FITE: Future Integrated Testing Environment

Patrice Godefroid[1], Leonardo Mariani[2], Andrea Polini[3],
Nikolai Tillmann[1], Willem Visser[4], Michael W. Whalen[5]

[1] Microsoft Research
Redmond, WA, USA
{pg,nikolait}@microsoft.com

[2] Dipartimento di Informatica, Sistemistica e Comunicazione
Università degli Studi di Milano Bicocca
Viale Sarca, 336 – Milano – ITALY
mariani@disco.unimib.it

[3] Computer Science Division
School of Science and Technologies
Università degli Studi di Camerino
Via Madonna delle Carceri, 9 – Camerino (MC) – ITALY
andrea.polini@unicam.it

[4] Department of Mathematical Sciences
Computer Science Division
University of Stellenbosch
7602 Matieland – SOUTH AFRICA
willem@gmail.com

[5] University of Minnesota
Software Engineering Center
Minneapolis, MN – USA
whalen@cs.umn.edu

Dagstuhl Seminar 10111
Practical Software Testing: Tool Automation and Human Factors

1 Motivation and Background

It is a well known fact that the later software errors are discovered during the development process, the more costly they are to repair. Recently, automatic tools based on static and dynamic analysis have become widely used in industry to detect errors, such as null pointer dereferences, array indexing errors, assertion violations, etc. However, these techniques are typically applied late in the development cycle, and thus, the errors detected by such approaches are expensive to repair. Additionally, these techniques can suffer from scalability issues and produce imprecise results since it is applied for the first time at a late stage when the code base is too large. There are also human factor issues that come into play when analysis are run at a late stage, namely that the tools cannot show all possible errors since the volume of possible false errors will overwhelm the user and thus they ignore the results; the inverse also happens, where tool developers suppress too many real errors in an effort to reduce false warnings.

Dagstuhl Seminar Proceedings 10111
Practical Software Testing : Tool Automation and Human Factors
http://drops.dagstuhl.de/opus/volltexte/2010/2619

To address these issues we suggest that code should be continuously analyzed from an early stage of development, preferably as the code is written. This will allow developers to get instant feedback to repair errors as they are introduced, rather than later when it is more expensive. This analysis should also be incremental in nature to allow better scaling. Static analysis tools can produce more errors at an early development stage since the negative impact of false warnings will be mitigated due to the small increment of code being analyzed. Dynamic analysis, in particular testing, can also benefit since it can use the static analysis results (for the code increment) to produce tests to cover potential errors as well as give high code coverage.

2 FITE Vision

If code is to be analyzed as it is written it implies that the analysis should form part of the development environment. We thus propose the Future Integrated Test Environment (FITE): an IDE with a test-centric focus. FITE will continuously test and analyze the increments and will produce recommendations to the user to repair and test the code. To address scaling up from units, FITE will combine incremental analysis with compositional reasoning. Since the tool is based on interaction with a user, human factors will play a large role in the design of FITE. Many different analysis can be integrated into a tool like FITE. In order to not overwhelm the user with too many recommendations from the tool to improve the code and tests, we foresee a pluggable view-based approach, where the user selects the kinds of analysis it wants to perform on the code (such as security, performance, reliability or numerical precision analysis) and the tool produces only recommendations addressing this selection. For example, when selecting performance the tool will only show code paths with high worst-case execution times and for security might only focus on buffer overruns and information leaking.

3 From Dream to Reality

3.1 Compositional Analysis

We believe the key to make our vision a reality is to effectively engineer compositional reasoning and analysis of large programs, in order to bridge the gap from unit analysis to system analysis, and ultimately the gap between developers and testers.

Two key sub-problems need to be addressed:

1. how to decompose large programs into smaller sub-components by identifying *interfaces* where those sub-components can be decoupled.
2. next, how to generate *contracts* at those interfaces, in order to capture *input pre-conditions* and *output post-conditions* in the form of constraints that *may happen* or *must hold*.

We envision a semi-automated process to solve these two problems of *interface extraction* and *contracts generation*. Those contracts are code *annotations* that capture semantic information about possible behaviors of the program.

Initially, in order to bootstrap the process, a fully-automatic static analysis of the program could first infer candidates interfaces on how to decompose the system (e.g., using heuristics based on measuring the "complexity" of those interfaces) and suggest those to the user. Those interfaces could then be associated with pre-computed contracts of two types: *may contracts* inferred by static analysis (such as "input integer variable x may have any value" or "output return value y may be any integer") and *must contracts* inferred by dynamic analysis of executions obtained with existing or automatically generated test cases (such as "input pointer p must be non-NULL" or "output pointer q always points to an allocated struct of type blah").

Despite this fully-automatic default mode, we really envision a *interactive* (semi-automatic) usage of the FITE tool where the user can be continually involved by receiving and providing feedback. Think of it as *pair programming* where FITE is your coding and test buddy who interacts with you as you write code, test it, and explore its possible behaviors. By means of these may and must contracts which can be inserted anywhere in the code (not just at component interfaces), the user and the tool communicate with each other, enriching the raw code with annotations capturing the intent and correctness of the code. The tool also actively suggests annotations by prompting the user (e.g., "do you assume this input pointer is always non-null?" or "did you mean to return a pointer that points to sometime allocated memory (program path A) and sometimes NULL (program path B)?").

Compositional reasoning allows the automatic inference of properties of the whole system by combining properties of sub-units: may contracts (summaries) can be combined to prove that some bad things cannot happen (proofs) while must contracts can be used for automatic test generation of system tests and bug finding. The framework can be extended to functional properties and non-functional properties.

3.2 Non-functional Analysis

Non-functional properties, such as performance, can cause some of the most expensive and difficult to debug problems within applications. However, *which* non-functional properties are important often depends upon the type of application being written. For example, an authentication server is critically concerned with security, while an embedded system may have no security constraints but may require worst-case execution time bounds and guarantees of numeric precision.

The FITE architecture will support plug-ins that can perform a wide variety of specialized, non-functional analysis. The IDE itself will have the concept of an analysis load-set, which allows a developer or project manager to determine which non-functional analysis are available (and most important) for the class of application being created.

We consider a handful of non-functional analysis below. These are meant to be representative rather than exhaustive:

Worst-Case/Average Case Timing Analysis A standard area of concern for developers are possible performance bottlenecks within an application. Symbolic evaluation tools such as JPF, Pex allow examination of code paths to determine which paths are longest or are known to make "expensive" API calls.

A plug-in that could flag potential performance bottlenecks could involve predefined configuration data that catalogues the relative cost of system functions and integration of this data with either (1) a symbolic simulator to describe feasible paths through the code or (2) an abstract interpretation engine (e.g., AbsInt). The symbolic simulator may have an advantage in that it may be able to sum-across-paths to talk about variance between symbolic paths and approximate average case performance, while abstract interpretation may be better at providing conservative guarantees about worst-case execution time.

Security Security problems are endemic to modern software. Old problems such as buffer overflows, continue to plague a variety of applications. More generally, attacks involving authentication, back-doors, SQL injection attacks, input validation, and many other causes cost billions of dollars in direct costs (to find, fix and patch) and in indirect costs (e.g., identity theft).

Existing software tools, such as SAGE, can automate many buffer overflow checks. SAGE is designed to run on large binary programs. We believe that we can create more precise analysis by reducing the scale of programs to be analyzed, and to broaden the category of attacks that can be checked. For example, statements creating dynamic SQL queries should be flagged and analyzed to prevent SQL injection attacks. Automated suggestions for writing SQL-injection resistant code, such as using stored procedures with typed parameters. Using JPF, Pex it should be possible to provide test cases that demonstrate SQL-injection attacks.

Numeric Precision Floating point numbers do not exactly represent real numbers, and the imprecision between complex computations over the reals and over floats can become significant. For example, the failure of the Patriot missile system (resulting in the deaths of 28 American soldiers) was due to cumulative imprecision in a floating point timing routine. Determining the loss of precision is therefore a common analysis that must be performed, usually manually, over embedded systems code.

Using symbolic execution, it is possible to describe imprecision at a per-path level. A näive approach would take the symbolic path and concretize it and then use interval analysis to examine the imprecision. It is not enough to enumerate the paths, however: one must examine the range of concrete values that are possible instantiations of the path in order to bound the precision errors that are possible. FITE should include a plug in that can generate both a constraint describing worst-case precision errors and a test (or series of tests) that demonstrate imprecise paths.

Concurrency It is difficult to reason compositionally about concurrency using most programming languages. However, concurrency bugs are among the most expensive to detect and fix. Recent work, such as the Symbolic Deadlock Analysis work by Deshmukh, Emerson, and Sankaranaryanan, can describe contacts between libraries and clients that guarantee deadlock free execution. This work is scalable enough to analyze large systems (1M SLOC Java / hour). Once the contracts are known, it is possible to cheaply analyze violations of the contract and present the result as a test case.

However, the other main concurrency problem of data races currently has no simple analysis solution. A task that could execute in the background and do pairwise analysis of "likely" concurrent method calls that could involve data races using a

tool like CHESS could be extremely valuable. The research challenges here involve the scale of the analysis and determination of "likely" interacting methods.

3.3 Regarding the user interface

Static analysis can determine potential program errors. Today, a typical IDE shows error messages of the type checker and other static analysis tools, relating them to particular lines in the code. We propose to augment this information with information gathered from and related to test cases. This includes already existing test cases as well as new test cases generated in addition. The additional information is meant to guide the developer towards errors directly related to the code the developer is currently working on. It is important that the additional information does not distract the developer from the main objective of writing code. Only relevant information must be shown. If a test case exposes an error, the code editor will associate the corresponding line in the code with the failing test case, also showing the stack trace of the failure. In addition to existing test cases, FITE generates new test cases, e.g. with (dynamic) symbolic execution. New test cases can be generated from scratch, or by "fuzzing" existing tests. When generating new tests, we will leverage the semi-automated analysis of interface boundaries and contracts. The analysis of existing test cases, and the generation of new test cases may happen continuously in the background, possibly on spare cores of modern multi-core machines, or the analysis may be delegated to the cloud. The developer can choose to include generated test cases into a regression test suite with a single button click, as for example realized in Pex:

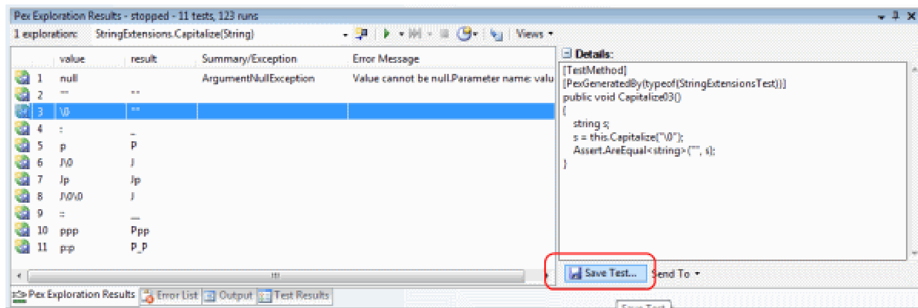

Since environment abstraction, i.e. automated generation of mock objects, may cause generation of test cases with spurious errors, i.e. errors which cannot occur in the integrated system, we propose a ranking of generated tests, and their failures. The result can be visualized by a "heat map" in the editor, which illustrates the points in the code at which generated tests cause failures, showing those failures which are most likely to reproduce in the integrated system in the most threatening color. When test cases cause a failure, an automated failure root cause analysis determines the failure condition, and suggests to the developer the addition of a precondition into his code, effectively raising the failure to the abstraction level of the code the developer is currently working on. This has already been realized in Pex:

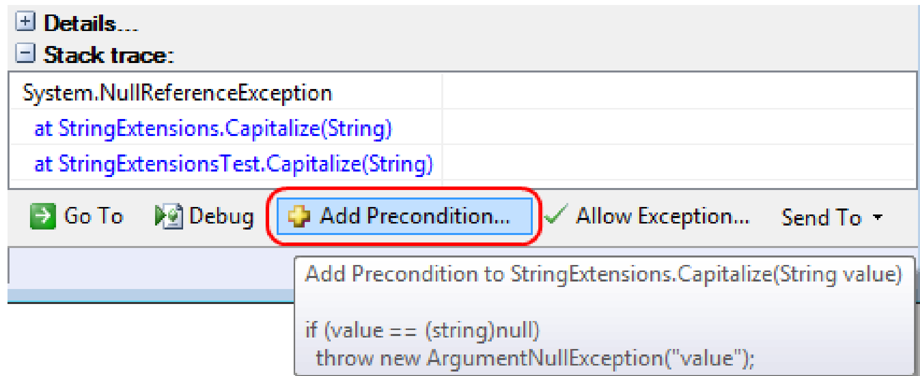

When the developer write new test cases, the editor will give suggestions what methods to call in a new test case. To this end, existing test cases and their code coverage are analyzed in the background to determine which methods of the product code or underrepresented in the existing test cases.

4 Process issues

So far the discussion has been mainly focused on unit analysis, from the developer's point of view. However, the approach and the environment we envision should assist developers and testers during the whole application development process. It would be particularly effective to anticipate possible integration issues at the time of developing each single module composing the entire system.

The FITE environment will base its analysis strategies and corresponding results taking into account also the other modules to which the module under development is interrelated. To do this the FITE environment will need to be implemented as a distributed and collaborative environment running on the cloud. In this phase particularly important are possible interactions with legacy modules to be integrated within the system. For such modules FITE will need to include an analysis step to investigate and highlight possible integration issues. The analysis is performed on-the-fly while the developer is coding the module, analyzing the consequences of the decisions on legacy module usage.

The analysis can successively be pushed even further permitting to derive integration test cases covering possible interaction sequences within the system when component are ready to be released. The integration steps will be supported by FITE also through the semi-automatic derivation of stubs for the different test cases.

4.1 From unit to integration / system testing

While a unit test targets a single isolated features, a system test spans multiple features. Via semi-automatically inferred interface boundaries/contracts, FITE is able to assist the developer locally with unit tests, where those parts of the system not currently under

test are mocked. In addition to traditional interface contracts, we propose to augment interface descriptions with a facility that allows to turn such mock instances into real instances. For example, when the database was mocked while testing a web application, then it should be possible to turn such a mock database state into an actual database state. This will in effect allow to turn unit tests into integration tests.