CITS5501 Software Testing and Quality Assurance Test Automation

Unit coordinator: Arran Stewart



- We looked at testing concepts failures, faults/defects, and erroneous states
- We looked at specifications and APIs these help us answer the question, "How do we know what to test against?" i.e., What is the correct behaviour for some piece of software?
- We have discussed what unit tests are, and what they look like

Questions

- How do we come up with tests?
- How do we know when we have enough tests?
- What are typical patterns and techniques when writing tests?
- How do we deal with difficult-to-test software?
 - (e.g. software components with many dependencies)
- What sorts of things can be tested?

Questions

- How do we come up with tests?
- How do we know when we have enough tests?

Both of these are covered in the next few lectures. We look at ways of grouping together different sorts of input so that we don't need to test exhaustively, and at ways of working out how much of the system we have tested (and ought to test).



• What are typical patterns and techniques when writing tests?

We look at data-driven tests (running "the same" test, but on different sets of input and expected output), and property-based tests (testing invariant properties of code or data).



How do we deal with difficult-to-test software?
 (e.g. software components with many *dependencies*)

We saw that unit tests should test things in isolation – what if something is hard to isolate?

(e.g. it uses a database)

We discuss the use of *mocks* to handle this.



• What sorts of things can be tested?

Not *just* the modules in your code!

We can also test examples and code fragments appearing in documents (e.g. user manuals), API documentation, and provided as example programs.

Coming up

- Testing is all about *running* software to see how it behaves.
- *Static* analysis of software consists of any way of inspecting or analysing software (or some other static artifact) *without* running it.
- We will look at:
 - Inspections (analysis by humans)
 - Static analysis and (late in the course) formal methods



• Then we will consider software quality more broadly (looking at processes and standards).

Test cases

Test automation

<ロト < (回)、 < (目)、 < (目)、 < (目)、 < (目)、 < (目)、 < (目)、 < (10)/94

Test cases

Testing frameworks

• In the last lecture, we saw some example unit tests.

- In the last lecture, we saw some example unit tests.
- We said that a unit test tests a *unit* of code (function, method, class, module) in isolation, and there are a few properties we would like it to have (e.g. run very quickly).

- In the last lecture, we saw some example unit tests.
- We said that a unit test tests a *unit* of code (function, method, class, module) in isolation, and there are a few properties we would like it to have (e.g. run very quickly).
- But there are other sorts of tests as well integration tests and system tests – which tend to run more slowly, and contain interacting sub-parts. How do we run those sorts of tests?

- Besides the fact that we can use them for unit tests, JUnit and the other xUnit frameworks are good examples of *testing frameworks*, or *test automation frameworks*.
- Amman and Offut define test automation as:

The use of software to control the execution of tests, the comparison of actual outcomes to predicted outcomes, the setting up of test preconditions, and other test control and test reporting functions.

• And a test framework as:

A set of assumptions, concepts, and tools that support test automation.

We can see that JUnit supports test automation in multiple ways.

• Does it let us control the execution of tests?

- Does it let us control the execution of tests?
 - Yes, we can run tests (or some subset of them), from the command line or an IDE

- Does it let us control the execution of tests?
 - Yes, we can run tests (or some subset of them), from the command line or an IDE
- Does it let us compare actual outcomes to predicated outcomes?

- Does it let us control the execution of tests?
 - Yes, we can run tests (or some subset of them), from the command line or an IDE
- Does it let us compare actual outcomes to predicated outcomes?
 - Yes, we saw that we can use assertions to compare (for instance) what we expect to be returned from a method, with what's actually returned.

- Does it let us control the execution of tests?
 - Yes, we can run tests (or some subset of them), from the command line or an IDE
- Does it let us compare actual outcomes to predicated outcomes?
 - Yes, we saw that we can use assertions to compare (for instance) what we expect to be returned from a method, with what's actually returned.
- Does it let us set up test preconditions?

- Does it let us control the execution of tests?
 - Yes, we can run tests (or some subset of them), from the command line or an IDE
- Does it let us compare actual outcomes to predicated outcomes?
 - Yes, we saw that we can use assertions to compare (for instance) what we expect to be returned from a method, with what's actually returned.
- Does it let us set up test preconditions?
 - Yes, we can write code in the body of a unit test that does this. (And we will see later that we can often pull common code for this out into *test fixtures*.)

(... continued):

- Does it provide other test control and reporting functions?
 - Yes, it provides multiple forms of output, that less us see whether our tests failed or succeeded.

Likewise, Python's unittest provides similar features:

- A structure for writing test drivers
- Assertions for testing expected results
- Test features for sharing common test data
- Test suites for easily organizing and running tests
- Graphical and textual test runners

- But they are just one sort of test automation tool.
- Test automation could be something as simple as a script which, once a day, compiles my team's project and sends an email to me if there are compilation errors.
- In general, if some sort of testing *can* be automated run without human intervention – then we should do so.
- We can imagine for unit tests, that we *could* have a person manually running each test and recording the outcome – but that is slow, subject to human error, and quickly becomes infeasible.

Some sorts of testing and QA activity can't be automated (... yet).

- A final system test for (say) a mobile phone might involve actually setting up and performing tasks with a real phone.
- Assessing the usability of an interface typically, how quickly a user can find and navigate to a particular item, can't be automated.

But what we can automate, we do - this helps reduce cost and improve the reliability of our testing.

The easier tests are to run, the more likely that they *will* be run, and the fewer 'manual handling' steps, the less chance for error.

So, why do test automation?

- Reduces cost
- Reduces human error
- Reduces variance in test quality from different individuals
- Significantly reduces the cost of regression testing

Test cases



< □ > < 圏 > < 直 > < 直 > < 直 > < 直 > ○ Q (~ 26 / 94

- One thing all testing frameworks do is let us run *test cases*.
- We can think of a test case as the smallest possible 'unit' of testing.
- A test case has multiple parts that make it up.
- Terminology is not standardised here, we follow the terminology of Amman and Offutt.

As an example, we'll consider a Python function addNumbers(m, n), which should add two numbers.

And we'll suppose we have a particular test case for it in mind:

• Test case 001: When passed the numbers 3 and 4, addNumbers should return 7.

Components that go to make up a test case:

• Test values or test case values.

These are the *inputs* to the program that we supply. In this case, the numbers 3 and 4.

• Expected results or expected values.

The values we expect to see if the program performs in accordance with specifications. In this case, the number 7.

Definitions:

- Test values: input values necessary to complete some execution of the software.
- Expected values: result to be produced iff program satisfies intended behaviour on a test case.

Sometimes, we will need to supply or execute other values or commands in order to execute a test case.

Prefix values.

Inputs necessary to put the software into the appropriate state to receive the test case values.

For instance, if addNumbers was actually a method of a class called Calculator – then before we can execute our test case, we'll need to construct an object of type Calculator.

The 'prefix values' are all values, commands etc. we supply in order to construct the Calculator object and get it ready.

Prefix values

'Prefix values' can be more complicated than just parameters passed or code executed. They can include all sorts of setup activity ...

• e.g. Suppose we're doing performance testing of a web site we start a server instances going, and send automated requests to it to see how it performs.

Then the 'prefix values' can be everything needed to prepare that test – the configuration for the server, the commands to run it, the configuration of the request-sending program, annd the commands needed to run it, etc.

A further possible component of a test case:

Postfix values:

Any inputs, values, commands etc that need to be sent to the software after the test case values have been sent.

- Verification Values : Values needed to see the results of the test case values
- Exit Commands : Values needed to terminate the program or otherwise return it to a stable state

Verification values

For instance, suppose the specification for addNumbers instead said that, when the function is passed 2 numbers, it should write the results to a file called "myresult.txt".

We run our test, and supply the input values 3 and 4. How can we tell if the test passed or failed?

We will need to send further commands, paramaters etc to find out what whether the file got written, and what was written to it.



If our software made use of, say, a database – then 'exit commands' might include commands needed to restore the database to a known, stable state.

Test cases in real-world code

Sometimes code that is packaged as a single "test" actually contains *multiple* test *cases*.

```
• e.g.:
def myTest():
    for n in range(0, 100):
        m = 3
        res = addNumbers(m, n)
        assertEquals(res, m+n)
```

• How many test cases do we have here?

Test fixtures

 One bit of terminology we saw previously is "test fixture": the commands/values needed to create test fixtures would all be examples of prefix values

Testing framework definitions

A few definitions relevant to testing frameworks:

Software Testability

The degree to which a system or component facilitates the establishment of test criteria and the performance of tests to determine whether those criteria have been met

- how hard it is to find faults in the software
- Testability is determined by two practical problems
 - How to provide the test values to the software
 - How to observe the results of test execution

Testing framework definitions

Observability and Controllability

How easy it is to observe the behavior of a program in terms of its outputs, effects on the environment and other hardware and software components

How easy it is to provide a program with the needed inputs, in terms of values, operations, and behaviors

- Observability
 - Software that affects hardware devices, databases, or remote files have low observability
- Controllability
 - Easy to control software with inputs from keyboards
 - Inputs from hardware sensors or distributed software is harder
- Some systems are very easy to observe and control, others less
 so.

- So, a testing framework is any set of assumptions, tools etc that assist in executing our test cases.
- They can range from the simple (a script that is automatically run, and emails us with a result) to the complex (like the JUnit and unittest frameworks).
- **Unit** testing frameworks are frameworks that just happen to be primarily intended for running unit tests.

They have that goal in mind, and thus tend to provide facilities for testing a unit code in isolation (e.g. the ability to 'mock' databases or other external systems, that we saw last lecture).

• We can actually use JUnit or unittest to test

• We can actually use JUnit or unittest to test

• interaction between several methods

- We can actually use JUnit or unittest to test
 - interaction between several methods
 - interaction between several objects

- We can actually use JUnit or unittest to test
 - interaction between several methods
 - interaction between several objects
 - setting up several systems/subsystems, and testing interaction between them.

- We can actually use JUnit or unittest to test
 - interaction between several methods
 - interaction between several objects
 - setting up several systems/subsystems, and testing interaction between them.
- The framework may not always provide good support for the sort of things we're doing, as it was set up with a different purpose in mind.

- We can actually use JUnit or unittest to test
 - interaction between several methods
 - interaction between several objects
 - setting up several systems/subsystems, and testing interaction between them.
- The framework may not always provide good support for the sort of things we're doing, as it was set up with a different purpose in mind.
- But we can still use its general features (e.g. report output) regardless.

- We can actually use JUnit or unittest to test
 - interaction between several methods
 - interaction between several objects
 - setting up several systems/subsystems, and testing interaction between them.
- The framework may not always provide good support for the sort of things we're doing, as it was set up with a different purpose in mind.
- But we can still use its general features (e.g. report output) regardless.
- We should probably separate these non-unit tests out from our unit tests, though, and document what their purpose is.

unittest

• The standard python module helps you write unit tests: import unittest

from my_script import is_palindrome

class KnownInput(unittest.TestCase):
 knownValues = (('lego', False),('radar', True))

def testKnownValues(self):
 for word, palin in self.knownValues:
 result = is_palindrome(word)
 self.assertEqual(result, palin)

Test fixtures

- Recall that test fixtures are things we need in order to get the system into a known state, ready for a test
- Often, multiple tests will share some requirements for what environment needs to be set up
- A typical approach in object-oriented languages is to group tests with shared fixture requirements into the same class
- And then to specify "setup" and "tear-down" methods for the class, which will be run before and after each test, respectively.
- Shared objects will be declared as *instance variables*

Test fixtures

```
class TestArithmeticOperations {
  Calculator myCalculator;
```

```
@Override // This is run before each test method
protected void setUp() throws Exception {
   System.out.println("Setting_things_up!");
   myCalculator = new Calculator();
}
```

```
@Override // This is run after each test method
protected void tearDown() throws Exception {
   System.out.println("Running_tearDown");
   myCalculator = null;
   assertNull(myCalculator);
}
```

```
@Test
void test1() {
    // ...
```

Fixtures in Python

import unittest

```
class FixturesTest(unittest.TestCase):
    def setUp(self):
        print('In setUp()')
        self.fixture = range(1, 10)
    def tearDown(self):
        print('In_tearDown()')
        del self.fixture
    def test(self):
        print('in_test()')
        self.assertEqual(self.fixture, range(1, 10))
if ___name___ == '___main___':
    unittest.main()
```

Some assertion methods

Common assertions assertTrue(x, msg=None) assertFalse(x, msg=None) assertIsNone(x, msg=None) assertIsNotNone(x, msg=None) assertEqual(a, b, msg=None) assertNotEqual(a, b, msg=None) assertIs(a, b, msg=None) assertIsNot(a, b, msg=None) assertIn(a, b, msg=None) assertNotIn(a, b, msg=None) assertIsInstance(a, b, msg=None) assertNotIsInstance(a, b, msg=None)

More assertion methods

Other assertions assertAlmostEqual(a, b, places=7, msg=None, delta=None) assertNotAlmostEqual(a, b, places=7, msg=None, delta=None) assertGreater(a, b, msg=None) assertGreaterEqual(a, b, msg=None) assertLess(a, b, msg=None) assertLessEqual(a, b, msg=None) assertRegex(text, regexp, msg=None) assertNotRegex(text, regexp, msg=None) assertCountEqual(a, b, msg=None) assertMultiLineEqual(a, b, msg=None) assertSequenceEqual(a, b, msg=None) assertListEqual(a, b, msg=None) assertTupleEqual(a, b, msg=None) assertDictEqual(a, b, msg=None)

イロト (雪) (日) (日) (日) (日)

Running Tests

• Given the script test_simple.py:

import unittest

class SimplisticTest(unittest.TestCase):
 def test(self):
 self.assertTrue(True)

if __name__ == '__main__': unittest.main()

Running Tests

- Run it with python3 test_simple.py:
 - \$ python3 test_simple.py

Ran 1 test **in** 0.000s

0K

.

Structuring test code

 As with any software system, we want to factor out common code – an example:

```
knownValues = (('lego', False),('radar', True)
```

...

```
for word, palin in self.knownValues:
    result = is_palindrome(word)
    self.assertEqual(result, palin)
```

- Follow the "DRY" principle Do not Repeat Yourself
- Question: what constitutes a "test case", in this code?
- This style of test is sometimes called a "data-driven unit test"

Data-driven unit tests

- Problem: Testing a function multiple times with similar values
 - How to avoid test code bloat?
- Simple example: Adding two numbers
 - Adding a given pair of numbers is just like adding any other pair
 - You really only want to write one test
- Data-driven unit tests call constructor for each logical set of data values
 - Same tests are then run on each set of data values

Structuring test code

- More broadly, how test cases are structured will depend somewhat on the conventions of the language and the framework being used.
 - in Java, typical to put source code in a directory called "src", and have a separate directory (e.g "test") for unit tests, with structure mirroring the main code.
 - in Python, most tests are put into a separate module.

Doubles

- Actors use doubles to replace them during certain scenes
 - Dangerous or athletic scenes
 - Skills the actor doesn't have, like dancing or singing
- Test doubles replace software components that cannot be used during testing

Reasons for Test Doubles

- Component has not been written
- The real component does something destructive that we want to avoid during testing (unrecoverable actions)
- The real component interacts with an unreliable resource
- The real component runs very slowly
- The real component creates a test cycle
 - A depends on B, B depends on C, C depends on A

A test double is a software component that implements partial functionality to be used during testing

Dependencies

• Very often, a class or function is not designed to work on its own, but in combination with other classes or functions -

e.g. an AddressBook class may make use of a Contact class

- or with other subsystems, or external systems:
 - dependency on a database for an HR system
 - dependency on a network, for an Internet chat system
 - dependency on particular hardware devices
- How do we deal with these?

Mocks, stubs and more

- Often, we'll use objects or function that mimic other ones for testing purposes. There does not seem to be any universally accepted term for these, but one author [Gerard Meszaros] uses the generic term "Test Double".
- Specific sorts of Test Double -
 - Dummy objects
 - Fake objects
 - Stubs
 - Spies
 - Mocks

[Fowler, in e.g. "Mocks Aren't Stubs", uses Meszaros's terminology.]

Dummy objects

- These are objects that are passed around but not used for instance, they may be used to fill parameter lists (in statically typed languages).
 - In languages with a null, Nil or undefined value, we might be able to use that value

(which also serves to document the fact that we don't care what it is)

Fake objects

- *Fake* objects actually do have working implementations, but for some reason are not suitable for production
 - An example of this is when we use an in-memory database, instead of an on-disk database



 Stubs (often, "stub methods") provide canned answers to calls made during the test –
i.e., the answers are usually fixed, and don't change in response to the parameters passed



- These are stubs that *record information* on how they were called.
- These are particularly useful for testing code that calls (e.g.) an object representing a server, such as a mail server, or which writes to a file-like object.

Spies – example

- In Java, we often write to files (or network sockets) using classes like BufferedWriter
- If we want to verify, in some unit test, what is written, we could use a "Spy" class that implements the java.io.Writer class but instead of writing to a file, it records whatever data would have been written
- In Python, we do not have static types, and any class with a "write()" method suffices.
- Making our code agnostic about what sort of thing it is writing to has the benefit that if we *do* decide to change it at a later date, we don't have to revise our tests



• *Mock* object are pre-programmed to expect particular calls, and respond with particular behaviour.



- *Mock* object are pre-programmed to expect particular calls, and respond with particular behaviour.
- Unlike the other types of test double, mock objects can verify things about the *behaviour* of a class.



- *Mock* object are pre-programmed to expect particular calls, and respond with particular behaviour.
- Unlike the other types of test double, mock objects can verify things about the *behaviour* of a class.
- For instance:



- *Mock* object are pre-programmed to expect particular calls, and respond with particular behaviour.
- Unlike the other types of test double, mock objects can verify things about the *behaviour* of a class.
- For instance:
 - Suppose our code uses a database; we know that to work correctly, it must call the connect() method of a database object, and can then call the query() method; but it is an error to call query() before connect().



- *Mock* object are pre-programmed to expect particular calls, and respond with particular behaviour.
- Unlike the other types of test double, mock objects can verify things about the *behaviour* of a class.
- For instance:
 - Suppose our code uses a database; we know that to work correctly, it must call the connect() method of a database object, and can then call the query() method; but it is an error to call query() before connect().
 - Our mock object can contain code that checks whether query() has been called before connect().

Mocks – another example

- We might have a order fulfilment system that is supposed to send an email when (for some reason) an order can't be fulfilled.
- The class that handles sending emails may need particular methods to be called, in a particular order; we can write a *mock* that tests that they are called in the right way.

Mocks in Python

- Python has the standard library unittest.mock
- MagicMock() lets us create methods that return specific results, or expect to be called a particular way, on the fly.
- > from unittest.mock import *
- > mock = MagicMock()

Mocks in Python (2)

- Once we have called our mock() object, the fact that it has been called is recorded.
- We then (before the test ends) *assert* what we expect to have happened
 - (e.g. that the method was called)
- If not, then an exception will be raised.
- Much more complex behaviour can be created check the API for details.

Testable documentation

- We have said that sometimes, tests are the best documentation of an API (since documentation often gets out of date)
- *Testable documentation* frameworks ensure that documentation is kept up to date with code tests are generated from the documentation of an API.
- One example, from the Python language, is the doctest library.
- A good API will often give *examples* of how methods are functions should be called, and the Python doctest module allows these examples to be extracted and run as tests.

Testable documentation vs unit testing

- The purpose of these is to ensure that the *documentation examples are still correct*.
- This is not the same as unit testing doctests will usually only exercise a small number of examples, and are not nearly as thorough as unit tests should be.

Doctest example

```
def square(x):
    """Return the square of x.
    >>> square(2)
    4
    >>> square(-2)
    4
    .....
    return x * x
if ___name__ == '___main___':
    import doctest
    doctest.testmod()
```

 Like xUnit, doctest has been ported to a great many other languages.
 (An encouraging feature of testing techniques is that they tend to be widely adopted if they work well.)

• Like xUnit, doctest has been ported to a great many other languages.

(An encouraging feature of testing techniques is that they tend to be widely adopted if they work well.)

Java has JDoctest

• Like xUnit, doctest has been ported to a great many other languages.

(An encouraging feature of testing techniques is that they tend to be widely adopted if they work well.)

- Java has JDoctest
- Haskell has a package simply called doctest

• Like xUnit, doctest has been ported to a great many other languages.

(An encouraging feature of testing techniques is that they tend to be widely adopted if they work well.)

- Java has JDoctest
- Haskell has a package simply called doctest
- Ruby has rdoctest

Property-based testing

• This sort of testing originates from the Haskell testing framework QuickCheck, and is sometimes called *generative testing*

Property-based testing

- This sort of testing originates from the Haskell testing framework QuickCheck, and is sometimes called *generative testing*
- Our tests are of the form:

for all data or parameters that are generated in a particular way, the function or method should produce the following results.

Simple example

 The "tail" function, applied to a list, returns everything but the first element – what invariants hold?

Simple example

- The "tail" function, applied to a list, returns everything but the first element –
 - what invariants hold?
- We know that if tail is called on a non-empty list, the length of the result is one less than the length of the list passed in.

Use for interfaces and sub-classes

• This can be particularly useful when testing interfaces and subclasses

Use for interfaces and sub-classes

- This can be particularly useful when testing interfaces and subclasses
- Our documentation states that all subclasses of a class should maintain some invariant; the property-based test checks whether it can find counterexamples.

Summary of test types so far

Unit tests provide a way of identifying ways in which a software component deviates from its specification.

Test doubles provide a way of testing a unit of code, even when it depends on other code.

Testable documentation provides a way of testing examples written in the *documentation* for a system, and making sure they still hold.

Property-based testing provides a way of finding counterexamples to any *invariants* we think should hold about a software component.

A wrinkle – user expectations

- In some cases, software may perform according to its specification, but still violate user expectations.
- For instance, users may expect a GUI system or mobile app to conform with the behaviour of familiar applications or may expect that a system will *not* do something (e.g., transmit their data to a third party)
- These are not *faults*, per se but they can be just as important for software quality.

Testing as a way of improving reliability

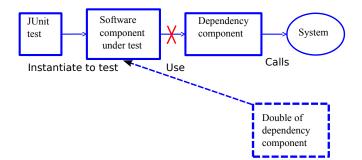
- Testing is one way of improving the reliability of a system. It aims to detect faults that have already found their way into a system.
- In general, techniques for improving reliability fall into three categories:
 - Fault avoidance try to prevent faults from ever being introduced into the system
 - Fault detection try to *detect* faults that *have* found their way into the system
 - Fault tolerance incorporate ways of recovering from faults in the system at runtime.

Examples of improving reliability

- Fault avoidance we can try to avoid introducing faults by our use of particular development methodologies, by statically analysing the system design, and through the use of formal methods.
- Fault detection we can try to detect *failures*, and use debugging and testing to identify the causes (the faults) that result in those failures.
- Fault tolerance we can introduce redundancy into the system.
 For instance, the Airbus flight control system actually contains multiple systems, and control switches to a backup if one becomes unavailable.

(Query – what sort of faults will this guard against? What sort might it not?)

Test Double Illustration





- Next question what test values to use, what test cases to write?
- This is test *design* ...