

Towards Understanding the Rhetoric of Small Source Code Changes

Ranjith Purushothaman
Server Operating Systems Group
Dell Computer Corporation
Round Rock, Texas 78682
ranjith_purush@dell.com

Dewayne E. Perry
Experimental Software Engineering Lab, ECE
The University of Texas at Austin
Austin, Texas 78712
perry@ece.utexas.edu

Abstract

Understanding the impact of software changes has been a challenge since software systems were first developed. With the increasing size and complexity of systems, this problem has become more difficult. There are many ways to identify the impact of changes on the system from the plethora of software artifacts produced during development and maintenance. We present the analysis of the software development process using change and defect history data. Specifically, we address the problem of small changes. The studies revealed that (1) there is less than 4 percent probability that a one-line change will introduce an error in the code; (2) nearly 10 percent of all changes made during the maintenance of the software under consideration were one-line changes; (3) the phenomena of change differs for additions, deletions and modifications as well as for the number of lines affected.

1. Introduction

Change is one of the essential characteristics of software systems [1]. The typical software development life cycle consists of requirements analysis, architecture design, low level design, coding, testing, delivery and finally, maintenance. Beginning with the coding phase and continuing with the maintenance phase, change becomes ubiquitous through the life of the software. Software may need to be changed to fix defects, to change executing logic, to make the processing more efficient, or to introduce new features and enhancements.

Despite its omnipresence, source code change is perhaps the least understood and most complex aspect of the development process. An area of concern is the issue of software code degrading through time as more and more changes are introduced to it – code decay [5]. While change itself is unavoidable, there are some aspects of change that we can control. One such aspect is the introduction of defects while making changes to software, thus preventing the need for fixing those errors.

A software change has different properties such as size, diffusion, type, duration, etc., and we are interested in studying the impact of the size and type of change on the likelihood of introducing faults.

Managing risk is one of the fundamental problems in building and evolving software systems. How we manage the risk of small changes varies significantly, even within the same company. We may take a strict approach and subject all changes to the same rigorous processes. Or we may take the view that small changes generally have small effects and use less rigorous processes for these kinds of changes. We may deviate from what we know to be best practices to reduce costs, effort or elapse times. One such common deviation is not to bother much about one line or other small changes at all. For example, we may skip investigating the implications of small changes on the system architecture; we may not perform code inspections for small changes; we may skip unit and integration testing for them; etc. We do this because our intuition tells us that the risk associated with small changes is also small.

However, we all know of cases where one line changes have been disastrous. Gerald Weinberg [9] documents an error that cost a company 1.6 billion dollars and was the result of changing a single character in a line of code. Gerard Holzmann in his keynote talk at FSE 2002[22] related similar one-line disasters.

In either case, innocuous or disastrous, we have very little actual data on small changes and their effects to support our decisions. We base our decisions about risk on experience, intuition and anecdotal evidence at best.

Our approach is different from most other studies that address the issue of software errors because we have based the analysis on the property of the change itself rather than the properties of the code that is being changed [7]. Change to software can be made by addition of new lines, modifying existing lines, or by deleting lines. We expect each of these different types of change to have different risks of failure.

Our first hypothesis is specific to one-line changes, namely that the probability of a one-line change resulting in an error is small. Our second hypothesis is that the failure probability is higher when the change involves adding new lines than either deleting or modifying existing lines of code.

To test our hypotheses, we used data from the change management and version management systems of a large scale software project. While historic data from project management systems have been used to analyze the various attributes affecting software development, the use of this data to study the impact of making one-line and other small changes to software has not been done before.

In the next section we provide a look at past research that has addressed issues related to our analysis. In section 3, we provide the background for this study, describing the change data and the methodology employed for our research. In section 4, we describe

our approach for the analysis of the changed lines and how we prepared the data. In section 5, we discuss the results of our analysis. We conclude the paper in section 6.

2. Literature Review

Software maintenance and evolution is the “final phase” of the software life cycle and is frequently viewed as a phase of lesser importance than the initial design and development phases. Quite the contrary, statistical data shows that maintaining two to ten year old software systems demands possibly as high as 40 percent to 70 percent of the total development effort [15]. We suspect that the number is actually much higher than that. Software maintenance and evolution still remains a difficult process to understand and manage.

Understanding the need for classification of the software changes, E. B. Swanson [12] proposed that changes be classified as belonging to three types of maintenance activities. The three types are corrective, adaptive, and perfective. As defined by Swanson, corrective maintenance is performed to correct defects that are uncovered after the software is brought to use. Adaptive maintenance is applied to properly interface with changes in the external processing environment and very often translates into new development and new features. Perfective maintenance is applied to improve non-functional characteristics – for example, to eliminate inefficiencies, enhance performance, increase fault tolerance and reliability, or improve maintainability.

Mockus and Votta [3] used the change history from the 5ESS™ switching software project to identify the reasons for software changes. In their analysis, changes were classified as corrective, adaptive, and perfective. They also introduced a fourth type of change classification: changes performed following inspections. Though the changes from inspections were mostly perfective and corrective changes, the number of such changes justified the introduction of a different type of change classification. In any systematic software development environment, code inspections and modifications of code following each inspection are standard procedures. Hence, for our results to be valid in such an environment and since our analysis was also based on the same data, we have retained the “inspection” type of change classification. Our research is based on Mockus’ and Votta’s [3] classification results. Indeed, we used their technique to automatically classify the types of changes and provide that change data for our study.

In his analysis, Les Hatton [17] related the defect frequency to file size. He stated that contrary to conventional wisdom that smaller components contain fewer faults, medium sized components are proportionally more reliable than small or large ones.

Analysts use both product measures such as the number of lines of code and process measures such as those obtained from the change history [10]. In their study looking for factors to predict fault incidence, Graves et al. [13] stated that, in general, process measures based on change history are more useful in predicting fault rates than product measures. They gave an example of how a

process metric such as the number of times the code has already been changed is a better indication of how many faults it will contain than its length (which is a product measure). Their study concluded that a module's expected number of faults is proportional to the number of times it has been changed.

Mockus and Weiss [7] studied the relation between the size of the change and probability of error and found that the failure probability increases with the number of changes, the number of lines of code added, and the number of subsystems touched. They also concluded that the probability of error is much higher for new development as compared to defect fixes because the change size associated with defect fixes tend to be much smaller in size. Dunsmore and Gannon [14] state that there is statistical evidence (Spearman $\rho = 0.56$ with $\alpha = .05$)¹ that shows a direct relationship between the amount of program changes and the error occurrences.

In the analysis done by Leszak et al. [2], the authors conclude that large changes to existing code are fault prone and provide statistical data to support their claim. They go a step further to propose that changes of more than 25 percent of existing code should be avoided and recommend recoding instead of modification. Basili and Perricone [18] categorize software modules based on their size (lines of code) and then checked for the errors at the module level. An interesting observation from their research was that, of the modules found to contain errors, 49 percent were categorized as modified and 51 percent as new modules.

Our primary contribution in this empirical research is an initial descriptive and relational study of small changes. As shown from our related research discussion above, we are the first to study this phenomenon. Another unique aspect of our research is that we have used a combination of product measures such as the lines of code and process measures such as the change history (change dependency) to analyze the data. In doing so, we have tried to gain the advantages of both measures while removing any bias associated with each of them.

While several papers discuss the classification of changes based on its purpose (corrective, adaptive, preventive) there is virtually no discussion based on the type of change: software can be changed by adding lines, deleting lines or by modifying existing lines. As a byproduct of our analyses, we have provided useful information that gives some insight into the impact that the type of change has on the software evolution process.

The 5ESS™ change and version history data that we use here has been used for various research purposes such as, for inferring change effort from configuration management databases [4], studying the impact of parallel changes in large scale software development projects [16], analyzing the challenges in evolving a large scale software product [6], identifying the reasons for

¹ Spearman's ρ is a correlation coefficient for data in rank form; 0.56 represents a medium effect.

software changes [3], and predicting fault incidence [13], to name a few. The wide range of studies that have used this particular change history data ensures good content validity for the results of the analysis based on this data.

3. The Context of the Study

The context of this study is Lucent Technologies' 5ESS™ telephone switch [22]. 5ESS itself is comprised of approximately 100,000,000 lines in C, C++ and a number of domain specific languages plus another 100,000,000 lines of header and make files. The system is organized into 50 subsystems, which are in turn organized into a total of 5000 modules (where a module is a directory of source code files). Each release consists of approximately 20,000,000 lines of code.

For this study we use the same 5ESS subsystem used in the Code Decay Project [5] – that is, the complete change and version management history from the first 15 years of that subsystem. This subsystem was built at a single development site within a single organization, the structure of which has varied over the years reaching a peak of 200 developers and eventually decreasing to the current 50 developers. The members of technical staff (MTS) who developed 5ESS typically have MS degrees. While their degrees range over a wide range, most are from scientific areas and all are subject to initial training courses in the telephony.

To coordinate this very large project, 5ESS has well-defined process available online with a process management organization responsible for their creation, evolution and improvement. The processes are described informally but in a highly structured manner and occupy over a 100 megabytes of disc space and are accessed over 500 times a day by the developers [23]. Some of these processes represent specific cases where small sized changes are exempt from the standard processes. For example, the architecture process is such a case: small changes are not given the full architectural review². In other cases, they have recognized that the size of features (the basic unit of development) is important enough to differentiate the design process according to size [24].

In this section, we describe the change process in the 5ESS software development project and also give an introduction to the product subsystem that we use for our analysis.

3.1. The 5ESS Change Process

² One of the authors was a member of the team at AT&T (before the split between Lucent and AT&T) defining the requirements for best in practice processes in which it was required that all changes regardless of size be given a full review. Those responsible for the 5ESS architecture process argued that the recommendation ought to be change to reflect their practice of not reviewing small changes to the system.

In the 5ESS, a *Feature* is the fundamental unit of system functionality to be added to the system. Each Feature is implemented by a set of Initial Modification Requests (IMRs) where each IMR represents a logical problem to be solved and as such are the basic unit of development. Features range in size from those requiring only one developer to those requiring development teams in the hundreds, and may be either software only or a combination of software and hardware. Features then may require anywhere from one IMR to hundreds of IMRs

Each IMR is implemented by a set of functionally distinct Modification Requests (MRs) where each MR represents a logical part of an IMR's solution (and a developer may further subdivide a solution into multiple MRs if they are seen as multiple logical sets of changes). Multiple developers may well be involved in solving an IMR; however, only one developer is assigned to (i.e. "owns") a specific MR. Of course, a developer may be responsible for more than one MR, and may easily work on more than one IMR. The solutions represented by MRs (to the problems described by the IMRs) are reified in terms of changes to source code by changes to one or more files. Thus there are three logical units of change: Feature, IMR and MR.

The change history [3,5,7] of the system is maintained using the IMR Tracking System (IMRTS), the Extended Change Management System (ECMS [25]) for initiating MRs and tracking changes and the Sources Code Control System (SCCS [26]) for managing different versions of the files. Features and their associated IMRs are managed by IMRTS and include information about the associated MRs. ECMS records information about each MR and its relationship to SCCS's versions. Each IMR has an extensive set of information (89 fields in all). Each MR includes such data as the parent IMR, dates, the affected files, and text describing the changes and their rationale.

The actual source code changes themselves are maintained by SCCS in the form of one or more *deltas* (the physical unit of change) depending on the way the changes are committed by the developer. Each delta provides information on the attributes of the change: the parent MR, lines added, lines deleted, lines unchanged, login of the developer, and the time and date of the change.

While it is possible to make all changes that are required to be made to a file by an MR in a single delta, developers often perform multiple deltas on a single file for an MR. Multiple deltas are common because of the amount of parallel changes in building a new release [16]. Because of the size of the system together with the fact that as the system ages an increasingly large portion of the system remains unchanged [27], there has been a shift away from code ownership in evolving the 5ESS system. The unit of development is now a Feature (on average about 500 new Features per release) and these proceed in parallel with multiple MRs affecting individual files. Since SCCS does not support parallel versions, quasi-parallel versions are represented as a series of sequential but interleaved deltas. A series of deltas for a single MR is thus an artifact of parallel changes in a version management system that does not support parallel versions. It is for this reason that we consider an MR to represent one logically physical change per file.

3.2. 5ESS Change Data

The 5ESS™ subsystem we use for this study contains 4550 modules that have a total of nearly 2 million lines of code. Over the last decade, the OA subsystem had 31884 MRs that changed nearly 4293 files. So nearly *95 percent of all files were modified* after first release of the product.

Change to software can be introduced and interpreted in many ways. However, our definition of change to software is driven by the historic data that we used for the analysis: A change is *any alteration to the software recorded in the change history database* [5]. In accordance with this definition, in our analysis the following were considered to be changes:

- One or more modifications to single/multiple lines;
- One or more new statements inserted between existing lines;
- One or more lines deleted; and,
- A modification to a single/multiple lines accompanied by insertion or/and deletion of one or more lines.

Since we are looking at the phenomena of small changes specifically in this study, the definition of a one-line changes is of critical importance since there are several ways in which we might view them.. For example, we consider one line changes with respect to a delta, a file or an MR. But as we have already discussed, multiple deltas are an artifact of the tools used that forces interleaving of deltas for an MR and hence we view the definition in terms of deltas as unsatisfactory. Similarly, one could make an argument for viewing the one line changes on a per file basis. There is some merit in this choice, but it ignores the fact that an MR represents a set of related changes for one specific solution. Hence we decided that an MR, a unit of logical change, is the most conservative choice. The following kinds of changes, then, would qualify as a one-line change for this study when an MR consists of either:

- One or more modifications to a single line;
- One or more lines replaced by a single line (i.e, multiple lines deleted, one line added);
- One new statement inserted between existing lines; or,
- One line deleted.

Previous studies such as [14] do not consider deletion of lines as a change. However, from preliminary analysis, we found that lines were deleted for fixing bugs as well as making modifications. Moreover, in the SCCS system, a line modification is tracked as a line deleted and a line added. Hence in our research, we have analyzed the impact of deleting lines of code on the software development process.

4. Approach

In this section, we document the steps we took to obtain useful information from our project database. We first discuss the preparation of the data for the analysis and then explain some of the categories into which the data is classified. The final stage of the analysis identifies the logical and physical dependencies that exist between files and MRs.

4.1 Data Preparation

The change history database provides us with a large amount of information. Since our research focuses on analyzing one-line changes and changes that were dependent on other changes, one of the most important aspects of the project was to derive relevant information from this data pool. As mentioned above, multiple deltas per MR is an physical artifact resulting from the version management system SCCS lacking support for parallel changes in terms of logically parallel versions and are merged when the changes are completed. The logically parallel versions are represented as a sequence of interleaved deltas.

In the change process hierarchy, an MR is the lowest logical level of change. Hence if the MR was created to fix a defect, all the modifications that are required by an MR would have to be implemented to fix the bug. Hence we were interested in change information for each effected file at the MR level not at the delta level. For example, in Table 1, the MR *oa101472pQ* changes two files. Note that the file *oaMID213* is changed in two steps. In one of the deltas, it modifies only one-line. However, this cannot be considered to be a one-line change since for the complete change, the MR changed 3 lines of the file. Thus we have had to process the delta data to merge multiple deltas into one logical change to each file in an MR

With nearly 32000 MRs that modified nearly 4300 files in the subsystem under investigation, the aggregation of the changes made to each file at the MR level gave us 72258 change records for analysis. We note that with this number of samples, what appears to be significantly different in the distributions shown below is likely to be significantly different. However, for completeness' sake, we give the results of χ^2 analyses to show the extent to which the differences are indeed significant.

DELTA relation				
MR	FILE	Add	Delet	Date
			e	
Oa101472pQ	oaMID213	2	2	9/3/1986
Oa101472pQ	oaMID213	1	1	9/3/1986
Oa101472pQ	oaMID90	6	0	9/3/1986
Oa101472pQ	oaMID90	0	2	9/3/1986

Table 1: Delta relation snapshot

4.2. Data classification

Change data can be classified based on the purpose of the change and also based on how the change was implemented. The classification of the MRs based on the change purpose was derived from the work by Mockus and Votta [3] (and our data classified with help from Mockus). They classified MRs based on the keywords in the textual abstract of the change. For example, if keywords like 'fix', 'bug', 'error', and 'fail' were present, the change was classified as corrective. In Table 2 we provide a summary of the change information classified based on its purpose. The naming convention is similar to the work done in their original paper. Perfective changes are typical made to improve performance, make a piece of code more maintainable, or generally improve the quality of the code. Adaptive changes are those in which new features and functionality are added to the system.

However, there were numerous instances when changes made could not be classified clearly. For example, certain changes were classified as 'IC' since the textual abstract had keywords that suggested changes from inspection (I) as well as corrective changes

(C). Though this level of information provides for better exploration and understanding, in order to maintain simplicity, we made the following assumptions:

- Changes with multiple 'N' were classified as 'N'
- Changes with multiple 'C' were classified as 'C'
- Changes containing at least one 'I' were classified as 'I'

ID	Change Purpose	Explanation
B	Corrective	Fix defects
C	Perfective	Enhance performance
N	Adaptive	New development
I	Inspection	Following inspection

Table 2: Change Purpose

Changes which had 'B' and 'N' combinations were left as 'Unclassified' since we did not want to corrupt the data. Classification of these as either a corrective or adaptive change would have introduced validity issues in the analysis. Based on the above rules, we were able to classify nearly 98 percent of all the MR into corrective, perfective, adaptive or inspection changes.

Figure 1 compares the distribution of one-line changes against the total set of changes classified according to purpose. For the complete set of changes, the largest set of changes is due to adaptation. This is not surprising when considering the fact that each release of the system incorporates a significant number of new features. The second largest set of changes is due to corrections. Again this is not surprising considering the official development and evolution processes begin tracking defects as soon as testing begins. Perfective and inspection sets of changes are respectively relatively small. Conversely for one-line changes, corrective changes are the most frequent while adaptive are slightly less frequent.




Figure 1: **Distribution of changes based on purpose (series 1: all changes; series 2: one-line changes)**

The distributions shown in Figure 1 of one-line changes versus all the changes are significant. In the χ^2 analysis shown in Table 3, the value of $\chi = 234.995$ with 4 degrees of freedom and with $p \leq .001$. One-line corrective changes are significantly higher than expected while one-line adaptive changes are significantly lower than expected. Thus, compared to the complete set of changes, one-line changes are used to fix faults significantly more often but are used for adaptive purposes significantly less often.

		all lines	one-line
corrective	expected	24049	1897
	observed	23687	2259
	chi-square	5.438	68.919
perfective	expected	5801	458

significance	chi-square
<.2	5.99
<.1	7.78
<.05	9.49
<.025	11.14

	observed	5764	495
	chi-square	0.24	3.038
adaptive	expected	34017	2684
	observed	34557	2144
	chi-square	8.203	108.599
inspection	expected	6612	522
	observed	6473	661
	chi-square	2.935	37.197
unclassified	expected	1779	140
	observed	1777	142
	chi-square	0.002	0.02

<.01	13.28
<.001	18.47

total chi-square	df
234.995	4

Table 3: χ^2 analysis of the purpose of one line changes versus the entire set of changes

Another way to classify changes is on the basis of the implementation method into insertion, deletion, or modification. But the SCCS system maintains records of only the number of lines inserted or deleted for the change and not the type of change. Modifications to the existing lines are tracked as old lines being replaced by new lines (delete and insert). However, for every changed file SCCS maintains an SCCS file that relates the MR to the insertions and deletions made to the actual module. Scripts³ were used to parse these SCCS files and categorize the changes made by the MR into inserts, deletes or modifications. Table 4 lists different types of changes based on their implementation method. Inserts and deletes are straightforward; modifications are accomplished by a combination of delete and inserts (An analog to an SCCS file is that produced by the `diff` command in UNIX™). The scripts parse and analyze the records into the various classifications listed in Table 4.

ID	Change Type	Description
----	-------------	-------------

³ The scripts made available for our research were those used in the Code Decay project at Bell Labs (Perry was a member of that project) and were used for a variety of studies in that project as well as other analysis and visualization (see for example [32]) studies in the Software Production Research Project, Bell Labs, Lucent Technologies. In this case the scripts were provided by Tom Ball, now of Microsoft Research.

C	Modify	Change existing lines
I	Insert	Add new lines
D	Delete	Delete existing lines
IC	Insert/Modify	Inserts and modifies lines
ID	Insert/Delete	Inserts and deletes lines
DC	Delete/Modify	Deletes and modifies lines
DIC	All of the above	Inserts, deletes and modifies lines

Table 4: ChangeType

Figure 2 compares the distribution of one-line changes against all the changes classified by change type – i.e., by the means of change implementation. Given that, for all changes, the most frequent purpose of changes to the system is adaptation, it is not surprising that the most frequent means of implementing changes is by inserting lines. Change and the various combinations are about equal in distribution. Only a very small set of changes were accomplished solely by deleting lines. This low percentage very likely stems from two reasons⁴. First, actual releases are configured from a large set of features and customized for specific uses, so there is little motivation to remove functionality. Second, very little functionality is removed for fear of breaking the system.

⁴ This is based on anecdotal evidence gathered by one of the authors in working with 5ESS developers




Figure 2: Distribution of changes based on type (series 1: all changes; series 2: one-line changes)

For one-line changes the distribution is significantly different from the entire set of changes. In the χ^2 analysis, the value of $\chi = 1615.098$ with 2 degrees of freedom and with p significantly less than .001. Modifications are used significantly more often in one-line changes than in the total set of changes, while inserts are used significantly less often. Lines are deleted significantly more often in one-line changes. The way one-line changes are done, then, is significantly different than the total set of changes.

		all lines	one-line
modify	expected	21592	2402
	observed	20286	3708

significance	chi-square
<.2	3.22
<.1	4.61

	chi-square	79.004	710.206		<.05	5.99
insert	expected	27535	3063		<.025	7.38
	observed	28974	1624		<.01	9.21
	chi-square	75.204	676.05		<.001	13.82
delete	expected	2365	263			
	observed	2232	396			
	chi-square	7.471	67.163			
				total chi-square		df
				1615.098		2

Table 5: χ^2 analysis of the types of one line changes versus all changes

4.3 Identifying file dependencies

Our primary concern was in isolating those changes that resulted in errors. To do so, we identified those changes that were *dependencies – changes to lines of code that were changed by an earlier MR*. If the latter change was a bug fix our assumption was that the original change was in error. There are two ways in which the original change might be in error: an error by commission – i.e., the change was incorrect in that it introduced an error with the change; an error by omission – i.e., the change was incorrect in that it didn't make sufficient changes to solve the problem.

The one argument against the validity of this assumption would be that the latter change might have fixed a defect that was introduced before the original change was made. However, in the absence of prima facie evidence to support either case, and since preliminary analysis of some sample data did not support the challenging argument, we ruled out this possibility. In this report, we will refer to those files in which changes were made to those lines that were changed earlier by another MR as *dependent files*.

The dependency, as we have defined earlier, may have existed due to bug fixes (corrective), enhancements (perfective), changes from inspection, or new development (adaptive). 2530 files in the subsystem we studied were found to have dependent changes. That is nearly 55 percent of all files in the subsystem and nearly 60 percent of all changed files. So, *in nearly 60 percent of cases, lines that are changed were changed again*. This kind of information can be very useful to the understanding of the maintenance phase of a software project. We had 51478 dependent change records and this data was the core of our analysis.

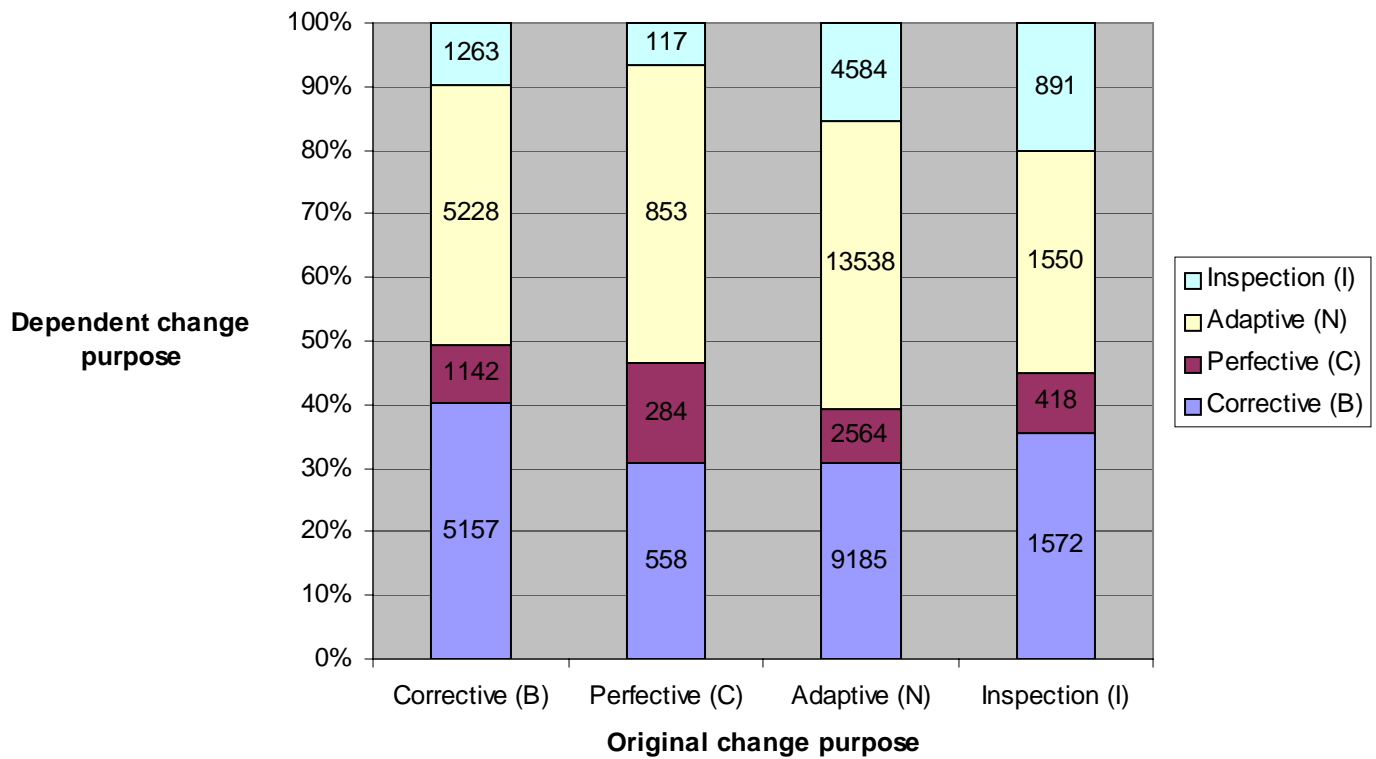


Figure 3: **Distribution of change classification on dependent files**

In Figure 3, we show the distribution of change classifications of the dependent files across the original files. The horizontal axis shows the types of changes made to the dependent files originally. In the vertical axis, we distribute the new changes based on their classification based on the purpose type. We note first that the bulk of dependent changes to all categories were adaptive and corrective. We note second that the largest percentage of bug fixes was made to code that was already changed by an earlier MR to fix bugs -- that is *roughly 40 percent of changes made to fix errors introduced more errors* (as shown by the fact that 40% of dependent changes made to corrective changes were corrective).

It is also interesting to note that nearly 40 percent of all the dependent changes were of the adaptive type and most perfective changes were made to lines that were previously changed for the same reason, i.e., enhancing performance or removing inefficiencies.

corrective	perfective	adaptive	inspection
------------	------------	----------	------------

significance	chi-square
--------------	------------

inspection	expected	1793	254	4187	621
	observed	1263	117	4584	891
	chi-square	156.568	73.888	37.624	117.28
adaptive	expected	5536	784	12930	1918
	observed	5228	853	13538	1550
	chi-square	17.178	6.007	28.569	70.621
perfective	expected	1153	163	2692	399
	observed	1142	284	2564	418
	chi-square	252.47	89.16	6.128	0.867
corrective	expected	4308	610	10061	1492
	observed	5157	558	9185	1572
	chi-square	167.331	4.486	76.313	4.239

<.2	12.24
<.1	14.68
<.05	16.93
<.025	19.02
<.01	21.67
<.001	27.88

total chi-square	df
856.36	9

Table 6: χ^2 analysis of dependent versus original change purpose.

As one would expect, the structure of the distribution is significant – i.e, the two populations are not independent. In the χ^2 analysis shown in Table 6, the value of $\chi = 856.36$ with 9 degrees of freedom and with $p \leq .001$. Interestingly, the number of dependent corrective changes made to corrective changes is significantly higher than would be expected, while it is significantly lower than expected for adaptive changes. The number of dependent perfective changes is significantly higher than expected for original perfective changes. Adaptive dependent changes were significantly higher than expected for adaptive changes, but significantly lower than expected for corrective changes. All dependent inspection changes were significant: lower than expected for corrective and perfective changes, higher than expected for adaptive and inspection changes.

The original corrective changes had significantly more than expected dependent corrective changes but fewer dependent adaptive and inspection changes. The original set of perfective changes had significantly more dependent perfective changes and fewer inspection changes than expected, while the original set of adaptive changes had significantly fewer dependent corrective changes and significantly more dependent adaptive and inspection changes. Original inspection changes had significantly fewer dependent adaptive changes and significantly more dependent inspection changes.

While not necessarily the largest magnitude in each case, each change purpose had more dependent changes of the same purpose than expected.

5. Analysis

The analysis of the data proceeds in several steps. We begin with an investigation of the software project based on the change size. We then look at the problem of changes introducing defects and complete our analysis by considering the issues surrounding change process metrics.

5.1. Change size

Change size is an effective way to estimate the change effort in a software development project. From our analysis, we were able to derive meaningful information that gives a measure of the number of lines that are changed as part of an MR. Figure 3 shows the distribution of the changed files based on the number of lines per MR that were changed. The vertical axis shows the percentage of changed files that changed the number of lines specified on the horizontal axis. Note that one line changes (i.e., one line in one MR) are of necessity changes to only one file. MRs with multiple lines changed may affect (and often do) affect more than one file. It is for this reason that it is useful to relate the number of lines changed per MR to the number of files changed.

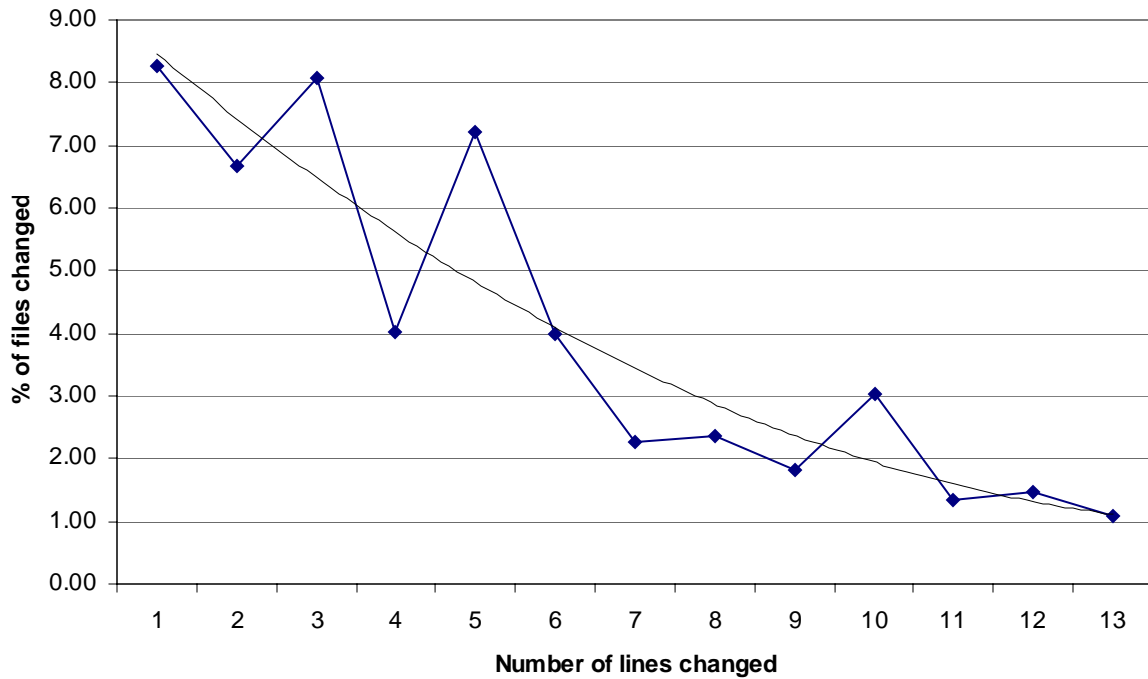


Figure 4: **Distribution of small changes**

Figure 4 plots the number of lines changed (horizontal axis) against the percentage of files changed with that number of lines (vertical axis). We can see that *nearly 10 percent of files changed involved changing only a single line of code*. Since the data fluctuated slightly, we did a second degree polynomial regression analysis of the data as shown by the regression line in the figure. From the regression line obtained, we can see that percentage of effected files reduces as the size of the change increases. *Nearly 50 percent of all changed files involved changing 10 lines of code or less.*

So, though the effort for changing one-line of code is generally smaller, the magnitude of these changes is very large in the software evolution process. However, it has been found that developers tend to give less priority to smaller changes and especially one-line changes. To illustrate further, Figure 5 shows the distribution of all the changed files in the subsystem under study across their change sizes. From this figure, we note that *nearly 95% of all changes were those that changed less than 50 lines of code.*

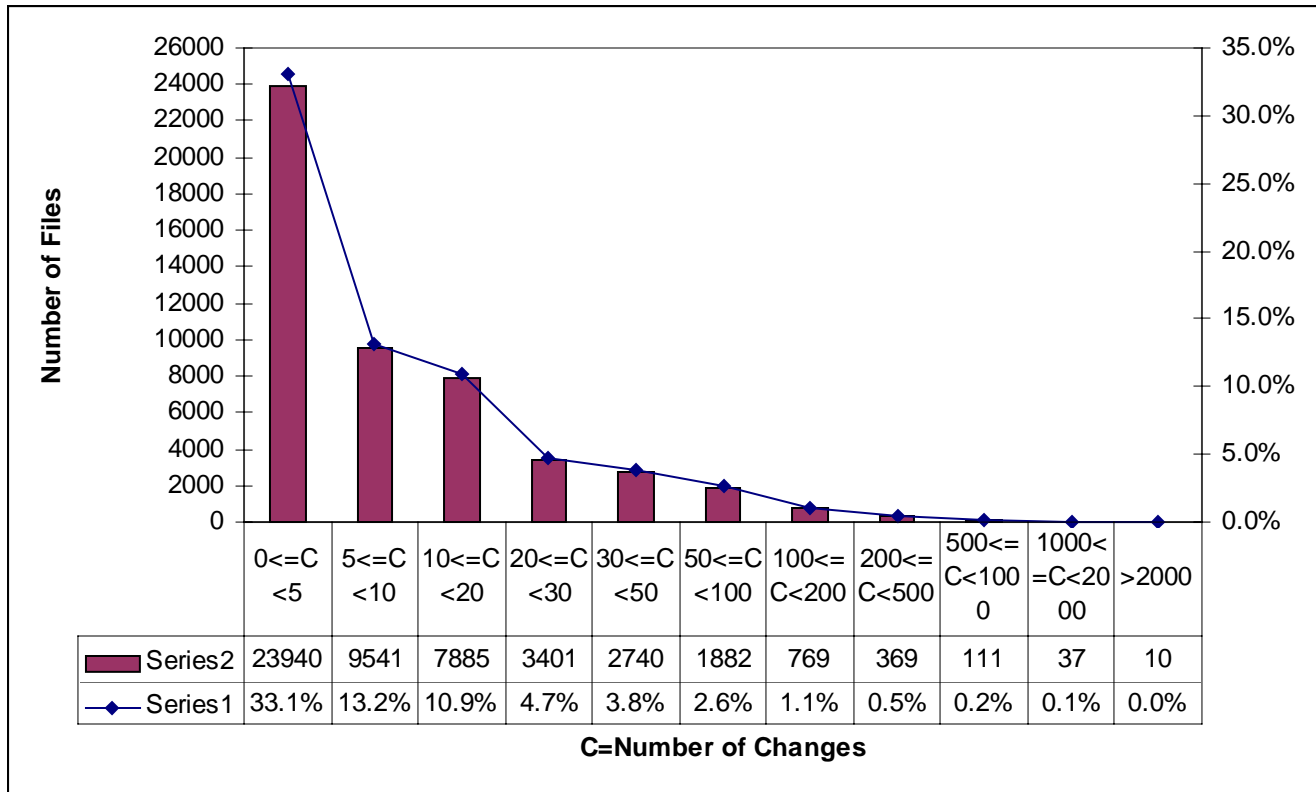


Figure 5: Change size distribution across files

5.2. Erroneous changes

We next analyze those changes that resulted in errors. In Figure 6, we present the data for erroneous changes that affected less than 10 lines of code. The vertical axis gives the percentage of changes that resulted in error out of the total changes that affected the number of lines specified in the horizontal axis. The data was derived from the change file dependencies that we had defined in an earlier section of this paper. This analysis also answers a very important question: What percentage of one-line changes result in error? *Less than 4 percent of one-line changes result in error* (see Figure 6).

It may also be noted that the changes tend to be more erroneous as the number of lines changed increases. One possible explanation to this behavior can be that as the number of lines that are changed increases, there are more avenues provided for the developer to make mistakes. These increased opportunities to introduce errors are likely due to an increase in the number of possible interactions with other lines of code.

We mentioned earlier the classification of changes based on their type into changes by insertion, deletion, and modification. We thought it would be a useful metric to analyze the distribution of erroneous changes based on the type of change. Figure 7 shows

the results of this analysis. Changes made by deletion of lines have been excluded since our analysis *did not produce any credible evidence that deletion of less than 10 lines of code resulted in errors.*




Figure 6: Errors introduced by change




Figure 7: Erroneous changes by type of change

From Figure 7, we note that while the probability that an insertion of a single line might introduce an error is 2 percent, there is nearly a 5 percent chance that a one-line modification will cause an error. It can also be seen that while modified lines cause more errors when less than 5 lines are changed; inserted new lines introduce more errors with larger change sizes.

To emphasize this behavior, in Figure 8, we have shown the distribution of the probability of error introduced by change over a wider range of change sizes. It may be noted that there is nearly 50 percent chance of at least one error being introduced if more than 500 lines of code are changed. The trend of the lines for change implemented by lines inserted and modified clearly shows that insertion of new lines generates a lot more errors when the change size is higher. One plausible explanation for this may be that developers tend to be more cautious when existing code has to be modified than when new development is done.




Figure 8: Erroneous changes versus change size

5.3. Change Process Metrics

How are the types of change related to change classifications? In Figure 9, the vertical axis categorizes changes based on their purpose and the horizontal axis classifies changes based on how the change was implemented. As expected, the largest number of lines was inserted for adaptive changes since new development involves addition of new lines of code. Modifications were made to existing lines of code equally for both adaptive and corrective changes.

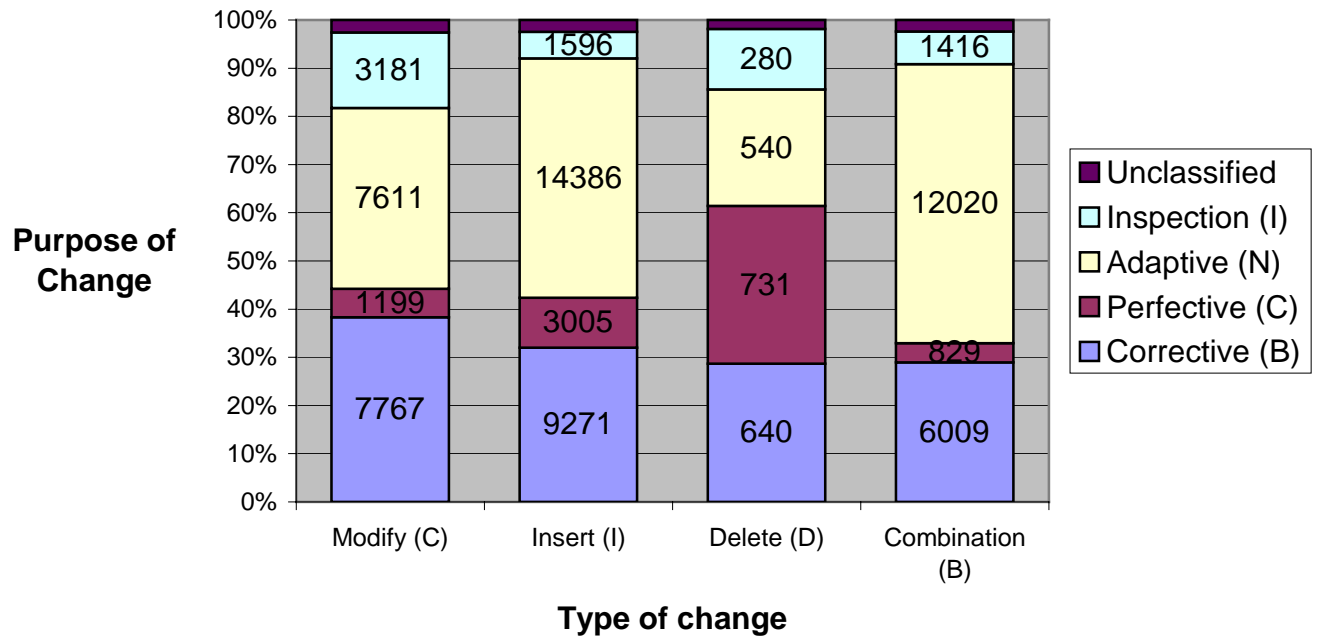


Figure 9: Relation between change classification and change type

		corrective	perfective	adaptive	inspection	unclassified	significance	chi-square
modify	expected	6762	1605	9622	1802	495	<.2	15.81
	observed	7767	1199	7611	3181	528	<.1	18.55
	chi-square	149.179	102.649	420.215	1054.68	2.231	<.05	21.03
insert	expected	9658	2292	13743	2574	707	<.025	23.34
	observed	9271	3005	14386	1596	716	<.01	26.22
	chi-square	15.538	221.648	30.128	371.698	0.123	<.001	32.91
delete	expected	744	177	1059	198	54		
	observed	640	731	540	280	41		
							total chi-square	df

	chi-square	14.546	1740.753	254.096	33.661	3.317
combination	expected	7122	1690	10134	1898	521
	observed	6609	829	12020	1416	492
	chi-square	36.992	438.898	350.986	122.512	1.627

5365.577	12
----------	----

Table7: χ^2 analysis of type versus pupose of change distribution

We can see that the Figure 9 holds no surprises except maybe that deletion of lines occurred pretty much uniformly for adaptive, corrective and perfective changes. Note, however, that there higher percentage of deleted lines than modified, inserted and combined in perfective evolution. In the χ^2 analysis shown in Table 7, the value of $\chi = 5365.58$ with 12 degrees of freedom and p significantly less than .0001. While the overall distribution is significant, the unclassified changes are fairly independent – i.e, the observed and expected are not significantly different.

Among the significant facts about the structure were the following. Corrective changes were modified more that would be expected while a combinations were applied less often then expected. Perfective changes had 30% more insertions and 400% more deletions than were expected and half the combinations and a quarter the modifications. Adaptive changes had fewer modifications and half the deletions expected but more combinations. Everything about inspection changes: ~75% more modifications than expected, half the insertions expected, more deletions, and fewer combinations.

Thus there are the following significant tendencies: to make more modifications to corrective and inspection changes and fewer to perfective and adaptive changes; to insert approximately the expected number of changes for corrective and adaptive changes, but insert more lines in perfective but fewer lines in inspection changes; to delete more lines than expected for perfective and inspection changes, fewer for perfective; to apply fewer combinations to corrective, perfective and inspection changes, and more to adaptive.

Figure 10 continues this discussion but restricts the change data to only one-line changes. The similarity of the data distribution in the two figures show that the behavior of one-line changes at least in regard to their distribution among the change types is representative of the behavior of changes irrespective of the size of the change. The only notable difference between the data in Figure 9 and Figure 10 is in the case when new single lines are inserted – less than 2.5 percent of one-line insertions were for perfective changes compared to nearly 10 percent of insertions towards perfective changes when all change sizes were considered.

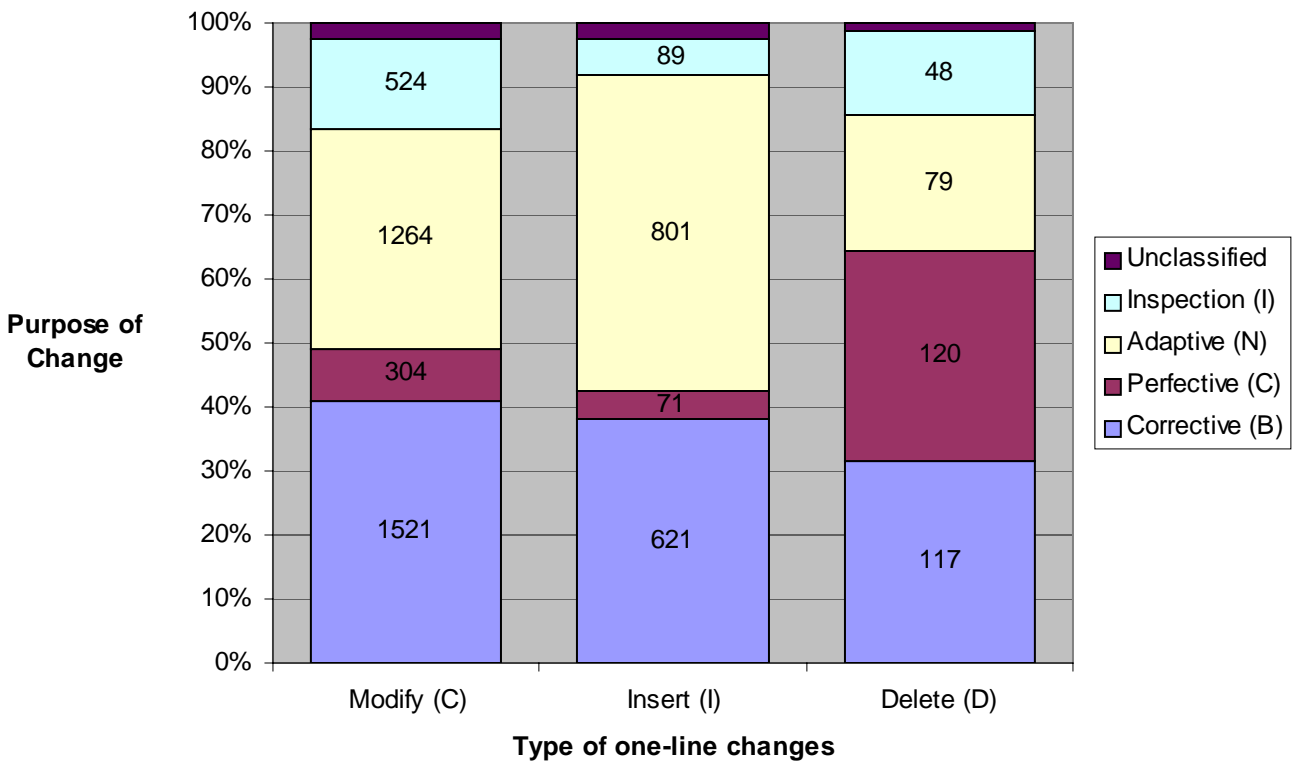


Figure 10: **Relation between various change types for one-line changes**

In the χ^2 analysis shown in Table 8, the value of $\chi = 458.54$ with 8 degrees of freedom and $p \leq .00$ showing a significant distribution. The one line changes made to unclassified were fairly independent (i.e., the expected and observed values were not significantly different). Modifications were close to expected for corrective, perfective and adaptive one lin changes, but more than expected for inspection changes. Insertions were about half of what was expected for perfective and inspections changes, but about 30% more than expected for adaptive changes. Deletions were significantly more for perfective (400%) and adaptive (300%). There was nothing significant about the type of changes to corrective changes and the types of changes were pretty much as expected.. Fewer lines were inserted than expected and more lines deleted than expected in perfective changes, while there were more lines inserted and deleted than expected in adaptive changes. For inspection changes, more lines were modified and fewer inserted than expected.

corrective	perfective	adaptive	inspection	unclassified	significance	chi-square
------------	------------	----------	------------	--------------	--------------	------------

modify	expected	1469	322	1394	430	92	<.2	11.03
	observed	1521	304	1264	524	95	<.1	13.36
	chi-square	1.821	1.001	12.21	20.586	0.076	<.05	15.51
insert	expected	644	141	611	188	40	<.025	17.53
	observed	621	71	801	89	42	<.01	20.09
	chi-square	36.48	34.759	59.267	52.361	0.059	<.001	26.12
delete	expected	146	32	139	43	9		
	observed	117	120	79	48	5		
	chi-square	5.837	241.49	25.745	0.636	1.911		
							total chi-square	df
							458.544	8

Table 9: χ^2 analysis of type versus purpose of change distribution of one-line changes

6. Analysis Summary

We have found that the probability that a one-line change would introduce one error is less than 4 percent. This result supports the typical risk strategy for one line changes and puts a bound on our search for catastrophic changes. Interestingly, this result is very surprising considering the anecdotal claim of a software guru that “*one-line changes are erroneous 50 percent of the time*”. This large deviation may be attributed to the structured programming practices, and development and evolution processes involving code inspections and walkthroughs that were practiced for the development of the project under study. Earlier research [9] shows that without proper code inspection procedures in place, there is a very high possibility that one-line changes could result in error.

We have also provided key insights that can be very useful for better understanding the software development and evolution process. In summary, some of the more interesting observations that we made during our analysis include:

- Nearly 95 percent of all files in the software project were maintained at one time or another. If the common header and constants files are excluded from the project scope, we can conclude that nearly 100 percent of files were modified at some point in time after the initial release of the software product.
- Nearly 40 percent of changes that were made to fix defects introduced one or more other defects in the software.

- Nearly 10 percent of changes involved changing only a single line of code; nearly 50 percent of all changes involved changing 10 or fewer lines of code; nearly 95% of all changes were those that changed fewer than 50 lines of code.
- Less than 4 percent of one-line changes result in error.
- The probability that the insertion of a single line might introduce a defect is 2 percent; there is nearly a 5 percent chance that a one-line modification will cause a defect. There is nearly a 50 percent chance of at least one defect being introduced if more than 500 lines of code are changed.
- Less than 2.5 percent of one-line insertions were for perfective changes, compared to nearly 10 percent of insertions towards perfective changes when all change sizes were considered.
- The maximum number of changes was made for adaptive purposes, and most changes were made by inserting new lines of code.
- There is no credible evidence that deletions of fewer than 10 lines of code resulted in defects.

To fully understand these effects of small changes in particular, and changes in general, this study should be replicated across systems in different domains and of different sizes.

7. Discussion of Threats to Validity

There are three types of validity that must be considered in this descriptive and relational study: construct, internal, and external validity.

Construct Validity. We divide construct validity into three parts [34]: intentional validity, representation validity and observational validity. In terms of intentional validity, our constructs are well understood and agreed upon in the general context in which this research has been done, and represent what we intend to study.

It is in terms of representational validity that we have a weakness with our change construct in the context of corrective dependent changes. There are three ways a change can be made to fix a fault: correct the line that was changed; change another line to bring the semantics of the change into full compliance with the originally intended change; or compensate for the fault by fixing another piece of code entirely. For dependent changes we measure only the first of the three representations of change. The remaining two cases would require more information about the relationship amongst changes that is either available or automatically inferable.

For our definition of dependent changes, our observational validity is very good. Our measure is completely accurate in examining and reporting the version delta data. However, our automated measure for determining the purpose of changes is not as reliable as we would like, but workable in this context.

Internal Validity. Most of the data we use in this study is explicit in the change and version management systems and can be considered to be completely reliable. The straightforward presentation of the data with a minimum of manipulation strengthens the internal validity for the study.

There is some slight concern about the magnitude of one-line changes because we count only one of the three ways of correcting a fault. However, we argue that while the net affect is that the number of one-line changes automatically determined may be less than the actual number of one-line changes, it does not significantly affect analyses. First, we believe that the number of compensatory changes are very few in number and hence insignificant in our study. Second, if we assume that one-line changes in error only require one line changes to fix them, then only 4.5% of the one line changes are not accounted for in our definition of dependent changes (5701 dependent one line changes from a total population of 5969).

Another threat to internal validity in the data we used is the reliability of the automatic classification mechanism. Given that this is always an issue even when using human judgments and that the reliability reported by Mockus and Votta is fairly good, we do not see this as a significant issue.

External Validity. Meltzoff [33] distinguishes between external validity and generalizability. While external validity and generalizability are often considered to be synonymous, there is merit in considering them to be separate concepts with external validity focused on the consistency of the purpose of, the intent of, or the claims made about a study and the supporting evidence provided by that study. Trying to show that something exists, or to understand and to describe a particular phenomenon (such as the study of small source code changes here) is different from claiming that results generalize to all possible domains of software systems. With respect to this view of external validity, we claim strong external validity:- namely, that our results strongly support the purpose and intent of our study. The purpose of this study was to illuminate the rhetoric of small changes in the context of 5ESS and to understand the relationships between the size of, the types of, the purpose of, and the qualitative results of these changes.

It is with respect to generalizability that we cannot make claims as strongly as we would like. The subsystem used for this study is representative of the various subsystems of 5ESS and thus can be used a surrogate for the entire system. (cf. [6, 16]). Further we claim that 5ESS is representative system for large-scale, highly reliable, real-time systems for the following reasons⁵. First, it is

⁵ Our generalizability claim here is consistent with a variety of studies published on 5ESS – cf. [2,3,4,5,6,7,13,16,24]

built with programming languages (C and C++) and a development environment (UNIX™) commonly used for building and evolving these kinds of systems. Second, these kinds of systems cannot be successfully built without the well-educated and well-trained staff, the well-defined processes, as well as the elaborate management, evolution and development environmental support found in the 5ESS project. A mature organization with mature processes and technology support is a necessary precondition to building and these kinds of systems. One factor about telephony that may not apply in these systems, however, is the amount of new features added with each release. This difference may diminish the generalizability of our results to reliable real-time systems.

The primary problem in generalizability lies in the fact that it is not clear whether these results apply to smaller systems and systems of different domains and applications. Given the size and complexity of the system, we can certainly argue that the problems found here are at least as severe as any found in smaller systems or systems in other domains. The limits of generalizability will only become clear as studies such as this are replicated across a wide variety of domains and systems.

8. Issues in Replication

In this study we used data gathered from two repositories: the change management system (ECMS) and the version management system (SCCS). Modern version or configuration management systems such as CVS [28] and ClearCase [29] go beyond the basic functionality of managing version histories and include some of the change management information functions we found in ECMS. The data comparable to what we used that is readily available in these modern version management systems include

- the standard version management data
 - version lineage: current version, source version
 - delta information: lines added, deleted, changed
 - time of the changes
- change management information
 - information about the person responsible for the change
 - text explaining the change and the rationale for it

The standard version management data enables the size and kinds of changes to lines of code to be calculated and classified. Moreover the version management data can be used to calculate the dependency data as we have done here. The text explaining the changes and the reasons for those changes can be used to analyze and to classify the purposes of those changes either automatically

or manually. One can use the mechanism of Mockus and Votta [3] derived data from and applied to 5ESS (as we did here). Alternatively one can use the more recent mechanism of Hassan and Holt [30,31] used on CVS-managed open source project data. The two mechanisms are comparable with the latter achieving a somewhat higher level of agreement with manual classification than the former.

Without an MR mechanism it would be difficult to relate logical changes to physical changes as we have done here. Using files as both the logical and physical unit of change instead of MRs would provide congruent data. Thus, exact replication would require the same logical data we used from ECMS. However, there are sufficient data available to perform studies sufficiently congruent to continue exploring the rhetoric of small changes in particular and the phenomena of evolution in general.

8. Recommendations

There are several types of recommendations that stem from this research and may be useful in practice for the building and evolving of software systems.. The first set of recommendations is for version management systems. 1) The notion of a logical set (provided in the case of 5ESS by MRs) of changes is critical for understanding both the rhetoric and the implications of changes. This may be achieved manually by disciplined and systematic commenting but is easier, systematic and consistent if wrapped in an interface around the notion of a logical change set on top of your version management system. This is especially important since there are significant numbers of changes made to code that was previously changed. 2) As more change management functions are added to versions management systems, thought should be given to adding several kinds of change classifications. For example, the simple classification as to the purpose of the change is useful and likely to be more reliable than the automated classification techniques used here and elsewhere. This will make it easier to argue for stronger internal validity in future studies as well.

The second set of recommendations is for evolution processes. 1) While it is tempting to take the risk of ignoring small changes as not having significant effects, there is also the fact that small changes (10 lines and under) in this study accounted for approximately 50% of the changes to the system. They do, however, have significantly less risk associated with them than larger changes. We recommend a more agile process to assess the small changes and their impact. Indeed it may be useful to have a changes size sensitive assessment process: one for under 10 lines (46% of the files changed and likelihood of generating an error of about 8%⁶), one for 20 to 50 lines (20% of the files changed with a likelihood of generating an error of about 19%) and one for 50 and over (4% of the files changed with a likely hood of generating an error of about 35%). We base this on the rhetoric shown in

⁶ We arrived at these likelihood estimates by taking the average faults rates for each set of change sizes from the data in Figure 5.

the interplay of inserted and modified lines relative to the size of the change. 2) As a matter of practical quality control, particular attention should be paid to modified lines in corrective changes and inserted lines in adaptive changes since they are the most frequently occurring in those two reasons for change, although it is probably easier just to inspect both types of changes in both cases.

9. Future Work

Very few studies have been done to understand the software development process by the analysis of changed lines. Most studies have focused on the analysis of changed files or system components at various granularities. While the software project we analyzed had modules varying in sizes from 50 lines of code to 50,000 lines of code, we did not consider the individual module sizes separately. Is there a relationship between the size of the module and the probability of error due to change? Our intuition is that changes (irrespective of change size) made to larger files will introduce more errors since the developer is less likely to have a full understanding of the larger modules.

In this analysis, we have only considered those defects that were introduced in the lines affected by the change. However, making a change to a part of the code could affect another part of the same module, either very close to the changed lines or in other parts of the program. In the future we intend to extend this research to study localization effects of making changes.

Finally, to understand fully the small set of changes that result in faults, some of them catastrophic, we need to investigate the context of those changes. Are there common characteristics in the code that is changed? For example, is it in abnormal rather than normal code – studies in interface faults by one of the authors showed that a significant number of faults occurred in error handling code [19][20]. Are there common characteristics in the changes themselves? Are there domain specific aspects to this set of changes or are they uniform across domains?

There is still much work to be done to understand the rhetoric of small changes, indeed of changes of any size.

Acknowledgements

We wish to thank Harvey Siy, Bell Laboratories, Lucent Technologies, for sharing his knowledge of the 5ESS change management process. We would also like to thank Audrus Mockus, Avaya Research Labs, and Tom Ball, Microsoft Research, for their contributions and suggestions.

References

- [1] Fred Brooks, "The Mythical Man-Month", Addison-Wesley, 1975
- [2] Marek Leszak, Dewayne E. Perry and Dieter Stoll, "Classification and evaluation of defects in a project retrospective", *Journal of Systems and Software* (2002) 61(3): 173-187
- [3] Audris Mockus, Lawrence G. Votta, "Identifying Reasons for Software Changes using Historic Databases", In *International Conference on Software Maintenance*, San Jose, California, October 14, 2000, Pages 120-130
- [4] Todd L Graves, Audris Mockus, "Inferring Change Effort from Configuration Management Databases", *Proceedings of the Fifth International Symposium on Software Metrics*, IEEE, 1998, Pages 267-273
- [5] Stephen G. Eick, Todd L. Graves, Alan F. Karr, J.S. Marron, Audris Mockus, "Does Code Decay? Assessing the Evidence from Change Management Data", *IEEE Transactions on Software Engineering*, Vol. 27, No. 1, January 2001
- [6] Dewayne E. Perry, Harvey P. Siy, "Challenges in Evolving a Large Scale Software Product", *Proceedings of the International Workshop on Principles of Software Evolution*, 1998 *International Software Engineering Conference*, Kyoto, Japan, April 1998
- [7] Audris Mockus, David M. Weiss, "Predicting Risk of Software Changes", *Bell Labs Technical Journal*, April-June 2000, Pages 169-180
- [8] Rodney Rogers, "Deterring the High Cost of Software Defects", Technical paper, Upspring Software, Inc.
- [9] G. M. Weinberg, "Kill That Code!", *Infosystems*, August 1983, Pages 48-49
- [10] David M. Weiss, Victor R. Basili, "Evaluating Software Development by Analysis of Changes: Some Data from the Software Engineering Laboratory", *IEEE Transactions on Software Engineering*, Vol. SE-11, No. 2, February 1985, Pages 157-168
- [11] Myron Lipow, "Prediction of Software Failures", *The Journal of Systems and Software*, 1979, Pages 71-75
- [12] Swanson. E. B., "The Dimensions of Maintenance", *Procedures of the Second International Conference on Software Engineering*, San Francisco, California, October 1976, Pages 492-497

- [13] Todd L. Graves, Alan F. Karr, J.S. Marron, Harvey Siy, "Predicting Fault Incidence Using Software Change History", IEEE Transactions on Software Engineering, Vol. 26, No. 7, July 2000, Pg 653-661
- [14] H.E. Dunsmore, J.D. Gannon, "Analysis of the Effects of Programming Factors on Programming Effort", The Journal of Systems and Software, 1980, Pages 141-153
- [15] Ie-Hong Lin, David A. Gustafson, "Classifying Software Maintenance", 1988 IEEE, Pages 241-247
- [16] Dewayne E. Perry, Harvey P. Siy, Lawrence G. Votta, "Parallel Changes in Large Scale Software Development: An Observational Case Study", ACM Transactions on Software Engineering and Methodology 10:3 (July 2001), pp 308-337.
- [17] Les Hatton, Programming Research Ltd, "Reexamining the Fault Density – Component Size Connection", IEEE Software, March/April 1997, Vol. 14, No. 2, Pages 89-97
- [18] Victor R. Basili, Barry T. Perricone, "Software Errors and Complexity: An Empirical Investigation", Communications of the ACM, January 1984, Vol 27, Number 1, Pages 42-52
- [19] Dewayne E. Perry and W. Michael Evangelist. "An Empirical Study of Software Interface Errors", *Proceedings of the International Symposium on New Directions in Computing*, IEEE Computer Society, August 1985, Trondheim, Norway, pages 32-38
- [20] Dewayne E. Perry and W. Michael Evangelist. "An Empirical Study of Software Interface Faults --- An Update", *Proceedings of the Twentieth Annual Hawaii International Conference on Systems Sciences*, January 1987, Volume II, pages 113-126.
- [21] Gerard J Holtzmann, "The Logic of Bugs", Proceedings of the Tenth ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE-10), Chraleston SC, November 2002. pp 81-87.
- [22] K. E. Martersteck and A. E. Spencer, "Introduction to the 5ESS™ Switching System", AT&T Technical Journal, 64:(6 part 2) (July-August (1985), pp 1305-1314
- [23] David C. Carr, Ashok Dandekar, Dewayne E. Perry: Experiments in Process Interface Descriptions, Visualizations and Analyses. EWSPT 1995: 119-137
- [24] Lawrence G. Votta, Mary L. Zajak: Design Process Improvement Case Study Using Process Waiver Data. ESEC 1995: 44-58

- [25] P. A. Tuscan, "Software Development environment for Large Switching Projects", in Proceedings of the 6th International Conference on Software Engineering, Tokyo, September 1992, pp 58-67.
- [26] Marc J. Rochkind, "The Source Code Control System", IEEE Transactions on Software Engineering, SE-1(4) (December 1975), pp 364-370.
- [27] M. M. Lehman, Dewayne E. Perry, Juan F. Ramil: Implications of Evolution Metrics on Software Maintenance. ICSM 1998: 208-
- [28] CVS – Concurrent Versions System. <http://www.cvshome.org>
- [29] David B. Leblang, "The CM Challenge: Configuration Management that Works", in Walter F. Tichy, editor, Configuration Management: Trends in Software, John Wiley & Sons, 1994.
- [30] Ahmed E. Hassan, Mining Software Repositories to Assist Developers and Support Managers, PhD Thesis, University of Waterloo, Department of Computer Science, February 2005.
- [31] Ahmed E. Hassan and Richard C. Holt, "Studying the Chaos of Code Development", in Proceedings of WCRE 2003: Working Conference on Reverse Engineering, Victoria BC, November 2003.
- [32] Stephen G. Eick, Joseph L. Steffen, Eric E. Sumner Jr, "Seesoft – A Tool for Visualizing Line Oriented Software Statistics", IEEE Transactions on Software Engineering, 18:11 (November 1992), pp 957-968.
- [33] Julian Meltzoff. Critical Thinking about Research: Psychology and Related Fields. Washington DC: Psychological Association, 1997 (see pages 45-47).
- [34] Dewayne E. Perry, "An Empirical Approach to Design Metrics and Judgements", Proceedings of the New Vision for Software Design and Production Workshop, Vanderbilt University, December 2001.