

Understanding Semantic Impact of Source Code Changes: an Empirical Study

Danhua Shao, Sarfraz Khurshid, and Dewayne E. Perry
Electrical and Computer Engineering, The University of Texas at Austin
{dshao, khurshid, perry}@ece.utexas.edu

Abstract

Since source code is the ultimate definition of the behavior of a software product, changes to source code become the critical factor in understanding behavioral changes and predicting faults. In studies on source code changes, text or syntactic approaches have been widely used. Textual analysis focuses on changed text fragments while syntactic analysis focuses on changed syntactic entities. Although both of them have demonstrated their advantages in experimental results, they have only focused on changed code. Because of semantic dependencies within programs, we believe that code impacted by changes is also helpful. Given a source code change, we identify its impact by program slicing along the variable def-use chains. To evaluate the effectiveness of change impacts in fault detection and prediction, we compare impacted code with changed code according to size and fault density. Our experiment on the change history of a successful industrial project shows that: for large changes, their impacts have relative small size and high fault density; while for small changes, change themselves have relative small size and high fault density. Our study suggests that: 1) change impacts are complementary to change themselves in detecting or predicting faults; 2) within the impact of a large change, a high degree of interference between impacts of different changed lines contributes to the high fault density; 3) high fault density in impacts aggravates the danger of large changes.

1. Introduction

During a software development process, change is a pervasive phenomenon. Changes happen in almost every phase: requirements, design, implementation, and maintenance. Changes in any phase will affect the behavior of the delivered software product. Thus, changes are significantly related to software product quality and may be used to predict faults.

Since source code is the ultimate definition of the behavior of a software product, changes in source code become an important factor in understanding

changes in behaviors and detecting, or predicting, faults.

In studies based on source code changes, textual or syntactic approaches have been widely used. Textual analysis treats the source code as just a piece of text and builds change rules according to the textual properties of changes. For example, [9] [16] [17] analyze how many lines added or deleted in a change and associate change size with faults. Syntactic analysis uses syntactic entities, such as class/method/field in [27], the number of incoming and outgoing function calls in [13], and the association or sequence of function calls in [3] [12]. They extract change rules or patterns on syntactic entities and relate them to faults.

Although the advantages of textual and syntactic approaches have been demonstrated in experimental results, both of them only focus on the code modified during changes. Neither of them considers the internal dependencies within a program. For example, changes to a statement assigning a value to a variable might impact subsequent statements using the value of that variable. We believe the impacted statements are also helpful to detect or predict faults.

To evaluate the efficiency of change impact in fault detection and prediction, we conduct a local data-flow analysis of the change history of a successful industry project, the 5ESS system. For each sampled change, we identify its impact by program slicing along the variable def-use chains. We group changes according to change size and compare the size and fault density of change impacts and changes themselves across groups. In this study, we want to address three questions: for small changes and large changes,

- 1) Is there any difference on the size of their impacts?
- 2) Is there any difference on the ratio of impact size to change size?
- 3) Is there any difference on the faults densities in change impacts and in change themselves?

Section 2 gives an overview of the identification of semantic change impact by local data-flow analysis. The context for this study is discussed in Section 3. Section 4 presents the experimental design and its results. Validity issues are discussed in Section 5. Related researches are introduced in Section 6. Finally, we summarize our study and propose future work in Section 7.

2. Semantic Impacts of Source Code Changes

To identify impacts of source code changes, we combine local data-flow analysis and program slicing. Local data-flow analysis discloses the internal dependencies between semantic structures in a program while program slicing identifies semantic structures impacted by changes. Figure 1 illustrates our semantic analysis of change impact.

Given two versions in a change, we first calculate the control flow graph for each version. From each control flow graph, we identify the dependencies based on variable def-use pairs. Each identified local data-flow dependency is represented as a triple (*var*: *def*, *use*), where *var* is the variable on which the dependence is built, *def* is the line that defines variable *var*, and the *use* line uses the variable defined at *def* line. The dependences in version v1 are {(a: 1, 3), (b: 2, 4), (i: 3, 5), (j: 4, 5)}.

With the variable def-use dependency analysis on the two versions, the semantic impact of a change can be calculated by forward slicing from the changed lines. In this example, the change from v1 to v2 modified Line 1 from “a = 0” to “a = 1”. According to the variable def-use chains, {(a: 1, 3), (i: 3, 5)}, Line 3 and 5 are impacted. So the semantic impact of change v1→v2 is {3, 5}.

3. Study Context

The data in this study come from the change and version management history of the Office

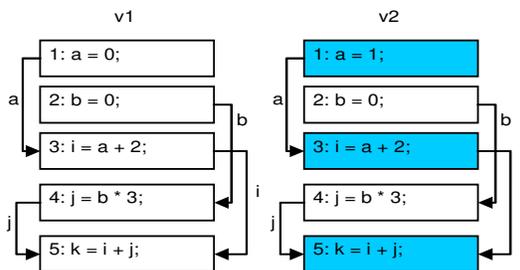


Figure 1. Semantic analysis on change: v1→v2

Automation (OA) subsystem in 5ESS™, a large-scale successful industrial project.

The 5ESS™ change history has been widely used for various studies, such as inferring change effort from configuration management database [7], automatically extracting reasons for changes [15], and studying the semantic impact of parallel changes and small changes [19] [20]. The widely used data source can increase the content validity of our study.

In the change and version management repositories of 5ESS™, two kinds of artifacts are related to our study: Modification Request (MR) and delta. A *delta* records the actual lines added or deleted during a change on a file. A MR is a group of deltas made by a developer for a purpose such as adding a new function or fixing a bug.

In 5ESS™, all deltas were recorded in SCCS [22] version management system. Information about MRs is kept in ECMS [24] database. For each MR, ECMS records the data such as the date it was opened and closed, its developer, and a short text abstract of the change purpose. For each delta, the ECMS database also recorded its date, the developer, and the MR it belongs to.

MRs can be classified according to their purpose. Bug-fixing MRs can be used to identify faults. The classification used in this study is derived from the work by Mockus and Votta[15]. They classified MRs according to keywords in the textual abstract of change purposes. For example, if keywords like “fix,” “bug,” “error,” and “fail” were present, the change was classified as corrective. There are four classes of MRs:

- Inspection changes: found in code inspections.
- Perfective changes: changes to improve performance, make a piece of code more maintainable, or generally improve the quality of the code.
- Adaptive changes: changes to add new features and functionality into the system.
- Corrective change: changes to fix faults.

Given a change on source code, or a *delta*, its class is the class of the MR it belongs to. For example, all the changes belong to a Corrective MR are Corrective changes.

In this study, data dependency calculation and program slicing is based on GrammaTech’s CodeSurfer [5]. The C compiler is Visual C++ 6.0. For the variables used in def-use analysis, we count both pointer variables and non-pointer variables. For

pointer analysis, we select CodeSurfer’s option that distinguishes individual fields in a referenced structure and use the pointer analysis offered by Codesurfer to estimate point-to sets.

4. Study and Results

To address the three questions mentioned before, we propose three hypotheses:

- 1) H1: Large changes impact more lines than small changes;
- 2) H2: Large changes have a lower impact size to change size ratio than small changes;
- 3) H3: For fault density of impacted code, large changes are higher than small changes; but for fault density of changed code, large changes are lower than small changes.

4.1. Constructs

Evaluating these three hypotheses needs to define constructs to represent *change*, *change impact*, and *fault*:

Change: Change has different meanings in different contexts. In this study, change is represented as the different lines between the old version and the new version. Given a change, *change size* is the number of the lines added, deleted, and modified in the change. But, SCCS system only records the number of lines inserted or deleted during a change. Modifications to the existing lines are tracked as old lines being replaced by new lines (delete and insert). We use the maximum of added and deleted lines as change size.

Change impact: Given a change, its impact is represented as the lines that have local data dependencies on changed lines. According to the variable def-use chains, the impact of a change is identified by forward slicing from the changed lines. *Impact size* is the number of lines in a change impact.

Fault: A fault is the collection of all lines changed in a Corrective type change.

4.2. Studies

In this study, 273 changes were sampled from four modules in 5ESS change repository. The size of each version is around 1000 lines. All the changes are grouped according to change size. To evaluate the three hypotheses, we performed three studies to compare these groups according to following three metrics: impact size, ratio of impact size to change size, and fault density.

Impact size

According to the sampled changes and their impacts, Figure 2 shows the distribution of impact size for each group. The x-axis is change size and the y-axis is impact size. For each group, the vertical line is the box-plot representation of the distribution of impact size within that change-size group. On each vertical line, the box shows the central 50% of the distribution. The upper and lower ends of the box mark the upper and lower quartiles. The bold dot shows the mean of the distribution. The solid horizontal line in the box shows the median of the distribution. The dashed vertical lines attached to the box indicate the tails of the distribution. The vertical lines above and below the box are the maximum and minimum. This figure shows that, with the increase of the change size, the change impact size also increases accordingly. For small changes that modified less than 5 lines, the median and mean of their impact size are 0 and 2.587413. For large changes modified more than 100 lines, the median and mean of their impact increase to 42 and 60.5.

For all sampled changes, statistical analysis on their change size and impact size shows that they are strongly correlated. Their correlation coefficient is $R = 0.75$, with a very high significance $p \leq 3.24e-48$. In the t-test for the significance of the coefficient, $t = 18.76$ with *Degree of Freedom* = 271.

The results from statistic analysis and Figure 2 support hypothesis H1: large changes impact more lines than small changes.

This result supports our previous study on small changes [20]. In the change history of a subsystem of 5ESS, less than 4% of small (one-line) changes result in faults. According to the result from H1, impacts of small changes are also small. Because there is little effect on the semantics of the program,

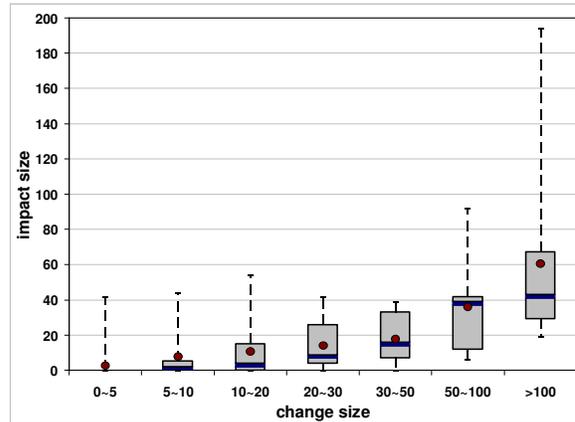


Figure 2. Distribution of impact size.

small changes are not likely to introduce faults.

Ratio of impact size to change size

For each group, we also calculate the ratio of impact size to change size. Figure 3 illustrates the results. The x-axis is the change size and the y-axis is the ratio of impact size to change size. The x and y coordinate of each point represents the average of change size and ratio for each group. The trend shows that, with the increase of change size, the ratio goes down quickly. For changes less than 10 lines, their impact size is larger than change size. For changes around 70 lines, the impact size is only half the change size. Statistic analysis shows a nearly linear relation between change size (on a logarithmic scale) and the ratio of impact size to change size. Their *correlation coefficient* is $R = -0.951$, $p \leq 0.00103$. In the t-test for the significance of the coefficient, $t = 6.874$ with *Degree of Freedom* = 5

The results of our statistical analysis and Figure 3 support hypothesis H2: Large changes have lower impact size to change size ratio than small changes. The relative small impact size suggests that, to detect or predict faults in large changes, checking change impact is cheaper than checking change itself.

This result also discloses the *internal interference* phenomena of large changes. For a large change, the impacts of different changed lines highly overlap with each other. In a program, a change on one line usually touches at least one variable def-use chain. With the increase of change size, more and more variable def-use chains are touched. However, the size of change impact, the union of all lines in all touched variable def-use chains, increases very slowly. For changes less than 5 lines, the impact size is larger than change size, and their impact-change ratio is 1.47. For changes less than 100 lines and greater than 50 lines, the impact size is much smaller than change size, their ratio is only 0.49. With the

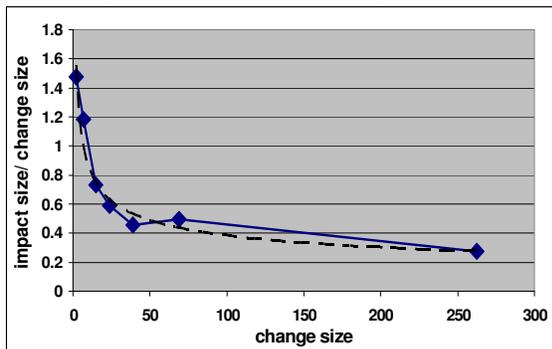


Figure 3. Ratio of impact size to change size

increase on change size, more and more lines are shared by the impacts of different changed lines. Since more and more affected variable def-use chains overlap or converge, their interference degree becomes higher and higher. According to the trend in Figure 3, the interference within large changes is very significant.

Fault density

To evaluate the effectiveness of change impact in fault prediction, we compare the fault density in impacted code with changed code across the seven groups.

Given a change, we search subsequent faults by looking for Corrective MRs subsequent (in time) to this change in ECMS. Then, in SCCS, we identify faults by the code fragments changed in deltas belong to these Corrective MRs.

For each change, we match its impact with the faults found in subsequent corrective changes. The ratio of fault size to impact size is the fault density of the impact.

For each change, we also match change itself with the faults found in subsequent corrective changes. The ratio of fault size to change size is the fault density of the change itself.

Figure 4 illustrates the fault density of change impact and change itself according to groups. The x-axis is the change size while the y-axis is fault density. With the increase of change size, the fault density in change impact increases and fault density in change itself decrease. For changes less than 5 lines, their average change size is 1.5 lines and the fault density in their change impact is only 3.7%, while fault density in change itself is about 9.1%. For changes less than 100 lines and greater than 50 lines, their average change size is 70 lines. Their fault density in change impact increases to 7.1%. But the fault density in change itself drops to 4.6%.

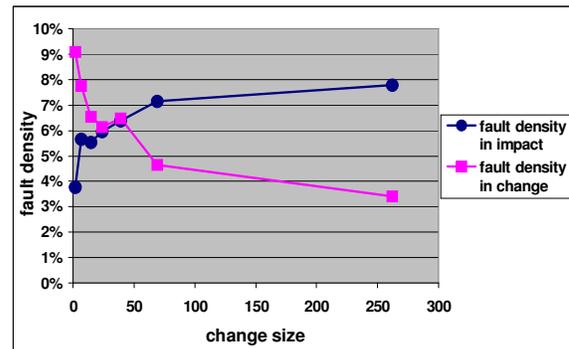


Figure 4. Fault density in change and impact.

We performed statistical analysis on the correlation between change size (on a logarithmic scale) and the ratio of impact fault density to change fault density. Their *correlation coefficient* is $R = 0.92$, $p \leq 0.003339$. In the t-test for the significance of the coefficient, $t = 5.351$ with *Degree of Freedom* = 5

The results of our statistical analysis and Figure 4 support hypothesis H3: for fault density in change impacts, large changes are higher than small changes; but for fault density in the change themselves, large changes are lower than small changes. The high fault density in change impact suggests that, to detect or predict faults in large changes, checking change impact is more effective than checking change itself.

Combining hypothesis H2 and H3, we come to a conclusion: change impacts can be complementary to change themselves in detecting or predicting faults. For large changes, checking change impact is efficient because of relative small size and high density; for small changes, checking the change themselves is more efficient.

4.3. Result Analysis

Our study shows that large changes are dangerous, not only because of the large change scope, but also for the high fault density in their impacts. It is more severe that change impacts are not as easily identified as changes themselves.

Because of the high internal interference found in impacts of large changes, we suggest that internal interference contributes to the fault density of change impact. In a large change, modified lines are distributed across a large scope and their impacts highly overlap with each other. It is difficult for a developer to identify *all* of the impacts and their high degree interference.

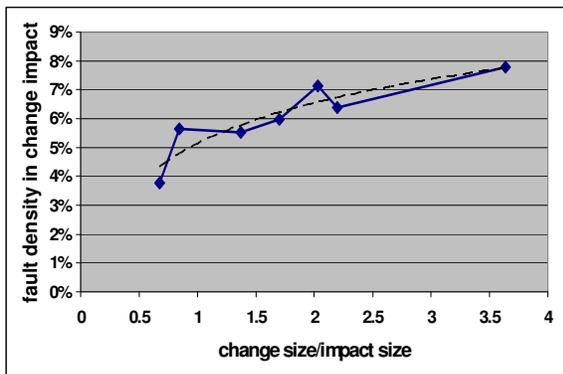


Figure 5. Relation between internal interference and fault density in change impacts.

Statistical analysis on the relation between internal interference and fault density in change impact also supports our suggestion. We use the ratio between change size to impact size as the metric for internal interference. A larger ratio value means higher internal interference. Figure 5 shows the result. The x-axis is the ratio of change size to impact size and the y-axis is the fault density in change impact. The x and y coordinate of each point represents the average of change-impact ratio and fault density of each group. The trend shows that, with the increase change-impact ratio (interference), the fault density in change impact increases. For low interference where the ratio is 0.67, the fault density is less than 4%. For high interference where the ratio is 3.6, the fault density is almost 8%. The correlation analysis between change-impact ratio and fault density shows that their *correlation coefficient* is $R = 0.877$, $p \leq 0.01149$. In t-test for the significance of the coefficient, $t = 4.081$, with *Degree of Freedom* = 5.

5. Validity Analysis

To analyze the soundness of this experiment, we discuss its construct, internal, and external validity.

Construct validity concerns in our studies center around the identification of changes and faults. In this experiment, changes are identified by textual comparison and change impacts are identified by program slicing along the variable def-use chains. Both of them are widely used constructs in program analysis. Identification of changes by textual comparison can introduce false positives. For example, breaking a statement from one line into two lines is a change according to text comparison. But such changes have no syntactic or semantic affect. However, we found such changes often when we manually checked our change impact analysis results. Faults are identified by the MR classification tools [15] with reasonable reliability. Previous studies [19] [20] showed that such classification is reasonable.

The data used in our study comes from a real industrial project and can be considered completely reliable. All the data in our study comes are from an original history without any artificial instrumentation.

With respect to internal validity, we believe that our results are consistent with our previous studies. For the semantic impact of changes, one would intuitively expect that small changes usually touch few semantic entities and affect only a small part of a program. In this study, the statistical analysis of correlation coefficients, significant levels, and t-tests

shows a strong correlation between change size and change impact size. The matching between change impact and faults has already been successfully used in the previous study on change interference [23]. These consistencies are strong support for the internal validity.

Although our study is based on the history data in a pessimistic version control system, SCCS, this approach can be easily extended to optimistic version control systems, such as Concurrent Versions System (CVS), which is widely used in open source projects. CVS can supply the same kinds of data as SCCS for our change impact analysis algorithm.

A threat to the validity of our study is that 5ESS is a very large-scale real-time project with a large number of developers who are geographically distributed. We argue, however, that the subsystem we studied is perhaps by itself more representative of a typical large project. It is developed with programming languages (C and C++) under UNIX environment. Both of them are very common in software developments. For a large project like 5ESS, well-trained developers and well-organized development processes, organization, and management are required.

Replicating our study is not difficult under other change or version management systems. This study needs two kinds of data: versions and delta information (lines added, deleted or changed) and change classification. Our data comes from two repositories: the version management system (SCCS) and change management system (ECMS). Like SCCS, current version management systems, such as CVS and Subversion, can support versions and delta information. To classify change according to purposes, the text explanations of current version management systems can be used. Automatic classification approach in [15] can also be applied on these new version management systems. A widely used tracking system, such as BugZilla, can also be used to classify changes, especially the corrective changes for bug-fixing. In our study, the classification is based on MR, a group of changes made together for a same purpose. Although MRs are not supported in CVS or Subversion, there are still other approaches to group changes. Subversion supports change transactions. Although CVS does not support transactions, other effective approaches, such as sliding windows, have been used to identify transactions. With the help from current version and change management systems, our study can be easily replicated.

6. Related Work

In studies on semantic program differencing, dependence graphs and program slicing are widely used. [10] performs semantic program differencing by partitioning the components of the two programs according to equivalent behaviors. Algorithms based on [10] can provide high precision semantic differencing. [1] uses system dependence graph and inter-procedural slicing to identify impacted part of source code change. And [26] increased the soundness of semantic conflict detection with semantic preserving transformations. These semantic differencing approaches are based on the comparison on the whole dependency graph. Their workload is heavy and difficult to be used in practice. While performing analysis on whole programs, the false positive in one procedure will be propagated and amplified. Our change impact identification approach only focuses on data-flow analysis. The cost is not that expensive. Limit our analysis scope to local can reduce the false positive rate. Such compromise can reduce cost and improve applicability in large programs. The benefits of our approach have been demonstrated in [23].

Dynamic slicing is also an important approach in impact analysis. Program chopping [8] can minimize possible fault inducing code fragments. Compared with static program slicing we used in this study, dynamic slicing can improve the precision for pointer analysis and reduce false positives in semantic interference detection. However executable versions are required beforehand. This imposes a significant build overhead in a large system like 5ESS.

Abstract syntax trees based approaches were proposed by [4] [21] [25]. Given two programs, their abstract syntax trees are compared with tree-matching algorithms. These approaches identify changes on syntax entities, but do not calculate impacts from changes.

Similar with our work, [11] checks the variable def-use dependency with a program. However changes are identified only by added or deleted dependencies. Impacted dependencies are not counted.

In empirical evaluation to semantic differencing algorithms, [2] applied the system dependence graph and inter-procedural slicing based semantic differencing algorithm on 10 programs. On average, the identified impact is 37% less than the whole

program. The relative small size of change impact supports the application of semantic differencing in practical use. However, this work does not consider the affect of change size on change impact, and does not relate change impact with faults.

Local data-flow analysis has already been used in our previous work and interference also been detected [23]. But change impact and its relation to faults were studied according to change interval, rather than change size in current study. The interference identified during our previous study is *external* interference, happening *between* changes. Given two changes on the same file, they have interference if their impacts have overlap. The interference found in current study is *internal* interference. It happens *within* the impact of a single change.

Although external interference and internal interference come from difference sources, both of them are related to faults. External interference is related to faults because of the short interval between changes. It is difficult for a developer to fully understand the impact of *all* the previous changes made by others in a short interval. So the interference between short interval changes becomes dangerous and likely to introduce faults. Internal interference is related to faults because of large change size. Modified lines are distributed across a large scope and their impact highly overlapped with each other. It is difficult for a developer to identify *all* of impacts and their interference. In large changes, the internal interference is dangerous and is very likely to introduce faults.

The relations between change size and faults have been studied in [6] [9] [16] [17] [18] [20]. The risks of large changes were also observed in these studies. But all of them are based on changes themselves, rather than impact of changes.

7. Conclusions and Future Work

Source code change histories provide abundant resources to understand the evolution of software systems. Based on source code involved in changes, many studies have been performed and interesting results have been generated. However, most of them are at the textual level or syntactic level, and do not consider internal dependencies at the semantic level. To apply semantic analysis to source code change histories, we performed an empirical study to use local data-flow analysis on the history of a successful industry project. Our studies show that:

- Large changes impact, not surprisingly, more lines than small changes; but
- Large changes have smaller impact size to change size ratio than that of small changes; and
- For fault density of impacted code, large changes are higher than small changes; but for fault density of changed code, large changes are lower than small changes.

According to these results, change impacts and change themselves are complementary in fault detection or prediction. Change impacts have relative small size and high fault density in large changes while change themselves have relative small size and high fault density in small changes. To efficiently detect or predict faults, porting approaches successfully used on changes themselves to the impacts of large changes should be helpful.

For large changes, our experiment found a high degree of interference within their impacts. It is difficult for a developer to identify *all* of the impacts and their high degree interference. We suggest that the high degree internal interference contributes to the high fault density in the impact of large changes.

The high fault density found in impacts aggravates the danger of making large changes. For large changes, the fault density in changes themselves has already become very dangerous. [14] suggests that a component should be rewritten if more than 25% of the code is changed. The high fault density in impacts intensifies the danger of large changes because change impacts are not easily identified as changes themselves.

The high fault density in the impact of large changes suggests some future research topics. To detect or predict more faults, efficient calculation of the semantic impact of large change becomes important. To disclose the internal mechanisms of high fault density, more semantic analysis on change impact is required, especially on large changes. Dynamic analysis techniques, such as dynamic slicing, might also be introduced to provide more precious impact identification and fault prediction.

8. Acknowledgements

This work was supported in part by NSF CISE Grant IIS-0438967.

9. References

- [1] D. Binkley, "Using semantic differencing to reduce the cost of regression testing", *Proc. of IEEE International Conference on Software Maintenance (ICSM'92)*, Orlando, FL, USA, Nov. 1992, 41-50.
- [2] D. Binkley, R. Capellini, L.R. Raszewski, and C. Smith, "An implementation of and experiment with semantic differencing", *Proc. of IEEE International Conference on Software Maintenance (ICSM'01)*, Florence, Italy, Nov. 2001, 82-91.
- [3] S. Breu, T. Zimmermann, "Mining Aspects from Version History", *21st IEEE International Conference on Automated Software Engineering (ASE'06)*, 2006, 221-230.
- [4] B. Fluri, M. Wursch, M. Pinzger, and H.C. Gall, "Change Distilling: Tree Differencing for Fine-Grained Source Code Change Extraction", *IEEE Transactions on Software Engineering, Special Issue on Mining Software Repositories*, Vol. 33, Issue 11, Nov. 2007, 725-743.
- [5] GrammaTech, Inc. <http://www.grammatech.com/>
- [6] T.L. Graves, A.F. Karr, J.S. Marron and H. Siy, "Predicting Fault Incidence Using Software Change History," *IEEE Transactions on Software Engineering*, Vol. 26, No. 7, July 2000, 653-661.
- [7] T.L. Graves, A. Mockus, "Inferring Change Effort from Configuration Management Databases", *Proc. of the Fifth International Symposium on Software Metrics, IEEE*, 1998, 267-273.
- [8] N. Gupta, H. He, X. Zhang, and R. Gupta, "Locating Faulty Code Using Failure-Inducing Chops", *Proc. of the 20th IEEE/ACM International Conference on Automated Software Engineering (ASE 2005)*, Long Beach, California, Nov. 2005, 263-272.
- [9] A.E. Hassan, and R.C. Holt, "The top ten list: Dynamic fault prediction", *Proc. of IEEE International Conference on Software Maintenance (ICSM'05)*, Budapest, Hungary, Sep. 2005, 263-272.
- [10] S. Horwitz, "Identifying the Semantic and Textual Differences between Two Versions of a Program", *Proc. of the ACM SIGPLAN conference on Programming language design and implementation (PLDI'90)*, New York, June 1990, 234-245.
- [11] D. Jackson, and D.A. Ladd, "Semantic Diff: A Tool for Summarizing the Effects of Modifications", *Proc. of IEEE International Conference on Software Maintenance (ICSM'94)*, Nov. 1994, pages 243-252.
- [12] M. Kim, D. Notkin, D. Grossman, "Automatic Inference of Structural Changes for Matching across Program Versions", *Proc. of the International Conference on Software Engineering (ICSE'07)*, 2007, 333-343.
- [13] P. Knab, M. Pinzger, A. Bernstein, "Predicting defect densities in source code files with decision tree learners", *Proc. of the 2006 international workshop on Mining software repositories (ICSE'06)*, Shanghai, China, 119 - 125.
- [14] M. Leszak, D.E. Perry, and D. Stoll, "Classification and evaluation of defects in a project retrospective", *Journal of Systems and Software*. Vol. 61, Issue 3, April 2002, 173-187.
- [15] A. Mockus, and L.G. Votta, "Identifying Reasons for Software Changes Using Historic Databases", *Proc. of IEEE International Conference on Software Maintenance (ICSM'00)*, San Jose, CA, USA, October. 2000, 120-130.
- [16] A. Mockus, and D.M. Weiss, "Predicting risk of software changes", *Bell Labs Technical Journal* 5, 2, 2000, 169-180.
- [17] N. Nagappan, and T. Ball, "Use of relative code churn measures to predict system defect density", *Proc. of International Conference on Software Engineering (ICSE'05)*, Saint Louis MO, May 2005, 284- 292.
- [18] T.J. Ostrand, E.J. Weyuker, R.M. Bell, "Predicting the location and number of faults in large software systems" *IEEE Transactions on Software Engineering*, Vol. 31, Issue 4, April 2005, 340-355
- [19] D.E. Perry, H.P. Siy, and L.G. Votta, "Parallel Changes in Large Scale Software Development: An Observational Case Study", *ACM Transactions on Software Engineering and Methodology*, Vol. 10, No. 3, July, 2001, 308-337.
- [20] R. Purushothaman, and D.E Perry, "Toward Understanding the Rhetoric of Small Source Code Changes", *IEEE Transactions on Software Engineering, Special Issue on Mining Software Repositories*, Vol. 31, No. 6, June 2005, 511-526.
- [21] S. Raghavan, R. Rohana, D. Leon, A. Podgurski, V. Augustine, "Dex: A Semantic Graph Differencing Tool for Studying Changes in Large Code Bases", *Proc. of IEEE International Conference on Software Maintenance (ICSM'04)*, September 2004, 188-197.
- [22] M.J. Rochkind, "The Source Code Control System", *IEEE Transactions on Software Engineering*, Vol. SE-1, No. 4, December 1975, 364-370.
- [23] D. Shao, S. Khurshid, and D.E. Perry, "Evaluation of semantic interference detection in parallel changes: an exploratory experiment", *Proc. of the 23rd IEEE International Conference on Software Maintenance (ICSM'07)*, Paris, France, October 2007, 74-83.
- [24] P. A. Tuscany, "Software Development environment for Large Switching Projects", *Proc. of the 6th International Conference on Software Engineering*, Tokyo, September 1992, 58-67.
- [25] W. Yang, "Identifying syntactic differences between two programs", *Software - Practice & Experience*, Vol. 21, Issue 7, July 1991, 739 - 755.
- [26] W. Yang, S. Horwitz, and T. Reps, "A program integration algorithm that accommodates semantics-preserving transformations", *ACM Transactions on Software Engineering and Methodology*, Vol. 1, No. 3, July 1992, 310-354.
- [27] T. Zimmermann, A. Zeller, P. Weissgerber, S. Diehl, "Mining version histories to guide software changes", *IEEE Transactions on Software Engineering*, Vol. 31, Issue 6, June 2005, 429 - 445.