

Defect Density Estimation Through Verification and Validation

Mark Sherriff and Laurie Williams
North Carolina State University
Raleigh, NC, USA 27695
{mssherr, lawilli3}@ncsu.edu

Abstract

In industry, information on defect density of a product tends to become available too late in the software development process to affordably guide corrective actions. Our research objective is to build a parametric model which utilizes a persistent record of the validation and verification (V&V) practices used with a program to estimate the defect density of that program. The persistent record of the V&V practices are recorded as certificates which are automatically recorded and maintained with the code. To date, we have created a parametric modeling process with the help of the Center for Software Engineering at the University of Southern California and have created the second version of an Eclipse plug-in for recording V&V certificates.

1. Introduction

In industry, post-release defect density of a software system cannot be measured until the system has been put into production and has been used extensively by end users. The defect density of a software system is calculated by measuring the number of failures divided by the size of the system, using a size measure such as lines of code. Actual post-release defect density information becomes available too late in the software lifecycle to affordably guide corrective actions to software quality. Correcting software defects is significantly more expensive when the defects are discovered by an end user compared with earlier in the development process [4].

Because of this increasing cost of correcting defects, software developers can benefit from early estimates of the defect density of their product. If developers can be presented with this defect density information during the development process in the environment where they are creating the system, more affordable corrective action can be taken to rectify defect density concerns as they appear.

A development team will use several different methods to ensure that a system is of high-assurance [16]. However, the verification and validation (V&V) practices used to make a system reliable might not always be documented or this documentation may not be maintained. This lack of documentation can hinder other developers from knowing what V&V practices have been performed on a given section of code. Further, if code is being reused from an earlier project or code base, developers might spend extra time re-verifying a section of code that has already been verified thoroughly.

Research has shown that parametric models [6] using software metrics can be an effective means to predict product quality. *Our research objective is to build a parametric model which utilizes a persistent record of the V&V practices used with a program to estimate the defect density of that program.* To accomplish this objective, we are developing a method called Defect Estimation with V&V Certificates on Programming (DevCOP). This method includes: a mechanism for creating a persistent record of V&V practices as *certificates* stored with the code base; a parametric model to provide an estimate of defect density; and tool support to make this method accessible for developers. A DevCOP certificate is used to track and maintain the relationship between code and the evidence of the V&V technique used. We will build the parametric model using a nine-step systematic methodology for building software engineering parametric models [3], which has been used to build other successful parametric models [2, 11, 14].

In this paper, we describe our current work in developing and validating the DevCOP parametric model and the DevCOP Eclipse plug-in to support the creation and maintenance of DevCOP certificates. Section 2 presents background research. Section 3 provides an overview of the DevCOP model, while section 4 describes its limitations and section 5 discusses tool support through the DevCOP Eclipse plug-in. Finally, section 6 describes our conclusions and future work.

2. Background

In this section, we will discuss the relevant background work and methodologies used during our research. It includes descriptions of research regarding other metric-based defect density estimation, V&V techniques, and parametric modeling in software engineering.

2.1 Parametric Modeling

Parametric models relate dependent variables to one or more independent variables based on statistical relationships to provide an estimate of the dependent variable with regards to previous data [6]. The general purpose of creating a parametric model in software engineering is to help provide an estimated answer to a software development question early in the process so that development efforts can be directed accordingly. The software development question could relate to what the costs are in creating a piece of software, how reliable a system will be, or any number of other topics.

Parametric modeling has been recognized by industry and government as an effective means to provide an estimate for project cost and software reliability. The Department of Defense, along with the International Society of Parametric Analysts, acknowledges the benefit of using parametric analysis, and encourages their use when creating proposals for the government [6]. The Department of Defense claims that parametric modeling has reduced government costs and also improved proposal evaluation time [6]. Boehm developed the Constructive Cost Model (COCOMO) [2] to estimate project cost, resources, and schedule. Further, the Constructive Quality Model (COQUALMO) added defect introduction and defect removal parameters to the COCOMO to help predict potential defect density in a system. Nagappan [11] created a parametric model with his Software Testing Reliability Early Warning (STREW) metric suite to create an estimate of failure density based on a set of software testing metrics.

2.2 Parametric Modeling Process for Software Engineering

During our research, we are building a parametric model to estimate defect density based upon V&V certificates recorded with the code. To help facilitate this process, we worked together with the Center for Software Engineering at the University of Southern California [3] to create a parametric modeling process specifically for software engineering research [13]. This process, illustrated in Figure 1, shows the steps that can be followed to create an effective parametric model. More information on the individual steps can be found in [13].

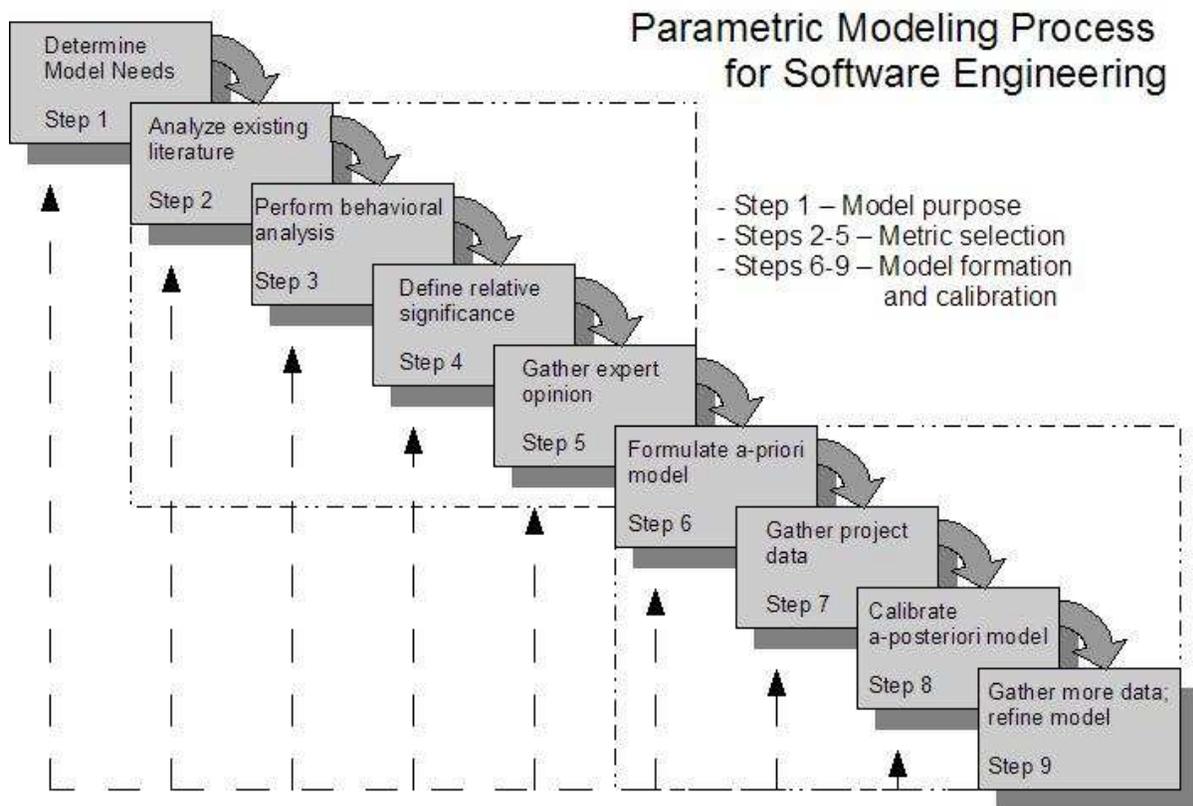


Figure 1. Parametric Modeling Process for Software Engineering

2.3 Verification and Validation Techniques

During the creation of software, a development team can employ various V&V practices to improve the quality of the software [1]. For example, different forms of software testing could be used to validate and verify various parts of a system under development. Sections of code can be written such that they can be automatically proven correct via an external

theorem prover [16]. A section of a program that can be logically or mathematically proven correct could be considered more reliable than a section that has “just” been tested for correctness.

Other V&V practices and techniques require more manual intervention and facilitation. For instance, formal code inspections [5] are often used by development teams to evaluate, review, and confirm that a section of code has been written properly and works correctly. Pair programmers [17] benefit from having another person review the code as it is written. Some code might also be based on technical documentation or algorithms that have been previously published, such as white papers, algorithms, or departmental technical reports. These manual practices, while they might not be as reliable as more automatic practices due to the higher likelihood of human error, still provide valuable input on the reliability of a system. Different V&V techniques will provide a different level of assurance as to how reliable a section of code is.

The extent of V&V practices used in a development effort can provide information about the estimated defect density of the software prior to product release. The Programatica team at the Oregon Graduate Institute at the Oregon Health and Science University (OGI/OHSU) and Portland State University (PSU) has developed a method for high-assurance software development [7, 16]. Programmers can create different types of *certificates* on sections of code based on the V&V technique used by the development on that section of the code. Certificates are used to track and maintain the relationship between code and the evidence of the V&V technique used. Programatica certificates are provided and maintained by *certificate servers* [7]. Each certificate server provides a different type of V&V evidence that is specialized for that server. For example, an external formal theorem prover called Alfa can provide a formal proof certificate for methods in a Haskell system [16]. These certificates are used as evidence that V&V techniques were used to make a high-assurance system [16]. We propose an extension of Programatica’s certificates for defect density estimation whereby the estimate is based upon the effectiveness of the V&V practice (or lack thereof) used in code modules.

2.4 Metrics to Predict Defect Density

Operational profiles have been shown to be effective tools to guide testing and help ensure that a system is reliable [9]. An operational profile is “the set of operations [available in a system] and their probabilities of occurrence” as used by a customer in the normal use of the system [10]. However, operational profiles are perceived to add overhead to the software development process as the development team must define and maintain the set of operations and their probabilities of occurrence. Rivers and Vouk recognized that operational profile testing is not always performed when modern constraints on market and cost-driven constraints are introduced [12]. They performed research on evaluating non-operational testing and found that there is a positive correlation between field quality and testing efficiency. Testing efficiency describes the potential for a given test case to find faults at a given point during testing. We will utilize non-operational methods in predicting software defect density to minimize developer overhead as much as possible.

Nagappan [11] performed research on estimating failure density without operational profiles by calibrating a parametric model which uses in-process, static unit test metrics. This estimation provides early feedback to developers so that they can increase the testing effort, if necessary, to provide added confidence in the software. The STREW metric suite consists of static measures of the automated unit test suite and of some structural aspects of the implementation code. A two-phase structured experiment was carried out on 22 projects from junior/senior-level software engineering students from the fall of 2003 [11], which helped to refine the STREW-J metric suite. The refined suite was then used 27 open source Java projects found on SourceForge¹, an open-source development website, and five projects from a company in the United States [11]. The research from these case studies indicates that the STREW-J metrics can provide a means for estimating software reliability when testing reveals no failures.

Another version of the STREW metric suite was developed specifically for the Haskell programming language, STREW-H. STREW-H was similarly built and verified using case studies from open-source and industry. An open-source case study [15] provided guidance to refine the metric suite for its use on an industry project with Galois Connections, Inc [14]. These two STREW projects demonstrate that in-process metrics can be used as an early indicator of software defect density.

3. Defect Estimation from V&V Certification on Programming

We are actively developing a parametric model which uses non-operational metrics to estimate defect density based upon records of which V&V practices were performed on sections of code. We are integrating our estimation directly into the development cycle so that corrective action to reduce defect density can take place early in the development process. We call this method the Defect Estimation from V&V Certification on Programming (DevCOP) method. A V&V certificate contains information on the V&V technique that was used to establish the certificate. Different V&V techniques will provide a different level of assurance as to how reliable a section of code is. For example, a desk check of code would be, in general, less effective than a formal proof of the same code.

¹ <http://www.sf.net/>

To build and verify our parametric model of our DevCOP method, we have adapted and refined a nine-step modeling methodology to create a parametric modeling process for software engineering [13], as described in Section 2.2. The first step in the process is to define the goal of the parametric model. *The goal of the model is to provide an estimate of defect density based on V&V certificates and the coefficient weights.* We anticipate that a model would need to be developed for each programming language we would study. Our current work involves the Java (object-oriented) and Haskell (functional) languages.

The second step of the parametric modeling process is to analyze existing literature to determine categories of V&V techniques and empirical findings on the defect removal efficacy of each V&V practice. Balci categorized V&V techniques with some regard to their general effectiveness as to finding defects in a system [1]. For the purposes of our scale, we began with Balci's categorization of the V&V techniques:

- Manual – includes all manual checking, such as pair programming [17] and code inspections [5];
- Static – includes automatic checking of code before run-time, such as syntax and static analysis;
- Dynamic – includes all automatic checking that takes place during execution, such as black-box testing;
- Formal – includes all strictly mathematical forms of checking, such as lambda calculus and formal proofs [16].

Assigning proper relative significance to certificates to place them on a single scale of relative effectiveness is a significant challenge in our research. Each of these V&V categories provides different evidence as to how reliable a system is [1]. For example, static V&V techniques can provide information as to whether the structure of the code is correct, while manual V&V techniques can provide information about both the structure of the code and if the code is providing the functionality requested by the customer(s). We will perform a causal analysis with our industry partners on our initial data to help build our V&V rubric. The causal analysis will provide us with more information about the efficacy of certain techniques under particular circumstances.

We must also determine the proper granularity for the model. Code certificates could potentially be associated with modules, classes, functions, or individual lines of code. Each level of granularity offers potentially different information about the defect density of the system, and also different challenges in gathering data. Currently, we are analyzing certificates at the function level and are involved in on-going analysis with our industry partners on this decision.

We envision the defect density parametric model to take the form of Equation 1. For each certificate type, we would sum the product of a size measure (number of functions/methods) and a coefficient produced via regression analysis of historical data. The calibration step of the regression analysis would yield the constant factor (a) and a coefficient weighting (c_j) for each certificate type, indicating the importance of a given V&V technique to an organization's development process.

$$DefectDensity = a + \sum_{j=1}^{certificate_type} (c_j * Size_j) \quad (1)$$

We are working with industry partners to gather expert opinion and our initial data sets to build and verify our early parametric model. Developers on a small Java team using Eclipse are recording their V&V efforts using the DevCOP plug-in (described in Section 6) as the project progresses. During defect removal and bug fixes, the team will also record these efforts as a different type of certificate. Proceeding through steps 6-9 of the parametric modeling methodology will require a significant number of projects for each language we work with. Also, we are using the DevCOP plug-in to record certificates on the DevCOP plugin code base as development progresses, along with other student projects.

4. Limitations

In the creation of certificates, we are not assigning more importance to certain functions or sections of code over others, as is done with operational profile means of estimation. Nor are we using the severity of defects detected to affect the importance of some certificates over another. While this level of granularity could be beneficial, one of our initial goal's is to make this method easy to use during development, and at this time, we think that adding this level of information could be a hindrance. Another granularity limitation is the granularity of certificates. Based on the Programatica Team's work and expert opinion, it was decided that functions would be the proper level of granularity for certificates. As previously discussed, we recognize that being able to record certificates at a line of code level could be beneficial, but method-level recording seemed to be the best course of action for the initial validation of the methodology.

5. Tool Support

We have automated the DevCOP method as an Eclipse² plug-in to facilitate the creation and maintenance of V&V certificates with as little additional overhead for developers as possible. Ease of use, along with the added benefit of being able to

² For more information, go to <http://www.eclipse.org/>.

calculate V&V and defect information with a defect density estimate, should make the DevCOP method practical for practicing engineers. The initial DevCOP Eclipse plug-in was released in Spring 2005 as a beta version to several Java development teams for evaluation. The plug-in allows developers to create certificates during the development process within the integrated development environment (IDE) so that this information can be utilized throughout the code's lifetime. Anecdotal reports from the teams indicated that the initial version of the plug-in was not robust enough yet to warrant inclusion into their development cycle due to concerns about distributed code development along with tighter integration with the Eclipse IDE. Further, the developers indicated that detailed metrics about the coverage of the different V&V techniques (e.g. how much of the code was pair programmed vs. solo programmed) would be extremely useful in their development efforts.

To address the concerns of our test Java developers, we have released the second version of an DevCOP Eclipse plug-in to handle the creation and management of V&V certificates during the development process³ [15]. Figure 2 shows a screenshot of the Eclipse plug-in for recording V&V certificates. Major changes in this version of the plug-in include external database support for certificate storage, tighter integration with Eclipse's refactoring plug-in, and V&V certificate coverage metrics.

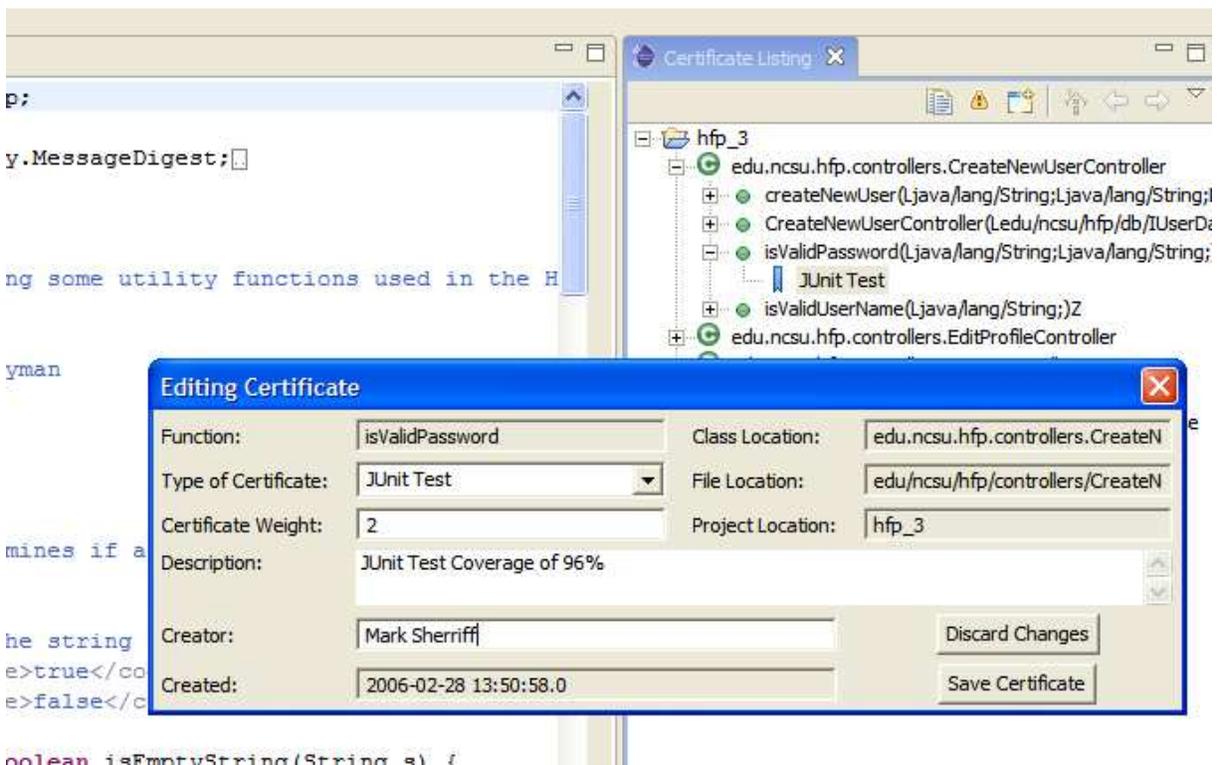


Figure 2. Screenshot of the DevCOP Eclipse plug-in for recording V&V certificates.

Our objective is to make the certificate creation process as easy and transparent as possible, and will continue to improve it in later iterations as we receive more developer feedback. The simplest way to create a certificate for a function is to right-click on a function or set of functions in Eclipse's Package Explorer or Outline view. The developer can then add in descriptive information as to what V&V technique was used to certify that function. The DevCOP plugin then records a hash of the function's abstract syntax tree (AST) at the time of certification. As long as the certificate's hash matches that of the current code base, the certificate is considered valid. If, however, the code is modified in such a way that its AST changes, the certificate is considered invalid. The developer can revalidate the certificate if he or she feels that the V&V technique still applies after the code base change.

Developers, however, may wish to add multiple certificates at once, across several files or even projects, such as during a long pair programming session or code review. The way we accomplish this method of certificate creation is through what

³ The plug-in is available at <http://agile.csc.ncsu.edu/devcop/>.

we call active certificates. An *active certificate* is a means by which Eclipse will automatically identify changed code during a programming session to be certified by the developer.

For example, if two programmers were about to start pair programming on a piece of code, they would click the active certificate button before they began. Eclipse would then actively record non-trivial changes to the system (i.e. changes to the abstract syntax tree of the code, not commenting or formatting changes) and will present the affected functions to the developers for certification at the end of the pair programming session. The concept of active certificates can extend to several different types of V&V activity, such as code inspections or bug fixes. Active certificates allow developers to write or modify code normally, without increasing their work overhead.

We have also begun incorporating external V&V tools and plug-ins into the DevCOP plug-in to utilize their ability to verify code. The first external tool that has been successfully integrated is the jcoverage GPL library⁴. Using the jcoverage library, the DevCOP plug-in can gather unit test case code coverage to use in creating certificates. Developers can choose the code coverage threshold needed in order for a function to be considered certified by their test case runs. We will incorporate other tools, such as ESC/Java2⁵ and other static code checkers.

6. Conclusions and Future Work

We have created and are validating a method for managing and leveraging the effort put into V&V by a development team to provide an estimate of software defect density in-process. Due to the high costs of fixing software defects once a product has reached the field, any information that can be provided to developers in-process and can give an indication of software defect density is invaluable. If corrective actions can be taken earlier in the software development life cycle to isolate and repair software defects, overall maintenance costs can decrease.

The DevCOP method that we are developing will help with this problem in several ways. First, after a set of certificates has been created, an overall estimate of defect density can be created based on the V&V weightings using a parametric model. Research has shown that parametric models using software metrics can be an effective method for predicting defect density [11]. We are gathering data from different industrial programming teams to calibrate our method to the general case.

DevCOP also allows developers to manage the effort that is put into V&V in a place where all developers can see what measures have been taken to ensure a piece of code is reliable and to treat it accordingly. The DevCOP method assists developers in identifying and analyzing sections of code that have not yet been certified, or to concentrate their efforts on a particularly critical section of code.

In addition to providing a defect density estimate, DevCOP information can be used to provide a V&V history for particular code segments. Development teams can see what efforts were used to verify the code, even if a different team was working on the system or if poor documentation was available. If the code is found to be error-prone, the certificate information can provide guidance as to what techniques might need to be improved in the organization. During system maintenance, certificates can be referenced to see what types of V&V techniques were performed on a given section of code. If the code is found to be trustworthy, the certificate information with this code could provide evidence that this code is reliable for reuse in future systems.

Another potential use for this V&V information is to build the certificates into the compiled program itself, allowing it to be referenced at runtime by other systems. One possible way of including certificate information with a system is to instrument the code with the certificates, thus storing this V&V effort in a manner that can be reference at runtime. This stored certificate information could prove to be useful information for numerous types of systems, from trust management for personal computers to systems that require load balancing or job distribution. Systems that have information that shows that effort was put in to make the system reliable could receive a greater share of distributed jobs. This potential use of DevCOP is similar to that of security certificates over the Internet [8], which show that effort has been put forth to ensure that the connection or website is secure. The V&V information could also be made available at the time that the system is delivered, to help show the customer that the proper techniques were used to ensure the quality of the product.

7. References

- [1] Balci, O., "Verification, Validation, and Accreditation of Simulation Models," Winter Simulation Conference, 1997, pp. 125-141.
- [2] Boehm, B. W., Horowitz, E., Madachy, R., Reifer, D., Clark, B., Steece, B., Brown, A. W., Chulani, S., and Abts, C., *Software Cost Estimation with COCOMO II*. Upper Saddle River, NJ: Prentice Hall, 2000.
- [3] Boehm, B. W., "Building Parametric Models," International Advanced School of Empirical Software Engineering, Rome, Italy, September 29, 2003.

⁴ <http://www.jcoverage.com/downloads.html>

⁵ <http://secure.ucd.ie/products/opensource/ESCJava2/>

- [4] Dolbec, J. and Shepard, T., "A Component Based Software Reliability Model," Conference of the Centre for Advanced Studies on C, 1995.
- [5] Fagan, M., "Design & Code Inspections to Reduce Errors in Program Development," *IBM Systems Journal*, vol. 15, no. 3, pp. 182-211, 1979.
- [6] International Society of Parametric Analysts, "Parametric Estimating Handbook," Department of Defense, Online Handbook, April 16, 2005. Available online: http://www.ispa-cost.org/PEIWeb/Third_edition/newbook.htm
- [7] Jones, M., "Evidence Management in Programatica," Workshop on Software Certificate Management, Palm Beach, California, 2005.
- [8] Kent, S., "Evaluating certification authority security," IEEE Aerospace Conference, Aspen, CO, March 21-28, 1998, pp. 319-327.
- [9] Musa, J., "Theory of Software Reliability and its Applications," *IEEE Transactions on Software Engineering*, pp. 312-327, 1975.
- [10] Musa, J., *Software Reliability Engineering*: McGraw-Hill, 1998.
- [11] Nagappan, N., "A Software Testing and Reliability Early Warning (STREW) Metric Suite," PhD Dissertation, North Carolina State University, 2005.
- [12] Rivers, A. T., Vouk, M.A., "Resource-Constrained Non-Operational Testing of Software," International Symposium on Software Reliability Engineering, Paderborn, Germany, 1998, pp. 154-163.
- [13] Sherriff, M., Boehm, B. W., Williams, L., and Nagappan, N., "An Empirical Process for Building and Validating Software Engineering Parametric Models," North Carolina State University CSC-TR-2005-45, October 19 2005.
- [14] Sherriff, M., Nagappan, N., Williams, L., and Vouk, M. A., "Early Estimation of Defect Density Using an In-Process Haskell Metrics Model," First International Workshop on Advances in Model-Based Software Testing, St. Louis, MO, May 15-21, 2005.
- [15] Sherriff, M., Williams, L., "Tool Support For Estimating Software Reliability in Haskell Programs," Student Paper, IEEE International Symposium on Software Reliability Engineering, St. Malo, France, 2004, pp. 61-62.
- [16] The Programatica Team, "Programatica Tools for Certifiable, Auditable Development of High-Assurance Systems in Haskell," High Confidence Software and Systems, Baltimore, MD, 2003.
- [17] Williams, L. and Kessler, R., *Pair Programming Illuminated*. Boston: Addison-Wesley, 2002.