

A Firmware Development Standard

Version 1.3, Updated Sep 2005

Jack G. Ganssle
jack@ganssle.com

The Ganssle Group
PO Box 38346
Baltimore, MD 21231
fax (647) 439-1454

Table of Contents

Scope

This document defines the standard way all programmers will create embedded firmware. Every programmer is expected to be intimately familiar with the Standard, and to understand and accept these requirements. All consultants and contractors will also adhere to this Standard.

The reason for the Standard is to insure all Company-developed firmware meets minimum levels of readability and maintainability. Source code has two equally-important functions: it must *work*, and it must clearly *communicate how it works* to a future programmer or the future version of yourself. Just as a standard English grammar and spelling makes prose readable, standardized coding conventions ease readability of one's firmware.

Part of every code review is to insure the reviewed modules and functions meet the requirements of the Standard. Code that does not meet this Standard will be rejected.

We recognize that no Standard can cover every eventuality. There may be times where it makes sense to take exception to one or more of the requirements incorporated in this document. Every exception must meet the following requirements:

- *Clear Reasons* - Before making an exception to the Standard, the programmer(s) will clearly spell out and understand the reasons involved, and will communicate these reasons to the project manager. The reasons must involve clear benefit to the project and/or Company; stylistic motivations, or programmer preferences and idiosyncrasies are not adequate reasons for making an exception.
- *Approval* - The project manager will approve all exceptions made
- *Documentation* - The effected module or function will have the exception clearly documented in the comments, so during code reviews and later maintenance the current and future technical staff understand the reasons for the exception, and the nature of the exception.

Projects

Directory Structure

To simplify use of a version control system, and to deal with unexpected programmer departures and sicknesses, every programmer involved with each project will maintain identical directory structures for the source code associated with the project.

The general “root” directory for a project takes the form:

`/projects/project-name/rom_name`

where

- “`/projects`” is the root of all firmware developed by the Company. By keeping all projects under one general directory version control and backup is simplified; it also reduces the size of the computer’s root directory.
- “`/project-name`” is the formal name of the project under development.
- “`/rom_name`” is the name of the ROM the code pertains to. One project may involve several microprocessors, each of which has its own set of ROMs and code. Or, a single project may have multiple binary images, each of which goes into its own set of ROMs.

Required directories:

`/projects/project-name/tools` - compilers, linkers, assemblers used by this project. All tools will be checked into the VCS so in 5 to 10 years, when a change is required, the (now obsolete and unobtainable) tools will still be around. It’s impossible to recompile and retest the project code every time a new version of the compiler or

assembler comes out; the only alternative is to preserve old versions, forever, in the VCS.

- `/projects/project-name/rom_name/headers` - all header files, such as .h or assemble include files, go here.
- `/projects/project-name/rom_name/source` - source code. This may be further broken down into header, C, and assembly directories. The MAKE files are also stored here.
- `/projects/project-name/rom_name/object` - object code, including compiler/assembler objects and the linked and located binaries.
- `/projects/project-name/rom_name/test` - This directory is the one, and only one, that is not checked into the VCS and whose subdirectory layout is entirely up to the individual programmer. It contains work-in-progress, which is generally restricted to a single module. When the module is released to the VCS or the rest of the development team, the developer must clean out the directory and eliminate any file that is duplicated in the VCS.

Version File

Each project will have a special module that provides firmware version name, version date, and part number (typically the part number on the ROM chips). This module will list, in order (with the newest changes at the top of the file), all changes made from version to version of the released code.

Remember that the production or repair departments may have to support these products for years or decades. Documentation gets

lost and ROM labels may come adrift. To make it possible to correlate problems to ROM versions, even after the version label is long gone, the Version file should generate only one bit of “code” - a string that indicates, in ASCII, the current ROM version. Some day in the future a technician - or yourself! - may then be able to identify the ROM by dumping the ROM’s contents. An example definition is:

```
# undef VERSION
# define VERSION "Version 1.30"
```

The Version file also contains the Abbreviations Table. See the section under “Variables” for more detail.

Example:

```
/******
 * Version Module - Project SAMPLE
 *
 * Copyright 1997 Company
 * All Rights Reserved
 *
 * The information contained herein is confidential
 * property of Company. The use, copying, transfer or
 * disclosure of such information is prohibited except
 * by express written agreement with Company.
 *
 * 12/18/97 - Version 1.3 - ROM ID 78-130
 *          Modified module AD_TO_D to fix scaling
 *          algorithm; instead of y=mx, it now
 *          computes y=mx+b.
 * 10/29/97 - Version 1.2 - ROM ID 78-120
 *          Changed modules DISPLAY_LED and READ_DIP
 *          to incorporate marketing’s request for a
 *          diagnostics mode.
 * 09/03/97 - Version 1.1 - ROM ID 78-110
 *          Changed module ISR to properly handle
 *          non-reentrant math problem.
 * 07/12/97 - Version 1.0 - ROM ID 78-100
 *          Initial release
 *****/
# undef VERSION
# define VERSION "Version 1.30"
```

Make and Project Files

Every executable will be generated via a MAKE file, or the equivalent supported by the tool chain selected. The MAKE file includes all of the information needed to automatically build the entire ROM image. This includes compiling and assembling source files, linking, locating (if needed), and whatever else must be done to produce a final ROM image.

An alternative version of the MAKE file may be provided to generate debug versions of the code. Debug versions may include special diagnostic code, or might have a somewhat different format of the binary image for use with debugging tools.

In integrated development environments (like Visual C++) specify a PROJECT file that is saved with the source code to configure all MAKE-like dependencies.

In no case is any tool *ever* to be invoked by typing in a command, as invariably command line arguments “accumulate” over the course of a project... only to be quickly forgotten once version 1.0 ships.

Startup Code

Most ROM code, especially when a C compiler is used, requires an initial startup module that sets up the compiler’s runtime package and initializes certain hardware on the processor itself, including chip selects, wait states, etc.

Startup code generally comes from the compiler or locator vendor, and is then modified by the project team to meet specific needs of the project. It is invariably compiler- and locator-specific. Therefore, the first modification made to the startup code is an

initial comment that describes the version numbers of all tools (compiler, assembler, linker, and locator) used.

Vendor-supplied startup code is notoriously poorly documented. To avoid creating difficult-to-track problems, *never* delete a line of code from the startup module. Simply comment-out unneeded lines, being careful to put a note in that you were responsible for disabling the specific lines. This will ease re-enabling the code in the future (for example, if you disable the floating point package initialization, one day it may need to be brought back in).

Many of the peripherals may be initialized in the startup module. Be careful when using automatic code generation tools provided by the processor vendor (tools that automate chip select setup, for example). Since many processor boot with RAM chip selects disabled, always include the chip select and wait state code in-line (not as a subroutine). Be careful to initialize these selects at the very top of the module, to allow future subroutine calls to operate, and since some debugging tools will not operate reliably until these are set up.

Stack and Heap Issues

Always initialize the stack on an *even* address. Resist the temptation to set it to a odd value like 0xffff, since on a word machine an odd stack will cripple system performance. Since few programmers have a reasonable way to determine maximum stack requirements, always assume your estimates will be incorrect. For each stack in the system, make sure the initialization code fills the entire amount of memory allocated to the stack with the value 0x55. Later, when debugging, you can view the stack and detect stack overflows by seeing no blocks of 0x55 in that region. Be sure, though, that the code that fills the

stack with 0x55 automatically detects the stack's size, so a late night stack size change will not destroy this useful tool.

Embedded systems are often intolerant of heap problems. Dynamically allocating and freeing memory may, over time, fragment the heap to the point that the program crashes due to an inability to allocate more RAM. (Desktop programs are much less susceptible to this as they typically run for much shorter periods of time).

So, be wary of the use of the malloc() function. When using a new tool chain examine the malloc function, if possible, to see if it implements garbage collection to release fragmented blocks (note that this may bring in another problem, as during garbage collection the system may not be responsive to interrupts). *Never* blindly assume that allocating and freeing memory is cost- or problem-free.

If you chose to use malloc(), *always* check the return value and safely crash (with diagnostic information) if it fails.

When using C, if possible (depending on resource issues and processor limitations), always include Walter Bright's MEM package (<http://c.snippets.org/browser.php>) with the code, at least for the debugging.

MEM provides:

- ISO/ANSI verification of allocation/reallocation functions
- Logging of all allocations and frees
- Verifications of Frees
- Detection of pointer over- and under-runs.
- Memory leak detection
- Pointer checking
- Out of memory handling

Modules

General

A *Module* is a single file of source code that contains one or more functions or routines, as well as the variables needed to support the functions.

Each module contains a number of *related* functions. For instance, an A/D converter module may include all A/D drivers in a single file. Grouping functions in this manner makes it easier to find relevant sections of code, and allows more effective encapsulation.

Encapsulation - hiding the details of a function's operation, and keeping the variables used by the function local - is absolutely essential. Though C and assembly language don't explicitly support encapsulation, with careful coding you can get all of the benefits of this powerful idea as do people using OOP languages.

In C and assembly language you can define all variables and RAM inside the modules that use those values. Encapsulate the data by defining each variable for the scope of the functions that use these variables only. Keep them private within the function, or within the module, that uses them.

Modules tend to grow large enough that they are unmanageable. Keep module sizes under 1000 lines to insure tools (source debuggers, compilers, etc.) are not stressed to the point they become slow or unreliable, and to enhance clarity.

Templates

To encourage a uniform module look and feel, create module templates named "module_template.c" and "module_template.asm", stored in the source directory, that becomes part of the code base maintained by the VCS. Use one of these files as the base for all new modules. The module template includes a standardized form for the header (the comment block preceding all code), a standard spot for file includes and module-wide declarations, function prototypes and macros. The templates also include the standard format for functions.

Here's the template for C code:

```

/*****
* Module name:
*
* Copyright 1997 Company as an unpublished work.
* All Rights Reserved.
*
* The information contained herein is confidential
* property of Company. The user, copying, transfer or
* disclosure of such information is prohibited except
* by express written agreement with Company.
*
* First written on xxxxx by xxxx.
*
* Module Description:
* (fill in a detailed description of the module's
* function here).
*
*****/
/* Include section
* Add all #includes here
*
*****/
/* Defines section
* Add all #defines here

```

```

*
*****/
/* Function Prototype Section
* Add prototypes for all functions called by this
* module, with the exception of runtime routines.
*
*****/

```

The template includes a section defining the general layout of functions, as follows:

```

/*****
* Function name      : TYPE foo(TYPE arg1, TYPE arg2...)
*   returns         : return value description
*   arg1            : description
*   arg2            : description
* Created by        : author's name
* Date created      : date
* Description       : detailed description
* Notes            : restrictions, odd modes
*****/

```

The template for assembly modules is:

```

;*****
; Module name:
;
; Copyright 1997 Company as an unpublished work.
; All Rights Reserved.
;
; The information contained herein is confidential
; property of Company. The user, copying, transfer or
; disclosure of such information is prohibited except
; by express written agreement with Company.
;
; First written on xxxxx by xxxxx.
;
; Module Description:
; (fill in a detailed description of the module
; here).
;
;*****
; Include section
; Add all "includes" here

```

```

;*****
The template includes a section defining the general
layout of functions, as follows:
;*****
; Routine name      : foobar
;   returns         : return value(s) description
;   arg1            : description of arguments
;   arg2            : description
; Created by        : author's name
; Date created      : date
; Description       : detailed description
; Notes            : restrictions, odd modes
;*****

```

Module Names

Never include the project's name or acronym as part of each module name. It's much better to use separate directories for each project.

Big projects may require many dozens of modules; scrolling through a directory listing looking for the one containing function main() can be frustrating and confusing. Therefore store function main() in a module named main.c or main.asm.

Filenames will be all lower case to enhance portability between Windows and Linux/Unix systems.

File extensions will be:

C Source Code	filename.c
C Header File	filename.h
Assembler files	filename.asm
Assembler include files	filename.inc
Object Code	filename.obj
Libraries	filename.lib
Shell Scripts	filename.bat
Directory Contents	README

Build rules for make project.mak

Variables

Names

Regardless of language, use long names to clearly specify the variable's meaning. If your tools do not support long names, get new tools.

Separate words within the variables by underscores. Do not use capital letters as separators. Consider how much harder `IcantReadThis` is on the eyes versus `I_can_read_this`.

Variable and function names are defined with the first words being descriptive of broad ideas, and later words narrowing down to specifics. For instance:

`Universe_Galaxy_System_Planet`. Consider the following names: `Timer_0_Data`, `Timer_0_Overflow`, and `Timer_0_Capture`. This convention quickly narrows variables to particular segments of the program. Never assume that a verb must be first, as often seen when naming functions.

`Open_Serial_Port` and `Close_Serial_Port` do a much poorer job of grouping than the better alternative of `Serial_Port_Open` and `Serial_Port_Close`.

Acronyms and abbreviations are not allowed as parts of variable names unless:

- 1) defined in a special Abbreviations Table which is stored in the Version file.
- 2) an accepted industry convention like LCD, LED and DSP

Clarity is our goal! An example Abbreviation Table is:

```
/* Abbreviation Table
* Dsply    == Display (the verb)
* Disp    == Display (our LCD display)
* Tot     == Total
* Calc    == Calculation
* Val     == Value
* Pos     == Position
*/
```

The ANSI C specification restricts the use of names that begin with an underscore and either an uppercase letter or another underscore (`_[A-Z][0-9A-Za-z_]`). Much compiler runtime code also starts with leading underscores. To avoid confusion, never name a variable or function with a leading underscore.

These names are also reserved by ANSI for its future expansion:

<code>E[0-9A-Z][0-9A-Za-z]*</code>	errno values
<code>is[a-z][0-9A-Za-z]*</code>	Character classification
<code>to[a-z][0-9A-Za-z]*</code>	Character manipulation
<code>LC_[0-9A-Za-z]*</code>	Locale
<code>SIG[_A-Z][0-9A-Za-z]*</code>	Signals
<code>str[a-z][0-9A-Za-z]*</code>	String manipulation
<code>mem[a-z][0-9A-Za-z]*</code>	Memory manipulation
<code>wcs[a-z][0-9A-Za-z]*</code>	Wide character manipulation

Global Variables

All too often C and especially assembly programs have one huge module with all of the variable definitions. Though it may seem nice to organize variables in a common spot, the peril is these are all then global in scope. Global variables are responsible for much

undebuggable code, reentrancy problems, global warming and male pattern baldness. Avoid them!

Real time code may occasionally require a few - and only a few - global variables to insure reasonable response to external events. *Every global variable must be approved by the project manager.*

When globals are used, put all of them into a single module. They are so problematic that it's best to clearly identify the sin via the name `globals.c` or `globals.asm`.

Portability

Avoid the use of "int" and "long", as these declarations vary depending on the machine. Create typedefs as follows:

	signed	unsigned
8 bit:	<code>int8_t</code>	<code>uint8_t</code>
16 bit:	<code>int16_t</code>	<code>uint16_t</code>
32 bit:	<code>int32_t</code>	<code>uint32_t</code>
64 bit:	<code>int64_t</code>	<code>uint64_t</code>

Don't assume that the address of an int object is also the address of its least-significant byte. This is not true on big-endian machines.

Functions

Regardless of language, *keep functions small!* The ideal size is less than a page; in no case should a function ever exceed two pages. Break large functions into several smaller ones.

The only exception to this rule is the very rare case where real time constraints (or sometimes stack limitations) mandate long sequences of in-line code. The project manager must approve all such code... but first look hard for a more structured alternative!

Explicitly declare every parameter passed to each function. Clearly document the meaning of the parameter in the comments.

Define a prototype for every called functions, with the exception of those in the compiler's runtime library. Prototypes let the compiler catch the all-too-common errors of incorrect argument types and improper numbers of arguments. They are cheap insurance.

In general, function names should follow the variable naming protocol.

Interrupt Service Routines

ISRs, though usually a small percentage of the code, are often the hardest bits of firmware to design and debug. Crummy ISRs will destroy the project schedule!

Decent interrupt routines, though, require properly designed hardware. Sometimes it's tempting to save a few gates by letting the external device just toggle the interrupt line for a few microseconds. This is unacceptable. Every interrupt must be latched until acknowledged, either by the processor's interrupt-acknowledge cycle (be sure the hardware acks the proper interrupt source), or via a handshake between the code and the hardware.

Use the non-maskable interrupt only for catastrophic events, like the apocalypse or imminent power failure. Many tools cannot properly debug NMI code. Worse, NMI is guaranteed to break non-reentrant code.

If at all possible, design a few spare I/O bits in the system. These are tremendously useful for measuring ISR performance.

Keep ISRs short! Long (too many lines of code) and slow are the twins of ISR disaster. Remember that *long* and *slow* may be disjoint; a five line ISR with a loop can be as much of a problem as a loop-free 500 line routine. When an ISR grows too large or too slow, spawn another task and exit. Large ISRs are a sure sign of a need to include a RTOS.

Budget time for each ISR. Before writing the routine, understand just how much time is available to service the interrupt. Base all

of your coding on this, and then *measure* the resulting ISR performance to see if you met the system's need. Since every interrupt competes for CPU resources, that slow ISR that works is just as buggy as one with totally corrupt code.

Never allocate or free memory in an ISR unless you have a clear understanding of the behavior of the memory allocation routines. Garbage collection or the ill-behaved behavior of many runtime packages may make the ISR time non-deterministic.

On processors with interrupt vector tables, fill every entry of the table. Point those entries not used by the system to an error handler, so you've got a prayer of finding problems due to incorrectly-programmed vectors in peripherals.

Though non-reentrant code is always dangerous in a real time system, it's often unavoidable in ISRs. Hardware interfaces, for example, are often non-reentrant. Put all such code as close to the beginning of the ISR as possible, so you can then re-enable interrupts. Remember that as long as interrupts are off the system is not responding to external requests.

Comments

Code *implements* an algorithm; the comments *communicate* the code's operation to yourself and others. Adequate comments allow you to understand the system's operation without having to read the code itself.

Write comments in *clear English*. Use the sentence structure Miss Grandel tried to pound into your head in grade school. Avoid writing the Great American Novel; be concise yet explicit... but be complete.

Avoid long paragraphs. Use simple sentences: noun, verb, object. Use active voice: "Motor_Start actuates the induction relay after a 4 second pause". Be complete. Good comments capture everything important about the problem at hand.

Use proper case. Using all caps or all lower case simply makes the comments harder to read and makes the author look like an illiterate moron.

Enter comments in C at block resolution and when necessary to clarify a line. Don't feel compelled to comment each line. It is much more natural to comment groups of lines which work together to perform a macro function. However, never assume that long variable names create "self documenting code". Self documenting code is an oxymoron, so add comments where needed to make the firmware's operation crystal clear. It should be possible to get a sense of the system's operation by reading only the comments.

Explain the meaning and function of every variable declaration. Every single one. Explain the return value, if any. Long variable names are merely an *aid* to understanding; accompany the descriptive name with a deep, meaningful, prose description.

Explain the parameters during the function definition, as follows:

```
type function_name(type parameter1 /* comment */
                  type parameter2 /* comment */)
```

Comment assembly language blocks, and any line that is not crystal clear. The worst comments are those that say "move AX to BX" on a MOV instruction! Reasonable commenting practices will yield about one comment on every other line of assembly code.

Acronyms and abbreviations are not allowed unless defined in the Abbreviation Table stored in the Version file (see more about this under variable names). While "DSP" might mean "Display" to you, it means "Digital Signal Processor" to me. Clarity is our goal!

Though it's useful to highlight sections of comments with string's of asterisks, never have characters on the right side of a block of comments. It's too much trouble to maintain proper spacing as the comments later change. In other words, this is not allowed:

```
/******
 * This comment incorrectly uses right-hand *
 * asterisks                               *
```

```
***** /
```

The correct form is:

```
/*  
 * This comment does not use right-hand  
 * asterisks  
*/
```

Coding Conventions

General

No line may ever be more than 80 characters.

Don't use absolute path names when including header files. Use the form `#include <module/name>` to get public header files from a standard place.

Never, ever use “magic numbers”. Instead, first understand where the number comes from, then define it in a constant, and then document your understanding of the number in the constant's declaration.

Spacing and Indentation

Put a space after every keyword, unless a semicolon is the next character, but never between function names and the argument list.

Put a space after each comma in argument lists and after the semicolons separating expressions in a `for` statement.

Put a space before and after every binary operator (like `+`, `-`, etc.). Never put a space between a unary operator and its operand (e.g., unary minus).

Put a space before and after pointer variants (star, ampersand) in declarations. Precede pointer variants with a space, but have no following space, in expressions.

Indent C code in increments of two spaces. That is, every indent level is two, four, six, etc. spaces.

Always place the `#` in a preprocessor directive in column 1.

C Formatting

Never nest IF statements more than three deep; deep nesting quickly becomes incomprehensible. It's better to call a function, or even better to replace complex ifs with a SWITCH statement.

Place braces so the opening brace is the last thing on the line, and place the closing brace first, like:

```
if (result > a_to_d) {
    do a bunch of stuff
}
```

Note that the closing brace is on a line of its own, except when it is followed by a continuation of the same statement, such as:

```
do {
    body of the loop
} while (condition);
```

When an `if-else` statement is nested in another `if` statement, always put braces around the `if-else` to make the scope of the first `if` clear.

When splitting a line of code, indent the second line like this:

```
function (float arg1, int arg2, long arg3,
         int arg4)
```

or,

```
if (long_variable_name && constant_of_some_sort == 2
    && another_condition)
```

Use too many parenthesis. Never let the compiler resolve precedence; explicitly declare precedence via parenthesis.

Never make assignments inside `if` statements. E.g., don't write:

```
if ((foo = (char *) malloc (sizeof *foo)) == 0)
    fatal ("virtual memory exhausted");
```

instead, write:

```
foo = (char *) malloc (sizeof *foo);
if (foo == 0)
    fatal ("virtual memory exhausted")
```

If you use `#ifdef` to select among a set of configuration options, add a final `#else` clause containing a `#error` directive so that the compiler will generate an error message if none of the options has been defined:

```
#ifdef sun
#define USE_MOTIF
#elif hpux
#define USE_OPENLOOK
#else
#error unknown machine type
#endif
```

Assembly Formatting

Tab stops in assembly language are as follows:

- Tab 1: column 8
- Tab 2: column 16
- Tab 3: column 32

Note that these are all in increments of 8, for editors that don't support explicit tab settings. A large gap - 16 columns - is between the operands and the comments.

Place labels on lines by themselves, like this:

```
label:
    mov    r1, r2        ; r1=pointer to I/O
```

Precede and follow comment blocks with semicolon lines:

```
;
; Comment block that shows how comments stand
; out from the code when preceded and followed by
; "blank" lines.
;
```

Never run a comment between lines of code. For example, do not write like this:

```
mov    r1, r2; Now we set r1 to the value
add    r3, [data] ; we read back in read_ad
```

Instead, use either a comment block, or a line without an instruction, like this:

```
mov    r1, r2; Now we set r1 to the value
; we read back in read_ad
add    r3, [data]
```

Be wary of macros. Though useful, macros can quickly obfuscate meaning. Do pick very meaningful names for macros.