

Continuous Self-Evaluation for the Self-Improvement of Software

Lori A. Clarke and Leon J. Osterweil*
Department of Computer Science

University of Massachusetts
Amherst, MA 01003
USA

*This work is being carried out as a collaboration among Leon J. Osterweil, Lori A. Clarke and George Avrunin at the University of Massachusetts, Debra Richardson at the University of California at Irvine, and Michael Young at the University of

1.0 INTRODUCTION:

The purpose of evaluation is to determine the extent to which a software system is meeting its requirements, and to suggest, as much as possible, the software modifications that should be expected to help improve its ability to meet its requirements. Our interest is in improving software. Thus in this paper we describe a self-evaluation approach, called perpetual testing and analysis, that is a part of a larger activity of self-improvement.

In a manual, or human-driven, software improvement process humans take the responsibility for carrying out the testing and evaluation of software, humans infer the changes that should be made, humans effect the changes, and humans reinstall the modified software. In this paper we describe a software self-improvement process in which some or all of these activities are to be assisted, or entirely carried out by the software itself, rather than being done solely by humans. There is an interesting research required in order to reduce the dependence upon humans in carrying out these activities. In this paper we describe how we propose to transition the responsibility for testing and evaluation of software from humans and onto software tools and processes.

The essence of our idea is that deployed software should be under continuous testing, analysis, and evaluation, drawing heavily upon the computational resource utilization patterns found in the deployment environment. As the amount of software that can be tested and analyzed are virtually limitless, there is value in carrying out these activities perpetually. It is important, however, that the activities not be undirected. We propose that the process of perpetual testing and analysis be carried out in such a way as to produce increasingly sharp and specific suggestions about what modifications to the software seem most likely to result in improvements. Thus, our perpetual testing and analysis proposal is aimed

reducing the need for human involvement in evaluation, but also at assist in proposing improvements.

There are substantial technical challenges in doing this. We propose that testing and static analysis go on essentially endlessly, but that results from these activities be used to focus and sharpen the other's activity. Research that makes these two complementary evaluation approaches most synergistic is not yet in vision. Orchestrating this synergy through the use of precisely defined tools and processes there is considerable research to be done in this area as well.

2.0 APPROACH

While self-modifying code has existed since the earliest days of computing, there is ample evidence that it poses clear dangers. Of most concern to us is that modifying code is generally difficult or impossible to analyze, and therefore its possible behaviors is difficult or impossible to bound. For this reason, we envision in which software can modify itself, without necessitating the need for a programmer to do it itself. Our approach to this problem is to view a software product as a collection of diverse types of software artifacts, including such types of components as a requirements specification, an architecture, low level design specifications, code, test cases, and analysis results. We also suggest that the software product also include the process by which various components of the software product (code, in particular) are to be modified. By doing so we can assure that no component of the software product need modify itself. Rather, a separate component, the modification process, is responsible for modifying other components, such as the code.

Going further, we suggest that the software product is also characterized by constraints that specify the way in which the various components should be related to each other. The constraints are particularly important because they determine when and how the components of the product need to be modified. For example, when test results derived through execution of testcases are inconsistent with the requirements to which they should be related, this is a signal that a modification is needed.

The foregoing suggests that a collection of tools should also be considered as part of the software product. These tools are the devices that are used to help analyze the components of the product and to determine the degree to which the product is consistent. Thus compilers, design tools, testcase monitors, and static analysis tools are examples of tools that should be considered to be part of the overall software product.

In classical software development, the tools are applied in order to build the software product and its components right up until deployment. But at deployment the code is separated from the rest of the components, and from its constraints, and is placed in the deployed environment. This complicates the modification process substantially.

We propose that deployed software product code remain tethered to the other components (including the software product modification process), as well as the constraint and tool sets that comprise a complete product. By doing this, it is possible for tools to continue to evaluate the consistency of the code with its environment and to effect modifications. This is the more precise sense in which continuous testing suggests how to effect the self-modification of software, without necessitating the modification of the code component.

Thus the perpetual testing approach implies that software code be perpetually accessible by access to an environment that supports its evaluation and improvement (the development environment, even though it will persist past development), and takes a pro-active role in assuring that evaluation contributes positive contributions to the improvement of the software. Coordinating numerous types of testing and analysis artifacts that exist in the development and deployed environments is a daunting job that is prohibitively expensive if carried out manually. Instead, a highly automated testing and analysis process using testing and analysis tools and artifacts, using the available computing resources, can be acquired at any given time. Although this process is considered to be external to the product, it is not required to be resident with the code in the development environment.

3.0 TECHNICAL AND RESEARCH CHALLENGES

3.1 DEPLOYMENT CONSIDERATIONS

It is important to emphasize that the nature of the interconnection between the development environment and the deployed code that is (perpetual testing) will vary considerably, depending upon the variation in deployment situations. Thus, for example, it may be quite reasonable to suggest that non-critical prototype research software deployed in a research setting may be under continuous evaluation, and continual interaction with its development environment. On the other hand, it is unthinkable that mission critical realtime software code deployed in a secure mission environment will have direct contact with its original development environment while it is deployed and in service. There is a large spectrum of circumstances between these two extreme situations. It is our belief, however, that we can fashion a corresponding spectrum of approaches to supporting the realization of the benefits of perpetual testing. Certainly the degree to which we connect deployed code to the rest of the product will dictate the degree of challenges we are likely to have with self-evaluation and self-improvement.

Clearly the furnishing of parameterized monitoring capabilities, coordination of communication between deployment and development environments, and the rehosting of major portions of the development environment all pose significant technical challenges.

3.2 ANALYSIS AND TESTING RESEARCH

3.2.1 INCREMENTAL RETESTING AND REANALYSIS

Testing and analysis techniques should be applicable continuously from initial software development, where it should be applied to incoming systems, and continuously as software modifications are considered and carried out. As modifications are being considered, it is highly desirable to repeat analysis and testing for unchanged aspects of the software. This entails the ability to carry out careful analysis of the impact of modifications, determination of which past analytic results remain valid, determination of the most pressing retesting/reanalysis needs, and an optimized process for addressing the most pressing needs first.

Incremental reanalysis and testing will require technologies for a web of summary information about the code and its various subcomponents, deciding what parts are partly or fully reusable, and optimizing portions that changes have rendered invalid. Incremental approaches have been developed for compilation, based on interprocedural data-flow techniques. Incremental techniques developed for testing and analysis are conceptually similar, but will involve a much more complex set of relations among multiple sources and kinds of information. The nature of these relations is indicated briefly below.

3.2.2 INTEGRATION OF TECHNIQUES

The prospect of Perpetual Testing suggests that diverse analysis techniques will have to be synergistically integrated and reintegrated dynamically and in diverse ways. Previous research has described how different analysis techniques could complement one another, but the proposed fixed integration strategies consistent with the assumption that testing was a phase of fixed duration. These activities have been expensive, but usually less precise, analysis techniques with more human intervention. For example, static dataflow analysis scans and sharpens dynamic testing regimens, and can be used to iteratively solve successions of dataflow analyses, that iteratively sharpen results. These prototype integrations indicate many opportunities for synergy, and the clear feasibility of effecting new integrations.

Verification systems have traditionally maintained constraint relations, lemmas and theorems, but have had an overly simplistic view that a property is either proved (when all its relations are proved), or is not supported. Testing systems have maintained records of thoroughness (notions of coverage), as well as sets of properties in the form of analysis tools have been limited to a distinction between verified (must results) and inconclusive (may) results.

Integration of analysis and testing techniques, and particularly post-deployment usage information, provides an opportunity for a set of constraints among properties, techniques, and levels of as

example, a property may be verified by a static dataflow analysis dependent on the absence of an aliasing relation that is monitored. Each asserted property of a software system can be supported by a set of analyses and assurances of differing strengths, which are propagated through the web of relations and constraints. Monitoring of deployed software is critical in this regard, and explicit constraint links monitored in the deployed software provide a way to calibrate and assurances established in the development environment.

3.2.3 SPECIFICATION-BASED ANALYSIS

Testing is a human intensive process. Testers must develop test cases and then evaluate the results. The latter task can be consuming, tedious, and error prone. Specification-based testing are being developed so that humans no longer have to play the role. The use of specifications needs to be further developed, not only but for a range of analysis techniques.

Specifications are the driving impetus for a number of analysis approaches. Some modern dataflow analyzers use a quantified form of regular expressions for expressing properties that are to be validated, others use test specifications, while still others use graphical interval logic. What is actually done and the validity of the system being evaluated against the specification. Needless to say the quality of the analysis depends on the quality of the specification. For example, data flow analysis can determine the validity of a property unless that property is captured in the specification.

An advantage of specification-based analysis and testing is that the specification base can grow. Over time, as users have more experience with the software product, new specifications can be formulated. Errors discovered after deployment should be captured by a new specification intended behavior so that future modifications of the software product are evaluated against this specification as well. As the specification grows, more and more of the reanalysis and retesting will be driven by the specification. This should improve the quality of the testing and analysis, but should also improve the quality of the software product, reduce the amount of intervention, and greatly reduce the time devoted to testing and

3.2.4 PREDICTIVE-DRIVEN ANALYSIS AND TESTING

A perpetual testing framework provides a unique opportunity to gain insight about the sequence of modifications to the software product and to use metrics to predict the most appropriate analysis or testing techniques for the evaluation of subsequent modifications. For example, a code segment that has a history of containing faults might be given a high priority for reanalysis if it is modified. Information about the kinds of faults discovered in the past and metrics about the component itself, such as its use of concurrency or its complicated data structures, would impact

of analysis approach to employ. Also, past execution costs could predict computing resources that would be needed to complete the a timely fashion. Finally, if one technique proved to require sub interaction in the past for problems with a similar metric footprint consumptive techniques would be considered.

Recent work on testing and analysis has just started to address u to drive the choice of analysis techniques. With perpetual testi be able to be more complete and effective in gathering informatio software being evaluated, the kinds of faults discovered (or prop verified) by the analysis and testing techniques, the computing r and the amount of human intervention required. On the basis of th information, we should be able to develop a predictive model of w to be the best testing and analysis process. This model would its subject of evaluation and would continue to be modified as we gat experimental evidence. This should lead to a predictive meta-mod incorporated as part of the overall software product. As more inf learned about a particular software product and its modification information would be used by the meta model to evolve a predictiv drive the testing and analysis process for that software.

3.2.5 PERPETUAL TESTING PROCESS DEVELOPMENT

Analysis and testing activities should be viewed as processes to as software is, from requirements through coding, evaluation, and This is particularly important for perpetual testing processes. B to be indefinitely ongoing processes, it is essential that clear beforehand, so that progress towards those goals be continually m and so that revisions to either goals, of processes, or both can continually. Thus, perpetual testing can reasonably be viewed as integration of sequences of applications of both testing and anal demonstrable support of precisely articulated analysis and testin propose to demonstrate the value of articulating requirements for testing processes, and then also demonstrate perpetual testing pr architectural designs.

It seems particularly appropriate to consider, in addition, the d code for analysis and testing processes. For example, dynamic reg testing entails iteration that is comfortably expressed with trad Deciding when and whether to follow coarse-grained, static datafl with more precise, sharpened dataflow analyses is expressible wit case and if-then-else constructs. In both cases, the fine scale products must also be specified.

The testing and analysis process we envisage will also require re mechanisms. It is often important to program immediate reactive to such testing failures as incorrect results and system crashes. analysis results may need to trigger fixed standard subsequent re

reporting of analysis results obtained at deployment sites might be by timer events, or by the very action of their being completed.

All of this argues for the specification of actual executable programs to guide the execution of these perpetual testing processes, and the continuous self-evaluation of software products. In our work we demonstrated such perpetual testing process code, and now propose it to be used as the engine for driving software product self-improvement.

4.0 SUMMARY

The perpetual testing and analysis approach promises to enable the continuous self-evaluation of software throughout the sequence of modifications during its entire lifetime, and thereby to enable software self-improvement to be measured and evaluated. Key to doing this is to perpetually integrate software code to the rest of the overall software product's component constraints, tools, and the perpetual testing process itself. This will increase the confidence that people will have in their software products. As software products become larger and more complex, they will become increasingly difficult to evaluate and improve, trust, and predict unless an approach to perpetual testing is explored.

5.0 ACKNOWLEDGEMENTS

This research was partially supported by the Air Force Research Laboratory and the Defense Advanced Research Projects Agency under Contract F30600-92-2-0032, and by U.S. Department of Defense/Army and the Defense Advanced Research Projects Agency under Contract DAAH01-00-C-R231, The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied of the Defense Advanced Research Projects Agency, the Air Force Research Laboratory/IFTD, the U.S. Department of Defense, the U. S. Army, the U.S. Government, the National Science Foundation, or of IBM.