

Continue



C unit testing

This article is the series on Unit testing in C and carries the discussion on Unit Testing and its implementation. The aim of this series is to provide easy and practical examples that anyone can understand. This is the Introduction of Unit testing - Unit testing in C tutorial Part 1. You can also read ceedling installation, Unity, Cmock, and stringizing, token pasting in C, Unit Testing in C - Introduction Anyone who has been involved in the software development life cycle (SDLC) for a while, will have encountered some form of testing. Software testing is an acceptance mechanism for discovering how well a software works according to the specified requirements. Although the aim of testing is to find bugs, it cannot guarantee the absence of other faults, no matter how creative the test cases have been designed. Unit testing enables a more thorough level of acceptance testing. What is Unit Testing? A unit is simply a small piece of code for any single function. So when we test those units, it is called a unit test. The unit test is a short script or piece of code designed to verify the behavior of a particular unit independently to produce a pass or fail result. Unit Testing is performed during the application development phase. Unit testing is performed usually by developers. In V-model, SDLC, and STLC the unit testing is the first phase of testing before integration testing. It is a white box testing technique and QA engineers can also perform Unit Testing if required. However, it can sometimes be quite difficult to write a good unit test for a particular piece of code. Having difficulty testing their own or someone else's code, developers often think that their struggles are caused by a lack of some fundamental testing knowledge or secret unit testing techniques. In this unit testing tutorial series, I intend to demonstrate that unit tests are quite easy; the real problems that complicate unit testing. Why is Unit Testing necessary and its use? Sometimes, developers skip out unit testing due to lack of time. Skipping out unit testing leads to more defect fixing costs during integration, Beta, and system testing. Proper unit testing at the time of application development, saves time and money. Here are the key reasons to perform Unit Testing. The defects can be fixed early at the development stage, and it saves time and costs. It helps developers to understand the code base and make changes quickly that you think of. Good unit tests generally serve as the project documentation. Unit tests can be reused and migrated to the new project quickly when required. You should tweak the code a bit to run again. Better Design - When developers write unit tests, their emphasis is on thinking about how their code will be used throughout the system, which generally results in better design. Unit Test on Embedded Software/Firmware Unit tests can help you write a better embedded software. To allow unit testing for a software project, the R&D team must write a testable, modular code - code that can be divided into self-contained units that can be tested. On top of making the code testable, embedded software developers must make sure their code is portable. The unit test won't test the functionality like how it is running in hardware. Do we need hardware (target) to run the unit tests? Not really, if you use ceedling framework for unit tests, then you don't need your target hardware. You can use your PC to run the unit test. In our full unit test series, we are going to discuss ceedling only. So your PC machine is enough to learn. The misconception of Unit Test Our firmware is very simple, we don't need unit testing. We can't unit test microcontroller code, it's too close to the hardware. Unit testing will add an unnecessary cost burden to our project. What is the unit testing framework? A unit test framework is just some code/application that makes it easier to run, test the code which we have written, and recorded the results of unit tests. Frameworks used for Unit Test So, now we know that unit testing is a valuable activity for embedded firmware. There are many frameworks that are available for each programming language. We have listed some tools for our C and Embedded platform. You can use anyone for unit testing. Ceedling Embunit MinUnit Criterion LCUt etc. In our next tutorial, we will see the code coverage and its types. You can also read the below tutorials. Embedded Software | Firmware | Linux Devic Driver | RTOS Hi, I am a tech blogger and an Embedded Engineer. I am always eager to learn and explore tech-related concepts. And also, I wanted to share my knowledge with everyone in a more straightforward way with easy practical examples. I strongly believe that learning by doing is more powerful than just learning by reading. I love to do experiments. If you want to help or support me on my journey, considering sharing my articles, or Buy me a Coffee. Thank you for reading my blog! Happy learning! Hits (since 1 July 2022) - 32,637 On-demand pricing techscalex with your first-class support for all major languages and frameworks using lightweight SDKs and agentsNative integration with Slack, GitHub, Jira, Bitbucket, and more of your favorite tools When it comes to QA, you're not alone. We work side-by-side with your team to ensure your testing process is smooth, efficient, and impactful. Share - copy and redistribute the material in any medium or format for any purpose, even commercially. Adapt - remix, transform, build upon the material for any purpose, even commercially. The licensor cannot revoke these freedoms as long as you follow the license terms. Attribution - You must give appropriate credit - provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use. ShareAlike - If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original. No additional restrictions - You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits. You do not have to comply with the license for elements of the public domain or where your use is permitted by an applicable exception or limitation . No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material. When you find yourself (or your company) with more code than anyone could ever test by hand, what can you do? Well, unit testing has always been the perfect solution, as you can run tests that check more data than a person could in a day in a matter of milliseconds. So today I'll take a look into a few popular C# unit testing frameworks and try them out first hand so you can choose which one best suits your project. Unit tests can be run as often as you want, on as many different kinds of data as you want and with next to no human involvement beyond once the tests are written. Not only that, but using code to test code will often result in you noticing flaws with your program that would have been very difficult to spot from a programmer's viewpoint. Popular C# unit testing frameworks The unit testing frameworks I'll be testing are: NUnit XUnit Built-in Visual Studio testing tools All of these unit testing frameworks offer a similar end goal, to help make writing unit tests faster, simpler and easier to learn. But there are still a few key differences between them. Some are more geared towards powerful, complex test cases, while others rank in simplicity and usability as a higher priority. In most versions since 2005, Visual Studio has come with a built in testing framework supported by Microsoft. This framework certainly was the most popular for its installation. Though if your copy of Visual Studio doesn't come with it already included you are going to have to jump through a few hoops to get it going. (We wrote a review of the 2017 version of Visual Studio here.) This framework is the simplest of the three, and uses an easy-to-understand method attribute structure (much like most testing frameworks) where you are able to add tags such as [TestClass] and [TestMethod] to your code in order to get testing. Visual Studio even has a UI panel dedicated to visualizing your tests, which can be found under Test > Windows > Test Explorer. Now before we dive into trying out this testing framework let's introduce our example classes that need testing. First we have a Raygun, which we can fire and recharge. The only thing we need to keep track of with our Raygun is it's ammo, which can run out. (We also have a bug, which we can shoot at with our Raygun. But we have the ability to dodge our attempts to shoot it. If we shoot at a bug after it has just dodged, we will miss. Though if we hit the bug square on, it's safe to assume that it will be dead. These two classes are defined as follows: public class Raygun { private int ammo = 3; public void FireAt(Bug bug) { if (HasAmmo()) { if (bug.IsDodging()) { bug.Miss(); } else { bug.Hit(); } ammo--; } } public void Recharge() { ammo = 3; } public bool HasAmmo() { return ammo > 0; } } public class Bug { private bool dodging; private bool dead; public void Dodge() { dodging = true; } public void Hit() { dead = true; } public void Miss() { dodging = false; } public bool IsDodging() { return dodging; } public bool IsDead() { return dead; } } Seems simple enough, but we need to make sure that our Rayguns and bugs behave as we want them to. So then it's time to write some unit tests! (We wrote about how to write robust unit tests in C# here.) First up let's try a simple test where we want to shoot at, and hit, a bug. We'll start with the Raygun class and see if we can get it to just shoot a bit of ammo left in it. Well, let's see if we are right. [TestClass] public
class Test { [TestMethod] public void TestShootDodgeBug() { Bug bug = new Bug(); Raygun gun = new Raygun(); gun.FireAt(bug); Assert.IsTrue(bug.IsDead()); Assert.IsTrue(gun.HasAmmo()); } } The two new things you will notice in this snippet of code is the [TestClass] and [TestMethod] tags, which certainly don't just float around in normal code. These tags are what allow Visual Studio's built in testing framework to recognize this particular class as a class that contains unit tests, and to treat the method TryShootBug() as a test case, instead of just an ordinary method. Since these tools are built for Visual Studio, running your tests from within Visual Studio is very simple. Just right click on any [TestMethod] tags as shown: And would you look at that, the test passed. Looks like our Raygun can at least hit a stationary bug. Of course this is only showing the bare basics of what Visual Studio's testing tools can do. Some other very useful tags you will surely be using are the [TestInitialize] and [TestCleanup] tags. These tags allow you to specify code that is run before (initialize) and after (cleanup) every individual test is run. So if you want to reload your Raygun after every encounter like a stylish gunslinger, then this should do the trick: [TestInitialize] public void Initialize() { gun = new Raygun(); } [TestCleanup] public void Cleanup() { gun.Recharge(); } Stylish. While we are still talking about the Visual Studio testing tools I'll quickly mention the [ExpectedException] tag, which is incredibly useful for when you want to deliberately cause an exception in your test (which you will certainly want to do at some point to make sure your program isn't accepting data it shouldn't). Here's a quick example of how you would write a test that results in an exception: [TestMethod] [ExpectedException(typeof(System.IndexOutOfRangeException))] public void TryMakingHolesOfGuns() { Raygun[] guns = new Raygun[]5; Bug bug = new Bug(); guns[5].FireAt(bug); } Don't worry about the index out of bounds exception, thanks to the [ExpectedException] tag the exception will be treated as a success and your test will pass. On the contrary if the exception isn't thrown then the test will fail. Anyway, it's about time that you can write a test case with arguments, then easily with a range of unit data. This means we need to write unique test cases for every set of arguments you want to test. Here's a quick example test case we could use to make sure our Raygun was actually running out of ammo at the right time, in a much smarter way than before: [TestCase(1)] [TestCase(2)] [TestCase(3)] [TestCase(4)] public void FireMultipleTimes(int fireCount) { Bug bug = new Bug(); Raygun gun = new Raygun(); for(int i = 0; i < fireCount; i++) { gun.FireAt(bug); } if (fireCount == 3) { Assert.IsFalse(gun.HasAmmo()); } else { Assert.IsTrue(gun.HasAmmo()); } } Excellent, with this one test case we were able to make sure a Raygun which has fired two shots still has ammo, while one that has fired three is empty. And thanks to the [TestCase] tag we were easily able to test a whole bunch of other values while we were at it! Overall NUnit is an excellent testing framework, and as you delve deeper into what it can offer, it surely exceeds what Microsoft's built in testing can offer. Anyway, let's look at our last testing framework, and our last attempt at shooting bugs with Rayguns! If you like the sound of Facts and Theories, then it's time to take look at XUnit XUnit is an open source testing platform with a larger focus in extensibility and flexibility. XUnit follows a more community minded development structure and focuses on being easy to expand upon. XUnit actually refers to a grouping of frameworks, but we will be focusing on the C# version. Other versions include JUnit, a very well known testing framework for Java. XUnit also uses a more modern and unique style of testing, by doing away with the standard [Test] [TestFixture] terminology and using new fancy tags like Facts and Theories. NUnit and XUnit are actually quite similar in many ways, as NUnit serves as a base for a lot of the new features XUnit brings forward. Note that XUnit is also installed via a NuGet package much like NUnit, which you can search for within Visual Studio. The packages I've used for this example are XUnit and XUnit.ConsoleRunner, though you also have the option of installing a GUI-based plugin for Visual Studio. Much like the [TestClass] tag in NUnit, XUnit has its own solution to providing parameters to a test case. To do so we will be using the new [InlineData] tag and Theories. In general, a test case that has no parameters (so it doesn't rely on any changing data) is referred to as a Fact in XUnit, meaning that it will always execute the same (so "Fact" suits it pretty well). On the other hand, we have Theories, which refers to a test case that can take data directly from [InlineData] tags or even from an Excel spreadsheet. So with all these new fancy keywords in mind, let's write a test in XUnit that uses a theory to test our bug dodge ability: [Theory][InlineData(true, false)] [InlineData(false, true)] public void TestBugDodges(bool didDodge, bool shouldBeDead) { Bug bug = new Bug(); Raygun gun = new Raygun(); if (didDodge) { bug.Dodge(); } gun.FireAt(bug); if (shouldBeDead) { Assert.True(bug.IsDead()); } } else { Assert.False(bug.IsDead()); } } This test covers both cases at once, where the bug dodges and survives, or doesn't dodge and gets hit. Lovely! Now, last step, lets run the XUnit test runner from the command line (note that much like NUnit, XUnit also has a GUI based visual studio plugin available for you to run tests with). First you must make sure you are in your project's test directory, just like NUnit (e.g. C:\Users\yourUserName\Documents\Visual Studio 2015\Projects\YourProjectName) and then enter the following command in a new cmd window: packages\xunit.runner.console.2.1.0\tools\xunit.console.exe YourProjectName\bin\Debug\YourProjectName.dll Assuming everything is set up properly, the XUnit console runner will run all the tests in your project and let you know how your tests turned out. Looks like our dodging tests passed! Overall XUnit acts as the more contemporary version of NUnit, offering flexible and usable testing with a fresh coat of paint. In conclusion... Regardless of which of the unit testing frameworks you use, you're going to be getting all the basics. However, there are a few differences between them that I hope I've highlighted so you can choose the right one for your project. Whether it's the convenience of Microsoft's built in unit testing framework, the solid and well proven status of NUnit, or the modern take on unit testing that XUnit provides, theres always something out there that will give you exactly what you need! Want to add an extra layer of protection for your code? Catch the errors that fall through the cracks with Raygun. Take a free trial here. Complex embedded projects have thousands and frequently tens of thousands lines of code. The majority of that code is entirely software (rather than "firmware"), and software in every industry is typically unit tested. However, in the embedded and firmware industry, unit testing is typically an after-thought or a task that is begun after working on a project for months or even years. Today's firmware projects require filesystems, BLE and Wi-Fi stacks, specialized data structures (both in-memory and in-flash), and complex algorithms, such as those interpreting accelerometer and gyroscope data. All of these items can be easily unit tested after becoming acquainted with best practices and writing a few tests of your own. In this post, we go into detail on how to properly build abstractions to stub, fake, and mock out implementations of low level embedded software and provide a full real-world example of a unit test using the CppUTest 3.8 unit test framework. This is the second post in our Building Better Firmware series, following the post about Continuous Integration for firmware projects, which is a wonderful pre-cursor to this post. Unit Testing Overview Unit testing is a method of testing software where individual software components are isolated and tested for correctness. Ideally, these unit tests are able to cover most if not all of the code paths, argument bounds, and failure cases of the software under test. Through proper use of unit tests, and especially while using practices from Test Driven Development (TDD1), the time it takes to stabilize embedded software can decrease dramatically, making individuals and teams more productive and firmware less likely to experience functional bugs, control flow bugs, and even fatal issues, such as memory leaks and (gasp!) bootloops. Life Before Unit Testing Here are a few examples that I've experienced in the past that were alleviated by the team doubling down on unit testing the firmware. You find testing on target hardware slow and inconvenient, especially when multiple devices (e.g. a mobile phone) or prior setup (e.g. a factory reset) is required Bugs and regressions occur repeatedly in a single piece of software. Deadlocks, HardFaults and Memory Leaks are in the norm and have become accepted unit testing examples for preventing deadlocks (included below). The amount of time spent debugging and testing firmware is 60% or more2. The fix instinct when starting a new software module is to write a chunk of code and test on hardware. Life After Unit Testing (Possibly) At a previous company, after scraping most legacy code and writing new modules with 90%+ code coverage and through the use of TDD, this is what development felt like sometimes. You write a new file, maybe an in-flash log storage module, and it works the first time when flashed on the device (no better file testing). Regressions are caught immediately when running tests locally or in CI. Memory leaks are raised as errors in unit tests. Testing a majority of the
firmware only takes a minute. The overall codebase has better structure and cleaner boundaries between modules. Disclaimers Unit tests in the embedded space is a controversial topics, so I want to clear set up expectations up front. This post covers how to test embedded software. Testing firmware drivers and hardware is very different and time is best spent writing functional and integration tests that run on target to validate hardware components. As soon as drivers are written and stable, switch to the unit test approaches provided in this post. I do not suggest rewriting all of your code to accommodate unit tests, or writing tests for the current code base, but I heavily suggest writing unit tests for most new modules and heavily suggesting them in code reviews. Integration tests and on-target tests have their place. This infrastructure is a huge time and money investment, and the tests run in minutes and hours. Keep these to a minimum at first to ensure hardware stability, and leave software stability to unit tests. If time allows, then build these types of tests. Like Interrupt? Subscribe to get our latest posts straight to your inbox. Framework-less Unit Tests It is very common to initially write unit tests using one-off .c files. Below is an example of a test that is commonly found in firmware projects or written by the author of a piece of firmware code. #include < In my sum.c int my sum(int a, int b) { return a + b; } // In test my sum.c int main(int argc, char *argv[]) { assert(2 == my sum(1, 1)); assert(2 == my sum(1, -1)); assert(0 == my sum(0, 0)); // ... return(0); } This works for a short period of time, but as a firmware project grows in complexity, lines of code, and number of developers, there are a few things that become requirements. Running unit tests in Continuous Integration Reporting results failed, run time duration, etc. Reporting code coverage to give insight into how much code coverage is in place. Ability for a developer to create a new unit test easily and quickly. The most scalable way to write unit tests in C is using a unit testing framework, such as: CppUTest Unity Google Test Even though CppUTest and Google Test are written in C++, they can be used to test C source code, as long as the C headers/files are wrapped with extern "C" { #include "my_sum.h" } Minimal Unit Test Example Let's come up with a bare bones unit test to instrument our simple my_sum module. NOTE: Our examples use the CppUTest framework. If you want to follow along, check out the Setting Up CppUTest section first. The source code for the my_sum.c module is as follows: #include "my_sum.h" int my sum(int a, int b) { return (a + b); } A unit test generally contains the following pieces: Setup and Teardown functions, which run before and after each test respectively. Individual tests that test logical components or paths of a module. Many checks, such as LONGS_EQUAL which compares integer values and STRCMP_EQUAL which would compare string values. Our basic unit test is as follows: #include "CppUTest/TestHarness.h" extern "C" { #include "my_sum.h" } TEST_GROUP(TestMySum) { void setup() { // This gets run before every test } void teardown() { // This gets run after every test } ; TEST(TestMySum, Test MySumBasic) { LONGS_EQUAL(7, my sum(3, 4)); } Although the example is basic, let's go over what is happening here. We import my_sum.h inside of the extern "C" { section so that it is compiled as C instead of C++. We have empty setup() and teardown() functions since the modules we are testing don't require any initial setup or cleanup routines. We have a single LONGS_EQUAL state, which compares == after the my_sum function is called. We did not include any fakes or stubs, as our module didn't have any dependencies. If this test passes, we get something like: Running build/sum/sum tests . OK (1 tests, 1 ran, 1 checks, 0 ignored, 0 filtered out, 0 ms) And if the test fails (for example, change 7 to 6): Running build/sum/sum tests src/test my_sum.cpp:17: error: Failure in TEST(TestMySum, Test MySumBasic) expected but: Errors (1 failures, 1 tests, 1 ran, 1 checks, 0 ignored, 0 filtered out, 1 ms) To build and run this unit test, we give the unit harness the test binary, the list of files to compile into the test binary, and any extra compilation flags necessary. COMPONENT_NAME= sum, SRC_FILES = \$(PROJECT_SRC_DIR)/my_sum.c \TEST_SRC_FILES = \SUMUNITEST_SRC_DIR/test_my_sum.c Here, we have SRC_FILES, which would contain any sources file used by the test, and TEST_SRC_FILES which contains the test files that contain the tests themselves. Unit Testing Best Practices: The "Minimal Example" is a contrived example and very rarely will there be a test with no other dependencies. Firmware is naturally coupled with other parts of hardware, and that makes it difficult at first to set up a unit test. For example, a flash storage module may call an analytics inc() function to record the number of writes, a watchdog feed() function during a large flash erase operation, and timer_schedule() to help defragment the flash later in the future. If we are testing only the flash key/value store, we do not want to include the analytics, watchdog, and timer source files into our unit test. That brings us to a few best practices to follow, especially when writing unit tests for complex and entangled code. Each TEST() within a unit test file should ideally test a single path or feature of the module. A test called TestEverything is an anti-pattern. Each test should be quick. A few milliseconds is ideal, and one second is the worst case run time. Each unit test should ideally include one real implementation of a module. The rest should be stubbed or fake versions of the modules not under test. Those stubbed and fake versions of modules should be written early, reused, and shared. Which brings us to explaining what are stubs, fakes, and mocks? Stubs, Fakes, and Mocks When starting to write unit tests, it is common to write alternate implementations to modules that make sense for a particular unit test. Since unit tests will be run on the host machine, they won't have hardware, such as an LED. But if a module within a unit test calls Enable_LED(), we could instead have a virtual LED and a state boolean value saving whether the LED is on or off. These alternate implementations of modules have different types. Let's explain them. Fakes are a working implementation, but will usually substitute with something simpler and easier to implement. Example: an in-memory key/value store vs a NOR Flash based Key/Value store. Stubs are a trivial implementation that returns canned values. They are always returning valid or invalid values. Mocks are an implementation that is controlled by the unit test. They can be pre-programmed with return values, check values of arguments, and help verify that functions are called. After having worked at two software oriented hardware companies with 20+ firmware engineers each, my favorite way to organize the test directory is as follows: |--- header overrides |--- string.h |--- error codes.h |--- fakes |--- fake analytics.c |--- fake analytics.h |--- fake kv_store.c |--- fake mutex.c |--- fake mutex.h |--- stubs |--- stub analytics.h |--- stub kv_store.h |--- stub mutex.h |--- mocks |--- mock analytics.h |--- mock kv_store.h |--- src |--- AllTests.cpp |--- test kv_store.cpp |--- test littlefs_basic.cpp |--- test littlefs_format.cpp |--- makefiles |--- Makefile littlefs_basic.mk |--- Makefile littlefs_format.mk |--- Makefile settings_file.mk where header overrides/ headers that override complex, auto-generated, or target specific headers. src/ - The unit tests themselves. makefiles/ - Makefiles to build the individual unit tests. Stubs These are used when the implementation of specific functions or their return values do not matter to the module under test. They are primarily used to fix the linker's ld: symbol(s) not found errors. These should generally have only a return statement that always returns true, false, 0, NULL, or whatever makes sense in the context of the module. If there is anything more complex than a return statement, consider implementing a Fake instead. Examples: Hardware or peripheral initialization functions since they have little relevance in testing on the host (x86 machine). A time module which returns the time of day (just return a random time). Mutex stubs when the locking or unlocking isn't being tested. (shown below) #include "mutex/mutex.h" // Stubs Mutex "mutex create(void) { return NULL; } void mutex lock(Mutex "mutex" { return; } void mutex unlock(Mutex "mutex" { return; } Fakes A Fake is commonly used in firmware where it is impractical to use the real implementation for reasons such as: It requires specific hardware (flash chip, peripherals, LEDs, etc.) Examples: A mutex module which checks at the end of that all mutexes were properly unlocked (example provided later in the post). A RAM based NOR flash implementation, which flips bits from 1 to 0 when writing and requires flipping bits back to 1 before writing new data. A RAM based key-value store (example shown below). #include "kv_store.h" #include #include typedef struct { char *key; void *val; uint32_t len; } KvEntry; static KvEntry s_kv_store[256]; bool kv_store_write(const char *key, const void *val, uint32_t len) { // Write key/value into RAM store } bool kv_store_read(const char *key, void *buf, uint32_t buf_len, uint32_t *len_read) { // Read key/value from RAM store into buffer } bool kv_store_delete(const char *key) { // Delete key/value from RAM store } Mocks Mocks are incredibly useful if you want to declare each and every return value of a given function. Using many mocks in a single unit test is also the easiest way to test every single code path of the module under test, as you can force any function to return error codes, NULL values, and invalid pointers. These are the most powerful, provide the programmer the most control, and
isolate the module under test the best, but they are also the most cumbersome and verbose to use, as every return value has to be pre-programmed. Common mocking frameworks include: We are not going to cover examples of mocks and how to implement them (the topic is big enough for another post), but some pseudo code is shown below to give an understanding: Learn more about mocks in our separate post, Unit Testing with Mocks. TEST(TestKvStore, Test_InitMockCreated) { // On the next call to 'my_malloc', return the value 'NULL'. MOCK my_malloc.return_value = NULL; // This calls 'my_malloc' in its implementation. void *buf = allocate_buffer(); // Ensure that 'my_malloc' was called once and only once. LONGS_EQUAL(1, MOCK my_malloc.call_count); // Ensure that the buffer returned was indeed 'NULL' since 'my_malloc' returned 'NULL'. POINTERS_EQUAL(NULL, buf); } Examples: A malloc implementation that can be pre-programmed with return values (return read buffers vs NULL). A mock flash driver which returns error codes and forces different paths in a higher level module. A Bluetooth socket implementation which is fed artificially crafted packed data after scraping most legacy code and writing new modules with 90%+ code coverage and through the use of TDD, this is what development felt like sometimes. You write a new file, maybe an in-flash log storage module, and it works the first time when flashed on the device (no better file testing). Regressions are caught immediately when running tests locally or in CI. Memory leaks are raised as errors in unit tests. Testing a majority of the firmware only takes a minute. The overall codebase has better structure and cleaner boundaries between modules. Disclaimers Unit tests in the embedded space is a controversial topics, so I want to clear set up expectations up front. This post covers how to test embedded software. Testing firmware drivers and hardware is very different and time is best spent writing functional and integration tests that run on target to validate hardware components. As soon as drivers are written and stable, switch to the unit test approaches provided in this post. I do not suggest rewriting all of your code to accommodate unit tests, or writing tests for the current code base, but I heavily suggest writing unit tests for most new modules and heavily suggesting them in code reviews. Integration tests and on-target tests have their place. This infrastructure is a huge time and money investment, and the tests run in minutes and hours. Keep these to a minimum at first to ensure hardware stability, and leave software stability to unit tests. If time allows, then build these types of tests. Like Interrupt? Subscribe to get our latest posts straight to your inbox. Framework-less Unit Tests It is very common to initially write unit tests using one-off .c files. Below is an example of a test that is commonly found in firmware projects or written by the author of a piece of firmware code. #include < In my sum.c int my sum(int a, int b) { return a + b; } // In test my sum.c int main(int argc, char *argv[]) { assert(2 == my sum(1, 1)); assert(2 == my sum(1, -1)); assert(0 == my sum(0, 0)); // ... return(0); } This works for a short period of time, but as a firmware project grows in complexity, lines of code, and number of developers, there are a few things that become requirements. Running unit tests in Continuous Integration Reporting results failed, run time duration, etc. Reporting code coverage to give insight into how much code coverage is in place. Ability for a developer to create a new unit test easily and quickly. The most scalable way to write unit tests in C is using a unit testing framework, such as: CppUTest Unity Google Test Even though CppUTest and Google Test are written in C++, they can be used to test C source code, as long as the C headers/files are wrapped with extern "C" { #include "my_sum.h" } Minimal Unit Test Example Let's come up with a bare bones unit test to instrument our simple my_sum module. NOTE: Our examples use the CppUTest framework. If you want to follow along, check out the Setting Up CppUTest section first. The source code for the my_sum.c module is as follows: #include "my_sum.h" int my sum(int a, int b) { return (a + b); } A unit test generally contains the following pieces: Setup and Teardown functions, which run before and after each test respectively. Individual tests that test logical components or paths of a module. Many checks, such as LONGS_EQUAL which compares integer values and STRCMP_EQUAL which would compare string values. Our basic unit test is as follows: #include "CppUTest/TestHarness.h" extern "C" { #include "my_sum.h" } TEST_GROUP(TestMySum) { void setup() { // This gets run before every test } void teardown() { // This gets run after every test } ; TEST(TestMySum, Test MySumBasic) { LONGS_EQUAL(7, my sum(3, 4)); } Although the example is basic, let's go over what is happening here. We import my_sum.h inside of the extern "C" { section so that it is compiled as C instead of C++. We have empty setup() and teardown() functions since the modules we are testing don't require any initial setup or cleanup routines. We have a single LONGS_EQUAL state, which compares == after the my_sum function is called. We did not include any fakes or stubs, as our module didn't have any dependencies. If this test passes, we get something like: Running build/sum/sum tests . OK (1 tests, 1 ran, 1 checks, 0 ignored, 0 filtered out, 0 ms) And if the test fails (for example, change 7 to 6): Running build/sum/sum tests src/test my_sum.cpp:17: error: Failure in TEST(TestMySum, Test MySumBasic) expected but: Errors (1 failures, 1 tests, 1 ran, 1 checks, 0 ignored, 0 filtered out, 1 ms) To build and run this unit test, we give the unit harness the test binary, the list of files to compile into the test binary, and any extra compilation flags necessary. COMPONENT_NAME= sum, SRC_FILES = \$(PROJECT_SRC_DIR)/my_sum.c \TEST_SRC_FILES = \SUMUNITEST_SRC_DIR/test_my_sum.c Here, we have SRC_FILES, which would contain any sources file used by the test, and TEST_SRC_FILES which contains the test files that contain the tests themselves. Unit Testing Best Practices: The "Minimal Example" is a contrived example and very rarely will there be a test with no other dependencies. Firmware is naturally coupled with other parts of hardware, and that makes it difficult at first to set up a unit test. For example, a flash storage module may call an analytics inc() function to record the number of writes, a watchdog feed() function during a large flash erase operation, and timer_schedule() to help defragment the flash later in the future. If we are testing only the flash key/value store, we do not want to include the analytics, watchdog, and timer source files into our unit test. That brings us to a few best practices to follow, especially when writing unit tests for complex and entangled code. Each TEST() within a unit test file should ideally test a single path or feature of the module. A test called TestEverything is an anti-pattern. Each test should be quick. A few milliseconds is ideal, and one second is the worst case run time. Each unit test should ideally include one real implementation of a module. The rest should be stubbed or fake versions of the modules not under test. Those stubbed and fake versions of modules should be written early, reused, and shared. Which brings us to explaining what are stubs, fakes, and mocks? Stubs, Fakes, and Mocks When starting to write unit tests, it is common to write alternate implementations to modules that make sense for a particular unit test. Since unit tests will be run on the host machine, they won't have hardware, such as an LED. But if a module within a unit test calls Enable_LED(), we could instead have a virtual LED and a state boolean value saving whether the LED is on or off. These alternate implementations of modules have different types. Let's explain them. Fakes are a working implementation, but will usually substitute with something simpler and easier to implement. Example: an in-memory key/value store vs a NOR Flash based Key/Value store. Stubs are a trivial implementation that returns canned values. They are always returning valid or invalid values. Mocks are an implementation that is controlled by the unit test. They can be pre-programmed with return values, check values of arguments, and help verify that functions are called. After having worked at two software oriented hardware companies with 20+ firmware engineers each, my favorite way to organize the test directory is as follows: |--- header overrides |--- string.h |--- error codes.h |--- fakes |--- fake analytics.c |--- fake analytics.h |--- fake kv_store.c |--- fake mutex.c |--- fake mutex.h |--- stubs |--- stub analytics.h |--- stub kv_store.h |--- stub mutex.h |--- mocks |--- mock analytics.h |--- mock kv_store.h |--- src |--- AllTests.cpp |--- test kv_store.cpp |--- test littlefs_basic.cpp |--- test littlefs_format.cpp |--- makefiles |--- Makefile littlefs_basic.mk |--- Makefile littlefs_format.mk |--- Makefile settings_file.mk where header overrides/ headers that override complex, auto-generated, or target specific headers. src/ - The unit tests themselves. makefiles/ - Makefiles to build the individual unit tests. Stubs These are used when the implementation of specific functions or their return values do not matter to the module under test. They are primarily used to fix the linker's ld: symbol(s) not found errors. These should generally have only a return statement that always returns true, false, 0, NULL, or whatever makes sense in the context of the module. If there is anything more complex than a return statement, consider implementing a Fake instead. Examples: Hardware or peripheral initialization functions since they have little relevance in testing on the host (x86 machine). A time module which returns the time of day (just return a random
time). Mutex stubs when the locking or unlocking isn't being tested. (shown below) #include "mutex/mutex.h" // Stubs Mutex "mutex create(void) { return NULL; } void mutex lock(Mutex "mutex" { return; } void mutex unlock(Mutex "mutex" { return; } Fakes A Fake is commonly used in firmware where it is impractical to use the real implementation for reasons such as: It requires specific hardware (flash chip, peripherals, LEDs, etc.) Examples: A mutex module which checks at the end of that all mutexes were properly unlocked (example provided later in the post). A RAM based NOR flash implementation, which flips bits from 1 to 0 when writing and requires flipping bits back to 1 before writing new data. A RAM based key-value store (example shown below). #include "kv_store.h" #include #include typedef struct { char *key; void *val; uint32_t len; } KvEntry; static KvEntry s_kv_store[256]; bool kv_store_write(const char *key, const void *val, uint32_t len) { // Write key/value into RAM store } bool kv_store_read(const char *key, void *buf, uint32_t buf_len, uint32_t *len_read) { // Read key/value from RAM store into buffer } bool kv_store_delete(const char *key) { // Delete key/value from RAM store } Mocks Mocks are incredibly useful if you want to declare each and every return value of a given function. Using many mocks in a single unit test is also the easiest way to test every single code path of the module under test, as you can force any function to return error codes, NULL values, and invalid pointers. These are the most powerful, provide the programmer the most control, and isolate the module under test the best, but they are also the most cumbersome and verbose to use, as every return value has to be pre-programmed. Common mocking frameworks include: We are not going to cover examples of mocks and how to implement them (the topic is big enough for another post), but some pseudo code is shown below to give an understanding: Learn more about mocks in our separate post, Unit Testing with Mocks. TEST(TestKvStore, Test_InitMockCreated) { // On the next call to 'my_malloc', return the value 'NULL'. MOCK my_malloc.return_value = NULL; // This calls 'my_malloc' in its implementation. void *buf = allocate_buffer(); // Ensure that 'my_malloc' was called once and only once. LONGS_EQUAL(1, MOCK my_malloc.call_count); // Ensure that the buffer returned was indeed 'NULL' since 'my_malloc' returned 'NULL'. POINTERS_EQUAL(NULL, buf); } Examples: A malloc implementation that can be pre-programmed with return values (return read buffers vs NULL). A mock flash driver which returns error codes and forces different paths in a higher level module. A Bluetooth socket implementation which is fed artificially crafted packed data after scraping most legacy code and writing new modules with 90%+ code coverage and through the use of TDD, this is what development felt like sometimes. You write a new file, maybe an in-flash log storage module, and it works the first time when flashed on the device (no better file testing). Regressions are caught immediately when running tests locally or in CI. Memory leaks are raised as errors in unit tests. Testing a majority of the firmware only takes a minute. The overall codebase has better structure and cleaner boundaries between modules. Disclaimers Unit tests in the embedded space is a controversial topics, so I want to clear set up expectations up front. This post covers how to test embedded software. Testing firmware drivers and hardware is very different and time is best spent writing functional and integration tests that run on target to validate hardware components. As soon as drivers are written and stable, switch to the unit test approaches provided in this post. I do not suggest rewriting all of your code to accommodate unit tests, or writing tests for the current code base, but I heavily suggest writing unit tests for most new modules and heavily suggesting them in code reviews. Integration tests and on-target tests have their place. This infrastructure is a huge time and money investment, and the tests run in minutes and hours. Keep these to a minimum at first to ensure hardware stability, and leave software stability to unit tests. If time allows, then build these types of tests. Like Interrupt? Subscribe to get our latest posts straight to your inbox. Framework-less Unit Tests It is very common to initially write unit tests using one-off .c files. Below is an example of a test that is commonly found in firmware projects or written by the author of a piece of firmware code. #include < In my sum.c int my sum(int a, int b) { return a + b; } // In test my sum.c int main(int argc, char *argv[]) { assert(2 == my sum(1, 1)); assert(2 == my sum(1, -1)); assert(0 == my sum(0, 0)); // ... return(0); } This works for a short period of time, but as a firmware project grows in complexity, lines of code, and number of developers, there are a few things that become requirements. Running unit tests in Continuous Integration Reporting results failed, run time duration, etc. Reporting code coverage to give insight into how much code coverage is in place. Ability for a developer to create a new unit test easily and quickly. The most scalable way to write unit tests in C is using a unit testing framework, such as: CppUTest Unity Google Test Even though CppUTest and Google Test are written in C++, they can be used to test C source code, as long as the C headers/files are wrapped with extern "C" { #include "my_sum.h" } Minimal Unit Test Example Let's come up with a bare bones unit test to instrument our simple my_sum module. NOTE: Our examples use the CppUTest framework. If you want to follow along, check out the Setting Up CppUTest section first. The source code for the my_sum.c module is as follows: #include "my_sum.h" int my sum(int a, int b) { return (a + b); } A unit test generally contains the following pieces: Setup and Teardown functions, which run before and after each test respectively. Individual tests that test logical components or paths of a module. Many checks, such as LONGS_EQUAL which compares integer values and STRCMP_EQUAL which would compare string values. Our basic unit test is as follows: #include "CppUTest/TestHarness.h" extern "C" { #include "my_sum.h" } TEST_GROUP(TestMySum) { void setup() { // This gets run before every test } void teardown() { // This gets run after every test } ; TEST(TestMySum, Test MySumBasic) { LONGS_EQUAL(7, my sum(3, 4)); } Although the example is basic, let's go over what is happening here. We import my_sum.h inside of the extern "C" { section so that it is compiled as C instead of C++. We have empty setup() and teardown() functions since the modules we are testing don't require any initial setup or cleanup routines. We have a single LONGS_EQUAL state, which compares == after the my_sum function is called. We did not include any fakes or stubs, as our module didn't have any dependencies. If this test passes, we get something like: Running build/sum/sum tests . OK (1 tests, 1 ran, 1 checks, 0 ignored, 0 filtered out, 0 ms) And if the test fails (for example, change 7 to 6): Running build/sum/sum tests src/test my_sum.cpp:17: error: Failure in TEST(TestMySum, Test MySumBasic) expected but: Errors (1 failures, 1 tests, 1 ran, 1 checks, 0 ignored, 0 filtered out, 1 ms) To build and run this unit test, we give the unit harness the test binary, the list of files to compile into the test binary, and any extra compilation flags necessary. COMPONENT_NAME= sum, SRC_FILES = \$(PROJECT_SRC_DIR)/my_sum.c \TEST_SRC_FILES = \SUMUNITEST_SRC_DIR/test_my_sum.c Here, we have SRC_FILES, which would contain any sources file used by the test, and TEST_SRC_FILES which contains the test files that contain the tests themselves. Unit Testing Best Practices: The "Minimal Example" is a contrived example and very rarely will there be a test with no other dependencies. Firmware is naturally coupled with other parts of hardware, and that makes it difficult at first to set up a unit test. For example, a flash storage module may call an analytics inc() function to record the number of writes, a watchdog feed() function during a large flash erase operation, and timer_schedule() to help defragment the flash later in the future. If we are testing only the flash key/value store, we do not want to include the analytics, watchdog, and timer source files into our unit test. That brings us to a few best practices to follow, especially when writing unit tests for complex and entangled code. Each TEST() within a unit test file should ideally test a single path or feature of the module. A test called TestEverything is an anti-pattern. Each test should be quick. A few milliseconds is ideal, and one second is the worst case run time. Each unit test should ideally include one real implementation of a module. The rest should be stubbed or fake versions of the modules not under test. Those stubbed and fake versions of modules should be written early, reused, and shared. Which brings us to explaining what are stubs, fakes, and mocks? Stubs, Fakes, and Mocks When starting to write unit tests, it is common to write alternate implementations to modules that make sense for a particular unit test. Since unit tests will be run on the host machine, they won't have hardware, such as an LED. But if a module within a unit test calls Enable_LED(), we could instead have a virtual LED and a state boolean value saving whether the LED is on or off. These alternate implementations of modules have different types. Let's explain them. Fakes are a working implementation, but will usually substitute with something simpler and easier to implement. Example: an in-memory key/value store vs a NOR Flash based Key/Value store. Stubs are a trivial implementation that returns canned values. They are always returning valid or invalid values. Mocks are an implementation that is controlled by the unit test. They can be pre-programmed with
return values, check values of arguments, and help verify that functions are called. After having worked at two software oriented hardware companies with 20+ firmware engineers each, my favorite way to organize the test directory is as follows: |--- header overrides |--- string.h |--- error codes.h |--- fakes |--- fake analytics.c |--- fake analytics.h |--- fake kv_store.c |--- fake mutex.c |--- fake mutex.h |--- stubs |--- stub analytics.h |--- stub kv_store.h |--- stub mutex.h |--- mocks |--- mock analytics.h |--- mock kv_store.h |--- src |--- AllTests.cpp |--- test kv_store.cpp |--- test littlefs_basic.cpp |--- test littlefs_format.cpp |--- makefiles |--- Makefile littlefs_basic.mk |--- Makefile littlefs_format.mk |--- Makefile settings_file.mk where header overrides/ headers that override complex, auto-generated, or target specific headers. src/ - The unit tests themselves. makefiles/ - Makefiles to build the individual unit tests. Stubs These are used when the implementation of specific functions or their return values do not matter to the module under test. They are primarily used to fix the linker's ld: symbol(s) not found errors. These should generally have only a return statement that always returns true, false, 0, NULL, or whatever makes sense in the context of the module. If there is anything more complex than a return statement, consider implementing a Fake instead. Examples: Hardware or peripheral initialization functions since they have little relevance in testing on the host (x86 machine). A time module which returns the time of day (just return a random time). Mutex stubs when the locking or unlocking isn't being tested. (shown below) #include "mutex/mutex.h" // Stubs Mutex "mutex create(void) { return NULL; } void mutex lock(Mutex "mutex" { return; } void mutex unlock(Mutex "mutex" { return; } Fakes A Fake is commonly used in firmware where it is impractical to use the real implementation for reasons such as: It requires specific hardware (flash chip, peripherals, LEDs, etc.) Examples: A mutex module which checks at the end of that all mutexes were properly unlocked (example provided later in the post). A RAM based NOR flash implementation, which flips bits from 1 to 0 when writing and requires flipping bits back to 1 before writing new data. A RAM based key-value store (example shown below). #include "kv_store.h" #include #include typedef struct { char *key; void *val; uint32_t len; } KvEntry; static KvEntry s_kv_store[256]; bool kv_store_write(const char *key, const void *val, uint32_t len) { // Write key/value into RAM store } bool kv_store_read(const char *key, void *buf, uint32_t buf_len, uint32_t *len_read) { // Read key/value from RAM store into buffer } bool kv_store_delete(const char *key) { // Delete key/value from RAM store } Mocks Mocks are incredibly useful if you want to declare each and every return value of a given function. Using many mocks in a single unit test is also the easiest way to test every single code path of the module under test, as you can force any function to return error codes, NULL values, and invalid pointers. These are the most powerful, provide the programmer the most control, and isolate the module under test the best, but they are also the most cumbersome and verbose to use, as every return value has to be pre-programmed. Common mocking frameworks include: We are not going to cover examples of mocks and how to implement them (the topic is big enough for another post), but some pseudo code is shown below to give an understanding: Learn more about mocks in our separate post, Unit Testing with Mocks. TEST(TestKvStore, Test_InitMockCreated) { // On the next call to 'my_malloc', return the value 'NULL'. MOCK my_malloc.return_value = NULL; // This calls 'my_malloc' in its implementation. void *buf = allocate_buffer(); // Ensure that 'my_malloc' was called once and only once. LONGS_EQUAL(1, MOCK my_malloc.call_count); // Ensure that the buffer returned was indeed 'NULL' since 'my_malloc' returned 'NULL'. POINTERS_EQUAL(NULL, buf); } Examples: A malloc implementation that can be pre-programmed with return values (return read buffers vs NULL). A mock flash driver which returns error codes and forces different paths in a higher level module. A Bluetooth socket implementation which is fed artificially crafted packed data after scraping most legacy code and writing new modules with 90%+ code coverage and through the use of TDD, this is what development felt like sometimes. You write a new file, maybe an in-flash log storage module, and it works the first time when flashed on the device (no better file testing). Regressions are caught immediately when running tests locally or in CI. Memory leaks are raised as errors in unit tests. Testing a majority of the firmware only takes a minute. The overall codebase has better structure and cleaner boundaries between modules. Disclaimers Unit tests in the embedded space is a controversial topics, so I want to clear set up expectations up front. This post covers how to test embedded software. Testing firmware drivers and hardware is very different and time is best spent writing functional and integration tests that run on target to validate hardware components. As soon as drivers are written and stable, switch to the unit test approaches provided in this post. I do not suggest rewriting all of your code to accommodate unit tests, or writing tests for the current code base, but I heavily suggest writing unit tests for most new modules and heavily suggesting them in code reviews. Integration tests and on-target tests have their place. This infrastructure is a huge time and money investment, and the tests run in minutes and hours. Keep these to a minimum at first to ensure hardware stability, and leave software stability to unit tests. If time allows, then build these types of tests. Like Interrupt? Subscribe to get our latest posts straight to your inbox. Framework-less Unit Tests It is very common to initially write unit tests using one-off .c files. Below is an example of a test that is commonly found in firmware projects or written by the author of a piece of firmware code. #include < In my sum.c int my sum(int a, int b) { return a + b; } // In test my sum.c int main(int argc, char *argv[]) { assert