



Quantitative Analysis of Fault and Failure Using Software Metrics

By Shital V. Tate & S. Z. Gawali

Bharati Vidyapeeth Deemed University, College of Engineering, Pune

Abstract - It is very complex to write programs that behave accurately in the program verification tools. Automatic mining techniques suffer from 90–99% false positive rates, because manual specification writing is not easy. Because they can help with program testing, optimization, refactoring, documentation, and most importantly, debugging and repair. To concentrate on this problem, we propose to augment a temporal-property miner by incorporating code quality metrics. We measure code quality by extracting additional information from the software engineering process, and using information from code that is more probable to be correct as well as code that is less probable to be correct. When used as a preprocessing step for an existing specification miner, our technique identifies which input is most suggestive of correct program behaviour, which allows off-the-shelf techniques to learn the same number of specifications using only 45% of their original input.

GJCST-C Classification: D.2.8



QUANTITATIVE ANALYSIS OF FAULT AND FAILURE USING SOFTWARE METRICS

Strictly as per the compliance and regulations of:



RESEARCH | DIVERSITY | ETHICS

Quantitative Analysis of Fault and Failure Using Software Metrics

Shital V. Tate^α & S. Z. Gawali^α

Abstract - It is very complex to write programs that behave accurately in the program verification tools. Automatic mining techniques suffer from 90–99% false positive rates, because manual specification writing is not easy. Because they can help with program testing, optimization, refactoring, documentation, and most importantly, debugging and repair. To concentrate on this problem, we propose to augment a temporal-property miner by incorporating code quality metrics. We measure code quality by extracting additional information from the software engineering process, and using information from code that is more probable to be correct as well as code that is less probable to be correct. When used as a pre-processing step for an existing specification miner, our technique identifies which input is most suggestive of correct program behaviour, which allows off-the-shelf techniques to learn the same number of specifications using only 45% of their original input.

I. INTRODUCTION

Software remains buggy and testing is still the leading approach for detecting software errors. Incorrect and buggy behaviour in deployed software costs up to \$70 billion each year in the US[1]. Thus debugging, testing, maintaining, optimizing, refactoring, and documenting software, while time-consuming, remain significantly important. Such maintenance is reported to consume up to 90% of the total cost of software projects[2]. Maximum maintenance time is spent studying existing software since maintenance concern is incomplete documentation.

Consistently, however, verification tools require specifications that describe some aspect of program accuracy. Creating accurate specifications is difficult, time-consuming and error-prone. Verification tools can only point out disagreements between the program and the specification. Even assuming a sound and complete tool, a defective specification can still yield false positives by pointing out non-bugs as bugs or false negatives by failing to point out real bugs. Crafting specifications typically requires program-specific knowledge.

Specification mining can be compared to learning the rules of English grammar by reading essays written by high school students; we propose to focus on the essays of passing students and be doubtful of the essays of failing students. We claim that existing miners have high false positive rates in large part because they

treat all code equally, even though not all code is created equal. For example, consider an execution trace through a recently modified, rarely-executed piece of code that was copied-and-pasted by an inexperienced developer. We argue that such a trace is a poor guide to correct behaviour when compared with a well-tested, infrequently-changed, and commonly-executed trace.

Various pre-existing software projects are not yet formally specified[3]. Formal program specifications are difficult for humans to construct[4], and incorrect specifications are difficult for humans to debug and modify[5]. Accordingly, researchers have developed techniques to automatically infer specifications from program source code or execution traces[6],[7],[8],[9]. These techniques typically produce specifications in the form of finite state machines that describe legal sequences of program behaviours.

Unfortunately, these existing mining techniques are insufficiently precise in practice. Some miners produce large but approximate specifications that must be corrected manually [5]. As these large specifications are indefinite and difficult to debug, this article focuses on a second class of techniques that produce a larger set of smaller and more precise candidate specifications that may be easier to evaluate for correctness. These specifications typically take the form of two-state finite state machines that describe temporal properties, e.g. “if event a happens during program execution, event b must eventually happen during that execution.” Two-state specifications are limited in their expressive power; comprehensive API specifications cannot always be expressed as a collection of smaller machines[8].

Recognize and illustrate lightweight, automatically collected software features that fairly accurate source code quality for the purpose of mining specifications. In this approach explain how to lift code quality metrics to metrics on traces, and empirically measure the utility of our lifted quality metrics when applied to previous static specification mining techniques. To avoid false positives recommend two novel specification mining techniques that use our automated quality metrics to learn temporal safety specifications.

II. ON GOING METHODOLOGY

a) *Specification Mining With Few False Positive*

This methodology presents a new automatic specification miner that uses artifacts from software

Author α : Department of Information Technology, Bharati Vidyapeeth Deemed University, College of Engineering, Pune-46.

engineering processes to capture the reliability of its input traces.

The main contributions of this project are:

- A set of source-level features related to software engineering processes that capture the trustworthiness of code for specification mining. We analyze the relative analytical power of each of these features.
- Experimental evidence that our notions of trustworthy code serve as a basis for evaluating the trustworthiness of traces. We provide a characterization for such traces and show that off-the-shelf specification miners can learn just as many specifications using only 60% of traces.
- A novel automatic mining technique that uses our trust-capturing features to learn temporal safety specifications with few false positives in practice. We evaluate it on over 800,000 lines of code and explicitly compare it to two previous approaches. Our basic mining technique learns specifications that locate more safety-policy violations than previous miners (740 vs. 426) while presenting far fewer false positive specifications (107 vs. 567). When focused on precision, our technique obtains a low 5% false positive rate, an order-of-magnitude improvement on previous work, while still finding specifications that locate 265 violations. To our knowledge, this is the first specification miner that produces multiple candidate specifications and has a false positive rate under 90%.

i. Approach

In this approach present a specification miner that works in three stages:

1. Statically estimate the trustworthiness of each code fragment.
2. Lift that judgment to traces by considering the code visited along a trace.
3. Weight the contribution of each trace by its trustworthiness when counting event frequencies for specification mining.

The code is most trustworthy when it has been written by experienced Programmers who are familiar with the project at hand, when it has been well-tested, and when it has been mindfully written.

b) Mining Temporal Specification for Error Detection

If we use implicit language-based specifications (e.g., null pointers should not be dereferenced) or to reuse standard library specifications then it can reduce the cost of writing specifications. More recently, however, a variety of attempts have been made to conclude program-specific temporal specifications and API usage rules automatically. These specification mining techniques take programs (and possibly dynamic traces, or other hints) as input and produce

candidate specifications as output. Basically specifications could also be used for documenting, refactoring, testing, debugging, maintaining, and optimizing a program. Centre of attention is that finding and evaluating specifications in a particular context: given a program and a generic verification tool, what specification mining technique should be used to find bugs in the program and thereby improve software quality? Thus we are concerned both with the number of "real" and "false positive" specifications produced by the miner and with the number of "real" and "false positive" bugs found using those "real" specifications.

In this methodology propose a novel technique for temporal specification mining that uses information about program error handling. Our miner assumes that programs will generally adhere to specifications along normal execution paths, but that programs will likely violate specifications in the presence of some run-time errors or exceptional situations. Intuitively, error-handling code may not be tested as often or the programmer may be unaware of sources of run-time errors. Taking advantage of this information is more important than ranking candidate policies.

i. Contributions

- Propose a novel specification mining technique based on the observation that programmers often make mistakes in exceptional circumstances or along uncommon code paths.
- Present a qualitative comparison of five miners and show how some miner assumptions are not well-supported in practice.
- Finally, we give a quantitative comparison of our technique's bug-finding powers to generic "library" policies. For our domain of interest, mining finds 250 more bugs. We also show the relative unimportance of ranking candidate policies. In all, we find 69 specifications that lead to the discovery over 430 bugs in 1 million lines of code.

III. PROPOSED SYSTEM FOR QUANTITATIVE ANALYSIS OF FAULT AND FAILURE

In proposed system, aim to develop a system which can be used to measure the quality of the code considering different aspects affecting the quality of the code. The term quality of the code can be explained using different factors such as code clone, author rank, code churn, code readability, path feasibility etc.

To Present a new specification miner that works in three stages. First, it statically estimates the quality of source code fragments. Second, it lifts those quality judgments to traces by considering all code visited along a trace. Finally, it weights each trace by its quality when counting event frequencies for specification mining.

This system develops an automatic specification miner that balances true positives – as required behaviours –with false positives – non-required behaviours. We claim that one important reason that previous miners have high false positive rates is that they falsely assume that all code is equally likely to be correct. For example, consider an execution trace through a recently modified, rarely-executed piece of code that was copied and-pasted by an inexperienced developer. We believe that such a trace is a poor guide to correct behaviour, especially when compared with a well-tested, stable, and commonly-executed piece of code. Patterns of specification adherence may also be useful to a miner: a candidate that is violated in the high quality code but adhered to in the low quality code is less likely to represent required behaviour than one that is adhered to on the high quality code but violated in the low quality code. We assert that a combination of

lightweight, automatically collected quality metrics over source code can usefully provide both positive and negative feedback to a miner attempting to distinguish between true and false specification candidates.

Code quality information may be gathered either from the source code itself or from related artefacts, such as version control history. By augmenting the trace language to include information from the software engineering process, we can evaluate the quality of every piece of information supporting a candidate specification (traces that adhere to a candidate as well as those that violate it and both high and low quality code) on which it is followed and more accurately evaluate the likelihood that it is valid.

The system architecture of the system is as in following figure, which explains the modules to be generated.

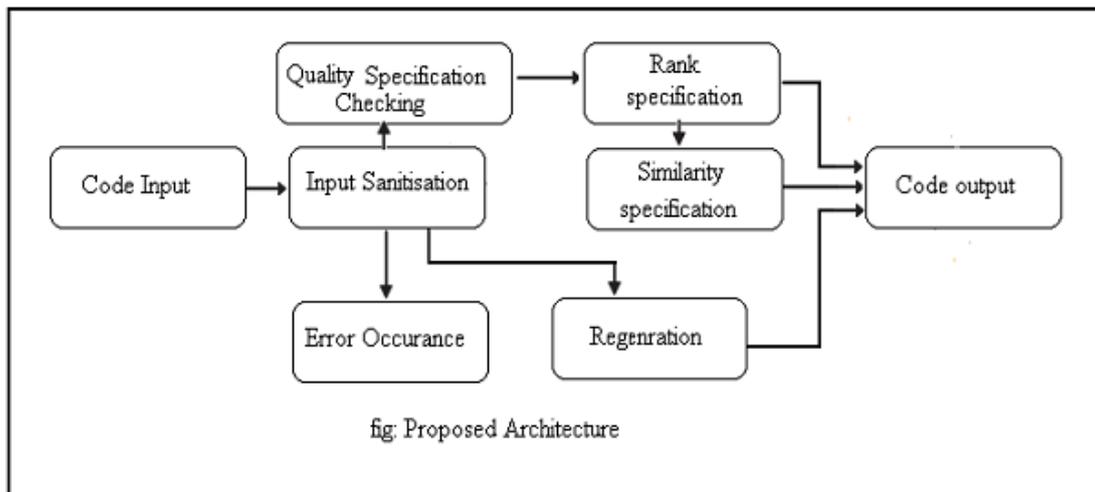


Figure 1 :

a) Description of proposed system

Proposed system for quantitative analysis of and fault and failure using software metrics uses the following stages-

1. Accept input in the form of computer program code.
2. Perform input sanitization.
3. Check for error occurrence in the code.
4. Check for the quality specification regarding the given code.
5. Specify the rank for the different condition, using calculated result.
6. Generate output in the form of quality report.

IV. CONCLUSION

Testing, maintenance, optimization, refactoring, documentation, and program repair these are the various applications of formal specification. Though human programmers should not produce and verify such specification manually. These technique is also problematic since it treat all parts of program as equally indicative as correct behaviour. We encode this intuition

using dependability metrics such as analytical execution frequency, copy paste code measurements, code duplication software readability or path feasibility. We compare the bug finding power of various miners. This technique improves the performance of existing trace based miners by focusing on high quality traces. Our technique is also useful to improve the quality of code through specification mining.

REFERENCES RÉFÉRENCES REFERENCIAS

1. National Institute of Standards and Technology, "The economic impact of inadequate infrastructure for software testing," Tech. Rep. 02-3, may 2002.
2. R. C. Seacord, D. Plakosh, and G. A. Lewis, Modernizing Legacy Practices, 2003.
3. M. Das, "Formal specifications on industrial-strength code from myth to reality," in Computer-Aided Verification, 2006, p. 1.

4. H. Chen, D. Wagner, and D. Dean, "Setuid demystified," in USENIX Security Symposium, 2002, pp. 171–190.
5. G. Ammons, D. Mandelin, R. Bodík, and J. R. Larus, "Debugging temporal specifications with concept analysis," in Programming Language Design and Implementation, 2003, pp. 182–195.
6. G. Ammons, R. Bodik, and J. R. Larus, "Mining specifications," in Principles of Programming Languages, 2002, pp. 4–16.
7. D. R. Engler, D. Y. Chen, and A. Chou, "Bugs as inconsistent behaviour: A general approach to inferring errors in systems code," in Symposium on Operating System Principles, 2001, pp. 57–72.
8. M. Gabel and Z. Su, "Symbolic mining of temporal specifications," in ICSE, 2008, pp. 51–60.
9. J. Whaley, M. C. Martin, and M.S. Lam, "Automatic extraction of object-oriented component interfaces," in ISSTA, 2002.
10. Claire Le Goues, Westely Weimer "Measuring code quality to improve specification mining" IEEE Trans. Software Eng.
11. Mohammed Kayed and Chia-Hui Chang, "FiVaTech: Page-Level Web Data Extraction from Template Pages" IEEE Transactions On Knowledge And Data Engineering, Vol. 22, No. 2, February 2010
12. S. R. Chidamber and C. F. Kemerer, "A metrics suite for object oriented design," IEEE Trans. Softw. Eng., vol. 20, no. 6, pp. 476–493, 1994.
13. D. Detlefs, G. Nelson, and J. B. Saxe, "Simplify: a theorem prover for program checking," J. ACM, vol. 52, no. 3, pp. 365–473, 2005.
14. M. Di Penta and D. M. German, "Who are source code contributors and how do they change?" in Working Conference on Reverse Engineering. IEEE Computer Society, 2009, pp. 11–20.
15. C. Kapsner and M. W. Godfrey, ""Cloning Considered Harmful" in WCRE, 2006, pp. 19–28.
16. J. Krinke, "A study of consistent and inconsistent changes to code clones," in WCRE. IEEE Computer Society, 2007, pp. 170–178.
17. C. Le Goues and W. Weimer, "Specification mining with few false positives." In TACAS, 2009, pp. 292–306.
18. T. J. McCabe, "A complexity measure," IEEE Trans. Software Eng., vol. 2, no. 4, pp. 308–320, 1976.
19. N. Nagappan and T. Ball, "Using software dependencies and churn metrics to predict field failures: An empirical case study," in ESEM, 2007, pp. 364–373.
20. J. C. Sanchez, L. Williams, and E. M. Maximilien, "On the Sustained Use of a Test Driven Development Practice at IBM," in Agile 2007. IEEE Computer Society, August 2007, pp. 5–14.
21. W. Weimer and N. Mishra, "Privately finding specifications," IEEE Trans. Software Eng., vol. 34, no. 1, pp. 21–32, 2008.
22. W. Weimer and G. C. Necula, "Mining temporal specifications for error detection," in TACAS, 2005, pp. 461–476.
23. D. R. Engler, D. Y. Chen, and A. Chou. "Bugs as inconsistent behavior: A general approach to inferring errors in systems code". In Symposium on Operating Systems Principles, pages 57–72, 2001.