

# Prioritization of Regression Tests using Singular Value Decomposition with Empirical Change Records

Mark Sherriff<sup>1,2</sup>, Mike Lake<sup>1</sup>, and Laurie Williams<sup>2</sup>

<sup>1</sup>IBM, <sup>2</sup>North Carolina State University

mark.sherriff@ncsu.edu, johnlake@us.ibm.com, williams@csc.ncsu.edu

## Abstract

*During development and testing, changes made to a system to repair a detected fault can often inject a new fault into the code base. These injected faults may not be in the same files that were just changed, since the effects of a change in the code base can have ramifications in other parts of the system. We propose a methodology for determining the effect of a change and then prioritizing regression test cases by gathering software change records and analyzing them through singular value decomposition. This methodology generates clusters of files that historically tend to change together. Combining these clusters with test case information yields a matrix that can be multiplied by a vector representing a new system modification to create a prioritized list of test cases. We performed a post hoc case study using this technique with three minor releases of a software product at IBM. We found that our methodology suggested additional regression tests in 50% of test runs and that the highest-priority suggested test found an additional fault 60% of the time.*

## 1. Introduction

During development and testing, changes made to a system to repair a detected fault can often inject a new fault into the code base. These injected faults may or may not be in the same file(s) that were just changed, since the effects of a change in the code base can have ramifications in other parts of the system. New faults introduced during testing and maintenance can be isolated by impact analysis and regression testing techniques. Impact analysis and regression testing techniques exist that find changes in a system based upon a modification [11, 15]. However, techniques that utilize call graphs or other dynamic means can be computationally intensive, and may not be cost-effective for use in industry [16]. Often regression test selection techniques do not include files that are not part of the source code, such as properties files, help files, and configuration files. Additionally, current impact analysis and regression techniques are based upon semantic analysis and may not

consider trends in actual usage and/or the fault-proneness of the set of files impacted [13, 14].

Insight into the injected faults due to fault removal techniques can be gathered from the information generated during the development process itself. During development, programming teams will produce a variety of software development artifacts. A *software development artifact* is an intermediate or final product that is the result or by-product of software development [8, 9]. Development artifacts can also provide information about underlying structures within a system that would normally not be apparent [2, 3, 6, 10, 21, 22].

One such development artifact is change records. Change records can show how components interact with one another or how a system is evolving during development by examining what areas of the system change together and where new sets of changes are emerging [2, 3]. We refer to a group of files that have been changed together to address a fault as a *track*. Testers could use historical information about tracks to direct their testing efforts when a new track is created.

Change records can be used to identify *association clusters* in a software system. An *association cluster* consists of a set of files that exhibit a specific relationship with the other files in the cluster with regards to particular development artifact [3]. If a number of files repeatedly change together to fix a set of faults, then those files would be identified as an association cluster. Association clusters are based upon historical patterns in the composition of tracks. Quantification of the frequency of occurrence (indicative of fault-proneness) of tracks can be used to rank the “strength” of an association cluster.

*The goal of this research is to reduce the number of regression faults that are delivered to the field.* The methodology proposed in this research to address our goal is called SVD-Based Regression Test Prioritization (SRTP). SRTP provides a framework for gathering change records from testing and field failures, generating association clusters, and leveraging those clusters to guide regression test prioritization. The data from change records are compiled into a matrix that portrays the historically-based usage relationship between two files. A singular value decomposition (SVD) [5] is performed on the matrix to generate the association clusters. The results of the SVD can then be utilized to guide

regression test prioritization. *Our hypothesis is that a methodology based upon SVD using historical change records can be can prioritize regression tests to reduce the number of regression faults released to the field.* To examine the efficacy of SRTP, an empirical case study was conducted with an industrial project at IBM. During this case study, we investigated the following research question: *Do the association clusters produced by SVD accurately surface sets of files that are indirectly impacted by a set of changes?* Change records initiated from fault removal efforts were gathered on three consecutive minor releases of an IBM product. A *minor release* is defined as an “update and fix” pack to a final release of a product.

The rest of this paper is organized as follows. Section 2 provides information on background and related work. Section 3 describes the SRTP methodology in detail, while Section 4 describes our case study at IBM. Section 5 presents our summary and future work.

## 2. Related Work

In this section, we will discuss related research and background literature in regression test selection, uses of clustering change records, and the singular value decomposition.

### 2.1 Regression test selection

*Regression testing* is the process of retesting a system or component to verify that changes made to the system code have not caused unintended effects and that the system is still compliant with the specified requirements [1]. The goal of regression test selection (RTS) techniques is to isolate the tests most likely to uncover injected faults after a system modification. The simplest RTS technique is the *retest-all* strategy, wherein the entire test suite is exercised. However, techniques have been developed that can minimize the set of tests that need to be rerun while still maintaining the overall effectiveness of a retest-all strategy [15]. A RTS strategy is considered *safe* if the subset of tests identified by that strategy contain all the test cases from the test suite capable of finding faults in the system based upon a system modification [19].

One safe RTS strategy was developed by Rothermel and Harrold [15]. Their RTS strategy analyzes changes in the control dependence graph of the software system and changes to the test suite to determine the best subset of test cases. Rothermel et al. have shown that this technique is considered safe, identifying tests that cover not only unchanged lines but also newly added lines of code. Further studies from Rothermel have shown that regression test selection and test case prioritization in

general can be effective at reducing the amount of work needed to be done in testing after a system modification, but that the cost-effectiveness of these techniques can vary [16, 17].

One aspect of the technique developed by Rothermel et al. is that it requires source code to be effective. However, faults can be injected into non-source files in a system that might not be detected through a technique that solely examines source code. Consider a system that utilizes static properties files that are read in by different modules within the system. If a fault is injected into these properties files, there would be no reflection of this change in the source code, yet the fault could be as severe as any fault in the source itself [7].

Our technique adds to the body of knowledge of RTS strategies by addressing the problem of test selection when non-source files are involved. For example, 35% of all faults involved in one system examined in our study contained a non-source file. Further, we are trying to improve the overall cost-effectiveness of RTS strategies by using historical trends over time as opposed to processing dynamic call information.

### 2.2 Clustering files based upon change records

Research is currently being performed in gathering and analyzing data from source control systems to identify core components in a software system for use in impact analysis [2, 3, 6, 10, 21, 22]. Ren et al. has created an Eclipse plug-in to predict the impact of code changes for developers to use in-process through white-box techniques [13]. Their plug-in, called Chianti, works by capturing atomic-level changes in the code base. Dependencies are then calculated between these atomic changes to predict what other areas of the code might be affected by a change through the use of call graphs. Ren performed two case studies on a 100 KLOC system and found that Chianti was able to reduce the number of regression tests depending on the degree of the change implemented. The primary difference between the impact analysis technique used in Chianti and our technique is that Chianti is based upon semantically-based methods in which all associations are created equal regardless of actual usage. The association clusters created in our technique are based upon historical data and, therefore, might be better for prioritization.

Canfora and Cerulo use the descriptions of faults and change records from developers to determine the effect of a change [4]. Their technique compares similarities in the description of a new change to previous changes to identify possible areas that have been affected. If the description of a new fault matches keywords in previous faults, then those files identified by the previous faults may be affected by this new change. Canfora and Cerulo found that, if fault records used consistent keywords and

phrases, a recall rate of nearly 98% was possible, with a precision rate around 85% [4].

While Ren’s works focused on guiding developer efforts, other research uses the same sets of change records to improve program comprehension. Further work expanded on the idea of gathering change records to isolate clusters of files within a software system to drive program comprehension. Beyer and Noack’s work analyzed clusters of files using a co-change graph to plot files onto a two- or three-dimensional graph using an energy-based graph layout [3]. Files that contained more edges between them were closer together on the graph, thus creating clusters of files. These clusters of files could be directly identified and related to functional requirements or third party components within the system [3]. Other research by Gall also generated association clusters of files within a system for program comprehension [6]. His method used a set of commonalities that could be detected from change logs (i.e. files that were edited on the same day by the same person) to create his sets of sub-modules.

### 2.3 Singular value decomposition

Singular value decomposition (SVD) is a linear algebra technique that decomposes a given matrix into three component matrices [5]: (1) the left singular vectors; (2) a set of singular values; and (3) and right singular vectors. The two matrices that are made up of singular vectors provide information about the structure of the original matrix. The singular values describe the strength of the given components of the original matrix. The SVD theorem [5] states that given a matrix  $M$ , then there exists a decomposition of  $M$  such that  $M = USV^T$ .

The SVD of a matrix can also be described geometrically. The SVD shows that the values of any matrix  $M$  can be reconstructed by a rotation ( $U$ ), followed by increasing the matrix values ( $S$ ), followed by another rotation ( $V$ ) [20]. For example, if  $M$  represented coordinates that generated a three-dimensional shape, then that shape could be constructed from the rotational information in  $U$  and  $V$ , along with stretching the shape out to its proper size with the information in  $S$  [20]. This type of decomposition can be important and useful in that the rotational matrices isolate the key components of the original matrix, finding relationships between the various data points, while the strength matrix indicates which of the key components illuminated in the rotational matrices are the most important [5, 20]. In our research, this core idea of isolating key components of the original matrix is the basis for using the SVD with SRTP. When the matrix is comprised of change records, fault information, or some other data from the development process, these key components highlight underlying structures in the code base.

The SVD has other uses in computing. For instance, this technique can be used in image and signal compression. A gray scale image could be represented as a two-dimensional matrix made up of intensity values, indicating the darkness of a particular pixel. In this instance, we could treat the image matrix itself as the original  $M$  matrix and perform a SVD on it. Once the decomposition is completed, the resulting matrices,  $USV^T$ , can also be represented as the sum of component matrices, as shown in Equation 1:

$$M = u_1 s_{1,1} v_1^T + u_2 s_{2,2} v_2^T + \dots + u_k s_{k,k} v_k^T \quad (1)$$

where  $k$  is less than or equal to the rank of matrix  $M$ . This provides a rank- $k$  approximation of the matrix. The first component of this factorization indicates the component that has the largest singular value and contributes the most variability to the overall matrix. As subsequent components are added together, the matrix, and thus the image itself, begins to re-take its original form [5].

Osinski et al. created a clustering algorithm based upon SVD to improve search queries on a set of documents [12]. They built an original matrix based upon keywords in the document set. The SVD was performed on this matrix to generate clusters of documents that were similar based on their keywords. Enough clusters were gathered to account for 90% of the variability of the original matrix, with the remaining clusters discarded as signal noise [12]. The documents were then assigned to clusters based upon which cluster they had the closest association with. Anecdotal evidence from users who were presented with the clusters generated with this study found that 70-80% of the clusters were useful and over 75% of identified cluster labels were correct [12].

### 3. SVD-based Regression Test Prioritization

SRTP provides a methodology that derives associations using SVD based upon a set of change records from testing and field failures. These association clusters of files portray an underlying structure in the system indicating how files tend to be executed, tested, and changed together [18]. SVD is used to leverage its ability to illuminate underlying structures in a data set in which the data could be associated in multiple ways. In this section, we will describe the process of SRTP, which includes deriving the association clusters from change records and producing a reduced set of regression tests. Figure 1 outlines the steps of the SRTP method that will be further described in the remaining subsections. Table 1 summarizes the matrices and data sets we will introduce in this section.

```

1 Create matrix M where the values in the
matrix indicate the number of times two
files have changed together.
2 Create matrix T where the values in the
matrix indicate whether a file is affected
by a particular fault/test case.
3 [U, S, V] = svd(M);
4 for i:size of U
5   Gather cluster i information
6   for j:size of U
7     if |U(j, i)| > threshold
8       Place element of cluster i into R
9     end
10  end
11 end
12 Y = (R * R') * T;
13 Represent new system modification x as a
vector in which the value indicates whether
the file has been affected in this change.
14 Y * x yields the prioritization of each
test case based upon the similarity of the
new system modification x to previous
changes in the system.

```

**Figure 1. Algorithm for SVD-based RTP.**

**Table 1. Matrix Summary.**

Matrix ID	Figure 1 Step ID	Description
$M$	1	Describes how files change together
$U, V$	3	Created from SVD – gives cluster structure
$S$	3	Created from SVD – gives cluster strength
$R$	8	Expanded cluster list by separating the clusters from the $U$ matrix
$T$	2	Which files are affected by a test
$P$		Temporary matrix to show how tests can be associated with clusters of files
$x$	12	Indicates whether a file has changed in a new system modification
$Y$	12	Shows how files are linked to tests through the cluster structure – multiplying a change vector $x$ by this matrix gives the test case prioritization

### 3.1 Identifying data sources

Source control systems are the primary source for gathering change records. When a developer checks a file in to a source control system, the system typically records the time of the check-in along with information about the developer and the nature of the change. Individual changes can be often be linked together into tracks, either through a specific mechanism in the source control system that records that information or through the examination of change record check-in information. With information regarding tracks, we are able to ascertain how files change together.

Some more complex source control systems are also integrated with a fault tracking system. With these more complex systems, tracks can be associated directly with the fault record that the changes are addressing, providing detail about how tracks are linked together. Information from fault tracking systems allows us to isolate tracks to those made under specific circumstances. For example, changes derived from faults found during system test could be compared to changes derived from field failures discovered by customers.

### 3.2 Gathering software development artifact data

After appropriate data sources have been identified, an analysis matrix can be generated that contains the systems files along each axis. The values within the analysis matrix show how the files are connected through change records. For illustrative purposes on how to build the analysis matrix, we will use a set of sample data to generate this analysis matrix as an example. Table 2 shows a small sample of the set of the change records that were used to create our example analysis matrix. This example uses a small system consisting of five files.

**Table 2. Sample Change Record Information.**

Test Case ID	Fault ID	Track ID	Files Changed
T1	A1	988	1
T2	A2	989	2, 3
T3	B1	990	4, 5
T3	B2	991	4, 5
T1	B3	992	1, 2
T4	C1	993	2, 3
...	...	...	...

We have built an example analysis matrix  $M$ , shown below in Equation 2. The values in the matrix represent the number of times that each file appeared in a track with another file. Thus, File 2 has appeared in a track 10 times together with File 1, 21 times together with File 3, and 0 times by itself (since  $M(2,2) = M(2,1) + M(2,3)$ ).

Similarly, File 3 has changed 21 times with File 2 and 3 times by itself.

$$M = \begin{matrix} & \begin{matrix} F1 & F2 & F3 & F4 & F5 \end{matrix} \\ \begin{matrix} F1 \\ F2 \\ F3 \\ F4 \\ F5 \end{matrix} & \begin{bmatrix} 25 & 10 & 0 & 0 & 0 \\ 10 & 31 & 21 & 0 & 0 \\ 0 & 21 & 24 & 0 & 0 \\ 0 & 0 & 0 & 15 & 12 \\ 0 & 0 & 0 & 12 & 17 \end{bmatrix} \end{matrix} \quad (2)$$

Upon initial examination of this matrix, we note that Files 4 and 5 change together or by themselves. Based on this, it appears that Files 4 and 5 are strongly linked in isolation from the rest of the system. Similarly, Files 1, 2, and 3 are also linked, with Files 2 and 3 having the strongest bond of the three.

### 3.3 Perform the singular value decomposition

To determine the strength of the associations between files and to generate the association clusters, we perform a SVD of this matrix. The strength of the association is determined by the frequency of time the files changed together. A SVD of  $M$  provides the following matrices, shown in Equations 3 and 4:

$$U = V = \begin{bmatrix} -.29 & 0 & .9 & .31 & 0 \\ -.76 & 0 & -.02 & -.56 & 0 \\ -.59 & 0 & -.43 & .69 & 0 \\ 0 & -.68 & 0 & 0 & -.74 \\ 0 & -.74 & 0 & 0 & .68 \end{bmatrix} \quad (3)$$

$$S = \begin{bmatrix} 51.1 & 0 & 0 & 0 & 0 \\ 0 & 28.4 & 0 & 0 & 0 \\ 0 & 0 & 24.8 & 0 & 0 \\ 0 & 0 & 0 & 4.1 & 0 \\ 0 & 0 & 0 & 0 & 3.9 \end{bmatrix} \quad (4)$$

The  $U$  and  $V$  matrices provide information as to the structure of the association clusters. The singular values from the  $S$  matrix represent the amount of variability each association cluster contributes to the original analysis matrix. Note that  $U$  and  $V$  are equal, due to  $M$  being a symmetric matrix.

A cluster's strength, represented by the size of the singular value coupled with it, indicates the amount of variability that the association cluster provides to the original analysis matrix [20]. Dividing a cluster's singular value by the sum of all the singular values provides the percentage of how representative the cluster is of the original matrix. In this example, a high singular value indicates that that association cluster is more

prominent in the analysis matrix, due to a greater number of changes that have occurred to that set of files. A high singular value could be indicative of a particularly problematic section of code or a new feature that has just been introduced into the system and is experiencing its first rigorous testing.

### 3.4 Gather the association clusters

The values in the  $U$  matrix correspond to the composition of each association cluster. In this example, there are five association clusters because the rank of  $M$  is five. The first column of  $U$ , representing the structure of the first association cluster, is coupled with the first singular value in  $S$ , representing the strength of that association cluster. Since it is coupled with the largest singular value, the first association cluster represents the greatest amount of variability in the original analysis matrix and is the most prominent association cluster. From the  $U$  matrix, we see that the first association cluster is comprised of Files 1, 2, and 3, indicated by the fact that the three files all have values with a similar sign. Further, each of these values has a larger magnitude than .1, the threshold we used in our research. A threshold is used when selecting cluster members so that only files with a strong association to the other files are included in the cluster. When a file's value is greater than the threshold, we can add that value to the matrix  $R$ , which will contain the final set of clusters. This is similar to the threshold that Osinski used in his algorithm [12]. In the third cluster, we see that File 1 is its own cluster that can, at times, change without Files 2 and 3. So, in effect, we get two associations out of the third cluster, one with File 1 by itself and one with Files 2 and 3 together. Matrix  $R$  from this  $U$  matrix is shown in Equation 5:

$$R = \begin{matrix} & \begin{matrix} C1 & C2 & C3 & C4 & C5 & C6 & C7 & C8 \end{matrix} \\ \begin{matrix} F1 \\ F2 \\ F3 \\ F4 \\ F5 \end{matrix} & \begin{bmatrix} .29 & 0 & .9 & 0 & .31 & 0 & 0 & 0 \\ .76 & 0 & 0 & 0 & 0 & .56 & 0 & 0 \\ .59 & 0 & 0 & .43 & .69 & 0 & 0 & 0 \\ 0 & .68 & 0 & 0 & 0 & 0 & .74 & 0 \\ 0 & .74 & 0 & 0 & 0 & 0 & 0 & .68 \end{bmatrix} \end{matrix} \quad (5)$$

Note that the values in each association cluster's column vector represents the degree to which that file is likely to change in that cluster. In this way, each file is weighted within that association cluster as to its degree of participation. For example, the first association cluster is primarily composed of File 2 and File 3 due to their higher values. File 1 is a minor participant in this association cluster. If we reexamine the original analysis matrix  $M$ , we can see the strong correlation between Files 2 and 3 with a somewhat looser correlation with File 1, since these files only tend to change together and not at

all with Files 4 and 5. The association cluster in the second column portrays the next most significant cluster, comprised of Files 4 and 5.

The singular values of these clusters found in the  $S$  matrix provide some indication as to how they should be analyzed. The first cluster represents 45% of the overall variability in the matrix, which can be determined by dividing the first singular value by the sum of all the singular values. Further, the third and fourth clusters collectively represent 22% (since these two clusters were split apart from the values in the third left singular vector) and the fifth and sixth represents 4%. These percentages show that the first cluster defines the majority of the information regarding these files. Clusters three, four, five, and six are, in effect, sub-clusters of the first cluster because they contain a similar set of files. At this step in our technique, the matrix  $R$  can provide information about the likelihood of a change in an association cluster based upon previous change information.

This technique is similar to the cluster rank algorithm used by Osinski et al. in their SVD-based search term clustering algorithm [12]. Osinski multiplied their document matrix by a modified  $U$  matrix from the SVD to derive the impact that each search term had on a given document. In this fashion, the values from the result vector were used to assign a document to its closest-matching search term cluster [12].

### 3.5 Generate reduced test suite

The next step is to use the current tracks in the system to determine how each test case in the system is related to a given association cluster. Using data from the source control and fault management system, we can map a track  $x$  directly to a particular fault in a one-to-one relationship. To accomplish this analysis, we can take the set of all tracks in the source control system and create a matrix  $T$ , with each row indicating a track and each column representing a particular file. An example matrix  $T$  has been constructed from the information in Table 1 in Equation 6.

$$T = \begin{matrix} & \begin{matrix} F1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} T1 \\ T2 \\ T3 \\ T4 \\ T5 \\ T6 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \end{bmatrix} \end{matrix} \quad (6)$$

Since a track is opened by a single distinct fault, we can say that each row represents a specific fault found in the system. If each test case created one fault, the matrix  $T$

would also represent test case traceability to the files affected by those tests. However, in all likelihood, a test case can find more than one fault in a system. In this instance, tracks in  $T$  that are opened by the same test case can be combined. Thus, the matrix  $T$  represents which files are affected by each particular test case.

Multiplying matrix  $T$  by matrix  $R$  yields another set of associations in matrix  $P$  in Equation 7:

$$T * R = [Tests \times Files] * [Files \times Clusters] = [Tests \times Clusters] = P \quad (7)$$

The  $P$  matrix that is generated from multiplying  $T$  and  $R$  together links a particular test case with its connection to a given association cluster in the same way that multiplying a single change vector  $x$  multiplied by  $R$  provided information on the impact of that track.

We can then relate the link between test cases and clusters back to individual files by multiplying  $P^T$  by another  $R$ , as shown in Equation 8:

$$R * P^T = [Files \times Clusters] * [Clusters \times Tests] = [Files \times Tests] = Y \quad (8)$$

The matrix  $Y$  that is produced now has associated test cases in the system to particular files after transforming their relationship through the association clusters. This is analogous to changing the basis of the original analysis matrix  $M$  by multiplying it by the results from the SVD.

This file association matrix  $Y$  is then used in conjunction with the fault information contained in the change vectors to provide an association between files and faults. If we multiply the matrix  $Y$  by a new change vector  $x$ , the output is a single vector where each column is a particular test case and the values represent how closely the new change vector matches with each test case. An example of this multiplication is shown in Equation 9:

$$Yx' = \begin{matrix} & \begin{matrix} F1 & F2 & F3 & F4 & F5 \end{matrix} \\ \begin{matrix} T1 \\ T2 \\ T3 \\ T4 \\ T5 \\ T6 \end{matrix} & \begin{bmatrix} .69 & .09 & .23 & 0 & 0 \\ .32 & .90 & .76 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ .78 & .60 & .62 & 0 & 0 \\ .32 & .90 & .76 & 0 & 0 \end{bmatrix} \end{matrix} * \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = [.32 \ 1.7 \ 0 \ 0 \ 1.2 \ 1.7] \quad (9)$$

The value in the result priority vector itself represent how close the new modification matches a given test case, taking into account files that tend to change together. For instance, a high-priority test case will not only contain files that were directly changed in the new modification, but will also contain files that tend to be affected when the newly modified files change. The values themselves are used as a relative prioritization among the test suite

for this modification and this modification alone as opposed to using the prioritization values as an absolute that can be compared against different modifications. For example, if more files have changed, the priority values can be higher since there is now an additional component to match with available tests. However, prioritization vectors can be additive if several system modifications are being examined simultaneously. In this instance, we find that T2 and T6 are equally appropriate to run in this instance, while T3 and T4 are not related at all to the changed files.

Faults that are associated with non-zero values in the result vector execute areas of code that have previous problems and are candidates for being rerun. These faults are ones that either directly involve the files that have been affected by the change or files that are in the same association clusters as the affected files. Thus, files that are closely tied to the affected files through previous system modifications, will have a greater impact on the values in the result vector.

### 3.6 Limitations

The association clusters are based upon the change records that may or may not be accurate. For example, the case study reported in this paper uses data from an IBM system. IBM's documented development process and interviews from developers and managers indicated that files that were changed together are related and are addressing a specific fault. The process in place in IBM helps to minimize opportunistic changes whereby a developer makes changes unrelated to an open fault while fixing that fault. If opportunistic changes occur during fault removal efforts, we cannot be certain that the set of files that change together are related to a particular fault. Conversely, opportunistic changes and aggregating multiple changes in one change record may be prevalent in open source projects

Another limitation is that our technique requires that some level of traceability exists between test cases and the faults that are discovered by the test cases. This traceability is required so that a matrix of files affected by test cases can be compiled. In our study, this data exists as an additional note in the description of the fault in the fault tracking system.

## 4. IBM Case Study

During the second half of 2006, research in SRTP was performed at IBM in Durham, NC. In this section, we will demonstrate the formative research that led to the evolution of SRTP v1.0. *Our hypothesis is that a methodology based upon SVD using historical change*

*records can be can prioritize regression tests to reduce the number of regression faults released to the field.* In this section, we will investigate the two research questions posed in Section 1.

### 4.1 Case study setup

Following Step 1 of our methodology outlined in Section 3, we began by examining available data sources. IBM's source control and fault management system generates detailed logs on tracks. The project that was selected had three consecutive minor releases. Each release of the project contained approximately 21,000 files with over one million lines of code.

### 4.2 SVD-based regression test prioritization

The prioritization approach through SRTP identifies related files and includes them in the regression test prioritization based upon the weighting values given to files in clusters by the SVD in the  $R$  matrix generated from the  $U$  matrix. The value a file has with a given cluster in the  $R$  matrix dictates its impact within a regression test prioritization as well. What this yields is a prioritization list in which the test cases are ordered as to not only how similar the track is to a given test case, but also the impact that a test could have on surrounding files in an association cluster as described in Section 3.6. A high prioritization value on a test indicates that: (1) the files exercised by the test closely resemble the set of changed files; and (2) the set of changed files have a relatively high change frequency, indicating fault-proneness.

We evaluated the efficacy of this methodology by using a random data splitting technique with the three minor releases of an IBM software system to address the our research question "*Do the association clusters produced by SVD accurately surface sets of files that are indirectly impacted by a set of changes?*".

We began by randomly selecting two-thirds of the tracks for each release as the "historical data" from which we generated the association clusters. From the remaining one-third of the tracks, twenty tracks were selected for use as "future changes" that are being introduced into the system and for which we are prioritizing regression tests. The remaining tracks in the one-third set represented the future set of faults that would be found in the system. For example, if simple retest suggests we retest File 1 and SRTP further suggests we test Files 1 and 2, if there are instances of Files 1 and 2 together in the "future set," we assume our SRTP retest would have prevented that fault.

For each of the twenty tracks selected to represent new changes introduced into the system, we converted these tracks into vectors. Each of these vectors was then

multiplied by the association matrix  $Y$  generated from the “historical data” set of tracks. Since the association matrix  $Y$  links test cases to files, the results of this multiplication yielded a regression test prioritization list. We compared the prioritization lists from the twenty selected “future change” tracks with the future set of faults to determine if the SRTP prioritization list identifies additional test cases that actually appears in the future set. If the suggested tests did appear in the one-third future set of data, then it was recorded as a positive match recommendation for the regression test prioritization model. We repeated the data splitting analysis six times. The results of our analysis are found in Table 3. Figure 2 shows the percentage of the total number of tests that were recommended for each “future change” track.

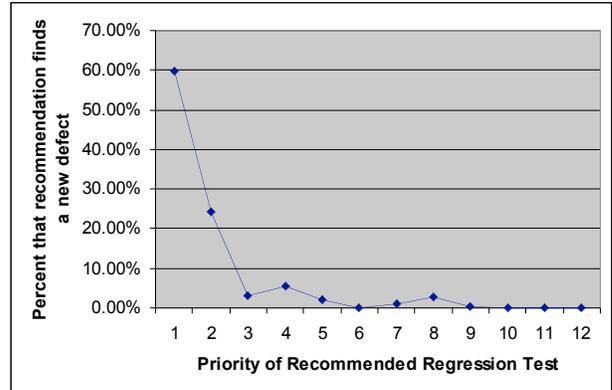
We found that 50% of the time the SRTP model did not recommend any additional tests beyond a simple retest of changed files. These instances are represented in Figure 2 by the points that appear close to 0%. If the model recommends the same tests as a simple retest, this indicates that the areas of the system that these tracks represent are either historically self-contained modules or that there is not enough change data among the files to make an accurate prediction. In the case of the self-contained modules, a regression set that is the same as a simple retest is a positive result, since testing can be expensive, and we do not want to incur additional cost with little extra benefit.

Conversely, if there is sparse change data in a section of the code, where numerous files have only changed once or twice together, a cohesive association cluster structure might not have formed yet. In these “loose” cluster areas of the system, our model recommends more tests with an overall lower confirmed true positive rate because of the number of recommended tests. These instances are represented by the spikes in the graph in Figure 2. In these instances, testers must rely more on the

prioritization of the list of tests rather than simply executing those tests that have non-zero values in the prioritization vector.

Upon examining the prioritization of the tests themselves, we found that the first test in each prioritization list was the most likely to uncover a future fault. Approximately 60% of the first test cases identified a future fault in the files included in a new set of changes.

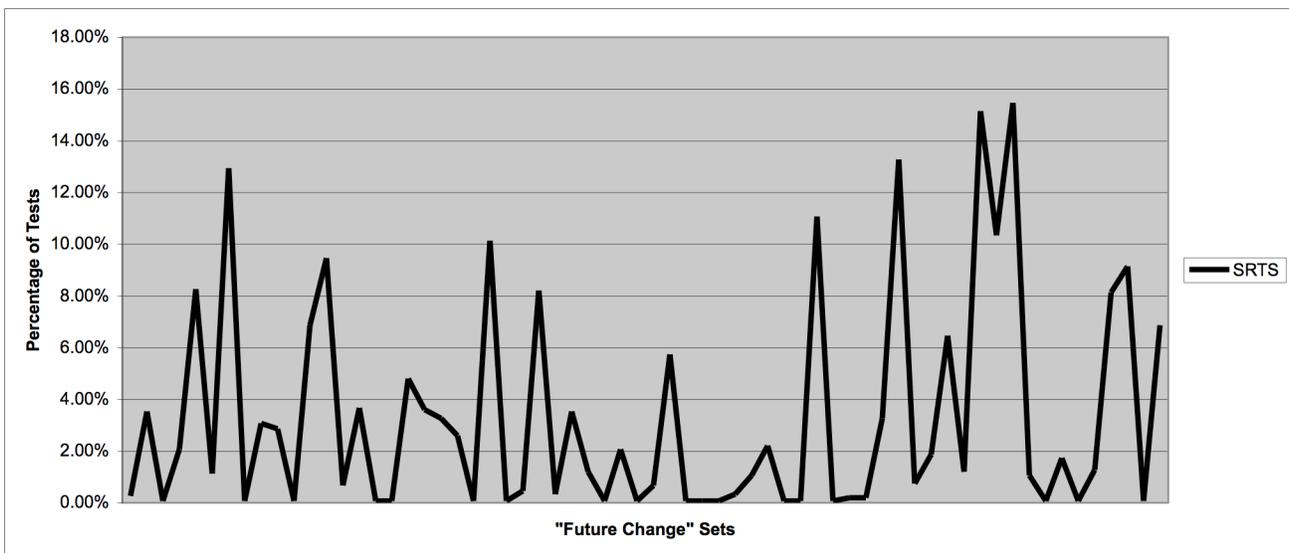
The percentage of the time that a prioritized test finds a fault is found in Figure 3.



**Figure 3. Percentage matches in future data by regression test priority.**

To further evaluate our technique, we utilized the framework for analyzing RTS techniques proposed by Rothermel and Harrold [14]. Rothermel and Harrold use four characteristics to describe and evaluate a RTS technique:

- *inclusiveness*: whether a RTS technique is considered safe or not;
- *precision*: the relative number of tests recommended for rerunning by the RTS technique that could not have found any newly-injected



**Figure 2. Percentage of Tests Recommended to be Rerun.**

**Table 3. Model Recommendations.**

Recommendations	Percentage Non-Source Files	Original test only	Confirmed True Positive
Original test only	35.2%	50%	N/A
15-50 additional tests	33.6%	19%	25%
1-14 additional tests	39.5%	31%	40%

faults,

- *efficiency*: the computational and overhead costs associated with running the RTS technique, and
- *generality*: the context for which the RTS technique can be used, including language, and program constructs.

Our technique is not safe because of our use of historical data. If a new system modification is introduced in which sets of files change together that never have before, such as a new feature being introduced, then our technique might omit some tests. The limitation of not being safe is one that is shared among all historically-based RTS techniques.

Rothermel and Harrold recognize that no RTS technique can be guaranteed as 100% precise [14]. The precision of our technique improves as the number of changes in the system increases. Fewer tests are identified for regression test and the number of matches that those tests have in the random one-third data set increase. A Pearson correlation between the number of changes of the files involved in a new track and the accuracy of the model ( $p = .003$ ) indicates that as more change data is collected on a given set of files, the better the regression test prioritization for those files. Thus, our technique's precision is only governed by how well future modifications match historical patterns. If new features are not being added to a system and new modifications follow historical patterns, then the precision of our technique could be high.

The most inefficient part of our technique is the computation of the SVD, assuming that there are automated procedures in place to gather change records to populate the  $M$  matrix. However, the SVD does not need to be calculated with each new system modification. We believe that it would be sufficient to compute a new SVD at the end of the day, automatically after no more modifications would be made for the day. In this way, the efficiency of our technique can be much higher than any dynamic RTS technique, since this would reduce our technique to a simple matrix times a linear vector for each RTS recommendation. In an organization where dynamic means are infeasible, our technique might be more appropriate.

The novel aspect of our RTS technique is its generality in that its context is all-inclusive. Any file that is managed through a change management system, whether that file is source code, media files, documentation, or anything else, our technique can recommend tests appropriate for those files based upon which other files

they have changed with. Faults that are found in non-source files can be as severe as those within source files and thus we believe that our technique for prioritizing tests with regard to all files in the system can provide some added insight into prioritizing regression test cases.

## 5. Summary and Future Work

In this paper, we propose an empirically-based regression test prioritization method based upon structures discovered through change records and singular value decomposition. SRTP makes use of the information in change records to discover and define relationships between files within the system. These relationships, called association clusters, can then be utilized within the development process identify areas of risk in the system that are currently under change.

To show the efficacy of SRTP, a case study was performed with three releases of a product from IBM. The generated association clusters from the analysis were identifiable and directly relatable to specific requirements for each release or for an identified internal system component. The association clusters specifically illuminated areas of the code base where cross-component dependencies existed and components that included files that would not normally be examined in an analysis that used execution-based files, such as help files and configuration files. We performed a post hoc case study using this technique with three minor releases of a software product. We found that our methodology suggested additional regression tests in 50% of test runs and that the highest-priority suggested test found an additional fault 60% of the time.

The illustrative example in Section 3 and the case study in Section 4 use just one possible set of development artifacts. In these particular examples, only changes derived from faults were used in the analyses. However, more types of analyses could be performed with various other sets of software development artifacts. For example, a similar analysis to the previous examples could be performed with only field failures as opposed to the entire set of reported faults and failures. Comparing the results from an analysis using changes from faults found in testing and the results from an analysis using changes from field failures would provide insight into what files were being tested internally versus what files were being changed due to customer-reported failures. If parts of the two analyses resulted in completely orthogonal association clusters from each other, then

there is a possibility of a lack of testing coverage in the areas identified by the field failure analysis. We will investigate the use of various subsets of change information along with other software development artifacts. By creating association clusters based on various types of artifacts, traceabilities could be achieved between field failures, test cases, and the identified change association clusters that they all touch. Through these traceabilities, developers and managers could have a detailed view of the risky areas of their system at any given point, helping them to make development and release decisions.

## 6. Acknowledgments

We would sincerely like to thank J.B. Baker at IBM for his input into this work. We would also like to thank the Realsearch reading group for critiquing early versions of this research. Partial funding was provided for NCSU authors by the National Science Foundation.

## 7. References

- [1] ANSI/IEEE, "IEEE Standard Glossary of Software Engineering Terminology, Standard 729," 1983.
- [2] Ball, T., Kim, J., Potter, A., and Siy, H., "If your version control system could talk," Workshop on Process Modeling and Empirical Studies of Software Engineering, 1997.
- [3] Beyer, D. and Noack, A., "Clustering Software Artifacts Based on Frequent Common Changes," 13th IEEE International Workshop on Program Comprehension, St. Louis, MO, 2005.
- [4] Canfora, G. and Cerulo, L., "Impact Analysis by Mining Software and Change Request Repositories," International Software Metrics Symposium, 2005.
- [5] Demmel, J., *Applied Numerical Linear Algebra*. Philadelphia: Society for Industrial and Applied Mathematics, 1997.
- [6] Gall, H., Jazayeri, M., and Krajewski, J., "CVS Release History Data for Detecting Logical Couplings," Sixth International Workshop on Principles of Software Evolution, 2003.
- [7] Jalote, P., *Software Project Management in Practice*. New York: Addison Wesley Professional, 2002.
- [8] Kroll, P. and Kruchten, P., *The Rational Unified Process Made Easy: A Practitioner's Guide to the RUP*. Boston: Addison Wesley, 2003.
- [9] Kruchten, P., *The Rational Unified Process: An Introduction*, Third ed. Boston: Addison Wesley, 2004.
- [10] Livshits, B. and Zimmermann, T., "DynaMine: Finding Common Error Patterns by Mining Software Revision Histories," European Software Engineering Conference and Symposium on the Foundations on Software Engineering, Lisbon, Portugal, 2005.
- [11] Orso, A., Apiwattanapong, T., and Harrold, M. J., "Leveraging field data for impact analysis and regression testing," Symposium on the Foundations of Software Engineering, Helsinki, Finland, 2003.
- [12] Osinski, S., Stefanowski, J., and Weiss, D., "Lingo: Search Results Clustering Algorithm Based on Singular Value Decomposition," *Advances in Soft Computing, Intelligent Information Processing and Web Mining*, Zakopane, Poland, 2004.
- [13] Ren, X., Shah, F., Tip, F., Ryder, B., and Chesley, O., "Chianti: a tool for change impact analysis of Java programs," Conference on Object-Oriented Programming, Systems, Languages, and Applications, Vancouver, Canada, 2004.
- [14] Rothermel, G. and Harrold, M., "Analyzing Regression Test Selection Techniques," *IEEE Transactions on Software Engineering*, vol. 22, pp. 529-551, 1996.
- [15] Rothermel, G. and Harrold, M. J., "A Safe, Efficient Algorithm for Regression Test Selection," International Conference on Software Maintenance, Montreal, Canada, 1993.
- [16] Rothermel, G. and Harrold, M. J., "Empirical Studies of a Safe Regression Test Selection Technique," *IEEE Transactions on Software Engineering*, vol. 24, pp. 401-418, 1998.
- [17] Rothermel, G., Untch, R. H., and Harrold, M. J., "Prioritizing Test Cases For Regression Testing," *IEEE Transactions on Software Engineering*, vol. 27, pp. 929-948, 2001.
- [18] Sherriff, M., Lake, M., and Williams, L., "Empirical Impact Analysis using Singular Value Decomposition," North Carolina State University, Raleigh, NC CSC-TR-2007-13, April 29 2007.
- [19] Tip, F., "A survey of program slicing techniques," *Journal of Programming Languages*, vol. 3, pp. 121-189, 1995.
- [20] Will, T., "Introduction to the Singular Value Decomposition," vol. 2006: UW-La Crosse, 1999.
- [21] Ying, A., Murphy, G., Ng, R., and Chu-Carroll, M., "Prediction Source Code Changes by Mining Change History," *IEEE Transactions on Software Engineering*, vol. 30, pp. 574-586, 2004.
- [22] Zimmermann, T., Diehl, S., and Zeller, A., "Mining Version Histories to Guide Software Changes," *IEEE Transactions on Software Engineering*, vol. 31, 2005.