

Distributed Expectation-Driven Residual Testing

Leila Naslavsky, Roberto S. Silva Filho, Cleidson R. B. de Souza¹, Marcio Dias,
Debra Richardson, David F. Redmiles
School of Information and Computer Science
University of California, Irvine
{lnaslavs,rsilvafi,cdesouza,mdias,djr,redmiles}@ics.uci.edu

Abstract

Software testing techniques are widely used to assure properties such as dependability, as well as to gather system usage information to support software evolution. However, in spite of such techniques, comprehensive testing of commercial-grade software is rarely performed, resulting in software that is released with neglected test obligations. The dissemination power of the Internet and the increased connectivity of systems have made the pressure to release software earlier even more apparent. Fortunately, at the same time, researchers have investigated using the infrastructure of the Internet for performing remote testing and analysis. One particular approach called residual testing allows the testing of software after it is deployed to beta testers or end users. This paper builds upon existing residual testing techniques, presenting an approach to integrate residual testing with a way to dynamically measure the behavioral conformity of the deployed system with its specifications. Our approach uses the facility of an extensible and configurable publish/subscribe infrastructure, dynamic instrumentation techniques and software tomography to address many of the issues faced by current remote residual testing strategies.

1. Introduction

As part of the quality assurance process (QA), software testing activities must always deal with the accuracy versus cost trade-off [1]. One can aim at proving system correctness, performing exhaustive testing or covering all the possible execution paths at the expense of infinite effort. Yet, time-to-market pressures, cost constraints, difficulty in predicting all possible combinations of hardware and software configurations, among other concerns impose practical limitations to the amount of testing performed. As a consequence, unforeseen problems are usually manifested in the client site, when software is exposed to different configurations and usage patterns. In this context, monitoring and identification of such problems provide useful information when reported back to the QA team, which facilitates fixing any defects identified or adapting and evolving the software to new configurations.

Residual testing techniques, as introduced by Pavlopoulou and Young [2], allow software developers to cope with this inevitable situation. These strategies test software released either to beta testers or to end-users with respect to its remaining obligations (the residue). They employ selective software instrumentation and monitoring techniques to collect usage data for the purpose of software quality assurance and evolution. Combining their techniques with the connectivity of the Internet allows the automatic collection and delivery of the monitored information from the client sites to the QA and development teams. Indeed, many techniques have been proposed that allow for monitoring of deployed software behavior aiming at isolating errors [3], increasing coverage with minimal degradation of system performance [4, 5], and distributing QA tasks to users [6], among other goals.

Clearly, ensuring that a part of the system was covered does not eliminate the need for ensuring that the results produced and the behavior of the system match the specification of the software. Thus, one of the major needs for the residual testing approaches is to add support for behavioral conformance. In this paper, we propose an approach to address this issue that allows QA and development teams to perform behavior validation of remotely deployed software instances. The main idea is to translate remote execution information into higher-level constructs, that can be combined with specification-based test oracles to check system correctness [7, 8]. System correctness can be verified with respect to the behavioral specification of the system, as expressed in many different ways such as: state machines, UML sequence diagrams, among others. Our approach integrates residual testing and behavior verification to dynamically and incrementally measure the conformity of the deployed system with its specifications. Our ultimate goal is not only identify specified behavior but also detect unpredicted or abnormal usage of the system. We argue that we can achieve this goal by combining different techniques, including specification-based test oracles, dynamic instrumentation approaches, software tomography [4], a publish/subscribe infrastruc-

¹ Also at: Departamento de Informática, Universidade Federal do Pará. Belém – PA – Brasil.

ture, client-side event filters, event correlation and visualization gauges.

2. Related Work

In addition to the work of Pavlopoulou and Young [2], several systems have been used to remotely monitor aspects of software after its deployment for testing purposes. In fact, most of them focus on low-level aspects such as statement or method coverage, hot-spots identification, error and exception detection, and so on. This section describes some of those systems.

The Gamma system [4] uses software tomography to achieve low-impact and minimally-intrusive monitoring. Software tomography consists on dividing the overall system monitoring task into different subtasks which are strategically instrumented and assigned to different deployed software copies. Thus, only small parts of the instance's code needs to be instrumented, which significantly reduces the individual runtime performance degradation. To obtain system's full coverage, individual subtasks' execution information are collected and sent back to the QA site, where they are combined to measure the overall monitoring coverage of the system. Additionally, the Gamma System permits dynamic adjustment to the monitoring by supporting on-site code modification/update of the instances.

The Skoll system [6] provides a client-server architecture for testing highly configurable software under different hardware and software configurations. This approach allocates and distributes tests to end user sites, where they are executed, and provides the QA team with information on test failures, incompatibilities and their context. The allocation of resources is based on information provided by the end user about his or her client platform. Different software configurations are defined by the steering agent component of the Skoll platform. This agent sends jobs to the Skoll client, installed in a client host, based on the client hardware/software configuration. The software being tested is then downloaded to some of the remote sites, locally compiled and automatically tested, with the results sent back to the agent for further analysis.

EDEM's (Expectation Driven Event Monitoring) [9] goal is to monitor the usage patterns of end-user interfaces. The monitoring is based on high-level specifications (called expectations), that define normal usage patterns of the software based on GUI events. EDEM detects expectation violations, generating reports to the QA team. Expectations can be dynamically deployed using agents built by using event correlation techniques [10]. When an expectation is violated, the agent sends e-mail messages to the QA team. Expectation agents can be defined and deployed at execution time, with the help of a special agent editor. Data collected is used by developers to improve the usability of the application.

3. Approach

We build upon the techniques described in the previous section in order to provide distributed residual testing based on behavioral specifications. Figure 1 provides an overall depiction of the approach..




Figure 1 Overall description of the monitoring steps

3.1 Approach Overview

From the overall system *behavioral specification*, the developer selects a partial specification he is interested in monitoring. Such *behavioral specifications* can be expressed in different ways (e.g. state machines, LTL, UML sequence diagrams). They represent high-level descriptions of the system behavior to be verified. To be monitored, however, this high-level specification needs to be *mapped* to concerns in the name space of the system execution (runtime execution events). The product of this *mapping* is a *monitoring configuration*, which is deployed and instrumented alongside to the software instances running on the end users' machines. In our approach, this configuration can be dynamically changed, thereby supporting evolution of the monitoring configuration according to the developers' needs. As the program instances execute, their execution data is optionally filtered and then published to the notification service. This data is then *collected* and *consolidated* in the QA team site, to obtain the overall monitoring information of the software. This execution data can be *visualized* and *analyzed* with respect to the expected system behavior, either at run-time or off-line, using the data collected.

Our approach differs from previous work in the following. The Gamma system collects coverage informa-

tion as opposed to behavioral information, its subtask assignment and reassignment is based on instances' usage, static and statistical analysis. We suggest such decision should be based on information such as end users' usage patterns, their machines' specific platforms and software configurations. The Skoll system uses user's resources to perform QA tasks, while our approach monitors actual system usage deployed to beta testers and developers. Finally, the EDEM system is limited to user interface events and cannot be used to monitor the overall system behavioral conformance.

3.2 Communication Infrastructure

Since the components of our approach are integrated using the publish/subscribe paradigm [11], we must explain it first. In this paradigm, also called implicit invocation, event producers publish information in the form of events (records, objects or attribute/value pairs), while consumers express interest on events using logical expressions called subscriptions. Subscriptions can include more advanced features such as event correlation expressions, allowing not only the filtering of information, but also the combination and abstraction of events into higher-level events [10]. For example, a sequence of events or a set of repetitive events can be combined and abstracted in higher-level events that express such pattern of occurrence. A notification server implements a distributed publish/subscribe paradigm, representing a logically centralized service that routes information from distributed publishers to the interested subscribers. This insulation between producers and consumers copes with dynamism and scalability of the system allowing publishers and subscribers to dynamically join and leave the distributed system.

3.3 Detailed Description

This section describes, in more detail, the steps and main components of our monitoring infrastructure as generally explained in section 3.1. It also highlights the technologies used in each phase of our process.

Behavioral specification

Our approach requires the adoption of a behavioral specification (e.g. state machines, LTL, UML sequence diagrams) to which the system execution must conform. For our current purposes, we focus on object level residual testing, so we are interested in gathering information about object method invocations. This coarse-grained level of abstraction is used, rather than a fine-grained (such as statement execution) for the following reasons. First, deployed applications should be tested at a higher level of abstraction so as to avoid undue performance degradation. Second, this level of abstraction allows observation and analysis of software behavior at a level closer to software design. For the same reason, we use UML sequence diagrams to express the system behavioral specifications.

Mapping

The mapping between high-level UML specifications into low-level monitoring configurations is performed by a parser component. A monitoring configuration describes where, in the deployed software, the probes need to be inserted, as well as which methods are to be monitored, together with specific *sequences* of methods to be detected (expressed in the form of event-based *subscriptions*). Since we are dealing with sequence diagrams, this mapping can associate one or more runtime method invocations with each object interaction in the UML description. The adopted publish/subscribe service uses those subscriptions to detect the event sequences (either in the client site, using a special event correlation component, or in the notification server). Both monitoring configurations and subscriptions are deployed with software. For example, consider a UML sequence diagram that defines a network protocol, with messages such as `open()`, `close()`, `read()` and `write()`. A typical usage pattern of such methods would be: an `open()` followed by several `read()` and `write()` operations, in a valid order, ending with a `close()`. In our approach, these method invocations are monitored and represented as events. So, whenever a method is invoked, a new event is generated and published to the publish/subscribe infrastructure, which detects specific event sequences and inform their occurrence to subscribers. Hence, considering that all events are coming from the same source, a subscription that detects this sequence can be defined as: "*openEvent THEN (readEvent OR writeEvent)* THEN closeEvent*", where *** means 0 or more times.

Deployment and Installation

In addition to defining which behavioral specifications should be verified, the developer must decide where to deploy the monitoring configuration by choosing which software instance to monitor. Such decision is based on information such as end-users' software usage patterns or the site hardware and software configurations (similarly to the Skoll project [6]). This information is obtained by either automatic usage profile detection or by direct user inquiry.

Once deployed, the monitoring configuration must be installed in order to start collecting the necessary events, and identifying the required event sequences. The steps undertaken here are dependent on the technology adopted for instrumentation. For example, in the Gamma approach, this is done by dynamically replacing class files. Our approach uses a modified class loader to identify and instrument the source necessary .class files, at load time. More details will be discussed in the implementation section.

Data collection and correlation

The data collection phase is concerned with the identification of sequence of events that need to be sent to the developers. These identified sequences might repre-

sent an expected or unexpected execution of the software instance with respect to the related behavioral specification. The identification of such sequences is performed by a software component that correlates events. The correlation is performed by the publish/subscribe infrastructure, which interprets the provided subscription and detects complete or incomplete event sequences, which are identified individually. These sequences might be summarized before being published to the QA and developer teams. In this phase, consolidation consists of identifying repetitive sequences and summarizing this information (e.g. a given sequence was identified x times), this activity can be supported by the publish/subscribe client, by the event notification server, or by a different application running at the QA and developer teams' site. Our approach allows such flexibility so that one can decide where in the distributed system the correlation is going to take place. This decision depends on factors such as the application being deployed, the sequence of event being monitored, the frequency of occurrence of such sequences, the overhead caused by monitoring the sequence of events, and so on.

The mapping process generates subscriptions to the set of events in a UML sequence diagram. These events may be generated, as specified, in the proper order, or in unexpected orders. Unexpected sequences are indicative of different sorts of problems: (1) The sequence is actually part of the overall system specification, but the developer did not include it in the partial specification to be monitored; (2) The sequence is not part of the overall system specification, so either the specification is incomplete and the system is still working properly, or (3) an erroneous behavior was actually identified. In either case, once an unexpected sequence of events is detected, the detailed sequence should be sent back to the developers for analysis. When a valid, expected, sequence is detected, the occurrence of such sequence is published for further analysis.

Consolidation, Analysis and Visualization

Since our approach uses the software tomography strategy (see section 2) to deploy monitoring configurations to different software instances in the client sites, this information needs to be further consolidated for the sake of analysis. In so doing, aggregated information about the effectiveness of the residual testing on multiply deployed systems yields more significant system usage profiles and measure of test adequacy than obtained during pre-deployment testing (during development) and also more than could be obtained from a single deployment. This consolidation and its respective interpretation provide feedback to the developers about system usage, as well as an insight into adjustments that can be made to the monitoring configuration. Consolidation is done according to the developer's focus of interest, for example, one can consider data gathered from all executions on all software instances but for a specific usage profile group (e.g. groups that exten-

sively use a specific kind of feature such as GUI, security, or database).

For the visualization, we use gauges and graphics to present the results obtained from the consolidation. This can be performed, at runtime, as the data is being received and consolidated, or afterwards, using local stored execution tracing information.

3.4 Implementation Issues

The proposed infrastructure is being implemented in Java due to its run-time characteristics, which facilitates the dynamic change of the monitoring configuration through different approaches being studied.

The first step of our approach is the *mapping* from UML sequence diagrams to monitoring configurations, which contain the sequence of methods to be monitored in the software instances. At this time, we are studying the use of off-the-shelf software development tools, such as Rational Rose, which perform forward engineering. In this case, the methods in the source-code generated from the sequence diagram will be the same methods to be monitored. Since we are using an in-house event-notification server, YANCESS [12], a tool will be developed to translate sequence diagrams to subscriptions in this infrastructure.

Currently, we are evaluating a strategy, based on a modified JVM class loader [13]. The loader controls the object activation and destruction, and intercepts method invocations and exception handlings. It can be configured to instrument specific methods to publish events whenever they are invoked. Class files are instrumented at load time allowing the change of the monitoring configurations at runtime. For example, the loader can be configured to monitor classes or methods in a specific java package, or to generate events only during object creation and destruction. We believe that, by using this modified class loader, we can minimize performance degradation, by dynamically instrumenting particular methods from a sequence diagram.

The *consolidation* phase is supported by the event correlation features provided by the event notification server YANCEES [12], which was chosen for its extensibility and configurability. This server can be configured with client and server-side plug-ins to process events. In particular, event correlation plug-ins are available, which allow event sequence detection and abstraction in both client and server sides.

Finally, for *visualization and analysis* tasks we are studying the use of execution time visualization tools such as the one proposed by Gogolla and Richters [14] which supports UML sequence diagrams visualization or off-line analysis tools such as used by EDEM [9], which summarizes the data in form of graphics and histograms.

4 Conclusions and Future Work

This paper presents an approach which extends ex-

isting residual testing techniques and tools such as expectation-driven monitoring, software tomography, dynamic program instrumentation, and publish/subscribe systems, to provide behavioral conformance checking. In our approach, deployed software is tested with respect to its conformance to dynamic high-level models such as UML sequence diagrams. The approach uses the facilities of a versatile publish/subscribe service that supports data collection, sequence detection and filtering. The software tomography strategy is used to disseminate monitoring configurations to specific clients coping with scalability. End-user usage profiles and hardware and software configurations are collected and used to support the software tomography task distribution. Expectation-driven monitoring is used to monitor execution patterns in the deployed software. Dynamic program instrumentation is used to provide flexibility to the approach, allowing the QA team to refine and deploy new tasks to the programs in the software on a client site. We are currently completing the implementation and integration of the different components of the system (notification server, the JVM monitor, and the UML sequence diagram specification parser).

As future work, we intend to validate our approach on applications with real users. We also plan to experiment with other behavioral specification models, and with residual testing in a coarser-grained level of abstraction. Finally, one important concern in our research is avoiding performance degradation during the monitoring of software instances. We plan to evaluate our approach in different scenarios to understand the trade-off between performing on-line vs. off-line sequence matching in the client side, and between executing the sequence detection in the server vs. client sites.

5 Acknowledgments

The authors thank CAPES (grant BEX 1312/99-5) for financial support. This effort was also sponsored in part by the National Science Foundation under grant numbers ITR-0083099, CCR-0205724, ITR-0326105. Funding was also provided by Conexant Systems, Inc. and the Intel Corporation. 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 any of the sponsoring agencies.

6 References

- [1] M. J. Harrold, "Testing: a roadmap," International Conference on Software Engineering, Future of Software Engineering Track, pp. 61-72, Limerick, Ireland, 2000.
- [2] C. Pavlopoulou and M. Young, "Residual Test Coverage Monitoring," International conference on Software engineering, Los Angeles, California, United States, 1999.
- [3] B. Liblit, et al. "Sampling User Executions for Bug Isolation," Workshop on Remote Analysis and Measurement of Software Systems (RAMSS'02), Portland, Oregon, 2003.
- [4] A. Orso, et al, "Gamma System: Continuous Evolution of Software after Deployment," ACM International Symposium on Software Testing and Analysis, Rome, Italy, 2002.
- [5] L. Naslavsky, et al, "Multiply-Deployed Residual Testing at the Object Level," IASTED International Conference on Software Engineering (SE2004), Innsbruck, Austria, 2004.
- [6] C. Yilmaz, et al, "Distributed Continuous Quality Assurance: The Skoll Project," Workshop on Remote Analysis and Measurement of Software Systems (RAMSS'02), Portland, Oregon, 2003.
- [7] D. J. Richardson, et al, "Specification-Based Test Oracles for Reactive Systems," 14th International Conference on Software Engineering, Melbourne, Australia, 1992.
- [8] L. Baresi and M. Young, "Test Oracles," University of Oregon, Department of Computer and Information Science CIS-TR-01-02, August 2001.
- [9] D. Hilbert and D. Redmiles, "An Approach to Large-scale Collection of Application Usage Data over the Internet," 20th International Conference on Software Engineering, Kyoto, Japan, 1998.
- [10] G. Liu, et al, "Composite Events for Network Event Correlation," presented at IFIP Symposium on Integrated Management (IM'99), Boston, MA, 1999.
- [11] P. T. Eugster, et al, "The many faces of publish / subscribe" *ACM Computing Surveys (CSUR)*, vol. 35, pp. 114 - 131, 2003.
- [12] R. S. Silva-Filho, et al, "The Design of a Configurable, Extensible and Dynamic Notification Service," Second International Workshop on Distributed Event-Based Systems, with The ACM SIGMOD/PODS Conference, San Diego, CA, 2003.
- [13] P. Dourish and J. Bytner, "A Visual Virtual Machine for Java Programs: Exploration and Early Experiences.," ICDMS Workshop on Visual Computing, Redwood City, CA, 2002.
- [14] M. Gogolla and M. Richters, "Aspect-Oriented Monitoring of UML and OCL Constraints," AOSD Modeling With UML Workshop, San Francisco, CA, USA, 2003.