

CEN 5016: Software Engineering

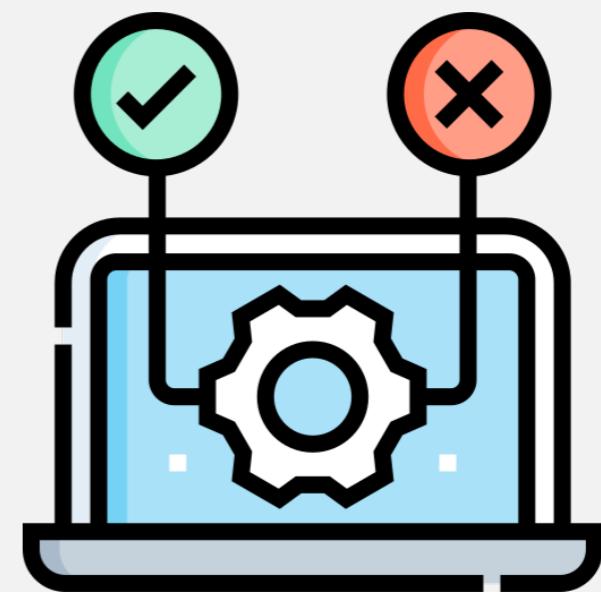
Spring 2026



University of
Central Florida

Dr. Kevin Moran

Week 4- Class I: Software Testing





- *Assignment 3 Posted*
 - Getting familiar with CI/CD and Deployment
 - All parts of the Assignment due Fri Feb 5th @ 11:59pm!
 - Get started today!!!
- *SDE Project Checkpoint 1 posted!*
 - We will go over this today.

Software Testing



Learning Goals



- Identify the scope and limitations of software testing
- Appreciate software testing as a methodology to use automation in improving software quality
- Describe the benefits of using continuous integration and deployment (CI/CD)
- Measure the quality of software tests and define test adequacy criteria
- Enumerate different levels of testing such as unit testing, integration testing, system testing, and testing in production
- Describe the principles of test-driven development
- Outline design principles for writing good tests
- Recognize and avoid testing anti-patterns

What is Testing Good For?



- What is testing?
 - Execution of code on sample inputs in a controlled environment
- Principle goals:
 - Validation: program meets requirements, including quality attributes.
 - Defect testing: reveal failures.

What is Testing Good For?



- Why should we test? What does testing achieve?
 - What does testing not achieve?
- When should we test?
 - And where should we run the tests?
- What should we test?
 - What CAN we test? (Software quality attributes)
- How should we test?
 - How many ways can you test the sort() function?
- How good are our tests?
 - How to measure test quality?

What Makes a Good Test?



What Makes a Good Test?



OpenBudgeteer

Manage your Budget with Buckets

Docker Image (pre-release) passing Docker Image (latest) passing

mentioned in awesome docker pulls 113k release v1.7

<https://github.com/TheAxelander/OpenBudgeteer>

What Makes a Good Test?



OpenBudgeteer [Home](#) [Account](#) [Transaction](#) [Bucket](#) [Import](#) [Report](#) [Rules](#) Database: openbudgeteer-qa Version: 1.3 (Change Log)

Icon	Income	Expenses	Icon	Month Balance	Icon	Bank Balance	
	3.761,17 €		-9.111,34 €		-5.350,17 €		20.793,14 €
Icon	Budget	Icon	Pending Want	Icon	Remaining Budget	Icon	Negative Bucket Balance
	484,89 €		0,00 €		484,89 €		0,00 €

[Create Bucket Group](#) [Distribute Budget](#) [Collapse All](#) [Expend All](#)

Bucket

Bucket	Balance	InOut	Want	In	Activity	Details
▼ Monthly Expenses						
Bucket 1	82,42	0,00	0	0,99	-0,99	
Bucket 2		0,00	0	23,44	-23,44	
Bucket 3		79,43	0	1.140,78	-1.114,59	
Bucket 4 (Inactive from: 01.04.2020)		0,00	0	754,17	-754,17	
▼ Non-Monthly Expenses						
Bucket 5	600,92	44,76	0	7,46	50%	89,50 until 2020-08
Bucket 6		48,47	0	45,59	9%	550,00 until 2021-01
Bucket 7		11,00	0		100%	11,00 until 2020-04
Bucket 8		0,00	0	17,50	-52,50	52,50 until 2020-02
▼ Living Expenses						
Car	336,57	132,91	0	50,00		
Groceries		203,66	0	250,00	-199,67	
Others		0,00	0			
Public Transport		0,00	0			
Treatment Expenses		0,00	0			

< Feb > 2020

<https://github.com/TheAxelander/OpenBudgeteer>

Why Write Tests at All?



- [Low bar] Ensure that our software meets requirements, is correct, etc.
- Preventing bugs or quality degradations from being accidentally introduced in the future -> *Regression Testing*
- Helps uncover unexpected behaviors that can't be identified by reading source code
- Increased confidence in changes ("will I break the internet with this commit?")
- Bridges the gap between a declarative view of the system (i.e., requirements) and an imperative view (i.e., implementation) by means of redundancy.
- Tests are executable documentation; increases code maintainability
- Forces writing testable code <-> checks software design

Testing Levels



- **Unit testing**
 - Code level, E.g. is a function implemented correctly?
 - Does not require setting up a complex environment
- **Integration testing**
 - Do components interact correctly? E.g. a feature that cuts across client and server.
 - Usually requires some environment setup, but can abstract/mock out other components that are not being tested (e.g. network)
- **System testing**
 - Validating the whole system end-to-end (E2E)
 - Requires complete deployment in a staging area, but fake data
- **Testing in production**
 - Real data but more risks

What are the Limitations of Testing?



- *"Testing shows the presence, not the absence of bugs."* - Edsger W. Dijkstra
- Testing doesn't really give any formal assurances
- Writing tests is hard, time consuming
- Knowing if your tests are good enough is not obvious
- Executing tests can be expensive, especially as software complexity and configuration space grows
 - Full test suite for a single large app can take several days to run

What can We Test for?





- “Oracles” are mechanisms that tell you when program execution seems abnormal or unexpected
- E.g. assert, segfault, exception
- Other examples: performance threshold, memory footprint, address sanitizer



- Obvious in some applications (e.g. “sort()”) but more challenging in others (e.g. “encrypt()” or UI-based tests)
- Lack of good oracles can limit the scalability of testing.
Easy to generate lots of input data, but not easy to validate if output (or other program behavior) is correct.
- Fortunately, we have some tricks.

Differential Testing



- If you have two implementations of the same specification, then their output should match on all inputs.
 - E.g. `mergeSort(x).equals(bubbleSort(x))` -> should always be true
 - Special case of a property test, with a free oracle.
- If a differential test fails, at least one of the two implementations is wrong.
 - But which one?
 - If you have $N > 2$ implementations, run them all and compare. Majority wins (the odd one out is buggy).
- Differential testing works well when testing programs that implement standard specifications such as compilers, browsers, SQL engines, XML/JSON parsers, media players, etc.
 - Not feasible in general e.g. for UCF's custom grad application system.

Regression Testing



- Differential testing through time (or versions, say V1 and V2).
- Assuming V1 and V2 don't add a new feature or fix a known bug, then $f(x)$ in V1 should give the same result as $f(x)$ in V2.
- *Key Idea:* Assume the current version is correct. Run program on current version and log output. Compare all future versions to that output.

When Should We Test?



Test Driven Development



- Tests first!
- Popular agile technique
- Write tests as specifications before code
- Never write code without a failing test
- **Claims:**
 - Design approach toward testable design
 - Think about interfaces first
 - Avoid unneeded code
 - Higher product quality
 - Higher test suite quality
 - Higher overall productivity

Common Bar for Contributions



Chromium

- **Changes should include corresponding tests.** Automated testing is at the heart of how we move forward as a project. All changes should include corresponding tests so we can ensure that there is good coverage for code and that future changes will be less likely to regress functionality. Protect your code with tests!

Firefox

Testing Policy

Everything that lands in mozilla-central includes automated tests by default. Every commit has tests that cover every major piece of functionality and expected input conditions.

Docker Conventions

Fork the repo and make changes on your fork in a feature branch:

- If it's a bugfix branch, name it XXX-something where XXX is the number of the issue
- If it's a feature branch, create an enhancement issue to announce your intentions, and name it XXX-something where XXX is the number of the issue.

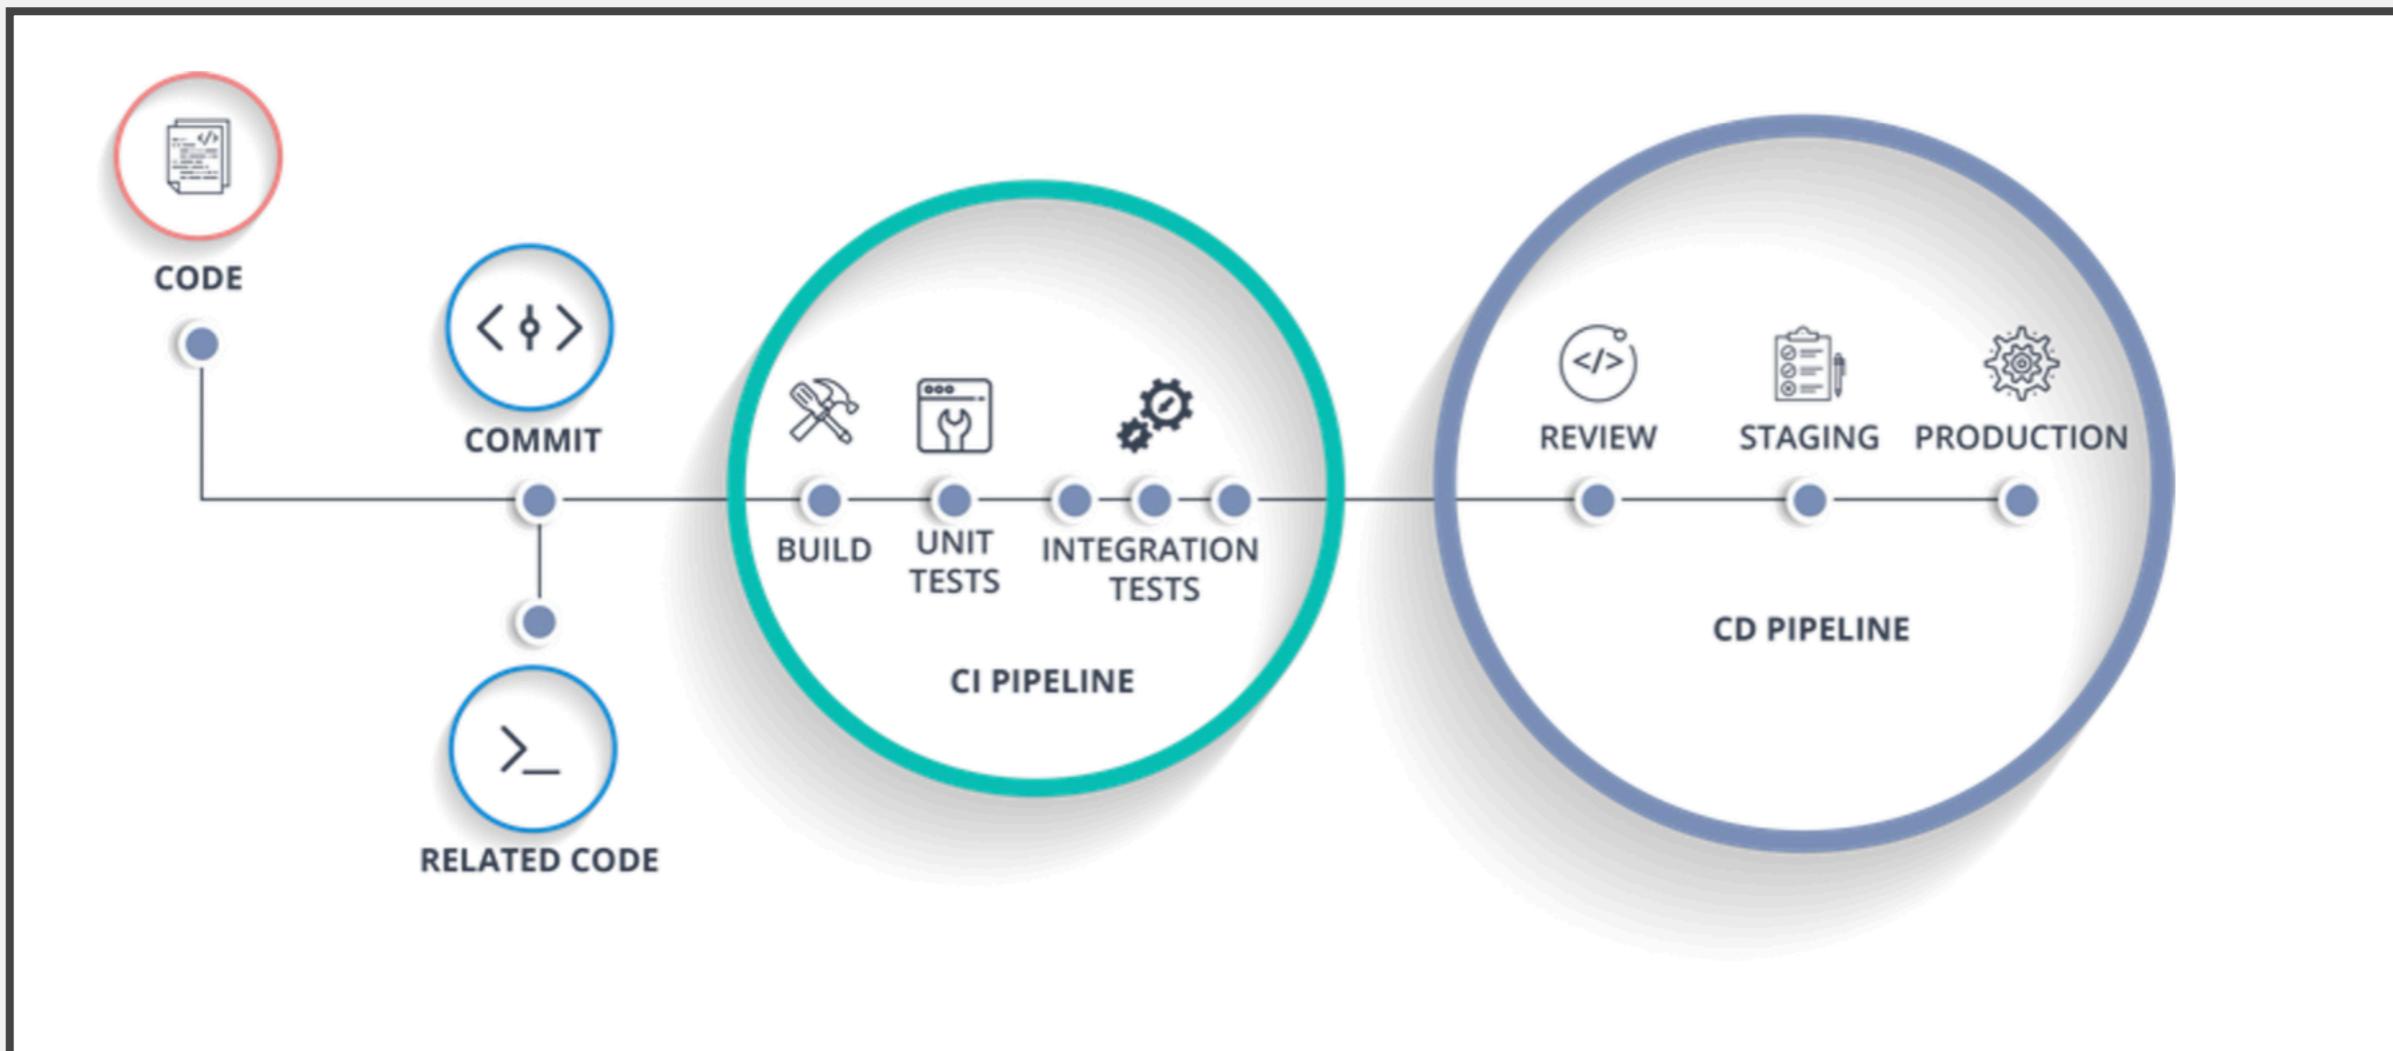
Submit unit tests for your changes. Go has a great test framework built in; use it! Take a look at existing tests for inspiration. Run the full test suite on your branch before submitting a pull request.

Regression Testing



- Usual model:
 - Introduce regression tests for bug fixes, etc.
 - Compare results as code evolves
 - $\text{Code1} + \text{TestSet} \rightarrow \text{TestResults1}$
 - $\text{Code2} + \text{TestSet} \rightarrow \text{TestResults2}$
 - As code evolves, compare TestResults1 with TestResults2 , etc.
- Benefits:
 - Ensure bug fixes remain in place and bugs do not reappear.
 - Reduces reliance on specifications, as $\langle \text{TestSet}, \text{TestResults1} \rangle$ acts as one.

Continuous Integration & Deployment



How Good Are Our Tests?



Code Coverage



- Line coverage
 - Statement coverage
 - Branch coverage
 - Instruction coverage
 - Basic-block coverage
 - Edge coverage
 - Path coverage
 - ...

Code Coverage



LCOV - code coverage report

Current view: [top level](#) - test

Test: coverage.info

Date: 2018-02-07 13:06:43

	Hit	Total	Coverage
Lines:	6092	7293	83.5 %
Functions:	481	518	92.9 %

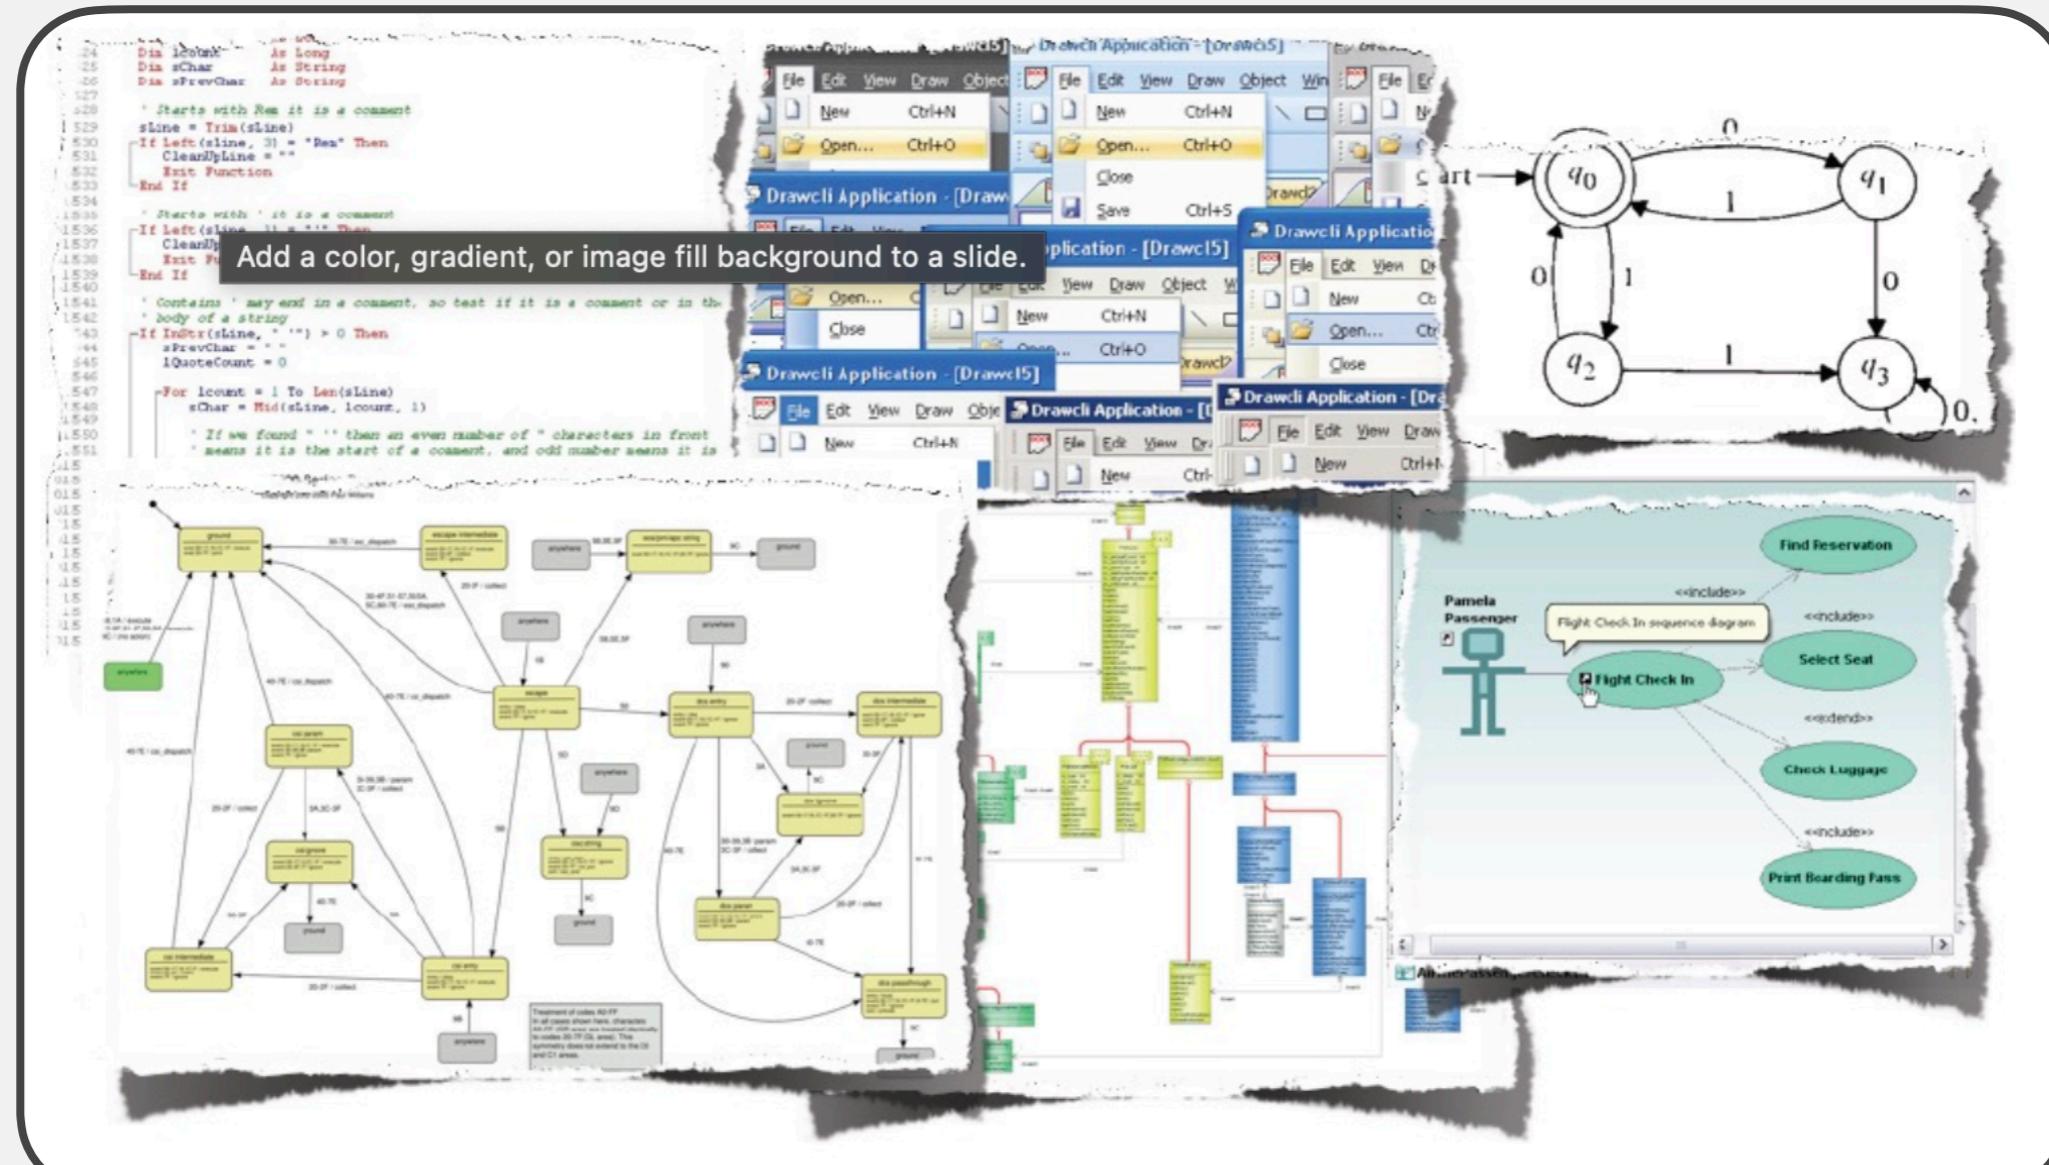
Filename	Line Coverage	Functions
asn1_string_table_test.c	58.8 %	20/34
asn1_time_test.c	72.0 %	72/100
bad_dtls_test.c	97.6 %	163/167
bftest.c	65.3 %	64/98
bio_enc_test.c	78.7 %	74/94
bntest.c	97.7 %	1038/1062
chache_internal_test.c	83.3 %	10/12
ciphername_test.c	60.4 %	32/53
crltest.c	100.0 %	90/90
ct_test.c	95.5 %	212/222
d2i_test.c	72.9 %	35/48
dancetest.c	75.5 %	123/163
dhatest.c	84.6 %	88/104
drgtest.c	69.8 %	157/225
dtls_mtu_test.c	86.8 %	59/68
dtltest.c	97.1 %	34/35
dtlsv1listtest.c	94.9 %	37/39
ecdsatest.c	94.0 %	140/149
enginetest.c	92.8 %	141/152
evp_extra_test.c	100.0 %	112/112
fatalerrtest.c	89.3 %	25/28
handshake_helper.c	84.7 %	494/583
hmacetest.c	100.0 %	71/71
ideatest.c	100.0 %	30/30
igetest.c	87.9 %	109/124
lhash_test.c	78.6 %	66/84
mdc2_internal_test.c	81.8 %	9/11
mdc2test.c	100.0 %	18/18
ocspapitest.c	95.5 %	64/67
packettest.c	100.0 %	248/248

```

97 1 / 1:    if (err = SSLHashSHA1.update(&hashCtx, &clientRandom)) != 0)
98 0 / 1:        goto fail;
99 1 / 1:    :
100 1 / 1:    :
101 1 / 1:    else {
102 0 / 1:        /* DSA, ECDSA - just use the SHA1 hash */
103 0 / 1:        dataToSign = &hashes[SSL_MD5_DIGEST_LEN];
104 1 / 1:        dataToSignLen = SSL_SHA1_DIGEST_LEN;
105 1 / 1:    :
106 1 / 1:    hashOut.data = hashes + SSL_MD5_DIGEST_LEN;
107 1 / 1:    hashOut.length = SSL_SHA1_DIGEST_LEN;
108 1 / 1:    if ((err = SSLFreeBuffer(&hashCtx)) != 0)
109 0 / 1:        goto fail;
110 1 / 1:    :
111 1 / 1:    if ((err = ReadyHash(&SSLHashSHA1, &hashCtx)) != 0)
112 0 / 1:        goto fail;
113 1 / 1:    if ((err = SSLHashSHA1.update(&hashCtx, &clientRandom)) != 0)
114 0 / 1:        goto fail;
115 1 / 1:    if ((err = SSLHashSHA1.update(&hashCtx, &serverRandom)) != 0)
116 0 / 1:        goto fail;
117 1 / 1:    if ((err = SSLHashSHA1.update(&hashCtx, &signedParams)) != 0)
118 0 / 1:        goto fail;
119 1 / 1:    if ((err = SSLHashSHA1.final(&hashCtx, &hashOut)) != 0)
120 1 / 1:        goto fail;
121 1 / 1:    :
122 1 / 1:    err = sslRawVerify(ctx,
123 1 / 1:        ctx->peerPubKey,
124 1 / 1:        dataToSign,
125 1 / 1:        dataToSignLen,
126 1 / 1:        signature,
127 1 / 1:        signatureLen);
128 1 / 1:    if (err) {
129 1 / 1:        sslErrorLog("SSLDecodeSignedServerKeyExchange: sslRawVerify "
130 1 / 1:                    "returned %d\n", (int)err);
131 1 / 1:        goto fail;
132 1 / 1:    }
133 1 / 1:    :
134 1 / 1:    fail:
135 1 / 1:    :
136 1 / 1:    SSLFreeBuffer(&signedHashes);
137 1 / 1:    SSLFreeBuffer(&hashCtx);
138 1 / 1:    return err;
139 1 / 1:    :
140 1 / 1:    :
141 1 / 1:    :

```

We Can Measure Coverage on Almost Anything



Be Aware of Coverage Chasing



- Recall: issues with metrics and incentives
 - Also: Numbers can be deceptive
- 100% coverage != exhaustively tested
 - “Coverage is not strongly correlated with suite effectiveness”
- Based on empirical study on GitHub projects
[Inozemtseva and Holmes, ICSE’14]
- Still, it’s a good low bar
 - Code that is not executed has definitely not been tested

Coverage of What?



- Distinguish code being tested and code being executed
- Library code >>> Application code
 - Can selectively measure coverage
- All application code >>> code being tested
- Not always easy to do this within an application

Coverage != Outcome



- What's better, tests that always pass or tests that always fail?
- Tests should ideally be falsifiable. Boundary determines
- specification
- Ideally:
 - Correct implementations should pass all tests
 - Buggy code should fail at least one test
 - Intuition behind mutation testing (we'll revisit this next week)
- What if tests have bugs?
 - Pass on buggy code or fail on correct code
- Even worse: flaky tests
 - Pass or fail on the same test case nondeterministically
- What's the worst type of test?

Test Design Principles



- Use public APIs only
- Clearly distinguish inputs, configuration, execution, and oracle
- Be simple; avoid complex control flow such as conditionals and loops
- Tests shouldn't need to be frequently changed or refactored
 - Definitely not as frequently as the code being tested changes



- Snoopy oracles
 - Relying on implementation state instead of observable behavior
 - E.g. Checking variables or fields instead of return values
- Brittle tests
 - Overfitting to special-case behavior instead of general principle
 - E.g. hard-coding message strings instead of behavior
- Slow tests
 - Self-explanatory(beware of heavy environments, I/O, and sleep())
- Flaky tests
 - Tests that pass or fail nondeterministically
 - Often because of reliance on random inputs, timing (e.g. sleep(1000)), availability of external services (e.g. fetching data over the network in a unit test), or dependency on order of test execution (e.g. previous test sets up global variables in certain way)

Takeaways



- Most tests that you will write will be muuuuuuch more complex than testing a sort function.
- Need to set up environment, create objects whose methods to test, create objects for test data, get all these into an interesting state, test multiple APIs with varying arguments, etc.
- Many tests will require mocks (i.e., faking a resource-intensive component).
- General principles of many of these strategies still apply:
 - Writing tests can be time consuming
 - Determining test adequacy can be hard (if not impossible)
 - Test oracles are not easy
 - Advanced test strategies have trade-offs (high costs with high returns)