

Follow the Path: Debugging Tools for Test-Driven Fault Navigation

Michael Perscheid and Robert Hirschfeld
Software Architecture Group
Hasso-Plattner-Institute
University of Potsdam, Germany
firstname.lastname@hpi.uni-potsdam.de

Abstract—Debugging failing test cases, particularly the search for failure causes, is often a laborious and time-consuming activity. Standard debugging tools such as symbolic debuggers and test runners hardly facilitate developers during this task because they neither provide advice to failure causes nor back-in-time capabilities.

In this paper, we present test-driven fault navigation as a debugging guide that integrates spectrum-based and state anomalies into execution histories in order to systematically trace failure causes back to defects. We describe and demonstrate our Path tools that implement our debugging method for the Squeak/Smalltalk development environment.

Index Terms—Debugging, Testing, Back-in-time, Spectrum-based Anomalies, Likely Invariants, Squeak, Smalltalk

I. INTRODUCTION

The correction of software failures tends to be very cost-intensive because their debugging is an often time-consuming development activity. During this activity, developers largely attempt to understand what causes failures: Starting with a test case that reproduces the observable failure, they have to follow failure causes along the infection chain back to the root cause (defect). This idealized procedure requires deep knowledge of the system because failures and defects can be far apart. Unfortunately, common debugging tools are inadequate for systematically investigating such infection chains in detail [1].

Our *test-driven fault navigation* [2] is a debugging guide that integrates spectrum- [3] and state-based [4] anomaly detection into a systematic breadth-first search for tracing failure causes back to defects. Starting with at least one test case that can reproduce the observable failure, we localize anomalies by comparing the method coverage and occurred state properties of all failed and passed test cases. For example, methods that are being executed by a large number of failing but only a few passing tests have a higher failure cause probability (anomaly) than methods that are being executed by less failing but many passing test cases. By integrating such anomalies into a back-in-time debugger [5], we highlight suspicious method calls, reveal corrupted state properties, and allow developers to distinguish between suspicious and expected run-time behavior. Thus, our execution histories include additional information about failure cause probabilities that gives developers helpful advice on how to follow infection chains back to their root causes.

Our systematic search consists of four specific navigation techniques that together support the creation, evaluation, and refinement of failure cause hypotheses. First, *structure navigation* localizes suspicious system parts and restricts the initial search space. Second, *team navigation* recommends experienced developers for helping with failure causes by focusing on authors of anomalies. Third, *behavior navigation* allows developers to follow emphasized infection chains back to root causes. Fourth, *state navigation* identifies corrupted state and reveals parts of the infection chain automatically.

In this paper, we present our *Path tools* that implement test-driven fault navigation for the Squeak/Smalltalk [6]. We demonstrate their application by debugging one example failure step by step until the root cause has been found.

The remainder of this paper is structured as follows: Section II introduces our Path tools that implement our debugging approach. Section III demonstrates a debugging session. Section IV presents related work. Section V concludes.

II. THE PATH TOOLS

Our Path tools as presented in Fig. 1 primarily consist of two debugging tools that realize test-driven fault navigation.

A. PathMap: Extended Test Runner Feedback

PathMap [7] is an extended unit test runner for implementing our structure, team, and state navigation. It does not only verify test cases but also localizes failure causes. Its integral components are from left to right a testing control panel (A), a compact tree map visualization of the software system (B), and several flaps for accessing various analysis techniques (C).

The testing control panel (A) provides nearly the same functionality as a standard test runner. Developers can choose from different test suites of the selected project and run them.

The tree map (B) in the middle presents the structure of the system under observation with its packages, classes, and methods. Each small method box can be colored to represent test case and analysis results. It is possible for developers to interact with the tree map: receive the name of a source code entity by hovering over a box; request additional information about colored metric values (as shown in the message box); or debug into a suspicious method execution with our lightweight back-in-time debugger called PathFinder. Furthermore, the test runner also presents a status bar on the top displaying a

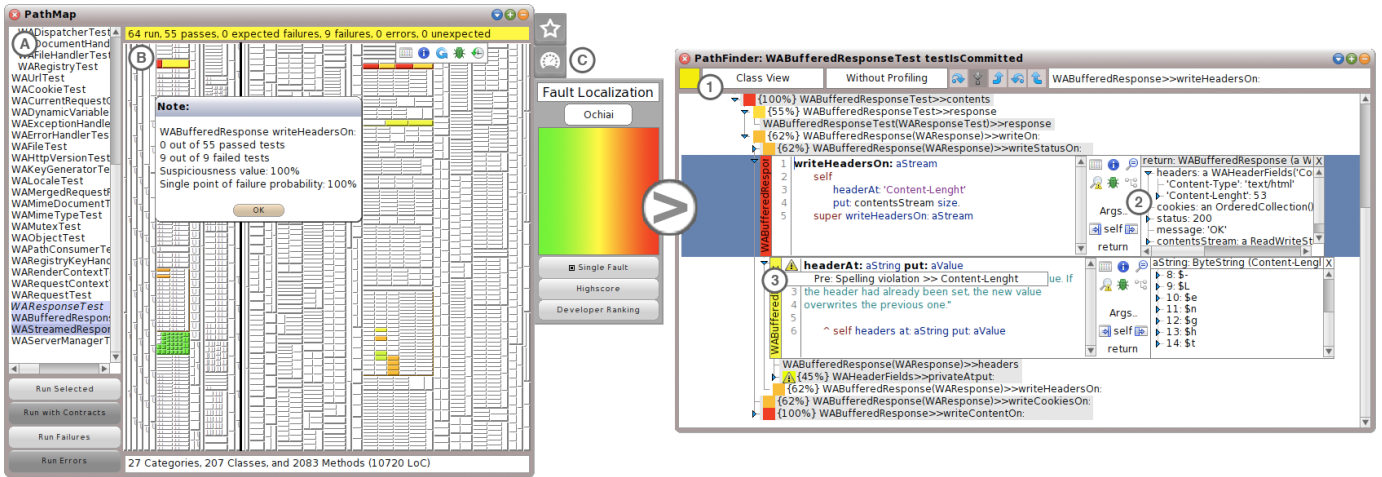


Fig. 1. PathMap as an extended test runner and PathFinder as a lightweight back-in-time debugger implement our test-driven fault navigation.

summary of the test suites execution and a status bar on the bottom displaying a summary of project metrics.

Flaps on the right (C) set PathMap into specific analysis modes for collecting valuable feedback during the execution of test cases. Without an opened flap, PathMap acts like a standard test runner and only colors test case results in the tree map. In Fig. 1 (left), the fault localization flap is open and allows developers to start the structure navigation [8]. If developers run selected tests, we automatically record their coverage, compute spectrum-based anomalies [3], and color the map with suspiciousness and confidence scores. Within the flap, developers can choose a proper spectrum-based metric, see a color legend, and enable filtering of partially covered methods. For our team navigation [8], pushing the “Developer Ranking” button presents a ranked list of expert developers by relating suspicious methods with their authors. Moreover, we offer other flaps to further reveal hidden test knowledge. For example, the inductive analysis flap allows developers to derive likely invariants [4] for our state navigation [9].

B. PathFinder: Lightweight Back-in-Time Debugger for Tests

PathFinder [10] is our lightweight back-in-time debugger [5] for exploring specific test case executions with a special focus on fault localization. Not only does it provide immediate access to run-time information [10], but also classifies traces with suspicious behavior [8] and anomalous state [9]. Developers start PathFinder at a chosen method in PathMap and follow the infection chain back to its root cause. Its main components are a control panel on the top and the test case execution history.

In Fig. 1 (right), PathFinder provides a control panel (1) to set up the dynamic analysis of test cases and to support the navigation through the large amount of run-time data. The yellow box presents the final test result and allows developers to rerun test cases. Views and profiling information enhance the results of initial call trees. While views record more details about called receiver objects such as names and identities, profiling precisely measures the required time for executing a test case. Finally, the subsequent buttons and a query engine

provide functionality for arbitrarily navigating through an execution history.

At the bottom of Fig. 1 (right), the visualized information primarily consists of a call tree that reflects a particular test case run. A call tree provides comprehensive information of the entire program execution and shows how methods call each other. From top to bottom, each node represents one method call and their subtrees describe its called methods. Each method call node consists of a colored box with a percentage value for its suspiciousness score and a name representing receiver class, implementation class in parentheses, and method name. We provide arbitrary navigation through method call trees and their state spaces (i. e. object properties). Developers can follow traces in both directions and expand and collapse subtrees interactively.

Some of the tree nodes have been expanded in Fig. 1 (right) to reveal details about the method implementation and the applied state. For example, an expanded method call node (2) shows its source code, a control panel for requesting details and refining run-time data, and the return value. Most notably, the control panel allows developers to start a source code editor and a symbolic debugger, obtain additional information about anomalies, refine coverage and spectrum-based fault navigation at statements, and explore object states. After indicating interest in a specific argument, receiver, or return object, we reexecute the test case, make a deep copy of the requested object, and present it in an object explorer on the right. Developers can explore all object properties and compare them to other method nodes in the execution history. Furthermore, our state navigation [9] maps violated contracts to traces by adding small exclamation marks (3) to methods. Developers can further inspect such state anomalies and receive detailed information about the exact violation.

III. DEBUGGING DEMONSTRATION

To demonstrate our test-driven fault navigation and its corresponding Path tools, we debug failure 4 of our previous

