

## Type qualifiers in C

### Abstract

Type qualifiers are part of the C language since 1989 (`const` and `volatile`) respectively 1999 (`restrict`). They are used to qualify types, modifying the properties of variables in certain ways. Since qualifiers are one of the lesser-understood features of the language, this article aims at experienced C programmers, and explains the reasoning behind these qualifiers.

### Introduction

Since 'Standard C' (C89), all variables are considered "unqualified" if none of the available qualifiers is used in the definition of the variable. Additionally, since all three qualifiers are completely independent from one another, for each unqualified simple type, we may have seven ( $2^3 - 1$ ) forms of qualified types.

Beware that qualifiers change the properties of their variables only for the scope and context in which they are used. A variable declared `'const'` is a constant only as far as the current scope and context is concerned. If we widen the scope (and regard, for instance, the callers of the function) or the context (for instance, other threads or tasks, interrupt service routines, or different autonomous systems), the variable may very well be not constant at all. For `'volatile'` and `'restrict'`, similar arguments exist. Consider calling `'memcpy'`, prototyped

```
void *memcpy(void *dest, const void *src, size_t len);
```

with a source pointer pointing to a regular (non-read-only) part of your memory.

### Const

The qualifier `'const'` is most often used in modern programs, and probably best understood. The addition of a `'const'` qualifier indicates that the (relevant part of the) program may not modify the variable. Such variables may even be placed in read-only storage (cf. section ""). It also allows certain kinds of optimizations, based on the premise that the variable's value cannot change. Please note, however, that "const-ness" may be cast away explicitly.

Since `'const'` variables cannot change their value during runtime (at least not within the scope and context considered), they must be initialized at their point of definition.

Example:

```
const int i = 5;
```

An alternate form is also acceptable, since the order of type specifiers and qualifiers does not matter:

```
int const i = 5;
```

Order becomes important when composite types with pointers are used:

```
int * const cp = &i; /* const pointer to int */
const int * ptc1; /* pointer to const int */
int const * ptc2; /* pointer to const int */
```

The pointer `cp` is itself const, i.e. the pointer cannot be modified; the integer variable it points to can.

The pointer `ptc1` can be modified, however, the variable it points to cannot.

Using `typedef` complicates the placement issue even more:

```
typedef int * ip_t;
const ip_t cp1 = &i; /* const pointer to int */
ip_t const cp2 = &i; /* const pointer to int!! */
```

Casting away 'const-ness' is possible, but considered dangerous. Modifying a const-qualified variable in that way is not only dangerous, but may even lead to run-time errors, if the values are placed in read-only storage:

```
const int * ptc1;
int *pti, i;
const int ci;

ptc1 = pti = &i;
```

```
ptci = &ci;
*ptci = 5;          /* Compiler error */
pti = &ci;          /* Compiler error */
pti = ptc;          /* Compiler error */
pti = (int *)&ci; /* OK, but dangerous */
*pti = 5;           /* OK, dangerous and potential runtime error */
```

PC-Lint and similar tools, as well as some compilers, will warn you about such dangerous situations, if you will let them.

## Placement

Placement of variables in actual memory is hardly standardized, because of the many requirements of specific compilers, processor architectures and requirements. But especially for `const`-qualified variables, it is a very interesting topic, and needs some discussion.

First of all, placement is compiler-specific. This means, that a compiler may specify how a programmer or system architect may direct the linker and loader as to where to place which variables or categories of variables. This may be done using extra configuration files, or using `#pragma`'s, or some other way. Refer to your compiler manual, especially when writing code for embedded systems.

If you revert to the compiler defaults, the compiler/linker/loader<sup>1</sup> may put `const`-qualified variables (not such combinations like 'pointer-to-const', since here the variable is a 'pointer' and non-const!) into read-only storage. If the compiler has no other indication, and can oversee the full scope of the variable (for instance, a `static const int const_int = 5;` at the global level in some C source file), it may even optimize in such a way, that the variable effectively disappears (replacing each occurrence with an immediate value), though not all compilers provide this kind of optimization.

If the compiler retains the variable as such (i.e. the variable is still present in the object-file), qualified with the property '`const`', the linker combines all corresponding references throughout all modules into one, complaining if the qualifications do not match, and the loader gets to decide, where the variable is placed in memory (dynamic linkers are even more complex, and disregarded here). If a memory area with read-only storage is available, `const`-qualified variables may end up there, at the discretion of the loader.

For details, consult your compiler manuals.

## Volatile

The qualifier '`volatile`' is normally avoided, understood only marginally, and quite often forgotten. It indicates to the compiler, that a variable may be modified outside the scope of the program. Such situations may occur for example in multitasking/-threading systems, when writing drivers with interrupt service routines, or in embedded systems, where the peripheral registers may also be modified by hardware alone.

The following fragment is a classical example of an endless loop:

```
int ready = 0;
while (!ready);
```

An aggressively optimizing compiler may very well create a simple endless loop (Microsoft Visual Studio 6.0, Release build with full optimization):

```
$L837:
; 5  :      int ready = 0;
; 6  :      while (!ready);

0000 eb fe      jmp     SHORT $L837
```

If we now add '`volatile`', indicating that the variable may be changed out of context, the compiler is not allowed to eliminate the variable entirely:

<sup>1</sup> In most compiler tool chains, the loader is an integrated part of the linker. For several embedded systems, however, the linker only produces relocatable code segments, to be effectively placed in memory by the loader.

```
volatile int ready = 0;
while (!ready);
```

becomes (using the option "favor small code"):

```
; 5      :      volatile int ready = 0;
00004 33 c0      xor     eax, eax
00006 89 45 fc      mov     DWORD PTR _ready$[ebp], eax
$L845:
; 6      :      while (!ready);
00009 39 45 fc      cmp     DWORD PTR _ready$[ebp], eax
0000c 74 fb      je     SHORT $L845
```

As you can see, even with aggressive, full optimization, the code still checks the variable every time through the loop.

Most compilers do not optimize this aggressively by default, but it is good to know that it is possible. The ordering issues as discussed in the section for the qualifier '`const`' also apply for '`volatile`'; if in doubt, refer back to page 1.

### ***When do you need to use 'volatile'?***

The basic principle is simple: Every time when a variable is used in more than one context, qualify it with '`volatile`':

- Whenever you use a common variable in more than one task or thread;
- Whenever you use a variable both in a task and one or more interrupt service routines;
- Whenever a variable corresponds to processor-internal registers configured as input (consider the processor or external hardware to be an extra context).

### ***Does it hurt to use 'volatile' unnecessarily?***

Well, yes and no. The functionality of your code will still be correct. However, the timing and memory footprint of your application will change: Your program will run slower, because of the extra read operations, and your program will be larger, since the compiler is not allowed to optimize as thoroughly, although that would have been possible.

### ***Why don't we declare all variables 'volatile'?***

Well, we partially do: On DEBUG-builds, all optimization is usually disabled. This is not quite the same, since the read operations needed extra are not necessarily inserted, but for most practical purposes, no optimizations involving the (missing) '`volatile`' qualification are executed. This can be considered at least partially equivalent.

We don't, however, deliver DEBUG-builds to the customer: They usually are too big and too slow (among a few other properties), just like if we declare all variables to be '`volatile`'.

But, whenever in doubt, it is better to use '`volatile`' unnecessarily, than to forget it when really necessary.

## **Restrict**

A discussion of the qualifier '`restrict`' is postponed until later, for multiple reasons:

- It is a new addition for C99 (the standard from 1999), hardly available in older compilers;
- The optimizations enabled by using the '`restrict`' qualification have been present in most commercial compilers, however, only on a global basis (compiler options); the extra gain to be expected is not quite as large, therefore, the necessity of using this qualifier can be considered minimal;
- My personal experience with this qualifier is still minimal.

One hint: If your compiler doesn't implement 'restrict' (yet), and doesn't even reserve the keyword, make sure to define a high-level macro, in order to prevent programmers using the name for a variable or similar:

```
#define restrict /* Reserved word */
```

## Combining qualifiers

In the introduction the existence of seven different qualified types for each (simple) unqualified one has been stated. Using three qualifiers, this means that qualifiers can be combined.

In C89, each qualifier may only be used once, in C99, multiple occurrences of each single qualifier are explicitly allowed and silently ignored.

But now, take 'volatile' and 'const': What does it mean to have a variable qualified with both:

```
const volatile unsigned int * const ptcvi = 0xFFFFFFFFCAUL;
```

OK, the initialization value is hexadecimal, unsigned long, and taken from an imaginary embedded processor. The pointer is const, so I cannot change the pointer. And the value it points to is "const volatile unsigned int": I cannot change the value (within my current scope and context), and the value may be changed out of context.

So, imagine a free running counter, counting upwards from 0 to 65535 (hexadecimal 0xFFFF or 16 bit), and rolling over again to 0. If this counter is automatically started by the hardware, or started by the (assembly-coded) startup-routine (outside the scope of the C program), is never stopped, and only used for relative time measurements, we have exactly this situation: The counter is read-only, so I want the compiler to supervise all programmers, that they do not try to write the counter register.

At the same time, the value is constantly changed, so if I want to use the value, the compiler better make sure to re-read the value in every single case.

You can also imagine a battery backed-up clock chip, running autonomously, with values for the current date and time memory-mapped into the processors virtual memory space.

OK, such situations will not occur every day, and for many programmers they will never occur. But it is not unimaginable. And now go out and ask the most experienced C programmer you know, whether it is possible, allowed and/or useful. You'll be amazed about the answers you'll get (or maybe not).

Literature:

C A Reference Manual – (5<sup>th</sup> edition); Samuel P. Harbison III, Guy L. Steele Jr; Prentice hall, 2002; ISBN 0-13-089592X.

Comments, remarks and criticism are welcome.

Johan Bezem  
j.bezem@computer.org  
<http://www.bezem.de>