

Software Dependencies, Work Dependencies, and Their Impact on Failures

Marcelo CATALDO, Audris MOCKUS, Jeffrey A. ROBERTS, and James D. HERBSLEB

Abstract—Prior research has shown that customer reported software faults are often the result of violated dependencies that are not recognized by developers implementing software. Many types of dependencies and corresponding measures have been proposed to help address this problem. The objective of this research is to compare the relative performance of several of these dependency measures as they relate to customer reported defects. Our analysis is based on data collected from two projects from two independent companies. Combined, our data set encompasses eight years of development activity involving 154 developers. The principal contribution of this study is the examination of the relative impact that syntactic, logical and work dependencies have on the failure proneness of a software system. While all dependencies increase the fault proneness, the logical dependencies explained most of the variance in fault proneness, while workflow dependencies had more impact than syntactic dependencies. These results suggest that practices such as re-architecting, guided by the network structure of logical dependencies, holds promise for reducing defects.

Index Terms — Distribution / maintenance / enhancement, metrics / measurement, organizational management and coordination, quality analysis and evaluation.

15

I. INTRODUCTION

It has long been established that many software faults are caused by violated dependencies that are not recognized by developers designing and implementing a software system [12, 26]. The failure to recognize these dependencies could stem from technical properties of the dependencies themselves as well as from the way development work is organized. In other words, two dimensions are at play – technical and organizational.

On the technical side, the software engineering literature has long recognized call and data-flow syntactic relationships as an important source of error [4, 29, 40]. Research in the software evolution literature has introduced a new view on technical dependencies among software modules. Gall and colleagues [21] introduced the idea of “logical” coupling (or dependencies) by showing that source code files that are changed together can uncover dependencies among those files that are not explicitly identified by traditional syntactic approaches. Past work has also examined aspects of the relationship between logical dependencies and failures in software systems. Eick and colleagues [15] used increases of such logical coupling as an indicator of “code

1decay”. Graves and colleagues [23] showed that past changes are good predictors of future
2faults, and Mockus and Weiss [32] found that the spread of a change over subsystems and files is
3a strong indicator that the change will contain a defect.

4 Human and organizational factors can also strongly affect how dependencies are handled, po-
5tentially affecting the quality of a software system. Research has shown that the level of interde-
6pendency between tasks tends to increase the level of communication and coordination activities
7among workers [20, 46]. Recent studies suggest however, that the identification and manage-
8ment of technical dependencies is a challenge in software development organizations, particular-
9ly when those dependencies are semantic rather than syntactic [7, 12, 24, 27]. Appropriate levels
10of communication and coordination may not occur, potentially decreasing the quality of a system
11[11, 26]. Consequently, it is important to understand how work dependencies (i.e., the way de-
12pendencies are manifested in development tasks) impact failure proneness.

13 In contrast with research on fault prediction models [35, 36, 48], our work focuses on evaluat-
14ing several potential causes of defects, rather than formulating a predictive model. The principal
15contribution of this study is the examination of the *relative* impact that syntactic, logical and
16work dependencies have on the failure proneness of software systems. While all these factors
17are shown to be related to failures, the strength of the relationships varies dramatically. Under-
18standing the relative impact is critical for determining where to focus research, tools, and process
19improvement. In addition, we also sought to improve the external validity of the study by repli-
20cating the analysis over multiple releases of two distinct projects from two unrelated companies.

21 The remainder of the paper is organized as follows. The next two sections elaborate on how
22syntactic, logical, and work-related dependencies relate to a software system’s failure proneness.
23Sections 4, 5 and 6 describe the study methodology, preliminary analyses and the results, respec-

1tively. We conclude the paper with a discussion of the contributions, limitations, and future
2work.

3

II. SOFTWARE DEPENDENCIES AND FAILURE PRONENESS

4 The traditional syntactic view of software dependency had its origins in compiler optimiza-
5tions, and focused on control and dataflow relationships [28]. This approach extracts relational
6information between specific units of analysis such as statements, functions or methods, and
7source code files. Dependencies are discovered, typically, by analysis of source code or from an
8intermediate representation such as bytecodes or abstract syntax trees. These relationships can be
9represented either by a data-related dependency (e.g. a particular data structure modified by a
10function and used in another function) or by a functional dependency (e.g. method A calls
11method B).

12 The work by Hutchens and Basili [29] and Selby and Basili [40] represents the first use of de-
13pendency data in the context of a system's propensity for failure. Building on the concepts of
14coupling and cohesion proposed by Stevens, Myers and Constantine [43], Hutchens and Basili
15[29] presented metrics to assess the structure of a system in terms of data and functional relation-
16ships, which were called bindings. The authors used clustering methods to evaluate the modular-
17ization of a particular system. Selby and Basili [40] used the data binding measure to relate sys-
18tem structure to errors and failures. They found that routines and subsystems with lower coupling
19were less likely to exhibit defects than those with higher levels of coupling. Similar results have
20been reported in object-oriented systems. Chidamber and Kemerer [9] proposed a set of mea-
21sures that captures different aspects of the system of relationships between classes. Briand and
22colleagues [4] found that the measures of coupling proposed by Chidamber and Kemerer were
23positively associated with failure proneness of classes of objects.

1 More recently, models focused on the prediction of failure proneness have been explored using
2 various concepts to organize (or group) software artifacts into various units of analysis. These or-
3 ganizing concepts include architectural, graph-theoretic, and “concerns” perspectives. Measures
4 such as network, syntactic dependency, and complexity metrics are used to explore the associa-
5 tion between the artifact groups and post-release defects. Eaddy and colleagues [14] explored de-
6 fects using concerns (i.e., features or requirements) to organize software artifacts for analysis.
7 Here, the authors found that dispersion of a concern’s implementation (“scatter”) was associated
8 with software defects. Nagappan and Ball [35] explored software failures using two architectural
9 levels within Microsoft Windows to establish their unit of analysis. The authors found that syn-
10 tactic dependencies and source-code change metrics (“churn”) calculated within and between
11 components (binaries or DLLs) and higher level application areas (e.g. the Internet Explorer
12 area) were predictive of post-release failures. Zimmerman and Nagappan [48] applied a graph
13 theoretic lens to classify and calculate network measures for Windows binaries. In this work, the
14 authors demonstrated that orthogonal linear combinations of network, syntactic dependency, and
15 complexity metrics could be used to predict post-release defects.

16 In contrast to the previously discussed research, an alternative view of dependency has been
17 developed in the software evolution literature. This approach focuses on deducing dependencies
18 between the source code files of a system that are changed together as part of the software devel-
19 opment effort and it was first discussed in the literature as “logical coupling” by Gall and col-
20 leagues [21]. Unlike traditional syntactic dependencies, this approach identifies indirect or se-
21 mantic relationships between files that are not explicitly deducible from the programming lan-
22 guage constructs [21]. There are several cases where logical dependencies provide more valuable
23 information than syntactic dependencies. Remote procedure calls (RPCs) represent a simple ex-

1ample. Although the syntactic dependency approach would provide the necessary information to
2relate a pair of modules, such information would be embedded in a long path of connections
3from the RPC caller through the RPC stubs all the way to the RPC server module. On the other
4hand, when the module invoking the RPC and the module implementing the RPC server are
5changed together a logical dependency is created showing, a direct dependency between the af-
6fected source code files. The logical dependency approach is even more valuable in cases such as
7publisher-subscriber or event-based systems where the call-graph approach would fail to relate
8the interdependent modules since no syntactically visible dependency would exist between, for
9instance, a module that generates an event and a module that registers to receive such event.

10 Not only does the logical dependency approach have the potential to identify important depen-
11dencies not visible in syntactic code analyses, it may also filter out syntactic dependencies that
12are unlikely to lead to failures. For example, in the case of basic libraries (e.g. memory manage-
13ment, printing functionality, etc.) the syntactic dependencies approach would highlight these
14highly coupled files. Yet, they tend to be very stable and unlikely to fail despite a high level of
15coupling. The logical dependency approach eliminates these problems as the likelihood of
16change in files that implement these basic functions is very low, hence, a logical dependency
17would not be established.

18 It is difficult to know if the logical dependency approach actually realizes these potential ad-
19vantages. Only limited work has focused on the relationship between logical dependencies and
20failure proneness of a system. Mockus and Weiss [32] found that in a large switching software
21system, the number of subsystems modified by a change is an excellent predictor of whether the
22change contains a fault. Nagappan and Ball [35] found that architecturally based logical coupling
23metrics are correlated with post-release failure proneness of programs. However, the authors

1 computed metrics at the level of component and program areas, a coarse-grain approach resulting
2 in measures too highly correlated to allow the authors to assess each metric's relative impact on
3 failure proneness.

4 In sum, the extant research exploring the relationship between failure proneness of software
5 with regard to dependencies has focused on a single dependency type (syntactic or logical) and
6 has not examined the relative contribution of each of these types. One implication of this limita-
7 tion is that decisions regarding the focus of quality improvement efforts may be misplaced. Ad-
8 ditionally, research in this area has examined only a single project limiting the external validity
9 of results. This leads to our first research question:

10 **RQ 1: What is the relative impact of syntactic and logical dependencies on the failure**
11 **proneness of a software system?**

12

III. WORK DEPENDENCIES AND FAILURE PRONENESS

13 The literature on failure proneness has only recently begun to look at the impact of human and
14 organizational factors on the quality of such systems. The work on coordination in software de-
15 velopment suggests that identification and management of work dependencies is a major chal-
16 lenge in software development organizations [12, 24, 27]. Modularization is the traditional ap-
17 proach used to cope with dependencies in product development. In software engineering, Parnas
18 [37] was the first to articulate the idea of modular software design introducing the concept of in-
19 formation hiding. Parnas argued that modules be considered work items, not just a collection of
20 subprograms. The idea being that development on one module can proceed independently of the
21 development of another. Baldwin and Clark [2], in the product development literature, argued
22 that modularization makes complexity manageable, enables parallel work and tolerates uncer-
23 tainty. Like Parnas, Baldwin and Clark argued that a modular design structure leads to an equiva-

1lent modular work structure.

2 The modularization argument assumes a simple and obvious relationship between product
3 modularization and task modularization – reducing the technical interdependencies among mod-
4 ules also reduces the interdependencies among the tasks involved in producing those modules.
5 In addition, the modular design approach assumes that reducing dependencies reduces the need
6 for work groups to communicate. Unfortunately, there are several problems with these assump-
7 tions. Recent empirical evidence indicates that the relationship between product structure and
8 task structure is not as simple as previously assumed [6]. Moreover, promoting minimal commu-
9 nication between teams responsible for related modules is problematic because it significantly in-
10 creases the likelihood of integration problems [13, 24]. Herbsleb and colleagues [26] theorized
11 that the irreducible inter-dependence among software development tasks can be thought of as a
12 distributed constrain satisfaction problem (DSCP) where coordination is a solution to the DSCP.
13 Within that framework, the authors argued that the patterns of task interdependence among the
14 developers as well as the density of the dependencies in the constraint landscape are important
15 factors affecting coordination success and, by extension, the quality of a software system and the
16 productivity of the software development organization.

17 More recently, Nagappan and colleagues [36], Pinzger and colleagues [38], and Meneely and
18 colleagues [32] investigated a series of organizational metrics as predictors of failure proneness
19 in Windows components and other software. All of the above studies share important limitations
20 with respect to understanding the impact of organizational and social factors in failure proneness.
21 First, they focus on failure prediction models and contain no analysis of the relative importance
22 of the measures in predicting software defects. Furthermore, the proposed measures do not
23 specifically capture work dependencies per se but rather they are proxies for numerous phenome-

1na not necessarily related to the issue of work dependencies. For instance, the measure “number
2of unique engineers who have touched a binary” in [36, pg. 524] could be capturing different
3sources of failures such as difficulties stemming from disparities in engineers' experience and or-
4ganizational processes rather than capturing issues of coordination [36]. In sum, there is a need
5to better understand how the quality of a software system is affected by the ability of the devel-
6opers to identify and manage work dependencies. This leads to our second research question:

**7RQ 2: Do higher levels of work dependencies lead to higher levels of failure proneness of a
8software system?**

9

IV. METHODS

10 We examined our research questions using two large software development projects. One
11project was a complex distributed system produced by a company operating in the computer stor-
12age industry. The data covered a period of approximately three years of development activity and
13the first four releases of the product. The company had one hundred and fourteen developers
14grouped into eight development teams distributed across three development locations. All the de-
15velopers worked full time on the project during the time period covered by our data. The system
16was composed of approximately 5 million lines of code distributed in 7737 source code files in C
17language with a small portion of 117 files, in C++ language.

18 The second project was an embedded software system for a communications device developed
19by a major telecommunications company. Forty developers participated in the project over a pe-
20riod of five years covering six releases of the product. All but one developer worked in the same
21location. The system had more than 1.2 million lines of C and C++ code in 1224 files with 427
22files written using in C++. We will refer to the distributed system as “project A” and to the em-
23bedded system as “project B”.

1 In both development organizations, every change to the source code was controlled by modifi-
2 cation requests. A modification request (MR) is a development task that represents a conceptual
3 change to the software that involves modifications to one or more source code files by one or
4 more developers [33]. The changes could represent the development of new functionality or the
5 resolution of a defect encountered by a developer, the quality assurance organization, or reported
6 by a customer. We refer to latter type of defects as “field” defects. A similar process was associ-
7 ated with each modification request in both projects. Upon creation, the MR is in *new* state, it is
8 then assigned to a particular development team by a group of managers performing the role of a
9 change control board. Commits to the version control systems were not allowed without modifi-
10 cation request identifier. This characteristic of the process allowed the organizations to have a re-
11 liable mechanism of associating the modification request reports with the actual changes to the
12 software code. As soon as all the changes associated with a modification request are completed,
13 the MR is set to *review required* state and a reviewer is assigned. Once the review is passed and
14 the changes are integrated and tested, the modification request is set to *closed* state. In project A,
15 we collected data corresponding to a total of 8257 resolved MRs belonging to the first four re-
16 leases of the product. We collected the data associated with more than 3372 MRs in project B. In
17 the remainder of this section, we describe the measures and the statistical models used in this re-
18 search.

19 *A. Descriptions of the Data and Measures*

20 We used three main sources of data in both projects A and B. First, the MR-tracking system
21 data was used to collect the modification requests included in our analysis. Secondly, the version
22 control systems provided the data that captured the changes made to the system’s source code.
23 Finally, the source code itself. Using the above data sources, we constructed our dependent and
24 independent measures that are described in the following paragraphs.

1) *Measuring Failure*

We chose to investigate failure proneness at the file level. Our dependent variable, *File Buggy-ness*, is a binary measure indicating whether a file has been modified in the course of resolving a field defect. For each file, we determined if it was associated with a field defect in any release of the product covered by our data. We used the logistic regression model shown in Equation 1 in order to model the binary dependent variable and assess the effect of syntactic, logical and work dependencies.

$$\begin{aligned}
 \text{FileBuggy}ness = & \sum_i \beta_i * \text{SyntacticDependenciesMeasure}_i + \\
 & \sum_j \chi_j * \text{LogicalDependenciesMeasure}_j + \\
 & \sum_n \delta_n * \text{WorkDependenciesMeasure}_n + \\
 & \sum_k \varphi_k * \text{Additional Measure}_k + \varepsilon
 \end{aligned} \tag{1}$$

2) *Syntactic Dependencies*

We obtained syntactic dependency information using a modified version of the C-REX tool [25] to identify programming language tokens and references in each entity of each source code file.¹ For all revisions of both systems, a separate syntactic dependency analysis was performed for a snapshot of all source code associated with that revision. Each source code snapshot was created at the end of the quarter in which the release took place. Using the resulting data, we computed syntactic dependencies between source code files by identifying data, function and method references crossing the boundary of each source code file. Let D_{ij} represent the number of data/function/method references that exist from file i to file j . We refer to data references as *data dependencies* and function/method references as *functional dependencies*.

Arguably, data and functional syntactic dependencies could impact failure proneness different-

¹ We were not able to utilize common object oriented coupling measures as both systems are predominantly written using the C programming language.

1ly. Functional dependencies provide explicit information about the relationship between a caller
2and a callee. On the other hand, data relationships are not quite obvious, particularly, in terms of
3understanding the modification sequences of data objects such as a global variable. Such under-
4standing, typically, requires the usage of a tool such as a debugger. Consequently, we collected
5four syntactic dependencies measures: inflow and outflow data relationships and inflow and out-
6flow functional dependencies. Each of those four measures capture the number of syntactic de-
7pendencies of such type exhibited by each file i .

8 3) Logical Dependencies

9 Logical dependencies relate source code files that are modified together as part of an MR. If an
10MR can be implemented by changing only one file, it provides no evidence of any dependencies
11among files. However, when an MR requires changes to more than one file, we assume that de-
12cisions about the change to one file depend in some way on the decisions made about changes to
13the other files involved in the MR. The concept of logical dependencies is equivalent to Gall and
14colleagues [21] idea of logical coupling.

15 In both projects, modification requests contained information about the commits made in the
16version control system. As described earlier, such information was reliably generated as part of
17the submission procedures established in the development organizations. Such data allowed us to
18identify the relationship between development tasks and the changes in the source code associat-
19ed with such tasks. Using this information, we constructed a logical dependency matrix. The
20logical dependency matrix is a symmetric matrix of source code files where C_{ij} represents the
21sum, across all releases, of the number of times files i and j were changed together as part of an
22MR. We accumulate the data across releases as files that are changed together in an MR provide
23mounting evidence of the existence of a logical dependency. The longer the period of time con-

1sidered, the more changes take place, increasing accuracy of the identified logical dependencies.

2 Although the association between MRs and changes in the code was enforced by processes and
3tools, there are other sources of potential errors that might impact the quality of the data repre-
4sented in the logical dependency matrix. For instance, a developer could commit a single change
5to two files where one contained a fix to one MR and the second file had an unrelated change to
6a second MR. We performed a number of analyses to assess the quality of our MR-related data
7and minimize measurement error. We compared the revisions of the changes associated with the
8modification requests and we did not find evidence of such type of behavior. We also grouped
9version control commits that might have been associated with modification requests that were
10marked as duplicates under a single MR. Finally, we examined random samples of modification
11requests to determine if developers have work patterns that could impact the quality of our data
12such as the example described above. For instance, during the data collection process of project
13A, one of the authors and a senior developer from the project examined a random sample of 90
14modification requests. None of the commits contained changes to the code that were not associ-
15ated with the task represented in the modification requests.

16 Two file-level measures were extracted from the logical dependency matrix – *Number of Logi-*
17*cal Dependencies* and *Clustering of Logical Dependencies*. The *Number of Logical Depend-*
18*encies* measure for file i was computed as the number of non-zero cells on column i of the matrix.²
19Since the logical dependencies matrix is symmetric, this measure is equivalent to the degree of a
20node in undirected graph, excluding self-loops. The difference in the nature of the technical de-
21pendencies captured by the syntactic and logical approaches is evidenced by the limited overlap
22between those two types of dependencies. In project A, 74.3% of the syntactic dependencies
23were not identified as logical relationships between a pair of source of files while in project B

2 ² The diagonal of the matrix indicates the number of times a single file was modified and can be disregarded from further analysis.

1 such difference was 97.3%.

2 Herbsleb and colleagues [26] argued that the density of dependencies increases the likelihood
 3 of coordination breakdowns. Building on that argument, we constructed a second measure from
 4 the logical dependency matrix that we called *Clustering of Logical Dependencies*. Unlike the
 5 *Number of Logical Dependencies*, this measure captures the degree to which the files that have
 6 logical dependencies to the focal file have logical interdependencies among themselves. Formal-
 7 ly and in graph theoretic terms, the *Clustering of Logical Dependencies* measure for file i is com-
 8 puted as the density of connections among the direct neighbors of file i . This measure is equiva-
 9 lent to Watt’s [47] local clustering measure and it is mathematically represented by equation 2
 10 where k_i is the number of files or “neighbors” that a particular file i is connected to through logi-
 11 cal dependencies and e_{jk} is a link between files j and k which are neighbors of file i . The values of
 12 this measure range from 0 to 1.

$$13 \quad CLD(f_i) = \frac{E|\{e_{jk}\}|}{k_i(k_i - 1)} \quad (1)$$

14

15 4) *Work Dependencies*

16 We constructed two different measures of work dependencies – *Workflow Dependencies* and
 17 *Coordination Requirements*. *Workflow Dependencies* capture the temporal aspects of the devel-
 18 opment effort while *Coordination Requirements* capture the intra-developer coordination re-
 19 quirements.

20 Workflow Dependencies: As described previously, both projects used MR-tracking systems to
 21 assess the progress of development tasks. Each modification request followed a set of states from
 22 creation until closure. Those transitions represent a MR workflow where particular members of
 23 the development organization had work-related responsibilities associated with such MR at some

1 point in time during its lifecycle. Such workflow constitutes the traditional view of work depen-
2 dencies where individuals are sequentially interdependent on a temporal basis [45]. More specifi-
3 cally, two developers i and j are said to be interdependent if the MR was transferred from devel-
4 oper i to developer j at some point between the creation and closure of the MR. For instance,
5 suppose a MR requires changes to two subsystems with the changes to the second relying on
6 changes to the first. Developer i completes the work on subsystem one and then he/she transfers
7 the development task to developer j to finish the work on the subsystem two.

8 Grouping the workflow information of all the MRs associated with a particular release of the
9 products, we constructed a developer-to-developer matrix where a cell c_{ij} represents the number
10 of work dependencies developer i has on developer j . The information in such a matrix captures
11 the web of work-related dependencies that each developer was embedded during a particular re-
12 lease of the product. Such developer-to-developer relationships can be examined through the
13 lenses of social network analysis which provides the relevant theoretical background and
14 methodological framework [30, 46]. A traditional result in the social network literature is that in-
15 dividuals centrally located in the network (i.e., have, on average, a larger number of relationships
16 to other individuals) tend to be more influential because they control the flow of information [5,
17 30]. On the negative side, a high number of linkages requires a significant effort on the part of
18 those individuals in order to maintain the relationships [5, 30]. This latter point is particularly
19 important in the context of the workflow dependencies because it argues that centrally located
20 developers are more likely to be overloaded because of the effort associated with managing the
21 work dependencies, increasing the likelihood for communication break downs and thus the quali-
22 ty of software produced could be expected to diminish.

23 Degree centrality [19] is a traditional measure used in the social network literature to identify

1 central individuals based on the number of ties to other actors in the network. Formally, degree
 2 centrality is defined as $DC(n_i, M) = d(n_i)$, where $d(n_i)$ is the number of connections of node n_i in
 3 matrix M . The values of this measure range from 0 to $n-1$ where 0 indicates the node is an isolate
 4 (i.e., not connected to any other node) and $n-1$ indicates that the node i has a ties to all other $n-1$
 5 nodes. Building on the theoretical argument outlined in the previous paragraph and on the con-
 6 cept of degree centrality, the *Workflow Dependencies* measure was constructed as follows. For
 7 each file i , we identified the developer j that worked on the file and was linked to the greatest
 8 number of individuals in the developer-to-developer workflow network for each release. That is,
 9 the developer exhibiting the highest degree centrality. Then, as discussed earlier, such individu-
 10 als are the most likely to introduce an error due to higher levels effort those individuals face in
 11 managing a higher number of work dependencies. Equation 3 formally describes the *Workflow*
 12 *Dependencies* measure. We also considered the average of the number of linked developers over
 13 the set of developers that worked on each file. However, this measure was highly correlated with
 14 our other independent measures and thus excluded from further analysis.

$$15 \quad WD(f_i) = \max \{DC(dev_j, WD) \mid j \in \{developers \text{ that changed } f_i\}\} \quad (3)$$

16

17 Coordination Requirements: Workflow dependencies relate developers through the temporal
 18 evolution of modification requests and the developers' involvement in those MR. There are addi-
 19 tional work-related dependencies that emerge as development work is done in different parts of a
 20 system. For instance, two developers could work on two different modification requests involv-
 21 ing files that are syntactically or logically interdependent, then, the modifications made by both
 22 developers could impact each other's work. There types of work-related dependencies are more
 23 subtle in nature and require more effort on the part of the developers to identify and manage

1them. Cataldo and colleagues [6] proposed a framework for examining the relationship between
 2the technical dependencies of a software system and the structure of the development work to
 3construct such system. Coordination requirements, an outcome of that framework, represent a
 4developer-by-developer matrix (C_R) where each cell $C_{R\ ij}$ represents the extent to which develop-
 5er i needs to coordinate with developer j given the assignments of development tasks and techni-
 6cal dependencies of the software system. More formally, Cataldo and colleagues [6] defined the
 7 C_R matrix with the following product:

$$8 \quad C_R = T_A * T_D * T_A^T \quad (4)$$

9where, T_A is the *Task Assignments* matrix, T_D is the *Task Dependencies* matrix and T_A^T is the
 10transpose of the *Task Assignments* matrix. In the context of our study, the T_A and T_D matrices
 11were constructed using data from the MR reports and the version control system in the following
 12way. A MR report provides the “developer i modified file j ” relationship. We grouped such in-
 13formation across all modification requests in a particular release to construct the *Task Assign-*
 14*ment* matrix which is a developer-to-file matrix. The *Task Dependency* matrix was a file-to-file
 15matrix and it was constructed using the same approach described in the computation of the logi-
 16cal dependencies measures. In other words, each cell c_{ij} of the *Task Dependency* matrix repre-
 17sents the number of times a particular pair of source code files changed together as part of the
 18work associated with the MRs. Finally, using equation 4, we computed the C_R matrix. Following
 19the theoretical argument and the process presented in the previous section (description of work-
 20flow dependencies), the *Coordination Requirements* measure captures for each file i , the degree
 21centrality of the most central developer in the C_R matrix (a developer-to-developer matrix) that
 22worked on the file i . Equation 5 formally describes the *Coordination Requirements* measure.

$$23 \quad CR(f_i) = \max \{DC(dev_j, C_R) \mid j \in \{developers \text{ that changed } f_i\}\} \quad (5)$$

1 5) *Additional Control Factors*

2 The objective of this study is to examine the relative impact that important conceptual factors
3 such as technical and work dependencies have on failure. In order to account for the effects of
4 potentially confounding influences however, our analysis must include factors that past research
5 has found to be associated with failures. Numerous measures have been used to predict failures
6 [14, 18, 23, 35, 36, 48]. As suggested by Graves and colleagues [23], such measures can be clas-
7 sified as either process or product measures. Process measures such as number of changes, num-
8 ber of deltas, and age of the code (i.e., churn metrics) have been shown to be very good predic-
9 tors of failures [23, 35]. Accordingly, we control for the *Number of MRs*, which is the number of
10 times the file was changed as part of a past defect or feature development. We also control for
11 the *Average Number of Lines Changed* in a file as part of MRs.

12 In contrast, product measures such as code size and complexity measures have produced some-
13 what contradictory results as predictors of software failures. Some researchers have found a posi-
14 tive relationship between lines of code and failures [4, 23], while others have found a negative
15 relationship [3]. Our collective experience regarding the relationship between product measures
16 and software defects has been that such measures are associated with increased software failure.
17 Thus, we expect that product measures will be positively associated with software defects. We
18 measure size of the file (*LOC*) as the number of non-blank non-comment lines of code.

19

V. PRELIMINARY ANALYSIS

20 Our four dependency measures (syntactic, logical, workflow and coordination requirements)
21 capture different characteristics of the technical and work-related dependencies that emerge in
22 the development of software systems. Table I presents a comparative summary of our dependen-
23 cy measures. Syntactic and Workflow dependencies are explicit in nature, therefore, easier to

1 identify and manage by developers or other relevant stakeholders in software development
 2 projects. On the other hand, the Logical and Coordination Requirement dependency measures
 3 capture more less explicit and subtle relationships among software artifacts and developers, re-
 4 spectively. The implicit nature of those dependencies makes identification and management of
 5 such relationship more challenging. In sum, our measures assess explicit and implicit dependen-
 6 cies that emerge in the technical and work-related dimensions of software projects.

TABLE I
 COMPARATIVE SUMMARY OF DEPENDENCY MEASURES

	Dimension	Identifiability	Manageability
<i>Syntactic Dependencies</i>	Technical	Captures explicit relationships between source code files.	A host of tools can aid developers in the management of this type of dependencies.
<i>Logical Dependencies</i>	Technical	Captures semantic or implicit relationships between source code files, in addition to some explicit relationships.	Dependence on historical data, attributes of the tools (e.g. version control system) and consistent processes over time limits the developers' ability to manage these type of dependency.
<i>Workflow Dependencies</i>	Work / Social	Captures explicit relationships among project members based on workflows and/or processes	Traditional tools (e.g. ClearQuest or Bugzilla) facilitate significantly the management of these dependencies.
<i>Coordination Requirement Dep.</i>	Work / Social	Captures less explicit relationships among project members based on their past contributions to the development effort and the technical dependencies of the system under development.	Dependence on historical data, attributes of the tools (e.g. version control system) and consistent processes over time limits the developers' ability to manage these type of dependency.

7

8 Table II summarizes the descriptive statistics of all the measures described in the previous sec-
 9 tions. Due to a moderate degree of skewness, we applied a log-transformation to each of the in-
 10 dependent variables. Table III reports the pair-wise correlations of all our measures. Overall, the
 11 pair-wise correlations are relatively similar across projects **indicating** that the phenomena reflect-
 12 ed by these measures may be common in both projects. There are, however, several high correla-
 13 tions that deserve attention. For instance, the *Number of MRs* (past changes) variable is highly
 14 correlated with *LOC*, *Average Lines Changed* and our measure of logical dependencies, particu-
 15 larly in project B. In addition, the syntactic dependencies measures are also highly correlated
 16 among themselves and with other measures such as *LOC* and *Number of MRs*. We computed

1variance inflation factors and tolerances to further examine potential issues due to multicollinear-
2ity among our independent variables. A tolerance close to 1 indicates little multicollinearity,
3whereas a value close to 0 suggests that multicollinearity may be a significant threat. Variance
4inflation factor (VIF) is defined as the reciprocal of the tolerance.

TABLE II
DESCRIPTIVE STATISTICS

Project A: Distributed System						
	Mean	SD	Min	Max	Skew	Kurtosis
<i>File Buggyness</i>	0.49	0.500	0	1	0.011	1.001
<i>LOC</i>	481.9	836.1	0	17853	4.931	47.24
<i>Avg. Lines Changed</i>	10.85	32.67	0	738	8.512	108.9
<i>In-Data Syntactic Dep.</i>	4.57	58.94	0	1741	24.40	647.6
<i>Out-Data Syntactic Dep.</i>	8.90	9.243	0	53	0.792	3.050
<i>In-Functional Syntactic Dep.</i>	20.36	71.49	0	951	5.701	42.78
<i>Out-Functional Syntactic Dep.</i>	25.96	68.42	0	543	5.241	32.57
<i>Num. Logical Dep.</i>	87.27	99.54	0	836	1.856	7.584
<i>Clustering Logical Dep.</i>	0.72	0.316	0	1	-1.024	3.011
<i>Workflow Dep.</i>	22.53	12.76	0	44	-0.013	1.878
<i>Coordination Req.</i>	0.14	0.121	0	0.62	2.655	11.91
Project B: Embedded System						
	Mean	SD	Min	Max	Skew	Kurtosis
<i>File Buggyness</i>	0.14	0.35	0	1	2.026	5.105
<i>LOC</i>	750.8	2874.3	0	65542	18.24	389.6
<i>Avg. Lines Changed</i>	19.18	52.53	0	987	9.617	135.7
<i>In-Data Syntactic Dep.</i>	10.61	85.60	0	1805	16.18	287.1
<i>Out-Data Syntactic Dep.</i>	7.85	14.41	0	173	207.9	27.07
<i>In-Functional Syntactic Dep.</i>	9.17	29.09	0	612	11.11	180.4
<i>Out-Functional Syntactic Dep.</i>	15.84	29.08	0	238	3.396	18.01
<i>Num. Logical Dep.</i>	38.61	41.61	0	370	3.152	18.61
<i>Clustering Logical Dep.</i>	0.52	0.19	0	0.69	-1.241	4.010
<i>Workflow Dep.</i>	28.41	15.60	1	72	0.253	2.461
<i>Coordination Req.</i>	0.85	0.14	0	1	-2.956	15.29

TABLE III
PAIR-WISE CORRELATIONS (* $p < 0.01$) FOR LAST RELEASE IN EACH DATASET

Project A: Distributed System						
	1	2	3	4	5	6
1. <i>FileBugyness</i>	-					
2. <i>LOC (log)</i>	0.28*	-				
3. <i>Number MRs (log)</i>	0.37*	0.24*	-			
4. <i>Avg. Lines Changed (log)</i>	0.18*	0.27*	0.30*	-		
5. <i>In-Data Dep. (log)</i>	0.06*	0.01	0.08*	0.03	-	
6. <i>Out-Data Dep. (log)</i>	0.18*	0.47*	0.19*	0.19*	-0.26*	-
7. <i>In-Functional Dep. (log)</i>	0.04*	0.27*	0.09*	0.09*	-0.10*	0.37*
8. <i>Out-Functional Dep. (log)</i>	0.11*	0.43*	0.15*	0.16*	-0.24*	0.78*
9. <i>Num Logical Dep. (log)</i>	0.49*	0.33*	0.45*	0.16*	0.04*	0.23*
10. <i>Clustering Logical Dep. (log)</i>	-0.32*	-0.21*	-0.29*	-0.13*	-0.06*	-0.17*
11. <i>Workflow Dep. (log)</i>	0.33*	0.06*	0.33*	0.12*	0.02	0.07*
12. <i>Coordination Req. Dep. (log)</i>	0.04*	-0.06*	-0.15*	-0.06*	-0.01	-0.03
	7	8	9	10	11	12
8. <i>Out-Functional Dep. (log)</i>	0.44*	-				
9. <i>Num Logical Dep. (log)</i>	0.06*	0.19*	-			
10. <i>Clustering Logical Dep. (log)</i>	-0.10*	-0.14*	-0.05*	-		
11. <i>Workflow Dep. (log)</i>	-0.07*	-0.03	0.31*	-0.12*	-	
12. <i>Coordination Req. Dep. (log)</i>	-0.07*	-0.05*	0.02	0.12*	0.15*	-
Project B: Embedded System						
	1	2	3	4	5	6
1. <i>FileBugyness</i>	-					
2. <i>LOC (log)</i>	0.28*	-				
3. <i>Number MRs (log)</i>	0.55*	0.41*	-			
4. <i>Avg. Lines Changed (log)</i>	0.19*	0.42*	0.35*	-		
5. <i>In-Data Dep. (log)</i>	0.22*	0.33*	0.26*	0.19*	-	
6. <i>Out-Data Dep. (log)</i>	0.26*	0.60*	0.34*	0.35*	0.49*	-
7. <i>In-Functional Dep. (log)</i>	0.19*	0.36*	0.25*	0.19*	0.47*	0.54*
8. <i>Out-Functional Dep. (log)</i>	0.28*	0.59*	0.38*	0.39*	0.43*	0.88*
9. <i>Num Logical Dep. (log)</i>	0.29*	0.26*	0.62*	0.25*	0.13*	0.20*
10. <i>Clustering Logical Dep. (log)</i>	-0.28*	-0.15*	-0.34*	-0.10*	-0.17*	-0.21*
11. <i>Workflow Dep. (log)</i>	0.26*	0.09*	0.38*	0.01	0.19*	0.10*
12. <i>Coordination Req. Dep. (log)</i>	0.17*	-0.03	0.26*	-0.05	0.14*	0.02
	7	8	9	10	11	12
8. <i>Out-Functional Dep. (log)</i>	0.52*	-				
9. <i>Num Logical Dep. (log)</i>	0.12*	0.22*	-			
10. <i>Clustering Logical Dep. (log)</i>	-0.19*	-0.20*	0.17*	-		
11. <i>Workflow Dep. (log)</i>	0.08	0.10*	0.29*	-0.18*	-	
12. <i>Coordination Req. Dep. (log)</i>	0.07	0.04	0.24*	-0.12*	0.75*	-

3 Table IV reports the variance inflation factor and tolerance associated with each of our mea-
4sures. We start our multicollinearity diagnostic with model I that contains all our independent
5measures. We observe that for both projects A and B, the measures *Out-Data Syntactic Depen-*

Dependencies and *Out-Functional Syntactic Dependencies* have a VIF significantly higher (or a tolerance significantly lower) than the other measures. We removed those two variables and the re-computed VIF and tolerances values for the remaining measures are reported in model II in Table IV. We observe that *Number of MRs* has a lower tolerance than the rest of the measures, particularly in project B's data. Consequently, we removed it and the resulting VIFs and tolerances are reported in model III. In this case, the data for project A does not show signs of multicollinearity, with the tolerances of all measures above 0.70.

TABLE IV
COLLINEARITY DIAGNOSTICS

Project A: Distributed System			
	Model I	Model II	Model III
	VIF (Tolerance)	VIF (Tolerance)	VIF (Tolerance)
<i>Number of MRs (log)</i>	1.59 (0.6289)	1.59 (0.6297)	---
<i>LOC (log)</i>	1.53 (0.6530)	1.32 (0.7564)	1.32 (0.7564)
<i>Avg. Lines Changed (log)</i>	1.16 (0.8596)	1.16 (0.8625)	1.11 (0.9035)
<i>In-Data Dep. (log)</i>	1.13 (0.8867)	1.02 (0.9793)	1.02 (0.9825)
<i>Out-Data Dep. (log)</i>	2.85 (0.3503)	---	---
<i>In-Functional Dep. (log)</i>	1.26 (0.7916)	1.11 (0.9007)	1.11 (0.9031)
<i>Out-Functional Dep. (log)</i>	2.79 (0.3587)	---	---
<i>Num Logical Dep. (log)</i>	1.47 (0.6825)	1.45 (0.6880)	1.26 (0.7950)
<i>Clustering Logical Dep. (log)</i>	1.16 (0.8584)	1.16 (0.8628)	1.09 (0.9152)
<i>Workflow Dep. (log)</i>	1.26 (0.7921)	1.24 (0.8040)	1.18 (0.8487)
<i>Coordination Req. Dep. (log)</i>	1.09 (0.9213)	1.08 (0.9218)	1.05 (0.9523)
Project B: Embedded System			
	Model I	Model II	Model III
	VIF (Tolerance)	VIF (Tolerance)	VIF (Tolerance)
<i>Number of MRs (log)</i>	2.82 (0.3547)	2.80 (0.3573)	---
<i>LOC (log)</i>	1.83 (0.5467)	1.49 (0.6689)	1.45 (0.6897)
<i>Avg. Lines Changed (log)</i>	1.34 (0.7469)	1.30 (0.7687)	1.28 (0.7826)
<i>In-Data Dep. (log)</i>	1.47 (0.6787)	1.38 (0.7244)	1.38 (0.7260)
<i>Out-Data Dep. (log)</i>	4.91 (0.2038)	---	---
<i>In-Functional Dep. (log)</i>	1.58 (0.6344)	1.39 (0.7181)	1.39 (0.7184)
<i>Out-Functional Dep. (log)</i>	4.75 (0.2105)	---	---
<i>Num Logical Dep. (log)</i>	2.32 (0.4316)	2.31 (0.4321)	1.33 (0.7528)
<i>Clustering Logical Dep. (log)</i>	1.61 (0.6223)	1.60 (0.6251)	1.19 (0.8435)
<i>Workflow Dep. (log)</i>	2.56 (0.3913)	2.55 (0.3927)	2.50 (0.4003)
<i>Coordination Req. Dep. (log)</i>	2.38 (0.4201)	2.37 (0.4228)	2.36 (0.4230)

8

9 On the other hand, in project B, the low tolerance values for the two measures of work depen-
10dencies suggest some potential multicollinearity problems. Removing the *Coordination Require-*
11ment Dependencies measure from model III results in an improvement of the VIF associated

1with Workflow dependencies down to 1.20 (tolerance = 0.8304). In addition, the tolerances of all
2remaining variables increased with the minimum value being 0.7028 for the *LOC* measure. In
3section VI, we revisit this issue when discussing the results from our regression analyses.

4

VI. RESULTS

5 We approached the analysis in two stages. In the first stage, we focused on examining the rela-
6tive impact of each dependency type on failure proneness of source code files. The data corre-
7sponding to the last release from each project was used in this analysis. In the second stage, we
8verified the consistency of the initial results by conducting a number of confirmatory analyses
9for each project. These analyses included re-estimating our logistic regression models for each
10release as well as estimating a single longitudinal model comprising all releases. The detailed re-
11sults of each stage are discussed in turn.

12 *A. The Impact of Dependencies*

13 We constructed several logistic regression models to examine the relative impact of each class
14of independent variable on the failure proneness of a software system using the data from the last
15release of each project. Following a standard hierarchical modeling approach, we started our
16analysis with a baseline model that contains only the traditional predictors. In subsequent mod-
17els, we added the measures for syntactic, logical and work dependencies described in the previ-
18ous sections. We assessed the goodness-of-fit of the model to evaluate the impact of each class
19of dependency measures on failure. For each statistical model, we report the χ^2 of the model, the
20percentage of deviance explained by the model as well as the statistical significance of the differ-
21ence between a model that adds new factors and the previous model without the new measures.
22Deviance is defined as -2 times the log-likelihood of the model. The percentage of the deviance

1 explained is a ratio of the deviance of the null model (containing only the intercept), and the de-
2viance of the final model. Model parameters were estimated, as is customary in logistic regres-
3sion, using a maximum-likelihood method. In order to simplify the interpretation of the results,
4we report the odds ratios associated with each measure instead of reporting the regression coeffi-
5cients. Odds ratios larger than 1 indicate a positive relationship between the independent and de-
6pendent variables whereas an odds ratio less than 1 indicates a negative relationship. For exam-
7ple, an odds ratio of two for a binary factor doubles the probability of a file having a costumer re-
8ported defect when the remaining factors in the model are at their lowest values. The presented
9odds ratio is the exponent of the logistic regression coefficient.

10 Table V and VI report the odds ratios of the various logistic regression models using the data
11from project A and project B, respectively. In both tables, model I includes the *LOC* and *Avg.*
12*Lines Changed* measures. As discussed in section V, the *Number of MRs* measure (a proxy for
13past changes) was not included in the analyses due to multicollinearity concerns. Model I, in ta-
14bles V and VI, shows that *LOC* is positively associated with failure proneness. These results
15agree with those found by Briand and colleagues [4], in contrast with earlier findings [3, 34].
16*Avg. Lines Changed* is also positively related to failure proneness in both projects, **indicating** that
17the more modifications to a file, the higher the likelihood of encountering a field defect associat-
18ed with that file. Specifically, a unit change in the log-transformed *Avg. Lines Changed* measure
19(or a change from 1 to 2.7 lines per MR in untransformed units), increases the odds of a field de-
20fect by 20% for project A (Table V – Model I) and 25% in the case of project B (Table VI –
21Model I).

22

23

1

2

TABLE V
ODDS RATIOS FROM LOGISTIC REGRESSION ON PROJECT A (DISTRIBUTED SYSTEM) DATA

	Model I	Model II	Model III	Model IV	Model V
<i>LOC (log)</i>	1.392**	1.418**	1.119**	1.142**	1.150**
<i>Avg. Lines Changed (log)</i>	1.203**	1.200**	1.138**	1.114**	1.126**
<i>In-Data Dep. (log)</i>		1.166**	1.103*	1.105*	1.112*
<i>In-Functional Dep. (log)</i>		0.949*	0.953+	0.982	0.989
<i>Num Logical Dep. (log)</i>			2.277**	2.079**	2.108**
<i>Clustering Logical Dep. (log)</i>			0.009**	0.012**	0.009**
<i>Workflow Dep. (log)</i>				2.011**	1.905**
<i>Coordination Req. Dep. (log)</i>					2.801**
Model χ^2 (p-value)	388.87 (p < 0.01)	412.21 (p < 0.01)	1621.31 (p < 0.01)	1737.52 (p < 0.01)	1763.18 (p < 0.01)
Deviance Explained	7.1%	7.5%	29.5%	31.6%	32.1%
Model Comparison χ^2 (p-value)	--	23.34 (p < 0.01)	1209.10 (p < 0.01)	116.21 (p < 0.01)	25.67 (p < 0.01)

(+ p < 0.10; * p < 0.05; ** p < 0.01)

3

4 Model II introduces the syntactic dependency measures *Inflow Data* and *Inflow Functional*.

5 The results of the logistic regression show that the impact of data syntactic dependencies are

6 marginally consistent with previous research, particularly, as the other factors are included in the

7 regression model (see models III, IV and V in tables V and VI). In the case of project A, data

8 syntactic dependencies are statistically significant across the various models and with the expect-

9 ed direction in their impact on failure proneness. On the other hand, the impact of the functional

10 syntactic dependencies measure, unexpectedly, has the opposite direction. However, once the

11 models include logical and work dependencies, the functional syntactic dependency measure no

12 longer has statistical significance **indicating** that this type of syntactic relationship does not im-

13 pact failure proneness. This latter pattern is also reflected in the data for project B where both

14 syntactic dependency measures become irrelevant once the logical and work dependency mea-

15 sures enter the models (see table VI, models III, IV and V). Given the limited impact of the syn-

1tactic dependencies on failure proneness it is not surprising to see a relatively modest improve-
 2ment in the explanatory power of model II over model I (e.g. in project A deviance improves
 3from 7.1% to 7.5%). We do note however, that while improvement in the explanatory power is
 4modest, the addition of the syntactic dependency measures does do provide a statistically signifi-
 5cant improvement in model fit as indicated by the model comparison χ^2 (project A: 23.34 – p <
 60.01; project B: 14.41 – p < 0.01).

TABLE VI
 ODDS RATIOS FROM LOGISTIC REGRESSION ON PROJECT B (EMBEDDED SYSTEM) DATA

	Model I	Model II	Model III	Model IV	Model V
<i>LOC (log)</i>	1.800**	1.638**	1.497**	1.493**	1.499**
<i>Avg. Lines Changed (log)</i>	1.247**	1.253**	1.115	1.178	1.184
<i>In-Data Dep. (log)</i>		1.207*	1.124	1.046	1.142
<i>In-Functional Dep. (log)</i>		1.131	1.013	1.002	0.996
<i>Num Logical Dep. (log)</i>			2.303**	1.822**	1.803**
<i>Clustering Logical Dep. (log)</i>			0.005**	0.013**	0.014**
<i>Workflow Dep. (log)</i>				6.527**	4.899**
<i>Coordination Req. Dep. (log)</i>					37.616
Model χ^2 (p-value)	86.01 (p < 0.01)	100.42 (p < 0.01)	218.13 (p < 0.01)	239.27 (p < 0.01)	240.02 (p < 0.01)
Deviance Explained	11.8%	13.8%	30.1%	32.9%	33.0%
Model Comparison χ^2 (p-value)	--	14.41 (p < 0.01)	117.71 (p < 0.01)	21.14 (p < 0.01)	0.75 (p=0.387)

(+ p < 0.10; * p < 0.05; ** p < 0.01)

7

8 Model III also considers the logical dependencies measures. As Table V and VI show, the odds
 9ratios associated with each of the logical dependencies measures in the logistic regression are
 10greater than one, indicating that higher numbers of logical dependencies are related to an in-
 11crease in the likelihood of failure. In particular, a unit increase in the log-transformed *Number of*
 12*Logical Dependencies* measure, increases the odds of a failure 2.272 times higher for project A
 13(Table V – Model III) and 2.277 times higher for project B (Table VI – Model III). The analyses
 14reported in section V showed relatively low levels of correlation between syntactic and logical
 15dependency measures. Thus, the results reported in Tables V and VI suggest the effect of logical

1dependencies on failure proneness is complementary and significantly more important than the
2impact of syntactic dependencies. In addition, the levels of explained deviance for model III in
3both projects clearly shows that the contribution of the logical dependencies measures to the ex-
4planatory power of the model is much higher than the impact of the syntactic dependencies mea-
5sure.

6 The results reported in Model III in Tables V and VI also indicate that increases in the *Cluster-*
7ing of Logical Dependencies significantly reduce the likelihood of failures. This result may sug-
8gest that the clustering is a symptom of good, consistent modular design. Alternatively, it may
9be that as clusters of consistently interrelated files emerge, developers become more cognizant of
10such relationships and know where to look to make sure that changes to one part of the system
11do not introduce problems elsewhere.

12 In both Tables V and VI, model IV includes the first of our work dependency measures –
13workflow dependencies. The results are consistent across both projects. Higher number of work-
14flow dependencies increases the likelihood that source code files contain field defects. In particu-
15lar, a unit increase in the log-transformed number of *Workflow Dependencies* measure, increases
16the odds of a failure 2.011 times higher for project A (Table V – Model IV) and 6.527 times
17higher for project B (Table VI – Model IV). Model V shows the impact of the second work de-
18pendency measure – coordination requirements. In project A, the impact of the *Coordination Re-*
19*quirement* measure is statistically significant and with an odds ratio of 2.801, its impact is higher
20than the impact of the Workflow Dependencies. On the other hand, in project B, its effect is not
21statistically significant. As discussed in section V, there is high collinearity between the two
22work dependency measures in project B's data (Table III: correlation is 0.75; Table IV: VIFs >
232), consequently, the regression results were expected. In this paper, we set out to examine the

1 relative impact of syntactic, logical and work-related classes of dependencies on failure prone-
2ness. The results presented in this section showed that all types of dependencies affect failures in
3a software system. More importantly, their role is complementary suggesting the various types of
4dependencies capture different relevant aspects of the technical properties of a software system
5as well as elements of the software development process. Logical and work dependencies have a
6significantly higher impact on failure proneness as their associated odds ratios indicate. For in-
7stance, a unit increase in the log-transformed measures of *Number of Logical Dependencies* and
8*Workflow dependencies* increase the odds of post-release defects 2 times more than syntactic de-
9pendencies in the case of project A and 2 times and 6 times, respectively, for the case of project
10B.

11 B. Stability Analysis

12 In the previous section, we showed that the different types of dependencies affected failure
13proneness in the last release of each project. It is also critical to examine whether our results are
14robust across the various releases of the products covered by our data. Accordingly, we ran the
15same logistic regression models on the data from the first three releases from project A and the
16additional five releases from project B. Table VII reports the odd ratios for all the measures from
17the logistic regression using the data from project A. Table VIII reports the odd ratios for the
18measures from the logistic regression using the data from project B. As discussed in the previous
19section, we did not include the *Coordination Requirement Dependencies* measures in the analy-
20sis of project B because of the high correlation of that measure with the *Workflow Dependencies*
21measure. We observe that the results are mostly consistent with those reported in the previous
22section for both project A and B. However, there is one exception. The results for the measure
23of *Workflow Dependencies* are not consistent across releases in the data from project A. One pos-

1sible explanation is the changing nature of the development work associated with each release.
 2For instance, release 1 of project A was in fact the first release of the product. The development
 3effort associated with subsequent releases involved an increasing amount of work related to fix-
 4ing defects reported against previous releases and a decreasing amount of development effort on
 5new features. In the case of project B, the impact of the *Workflow Dependencies* measure is con-
 6sistent across all five releases. However, the coefficient for release 1 is not statistically signifi-
 7cant.

TABLE VII
 IMPACT OF DEPENDENCIES ACROSS RELEASES IN PROJECT A

	Release 1	Release 2	Release 3
<i>LOC (log)</i>	1.211**	1.087**	1.201**
<i>Avg. Lines Changed (log)</i>	1.122**	1.083*	1.048
<i>In-Data Dep. (log)</i>	1.243**	1.207*	1.125*
<i>In-Functional Dep. (log)</i>	0.985	1.041	1.013
Num Logical Dep. (log)	1.411**	1.949**	1.806**
<i>Clustering Logical Dep. (log)</i>	0.064**	0.023**	0.017**
<i>Workflow Dep. (log)</i>	1.287**	0.850**	1.448**
<i>Coordination Req. Dep. (log)</i>	1.007	10.852**	3.901**
Model χ^2 (p-value)	514.53 (p < 0.01)	821.61 (p < 0.01)	1121.96 (p < 0.01)
Deviance Explained	13.7%	191%	22.2%

(+ p < 0.10; * p < 0.05; ** p < 0.01)

8

9The results reported in Tables VII and VIII showed overall consistent effects of our predictors
 10across the different releases covered by our data. However, the development effort associated
 11with each release might have a temporal relationship. For instance, the technical or work depen-
 12dencies from release 2 could influence the measures from release 3. More formally, the various
 13measures associated with each of the releases could exhibit autocorrelation. Therefore, we ran an
 14additional confirmatory analysis using a longitudinal (random effects) model that considers the
 15data from all releases in each project simultaneously. Using this procedure, we accounted for
 16any potential temporal factors that might affect the estimation of the coefficients that represent
 17the impact of our measures on failure proneness. Overall, the results of the random effects model

1 were consistent with those reported in Tables V, VI, VII and VIII.

2

3

TABLE VIII
IMPACT OF DEPENDENCIES ACROSS RELEASES IN PROJECT B

	Release 1	Release 2	Release 3	Release 4	Release 5
<i>LOC (log)</i>	1.642*	1.823*	1.713*	1.447**	1.477**
<i>Avg. Lines Changed (log)</i>	0.984	0.816	0.892	1.116	1.171
<i>In-Data Dep. (log)</i>	1.126	0.905	0.948	0.981	1.057
<i>In-Functional Dep. (log)</i>	0.619	1.153	0.978	1.016	1.001
Num Logical Dep. (log)	3.964**	3.187**	2.166**	1.771**	1.865**
<i>Clustering Logical Dep. (log)</i>	0.001**	0.007**	0.008**	0.012**	0.013**
<i>Workflow Dep. (log)</i>	1.101	1.870*	1.711*	3.936**	3.904**
<i>Coordination Req. Dep. (log)</i>	---	---	---	---	---
Model χ^2 (p-value)	103.44 (p < 0.01)	150.09 (p < 0.01)	159.31 (p < 0.01)	201.63 (p < 0.01)	213.99 (p < 0.01)
Deviance Explained	42.1%	40.5%	29.4%	30.6%	30.8%

(+ p < 0.10; * p < 0.05; ** p < 0.01)

4

5

VII. DISCUSSION

6 The observed relative contributions of different types of dependencies on failure proneness in
7 two unrelated projects have consequences of both theoretical and practical interest. All three
8 types of dependencies are relevant and their impact is complementary showing their independent
9 and important role in the development process. These results suggest that quality improvement
10 efforts could be tailored to ameliorate the negative effects of particular types of dependencies
11 with emphasis on areas that have the largest impact on project quality.

12 Past research [4, 29, 40] has shown that source code files with higher number of syntactic de-
13 pendencies were more prone to failure. Our analyses indicate that such impact is limited. On the
14 other hand, our results suggest logical dependencies and work dependencies are significantly
15 more important factors impacting the likelihood of source code files to exhibit field defects. In
16 addition, this study is the first analysis that highlights the importance of the structure of the logi-

1cal relationships – source code files with logical dependencies to other files that are also highly
2interdependent among themselves were less likely to exhibit customer-reported defects. We can
3view these groups of files as a unit where the structure of the technical dependencies in the unit
4influences its quality. These results suggest a new view of product dependencies with significant
5implications regarding how we think about modularizing the system and how development work
6is organized. The effect of the structure of the network of product dependencies elevates the idea
7of modularity in a system to the level of “clusters” of source code files. These highly inter-relat-
8ed sets of files become the relevant unit to consider when development tasks and responsibilities
9are assigned to organizational groups.

10 The second significant contribution of this study is the recognition and the assessment of the
11impact the engineers’ social network has on the software development process. Nagappan and
12colleagues [36] have examined the impact on failure proneness of structural properties of the for-
13mal organization (e.g. organizational chart). However, the informal organization which emerges
14as part of personal relationships is significantly more important for performing tasks in organiza-
15tions [30]. Similarly, Meneely et al. [31] looked at the relationship among developers based on
16file-touched network that may to some extent reflect social relationships among the developers
17that are more directly captured using workflow measures. Our measures of work dependencies
18capture the important elements of the informal organization in the context of software develop-
19ment tasks. Our results showed that individuals that exhibited a higher number of workflow de-
20pendencies and coordination requirements were more likely to have defects in the files they
21worked on. These findings suggest the difficulty of needing to receive work from or coordinate
22with multiple people and manage those relationships appropriately in order to perform the tasks.
23 This study has an additional characteristic worthy of note. The empirical analyses were repli-

located across two distinct projects from two unrelated companies obtaining consistent results. This replication provides us with unusually good external validity that is not easily achieved given proprietary concerns, etc.³ We believe this study provides a proof of concept that such analyses are possible, and given the improved external validity, we think such an approach should be adopted (wherever logistics permit) as a standard of validity for industry studies.

6 *A. Threats to Validity and Limitations*

7 First, it is important to highlight some potential concerns for construct validity, particularly regarding work dependencies. Over the years, there have been many efforts to measure task interdependencies in the context of software development. However, most of the approaches have focused on stylized representations of work dependencies, particularly in organizational studies (e.g. [10, 42]). Our study proposed two measures that capture the fine-grained dependencies that exist in software development and emerge over time as technical decisions are implemented. Certainly, there might be other potentially superior measures of development work dependencies, however, little is known about how to develop such measures.

15 Operationalization of software dependency measures is fraught with difficulties as projects produce products for different domains, using different tools and disparate practices making it difficult to design measures that capture aspects of the same phenomena across unrelated projects. Therefore, we felt it was important to replicate the entire measurement and analysis process on two unrelated projects each using different sets of tools and practices. Furthermore, we investigated the stability of the results by analyzing individual releases and using random effects models to account for potential autocorrelation.

22 The work reported in this study has several limitations. First, our analysis cannot claim causal

³ In our case, it required a strategy in which data extraction was performed on machines inside company firewalls, to ensure that only anonymized data is provided for statistical modeling.

1 effects. For example, even though dependencies in workflow are related to customer reported de-
2fects, it may be possible that the defects somehow increase the dependencies in the workflow.
3 Secondly, our results on the role of syntactic dependencies is based on two projects where the
4 software was developed in two programming languages (C and C++) that are somewhat similar
5 in terms of how technical dependencies are represented. Projects that involve programming lan-
6 guages with very distinct technical properties might exhibit a different impact of syntactic depen-
7 dencies on failure proneness.

8 *B. Applications*

9 *1) Enhancing Dependency Awareness*

10 We observed that logical dependencies were considerably more relevant than syntactic depen-
11 dencies in relation to the failure proneness of a software system. They may also be less apparent
12 to developers, since they are not as easily discovered by tracing function calls, value assign-
13 ments, or other things locally visible in the code.

14 Tools such as TUKAN [41], Palantir [39] and Ariadne [44] provide visualization and aware-
15 ness mechanisms to aid developers coordinate their work. Those tools achieve their goal by mon-
16 itoring concurrent access to software artifacts, such as source code files, and by identifying syn-
17 tactic relationships among source code files. This information is visualized to assist the develop-
18 ers in resolving potential conflicts in their development tasks. Using the measures proposed in
19 this paper, new tools or extensions to those tools could be developed to provide an additional
20 view of product dependencies using logical dependencies. These new tools would then be in a
21 position to provide complementary product dependency information to the developers which
22 could be more valuable in terms of raising awareness among developers about the potential im-
23 pact of their changes in the software system. **Moreover, since logical dependencies might be of**

1 different types such as implicit relationships (e.g. events), cascading function calls or time-relat-
2 ed relationships, tools could leverage such a categorization to provide more selective awareness
3 information for particular user needs or work contexts. Secondly, these new tools could also pro-
4 vide a more precise view of coordination needs among developers using the work dependencies
5 measures presented in this paper. For instance, the coordination requirements measure goes be-
6 yond identifying such dependencies, allowing developers to identify those files that have depen-
7 dencies among themselves when those dependencies are not explicitly determined. It is impor-
8 tant to also highlight that the development of future tools that use logical and coordination re-
9 quirements dependencies is faced with important challenges such as the identification of the
10 most relevant subset of dependencies for a particular work context and the presentation of such
11 information to improve awareness and limit “play the system” behavior. There are also some mi-
12 nor but quite relevant process related issues that require attention such as difficulty of maintain-
13 ing consistent data about modification requests and version control changes over time and auto-
14 mation of the collection and processing of the data.

15 2) *Reducing and Coping with Dependencies*

16 Once developers, architects or other relevant stakeholders become aware of particular patterns
17 of technical dependencies, they could be in a position to utilize specific techniques to reduce
18 those dependencies, in particular logical relationships. For instance, **system re-architecting** is a
19 promising technique to reduce logical dependencies and in a large system it was demonstrated to
20 relate to quality improvements [22]. Other code reorganizations techniques that make the struc-
21 ture of the systems more suitable for geographically distributed software development organiza-
22 tions could also focus their attention on logical dependencies. Such is the case of the globaliza-
23 tion by chunking approach [33] that provides a way to select tightly clustered groups of source

code files (in terms of logical dependencies) that exhibit few logical dependencies with the rest of the system. Alternatively, methods to make logical dependencies more explicit by, for example, introducing syntactic dependencies where only logical dependencies exist could be explored given the important difference between the role of logical and syntactic dependencies suggested by our results.

6 In recent years, a number of tools that either implement some of the code re-organization approaches described in the previous paragraph or provide new mechanisms for coping with technical dependencies have been proposed. For instance, tools that highlight and filter changes from different releases helping to cope with interdependencies between changes in subsequent releases have been shown to improve productivity [1]. The results of this study provide valuable information to allow this type of tools to focus on those dependencies that are most relevant.

12 3) *Guiding Future Research*

13 While it seems clear that logical dependencies play a major role in software failures, we do not yet have a clear idea of the precise nature of these dependencies. Research and practices focused on syntactic dependencies, as found in strongly typed languages for example, are likely responsible for weakening the relationship between such dependencies and fault proneness. We suggest that an emphasis on understanding the precise nature of logical dependencies is a fertile area for future research. Such research could, for example, examine the code that is changed together to understand if it represents cascading function calls, or semantic dependencies, platform evolution, or other types of relationships. A more detailed understanding of the bases of logical dependencies is an important future direction with implications in research areas such as software quality and development tools.

23 In particular, we suggested adding syntactic dependencies where logical dependencies exist.

1 This could be done, for example, in case of two implementations, by templetizing the function
 2 call. We also highlighted process-related challenges such as the difficulty of maintaining
 3 MR/VCS data consistent and available for automatic collection/processing and the host of chal-
 4 lenges associated with using logical dependencies or coordination requirements as awareness en-
 5 hancers.

6

ACKNOWLEDGMENT

7 We gratefully acknowledge support by the National Science Foundation under Grants IIS-
 8 80414698, IIS-0534656 and IGERT 9972762, the Software Industry Center at Carnegie Mellon
 9 University and its sponsors, especially the Alfred P. Sloan Foundation, and the Software Engi-
 10 neering Institute grant for “Improving Architectural Design through Organizational Considera-
 11 tions”. The authors also gratefully thank A. Hassan and R. Holt for providing the source code for
 12 their C-REX tool.

13

REFERENCES

- 14[1] Atkins, D. Ball, T., Graves, T. and Mockus, A. Using version control data to evaluate the impact of software tools: A case study of the
 15 version editor. *IEEE Trans. on Soft. Eng.*, 28, pp. 625-637, 2002.
- 16[2] Baldwin, C.Y. and Clark, K.B. *Design Rules: The Power of Modularity*. MIT Press, 2000.
- 17[3] Basili, V.R. and Perricone, B.T. Software Errors and Complexity: An Empirical Investigation. *Comm. of the ACM*, 12, pp. 42-52, 1984.
- 18[4] Briand, L.C., Wust, J., Daly, J.W. and Porter, D.V. Exploring the Relationships between Design Measures and Software Quality in Object-
 19 Oriented Systems. *The Journal of Systems and Software*, 51, pp. 245-273, 2000.
- 20[5] Burt, R.S. *Structural Holes: The Social Structure of Competition*. Harvard University Press, 1992
- 21[6] Cataldo, M., Wagstrom, P., Herbsleb, J.D. and Carley, K.M. Identification of Coordination Requirements: Implications for the Design of
 22 Collaboration and Awareness Tools. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'06)*, 2006, pp.
 23 353-362.
- 24[7] Cataldo, M. Dependencies in *Geographically Distributed Software Development: Overcoming the Limits of Modularity*. Ph.D. dissertation,
 25 Institute for Software Research, School of Computer Sciences, Carnegie Mellon University, 2007.
- 26[8] Cataldo, M., Bass, M., Herbsleb, J.D. and Bass, L. On Coordination Mechanism in Global Software Development. In *Proceedings of the*
 27 *International Conference on Global Software Engineering (ICGSE '07)*, 2007, pp. 71-80.
- 28[9] Chidamber, S.R. and Kemerer, C.F. A Metrics Suite for Object-Oriented Design. *IEEE Trans. on Soft. Eng.*, 20, pp. 476-493, 1994.
- 29[10] Crowston, K.C. *Toward a Coordination Cookbook: Recipes for Multi-Agent Action*. Ph.D. Dissertation, Sloan School of Management, MIT,
 30 1991.
- 31[11] Curtis, B., Kransner, H. and Iscoe, N. A field study of software design process for large systems. *Comm. of ACM*, 31, pp. 1268-1287, 1988.
- 32[12] de Souza, C.R.B. *On the Relationship between Software Dependencies and Coordination: Field Studies and Tool Support*. Ph.D.
 33 dissertation, Donald Bren School of Information and Computer Sciences, University of California, Irvine, 2005.

- 1[13] de Souza, C.R.B., Redmiles, D., Cheng, L., Millen, D. and Patterson, J. How a Good Software Practice Thwarts Collaboration – The
 2 multiple roles of APIs in Software Development. In *Proceedings of the Conference on Foundations of Software Engineering (FSE '04)*, pp.
 3 221-230, 2004.
- 4[14] Eaddy, M., Zimmermann, T., Sherwood, K.D., Garg, V., Murphy, G.C., Nagappan, N., Aho, A.V. 2008. Do Crosscutting Concerns Cause
 5 Defects? *IEEE Trans. on Soft. Eng.*, 34, pp. 497-515, 2008.
- 6[15] Eick, S.G., Graves, T.L., Karr, A.F., Mockus, A. and Schuster, P. Visualizing Software Changes. *IEEE Trans. on Soft. Eng.*, 28, pp. 396-
 7 412, 2002.
- 8[16] Eppinger, S.D., Whitney, D.E., Smith, R.P. and Gebala, D.A. A Model-Based Method for Organizing Tasks in Product Development.
 9 *Research in Eng. Design*, 6, pp. 1-13, 1994.
- 10[17] Faraj, S. and Xiao, Y. Coordination in Fast-Response Organization. *Management Science*, 52, 8, pp. 1155-1169, 2006
- 11[18] Fenton, N.E. and Neil, M. A Critique of Software Defect Prediction Models. *IEEE Trans. on Soft. Eng.*, 25, pp. 675-689, 1999.
- 12[19] Freeman, L.C. Centrality in Social Networks: I. Conceptual Clarification. *Social Networks*, 1, pp. 215-239, 1979.
- 13[20] Galbraith, J.R. *Designing Complex Organizations*. Addison-Wesley Publishing, 1973.
- 14[21] Gall, H. Hajek, K. and Jazayeri, M. Detection of Logical Coupling Based on Product Release History. In *Proceedings of the International
 15 Conference on Software Maintenance (ICSM '98)*, pp. 190-198, 1998.
- 16[22] Geppert, B., Mockus, A. and Rößler, F. Refactoring for changeability: A way to go? In *Proceedings of the 11th International Symposium on
 17 Software Metrics (METRIC '05)*, pp. 35-48, 2005.
- 18[23] Graves, T.L., Karr, A.F., Marron, J.S. and Siy, H. Predicting Fault Incidence Using Software Change History, *IEEE Trans. on Soft. Eng.*,
 19 26, pp. 653-661, 2000.
- 20[24] Grinter, R.E., Herbsleb, J.D. and Perry, D.E. The Geography of Coordination Dealing with Distance in R&D Work. In *Proceedings of the
 21 Conference on Supporting Group Work (GROUP '99)*, 1999, pp. 306-315.
- 22[25] Hassan, A.E. and Holt, R.C. C-REX: An Evolutionary Code Extractor for C. Presented at *CSER Meeting*, Canada, 2004.
- 23[26] Herbsleb, J.D., Mockus, A. and Roberts, J.A. Collaboration in Software Engineering Projects: A Theory of Coordination. Presented at the
 24 *International Conference on Information Systems (ICIS'06)*, 2006.
- 25[27] Herbsleb, J.D. and Mockus, A. An Empirical Study of Speed and Communication in Globally Distributed Software Development. *IEEE
 26 Trans. on Soft. Eng.*, 29, pp. 481-494, 2003.
- 27[28] Horwitz, S., Reps, T., and Binkley, D. Interprocedural slicing using dependence graphs. *ACM Trans. on Programming Languages and
 28 Systems*, 22, pp. 26-60, 1990.
- 29[29] Hutchens, D.H. and Basili, V.R. System Structure Analysis: Clustering with Data Bindings. *IEEE Trans. on Soft. Eng.*, 11, pp. 749-757,
 30 1985.
- 31[30] Krackhardt, D. and Brass, J.D. Intra-organizational Networks: The Micro Side. In pp. 207-229, 1992.
- 32[31] Meneely, A., Williams, L., Snipes, W., Osborn, J. Predicting Failures with Developer Networks and Social Network Analysis. In
 33 *Proceedings, Foundations of Software Engineering (FSE '08)*, 2008.
- 34[32] Mockus, A. and Weiss, D. Predicting risk of software changes. *Bell Labs Tech. Journal*, 5, pp. 169-180, 2000.
- 35[33] Mockus, A. and Weiss, D. Globalization by chunking: a quantitative approach. *IEEE Software*, 18, pp. 30-37, 2001.
- 36[34] Moeller, K.H. and Paulish, D. An Empirical Investigation of Software Fault Distribution. In *Proceedings of the International Software
 37 Metrics Symposium*, IEEE CS Press, pp. 82-90, 1993.
- 38[35] Nagappan, N. and Ball, T. Using Software Dependencies and Churn Metrics to Predict Field Failures: An Empirical Case Study. In
 39 *Proceedings of the 1st International Symposium on Empirical Software Engineering and Measurement (ESEM'07)*, 2007, pp. 363-373.
- 40[36] Nagappan, N., Murphy, B., Basili, V.R. The Influence of Organizational Structure on Software Quality: An Empirical Case Study. In
 41 *Proceedings of the International Conference on Software Engineering (ICSE'08)*, 2008, pp. 521-530.
- 42[37] Parnas, D.L. On the criteria to be used in decomposing systems into modules. *Comm. of ACM*, 15, pp. 1053-1058, 1972.
- 43[38] Pinzger, M., Nagappan, N., Murphy, B. Can Developer-Module Networks Predict Failures? In *Proceedings, Foundations of Software
 44 Engineering (FSE '08)*, 2008.
- 45[39] Sarma, A., Noroozi, Z. and van der Hoek, A. Palantir: Raising Awareness among Configuration Management Workspaces. In *Proceedings
 46 of the International Conference on Software Engineering (ICSE'03)*, 2003, pp. 444-453.
- 47[40] Selby, R.W. and Basili, V.R. Analyzing Error-Prone System Structure. *IEEE Trans. on Soft. Eng.*, 17, pp. 141-152, 1991.
- 48[41] Schummer, T. and Haake, J.M. Supporting Distributed Software Development by Modes of Collaboration. In *Proceedings of the European
 49 Conference on Computer-Supported Collaborative Work (ECSCW '01)*, 2001, pp. 79-89.
- 50[42] Staudenmayer, N. *Managing Multiple Interdependencies in Large Scale Software Development Projects*. Unpublished Ph.D. Dissertation,
 51 Sloan School of Management, Massachusetts Institute of Technology, 1997.
- 52[43] Stevens, W.P., Myers, G.J. and Constantine, L.L. Structure Design. *IBM Systems Journal*, 13, pp. 231-256, 1974.

- 1[44] Trainer, E., Quirk, S., de Souza, C. and Redmiles, D. Bridging the Gap between Technical and Social Dependencies with Ariadne. In
2 *Proceedings of Workshop on the Eclipse Technology Exchange*, 2005, pp. 26-30.
- 3[45] Thompson, J.D. *Organizations in Action: Social Science Bases of Administrative Theory*. McGraw-Hill, New York, 1967
- 4[46] von Hippel, E. Task Partitioning: An Innovation Process Variable. *Research Policy*, 19, pp. 407-418, 1990.
- 5[47] Watts, D.J. *Small Worlds: The Dynamics of Networks between Order and Randomness*, Princeton University Press, Princeton, NJ, 1994.
- 6[48] Zimmermann, T. and Nagappan, N. The Predicting Defects using Network Analysis on Dependency Graphs. In *Proceedings of the*
7 *International Conference on Software Engineering (ICSE'08)*, 2008, pp. 531-540.