

FRUCTOSE – a C++ unit test framework

Rationale: Why yet another one?

Andrew Peter Marlow
44 Stanhope Road, North Finchley
London N12 9DT, England
`andrew@andrewpetermarlow.co.uk`

May 14, 2011

Contents

1	Introduction	1
2	Why another C++ unit test framework?	3
3	The FRUCTOSE approach	5
3.1	The Curiously Recurring Template Pattern (CRTP)	5
3.2	Named tests and command line options	7
3.3	To inherit or not to inherit	7
3.4	The assert macros	9
4	How FRUCTOSE works	11
5	FRUCTOSE coding techniques	13
6	FRUCTOSE and the xUnit architecture	15
6.1	What is the xUnit architecture	15
6.2	Comparing FRUCTOSE to xUnit	15
6.2.1	Test case	16
6.2.2	Test fixtures	16
6.2.3	Test suites	17
6.2.4	Test execution	17
6.2.5	Assertions	17
7	The problems with CppUnit	19
8	The problems with other frameworks	21
9	The advantages of FRUCTOSE	23
10	Possible future work	25
11	Conclusion	27

Chapter 1

Introduction

FRUCTOSE is a simple unit test framework for C++. It stands for FRamework for Unit testing C++ for Test Driven development of SoftwarE.

The purpose of this article is to describe the C++ techniques that were used to develop FRUCTOSE, why and how FRUCTOSE is different from other C++ unit test frameworks and the advantages this confers for Test Driven Development (TDD). It is based on an article I wrote for the ACCU (Association of C and C++ Users) when FRUCTOSE was first released.

For details on how to use FRUCTOSE, see the user guide.

FRUCTOSE is designed to be much smaller and simpler than frameworks such as CppUnit. It offers the ability to quickly create small standalone programs that produce output on standard out and that can be driven by the command line. The idea is that by using command line options to control which tests are run, the level of verbosity and whether or not the first error is fatal, it will be easier to develop classes in conjunction with their test harness. This makes FRUCTOSE an aid to Test Driven Development (TDD). This is a slightly different way of approaching unit test frameworks. Most other frameworks seem to assume the class has already been designed and coded and that the test framework is to be used to run some tests as part of the overnight build regime. Whilst FRUCTOSE can be used in this way, it is hoped that the class and its tests will be developed at the same time.

A simple unit test framework should be concerned with just the functions being tested (they do not have to be class member functions although they typically will be) and the collection of functions the developer has to write to get those tests done. These latter functions can be grouped into a testing class. This article shows by example how this testing class is typically written and what facilities are available.

There are certain kinds of assertion tests that are useful for any unit test framework to offer. This article only briefly mentions these. They are covered in more depth in the user guide.

FRUCTOSE has been developed as free software (ref [15]). This meant choosing an appropriate license for it. The intent is to have a license that is non-

product specific with an element of copyleft but that still allows FRUCTOSE to be used in proprietary software. The license chosen was the LGPL license (Lesser General Public License). The weaker copyleft provision of the LGPL allows FRUCTOSE to be used in proprietary software because the LGPL allows a program to be linked with non-free modules. A proprietary project can use FRUCTOSE to unit test its proprietary classes and does not have to release those classes under the GPL or LGPL. However, if such a project releases a modified version of FRUCTOSE then it must do so under the terms of LGPL. In the search for a suitable license the LGPL was the winner by default. It is the only license listed on the Free Software Foundation's web site that is non-product specific and a GPL-compatible free software license that has some copyleft provision. Version 3 of the LGPL was chosen.

FRUCTOSE does not follow the unit test framework pattern known as xUnit. This is to do with the aim of making FRUCTOSE cut down and much simpler than other frameworks. It is discussed in detail in chapter 6.

Chapter 2

Why another C++ unit test framework?

When I first started to look seriously at C++ unit testing I looked at CppUnit. Conventional wisdom said to build on the work of others rather than reinvent a framework, and here was an established one which was itself built on the work of JUnit. The feature list is impressive and it looked like it would be mature enough to give a trouble-free build and be usable right away. Unfortunately, this proved not to be the case. Since FRUCTOSE was written several other C++ unit test frameworks have been developed, many of which were written because of dis-satisfaction with CppUnit.

A quick search of sourceforge reveals that there are a few products for C++ unit testing. These also turn out to have issues of their own (discussed later).

In the October 2006 issue of *Overload* from the ACCU there is an article by Peter Sommerlad on a C++ unit testing framework called CUTE (ref [13]). This article also mentions that there are issues with using CppUnit and that some developers want something else that is smaller and simpler. I wondered if CUTE could be extended to address some of my concerns with alternatives like CppUnit. However, at that time CUTE was not available online and Peter was too busy to work on a collaboration. This provided the motivation for me to write something which is driven by the same need expressed in the CUTE article (smaller and simpler than CppUnit). There are some important differences between FRUCTOSE and CUTE. FRUCTOSE avoids two significant dependencies; one on Boost and the other on platform-specific RTTI.

Chapter 3

The FRUCTOSE approach

FRUCTOSE has a simple objective: provide enough functionality such that the developer can produce one standalone command line driven program per implementation file (usually one class per file). This means that most FRUCTOSE programs will use only two classes; the test class and the class being tested. This means that unlike other test frameworks, FRUCTOSE is not designed to be extensible. It was felt that the flexibility of other frameworks comes at the cost of increased size and complexity. Most other frameworks expect the developer to derive other classes from the framework ones to modify or specialise behaviour in some way, e.g. to provide HTML output instead of simple TTY output. They also provide the ability to run multiple suites. FRUCTOSE does not offer this. The test harness is expected to simply consist of a test class with its test functions defined inline, then a brief `main` function to register those functions with test names and run them with any command line options.

3.1 The Curiously Recurring Template Pattern (CRTP)

Having said that FRUCTOSE is smaller and simpler than other frameworks, I have to confess that when one writes a test class that is to be used with FRUCTOSE, the test class needs to inherit from a FRUCTOSE base class. Not all test frameworks require inheritance to be used but FRUCTOSE does. Also the style of inheritance employs a pattern that some people may not have seen before. It is known as the Curiously Recurring Template Pattern (CRTP) (ref [14]). CRTP is where one inherits from a template base class whose template parameter is the derived class. This section explains why.

There is a need for machinery that can add tests to a test suite and run the tests. This is all done by the test class inheriting from the FRUCTOSE base class, `test_base`. This base class provides, amongst other things, the function `add_test`, which takes the name of a test and a function that runs that test.

The functions that comprise the tests are members of the test class. So

`test_base` needs to define a function that takes a member function pointer for the test class. CRTP is used so that the base class can specify a function signature that uses the derived class.

Suppose our test class is called `simpletest`. Its declaration would start like this:

```
struct simpletest :
    public fructose::test_base<simpletest> {
    :
    :
```

Let's see how this works: `test_base` names its template parameter `test_container` (it is the class that contains the tests). `test_base` declares the typedef `test_case` expressed in terms of `test_container`:

```
typedef void (test_container::*test_case)
            (const std::string&);
```

This enables it to declare `add_test` to take a parameter of type `test_case`. `test_base` maintains a map of test case function pointers, keyed by test name. The declaration for this map is:

```
std::map<std::string, test_case> m_tests;
```

The declaration of the `add_test` function is:

```
void add_test(const std::string& name,
             test_case the_test);
```

Here is an example of a complete test harness, showing how the use of CRTP means the test cases are defined in the test class and registered using `add_test`:

```
#include "fructose/test_base.h"
const int neighbour_of_the_beast = 668;
struct simpletest :
    public fructose::test_base<simpletest> {
    void beast(const std::string& test_name) {
        fructose_assert(neighbour_of_the_beast == 668)
    }
};

int main(int argc, char* argv[]) {
    simpletest tests;
    tests.add_test("beast", &simpletest::beast);
    return tests.run();
}
```

3.2 Named tests and command line options

The example in the previous section shows that tests are named when they are registered but the example does not actually make any practical use of this. Using the `run()` function above, all registered tests are run. FRUCTOSE does provide a way to select which tests are run via command line options. Basically, the names of the tests are given on the command line. This is done by using the overloaded function `int run(int argc, char* argv[]);`.

The Open Source library TCLAP is used to parse the command line (ref [9]). The command line is considered to consist of optional parameters followed by an optional list of test names. During parsing, TCLAP sets various private data members according to the flags seen on the command line. These flags are available via accessors such as `verbose()`. This is another reason why the test class has to inherit from a base class. It provides access to these flags.

3.3 To inherit or not to inherit

The authors of some test frameworks feel that the requirement for a test class to have to inherit from anything is unreasonable. It is said that such a requirement makes the test framework hard to use and causes undesirable coupling between the test class and the framework. FRUCTOSE has several things to say in response to this:

- The most common kind of unit test framework is known as xUnit. This offers several facilities that are usually implemented via inheritance. For example, test cases and test fixtures are usually offered as base classes that the test program is expected to inherit from.
- Some test frameworks that employ inheritance may be hard to use (e.g. CppUnit) but it does not follow that inheritance is the cause. Some frameworks have just become large and complex during their evolution, and offer the developer a bewildering number of choices for the design of their test classes.
- CRTP can seem slightly daunting to those that have not seen it before. However, the use of CRTP by FRUCTOSE is quite simple and is there simply to enforce that the code that does the tests is in functions of the test class. If anyone knows of any other way to enforce this I would be most interested to hear from them.
- FRUCTOSE applications are intended to be run as command line programs with the ability to use various command line options that come as part of FRUCTOSE. The test class gets this capability by inheritance. If anyone knows of a better way to do this I would be most interested to hear from them.

- FRUCTOSE only makes a handful of functions available. There is `add_test` to register the tests and `run` to run them. The assertion testing macros (discussed later) rely on a function in the base class but that is an implementation detail that is of no concern to the programmer. Other FRUCTOSE functionality such as the command line options accessors is optional. Hence, the amount of coupling between the test class and the FRUCTOSE base class is actually quite low.
- The FRUCTOSE assert macros include the test name in any assertion failure message because the function `get_test_name()` is available to any class that inherits from `test_base`. This is part of what enables FRUCTOSE to have named tests. Without using this technique it is hard to see how user-friendly test names can be registered and used by the framework. This was a difficulty mentioned in the CUTE article. CUTE overcame the problem by using a compiler-specific demangle routine. Other frameworks just don't allow the user to name the tests at all.
- Some of the test frameworks I have examined avoid the need for the test class to inherit from a base class by their assert macros expanding to a large volume of code. The sort of code these macros expand to is the kind that FRUCTOSE places in the base class. FRUCTOSE takes the view that in general macros should be avoided in C++. However, FRUCTOSE does use macros for its asserts. It does this for two reasons: first it needs to get the assertion expression as a complete string much as the C assert facility does; and second, it uses the `__FILE__` and `__LINE__` macros to report the filename and line number at which the assert occurs. The macro definitions are small.
- Because the test class inherits from a base class to get all the functionality required, the writer of the test harness only needs to worry about two classes; the one being tested and the one doing the testing. This is felt to be a great simplification compared to other frameworks.
- It is a normal part of the xUnit architecture (see chapter 6) that the test case class is the one that all unit tests inherit from.

FRUCTOSE actually has two classes, `test_base` and `test_root`. The user sees the former since it must be inherited from but does not see the latter (it is an implementation detail). `test_base` inherits from `test_root`. This is a division of labour; `test_base` contains code that uses the template argument. `test_root` contains code that is not required to be in a template class. The reason for this is largely historical; during the early stages of FRUCTOSE design it came as a library that had to be built, then linked with. `test_root` was in the library whilst the template code was all in the headers and instantiated at compile time. When FRUCTOSE changed to be implemented entirely in header files (inspired by the same approach in TCLAP), the separation of classes was retained. It still provides a distinction between code that is required to be template code and that which does not.

`test_root` contains the following:

- private data members and associated accessors for the flags read from the command line.
- A private data member and associated accessor for the name of the current test.
- A count of the number of assertion errors, plus an associated accessor and mutator.
- Functions that implement most of the assertion code.
- Default `setup` and `teardown` functions.

3.4 The assert macros

The `FRUCTOSE` assert macros are compared with the classic C `assert` and the assert macros of other unit test frameworks.

First, there are a couple of minor style points about the naming of the `FRUCTOSE` macros.

1. `FRUCTOSE` uses the namespace `fructose` to scope all its externally visible symbols. Macros have global scope so they are prepended with `fructose_`. This makes the naming and scoping as consistent as possible. Too many unit test frameworks call their main assert macro `ASSERT`.
2. Generally one chooses uppercase for macros in order to tip the reader off that the token is a macro. However, where the macro is intended to look and feel like a function call macros are sometimes in lowercase. Well known examples include `assert` and `get_char`. Also, `toupper` and `tolower` are sometimes implemented as macros. The `FRUCTOSE` macros are expected to be used in a similar way to the standard C `assert` macros and are designed to look and behave as function calls. Hence they are in lowercase. During the early evolution of `FRUCTOSE` some developers asked for them to be in uppercase so both are provided. This document always uses the lowercase form.

The `fructose_assert` macro expands to a call to a function, `test_assert` which takes the boolean result of the assertion expression, the test name (obtained via the function call `get_test_name()`), the assertion expression as a string, the source filename and line number. Here is the definition:

```
#define fructose_assert(X) \  
    { fructose::test_root::test_assert((X), \  
        get_test_name(), #X, __FILE__, __LINE__);}
```

The `test_assert` function takes no action if the condition evaluates to true. However, when the condition is false it reports the name of the test that failed, the filename and line number and the assertion expression. It also increments the error report (for a final report at the end of all tests) and optionally halts (depends on whether or not the command line option has been specified to halt on first failure).

Note, unlike some other test frameworks, an assertion failure does not abort the function from which it was invoked. This is deliberate. Other test frameworks take the view that when a test fails all bets are off and the safest thing to do is to abort. FRUCTOSE takes a different view. FRUCTOSE assumes it is being used as part of a TDD effort where the developer will typically be developing the class and its test harness at the same time. This means that the developer will want to execute as many tests as possible so he can see which pass and which fail. It also means that the developer needs to be aware that code immediately after a FRUCTOSE assert will still be executed so it should not rely on the previous lines having worked. One style of writing test cases that is in keeping with this is to use a table of test data. Each row in the table contains the input and the expected output. This allows the test code to loop over the table performing lots of tests without a test having to rely on successful execution of the previous test.

Chapter 4

How FRUCTOSE works

The test class provides all the functions that do the testing. It must inherit from `test_base` using the curiously recurring template pattern (CRTP). It does this for several reasons.

1. FRUCTOSE maintains a map of function pointers for when it has to invoke those functions. The function pointer type needs to be declared which means establishing a calling convention. FRUCTOSE takes the view that the function should be a member function of the test class whose return type is void and that takes the test name as a `const std::string` reference. This is enforced at compile time by use of CRTP, which allows `test_base` to declare the function pointer type using the name of the test class.
2. `test_base` makes the function `bool verbose() const` available, which returns `true` if the verbose flag was given on the command line.
3. FRUCTOSE avoids users having to know about several classes by providing the test registration function `add_test` and the test runner function `run`, via the base class. It also provides an overloaded `run` function that parses the command line and runs the tests specified. These functions are not only convenient, they ensure that the user of the framework does not have to worry about the existence of any classes other than the one he is testing and the test class he is testing it with.

It is felt by some that when a framework forces the test class to inherit from anything this is an unreasonable requirement.

The argument says that because inheritance is very strong coupling between classes this arrangement couples the test case too tightly to the framework. The trouble is, there has to be some coupling between the test class and the machinery that invokes its functions. If nothing else, the invoking machinery must establish a call convention for the functions that it calls.

FRUCTOSE uses a convention of the test function having the void return type and taking a `const std::string` reference to the test name. As another

example, CUTE requires that the test function have the void return type and have no function arguments in order that a boost functor may be implicitly created when a test function is passed to the CUTE macro. The CUTE approach does eliminate the coupling by inheritance but does so at the cost of not being able to explicitly name the test. Also, there is no particular advantage in allowing the test functions to be unrelated. In fact one could argue that they should be all grouped in the same test class on the principle of grouping related things together.

FRUCTOSE has the idea that tests are registered by name. The name has to be explicitly given in the `add_test` function. It is not deduced from the names or functions or the use of RTTI. This allows the test to be referred to from the command line. The `test_case` class provides the registration and command line parsing functions.

There are pros and cons to the requirement to explicitly register tests by name.

- Pros
 - Tests can be selected by name when running the test program via the command line.
 - It encourages tests to be grouped into related areas by name.
 - It is clear from the test harness code what the names of the tests are and how they are invoked/registered. Compare this with implicit registration unit test frameworks such as TUT (Ref[18]).
- Cons
 - It is tedious to have to remember to call `add_test` and pick a name once the test method has been written. Forgetting to do this means the test never gets called. Some people prefer implicit test registration for this reason. However, FRUCTOSE does provide a code generation facility that gives some help in this area (see the User Guide).

Chapter 5

FRUCTOSE coding techniques

FRUCTOSE is designed to work in commercial environments as well as open source environments. Commercial environments place some constraints on the C++ coding techniques that can be used. In the commercial world ancient compilers are still very much alive and well (for various reasons). Also, multiple operating systems have to be supported (Microsoft Windows and Solaris are probably the most important). This means that the lowest common denominator approach has to be taken for the dialect of C++ chosen. Platform-dependent RTTI, use of complex template meta-programming, and other advanced C++ techniques were avoided. So were dependencies upon packages that use such techniques.

FRUCTOSE achieves what it needs by simple inheritance. True, it uses CRTP, but the main reason for this is so the function pointers it maintains have to be functions that belong to the test class.

The STL is also used. Usage is kept very simple. Strings are always of type `std::string`. The function pointers to run are held in a `std::map`, keyed by test name string. The list of named tests to run is held in a `std::vector`. The exception handling macros also use the standard exception header `stdexcept` for things such as catching exceptions by the standard base class, `std::exception`. The headers required for these uses of the STL and standard exceptions are automatically included by FRUCTOSE so there is no need for the test harness to include them again. All the functions are inlined, so there is no library to link against; just use the header files.

Chapter 6

FRUCTOSE and the xUnit architecture

6.1 What is the xUnit architecture

According to Wikipedia (ref [16]), all xUnit frameworks have the following basic components in common:

- Test case This is the most elemental class. All unit test are inherited from it.
- Test fixtures A test fixture (also known as a test context) is the set of preconditions or state needed to run a test. The developer should set up a known good state before the tests, and after the tests return to the original state.
- Test suites A test suite is a set of tests that all share the same fixture. The order of the test shouldn't matter.
- Test execution The execution of each test is preceded by a call to setup and followed by a call to teardown. These serve to initialise and cleanup the test fixtures.
- Assertions An assertion is a function or macro that verifies the behavior (or the state) of the unit under test. Failure of an assertion typically throws an exception, aborting the execution of the current test.

6.2 Comparing FRUCTOSE to xUnit

An easy way to assess how closely a C++ unit test framework conforms to the xUnit architecture is to consider how each component of the xUnit architecture is implemented in CppUnit, and compare tha against the framework under consideration.

6.2.1 Test case

In CppUnit one has to subclass the `TestCase` class, override the method `runTest`, and use the assert macros to check values/conditions. The inheritance is done like this:

```
class MyclassTest : public CppUnit::TestCase {
public:
    MyclassTest(std::string name) : CppUnit::TestCase(name) {}
    void runTest()
    {
        :
        :
    }
}
```

The idea is that the `MyclassTest` object is one single test case. To run the test the method `runTest` is used.

FRUCTOSE may seem very similar, since it has a base class, `test_base`, thus:

```
class MyclassTest : public fructose::test_base<MyclassTest> {
public:
    doSomeTestOrOther(const std::string& testName)
    {
        :
        :
    }
    :
    :
}
```

The `TestCase` of CppUnit maps to one test case. In contrast, the `test_base` of FRUCTOSE maps to a class that performs a number of tests, so it is not a test class as such. Rather it is a collection of test cases.

6.2.2 Test fixtures

A test fixture wraps a test case with setup and teardown methods. It is used to provide a common environment for a set of test cases.

In CppUnit, to define a test fixture, do the following:

- implement a subclass of `TestCase`
- the fixture is defined by instance variables
- initialize the fixture state by overriding `setUp` (i.e. construct the instance variables of the fixture)

- clean-up after a test by overriding `tearDown`.

Each test runs in its own fixture so there can be no side effects among test runs.

In CppUnit, `TestCase` and `TestFixture` are explicit classes. `TestCase` inherits from `TestFixture`. A project that uses CppUnit may have test harnesses that employ a mixture of test classes, some inheriting from `TestCase` and others inheriting from `TestFixture`. The former requires that `runTest` is overridden. The latter requires the programmer registers test methods into a `TestSuite` using a `TestCaller`.

In FRUCTOSE, the class that inherits from `test_base` also provides a setup/teardown mechanism that allows each test to run in its own environment. It also allows multiple test to share common setup/teardown code via the ctor/dtor for the class that inherits from `test_base`. So with FRUCTOSE there is no separately identifiable test fixture class. Test cases and test fixtures are combined in the one class. This doesn't mean that FRUCTOSE is missing a test class and test fixtures. The mechanisms are just not as explicit. This is deliberate because it is part of what makes FRUCTOSE smaller and simpler. FRUCTOSE requires that the user registers their test methods with a test suite always, but using a mechanism that is much simpler than that employed by CppUnit.

6.2.3 Test suites

A test suite is an object that runs a collection of test cases. CppUnit has a dedicated class that represents a test suite, `CppUnit::TestSuite`. In FRUCTOSE, the class that inherits from `test_base` is the test suite.

6.2.4 Test execution

The execution of an individual unit test proceeds as follows:

```
setup(); // First, we should prepare our 'world' to make
        //an isolated environment for testing.
...
// Body of test - Here we make all the tests.
...
tearDown(); // In the end, whether succeed or fail we should clean up
           // our 'world' to not disturb other tests or code.
```

The `setup()` and `tearDown()` methods serve to initialize and clean up test fixtures.

Initially, FRUCTOSE did not have this facility but it was added because it proved useful and relatively cheap to add.

6.2.5 Assertions

FRUCTOSE has a rich set of assertions, probably more than most other C++ unit test frameworks.

In conclusion, FRUCTOSE does not conform to the xUnit architecture, mainly on a point of technicality, namely that there is no test case class. Test fixtures and test suites are not as apparent as in other frameworks since in FRUCTOSE the same class is used for both. Hence to casual observer it could appear that FRUCTOSE is missing these xUnit components as well.

Chapter 7

The problems with CppUnit

When I first looked into providing a unit test environment for a commercial project I was working on, I was advised to look at Cppunit, so I did. I found that it would not build on the Solaris development environment we had (Forte 6.0). A port was in progress at the time but we needed something immediately. It was too much work to do a port ourselves when we were supposed to be using something off the shelf. We were also using a non-standard version of the STL and had to ensure that everything we built would build with that STL. The build procedure for CppUnit made this awkward. These problems are very specific to the development environment I was in. I made the recommendation that CppUnit not be used. Some of the reasons I gave were these environment-specific reasons but I also quoted the following reasons, which it seems have been the experience of others:

- Other people also find it hard to build. There is even a manual on the Wiki pages explaining how to build for various platforms. It should not be that complicated!
- CppUnit is very large for a unit test framework. The tarball is over 3MB once uncompressed. Admittedly, quite a bit of this is documentation and examples but no other C++ unit test framework I looked at was anywhere near this size. It takes quite a while to build it too (over four minutes on my dedicated Linux machine).
- It is hard to use. That's why there are lots of tutorials, lots of documentation and even a cookbook, discussion forums and an FAQ.
- CppUnit is too large and complex for many people's needs: I am sure that the reason for this volume of documentation is that CppUnit has alot of functionality to offer. However, in my opinion it does so at the cost of frightening off the developer that only needs something simple. The

number of C++ unit test frameworks that have sprung up, all with the goal of providing something smaller, cutdown and simpler are a testimony to the size and complexity of CppUnit.

- Even in the simple cases, CppUnit places too many complex requirements on the developer, particularly regarding new classes to be written and which base classes they are supposed to inherit from. In most of the alternatives to CppUnit there is only one class to inherit from (in the case of CUTE there is no need even for that). In CppUnit there is a choice of inheriting from `TestCase` or `TestFixture`. There are also `TestRunners`, `TestCallers` and `TestSuites` to worry about.

Chapter 8

The problems with other frameworks

My search of sourceforge revealed several C++ unit test frameworks. However, the problem was usually that it was either for a specific environment or it did not allow the test selection and verbose flag setting via the command line that are so useful when doing TDD. The packages considered include the following:

- `unit--` (ref [1]). The Unit test aid for C++. Judging from the examples and in the opinion of this author, `unit--` is too cut-down, not providing much of the basic functionality provided by the other packages.
- `csUnit` (ref [2]). It is for managed C++.
- Symbian OS C++ Unit Testing Framework (ref [3]). It is only for the Symbian operating system,
- `RapidoTest` (ref [4]). It is for Unix only with particular emphasis on Linux.
- `Mock Objects for C++` (ref [5]). It is a framework that builds on a framework; it provides mock objects by building on either `CppUnit` or `cxxunit`.
- `UnitTest++` (ref [6]). This actually comes quite close. It is much smaller and simpler than `CppUnit` and I did not encounter any build problems. It is multi-platform. However, there is no built-in control over which tests to run, neither can the test be identified during the run. This is fine for overnight regression testing but is not so good for TDD.
- `QuickTest` (ref [7]). This was discovered after `FRUCTOSE` was released. `QuickTest` only consists of one very small header file. There is documentation on the web site but it is not included with the distribution. `QuickTest` has a lot in common but `FRUCTOSE` is slightly richer in functionality, particularly with the test assertions than can be made and the command line and named test features. Again, this makes it more suitable for TDD. `FRUCTOSE` documentation and examples come with the distribution.

- CppTest (ref [8]). A slight wrinkle was found in the build procedure – doxygen is mandatory otherwise the configure script will not produce a Makefile. Apart from that this package looked to be quite good, more mature than QuickTest, also with better documentation. However, in common with QuickTest and UnitTest++ there does not seem to be a mechanism for selecting which tests to run, controlling verbosity and so on.
- Boost test. This was avoided because of the direct dependency on Boost. The test library has to be built in a similar way to Boost, which is well known for having a complex Unix-centric build procedure. Like other frameworks, Boost test has separate classes for tests and the ability to run the tests. It lacks the ability to name the tests and run them selectively.

Chapter 9

The advantages of FRUCTOSE

The main strength of FRUCTOSE compared to the packages above, is its emphasis on its use during code development. This is via its features for reporting in detail the test that failed, the ability to supplement this trace with diagnostics controlled by the verbose flag on the command line, and the ability to select tests by name and optionally fail at the first test failure. None of the other packages provide this; their focus seems to be on running batches of testing in an overnight run to detect regressions. FRUCTOSE will also do that but provides the command line flexibility as well.

Another strength of FRUCTOSE is that it only has one external package dependency. It uses TCLAP (ref [9]) for command line argument handling. This is a very small dependency, since TCLAP is quite small and is implemented entirely in header files. Some other frameworks have a larger set of dependency requirements and some of these dependencies are non-trivial. For example, some depend on Boost. Whilst Boost is recognised to be a fine set of high-quality libraries, some projects, particularly those in some commercial environments, do not want to depend on it. This is for several reasons:

1. Boost does not yet build out of the box in some commercial environments. This is the fault of the compilers, not the fault of Boost. But sometimes a commercial project has little choice of which compiler to use. This can be dictated by company policy, use of other closed-source third-party C++ libraries, and/or customer support obligations where the customer has an old compiler environment.
2. Boost is huge and complex. This turns off many projects/companies from looking at it and using it, even though there are many benefits.
3. If a project is already using Boost (and many are) then a unit test framework that also uses it is not a problem. But given the buildability issues with Boost and its size and complexity, some projects/companies would

be reluctant to be forced to use it just because the unit test framework requires it.

Chapter 10

Possible future work

FRUCTOSE deliberately does not provide any machinery for producing HTML test summaries, reports, or ways of running multiple test suites. Yet anything but the smallest projects will probably want this facility as part of the overnight build and test regime. One way to do this would be for FRUCTOSE to provide scripts, say in perl or python, that used some convention for naming and grouping the test harnesses so they can be run and the results organised into groups.

FRUCTOSE may provide some facility in the future to augment the command line options with additional options that a developer needs for their particular needs. For example, a harness that uses a database may wish to pass in the database parameters (database name, machine name, username, password). At the moment a FRUCTOSE harness would have to find some other way to receive these parameters.

FRUCTOSE does not provide any means to assess code coverage in its tests. There are separate tools to do this. For example, a FRUCTOSE test harness could be run with PureCoverage to assess how much code was exercised. Such tools do not often provide output or reports in way that lend themselves to brief reporting via such things as an overnight build and test run. One possible enhancement of FRUCTOSE would be to develop scripts that work with tools like PureCoverage to give a brief summary of which functions were called and which were not, and what the percentage code coverage was of the functions that were called.

Chapter 11

Conclusion

It is hoped that FRUCTOSE provides enough unit test machinery to enable projects to develop unit tests that can be used both in overnight regression tests and to help TDD. I welcome any feedback on the usefulness(or otherwise) of FRUCTOSE in other projects.

FRUCTOSE's simple implementation and cutdown approach mean that providing anything more complex will probably be permanently outside of the project's scope. However, this does not mean that FRUCTOSE is not open to changes. Any suggestions on how FRUCTOSE can be further simplified without reducing functionality will be gratefully received. It may be downloaded from sourceforge (ref [12]).

Bibliography

- [1] Unit— at <https://sourceforge.net/projects/unitmm>
- [2] csUnit at <https://sourceforge.net/projects/csunit>
- [3] Symbian OS C++ Unit Testing Framework at <https://sourceforge.net/projects/symbianosunit>
- [4] Rapido test at <https://sourceforge.net/projects/rapidotest>
- [5] Mock Objects for C++ at <https://sourceforge.net/projects/mockpp>
- [6] UnitTest++ at <https://sourceforge.net/projects/unittest-cpp>
- [7] QuickTest at <https://sourceforge.net/projects/quicktest>
- [8] CppTest at <https://sourceforge.net/projects/cppptest>
- [9] TCLAP Templatised C++ Command Line Parser Library at <http://tclap.sourceforge.net>.
- [10] The Art of Computer Programming, Volume 2, by Professor Don Knuth.
- [11] Software Engineering; a practioners approach, by Roger Pressman (4th Edition). Section 16.6.2.
- [12] FRUCTOSE at <https://sourceforge.net/projects/fructose>.
- [13] Overload October 2006, CUTE.
- [14] C++ Templates, The Complete Guide, section 16.3, by Vandevoorde and Josuttis.
- [15] The Free Software Foundation, <http://www.fsf.org>
- [16] Wikipedia - xUnit, <http://en.wikipedia.org/wiki/xUnit>
- [17] Hamcrest <http://code.google.com/p/hamcrest>
- [18] C++ Template Unit Test Framework <http://tut-framework.sourceforge.net>