

Representing and Manipulating Hardware in Standard C and C++
Embedded Systems Conference San Francisco 2002
Class 465

Dan Saks

Introduction

This paper discusses the basics of hardware manipulation using C and C++. It focuses on the common idioms for controlling devices that you can write in Standard C or Standard C++.

Both C and C++ provide the following features that aid embedded programming:

- bitwise operators and bitfields for packing data and manipulating individual bits in hardware registers
- the `const` qualifier for identifying ROMable data
- the `volatile` qualifier for identifying objects, such as memory-mapped ports, that may change due to events outside a program's control
- the type `sig_atomic_t` for declaring objects that a program can set and clear without interruption

C++ does not add any features specifically for embedded programming, but classes can be extremely helpful in hiding messy hardware details behind cleaner interfaces.

Touching Hardware

Nearly all computer architectures use at least one of the following methods for accessing the hardware of input/output (i/o) devices:

- Memory-mapped i/o maps device registers into the conventional data space. To a C or C++ programmer, memory-mapped i/o registers look much like ordinary data. Storing into a memory-mapped device register sends commands or data to a device. Reading from a memory-mapped i/o register retrieves status information or data from a device. This is the approach used in the Motorola 68K family of processors.
- Port i/o maps control and data registers into a separate (often small) data space. Port i/o is similar to memory-mapped i/o except that programs must use special instructions, such as the `in` and `out` instructions on the Intel x86 processors, to move data to or from the device registers. Neither the C nor C++ standard says anything spe-

cifically about port i/o, so it's almost impossible to write code for port i/o devices without using non-standard language or library features.

- Some architectures provide special i/o instructions dedicated to particular devices. As with port i/o, special i/o instructions are outside the C and C++ standards. Code for such devices usually requires using non-standard language or library features.

Since Standard C and C++ are best at memory-mapped i/o, we will explore that first.

Representing Hardware Device Registers

A given hardware device often communicates through a sequence of device registers. Some registers communicate control and status information; others communicate data. A given device may use separate registers for input and output. Registers may occupy bytes, or words, or whatever the architecture demands.

The simplest representation for a device register is as an object of the appropriately sized and signed integer type. For example, you might declare a one-byte register as a **char** or a two-byte register as a **short int**. Then you can receive data from a device by reading a value from its input data register, and send data to a memory-mapped device by storing a value into its output data register.

A typical control/status register is really not an integer-valued object – it's a collection of bits. One bit may indicate that the device is ready to perform an operation, while another bit might indicate whether interrupts have been enabled for that device. A given device might not use all the available bits in its control/status register.

Using the bitmask approach to manipulate control/status registers, you:

- use bitwise operators to set, clear and test bits in the registers
- use symbolic constants to represent masks for isolating bits

For example, here's the bitmask representation for a simple 16-bit control/status register:

```
typedef short int control;
#define ENABLE 0x0040      /* bit 6: enable interrupts */
#define READY  0x0080      /* bit 7: device is ready */
```

Neither Standard C nor Standard C++ lets you declare a variable so that it resides at a specified memory address. Rather, you locate a device register by declaring a pointer whose value is the specified address. For example, a program can define:

```
control *const pcr = (control *)0xFF70;
```

which uses a cast expression to initialize `pcr` to point to the address of a memory-mapped control/status register. Although Standard C and C++ both allow cast expressions that convert integers to pointers (such as the one above), the exact behavior of such casts vary across platforms.

C++ offers an alternative form for the cast syntax called a `reinterpret_cast`, as in:

```
control *const pcr = reinterpret_cast<control *>(0xFF70);
```

Of the two styles for casts, `reinterpret_cast` is now the preferred one in C++.

Once `pcr` points to a memory-mapped control/status register, the program can communicate with the device by testing or setting the value of the register via `pcr`. For example,

```
*pcr &= ~ENABLE;          /* clear enable bit */
```

clears the control register's enable bit, so the device will not trigger interrupts. A loop such as:

```
while ((*pcr & READY) == 0) /* wait until ready */  
    ;
```

repeatedly polls (tests) the control register's ready bit until that bit is non-zero. This loop makes the program wait until the device completes whatever it's doing.

In C, you can make a control register look more like an object by using a macro as follows:

```
#define cr (*pcr)
```

Then you can manipulate the device register by writing expressions such as:

```
cr &= ~ENABLE;
```

In C++, you can use a reference (rather than a macro), as in:

```
control &cr = *pcr;
```

Again, then you can write expressions such as:

```
cr &= ~ENABLE;
```

Many devices use a control/status and a data register in tandem. The declarations for such a pair might look like:

```
typedef short int control;
typedef short int data;

typedef struct port port;
struct port
{
    control cr;
    data dr;
};

port *const p = (port *)0xFF70;
```

The `typedef`:

```
typedef struct port port;
```

defines `port` as a type name that's equivalent to `struct port`. This allows you to refer to the structure type `port` without also writing the keyword `struct` in front of it. In C++, a `struct`, union or `enum` name is automatically a type name, so you don't need the `typedef` in C++. But it doesn't hurt. For the remainder of this paper, I will assume for convenience that a `struct`, union or `enum` name is a type name, even in C.

A bi-directional device (supporting both input and output) may use a pair of registers for input and another pair for output. The declaration for such a pair might look like:

```
struct ioport
{
    port in, out;
};
```

Using the declaration above,

```
ioport *const piop = (ioport *)0xFF70;
```

declares `piop` to point to an `ioport`, and:

```
piop->out.cr &= ~ENABLE;
```

disables output interrupts for that port. An assignment such as:

```
piop->out.dr = c;
```

sends the value of character **c** to the output device. It immediately clears the ready bit in the output control/status register to indicate that the device is busy. A loop such as:

```
while ((piop->out.cr & READY) == 0)
    ;
```

or:

```
while (!(piop->out.cr & READY))
    ;
```

waits until the output device is ready to receive another character. A function such as:

```
void put(char const *s, ioport *p)
{
    for (; *s != '\0'; ++s)
    {
        while ((p->out.cr & READY) == 0)
            ;
        p->out.dr = *s;
    }
}
```

sends the characters from null-terminated string **s** to the output device controlled by ***p**.

The Volatile Qualifier

The previous function might not work exactly as expected. Here's an example that illustrates why.

Using the previous declaration for **ioport**, the following sequence of code writes **'\r'** (carriage return) and **'\n'** (newline or line feed) to the output device of an **ioport** (based on Plauger [1988]):

```

ioport *const p = (ioport *)0xFFFFC4;
...
while (!(p->out.cr & READY))
    ;
p->out.dr = '\r';
while (!(p->out.cr & READY))
    ;
p->out.dr = '\n';

```

This code is supposed to wait until the output device's **READY** bit indicates that the device is ready to perform an operation. Then it writes a '**\r**' to the device's data register. The code waits again until the device is ready to perform another operation. Then it writes a '**\n**' to the device's data register.

The compiler has no way of knowing that **p->out.cr** is actually a device register whose value changes spontaneously in response to external events such as a change in state for a peripheral device. The compiler's optimizer might therefore conclude that the value of the **READY** bit in **p->out.cr** never changes. It's always 1 or it's always 0.

If the **READY** bit never changes, then there's no need to test the loop condition more than once. The program can test the bit once and then either loop forever or skip the loop entirely. Thus, the compiler's optimizer can transform:

```

while (!(p->out.cr & READY))
    ;

```

into:

```

if (!(p->out.cr & READY))
    for (;;)
        ;

```

In fact, the compiler can transform both loops in the same way. After this optimization, the driver code becomes:

```

if (!(p->out.cr & READY))
    for (;;)
        ;
p->out.dr = '\r';
if (!(p->out.cr & READY))
    for (;;)
        ;
p->out.dr = '\n';

```

Again, the compiler can't see any code that changes the **READY** bit. It concludes that the **READY** bit is either always set or always clear.

- If the **READY** bit is always set, the program always skips the first loop. When the execution reaches the second loop, the bit is still set and so the program skips the second loop as well.
- If the **READY** bit is always clear, the program enters the first loop and never leaves. It never gets a chance to execute the second loop.

In either case, the program never executes the second loop.

Therefore, the optimizer can eliminate the second if-statement entirely. The optimized driver code becomes:

```
if (!(p->out.cr & READY))
    for (;;)
        ;
p->out.dr = '\r';
p->out.dr = '\n';
```

Now, as far as the compiler can see, the first assignment writes '**r**' into a location, only to have the next assignment overwrite the same location with '**n**'. Therefore, only the last assignment is worth keeping. The final “optimized” code looks like:

```
if (!(p->out.cr & READY))
    for (;;)
        ;
p->out.dr = '\n';
```

It does the wrong thing, but much more efficiently than the original code!

The way to prevent this overly aggressive optimization is the use the volatile qualifier, as in:

```
ioport volatile *const p = (ioport *)0xFFFC4;
```

This declares **p** as a “**const** pointer to a **volatile ioport**”. The volatile qualifier indicates that an object, such as a memory-mapped port, may be change even though the program didn't change it. The volatile qualifier prevents the compiler from “optimizing away” references to an object that seems unchanging, but doesn't inhibit other optimizations.

In a declaration such as:

```
ioport volatile *const p = ... ;
```

volatile is not part of the type of **ioport**. This is appropriate only if some **ioports** are non-volatile. If volatility is inherent in **ioports**, **volatile** should be part of the **ioport** type.

Recall that each **ioport** is a pair of **ports**. It may be that only some registers within a **port** are volatile:

```
struct port
{
    control volatile cr;
    data dr;
};
```

If every member of a **port** is volatile, you must declare each member so, as in:

```
struct port
{
    control volatile cr;
    data volatile dr;
};
```

If you want to declare the entire **port** as a volatile type, this won't work:

```
volatile struct port
{
    control cr;
    data dr;
}; // error: missing declarator
```

To declare the entire **port** as a volatile type, you can use a **typedef**, as in:

```
typedef struct /* no tag */
{
    control cr;
    data dr;
} volatile port;
```

Just as pointer conversions can add but not remove a **const** qualifier, they can add but not remove a **volatile** qualifier:

```

T *p;
T volatile *pv;
pv = p;          // OK
p = pv;          // error
p = (T *)pv;    // OK

```

Believe it or not, the “new style” cast for removing **volatile** from a pointer type is **const_cast**. There is no such thing as a **volatile_cast**:

```

T *p;
T volatile *pv;
p = const_cast<T *>(pv);    // yes
p = volatile_cast<T *>(pv); // no

```

You must preserve the volatility of arguments in parameter passing as well:

```

ioport volatile *const piop = ...;
void put(char const *s, ioport *p);

put("...", piop);    // error

```

The call is an error because the compiler can’t convert **piop** from “pointer to **volatile ioport**” into “pointer to **ioport**”. You can eliminate the error by declaring **put**’s second parameter as **ioport volatile *p**, as in:

```

void put(char const *s, ioport volatile *p);

```

Objects can be both **const** and **volatile**, meaning “I promise I won’t change it, but I’ll assume that something else might.” This is appropriate for read-only device registers.

Representing Hardware Registers as Bitfields

Bitfields offer a slightly more elegant approach to modeling hardware registers. Using this approach, you:

- declare a register as a **struct** and its bits as bitfield members, and
- set, clear and test bits in registers by referring to each bitfield by name.

For example,

```

struct control
{
    unsigned int /* unused */ : 6;
    unsigned int enable : 1;
    unsigned int ready : 1;
    unsigned int /* unused */ : 8;
};
typedef short int data;

struct port
{
    control volatile cr;
    data volatile dr;
};

```

Using these declarations,

```
piop->out.cr.enable = 0;
```

disables output interrupts for that port and:

```
while (piop->out.cr.ready == 0)
    ;
```

waits until the output device is ready.

If you prefer, you can even think of each single-bit bitfield as a boolean, and write:

```
piop->in.cr.enable = false;
```

to disable output interrupts and:

```
while (!piop->in.cr.ready)
    ;
```

to waits until the output device is ready.

In C++, you can use the boolean constants **false** and **true**. In C, you must define these symbols appropriately.

A word of caution: Neither C nor C++ specifies whether a compiler must allocate bits starting with the most- or the least-significant bit. Therefore, although this technique is generally applicable across platforms, the code itself is not likely to be portable.

MIXING HARDWARE REGISTER METAPHORS

You can sometimes get the most compact and efficient code by employing different representations for a single register. For example, if `n` is an `int`, then:

```
if ((n & ~INT_MAX) == 0)
```

tests the high-order bit of `n` using a bitmask. (The symbol `INT_MAX` is a macro defined in `<limits.h>` as the largest positive value that an `int` can hold.) Typically, the size and speed of a bit masking instruction is independent of the position of the bit(s) within the mask.

You can also test the high-order bit of `n` using a signed comparison with zero:

```
if (n < 0)
```

On some architectures, this signed comparison generates a smaller or faster instruction than a bitwise-and. Thus, some hardware designers have a habit of placing the most often used bit of a control/status register as the sign bit so that software can use a signed comparison to test that bit. This encourages you to treat a control/status register as a **signed int** sometimes, and as a set of bits at other times.

A union can provide alternative views of a single device register:

```
struct byte_pair
{
    signed char lo, hi;
};

union control
{
    short int word;
    byte_pair byte;
};
```

This union assumes that `sizeof(short int)` is 2, which is not true on all platforms.

Anyway, using the union, here's a (possibly) efficient way to test the high-order bit of the high byte of a register:

```
control volatile *const pcr = (control *)0xFF70;
...
if (pcr->byte.hi < 0)
    ...
```

Again, you must be careful to adapt these coding techniques to match the byte ordering of your architecture. For example, the union above overlays `lo` with the low-order byte of `word` only if the machine is “little-endian”. On a “big-endian” machine, you should reverse the order of the bytes in the `byte_pair`, as in:

```
struct byte_pair
{
    signed char hi, lo;
};
```

Once you get the byte ordering correct, then for control `c`, the expressions `c.byte.hi` and `c.byte.lo` refer to the high- and low-order bytes of the register, respectively. `c.word` refers to the entire register.

Remember, no matter how you mix representations, you must rely on implementation-dependent storage representations.

Hiding Messy Details in a C++ Class

A C++ class can hide most of the messy details described above, such as the choice of bitmasks vs. bitfields (including any bitmask constants) or dependency on implementation-defined byte ordering. For example, you might define a class `control_status` for control/status registers as:

```
class control_status
{
public:
    void enable();
    void disable();
    bool enabled() const;
    bool ready() const;
```

```

private:
    enum { ENABLE = 0x0040, READY = 0x0080, ... };
    union
    {
        short int word;
        byte_pair byte;
    } u;
};

```

This class represents a control/status register as a union wherein you can view the register as either a (**signed**) **short int** or a pair of **signed char**. The class assumes that **sizeof(short int)** is 2, and that the low-order byte of a **short int** has the same address as the entire **short int**. Only the union member **u** occupies storage. The **enum** definition defines a type and constants, but no data that occupies space in **control_status** objects.

This example defines the constants **ENABLE** and **READY** as enumeration constants rather than as the macros used in earlier examples. The advantage that enumerations have over macros is that:

- Macros ignore the normal scope rules of C and C++. That is, they ignore { and }, so macros might substitute in places you don't want them to. Enumerations observe the scope rules.
- Some compilers don't preserve macro names among the symbols they pass on to their debuggers.

Class **control_status** defines its member functions as **inline** to avoid run-time performance penalties:

```

inline
void control_status::enable()
{
    u.word |= ENABLE;
}

inline
void control_status::disable()
{
    u.word &= ~ENABLE;
}

```

```

inline
bool control_status::enabled() const
{
    return u.word & ENABLE;
}

inline
bool control_status::ready() const
{
    return u.byte.lo < 0;
}

```

As is often the case with unions, we are not interested in the union itself, but rather only in the members in the union. That is, we never refer to `control_status` member `u` by itself. We always refer to `u.word` or `u.byte`. C++ lets you eliminate the name `u` by using an *anonymous union*, as in:

```

class control_status
{
public:
    ...
private:
    ...
    union
    {
        short int word;
        byte_pair byte;
    } /* no name here */;
};

```

This, in turn, simplifies the member functions. For example, using the anonymous union, the definition for `control_status::ready` simplifies from:

```

inline
bool control_status::ready() const
{
    return u.byte.lo < 0;
}

```

to:

```

inline
bool control_status::ready() const
{
    return byte.lo < 0;
}

```

Class `control_status` restricts access to hardware so that the rest of the program can manipulate that hardware only in proscribed ways. That is, all manipulation of a `control_status` must be through public member function calls, as in:

```

struct port
{
    control_status csr;
    data dr;
};

port *const p = ...;
...
p->csr.disable();
while (!p->csr.ready())
    ;

```

If `control_status` objects can be declared volatile, as in:

```

struct port
{
    control_status volatile csr;
    data volatile dr;
};

```

then `control_status` member functions such as `disable` must be declared `volatile`, as in:

```

class control_status
{
public:
    void enable() volatile;
    void disable() volatile;
    bool enabled() const volatile;
    bool ready() const volatile;
    ...
};

```

The keyword **volatile** also must appear in the corresponding member function definition:

```
inline
void control_status::enable() volatile
{
    u.word |= ENABLE;
}
```

A *volatile member function* treats the object it applies to (***this**) as a volatile object. You cannot apply a member function to a volatile object unless the member function is a volatile member function. A member function can be both const and volatile.

You can use **structs** and macros in C to organize hardware manipulation as if you were using a class. However, **structs** and macros cannot prevent erroneous operations on a device as well as a class can. Classes leave the program with fewer ways to misuse the hardware.

Write-Only Registers

Some machines have control/status registers that are write-only. Write-only registers pose a problem for functions such as:

```
inline
void control_status::enable() volatile
{
    word |= ENABLE;
}
```

The statement in the function body is shorthand for:

```
word = word | ENABLE;
```

The reference to **word** on the right-hand side of the = is a read access. If **word** is write-only, this function won't work properly.

For a write-only register, the program must maintain a read-write copy of the register. The place to keep that read-write copy is in a class.

The previous definition for **control_status** provides a class for objects that overlay bus addresses corresponding to hardware registers. The following **control_status**

class uses a different approach that's more appropriate when dealing with write-only registers. Here, a **control_status** object is not the hardware itself, but an object in memory that points to and monitors a register at a specified bus address. For example,

```
union csr_type
{
    short int word;
    byte_pair byte;
};

class control_status
{
public:
    typedef size_t address_type;
    control_status(address_type);
    void enable();
    ...
private:
    ...
    csr_type image;
    csr_type volatile *actual_csr;
};
```

Here, member **image** is a copy of the register pointed to by member **actual_csr**. The **control_status** constructor:

```
control_status::control_status(address_type at):
    actual_csr(reinterpret_cast<csr_type *>(at))
    {
        image.word = 0;
    }
```

sets **actual_csr** to point to the actual hardware register and clears the **image**. The corresponding **enable** function might look something like:

```
void control_status::enable()
{
    image |= ENABLE;
    *actual_csr = image;
}
```

Notice that the **control_status** member functions no longer need to be **volatile**. Now, **control_status** objects have a member that's a pointer to a **volatile csr_type**, but the **control_status** object is not **volatile**.

Port I/O

When using port i/o, you define the data structures for device registers pretty much as above; however, you cannot move data to and from the device registers by simple assignment. Since the C and C++ standards do not encompass port i/o, you must rely on language extensions for moving data to and from port i/o registers. Here are some examples of what your compiler might provide (from Plauger [1993]):

- special functions:

```
char c = in(0x40);
```

- inline assembly code:

```
char c;  
_asm in, 0x40  
_asm mov y, al
```

- unique address spaces and addressing modes:

```
@port char InBuf = 0x40;  
char c = InBuf;
```

Again, your best bet is to hide messy details in as little code as possible, preferably in member functions of a C++ class.

INTERRUPT HANDLING

Interrupt handling is largely outside the C and C++ standards. However, both languages provide type `sig_atomic_t` defined in `<signal.h>` as the integral type of an object that a program can access as an atomic entity, even in the presence of asynchronous interrupts.

Interrupt calling conventions rarely match normal C/C++ function calling conventions. When writing interrupt handlers in C or C++, you must either dip into assembly language, or use language extensions offered by some implementations as in the following example.

The keyword `volatile` and type `sig_atomic_t` are also useful in declaring data for communicating between asynchronous execution threads. For example, here's a small

real-mode MS-DOS™ program with an interrupt handler that responds to hitting the control-break key:

```
/*
 * ctrlbrk.c - a crude Ctrl-Break interrupt handler
 * (based on Duncan, [1986])
 */
#include <signal.h>

sig_atomic_t volatile cb_seen = 0;

void interrupt far cb_handler(void)
{
    cb_seen = 1; /* keep it simple ! */
}

/*
 * using the Ctrl-Break interrupt handler
 */
#define CB_VECTOR 0x1B

int main()
{
    int c;
    void (interrupt far *old_handler)(void);
    old_handler = _dos_getvect(CB_VECTOR);
    _dos_setvect(CB_VECTOR, cb_handler);
    for (;;)
    {
        ...
        if (cb_seen)
        {
            puts("\nCtrl-Break detected!\n");
            cb_seen = 0;
        }
        ...
    }
    _dos_setvect(CB_VECTOR, old_handler);
    return 0;
}
```

The keywords **far** and **interrupt** are common language extensions available on C and C++ compilers for PC platforms. **_dos_getvect** and **_dos_setvect** are DOS system calls. The call to **_dos_getvect** returns the current value of an interrupt vec-

tor. `_dos_setvect` stores a new value into an interrupt vector. The program stores the prior value of the interrupt vector in `old_handler` so it can put that value back at the end.

Note that `cb_seen` must be `volatile`; otherwise, the compiler might generate code on the assumption that it's always zero.

In Summary

- You can control memory-mapped i/o using only standardized language features (but they will still have platform-specific behavior).
- Port i/o, special i/o instructions and interrupt handling are outside the standard. You can do it only by using language extensions.
- Whatever you do, isolate language and hardware dependencies inside abstract types, such as a C++ class.

Acknowledgments

Some the material in this paper comes from articles that appeared in the *Programming Pointers* column of *Embedded Systems Programming* magazine, written and copyright © 1998 by Dan Saks. Also, thanks to P.J. Plauger for sharing his notes on this material.

References

- Duncan [1986]. Ray Duncan. *Advanced MS-DOS*. Microsoft Press.
- Plauger [1988]. P.J. Plauger. "Touching Memory: Standard C Makes The Act More Precise", *The C Users Journal*, Vol. 6, No. 4, May, 1988.
- Plauger [1993]. P.J. Plauger. "*Manipulating Hardware in C and C++*", Proceedings of the Embedded Systems Conference East, April, 1993. Miller Freeman.