We're all wide for %s and %ls!\n", mbs, ws);
10
11     double pi = 3.14159265358979;
12     wprintf(L"pi = %f\n", pi);
13 }
```

Output:

```
We're all wide for multibyte and wide!  
pi = 3.141593
```

## See Also

`printf()`, `vwprintf()`

---

## 31.3 `wscanf()` `fwscanf()` `swscanf()`

Scan a wide stream or wide string for formatted input

### Synopsis

```
#include <stdio.h> // for fwscanf()  
#include <wchar.h>  
  
int wscanf(const wchar_t * restrict format, ...);  
  
int fwscanf(FILE * restrict stream, const wchar_t * restrict format, ...);  
  
int swscanf(const wchar_t * restrict s, const wchar_t * restrict format, ...);
```

### Description

These are the wide variants of `scanf()`, `fscanf()`, and `sscanf()`.

See the `scanf()` page for all the details.

### Return Value

Returns the number of items successfully scanned, or `EOF` on some kind of input failure.

### Example

```
1  #include <stdio.h>  
2  #include <wchar.h>  
3  
4  int main(void)  
5  {  
6      int quantity;  
7      wchar_t item[100];  
8  
9      wprintf(L"Enter \"quantity: item\\n");  
10  
11     if (wscanf(L"%d:%99ls", &quantity, item) != 2)  
12         wprintf(L"Malformed input!\\n");  
13     else  
14         wprintf(L"You entered: %d %ls\\n", quantity, item);  
15 }
```

Output (input of 12: apples):

```
Enter "quantity: item"
12: apples
You entered: 12 apples
```

## See Also

`scanf()`, `vwscanf()`

---

## 31.4 `wprintf()` `vfwprintf()` `vswprintf()`

`wprintf()` variants using variable argument lists (`va_list`)

### Synopsis

```
#include <stdio.h>    // For vfwprintf()
#include <stdarg.h>
#include <wchar.h>

int wprintf(const wchar_t * restrict format, va_list arg);

int vswprintf(wchar_t * restrict s, size_t n,
              const wchar_t * restrict format, va_list arg);

int vfwprintf(FILE * restrict stream, const wchar_t * restrict format,
              va_list arg);
```

### Description

These functions are the wide character variants of the `vprintf()`, functions. You can refer to that reference page for more details.

### Return Value

Returns the number of wide characters stored, or a negative value on error.

### Example

In this example, we make our own version of `wprintf()` called `wlogger()` that timestamps output. Notice how the calls to `wlogger()` have all the bells and whistles of `wprintf()`.

```
1 #include <stdarg.h>
2 #include <wchar.h>
3 #include <time.h>
4
5 int wlogger(wchar_t *format, ...)
6 {
7     va_list va;
```

```

8     time_t now_secs = time(NULL);
9     struct tm *now = gmtime(&now_secs);
10
11     // Output timestamp in format "YYYY-MM-DD hh:mm:ss : "
12     wprintf(L"%04d-%02d-%02d %02d:%02d:%02d : ",
13            now->tm_year + 1900, now->tm_mon + 1, now->tm_mday,
14            now->tm_hour, now->tm_min, now->tm_sec);
15
16     va_start(va, format);
17     int result = vwprintf(format, va);
18     va_end(va);
19
20     wprintf(L"\n");
21
22     return result;
23 }
24
25 int main(void)
26 {
27     int x = 12;
28     float y = 3.2;
29
30     wlogger(L"Hello!");
31     wlogger(L"x = %d and y = %.2f", x, y);
32 }

```

Output:

```

2021-03-30 04:25:49 : Hello!
2021-03-30 04:25:49 : x = 12 and y = 3.20

```

## See Also

`printf()`, `vprintf()`

## 31.5 `wscanf()`, `vfwscanf()`, `vswscanf()`

`wscanf()` variants using variable argument lists (`va_list`)

### Synopsis

```

#include <stdio.h>    // For vfwscanf()
#include <stdarg.h>
#include <wchar.h>

int wscanf(const wchar_t * restrict format, va_list arg);

int vfwscanf(FILE * restrict stream, const wchar_t * restrict format,
             va_list arg);

```

```
int vswscanf(const wchar_t * restrict s, const wchar_t * restrict format,
            va_list arg);
```

## Description

These are the wide counterparts to the `vscanf()` collection of functions. See their reference page for details.

## Return Value

Returns the number of items successfully scanned, or `EOF` on some kind of input failure.

## Example

I have to admit I was wracking my brain to think of when you'd ever want to use this. The best example I could find was one on Stack Overflow<sup>1</sup> that error-checks the return value from `scanf()` against the expected. A variant of that is shown below.

```
1  #include <stdarg.h>
2  #include <wchar.h>
3  #include <assert.h>
4
5  int error_check_wscanf(int expected_count, wchar_t *format, ...)
6  {
7      va_list va;
8
9      va_start(va, format);
10     int count = vwscanf(format, va);
11     va_end(va);
12
13     // This line will crash the program if the condition is false:
14     assert(count == expected_count);
15
16     return count;
17 }
18
19 int main(void)
20 {
21     int a, b;
22     float c;
23
24     error_check_wscanf(3, L"%d, %d/%f", &a, &b, &c);
25     error_check_wscanf(2, L"%d", &a);
26 }
```

## See Also

`wscanf()`

<sup>1</sup><https://stackoverflow.com/questions/17017331/c99-vscanf-for-dummies/17018046#17018046>

## 31.6 getwc() fgetwc() getwchar()

Get a wide character from an input stream

### Synopsis

```
#include <stdio.h>    // For getwc() and fgetwc()
#include <wchar.h>

wint_t getwchar(void);

wint_t getwc(FILE *stream);

wint_t fgetwc(FILE *stream);
```

### Description

These are the wide variants of `fgetc()`.

`fgetwc()` and `getwc()` are identical except that `getwc()` might be implemented as a macro and is allowed to evaluate `stream` multiple times.

`getwchar()` is identical to `getwc()` with `stream` set to `stdin`.

I don't know why you'd ever use `getwc()` instead of `fgetwc()`, but if anyone knows, drop me a line.

### Return Value

Returns the next wide character in the input stream. Return `WEOF` on end-of-file or error.

If an I/O error occurs, the error flag is also set on the stream.

If an invalid byte sequence is encountered, `errno` is set to `ILSEQ`.

### Example

Reads all the characters from a file, outputting only the letter 'b's it finds in the file:

```
1  #include <stdio.h>
2  #include <wchar.h>
3
4  int main(void)
5  {
6      FILE *fp;
7      wint_t c;
8
9      fp = fopen("datafile.txt", "r"); // error check this!
10
11     // this while-statement assigns into c, and then checks against EOF:
12
13     while((c = fgetc(fp)) != WEOF)
14         if (c == L'b')
15             fputwc(c, stdout);
16
```

```
17     fclose(fp);
18 }
```

## See Also

fputwc, fgetws, errno

---

## 31.7 fgetws()

Read a wide string from a file

### Synopsis

```
#include <stdio.h>
#include <wchar.h>

wchar_t *fgetws(wchar_t * restrict s, int n, FILE * restrict stream);
```

### Description

This is the wide version of `fgets()`. See its reference page for details.

A wide `NUL` character is used to terminate the string.

### Return Value

Returns `s` on success, or a `NULL` pointer on end-of-file or error.

### Example

The following example reads lines from a file and prepends them with numbers:

```
1  #include <stdio.h>
2  #include <wchar.h>
3
4  #define BUF_SIZE 1024
5
6  int main(void)
7  {
8      FILE *fp;
9      wchar_t buf[BUF_SIZE];
10
11     fp = fopen("textfile.txt", "r"); // error check this!
12
13     int line_count = 0;
14
15     while ((fgetws(buf, BUF_SIZE, fp)) != NULL)
16         wprintf(L"%04d: %ls", ++line_count, buf);
17 }
```

```

18     fclose(fp);
19 }

```

Example output for a file with these lines in them (without the prepended numbers):

```

0001: line 1
0002: line 2
0003: something
0004: line 4

```

## See Also

`fgetwc()`, `fgets()`

---

## 31.8 putwchar() putwc() fputwc()

Write a single wide character to the console or to a file

### Synopsis

```

#include <stdio.h>    // For putwc() and fputwc()
#include <wchar.h>

wint_t putwchar(wchar_t c);

wint_t putwc(wchar_t c, FILE *stream);

wint_t fputwc(wchar_t c, FILE *stream);

```

### Description

These are the wide character equivalents to the ‘fputc()’ group of functions. You can find more information ‘in that reference section’.

`fputwc()` and `putwc()` are identical except that `putwc()` might be implemented as a macro and is allowed to evaluate `stream` multiple times.

`putwchar()` is identical to `putwc()` with `stream` set to `stdin`.

I don’t know why you’d ever use `putwc()` instead of `fputwc()`, but if anyone knows, drop me a line.

### Return Value

Returns the wide character written, or `WEOF` on error.

If it’s an I/O error, the error flag will be set for the stream.

If it’s an encoding error, `errno` will be set to `EILSEQ`.

### Example

Read all characters from a file, outputting only the letter ‘b’s it finds in the file:



```
1  #include <stdio.h>
2  #include <wchar.h>
3
4  int main(void)
5  {
6      FILE *fp;
7      wint_t c;
8
9      fp = fopen("datafile.txt", "r"); // error check this!
10
11     // this while-statement assigns into c, and then checks against EOF:
12
13     while((c = fgetc(fp)) != WEOF)
14         if (c == L'b')
15             fputwc(c, stdout);
16
17     fclose(fp);
18 }
```

## See Also

fgetwc(), fputc(), errno

---

## 31.9 fputws()

Write a wide string to a file

### Synopsis

```
#include <stdio.h>
#include <wchar.h>

int fputws(const wchar_t * restrict s, FILE * restrict stream);
```

### Description

This is the wide version of fputs().

Pass in a wide string and an output stream, and it will so be written.

### Return Value

Returns a non-negative value on success, or EOF on error.

## Example

```

1  #include <stdio.h>
2  #include <wchar.h>
3
4  int main(void)
5  {
6      fputws(L"Hello, world!\n", stdout);
7  }

```

## See Also

fputwc() fputs()

## 31.10 fwide()

Get or set the orientation of the stream

### Synopsis

```

#include <stdio.h>
#include <wchar.h>

int fwide(FILE *stream, int mode);

```

### Description

Streams can be either wide-oriented (meaning the wide functions are in use) or byte-oriented (that the regular multibyte functions are in use). Or, before an orientation is chosen, unoriented.

There are two ways to set the orientation of an unoriented stream:

- Implicitly: just use a function like `printf()` (byte oriented) or `wprintf()` (wide oriented), and the orientation will be set.
- Explicitly: use this function to set it.

You can set the orientation for the stream by passing different numbers to `mode`:

| mode | Description                  |
|------|------------------------------|
| 0    | Do not alter the orientation |
| -1   | Set stream to byte-oriented  |
| 1    | Set stream to wide-oriented  |

(I said `-1` and `1` there, but really it could be any positive or negative number.)

Most people choose the wide or byte functions (`printf()` or `wprintf()`) and just start using them and never use `fwide()` to set the orientation.

And once the orientation is set, you can't change it. So you can't use `fwide()` for that, either.

So what can you use it for?

You can *test* to see what orientation a stream is in by passing `0` as the `mode` and checking the return value.

## Return Value

Returns greater than zero if the stream is wide-oriented.

Returns less than zero if the stream is byte-oriented.

Returns zero if the stream is unoriented.

## Example

Example setting to byte-oriented:

```
1 #include <stdio.h>
2 #include <wchar.h>
3
4 int main(void)
5 {
6     printf("Hello world!\n"); // Implicitly set to byte
7
8     int mode = fwide(stdout, 0);
9
10    printf("Stream is %s-oriented\n", mode < 0? "byte": "wide");
11 }
```

Output:

```
Hello world!
Stream is byte-oriented
```

Example setting to wide-oriented:

```
1 #include <stdio.h>
2 #include <wchar.h>
3
4 int main(void)
5 {
6     wprintf(L"Hello world!\n"); // Implicitly set to wide
7
8     int mode = fwide(stdout, 0);
9
10    wprintf(L"Stream is %ls-oriented\n", mode < 0? L"byte": L"wide");
11 }
```

Output:

```
Hello world!
Stream is wide-oriented
```

---

## 31.11 ungetwc()

Pushes a wide character back into the input stream

### Synopsis

```
#include <stdio.h>
#include <wchar.h>

wint_t ungetwc(wint_t c, FILE *stream);
```

### Description

This is the wide character variant of `ungetc()`.

It performs the reverse operation of `fgetwc()`, pushing a character back on the input stream.

The spec guarantees you can do this one time in a row. You can probably do it more times, but it's up to the implementation. If you do too many calls without an intervening read, an error could be returned.

Setting the file position discards any characters pushed by `ungetwc()` without being subsequently read.

The end-of-file flag is cleared after a successful call.

### Return Value

Returns the value of the pushed character on success, or `WEOF` on failure.

### Example

This example reads a piece of punctuation, then everything after it up to the next piece of punctuation. It returns the leading punctuation, and stores the rest in a string.

```
1  #include <stdio.h>
2  #include <wctype.h>
3  #include <wchar.h>
4
5  wint_t read_punctstring(FILE *fp, wchar_t *s)
6  {
7      wint_t origpunct, c;
8
9      origpunct = fgetwc(fp);
10
11     if (origpunct == WEOF) // return EOF on end-of-file
12         return WEOF;
13
14     while (c = fgetwc(fp), !iswpunct(c) && c != WEOF)
15         *s++ = c; // save it in the string
16
17     *s = L'\0'; // nul-terminate the string
18
19     // if we read punctuation last, ungetc it so we can fgetc it next
20     // time:
21     if (iswpunct(c))
```

```
22     ungetwc(c, fp);
23
24     return origpunct;
25 }
26
27 int main(void)
28 {
29     wchar_t s[128];
30     wint_t c;
31
32     while ((c = read_punctstring(stdin, s)) != WEOF) {
33         wprintf(L"%lc: %ls\n", c, s);
34     }
35 }
```

Sample Input:

```
!foo#bar*baz
```

Sample output:

```
!: foo
#: bar
*: baz
```

## See Also

`fgetwc()`, `ungetc()`

---

## 31.12 `wcstod()` `wcstof()` `wcstold()`

Convert a wide string to a floating point number

### Synopsis

```
#include <wchar.h>

double wcstod(const wchar_t * restrict nptr, wchar_t ** restrict endptr);

float wcstof(const wchar_t * restrict nptr, wchar_t ** restrict endptr);

long double wcstold(const wchar_t * restrict nptr, wchar_t ** restrict endptr);
```

### Description

These are the wide counterparts to the `strtod()` family of functions. See their reference pages for details.

## Return Value

Returns the string converted to a floating point value.

Returns `0` if there's no valid number in the string.

On overflow, returns an appropriately-signed `HUGE_VAL`, `HUGE_VALF`, or `HUGE_VALL` depending on the return type, and `errno` is set to `ERANGE`.

On underflow, returns a number no greater than the smallest normalized positive number, appropriately signed. The implementation *might* set `errno` to `ERANGE`.

## Example

```

1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t *inp = L" 123.4567beej";
6      wchar_t *badchar;
7
8      double val = wcstod(inp, &badchar);
9
10     wprintf(L"Converted string to %f\n", val);
11     wprintf(L"Encountered bad characters: %ls\n", badchar);
12
13     val = wcstod(L"987.654321beej", NULL);
14     wprintf(L"Ignoring bad chars: %f\n", val);
15
16     val = wcstod(L"11.2233", &badchar);
17
18     if (*badchar == L'\0')
19         wprintf(L"No bad chars: %f\n", val);
20     else
21         wprintf(L"Found bad chars: %f, %ls\n", val, badchar);
22 }
```

Output:

```

Converted string to 123.456700
Encountered bad characters: beej
Ignoring bad chars: 987.654321
No bad chars: 11.223300
```

## See Also

`wcstol()`, `strtod()`, `errno`

## 31.13 `wcstol()` `wcstoll()` `wcstoul()` `wcstoull()`

Convert a wide string to an integer value

## Synopsis

```
#include <wchar.h>

long int wcstol(const wchar_t * restrict nptr,
                wchar_t ** restrict endptr, int base);

long long int wcstoll(const wchar_t * restrict nptr,
                     wchar_t ** restrict endptr, int base);

unsigned long int wcstoul(const wchar_t * restrict nptr,
                        wchar_t ** restrict endptr, int base);

unsigned long long int wcstoull(const wchar_t * restrict nptr,
                              wchar_t ** restrict endptr, int base);
```

## Description

These are the wide counterparts to the `strtol()` family of functions, so see their reference pages for the details.

## Return Value

Returns the integer value of the string.

If nothing can be found, 0 is returned.

If the result is out of range, the value returned is one of `LONG_MIN`, `LONG_MAX`, `LLONG_MIN`, `LLONG_MAX`, `ULONG_MAX` or `ULLONG_MAX`, as appropriate. And `errno` is set to `ERANGE`.

## Example

```
1  #include <wchar.h>
2
3  int main(void)
4  {
5      // All output in decimal (base 10)
6
7      wprintf(L"%ld\n", wcstol(L"123", NULL, 0));    // 123
8      wprintf(L"%ld\n", wcstol(L"123", NULL, 10));   // 123
9      wprintf(L"%ld\n", wcstol(L"101010", NULL, 2)); // binary, 42
10     wprintf(L"%ld\n", wcstol(L"123", NULL, 8));    // octal, 83
11     wprintf(L"%ld\n", wcstol(L"123", NULL, 16));   // hex, 291
12
13     wprintf(L"%ld\n", wcstol(L"0123", NULL, 0));    // octal, 83
14     wprintf(L"%ld\n", wcstol(L"0x123", NULL, 0));  // hex, 291
15
16     wchar_t *badchar;
17     long int x = wcstol(L" 1234beej", &badchar, 0);
18
19     wprintf(L"Value is %ld\n", x);                  // Value is 1234
20     wprintf(L"Bad chars at \"%ls\"\n", badchar);    // Bad chars at "beej"
21 }
```

Output:

```
123
123
42
83
291
83
291
Value is 1234
Bad chars at "beej"
```

## See Also

`wcstod()`, `strtol()`, `errno`, `wcstoimax()`, `wcstoumax()`

---

## 31.14 `wscpy()` `wscncpy()`

Copy a wide string

### Synopsis

```
#include <wchar.h>

wchar_t *wscpy(wchar_t * restrict s1, const wchar_t * restrict s2);

wchar_t *wscncpy(wchar_t * restrict s1,
                  const wchar_t * restrict s2, size_t n);
```

### Description

These are the wide versions of `strcpy()` and `strncpy()`.

They'll copy a string up to a wide NUL. Or, in the case of the safer `wscncpy()`, until then or until `n` wide characters are copied.

If the string in `s1` is shorter than `n`, `wscncpy()` will pad `s2` with wide NUL characters until the `n`th wide character is reached.

Even though `wscncpy()` is safer because it will never overrun the end of `s2` (assuming you set `n` correctly), it's still unsafe a NUL is not found in `s1` in the first `n` characters. In that case, `s2` will not be NUL-terminated. Always make sure `n` is greater than the string length of `s1`!

### Return Value

Returns `s1`.



## Example

```

1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t *s1 = L"Hello!";
6      wchar_t s2[10];
7
8      wcsncpy(s2, s1, 10);
9
10     wprintf(L"%ls\n", s2); // "Hello!"
11 }

```

## See Also

wmemcpy(), wmemmove() strcpy(), strncpy()

---

## 31.15 wmemcpy() wmemmove()

Copy wide characters

## Synopsis

```

#include <wchar.h>

wchar_t *wmemcpy(wchar_t * restrict s1,
                 const wchar_t * restrict s2, size_t n);

wchar_t *wmemmove(wchar_t *s1, const wchar_t *s2, size_t n);

```

## Description

These are the wide versions of memcpy() and memmove().

They copy `n` wide characters from `s2` to `s1`.

They're the same except that `wmemmove()` is guaranteed to work with overlapping memory regions, and `wmemcpy()` is not.

## Return Value

Both functions return the pointer `s1`.

## Example

```

1  #include <wchar.h>
2
3  int main(void)
4  {

```

```

5   wchar_t s[100] = L"Goats";
6   wchar_t t[100];
7
8   wmemcpy(t, s, 6);           // Copy non-overlapping memory
9
10  wmemmove(s + 2, s, 6);      // Copy overlapping memory
11
12  wprintf(L"s is \"%ls\"\n", s);
13  wprintf(L"t is \"%ls\"\n", t);
14 }

```

Output:

```

s is "GoGoats"
t is "Goats"

```

## See Also

`wscpy()`, `wcsncpy()`, `memcpy()`, `memmove()`

---

## 31.16 `wscat()` `wcsncat()`

Concatenate wide strings

### Synopsis

```

#include <wchar.h>

wchar_t *wscat(wchar_t * restrict s1, const wchar_t * restrict s2);

wchar_t *wcsncat(wchar_t * restrict s1,
                 const wchar_t * restrict s2, size_t n);

```

### Description

These are the wide variants of `strcat()` and `strncat()`.

They concatenate `s2` onto the end of `s1`.

They're the same except `wcsncat()` gives you the option to limit the number of wide characters appended.

Note that `wcsncat()` always adds a NUL terminator to the end, even if `n` characters were appended. So be sure to leave room for that.

### Return Value

Both functions return the pointer `s1`.

## Example

```

1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t dest[30] = L"Hello";
6      wchar_t *src = L", World!";
7      wchar_t numbers[] = L"12345678";
8
9      wprintf(L"dest before strcat: \"%ls\"\n", dest); // "Hello"
10
11     wscat(dest, src);
12     wprintf(L"dest after strcat: \"%ls\"\n", dest); // "Hello, world!"
13
14     wcsncat(dest, numbers, 3); // strcat first 3 chars of numbers
15     wprintf(L"dest after strncat: \"%ls\"\n", dest); // "Hello, world!123"
16 }

```

## See Also

strcat(), strncat()

## 31.17 wcsncmp(), wcsncmp(), wmemcmp()

Compare wide strings or memory

### Synopsis

```

#include <wchar.h>

int wcsncmp(const wchar_t *s1, const wchar_t *s2);

int wcsncmp(const wchar_t *s1, const wchar_t *s2, size_t n);

int wmemcmp(const wchar_t *s1, const wchar_t *s2, size_t n);

```

### Description

These are the wide variants of `memcmp()`, `strcmp()`, and `strncmp()`.

`wscmp()` and `wcsncmp()` both compare strings until a NUL character.

`wcsncmp()` also has the additional restriction that it will only compare the first `n` characters.

`wmemcmp()` is like `wcsncmp()` except it won't stop at a NUL.

The comparison is done against the character value (which might (or might not) be its Unicode code point).

### Return Value

Returns zero if both regions are equal.

Returns a negative number if the region pointed to by `s1` is less than `s2`.

Returns a positive number if the region pointed to by `s1` is greater than `s2`.

## Example

```

1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t *s1 = L"Muffin";
6      wchar_t *s2 = L"Muffin Sandwich";
7      wchar_t *s3 = L"Muffin";
8
9      wprintf(L"%d\n", wcscmp(L"Biscuits", L"Kittens")); // <0 since 'B' < 'K'
10     wprintf(L"%d\n", wcscmp(L"Kittens", L"Biscuits")); // >0 since 'K' > 'B'
11
12     if (wcscmp(s1, s2) == 0)
13         wprintf(L"This won't get printed because the strings differ\n");
14
15     if (wcscmp(s1, s3) == 0)
16         wprintf(L"This will print because s1 and s3 are the same\n");
17
18     // this is a little weird...but if the strings are the same, it'll
19     // return zero, which can also be thought of as "false". Not-false
20     // is "true", so (!wcscmp()) will be true if the strings are the
21     // same. yes, it's odd, but you see this all the time in the wild
22     // so you might as well get used to it:
23
24     if (!wcscmp(s1, s3))
25         wprintf(L"The strings are the same!\n");
26
27     if (!wcsncmp(s1, s2, 6))
28         wprintf(L"The first 6 characters of s1 and s2 are the same\n");
29 }
```

Output:

```

-1
1
This will print because s1 and s3 are the same
The strings are the same!
The first 6 characters of s1 and s2 are the same
```

## See Also

`wscoll()`, `memcmp()`, `strcmp()`, `strncmp()`

## 31.18 `wscoll()`

Compare two wide strings accounting for locale

## Synopsis

```
#include <wchar.h>

int wcscoll(const wchar_t *s1, const wchar_t *s2);
```

## Description

This is the wide version of `strcoll()`. See [that reference page](#) for details.

This is slower than `wcscmp()`, so only use it if you need the locale-specific compare.

## Return Value

Returns zero if both regions are equal in this locale.

Returns a negative number if the region pointed to by `s1` is less than `s2` in this locale.

Returns a positive number if the region pointed to by `s1` is greater than `s2` in this locale.

## Example

```
1  #include <wchar.h>
2  #include <locale.h>
3
4  int main(void)
5  {
6      setlocale(LC_ALL, "");
7
8      // If your source character set doesn't support "é" in a string
9      // you can replace it with '\u00e9', the Unicode code point
10     // for "é".
11
12     wprintf(L"%d\n", wcscmp(L"é", L"f")); // Reports é > f, yuck.
13     wprintf(L"%d\n", wcs.coll(L"é", L"f")); // Reports é < f, yay!
14 }
```

## See Also

`wcscmp()`, `wcsxfrm()`, `strcoll()`

---

## 31.19 `wcsxfrm()`

Transform a wide string for comparing based on locale

## Synopsis

```
#include <wchar.h>

size_t wcsxfrm(wchar_t * restrict s1,
               const wchar_t * restrict s2, size_t n);
```

## Description

This is the wide variant of `strxfrm()`. See that reference page for details.

## Return Value

Returns the length of the transformed wide string in wide characters.

If the return value is greater than `n`, all bets are off for the result in `s1`.

## Example

```
1  #include <wchar.h>
2  #include <locale.h>
3  #include <stdlib.h>
4
5  // Transform a string for comparison, returning a malloc'd
6  // result
7  wchar_t *get_xfrm_str(wchar_t *s)
8  {
9      int len = wcsxfrm(NULL, s, 0) + 1;
10     wchar_t *d = malloc(len * sizeof(wchar_t));
11
12     wcsxfrm(d, s, len);
13
14     return d;
15 }
16
17 // Does half the work of a regular wcscoll() because the second
18 // string arrives already transformed.
19 int half_wcscoll(wchar_t *s1, wchar_t *s2_transformed)
20 {
21     wchar_t *s1_transformed = get_xfrm_str(s1);
22
23     int result = wcscmp(s1_transformed, s2_transformed);
24
25     free(s1_transformed);
26
27     return result;
28 }
29
30 int main(void)
31 {
32     setlocale(LC_ALL, "");
33 }
```

```

34 // Pre-transform the string to compare against
35 wchar_t *s = get_xfrm_str(L"éfg");
36
37 // Repeatedly compare against "éfg"
38 wprintf(L"%d\n", half_wcscoll(L"fgh", s)); // "fgh" > "éfg"
39 wprintf(L"%d\n", half_wcscoll(L"àbc", s)); // "àbc" < "éfg"
40 wprintf(L"%d\n", half_wcscoll(L"âij", s)); // "âij" > "éfg"
41
42 free(s);
43 }

```

Output:

```

1
-1
1

```

## See Also

wscmp(), wcscoll(), strxfrm()

## 31.20 wcschr() wcsrchr()

Find a wide character in a wide string

### Synopsis

```

#include <wchar.h>

// Pre-C23:

wchar_t *wcschr(const wchar_t *s, wchar_t c);

wchar_t *wcsrchr(const wchar_t *s, wchar_t c);

wchar_t *wmemchr(const wchar_t *s, wchar_t c, size_t n);

// C23:

QWchar_t *wcschr(QWchar_t *s, wchar_t c);

QWchar_t *wcsrchr(QWchar_t *s, wchar_t c);

QWchar_t *wmemchr(QWchar_t *s, wchar_t c, size_t n);

```

### Description

These are the wide equivalents to `strchr()`, `strrchr()`, and `memchr()`.

They search for wide characters in a wide string from the front (`wcschr()`), the end (`wcsrchr()`) or for an arbitrary number of wide characters (`wmemchr()`).

## Return Value

All three functions return a pointer to the wide character found, or `NULL` if the character, sadly, isn't found.

## Example

```

1  #include <wchar.h>
2
3  int main(void)
4  {
5      // "Hello, world!"
6      //      ^  ^  ^
7      //      A  B  C
8
9      wchar_t *str = L"Hello, world!";
10     wchar_t *p;
11
12     p = wcschr(str, ',');           // p now points at position A
13     p = wcsrchr(str, 'o');         // p now points at position B
14
15     p = wmemchr(str, '!', 13);      // p now points at position C
16
17     // repeatedly find all occurrences of the letter 'B'
18     str = L"A BIG BROWN BAT BIT BEEJ";
19
20     for(p = wcschr(str, 'B'); p != NULL; p = wcschr(p + 1, 'B')) {
21         wprintf(L"Found a 'B' here: %ls\n", p);
22     }
23 }
```

Output:

```

Found a 'B' here: BIG BROWN BAT BIT BEEJ
Found a 'B' here: BROWN BAT BIT BEEJ
Found a 'B' here: BAT BIT BEEJ
Found a 'B' here: BIT BEEJ
Found a 'B' here: BEEJ
```

## See Also

`strchr()`, `strrchr()`, `memchr()`

## 31.21 wcssp() wscspn()

Return the length of a wide string consisting entirely of a set of wide characters, or of not a set of wide characters



## Synopsis

```
#include <wchar.h>

size_t wcsspncpy(const wchar_t *s1, const wchar_t *s2);

size_t wcscspncpy(const wchar_t *s1, const wchar_t *s2);
```

## Description

They are the wide character counterparts to `strspn()` (`#man-strspn`) and `strcspn()`.

They compute the length of the string pointed to by `s1` consisting entirely of the characters found in `s2`. Or, in the case of `wcscspncpy()`, the characters *not* found in `s2`.

## Return Value

The length of the string pointed to by `s1` consisting solely of the characters in `s2` (in the case of `wcsspncpy()`) or of the characters *not* in `s2` (in the case of `wcscspncpy()`).

## Example

```
1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t str1[] = L"a banana";
6      wchar_t str2[] = L"the bolivian navy on maneuvers in the south pacific";
7      int n;
8
9      // how many letters in str1 until we reach something that's not a vowel?
10     n = wcsspncpy(str1, L"aeiou");
11     wprintf(L"%d\n", n); // n == 1, just "a"
12
13     // how many letters in str1 until we reach something that's not a, b,
14     // or space?
15     n = wcsspncpy(str1, L"ab ");
16     wprintf(L"%d\n", n); // n == 4, "a ba"
17
18     // how many letters in str2 before we get a "y"?
19     n = wcscspncpy(str2, L"y");
20     wprintf(L"%d\n", n); // n = 16, "the bolivian nav"
21 }
```

## See Also

`wcschr()`, `wcsrchr()`, `strspn()`

---

## 31.22 wcsrbrk()

Search a wide string for one of a set of wide characters

### Synopsis

```
#include <wchar.h>

// Pre-C23:

wchar_t *wcsrbrk(const wchar_t *s1, const wchar_t *s2);

// C23:

Qwchar_t *wcsrbrk(Qwchar_t *s1, const wchar_t *s2);
```

### Description

This is the wide character variant of `strpbrk()`.

It finds the first occurrence of any of a set of wide characters in a wide string.

### Return Value

Returns a pointer to the first character in the string `s1` that exists in the string `s2`.

Or `NULL` if none of the characters in `s2` can be found in `s1`.

### Example

```
1  #include <wchar.h>
2
3  int main(void)
4  {
5      // p points here after wcsrbrk
6      //           v
7      wchar_t *s1 = L"Hello, world!";
8      wchar_t *s2 = L"dow!"; // Match any of these chars
9
10     wchar_t *p = wcsrbrk(s1, s2); // p points to the o
11
12     wprintf(L"%ls\n", p); // "o, world!"
13 }
```

### See Also

`wcschr()`, `wmemchr()`, `strpbrk()`

---

## 31.23 wcsstr()

Find a wide string in another wide string

## Synopsis

```
#include <wchar.h>

// Pre-C23:
wchar_t *wcsstr(const wchar_t *s1, const wchar_t *s2);

// C23:
QWchar_t *wcsstr(QWchar_t *s1, const wchar_t *s2);
```

## Description

This is the wide variant of `strstr()`.

It locates a substring in a string.

## Return Value

Returns a pointer to the location in `s1` that contains `s2`.

Or `NULL` if `s2` cannot be found in `s1`.

## Example

```
1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t *str = L"The quick brown fox jumped over the lazy dogs.";
6      wchar_t *p;
7
8      p = wcsstr(str, L"lazy");
9      wprintf(L"%ls\n", p == NULL? L"null": p); // "lazy dogs."
10
11     // p is NULL after this, since the string "wombat" isn't in str:
12     p = wcsstr(str, L"wombat");
13     wprintf(L"%ls\n", p == NULL? L"null": p); // "null"
14 }
```

## See Also

`wcschr()`, `wcsrchr()`, `wcsspn()`, `wcscspn()`, `strstr()`

---

## 31.24 wcstok()

Tokenize a wide string

## Synopsis

```
#include <wchar.h>
wchar_t *wcstok(wchar_t * restrict s1, const wchar_t * restrict s2,
                wchar_t ** restrict ptr);
```

## Description

This is the wide version of `strtok()`.

And, like that one, it modifies the string `s1`. So make a copy of it first if you want to preserve the original.

One key difference is that `wcstok()` can be threadsafe because you pass in the pointer `ptr` to the current state of the transformation. This gets initializers for you when `s1` is initially passed in as non-NULL. (Subsequent calls with a NULL `s1` cause the state to update.)

## Return Value

## Example

```
1  #include <wchar.h>
2
3  int main(void)
4  {
5      // break up the string into a series of space or
6      // punctuation-separated words
7      wchar_t str[] = L"Where is my bacon, dude?";
8      wchar_t *token;
9      wchar_t *state;
10
11     // Note that the following if-do-while construct is very very
12     // very very very common to see when using strtok().
13
14     // grab the first token (making sure there is a first token!)
15     if ((token = wcstok(str, L".,?! ", &state)) != NULL) {
16         do {
17             wprintf(L"Word: \"%ls\\n\"", token);
18
19             // now, the while continuation condition grabs the
20             // next token (by passing NULL as the first param)
21             // and continues if the token's not NULL:
22         } while ((token = wcstok(NULL, L".,?! ", &state)) != NULL);
23     }
24 }
```

Output:

```
Word: "Where"
Word: "is"
Word: "my"
Word: "bacon"
Word: "dude"
```

**See Also**`strtok()`

---

**31.25 wcslen()**

Returns the length of a wide string

**Synopsis**

```
#include <wchar.h>

size_t wcslen(const wchar_t *s);
```

**Description**

This is the wide counterpart to `strlen()`.

**Return Value**

Returns the number of wide characters before the wide NUL terminator.

**Example**

```
1  #include <wchar.h>
2
3  int main(void)
4  {
5      wchar_t *s = L"Hello, world!"; // 13 characters
6
7      // prints "The string is 13 characters long.":
8
9      wprintf(L"The string is %zu characters long.\n", wcslen(s));
10 }
```

**See Also**`strlen()`

---

**31.26 wcsftime()**

Formatted date and time output

## Synopsis

```
#include <time.h>
#include <wchar.h>

size_t wcsftime(wchar_t * restrict s, size_t maxsize,
                const wchar_t * restrict format,
                const struct tm * restrict timeptr);
```

## Description

This is the wide equivalent to `strftime()`. See that reference page for details.

`maxsize` here refers to the maximum number of wide characters that can be in the result string.

## Return Value

If successful, returns the number of wide characters written.

If not successful because the result couldn't fit in the space allotted, 0 is returned and the contents of the string could be anything.

## Example

```
1  #include <wchar.h>
2  #include <time.h>
3
4  #define BUFSIZE 128
5
6  int main(void)
7  {
8      wchar_t s[BUFSIZE];
9      time_t now = time(NULL);
10
11     // %c: print date as per current locale
12     wcsftime(s, BUFSIZE, L"%c", localtime(&now));
13     wprintf(L"%ls\n", s);    // Sun Feb 28 22:29:00 2021
14
15     // %A: full weekday name
16     // %B: full month name
17     // %d: day of the month
18     wcsftime(s, BUFSIZE, L"%A, %B %d", localtime(&now));
19     wprintf(L"%ls\n", s);    // Sunday, February 28
20
21     // %I: hour (12 hour clock)
22     // %M: minute
23     // %S: second
24     // %p: AM or PM
25     wcsftime(s, BUFSIZE, L"It's %I:%M:%S %p", localtime(&now));
26     wprintf(L"%ls\n", s);    // It's 10:29:00 PM
27
28     // %F: ISO 8601 yyyy-mm-dd
29     // %T: ISO 8601 hh:mm:ss
```

```

30 // %z: ISO 8601 time zone offset
31 wcsftime(s, BUFSIZE, L"ISO 8601: %FT%T%Z", localtime(&now));
32 wprintf(L"%ls\n", s); // ISO 8601: 2021-02-28T22:29:00-0800
33 }

```

## See Also

strftime()

## 31.27 btowc() wctob()

Convert a single byte character to a wide character

### Synopsis

```

#include <wchar.h>

wint_t btowc(int c);

int wctob(wint_t c);

```

### Description

These functions convert between single byte characters and wide characters, and vice-versa.

Even though `ints` are involved, don't let this mislead you; they're effectively converted to `unsigned chars` internally.

The characters in the basic character set are guaranteed to be a single byte.

### Return Value

`btowc()` returns the single-byte character as a wide character. Returns `WEOF` if `EOF` is passed in, or if the byte doesn't correspond to a valid wide character.

`wctob()` returns the wide character as a single-byte character. Returns `EOF` if `WEOF` is passed in, or if the wide character doesn't correspond to a value single-byte character.

See `mbtowc()` and `wctomb()` for multibyte to wide character conversion.

### Example

```

1 #include <wchar.h>
2
3 int main(void)
4 {
5     wint_t wc = btowc('B'); // Convert single byte to wide char
6
7     wprintf(L"Wide character: %lc\n", wc);
8 }

```

```

9     unsigned char c = wctob(wc); // Convert back to single byte
10
11     wprintf(L"Single-byte character: %c\n", c);
12 }

```

Output:

```

Wide character: B
Single-byte character: B

```

## See Also

`mbtowc()`, `wctomb()`

## 31.28 mbsinit()

Test if an `mbstate_t` is in the initial conversion state

### Synopsis

```

#include <wchar.h>

int mbsinit(const mbstate_t *ps);

```

### Description

For a given conversion state in a `mbstate_t` variable, this function determines if it's in the initial conversion state.

### Return Value

Returns non-zero if the value pointed to by `ps` is in the initial conversion state, or if `ps` is `NULL`.

Returns 0 if the value pointed to by `ps` is **not** in the initial conversion state.

### Example

For me, this example doesn't do anything exciting, saying that the `mbstate_t` variable is always in the initial state. Yay.

But if have a stateful encoding like 2022-JP, try messing around with this to see if you can get into an intermediate state.

This program has a bit of code at the top that reports if your locale's encoding requires any state.

```

1  #include <locale.h> // For setlocale()
2  #include <string.h> // For memset()
3  #include <stdlib.h> // For mbtowc()
4  #include <wchar.h>
5

```



```

6  int main(void)
7  {
8      mbstate_t state;
9      wchar_t wc[128];
10
11     setlocale(LC_ALL, "");
12
13     int is_state_dependent = mbtowc(NULL, NULL, 0);
14
15     wprintf(L"Is encoding state dependent? %d\n", is_state_dependent);
16
17     memset(&state, 0, sizeof state); // Set to initial state
18
19     wprintf(L"In initial conversion state? %d\n", mbsinit(&state));
20
21     mbrtowc(wc, "B", 5, &state);
22
23     wprintf(L"In initial conversion state? %d\n", mbsinit(&state));
24 }

```

## See Also

`mbtowc()`, `wctomb()`, `mbrtowc()`, `wcrtomb()`

---

## 31.29 `mbrlen()`

Compute the number of bytes in a multibyte character, restartably

### Synopsis

```

#include <wchar.h>

size_t mbrlen(const char * restrict s, size_t n, mbstate_t * restrict ps);

```

### Description

This is the restartable version of `mblen()`.

It inspects at most `n` bytes of the string `s` to see how many bytes in this character.

The conversion state is stored in `ps`.

This function doesn't have the functionality of `mblen()` that allowed you to query if this character encoding was stateful and to reset the internal state.

### Return Value

Returns the number of bytes required for this multibyte character.

Returns `(size_t)(-1)` if the data in `s` is not a valid multibyte character.

Returns `(size_t)(-2)` if the data in `s` is a valid but not complete multibyte character.

## Example

If your character set doesn't support the Euro symbol “€”, substitute the Unicode escape sequence `\u20ac`, below.

```
1  #include <locale.h>    // For setlocale()
2  #include <string.h>    // For memset()
3  #include <wchar.h>
4
5  int main(void)
6  {
7      mbstate_t state;
8      int len;
9
10     setlocale(LC_ALL, "");
11
12     memset(&state, 0, sizeof state); // Set to initial state
13
14     len = mbrlen("B", 5, &state);
15
16     wprintf(L"Length of 'B' is %d byte(s)\n", len);
17
18     len = mbrlen("€", 5, &state);
19
20     wprintf(L"Length of '€' is %d byte(s)\n", len);
21 }
```

Output:

```
Length of 'B' is 1 byte(s)
Length of '€' is 3 byte(s)
```

## See Also

`mblen()`

---

## 31.30 `mbrtowc()`

Convert multibyte to wide characters restartably

### Synopsis

```
#include <wchar.h>

size_t mbrtowc(wchar_t * restrict pwc, const char * restrict s,
               size_t n, mbstate_t * restrict ps);
```

### Description

This is the restartable counterpart to `mbtowc()`.

It converts individual characters from multibyte to wide, tracking the conversion state in the variable pointed to by `ps`.

At most `n` bytes are inspected for conversion to a wide character.

These two variants are identical and cause the state pointed to by `ps` to be set to the initial conversion state:

```
mbrtowc(NULL, NULL, 0, &state);
mbrtowc(NULL, "", 1, &state);
```

Also, if you're just interested in the length in bytes of the multibyte character, you can pass `NULL` for `pwc` and nothing will be stored for the wide character:

```
int len = mbrtowc(NULL, "€", 5, &state);
```

This function doesn't have the functionality of `mbtowc()` that allowed you to query if this character encoding was stateful and to reset the internal state.

## Return Value

On success, returns a positive number corresponding to the number of bytes in the multibyte character.

Returns `0` if the character encoded is a wide NUL character.

Returns `(size_t)(-1)` if the data in `s` is not a valid multibyte character.

Returns `(size_t)(-2)` if the data in `s` is a valid but not complete multibyte character.

## Example

If your character set doesn't support the Euro symbol “€”, substitute the Unicode escape sequence `\u20ac`, below.

```
1  #include <string.h> // For memset()
2  #include <stdlib.h> // For mbrtowc()
3  #include <locale.h> // For setlocale()
4  #include <wchar.h>
5
6  int main(void)
7  {
8      mbstate_t state;
9
10     memset(&state, 0, sizeof state);
11
12     setlocale(LC_ALL, "");
13
14     wprintf(L"State dependency: %d\n", mbrtowc(NULL, NULL, 0));
15
16     wchar_t wc;
17     int bytes;
18
19     bytes = mbrtowc(&wc, "€", 5, &state);
20
21     wprintf(L"L'%lc' takes %d bytes as multibyte char '€'\n", wc, bytes);
22 }
```

Output on my system:

```
State dependency: 0
L'€' takes 3 bytes as multibyte char '€'
```

## See Also

`mbtowc()`, `wrtomb()`

### 31.31 `wrtomb()`

Convert wide to multibyte characters restartably

#### Synopsis

```
#include <wchar.h>

size_t wrtomb(char * restrict s, wchar_t wc, mbstate_t * restrict ps);
```

#### Description

This is the restartable counterpart to `wctomb()`.

It converts individual characters from wide to multibyte, tracking the conversion state in the variable pointed to by `ps`.

The destination array `s` should be at least `MB_CUR_MAX`<sup>2</sup> bytes in size—you won't get anything bigger back from this function.

Note that the values in this result array won't be NUL-terminated.

If you pass a wide NUL character in, the result will contain any bytes needed to restore the conversion state to its initial state followed by a NUL character, and the state pointed to by `ps` will be reset to its initial state:

```
// Reset state
wrtomb(mb, L'\0', &state)
```

If you don't care about the results (i.e. you're just interested in resetting the state or getting the return value), you can do this by passing `NULL` for `s`:

```
wrtomb(NULL, L'\0', &state); // Reset state

int byte_count = wrtomb(NULL, "X", &state); // Count bytes in 'X'
```

This function doesn't have the functionality of `wctomb()` that allowed you to query if this character encoding was stateful and to reset the internal state.

<sup>2</sup>This is a variable, not a macro, so if you use it to define an array, it'll be a variable-length array.

## Return Value

On success, returns the number of bytes needed to encode this wide character in the current locale.

If the input is an invalid wide character, `errno` will be set to `EILSEQ` and the function returns `(size_t)(-1)`. If this happens, all bets are off for the conversion state, so you might as well reset it.

## Example

If your character set doesn't support the Euro symbol “€”, substitute the Unicode escape sequence `\u20ac`, below.

```

1  #include <string.h> // For memset()
2  #include <stdlib.h> // For mbtowc()
3  #include <locale.h> // For setlocale()
4  #include <wchar.h>
5
6  int main(void)
7  {
8      mbstate_t state;
9
10     memset(&state, 0, sizeof state);
11
12     setlocale(LC_ALL, "");
13
14     wprintf(L"State dependency: %d\n", mbtowc(NULL, NULL, 0));
15
16     char mb[10] = {0};
17     int bytes = wctomb(mb, L'€', &state);
18
19     wprintf(L"L'€' takes %d bytes as multibyte char '%s'\n", bytes, mb);
20 }
```

## See Also

`mbrtowc()`, `wctomb()`, `errno`

## 31.32 mbsrtowcs()

Convert a multibyte string to a wide character string restartably

### Synopsis

```

#include <wchar.h>

size_t mbsrtowcs(wchar_t * restrict dst, const char ** restrict src,
                 size_t len, mbstate_t * restrict ps);
```

### Description

This is the restartable version of `mbstowcs()`.

It converts a multibyte string to a wide character string.

The result is put in the buffer pointed to by `dst`, and the pointer `src` is updated to indicate how much of the string was consumed (unless `dst` is `NULL`).

At most `len` wide characters will be stored.

This also takes a pointer to its own `mbstate_t` variable in `ps` for holding the conversion state.

You can set `dst` to `NULL` if you only care about the return value. This could be useful for getting the number of characters in a multibyte string.

In the normal case, the `src` string will be consumed up to the NUL character, and the results will be stored in the `dst` buffer, including the wide NUL character. In this case, the pointer pointed to by `src` will be set to `NULL`. And the conversion state will be set to the initial conversion state.

If things go wrong because the source string isn't a valid sequence of characters, conversion will stop and the pointer pointed to by `src` will be set to the address just after the last successfully-translated multibyte character.

## Return Value

If successful, returns the number of characters converted, not including any NUL terminator.

If the multibyte sequence is invalid, the function returns `(size_t)(-1)` and `errno` is set to `EILSEQ`.

## Example

Here we'll convert the string "€5 ± π" into a wide character string:

```

1  #include <locale.h> // For setlocale()
2  #include <string.h> // For memset()
3  #include <wchar.h>
4
5  #define WIDE_STR_SIZE 10
6
7  int main(void)
8  {
9      const char *mbs = "€5 ± π"; // That's the exact price range
10
11     wchar_t wcs[WIDE_STR_SIZE];
12
13     setlocale(LC_ALL, "");
14
15     mbstate_t state;
16     memset(&state, 0, sizeof state);
17
18     size_t count = mbsrtowcs(wcs, &mbs, WIDE_STR_SIZE, &state);
19
20     wprintf(L"Wide string L\"%ls\" is %d characters\n", wcs, count);
21 }
```

Output:

```
Wide string L"€5 ± π" is 6 characters
```

Here's another example of using `mbsrtowcs()` to get the length in characters of a multibyte string even if the string is full of multibyte characters. This is in contrast to `strlen()`, which returns the total number of bytes in the string.

```

1  #include <stdio.h>    // For printf()
2  #include <locale.h>   // For setlocale()
3
4  #include <string.h>    // For memset()
5  #include <stdint.h>    // For SIZE_MAX
6  #include <wchar.h>
7
8  size_t mbstrlen(const char *mbs)
9  {
10     mbstate_t state;
11
12     memset(&state, 0, sizeof state);
13
14     return mbsrtowcs(NULL, &mbs, SIZE_MAX, &state);
15 }
16
17 int main(void)
18 {
19     setlocale(LC_ALL, "");
20
21     char *mbs = "€5 ± π"; // That's the exact price range
22
23     printf("\'%s\' is %zu characters...\n", mbs, mbstrlen(mbs));
24     printf("but it's %zu bytes!\n", strlen(mbs));
25 }

```

Output on my system:

```

"€5 ± π" is 6 characters...
but it's 10 bytes!

```

## See Also

`mbrtowc()`, `mbstowcs()`, `wcsrtombs()`, `strlen()`, `errno`

## 31.33 `wcsrtombs()`

Convert a wide character string to a multibyte string restartably

### Synopsis

```

#include <wchar.h>

size_t wcsrtombs(char * restrict dst, const wchar_t ** restrict src,
                 size_t len, mbstate_t * restrict ps);

```

## Description

If you have a wide character string, you can convert it to a multibyte character string in the current locale using this function.

At most `len` bytes of data will be stored in the buffer pointed to by `dst`. Conversion will stop just after the NUL terminator is copied, or `len` bytes get copied, or some other error occurs.

If `dst` is a `NULL` pointer, no result is stored. You might do this if you're just interested in the return value (nominally the number of bytes this would use in a multibyte string, not including the NUL terminator).

If `dst` is not a `NULL` pointer, the pointer pointed to by `src` will get modified to indicate how much of the data was copied. If it contains `NULL` at the end, it means everything went well. In this case, the state `ps` will be set to the initial conversion state.

If `len` was reached or an error occurred, it'll point one address past `dst+len`.

## Return Value

If everything goes well, returns the number of bytes needed for the multibyte string, not counting the NUL terminator.

If any character in the string doesn't correspond to a valid multibyte character in the currently locale, it returns `(size_t)(-1)` and `EILSEQ` is stored in `errno`.

## Example

Here we'll convert the wide string "€5 ± π" into a multibyte character string:

```

1  #include <locale.h> // For setlocale()
2  #include <string.h> // For memset()
3  #include <wchar.h>
4
5  #define MB_STR_SIZE 20
6
7  int main(void)
8  {
9      const wchar_t *wcs = L"€5 ± π"; // That's the exact price range
10
11     char mbs[MB_STR_SIZE];
12
13     setlocale(LC_ALL, "");
14
15     mbstate_t state;
16     memset(&state, 0, sizeof state);
17
18     size_t count = wcsrtombs(mbs, &wcs, MB_STR_SIZE, &state);
19
20     wprintf(L"Multibyte string \"%s\" is %d bytes\n", mbs, count);
21 }
```

Here's another example helper function that `malloc()`s just enough memory to hold the converted string, then returns the result. (Which must later be freed, of course, to prevent leaking memory.)



```

1  #include <stdlib.h> // For malloc()
2  #include <locale.h> // For setlocale()
3  #include <string.h> // For memset()
4  #include <stdint.h> // For SIZE_MAX
5  #include <wchar.h>
6
7  char *get_mb_string(const wchar_t *wcs)
8  {
9      setlocale(LC_ALL, "");
10
11      mbstate_t state;
12      memset(&state, 0, sizeof state);
13
14      // Need a copy of this because wcsrtombs changes it
15      const wchar_t *p = wcs;
16
17      // Compute the number of bytes needed to hold the result
18      size_t bytes_needed = wcsrtombs(NULL, &p, SIZE_MAX, &state);
19
20      // If we didn't get a good full conversion, forget it
21      if (bytes_needed == (size_t)(-1))
22          return NULL;
23
24      // Allocate space for result
25      char *mbs = malloc(bytes_needed + 1); // +1 for NUL terminator
26
27      // Set conversion state to initial state
28      memset(&state, 0, sizeof state);
29
30      // Convert and store result
31      wcsrtombs(mbs, &wcs, bytes_needed + 1, &state);
32
33      // Make sure things went well
34      if (wcs != NULL) {
35          free(mbs);
36          return NULL;
37      }
38
39      // Success!
40      return mbs;
41  }
42
43  int main(void)
44  {
45      char *mbs = get_mb_string(L"€5 ± π");
46
47      wprintf(L"Multibyte result: \"%s\\n\"", mbs);
48
49      free(mbs);
50  }

```

## See Also

`wcrtomb()`, `wcstombs()`, `mbsrtowcs()`, `errno`

## Chapter 32

# <wctype.h> Wide Character Classification and Transformation

| Function                 | Description                                              |
|--------------------------|----------------------------------------------------------|
| <code>iswalnum()</code>  | Test if a wide character is alphanumeric.                |
| <code>iswalpha()</code>  | Tests if a wide character is alphabetic                  |
| <code>iswblank()</code>  | Tests if this is a wide blank character                  |
| <code>iswcntrl()</code>  | Tests if this is a wide control character.               |
| <code>iswctype()</code>  | Determine wide character classification                  |
| <code>iswdigit()</code>  | Test if this wide character is a digit                   |
| <code>iswgraph()</code>  | Test to see if a wide character is a printable non-space |
| <code>iswlower()</code>  | Tests if a wide character is lowercase                   |
| <code>iswprint()</code>  | Tests if a wide character is printable                   |
| <code>iswpunct()</code>  | Test if a wide character is punctuation                  |
| <code>iswspace()</code>  | Test if a wide character is whitespace                   |
| <code>iswupper()</code>  | Tests if a wide character is uppercase                   |
| <code>iswxdigit()</code> | Tests if a wide character is a hexadecimal digit         |
| <code>towctrans()</code> | Convert wide characters to upper or lowercase            |
| <code>tolower()</code>   | Convert an uppercase wide character to lowercase         |
| <code>toupper()</code>   | Convert a lowercase wide character to uppercase          |
| <code>wctrans()</code>   | Helper function for <code>towctrans()</code>             |
| <code>wctype()</code>    | Helper function for <code>iswctype()</code>              |

This is like `<ctype.h>` except for wide characters.

With it you can test for character classifications (like “is this character whitespace?”) or do basic character conversions (like “force this character to lowercase”).

### 32.1 `iswalnum()`

Test if a wide character is alphanumeric.

## Synopsis

```
#include <wctype.h>

int iswalnum(wint_t wc);
```

## Description

Basically tests if a character is alphabetic (A-Z or a-z) or a digit (0-9). But some other characters might also qualify based on the locale.

This is equivalent to testing if `iswalpha()` or `iswdigit()` is true.

## Return Value

Returns true if the character is alphanumeric.

## Example

```
1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      wprintf(L"%ls\n", iswalnum(L'a')? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswalnum(L'B')? L"yes": L"no"); // yes
10     wprintf(L"%ls\n", iswalnum(L'5')? L"yes": L"no"); // yes
11     wprintf(L"%ls\n", iswalnum(L'')? L"yes": L"no"); // no
12 }
```

## See Also

`iswalpha()`, `iswdigit()`, `isalnum()`

---

## 32.2 `iswalpha()`

Tests if a wide character is alphabetic

## Synopsis

```
#include <wctype.h>

int iswalpha(wint_t wc);
```

## Description

Basically tests if a character is alphabetic (A-Z or a-z). But some other characters might also qualify based on the locale. (If other characters qualify, they won't be control characters, digits, punctuation, or spaces.)

This is the same as testing for `iswupper()` or `iswlower()`.

## Return Value

Returns true if the character is alphabetic.

## Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswalph(L'a')? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswalph(L'B')? L"yes": L"no"); // yes
10     wprintf(L"%ls\n", iswalph(L'5')? L"yes": L"no"); // no
11     wprintf(L"%ls\n", iswalph(L'?' )? L"yes": L"no"); // no
12 }
```

## See Also

`iswalnum()`, `isalpha()`

---

## 32.3 `iswblank()`

Tests if this is a wide blank character

## Synopsis

```

#include <wctype.h>

int iswblank(wint_t wc);
```

## Description

Blank characters are whitespace that are also used as word separators on the same line. In the “C” locale, the only blank characters are space and tab.

Other locales might define other blank characters.

## Return Value

Returns true if this is a blank character.

## Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswblank(L' ') ? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswblank(L'\t') ? L"yes": L"no"); // yes
10     wprintf(L"%ls\n", iswblank(L'\n') ? L"yes": L"no"); // no
11     wprintf(L"%ls\n", iswblank(L'a') ? L"yes": L"no"); // no
12     wprintf(L"%ls\n", iswblank(L'?') ? L"yes": L"no"); // no
13 }

```

## See Also

iswspace(), isblank()

---

## 32.4 iswcntrl()

Tests if this is a wide control character.

### Synopsis

```

#include <wctype.h>

int iswcntrl(wint_t wc);

```

### Description

The spec is pretty barren, here. But I'm just going to assume that it works like the non-wide version. So let's look at that.

A *control character* is a locale-specific non-printing character.

For the “C” locale, this means control characters are in the range 0x00 to 0x1F (the character right before SPACE) and 0x7F (the DEL character).

Basically if it's not an ASCII (or Unicode less than 128) printable character, it's a control character in the “C” locale.

Probably.

### Return Value

Returns true if this is a control character.

## Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswcntrl(L'\t')? L"yes": L"no"); // yes (tab)
9      wprintf(L"%ls\n", iswcntrl(L'\n')? L"yes": L"no"); // yes (newline)
10     wprintf(L"%ls\n", iswcntrl(L'\r')? L"yes": L"no"); // yes (return)
11     wprintf(L"%ls\n", iswcntrl(L'\a')? L"yes": L"no"); // yes (bell)
12     wprintf(L"%ls\n", iswcntrl(L' ')? L"yes": L"no"); // no
13     wprintf(L"%ls\n", iswcntrl(L'a')? L"yes": L"no"); // no
14     wprintf(L"%ls\n", iswcntrl(L'')? L"yes": L"no"); // no
15 }

```

## See Also

isctrl()

## 32.5 iswdigit()

Test if this wide character is a digit

## Synopsis

```

#include <wctype.h>

int iswdigit(wint_t wc);

```

## Description

Tests if the wide character is a digit (0-9).

## Return Value

Returns true if the character is a digit.

## Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v

```

```

8      wprintf(L"%ls\n", iswdigit(L'0')? L"yes": L"no");    // yes
9      wprintf(L"%ls\n", iswdigit(L'5')? L"yes": L"no");    // yes
10     wprintf(L"%ls\n", iswdigit(L'a')? L"yes": L"no");    // no
11     wprintf(L"%ls\n", iswdigit(L'B')? L"yes": L"no");    // no
12     wprintf(L"%ls\n", iswdigit(L'')? L"yes": L"no");    // no
13 }

```

## See Also

`iswalnum()`, `isdigit()`

---

## 32.6 `iswgraph()`

Test to see if a wide character is a printable non-space

### Synopsis

```

#include <wctype.h>

int iswgraph(wint_t wc);

```

### Description

Returns true if this is a printable (non-control) character and also not a whitespace character.

Basically if `iswprint()` is true and `iswspace()` is false.

### Return Value

Returns true if this is a printable non-space character.

### Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //          testing this char
7      //          v
8      wprintf(L"%ls\n", iswgraph(L'0')? L"yes": L"no");    // yes
9      wprintf(L"%ls\n", iswgraph(L'a')? L"yes": L"no");    // yes
10     wprintf(L"%ls\n", iswgraph(L'B')? L"yes": L"no");    // yes
11     wprintf(L"%ls\n", iswgraph(L'')? L"yes": L"no");    // yes
12     wprintf(L"%ls\n", iswgraph(L' ')? L"yes": L"no");    // no
13     wprintf(L"%ls\n", iswgraph(L'\n')? L"yes": L"no");  // no
14 }

```



## See Also

`iswprint()`, `iswspace()`, `isgraph()`

---

## 32.7 `iswlower()`

Tests if a wide character is lowercase

### Synopsis

```
#include <wctype.h>

int iswlower(wint_t wc);
```

### Description

Tests if a character is lowercase, in the range `a-z`.

In other locales, there could be other lowercase characters. In all cases, to be lowercase, the following must be true:

```
!iswcntrl(c) && !iswdigit(c) && !iswpunct(c) && !iswspace(c)
```

### Return Value

Returns true if the wide character is lowercase.

### Example

```
1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswlower(L'c')? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswlower(L'0')? L"yes": L"no"); // no
10     wprintf(L"%ls\n", iswlower(L'B')? L"yes": L"no"); // no
11     wprintf(L"%ls\n", iswlower(L'?' )? L"yes": L"no"); // no
12     wprintf(L"%ls\n", iswlower(L' ' )? L"yes": L"no"); // no
13 }
```

## See Also

`islower()`, `iswupper()`, `iswalphabeta()`, `towupper()`, `towlower()`

---

## 32.8 iswprint()

Tests if a wide character is printable

### Synopsis

```
#include <wctype.h>

int iswprint(wint_t wc);
```

### Description

Tests if a wide character is printable, including space (' '). So like `isgraph()`, except space isn't left out in the cold.

### Return Value

Returns true if the wide character is printable, including space (' ').

### Example

```
1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswprint(L'c')? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswprint(L'0')? L"yes": L"no"); // yes
10     wprintf(L"%ls\n", iswprint(L' ')? L"yes": L"no"); // yes
11     wprintf(L"%ls\n", iswprint(L'\r')? L"yes": L"no"); // no
12 }
```

### See Also

`isprint()`, `iswgraph()`, `iswcntrl()`

---

## 32.9 iswpunct()

Test if a wide character is punctuation

### Synopsis

```
#include <wctype.h>

int iswpunct(wint_t wc);
```

## Description

Tests if a wide character is punctuation.

This means for any given locale:

```
!isspace(c) && !isalnum(c)
```

## Return Value

True if the wide character is punctuation.

## Example

Results may vary based on locale.

```
1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswpunct(L',')? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswpunct(L'!')? L"yes": L"no"); // yes
10     wprintf(L"%ls\n", iswpunct(L'c')? L"yes": L"no"); // no
11     wprintf(L"%ls\n", iswpunct(L'0')? L"yes": L"no"); // no
12     wprintf(L"%ls\n", iswpunct(L' ')? L"yes": L"no"); // no
13     wprintf(L"%ls\n", iswpunct(L'\n')? L"yes": L"no"); // no
14 }
```

## See Also

`ispunct()`, `isspace()`, `iswalnum()`

## 32.10 iswspace()

Test if a wide character is whitespace

## Synopsis

```
#include <wctype.h>

int iswspace(wint_t wc);
```

## Description

Tests if `c` is a whitespace character. These are probably:

- Space (' ')
- Formfeed ('\f')

- Newline ('\\n')
- Carriage Return ('\\r')
- Horizontal Tab ('\\t')
- Vertical Tab ('\\v')

Other locales might specify other whitespace characters. `iswalnum()`, `iswgraph()`, and `iswpunct()` are all false for all whitespace characters.

## Return Value

True if the character is whitespace.

## Example

Results may vary based on locale.

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                      testing this char
7      //                      v
8      wprintf(L"%ls\\n", iswspace(L' ')? L"yes": L"no"); // yes
9      wprintf(L"%ls\\n", iswspace(L'\\n')? L"yes": L"no"); // yes
10     wprintf(L"%ls\\n", iswspace(L'\\t')? L"yes": L"no"); // yes
11     wprintf(L"%ls\\n", iswspace(L',')? L"yes": L"no"); // no
12     wprintf(L"%ls\\n", iswspace(L'!')? L"yes": L"no"); // no
13     wprintf(L"%ls\\n", iswspace(L'c')? L"yes": L"no"); // no
14 }
```

## See Also

`isspace()`, `iswblank()`

---

## 32.11 `iswupper()`

Tests if a wide character is uppercase

### Synopsis

```

#include <wctype.h>

int iswupper(wint_t wc);
```

### Description

Tests if a character is uppercase in the current locale.

To be uppercase, the following must be true:

```
!isctrl(c) && !isdigit(c) && !ispunct(c) && !isspace(c)
```

## Return Value

Returns true if the wide character is uppercase.

## Example

```
1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswupper(L'B')? L"yes": L"no"); // yes
9      wprintf(L"%ls\n", iswupper(L'c')? L"yes": L"no"); // no
10     wprintf(L"%ls\n", iswupper(L'0')? L"yes": L"no"); // no
11     wprintf(L"%ls\n", iswupper(L'?' )? L"yes": L"no"); // no
12     wprintf(L"%ls\n", iswupper(L' ' )? L"yes": L"no"); // no
13 }
```

## See Also

isupper(), iswlower(), iswalphat(), towupper(), towlower()

## 32.12 iswxdigit()

Tests if a wide character is a hexadecimal digit

## Synopsis

```
#include <wctype.h>

int iswxdigit(wint_t wc);
```

## Description

Returns true if the wide character is a hexadecimal digit. Namely if it's 0-9, a-f, or A-F.

## Return Value

True if the character is a hexadecimal digit.

## Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //                testing this char
7      //                v
8      wprintf(L"%ls\n", iswxdigit(L'B')? L"yes": L"no");    // yes
9      wprintf(L"%ls\n", iswxdigit(L'c')? L"yes": L"no");    // yes
10     wprintf(L"%ls\n", iswxdigit(L'2')? L"yes": L"no");    // yes
11     wprintf(L"%ls\n", iswxdigit(L'G')? L"yes": L"no");    // no
12     wprintf(L"%ls\n", iswxdigit(L'?'? L"yes": L"no");    // no
13 }

```

## See Also

`isxdigit()`, `iswdigit()`

---

## 32.13 `iswctype()`

Determine wide character classification

### Synopsis

```

#include <wctype.h>

int iswctype(wint_t wc, wctype_t desc);

```

### Description

This is the Swiss Army knife of classification functions; it's all the other ones rolled into one.

You call it with something like this:

```
if (iswctype(c, wctype("digit"))) // or "alpha" or "space" or...
```

and it behaves just like you'd called:

```
if (iswdigit(c))
```

The difference is that you can specify the type of matching you want to do as a string at runtime, which might be convenient.

`iswctype()` relies on the return value from the `wctype()` call to get its work done.

Stolen from the spec, here are the `iswctype()` calls and their equivalents:

| iswctype() call               | Hard-coded equivalent |
|-------------------------------|-----------------------|
| iswctype(c, wctype("alnum"))  | iswalnum(c)           |
| iswctype(c, wctype("alpha"))  | iswalpha(c)           |
| iswctype(c, wctype("blank"))  | iswblank(c)           |
| iswctype(c, wctype("cntrl"))  | iswcntrl(c)           |
| iswctype(c, wctype("digit"))  | iswdigit(c)           |
| iswctype(c, wctype("graph"))  | iswgraph(c)           |
| iswctype(c, wctype("lower"))  | iswlower(c)           |
| iswctype(c, wctype("print"))  | iswprint(c)           |
| iswctype(c, wctype("punct"))  | iswpunct(c)           |
| iswctype(c, wctype("space"))  | iswspace(c)           |
| iswctype(c, wctype("upper"))  | iswupper(c)           |
| iswctype(c, wctype("xdigit")) | iswxdigit(c)          |

See the `wctype()` documentation for how that helper function works.

## Return Value

Returns true if the wide character `wc` matches the character class in `desc`.

## Example

Test for a given character classification at when the classification isn't known at compile time:

```

1  #include <stdio.h>  // for fflush(stdout)
2  #include <wchar.h>
3  #include <wctype.h>
4
5  int main(void)
6  {
7      wchar_t c;      // Holds a single wide character (to test)
8      char desc[128];  // Holds the character class
9
10     // Get the character and classification from the user
11     wprintf(L"Enter a character and character class: ");
12     fflush(stdout);
13     wscanf(L"%lc %s", &c, desc);
14
15     // Compute the type from the given class
16     wctype_t t = wctype(desc);
17
18     if (t == 0)
19         // If the type is 0, it's an unknown class
20         wprintf(L"Unknown character class: \"%s\"\n", desc);
21     else {
22         // Otherwise, let's test the character and see if its that
23         // classification
24         if (iswctype(c, t))
25             wprintf(L"Yes! '%lc' is %s!\n", c, desc);
26         else
27             wprintf(L"Nope! '%lc' is not %s.\n", c, desc);
28     }

```

29

}

Output:

```
Enter a character and character class: 5 digit
Yes! '5' is digit!

Enter a character and character class: b digit
Nope! 'b' is not digit.

Enter a character and character class: x alnum
Yes! 'x' is alnum!
```

**See Also**`wctype()`**32.14 wctype()**Helper function for `iswctype()`**Synopsis**

```
#include <wctype.h>

wctype_t wctype(const char *property);
```

**Description**

This function returns an opaque value for the given `property` that is meant to be passed as the second argument to `iswctype()`.

The returned value is of type `wctype_t`.

Valid properties in all locales are:

```
"alnum"  "alpha"  "blank"  "cntrl"
"digit"  "graph"  "lower"  "print"
"punct"  "space"  "upper"  "xdigit"
```

Other properties might be defined as determined by the `LC_CTYPE` category of the current locale.

See the `iswctype()` reference page for more usage details.

**Return Value**

Returns the `wctype_t` value associated with the given `property`.

If an invalid value is passed for `property`, returns 0.



## Example

Test for a given character classification at when the classification isn't known at compile time:

```

1  #include <stdio.h> // for fflush(stdout)
2  #include <wchar.h>
3  #include <wctype.h>
4
5  int main(void)
6  {
7      wchar_t c; // Holds a single wide character (to test)
8      char desc[128]; // Holds the character class
9
10     // Get the character and classification from the user
11     wprintf(L"Enter a character and character class: ");
12     fflush(stdout);
13     wscanf(L"%lc %s", &c, desc);
14
15     // Compute the type from the given class
16     wctype_t t = wctype(desc);
17
18     if (t == 0)
19         // If the type is 0, it's an unknown class
20         wprintf(L"Unknown character class: \"%s\"\n", desc);
21     else {
22         // Otherwise, let's test the character and see if its that
23         // classification
24         if (iswctype(c, t))
25             wprintf(L"Yes! '%lc' is %s!\n", c, desc);
26         else
27             wprintf(L"Nope! '%lc' is not %s.\n", c, desc);
28     }
29 }

```

Output:

```

Enter a character and character class: 5 digit
Yes! '5' is digit!

Enter a character and character class: b digit
Nope! 'b' is not digit.

Enter a character and character class: x alnum
Yes! 'x' is alnum!

```

## See Also

`iswctype()`

## 32.15 `towlower()`

Convert an uppercase wide character to lowercase

## Synopsis

```
#include <wctype.h>

wint_t tolower(wint_t wc);
```

## Description

If the character is upper (i.e. `iswupper(c)` is true), this function returns the corresponding lowercase letter. Different locales might have different upper and lowercase letters.

## Return Value

If the letter `wc` is uppercase, a lowercase version of that letter will be returned according to the current locale. If the letter is not uppercase, `wc` is returned unchanged.

## Example

```
1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //              changing this char
7      //              v
8      wprintf(L"%lc\n", tolower(L'B')); // b (made lowercase!)
9      wprintf(L"%lc\n", tolower(L'e')); // e (unchanged)
10     wprintf(L"%lc\n", tolower(L'!')); // ! (unchanged)
11 }
```

## See Also

`tolower()`, `toupper()`, `iswlower()`, `iswupper()`

---

## 32.16 `toupper()`

Convert a lowercase wide character to uppercase

## Synopsis

```
#include <wctype.h>

wint_t toupper(wint_t wc);
```

## Description

If the character is lower (i.e. `iswlower(c)` is true), this function returns the corresponding uppercase letter. Different locales might have different upper and lowercase letters.

## Return Value

If the letter `wc` is lowercase, an uppercase version of that letter will be returned according to the current locale. If the letter is not lowercase, `wc` is returned unchanged.

## Example

```

1  #include <wchar.h>
2  #include <wctype.h>
3
4  int main(void)
5  {
6      //          changing this char
7      //          v
8      wprintf(L"%lc\n", towupper(L'B')); // B (unchanged)
9      wprintf(L"%lc\n", towupper(L'e')); // E (made uppercase!)
10     wprintf(L"%lc\n", towupper(L'!')); // ! (unchanged)
11 }
```

## See Also

`toupper()`, `tolower()`, `iswlower()`, `iswupper()`

## 32.17 towctrans()

Convert wide characters to upper or lowercase

## Synopsis

```

#include <wctype.h>

wint_t towctrans(wint_t wc, wctrans_t desc);
```

## Description

This is the Swiss Army knife of character conversion functions; it's all the other ones rolled into one. And by "all the other ones" I mean `towupper()` and `tolower()`, since those are the only ones there are.

You call it with something like this:

```
if (towctrans(c, wctrans("toupper"))) // or "tolower"
```

and it behaves just like you'd called:

```
toupper(c);
```

The difference is that you can specify the type of conversion you want to do as a string at runtime, which might be convenient.

`towctrans()` relies on the return value from the `wctrans()` call to get its work done.

| <code>towctrans()</code> call                 | Hard-coded equivalent   |
|-----------------------------------------------|-------------------------|
| <code>towctrans(c, wctrans("toupper"))</code> | <code>toupper(c)</code> |
| <code>towctrans(c, wctrans("tolower"))</code> | <code>tolower(c)</code> |

See the `wctrans()` documentation for how that helper function works.

## Return Value

Returns the character `wc` as if run through `toupper()` or `tolower()`, depending on the value of `desc`.

If the character already matches the classification, it is returned as-is.

## Example

```

1  #include <stdio.h> // for fflush(stdout)
2  #include <wchar.h>
3  #include <wctype.h>
4
5  int main(void)
6  {
7      wchar_t c;           // Holds a single wide character (to test)
8      char desc[128];      // Holds the conversion type
9
10     // Get the character and conversion type from the user
11     wprintf(L"Enter a character and conversion type: ");
12     fflush(stdout);
13     wscanf(L"%lc %s", &c, desc);
14
15     // Compute the type from the given conversion type
16     wctrans_t t = wctrans(desc);
17
18     if (t == 0)
19         // If the type is 0, it's an unknown conversion type
20         wprintf(L"Unknown conversion: \"%s\"\n", desc);
21     else {
22         // Otherwise, let's do the conversion
23         wint_t result = towctrans(c, t);
24         wprintf(L"%lc' -> %s -> '%lc'\n", c, desc, result);
25     }
26 }
```

Output on my system:

```

Enter a character and conversion type: b toupper
'b' -> toupper -> 'B'

Enter a character and conversion type: B toupper
'B' -> toupper -> 'B'

Enter a character and conversion type: B tolower
'B' -> tolower -> 'b'

Enter a character and conversion type: ! toupper
'!' -> toupper -> '!'

```

**See Also**

wctrans(), towupper(), towlower()

---

**32.18 wctrans()**

Helper function for towctrans()

**Synopsis**

```

#include <wctype.h>

wctrans_t wctrans(const char *property);

```

**Description**

This is a helper function for generating the second argument to towctrans().

You can pass in one of two things for the property:

- toupper to make towctrans() behave like towupper()
- tolower to make towctrans() behave like towlower()

**Return Value**

On success, returns a value that can be used as the desc argument to towctrans().

Otherwise, if the property isn't recognized, returns 0.

**Example**

```

1  #include <stdio.h>  // for fflush(stdout)
2  #include <wchar.h>
3  #include <wctype.h>
4
5  int main(void)
6  {
7      wchar_t c;      // Holds a single wide character (to test)

```

```
8     char desc[128];    // Holds the conversion type
9
10    // Get the character and conversion type from the user
11    wprintf(L"Enter a character and conversion type: ");
12    fflush(stdout);
13    wscanf(L"%lc %s", &c, desc);
14
15    // Compute the type from the given conversion type
16    wctrans_t t = wctrans(desc);
17
18    if (t == 0)
19        // If the type is 0, it's an unknown conversion type
20        wprintf(L"Unknown conversion: \"%s\"\n", desc);
21    else {
22        // Otherwise, let's do the conversion
23        wint_t result = towctrans(c, t);
24        wprintf(L"%lc' -> %s -> '%lc'\n", c, desc, result);
25    }
26 }
```

Output on my system:

```
Enter a character and conversion type: b toupper
'b' -> toupper -> 'B'
```

```
Enter a character and conversion type: B toupper
'B' -> toupper -> 'B'
```

```
Enter a character and conversion type: B tolower
'B' -> tolower -> 'b'
```

```
Enter a character and conversion type: ! toupper
'!' -> toupper -> '!'
```

## See Also

`towctrans()`

# Index

- `_Alignas()` alignment specifier, 178
- `_Alignof()` operator, 180
- `_Atomic` type qualifier, 190
- `_Atomic()` type specifier, 190
- `_Complex_I` macro, 26
- `_Exit()` function, 301
- `_Imaginary_I` macro, 26
- `__STDC_ENDIAN_BIG__` macro, 210–211
- `__STDC_ENDIAN_LITTLE__` macro, 210–211
- `__STDC_ENDIAN_NATIVE__` macro, 210–211
- `__STDC_NO_COMPLEX__`, 26
- `__STDC_VERSION__` macro, 210
- `__alignas_is_defined` macro, 178
- `__alignof_is_defined` macro, 178
  
- `abort()` function, 298
- `abs()` function, 308
- `acos()` function, 113, 340
- `acosf()` function, 113
- `acosh()` function, 119, 340
- `acoshf()` function, 119
- `acoshl()` function, 119
- `acosl()` function, 113
- Addition operator, *see* + addition operator
- `alignas()` alignment specifier, 178
- `aligned_alloc()` function, 293
- `alignof()` operator, 180
- `and` macro, 95
- `and_eq` macro, 95
- `asctime()` function, 397
- `asin()` function, 114, 340
- `asinf()` function, 114
- `asinh()` function, 120, 340
- `asinhf()` function, 120
- `asinhf()` function, 120
- `asinl()` function, 114
- `assert()` macro, 21
- `assert.h` header file, 21
- `at_quick_exit()` function, 299
- `atan()` function, 115, 340
- `atan2()` function, 115, 340
- `atan2f()` function, 115
- `atan2l()` function, 115
- `atanf()` function, 115
- `atanh()` function, 121, 340
- `atanhf()` function, 121
- `atanhl()` function, 121
- `atanl()` function, 115
- `atexit()` function, 299
- `atof()` function, 283
- `atoi()` function, 284
- `atol()` function, 284
- `atoll()` function, 284
- `atomic_bool` type, 190
- `ATOMIC_BOOL_LOCK_FREE` macro, 191
- `atomic_char` type, 190
- `atomic_char16_t` type, 190
- `ATOMIC_CHAR16_T_LOCK_FREE` macro, 191
- `atomic_char32_t` type, 190
- `ATOMIC_CHAR32_T_LOCK_FREE` macro, 191
- `ATOMIC_CHAR_LOCK_FREE` macro, 191
- `atomic_compare_exchange_*()` function, 202
- `atomic_exchange()` function, 201
- `atomic_fetch_*()` function, 204
- `atomic_flag` type, 191
- `atomic_flag_clear()` function, 208
- `atomic_flag_test_and_set()` function, 206
- `atomic_init()` function, 193
- `atomic_int` type, 190
- `atomic_int_fast16_t` type, 190
- `atomic_int_fast32_t` type, 190
- `atomic_int_fast64_t` type, 190
- `atomic_int_fast8_t` type, 190
- `atomic_int_least16_t` type, 190
- `atomic_int_least32_t` type, 190
- `atomic_int_least64_t` type, 190
- `atomic_int_least8_t` type, 190
- `ATOMIC_INT_LOCK_FREE` macro, 191
- `atomic_intmax_t` type, 190
- `atomic_intptr_t` type, 190
- `atomic_is_lock_free()` function, 197
- `atomic_llong` type, 190
- `ATOMIC_LLONG_LOCK_FREE` macro, 191
- `atomic_load()` function, 200
- `atomic_long` type, 190
- `ATOMIC_LONG_LOCK_FREE` macro, 191
- `ATOMIC_POINTER_LOCK_FREE` macro, 191
- `atomic_ptrdiff_t` type, 190
- `atomic_schar` type, 190
- `atomic_short` type, 190

- ATOMIC\_SHORT\_LOCK\_FREE macro, 191
- atomic\_signal\_fence() function, 196
- atomic\_size\_t type, 190
- atomic\_store() function, 199
- atomic\_thread\_fence() function, 194
- atomic\_uchar type, 190
- atomic\_uint type, 190
- atomic\_uint\_fast16\_t type, 190
- atomic\_uint\_fast32\_t type, 190
- atomic\_uint\_fast64\_t type, 190
- atomic\_uint\_fast8\_t type, 190
- atomic\_uint\_least16\_t type, 190
- atomic\_uint\_least32\_t type, 190
- atomic\_uint\_least64\_t type, 190
- atomic\_uint\_least8\_t type, 190
- atomic\_uintmax\_t type, 190
- atomic\_uintptr\_t type, 190
- atomic\_ullong type, 190
- atomic\_ulong type, 190
- atomic\_ushort type, 190
- ATOMIC\_VAR\_INIT() macro, 192
- atomic\_wchar\_t type, 190
- ATOMIC\_WCHAR\_T\_LOCK\_FREE macro, 191
  
- Bell, *see* \a operator
- bitand macro, 95
- bitor macro, 95
- bool macro, 227
- Boolean AND, *see* && operator
- Boolean NOT, *see* ! operator
- Boolean OR, *see* || operator
- bsearch() function, 304
- btowc() function, 445
  
- c16rtomb() function, 410
- c32rtomb() function, 410
- cabs() function, 40
- cabsf() function, 40
- cabsl() function, 40
- cacos() function, 27
- cacosf() function, 27
- cacosh() function, 33
- cacoshf() function, 33
- cacoshl() function, 33
- cacosl() function, 27
- call\_once() function, 344
- calloc() function, 294
- carg() function, 43, 340
- cargf() function, 43
- cargl() function, 43
- Carriage return, *see* \r operator
- casin() function, 28
- casinf() function, 28
- casinh() function, 34
- casinhf() function, 34
- casinhl() function, 34
- casinl() function, 28
- catan() function, 29
- catanf() function, 29
- catanh() function, 35
- catanhf() function, 35
- catanhhl() function, 35
- catanl() function, 29
- cbrt() function, 137, 340
- cbrtf() function, 137
- cbrtl() function, 137
- ccos() function, 30
- ccosf() function, 30
- ccosh() function, 36
- ccoshf() function, 36
- ccoshl() function, 36
- ccosl() function, 30
- ceil() function, 145, 340
- ceilf() function, 145
- ceilhl() function, 145
- cexp() function, 39
- cexpf() function, 39
- cexpl() function, 39
- char16\_t type, 406
- char32\_t type, 406
- CHAR\_BIT macro, 97
- CHAR\_MAX macro, 97
- CHAR\_MIN macro, 97
- cimag() function, 44, 340
- cimagf() function, 44
- cimagl() function, 44
- clearerr() function, 278
- clock() function, 389
- clog() function, 40
- clogf() function, 40
- clogl() function, 40
- CMPLX() macro, 45
- CMPLXF() macro, 45
- CMPLXL() macro, 45
- cnd\_broadcast() function, 345
- cnd\_destroy() function, 348
- cnd\_init() function, 349
- cnd\_signal() function, 351
- cnd\_timedwait() function, 353
- cnd\_wait() function, 355
- compl macro, 95
- complex.h header file, 25
- conj() function, 47, 340
- conjf() function, 47
- conjl() function, 47
- copysign() function, 156, 340
- copysignf() function, 156
- copysignl() function, 156



- `cos()` function, 116, 340
- `cosf()` function, 116
- `cosh()` function, 121, 340
- `coshf()` function, 121
- `coshl()` function, 121
- `cosl()` function, 116
- `cpow()` function, 41
- `cpowf()` function, 41
- `cpowl()` function, 41
- `cproj()` function, 48, 340
- `cprojf()` function, 48
- `cprojl()` function, 48
- `creal()` function, 49, 340
- `crealf()` function, 49
- `creall()` function, 49
- `csin()` function, 31
- `csinf()` function, 31
- `csinh()` function, 37
- `csinhf()` function, 37
- `csinhl()` function, 37
- `csinl()` function, 31
- `csqrt()` function, 42
- `csqrtf()` function, 42
- `csqrtl()` function, 42
- `ctan()` function, 32
- `ctanf()` function, 32
- `ctanh()` function, 38
- `ctanhf()` function, 38
- `ctanhl()` function, 38
- `ctanl()` function, 32
- `ctime()` function, 398
- `ctype.h` header file, 51
- `CX_LIMITED_RANGE` macro, 26
- 
- `DBL_DECIMAL_DIG` macro, 84
- `DBL_DIG` macro, 84
- `DBL_EPSILON` macro, 82
- `DBL_HAS_SUBNORM` macro, 83
- `DBL_MANT_DIG` macro, 81
- `DBL_MAX` macro, 81
- `DBL_MAX_10_EXP` macro, 81
- `DBL_MAX_EXP` macro, 81
- `DBL_MIN` macro, 82
- `DBL_MIN_10_EXP` macro, 81
- `DBL_MIN_EXP` macro, 81
- `DBL_TRUE_MIN` macro, 82
- `DECIMAL_DIG` macro, 81
- `difftime()` function, 390
- `div()` function, 309
- `div_t` type, 283
- Division operator, *see* / division operator
- `double` complex type, 26
- `double` imaginary type, 26
- `double_t` type, 108
- 
- Endianness, 210–211
- `erf()` function, 141, 340
- `erfc()` function, 142, 340
- `erfcf()` function, 142
- `erfcl()` function, 142
- `erff()` function, 141
- `erfl()` function, 141
- `errno` variable, 65
- `errno.h` header file, 65
- `exit()` function, 301
- `EXIT_FAILURE` macro, 283
- `EXIT_SUCCESS` macro, 283
- `exp()` function, 124, 340
- `exp2()` function, 125, 340
- `exp2f()` function, 125
- `exp2l()` function, 125
- `expf()` function, 124
- `expl()` function, 124
- `expm1()` function, 125, 340
- `expm1f()` function, 125
- `expm1l()` function, 125
- 
- `fabs()` function, 138, 340
- `fabsf()` function, 138
- `fabsl()` function, 138
- `false` macro, 227
- `fclose()` function, 242
- `fdim()` function, 161, 340
- `fdimf()` function, 161
- `fdiml()` function, 161
- `FE_ALL_EXCEPT` macro, 69
- `FE_DIVBYZERO` macro, 69
- `FE_INEXACT` macro, 69
- `FE_INVALID` macro, 69
- `FE_OVERFLOW` macro, 69
- `FE_UNDERFLOW` macro, 69
- `feclearexcept()` function, 70
- `fegetenv()` function, 76
- `fegetexceptflag()` function, 71
- `fegetround()` function, 74
- `feholdexcept()` function, 77
- `fenv.h` header file, 69
- `FENV_ACCESS` pragma, 70
- `fenv_t` type, 69
- `feof()` function, 278
- `feraiseexcept()` function, 72
- `ferror()` function, 278
- `fesetenv()` function, 76
- `fesetexceptflag()` function, 71
- `fesetround()` function, 74
- `fetestexcept()` function, 73
- `feupdateenv()` function, 78
- `fexcept_t` type, 69
- `fflush()` function, 243

- fgetc() function, 265
- fgetpos() function, 274
- fgets() function, 266
- fgetwc() function, 420
- fgetws() function, 421
- FILE\* type, 237
- float complex type, 26
- float imaginary type, 26
- float.h header file, 81
- float\_t type, 108
- floor() function, 146, 340
- floorf() function, 146
- floorl() function, 146
- FLT\_DECIMAL\_DIG macro, 84
- FLT\_DIG macro, 84
- FLT\_EPSILON macro, 82
- FLT\_EVAL\_METHOD macro, 81, 83, 108
- FLT\_HAS\_SUBNORM macro, 83
- FLT\_MANT\_DIG macro, 81
- FLT\_MAX macro, 81
- FLT\_MAX\_10\_EXP macro, 81
- FLT\_MAX\_EXP macro, 81
- FLT\_MIN macro, 82
- FLT\_MIN\_10\_EXP macro, 81
- FLT\_MIN\_EXP macro, 81
- FLT\_RADIX macro, 81
- FLT\_ROUNDS macro, 83
- FLT\_TRUE\_MIN macro, 82
- fma() function, 162, 340
- fmaf() function, 162
- fmal() function, 162
- fmax() function, 162, 340
- fmaxf() function, 162
- fmaxl() function, 162
- fmin() function, 162, 340
- fminf() function, 162
- fminl() function, 162
- fmod() function, 153, 340
- fmodf() function, 153
- fmodl() function, 153
- fopen() function, 245
- FP\_CONTRACT pragma, 109
- fpclassify() function, 110
- fprintf() function, 249
- fputc() function, 268
- fputwc() function, 422
- fputws() function, 423
- fread() function, 272
- free() function, 296
- freopen() function, 246
- frexp() function, 126, 340
- frexpf() function, 126
- frexpl() function, 126
- fscanf() function, 256
- fseek() function, 275
- fsetpos() function, 274
- ftell() function, 277
- fwide() function, 424
- fwprintf() function, 415
- fwrite() function, 273
- fwscanf() function, 416
- getc() function, 265
- getchar() function, 265
- getenv() function, 302
- gets() function, 266
- getwc() function, 420
- getwchar() function, 420
- gmtime() function, 399
- Hexadecimal, *see* 0x hexadecimal
- hypot() function, 139, 340
- hypotf() function, 139
- hypotl() function, 139
- I macro, 26
- ilogb() function, 128, 340
- ilogbf() function, 128
- ilogbl() function, 128
- imaxabs() function, 89
- imaxdiv() function, 90
- INT\_FASTN\_MAX macros, 233
- INT\_FASTN\_MIN macros, 233
- int\_fastN\_t types, 232
- INT\_LEASTN\_MAX macros, 233
- INT\_LEASTN\_MIN macros, 233
- int\_leastN\_t types, 232
- INT\_MAX macro, 97
- INT\_MIN macro, 97
- INTMAX\_C() macro, 234
- INTMAX\_MAX macros, 233
- INTMAX\_MIN macros, 233
- intmax\_t type, 233
- INTN\_C() macros, 234
- INTN\_MAX macros, 233
- INTN\_MIN macros, 233
- intN\_t types, 232
- INTPTR\_MAX macros, 233
- INTPTR\_MIN macros, 233
- intptr\_t type, 233
- inttypes.h header file, 88
- isalnum() function, 52
- isalpha() function, 52
- isblank() function, 54
- iscntrl() function, 54
- isdigit() function, 55
- isfinite() function, 111
- isgraph() function, 56
- isgreater() function, 163

- isgreater() function, 163
- isinf() function, 111
- isless() function, 163
- islessequal() function, 163
- islessgreater() function, 164
- islower() function, 57
- isnan() function, 111
- isnormal() function, 111
- iso646.h header file, 95
- isprint() function, 58
- ispunct() function, 59
- isspace() function, 60
- isunordered() function, 165
- isupper() function, 61
- iswalnum() function, 457
- iswalpha() function, 458
- iswblank() function, 459
- iswcntrl() function, 460
- iswctype() function, 468
- iswdigit() function, 461
- iswgraph() function, 462
- iswlower() function, 463
- iswprint() function, 464
- iswpunct() function, 464
- iswspace() function, 465
- iswupper() function, 466
- iswxdigit() function, 467
- isxdigit() function, 62
  
- kill\_dependency() function, 194
  
- labs() function, 308
- LDBL\_DECIMAL\_DIG macro, 84
- LDBL\_DIG macro, 84
- LDBL\_EPSILON macro, 82
- LDBL\_HAS\_SUBNORM macro, 83
- LDBL\_MANT\_DIG macro, 81
- LDBL\_MAX macro, 81
- LDBL\_MAX\_10\_EXP macro, 81
- LDBL\_MAX\_EXP macro, 81
- LDBL\_MIN macro, 82
- LDBL\_MIN\_10\_EXP macro, 81
- LDBL\_MIN\_EXP macro, 81
- LDBL\_TRUE\_MIN macro, 82
- ldexp() function, 129, 340
- ldexpf() function, 129
- ldexpl() function, 129
- ldiv() function, 309
- ldiv\_t type, 283
- lgamma() function, 143, 340
- lgammaf() function, 143
- lgammal() function, 143
- limits.h header file, 97
- llabs() function, 308
- lldiv() function, 309
- lldiv\_t type, 283
- LLONG\_MAX macro, 97
- LLONG\_MIN macro, 97
- llrint() function, 149, 340
- llrintf() function, 149
- llrintl() function, 149
- llround() function, 151, 340
- llroundf() function, 151
- llroundl() function, 151
- locale.h. header file, 100
- localeconv() function, 102
- localtime() function, 400
- log() function, 129, 340
- log10() function, 130, 340
- log10f() function, 130
- log10l() function, 130
- log1p() function, 131, 340
- log1pf() function, 131
- log1pl() function, 131
- log2() function, 132, 340
- log2f() function, 132
- log2l() function, 132
- logb() function, 133, 340
- logbf() function, 133
- logbl() function, 133
- logf() function, 129
- logl() function, 129
- long double complex type, 26
- long double imaginary type, 26
- LONG\_MAX macro, 97
- LONG\_MIN macro, 97
- longjmp() function, 169
- lrint() function, 149, 340
- lrintf() function, 149
- lrintl() function, 149
- lround() function, 151, 340
- lroundf() function, 151
- lroundl() function, 151
  
- malloc() function, 294
- math.h header file, 106
- MATH\_ERREXCEPT macro, 109
- math\_errhandling variable, 109
- MATH\_ERRNO macro, 109
- max\_align\_t type, 230
- MB\_CUR\_MAX macro, 283
- MB\_LEN\_MAX macro, 97
- mblen() function, 310
- mbrlen() function, 447
- mbrtoc16() function, 407
- mbrtoc32() function, 407
- mbrtowc() function, 448
- mbsinit() function, 446



- remove() function, 238
- remquo() function, 155, 340
- remquof() function, 155
- remquol() function, 155
- rename() function, 239
- rewind() function, 275
- rint() function, 148, 340
- rintf() function, 148
- rintl() function, 148
- round() function, 150, 340
- roundf() function, 150
- roundl() function, 150
  
- scalbln() function, 135, 340
- scalblnf() function, 135
- scalblnl() function, 135
- scalbn() function, 135, 340
- scalbnf() function, 135
- scalbnl() function, 135
- scanf() function, 256
- SCHAR\_MAX macro, 97
- SCHAR\_MIN macro, 97
- SCNdFASTn macros, 89
- SCNdLEASTn macros, 89
- SCNdMAX macro, 89
- SCNdn macros, 89
- SCNdPTR macro, 89
- SCNiFASTn macros, 89
- SCNiLEASTn macros, 89
- SCNiMAX macro, 89
- SCNin macros, 89
- SCNiPTR macro, 89
- SCNoFASTn macros, 89
- SCNoLEASTn macros, 89
- SCNoMAX macro, 89
- SCNon macros, 89
- SCNoPTR macro, 89
- SCNuFASTn macros, 89
- SCNuLEASTn macros, 89
- SCNuMAX macro, 89
- SCNun macros, 89
- SCNuPTR macro, 89
- SCNxFASTn macros, 89
- SCNxLEASTn macros, 89
- SCNxMAX macro, 89
- SCNx macros, 89
- SCNxPTR macro, 89
- setbuf() function, 248
- setjmp() function, 167
- setjmp.h header file, 167
- setlocale() function, 100
- SHRT\_MAX macro, 97
- SHRT\_MIN macro, 97
- SIG\_ATOMIC\_MAX macro, 234
- SIG\_ATOMIC\_MIN macro, 234
- signal() function, 172
- signal.h header file, 172
- signbit() function, 112
- sin() function, 117, 340
- sinf() function, 117
- sinh() function, 122, 340
- sinhf() function, 122
- sinhl() function, 122
- sinl() function, 117
- SIZE\_MAX macro, 234
- size\_t type, 230, 283, 406
- snprintf() function, 249
- sprintf() function, 249
- sqrt() function, 140, 340
- sqrtrf() function, 140
- sqrtrl() function, 140
- srand() function, 291
- sscanf() function, 256
- static\_assert() macro, 23
- stdalign.h header file, 178
- stdarg.h header file, 182
- stdatomic.h header file, 189
- stdbit.h header file, 210
- stdbool.h header file, 227
- stdc\_bit\_ceil() function, 225
- stdc\_bit\_floor() function, 224
- stdc\_bit\_width() function, 223
- stdc\_count\_ones() function, 221
- stdc\_count\_zeros() function, 220
- stdc\_first\_leading\_one() function, 217
- stdc\_first\_leading\_zero() function, 216
- stdc\_first\_trailing\_one() function, 219
- stdc\_first\_trailing\_zero() function, 218
- stdc\_has\_single\_bit() function, 222
- stdc\_leading\_ones() function, 213
- stdc\_leading\_zeros() function, 212
- stdc\_trailing\_ones() function, 215
- stdc\_trailing\_zeros() function, 214
- stddef.h header file, 229
- stderr standard error, 238
- stdin standard input, 238
- stdint.h header file, 232
- stdio.h header file, 236
- stdlib.h header file, 282
- stdnoreturn.h header file, 319
- stdout standard output, 238
- strcat() function, 324
- strchr() function, 330
- strcmp() function, 326
- strcoll() function, 327
- strcpy() function, 322
- strcspn() function, 332
- strdup() function, 323

- strerror() function, 337
- strftime() function, 401
- String, *see* char \*
- string.h header file, 320
- strlen() function, 339
- strncat() function, 324
- strncmp() function, 326
- strncpy() function, 322
- strndup() function, 323
- strpbrk() function, 333
- strrchr() function, 330
- strspn() function, 332
- strstr() function, 334
- strtod() function, 286
- strtof() function, 286
- strtoimax() function, 91
- strtok() function, 335
- strtol() function, 288
- strtold() function, 286
- strtoll() function, 288
- strtoul() function, 288
- strtoull() function, 288
- strtoumax() function, 91
- struct tm type, 388
- strxfrm() function, 328
- Subtraction operator, *see* - subtraction operator
- swprintf() function, 415
- swscanf() function, 416
- system() function, 303
  
- Tab (is better), *see* \t operator
- tan() function, 118, 340
- tanf() function, 118
- tanh() function, 123, 340
- tanhf() function, 123
- tanh1() function, 123
- tanl() function, 118
- Ternary operator, *see* ?: ternary operator
- tgamma() function, 144, 340
- tgammaf() function, 144
- tgamma1() function, 144
- tgmath.h header file, 340
- thrd\_create() function, 367
- thrd\_current() function, 368
- thrd\_detach() function, 370
- thrd\_equal() function, 371
- thrd\_exit() function, 373
- thrd\_join() function, 374
- thrd\_yield() function, 377
- threads.h header file, 343
- time() function, 394
- time.h header file, 388
- time\_t type, 388
- timegm() function, 393
- timespec\_get() function, 395
- tmpfile() function, 240
- tmpnam() function, 241
- tolower() function, 62
- toupper() function, 63
- towctrans() function, 473
- towlower() function, 471
- towupper() function, 472
- true macro, 227
- trunc() function, 152, 340
- truncf() function, 152
- truncl() function, 152
- tss\_create() function, 378
- tss\_delete() function, 381
- tss\_get() function, 383
- tss\_set() function, 385
  
- UCHAR\_MAX macro, 97
- UINT\_FASTn\_MAX macros, 233
- UINT\_LEASTn\_MAX macros, 233
- UINT\_MAX macro, 97
- UINTMAX\_C() macro, 234
- UINTMAX\_MAX macros, 233
- uintmax\_t type, 233
- UINTn\_MAX macros, 233
- UINTPTR\_MAX macros, 233
- uintptr\_t type, 233
- ULLONG\_MAX macro, 97
- ULONG\_MAX macro, 97
- ungetc() function, 270
- ungetwc() function, 426
- USHRT\_MAX macro, 97
  
- va\_arg() macro, 182
- va\_copy() macro, 183
- va\_end() macro, 185
- va\_list type, 182
- va\_start() macro, 187
- vfprintf() function, 262
- vfprintf() function, 264
- vfwprintf() function, 417
- vfwscanf() function, 418
- vprintf() function, 262
- vscanf() function, 264
- vsprintf() function, 262
- vsprintf() function, 262
- vsscanf() function, 264
- vswprintf() function, 417
- vswscanf() function, 418
- vwprintf() function, 417
- vwscanf() function, 418
  
- wchar.h header file, 413
- WCHAR\_MAX macro, 234
- WCHAR\_MIN macro, 234

wchar\_t type, 231, 283  
wctomb() function, 450  
wcscat() function, 432  
wcschr() function, 437  
wcscmp() function, 433  
wcscoll() function, 434  
wcscpy() function, 430  
wcscspn() function, 438  
wcsftime() function, 443  
wcslen() function, 443  
wsncat() function, 432  
wsncmp() function, 433  
wsncpy() function, 430  
wspbrk() function, 440  
wsrcchr() function, 437  
wstrtombs() function, 453  
wcssp() function, 438  
wcsstr() function, 440  
wcstod() function, 427  
wcstof() function, 427  
wcstoimax() function, 92  
wcstok() function, 441  
wcstol() function, 428  
wcstold() function, 427  
wcstoll() function, 428  
wcstombs() function, 316  
wcstoul() function, 428  
wcstoull() function, 428  
wcstoumax() function, 92  
wcsxfrm() function, 435  
wctob() function, 445  
wctomb() function, 313  
wctrans() function, 475  
wctype() function, 470  
wctype.h header file, 457  
WINT\_MAX macro, 234  
WINT\_MIN macro, 234  
wmemcmp() function, 433  
wmemcpy() function, 431  
wmemmove() function, 431  
wprintf() function, 415  
wscanf() function, 416  
  
xor macro, 95  
xor\_eq macro, 95