

Comparing four approaches for technical debt identification

Nico Zazworka · Antonio Vetro^{*} · Clemente Izurieta · Sunny Wong · Yuanfang Cai · Carolyn Seaman · Forrest Shull

Published online: 3 April 2013
© Springer Science+Business Media New York 2013

Abstract Software systems accumulate *technical debt* (TD) when short-term goals in software development are traded for long-term goals (e.g., quick-and-dirty implementation to reach a release date versus a well-refactored implementation that supports the long-term health of the project). Some forms of TD accumulate over time in the form of source code that is difficult to work with and exhibits a variety of anomalies. A number of source code analysis

N. Zazworka (✉)
Elsevier Information Systems GmbH, Frankfurt am Main, Germany
e-mail: zazworka@gmail.com

N. Zazworka · A. Vetro^{*} · C. Seaman · F. Shull
Fraunhofer CESE, College Park, MD, USA

F. Shull
e-mail: fshull@fc-md.umd.edu

A. Vetro^{*}
Automatics and Informatics Department, Politecnico di Torino, Turin, Italy
e-mail: antonio.vetro@polito.it

C. Izurieta
Department of Computer Science, Montana State University, Bozeman, MT, USA
e-mail: clemente.izurieta@cs.montana.edu

S. Wong
Siemens Healthcare, Malvern, PA, USA
e-mail: sunny.wong@siemens.com

Y. Cai
Department of Computer Science, Drexel University, Philadelphia, PA, USA
e-mail: yfcai@cs.drexel.edu

C. Seaman
Department of Information Systems, UMBC, Baltimore, MD, USA
e-mail: cseaman@umbc.edu

techniques and tools have been proposed to potentially identify the code-level debt accumulated in a system. What has not yet been studied is if using multiple tools to detect TD can lead to benefits, that is, if different tools will flag the same or different source code components. Further, these techniques also lack investigation into the symptoms of TD “interest” that they lead to. To address this latter question, we also investigated whether TD, as identified by the source code analysis techniques, correlates with interest payments in the form of increased defect- and change-proneness. Comparing the results of different TD identification approaches to understand their commonalities and differences and to evaluate their relationship to indicators of future TD “interest.” We selected four different TD identification techniques (code smells, automatic static analysis issues, grime buildup, and Modularity violations) and applied them to 13 versions of the Apache Hadoop open source software project. We collected and aggregated statistical measures to investigate whether the different techniques identified TD indicators in the same or different classes and whether those classes in turn exhibited high interest (in the form of a large number of defects and higher change-proneness). The outputs of the four approaches have very little overlap and are therefore pointing to different problems in the source code. Dispersed Coupling and Modularity violations were co-located in classes with higher defect-proneness. We also observed a strong relationship between Modularity violations and change-proneness. Our main contribution is an initial overview of the TD landscape, showing that different TD techniques are loosely coupled and therefore indicate problems in different locations of the source code. Moreover, our proxy interest indicators (change- and defect-proneness) correlate with only a small subset of TD indicators.

Keywords Technical debt · Software maintenance · Software quality · Source code analysis · Modularity violations · Grime · Code smells · ASA

1 Introduction

As Cunningham described (Cunningham 1992), in a software development project, rushing implementation to meet pending deadlines is “*like going into debt. A little debt speeds development so long as it is paid back promptly with a rewrite... The danger occurs when the debt is not repaid. Every minute spent on not-quite-right code counts as interest on that debt. Entire engineering organizations can be brought to a stand-still under the debt load of an unconsolidated implementation....*”

Nowadays “technical debt” (TD) has become a well-known metaphor indicating the possibly of significant economic consequences for such quick-and-dirty implementations. Accumulated technical debt may cause severe maintenance difficulty or even project cancellation. To prevent TD from accumulating, or to decide when, where, and how to pay off such debt, the first step is to make it explicit, that is, to identify TD. One important class of TD is manifested by problematic implementations in code. Many types of such code-based TD can be potentially detected automatically using static program analysis tools that find anomalies of various kinds in the source code.

There are a number of tools designed for this purpose. Some tools are designed to detect design problems such as code smells (Schumacher et al. 2010) or Modularity violations (Wong et al. 2011), some are designed to discover design pattern degradations (Izurieta and Bieman 2007), and some are intended to spot potential defects. From a tool user’s point of view, the relevant questions are: Which tool(s) should be used to inform the existence of TD under what circumstances? and Is it sufficient to use one of the tools, or can the usage of multiple tools lead to benefits in finding more TD?

Plus, not all problematic code detected by these tools is worth being fixed. Some detected source code problems are not likely to cause future maintenance problems or affect the overall quality of the system. In terms of the TD metaphor, the *TD principal* (i.e., the cost of fixing the debt) may be higher than the *TD interest* being paid on the debt (i.e., the probable future cost of not fixing it). Thus, another question is: Which tools reveal TD that is likely to incur interest?

TD interest is inherently difficult to estimate or measure. Given the data that we had available in this study, we chose to use two proxies for expected interest (hereafter referred to as “interest indicators”): defect- and change-proneness. These proxies are concrete manifestations of problematic code and are related to future maintenance cost, and therefore are useful, independent indicators of likely interest payments. The proxy measures are well established and have been previously used in the assessment of maintenance problems (Zazworka et al. 2011; Olbrich et al. 2010; Bieman et al. 2003; Khomh et al. 2009). However, they are not the only, and possibly not the best, indicators, since they do not capture other forms of TD interest, such as increasing effort to make changes. However, defect- and change-proneness are factors of concern to practitioners, and thus, it is relevant to determine whether they correlate with the TD indicators generated by the four approaches studied here. Thus, a lack of relationship between our selected interest indicators and a TD indicator cannot clearly be interpreted to mean that the associated TD identification approach is ineffective, but that instead it may identify TD that exhibits other forms of interest.

It should be noted that our aim is not to predict either defects or change-proneness in future instances of the code base, as has been done by many other researchers (e.g., Nagappan et al. 2006; Ball 2005; Muthanna et al. 2000; Fenton et al. 2007; Riaz et al. 2009). Rather, we are calculating, for a given version of the system, which classes are already exhibiting change- or defect-proneness, and using these indicators as proxies of a construct (TD interest) that is more difficult to measure.

So as a first attempt to compare and contrast different TD identification techniques, we conducted an empirical study to answer the following research questions:

RQ1: Considering a set of four representative TD detection techniques (resulting in a set of 25 “TD indicators”), do they report the same set of modules as problematic? If not, how much overlap is there?

RQ2: To what extent do any of the techniques for detecting TD in code happen to point to classes with high interest indicators (defect-proneness or change-proneness)?

We first compare the results of applying the different TD identification techniques to 13 versions of the Apache Hadoop open source software project.

Secondly, we investigate whether and how likely the problems detected by these different techniques are related to the two interest indicators we have chosen. This study is a first attempt to map out a “TD landscape” that illustrates the overlaps, gaps, and synergies between a variety of code analysis techniques with respect to their use in identifying TD.

Both research questions are answered through the computation and combination of the pairwise relationship between TD indicators (RQ1) and TD indicators versus interest indicators (RQ2). The aim is to understand which TD and interest indicators point to the same locations (i.e., Java classes).

We use four different statistical association measures: Pearson’s correlation, conditional probability, chance agreement, and Cohen’s Kappa. Each of the selected four measures assesses association from different perspectives that complement each other (see Sect. 4.3.1 for further details).

The TD indicators and interest indicators are computed for each class and each of the 13 selected versions of the Hadoop software: hence, a further aggregation is needed to combine measures from all versions and have an unique association measure for every possible combination of TD indicators (RQ1) or TD indicator versus interest indicator (RQ2).

The results in this paper have the potential to improve our understanding about how existing source code analysis approaches can be tailored toward identifying TD. The long-term vision of this work is to create practical approaches that software developers can use to make TD visible, to help assess the principal (i.e., value) and interest (i.e., long-term cost) of the debt, and to assist managers in making educated decisions on strategies for debt retirement.

2 Related work

Past research efforts into TD have been outlined in Brown et al. (2010) and focused on building techniques that independently identify TD through source code analysis. For instance, Gat and Heintz (2011) identified TD in a customer system using both dynamic (i.e., unit testing and code coverage) and static (computing rule conformance, code complexity, duplication of code, design properties) program analysis techniques.

Nugroho et al. (2011) also performed static analysis to identify TD. They first calculated lines of code, code duplication, McCabe's cyclomatic complexity, parameter counts, and dependency counts. After that, they assigned these metrics to risk categories to quantify the amount of interest owed in the form of estimated maintainability cost.

A CAST report (CAST) also presented the usage of static analysis as a way to identify technical debt. The proposed approach examines the density of static analysis issues on security, performance, robustness, and changeability of the code. The authors built a pricing model assuming that only a percentage of the issues are actually being fixed.

The Sonar tool (<http://www.sonarsource.org/>) is an open source application that has gained in popularity. It also uses static measurements against various source code metrics and attributes to assess the level of TD in a code base.

The approaches discussed thus far calculate TD holistically, that is, they yield an overall assessment of the total TD in a system, but do not point to specific problematic parts of the code base, or to specific remedies applicable to those parts. Another approach to TD identification, that attempts to yield more actionable information, is to use source code analysis to identify potentially problematic parts of the code and to use the results of that analysis to suggest specific changes to be made to that code. Examples of such approaches that have been partly evaluated to be valid TD indicators are code smells (Fowler et al. 1999), grime build up (Izurieta and Bieman 2007, 2012), and Modularity violations (Wong et al. 2011). We discuss these techniques in more detail in Sect. 4.

This work further evolves the study of these analysis techniques by investigating the amount of similarity between them. If it turned out that many, or even all, of the TD indicators point to the same code, one could propose to choose only one of the tools when searching for TD. Alternatively, if each of the techniques selects a unique subset of problems, the usage of multiple tools can be recommended, where a particular project must make intentional decisions about which types of TD they are most interested in and choose tools accordingly. The relationships between different TD identification approaches have not previously been addressed in the literature.

3 Goals and research questions

The objective of our research is twofold: the first goal is to compare the similarities and differences between four code analysis techniques in terms of TD identification. We are interested in understanding the degree of convergence and divergence of these techniques and their associations. The second goal is to understand how the problematic code identified by these four techniques relates to our chosen proxies for TD interest, defect- and change-proneness. As explained in Sect. 1, these proxies are well established and feasible to compute given the available data, but do not constitute a complete assessment of all TD types. We define the goals of our research according to the Goal Question Metric framework (Basili and Weiss 1984).

Goal 1 Characterizing the similarities and differences in the problematic classes reported by these four different TD detection approaches, in the context of an open source software project.

Goal 2 Comparing these four TD detection approaches in terms of their correlation to one subset of interest indicators, namely change-proneness and defect-proneness, in the context of an open source project.

We deduced from the above goals the following research questions:

RQ1: Considering a set of four representative TD detection techniques (resulting in a set of 25 “TD indicators”), do they report the same set of modules as problematic? If not, how much overlap is there?

RQ2: To what extent do any of the techniques for detecting TD in code happen to point to classes with high interest indicators (defect-proneness or change-proneness)?

4 Case study

The application studied is Apache Hadoop (<http://hadoop.apache.org>). Hadoop is a software library for the distributed processing of data across numerous computer nodes, based on the map-reduce processing model. It provides two key services: reliable data storage using the Hadoop Distributed File System (HDFS) and high-performance parallel data processing using a technique called MapReduce. Data are spread and replicated differently between all the nodes of the cluster, while operations are split so that each node works on its own piece of data and then sends results into a unified whole.

We selected Hadoop because it is a mature project (it has been released 59 times starting from April 2, 2006). We focused our analysis on the Java core packages of the system (`java/org.apache.hadoop.*`), which includes the common utilities that support the other Hadoop subprojects and provides access to the file systems supported by Hadoop. We focused the analysis from release 0.2.0 to release 0.14.0 (the latest release, at the time this paper was written, is 1.0.3). The system initially had 10.5 k NCSS (non-commented source statements) and 126 Java classes and grew to 37 k NCSS and 373 Java classes by release 0.14.0.

4.1 TD identification techniques selected

We selected four main techniques for identifying technical debt in source code: Modularity violations, grime buildup, code smells, and automatic static analysis (ASA). These

approaches have all been studied in previous work by the collaborating authors and institutions (Wong et al. 2011; Zazworka et al. 2011; Vetro' et al. 2010; Bieman 2008), and we chose to select these techniques as a proof of concept because the authors are most experienced with installing, using and interpreting the results of these four approaches. This work can be extended in the future to incorporate comparisons to other techniques. This section introduces the basic concepts of our selected techniques and reports on our and other related past work.

4.1.1 Modularity violations (tool: CLIO)

In large software systems, modules represent subsystems that are typically designed to evolve independently. During software evolution, components that evolve together though belonging to distinct modules represent a discrepancy. This discrepancy may be caused by side effects of a quick-and-dirty implementation, or requirements may have changed such that the original designed architecture could not easily adapt. When such discrepancies exist, the software can deviate from its designed modular structure, which is called a *Modularity violation*. Wong et al. (2011) have demonstrated the feasibility and utility of this approach. In their experiment using Hadoop, they identified 231 Modularity violations from 490 modification requests, of which 152 (65 %) violations were conservatively confirmed by the fact that they were either indeed addressed in later versions or recognized as problems in the developers' subsequent comments.

4.1.2 Design patterns and grime buildup

Design patterns are popular for a number of reasons, including but not limited to claims of easier maintainability and flexibility of designs, reduced number of defects and faults (Guéhéneuc and Albin-Amiot 2001), and improved architectural designs. Software designs decay as systems, uses, and operational environments evolve, and decay can involve design patterns. Classes that participate in design pattern realizations accumulate grime—non-pattern-related code. Design pattern realizations can also rot, when changes break the structural or functional integrity of a design pattern. Both grime and rot represent forms of TD. Izurieta and Bieman (Izurieta and Bieman 2007) introduced the notion of design pattern grime and performed a pilot study of the effects of decay on one small part of an open source system, JRefractory. They studied a small number of pattern realizations and found that coupling increased and namespace organization became more complex due to design pattern grime, but they did not find changes that “break” the pattern (design pattern rot). Izurieta and Bieman (2008) also examined the effects of design pattern grime on the testability of JRefractory, a handful of patterns were examined, and they found that there are at least two potential mechanisms that can impact testability: (1) the appearance of design anti-patterns (Brown et al. 1998) and (2) the increases in relationships (associations, realizations, and dependencies) that in turn increase test requirements. They also found that the majority of grime buildup is attributable to increases in coupling.

4.1.3 Code smells (tool: codevizard)

The concept of code smells (aka bad smells) was first introduced by Fowler et al. (1999) and describes choices in object-oriented systems that do not comply with widely accepted principles of good object-oriented design (e.g., information hiding, encapsulation, use of

inheritance). Code smells can be roughly classified into identity, collaboration, and classification disharmonies (Lanza and Marinescu 2006). Automatic approaches (detection strategies Marinescu 2004) have been developed to identify code smells. Schumacher et al.'s (2010) research focused on evaluating these automatic approaches with respect to their precision and recall, and their other work (Zazworka et al. 2011; Olbrich et al. 2010) evaluated the relationship between code smells (e.g., God classes) and the defect- and change-proneness of software components. This work showed that automatic classifiers for God classes yield high recall and precision when studied in industrial environments. Further, in these environments, God classes were up to 13 times more likely to be affected by defects and up to seven times more change prone than their non-smelly counterparts.

4.1.4 ASA issues (tool: FindBugs)

Automatic static analysis tools analyze source or compiled code looking for violations of recommended programming practices (“issues”) that might cause faults or might degrade some dimensions of software quality (e.g., maintainability, efficiency). Some issues can be removed through refactoring to avoid future problems. Vetro et al. (2010, 2011) analyzed the issues detected by FindBugs (Pugh 2004) on two pools of similar small programs (85 and 301 programs, respectively), each of them developed by a different student. Their purpose was to examine which issues detected by FindBugs were related to real defects in the source code. By analyzing the changes and test failures in both studies, they observed that a small percentage of detected issues were related to known defects in the code. Some of the issues identified as good/bad defect detectors by the authors in these studies were also found in similar studies with FindBugs, both in industry (Pugh 2010) and open source software (Ernst 2007). Similar studies have also been conducted with other tools (Moonen 2009; Ernst 2007) and the overall finding is a small set of ASA issues is related to defects in the software, but the set depends on the context and type of the software.

4.2 Data collection

We measured each class of every Hadoop version using the 30 class-level indicators listed in Table 1. We call each of the measurements a *data point*. Of all the 30 indicators, 25 (Table 1: indicators 1–25) were calculated by the four techniques described in Sect. 2. We included one additional size metric (Table 1: metric 26) to provide a point of comparison. We also considered four software interest indicators (Table 1: metrics 27–30), which reflect defect- and change-proneness, for studying our second goal.

For our analysis, we considered 13 Hadoop releases. We ignored the very first one (0.1.0) since CLIO's Modularity violation computation is based on the current and previous versions. Across the 13 Hadoop releases, from 0.2.0 to 0.14.0, and all 30 indicators over every class, the total size of our data set was 96,720 data points. Due to limitations in the tools used for TD identification, we excluded nested classes from our analysis. To understand the threat to validity, we inspected all versions of Hadoop and found that (depending on the version) 39–45 % of all classes were nested classes. We will discuss this threat in Sect. 6.

It should be noted that the range (i.e., possible values) of each indicator varies. As shown in Table 1, TD indicators that solely express the presence of TD (e.g., the presence of a Modularity violation or a code smell on class level) map to 0 (meaning no presence) or 1 (meaning the indicator is present). This is expressed by [0,1] in Table 1. For indicators

Table 1 Indicators used in our analysis

Technical debt indicators	Modularity violations	1. Presence of modularity violation <i>CLIO</i>	[0,1]			
	Grime	2. Presence of grime	[0,1]			
		3. Absence of design pattern	[0,1]			
		Code smells <i>Codevizard</i>	Class-level code smells	4. God class	[0,1]	
	5. Brain class			[0,1]		
	6. Refused parent bequest			[0,1]		
	7. Tradition breaker			[0,1]		
	8. Feature envy			[0,1]		
	9. Data class			[0,1]		
	Method level code smells		10. Brain method	[0..N]		
			11. Intensive coupling	[0..N]		
			12. Dispersed coupling	[0..N]		
			13. Shotgun surgery	[0..N]		
			ASA issues <i>FindBugs</i>	By priority	14. High	[0..N]
					15. Medium	[0..N]
					16. Low	[0..N]
	By category	17. Bad practice		[0..N]		
		18. correctness		[0..N]		
		19. experimental		[0..N]		
		20. I18 N (internationalization)		[0..N]		
		21. Malicious code		[0..N]		
			22. Multithread (MT) correctness	[0..N]		
			23. Performance	[0..N]		
			24. Security	[0..N]		
			25. Style	[0..N]		
Other metrics	Size	26. Number of methods <i>Eclipse Metrics Plugin</i>	[0..N]			
Interest indicators	Defect-proneness	27. Number of bug fixes affecting this version	[0..N]			
		28. Number of bug fixes fixed in this version	[0..N]			
		29. Number of bug fixes counting between affected and fixed in this version	[0..N]			
	Change-proneness	30. Change likelihood	[0.0...1.0]			

that can be identified multiple times in a Java class (e.g., code smells on the method level and ASA issues that can be repeatedly detected), the measure indicates how many times the indicator was identified. Table 1 shows this as [0..N].

We measured the presence of grime as well as the absence of design patterns. Even if we cannot be sure that the absence of design patterns is harmful, we included this information to investigate whether we can find interesting relationships. Therefore, classes not following design patterns received a value of “1” for the indicator. We collected issues reported by FindBugs (version 1.3.9) from the source code of each Hadoop version, considering all issues of any FindBugs category (Table 1: 17–25) and priority (Table 1: 14–16).




Fig. 1 Three ways of computing defect-proneness

4.2.1 Defect-proneness measurement

To link classes with defects, for a bug that was fixed and closed in a version v , we computed which classes were modified during the fix change (identifiable through Subversion repository by using bug links provided in commit comments, for example, HADOOP-123). The linkage between source code anomalies and their resulting defects is potentially stretched over time. For example, as illustrated in Fig. 1, a bug can be found and reported in version 0.3.0, but may not be fixed until version 0.5.0. We thus measure the defect-proneness of a class c in version v using the following three different ways, respectively:

1. The number of times class c is involved in fixing bugs that were *injected* in version v , that is, the version where the bugs were found and reported.
2. The number of times class c is involved in fixing bugs that were *resolved* in version v .
3. The number of times class c is involved in fixing bugs that were *alive* in version v ; that is, the bugs were reported before or in version v and were resolved after or in version v .¹

4.2.2 Change-proneness measurement

Following the work of Schumacher et al. (2010), we measure the change-proneness of class c in version v as the number of repository changes affecting class c divided by the total number of changes in the repository during the class' lifetime (e.g., creation to deletion date).

4.2.3 Size measurement

We chose the Number of Methods in each class as a measure of size. In Sect. 6, we discuss possible threats posed by this choice.

4.3 Analysis methodology

In order to investigate the two research questions proposed in Sect. 3, that is, the overlap between the results generated by these TD detection techniques (TD indicators), and their correlation with defect- and change-proneness (interest indicators), we designed a 5-step

¹ This approach is necessary since versions can overlap (in time) in the SVN repository and single revisions cannot always be clearly assigned to a single version.

methodology that reduces the complex set of indicator values on the large, multidimensional data set into a graph. The methodology is illustrated in Fig. 2:

- Step 1 Compute a set of association measures to examine how each TD indicator and each interest indicator are related to each other
- Step 2 Apply statistical and significance functions to filter only highly associated pairs of indicators
- Step 3 Combine the set of three significant association measures into one measure
- Step 4 Combine measures from each of the 13 Hadoop versions to one set of aggregate measures
- Step 5 Build visualization and data tables to provide insight into the most strongly associated indicators

To answer the first research question regarding the overlap between the results reported by these techniques, we examine the association measures between their respective TD indicators. To answer the second research question regarding interest indicators, we order the TD indicators (and the one-size measure) by their level of association with the four interest indicators introduced in Sect. 4.2.

4.3.1 Step1: compute statistical association measures

We identified different statistical techniques to quantify the relationship between pairs of TD indicators in each version of Hadoop. Because there are many choices for association measurement (e.g., Pearson’s correlation, conditional probability), we performed a sensitivity analysis to investigate how the results generated from these statistic models differ from each other. This analysis showed that different statistical analysis techniques result in different answers: from the top 50 list of the most highly associated pairs of indicators and

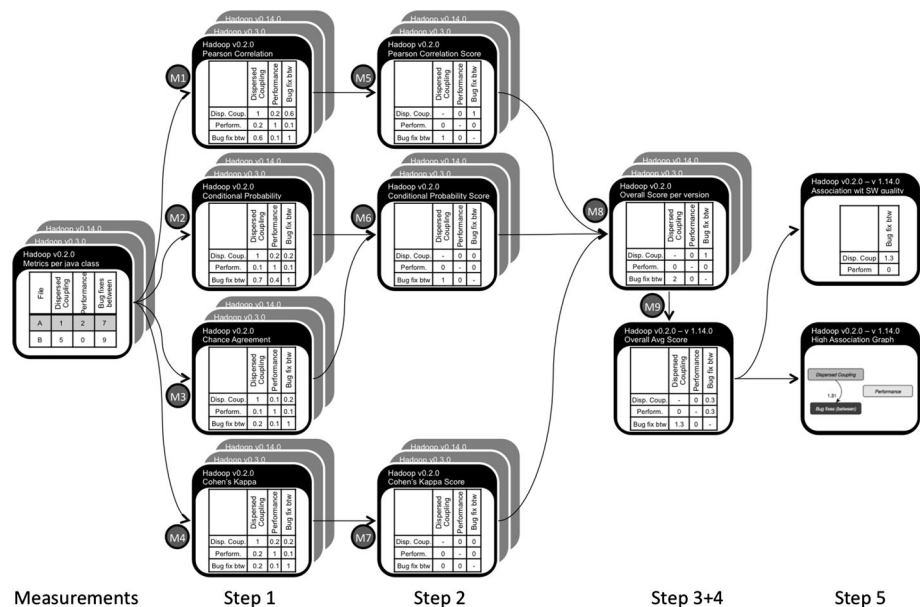


Fig. 2 Five-step analysis methodology

metrics generated by each of the statistical analysis techniques, only three pairs were common. Therefore, we concluded that using only one single measure of association would be inadequate because different statistic models assess association from different perspectives that complement each other. We thus apply 4 different association models to assess the relation between the 30×30 pairs of indicators. The different association measures will be combined in one unique measure in Step 3.

In order to illustrate our methodology, we use a pair of TD indicators, Dispersed Coupling (one type of code smell) and Performance issues (reported by FindBugs), as a running example. The four association techniques are described below:

1. The **Pearson's correlation** between the number of occurrences of Dispersed Coupling and the number of occurrences of Performance issues across all Java classes of one Hadoop version (Fig. 2: Step 1, M1). Pearson's correlation is widely used in the literature on defect prediction models (Nagappan et al. 2006), including the usage of ASA issues as an early indicator of defects (Ball 2005), and in maintainability prediction (Muthanna et al. 2000).
2. The **conditional probability** of a Java class having at least one occurrence of Dispersed Coupling given that the same class has at least one occurrence of Performance issue, and vice versa: $P(\text{Dispersed Coupling}|\text{Performance})$, $P(\text{Performance}|\text{Dispersed Coupling})$ (Fig. 2: Step 1, M2). The conditional probability has been used previously for software defects and maintainability predictions (Fenton et al. 2007; Riaz et al. 2009).
3. The **chance agreement**, which is the probability that a Java class holds an occurrence of Dispersed Coupling and Performance issue at the same time by chance. This probability is computed as: $P(\text{Performance}) * P(\text{Dispersed Coupling}) + P(\text{No Performance}) * P(\text{No Dispersed Coupling})$ (Fig. 2: Step 1, M3). Chance agreement is at the basis of Cohen's Kappa computation (Fleiss 1981).
4. **Cohen's Kappa**, which is an inter-rater agreement between Dispersed Coupling and Performance issues, indicates the strength of agreement and disagreement of two raters (e.g., TD indicators). Cohen's Kappa is appropriate for testing whether agreement exceeds chance levels for binary and nominal ratings. Therefore, Dispersed Coupling and Performance issues are in agreement if both occur at least once in the same Java class, or if both of them are not present. In any other case, they are in disagreement (Fig. 2: Step 1, M4). Cohen's Kappa has been used for the assessment of defect classification schemes (Wieczorek 1998) and in software process assessment (Jung 2003). Used in conjunction with other statistical measures, it prevents the misinterpretation of results possibly affected by prevalence and bias.

The computation of chance agreement, conditional probability, and Cohen's Kappa required data transformation for some measures: all metrics ranges $[0..N]$ were reduced to $[0,1]$, where "1" indicates that at least one occurrence of the TD indicator was found in the class and "0" otherwise. For example, for the Number of bug fixes (fixed) in version v , we assign "1" if the Java class was part of at least one defect fix during the analyzed version. The partial loss of measurement resolution is further discussed in Sect. 6.

For change likelihood, where the metric ranges from 0.0 to 1.0 (floating point numbers), we investigated its empirical distribution on each version and selected a threshold value to be used in the data transformation. This threshold was computed so as to guarantee that, on average, only the top 25 % of classes with high change likelihood obtained a value of "1" in the data transformation, "0" otherwise. The computed threshold was 0.01, indicating

that a Java class that is changed more often than in 1 out of 100 cases is considered *change prone*.

The same was done for the size metric, Number of Methods. We use a threshold of 11 to guarantee that only the top 25 % of classes were considered *large*. The computation of the four statistical analysis techniques produced four respective matrices shown in Fig. 2: M1, M2, M3, and M4. Each of the statistical measures offers a different perspective on the association between pairs of metrics.

The Pearson's correlation indicates whether two indicators increase/decrease together following a linear pattern. The conditional probability is useful in understanding the direction of the association, since it indicates whether an indicator is present in a class given that another indicator is present. The chance agreement, instead, is the probability that two indicators are either indicating or not indicating a problem in the same class randomly; in the following section, we will see how we use this measure with the conditional probability. Finally, the Cohen's Kappa is useful because the agreement takes into account not only when two indicators are present in the same class, but also when they are not present simultaneously.

4.3.2 Step 2: apply significance functions

After applying the different techniques in Step 1, the goal was to filter the more significant associations from the less significant ones. Therefore, we applied different significance functions Ω . The significance functions map the associations between each pair of indicators to $[0,1]$, expressing that the pairs are strongly associated ("1") or not associated strongly enough ("0"). The following three formulae define the functions applied to produce matrices M5, M6, and M7 for each of the 13 versions in Fig. 2:

Significant Pearson's correlations:

$$\Omega(M1_{i,j}) = \begin{cases} 1, & \text{if } (M1_{i,j}) \geq T1 \wedge p - \text{val} \leq 0.05 \wedge (i \neq j) \\ 0, & \text{otherwise} \end{cases}$$

Significant conditional probability: using M2 (Cond. Probability) and M3 (Chance agreement):

$$\Omega(M2_{i,j}, M3_{i,j}) = \begin{cases} 1, & \text{if } (M2_{i,j}) > (M3_{i,j}) \wedge (M2_{i,j}) \geq T2 \wedge (i \neq j) \\ 0, & \text{otherwise} \end{cases}$$

Significant Cohen's Kappa:

$$\Omega(M4_{i,j}) = \begin{cases} 1, & \text{if } (M4_{i,j}) > T3 \wedge (i \neq j) \\ 0, & \text{otherwise} \end{cases}$$

The goal of the significance function Ω is to discern significant relations from insignificant ones. Parameters $T1$, $T2$, and $T3$ are three specific thresholds of the respective significance functions that we chose based on the association strength levels found in the literature:

$T1$ (Correlation) = 0.60 We found two main correlation strength classifications, Cohen (Cohen 1988) and Evans (1996). We adopt Evan's strong definition because it is stricter than Cohen's definition.

$T2$ (Conditional probability) = 0.60 To our knowledge, existing literature does not provide a commonly accepted and generally applicable threshold for conditional probability. Therefore, we calculated the distribution of conditional probabilities in the different versions and we selected the threshold 0.60, which on average filtered out 80 % of all data.

Moreover, since this criterion is merged with the criterion conditional probability > chance agreement, we consider such threshold high enough to discriminate significant data.

$T3$ (Cohen's Kappa) = 0.60 Many tables of Kappa's strength of agreement can be found in the literature, the most relevant of which are (Landis and Koch 1977; Altman 1990; Fleiss 1981). We adopt a threshold value of 0.60. Thus, a constraint > 0.60 corresponds to a "good"/"substantial" agreement in all the proposed ranks.

4.3.3 Step 3: combine statistical association measures

The next step is an aggregation of the association measures. For each cell of a matrix (a pair of indicators), we compute the sum over all three matrices M5, M6, and M7. The resulting matrix (displayed as M8 in Fig. 2) contains the significance score for each pair and version that ranges from 0 (not significant in any of the three methods) to 3 (significant in all three methods).

4.3.4 Step 4: combine versions

The fourth and final computation step of the process is the aggregation of the significance score over versions. For every cell in Matrix M8, we compute the mean over the 13 versions of Hadoop, resulting in a single matrix (M9 in Fig. 2).

4.3.5 Step 5: visualize most significant outcomes

In the final step, we visualize the pairs of TD indicators with most significant associations as a graph (Fig. 3) and a list of TD indicators most related to interest indicators (Table 2).

5 Results

We made the following observations from the result. First, the value of TD indicators increases together with the size of Hadoop. The sum of all TD indicators increases in the evolution from release 0.2.0 to release 0.14.0. FindBugs issues ranged from 307 to 486 but the average number of issues per class does not expose the same monotone increasing trend, and the range of the average number of issues per class is [1.30–1.76].

Code smells in the last release (352) are more than twofold the number of code smells in the first release (143), and the average of code smells per class is [0.78–1.01]. Modularity violations have the sharpest increase: they were 8 in the first analyzed release and 37 in the last one (reaching a maximum of 38 in v. 0.13.0). The average number of Modularity violations per class ranged from 0.04 to 0.11. Moreover, we detected in each version two realizations of Singleton design pattern, two realizations of state pattern, and six of Abstract factory. However, none of the classes that are participating in these design patterns was affected by grime.

We do not observe a trend in any of the interest indicators. The sum of bug fixes collected with the defect-proneness strategy "inject" ranges from a minimum of 16 (v 0.8.0) to a maximum of 102 (v 0.3.0), while the range of the average number per class is [0.06–0.53]. Version 0.8.0 has no bug fixes (fixed), and version 0.7.0 is the version with the largest number (590). The average of all classes per version is in the range [0–2.63]. The

Table 2 Association of TD indicators with interest indicators (five strongest associations in bold)

	TD indicator	Bug fixes (between)	Bug fixes (inject)	Bug fixes (fixed)	Change likelihood
Mod viol	<i>Modularity violations</i>	1.23	0.23	0.54	1.38
Code smells	<i>God class</i>	1.00	0.23	0.77	0.85
	<i>Brain Class</i>	0.62	0.23	0.46	0.62
	<i>Tradition breaker</i>	0.69	0.31	0.38	0.69
	<i>Feature envy</i>	0.54	0.15	0.31	0.15
	<i>Brain Method</i>	0.77	0.23	0.54	0.46
	<i>Intensive Coupling</i>	0.54	0.00	0.08	0.15
	<i>Dispersed Coupling</i>	1.31	0.23	0.54	0.92
FindBugs issues	<i>Shotgun surgery</i>	0.31	0.00	0.08	0.08
	<i>High</i>	0.92	0.08	0.46	0.62
	<i>MT correctness</i>	1.00	0.08	0.46	0.69
	<i>Correctness</i>	0.92	0.15	0.46	0.62
	<i>Performance</i>	0.31	0.00	0.08	0.08
Size	<i>Style</i>	0.31	0.00	0.08	0.38
	<i>Number of Methods</i>	0.62	0.00	0.08	0.08

ranges of bug fixes (between) are [162–675] and [0.55–3.01], respectively, for their total by version and average of all classes by version. Finally, the change likelihood range is [0.006–0.017] per class.

Figure 3 shows the resulting directional graph of the TD pairs and interest indicators that were on average significant in more than one statistical measure (overall mean score in Matrix M9 >1). The nodes of the graph show the indicators. The color (or shade) of the node indicates which TD indicators were derived from the same TD detection technique. The directional edges are further labeled by their association strength (ranging from 1 to 3). And lastly, the direction indicates the conditional properties inherited from the conditional probability metric. For example, Modularity violations and bug fixes are associated, meaning that if a class has a Modularity violation, then it is also likely that such a class will have bug fixes. However, the reverse statement (i.e., classes containing bug fixes are not as likely to have Modularity violations at the same time) is not necessarily true and does not show in the graph.

Figure 3 represents the strongest associations found in our analysis. The graph contains 24 relationships (edges), one of them between TD indicators (nodes) belonging to different techniques (colors or shades), three of them between TD indicators and interest indicators or size, and the remaining 20 are between TD indicators detected by the same technique or between the interest indicators.

As for correlations between TD indicators and interest indicators, Dispersed Coupling points to classes that are more defect prone. Modularity violations do not strongly co-occur with code smells or ASA issues but are likely to point to defect- and change-prone classes.

Seven out of the twelve ASA/FindBugs issue types appear in the graph. The strongest associations (average score ≥ 2) are between style and low and bad practice and medium.




Fig. 3 Graph of top ranked pairs (average score > 1)

We also observe an association between a FindBugs issue (high) and a code smell (Intensive Coupling). Four out of ten code smells show up in the graph. Brain Class and Brain Method code smells are related in both directions, as well as Dispersed and Intensive Coupling (but only one direction).

Lastly, defect-prone classes also tend to be change-prone and vice versa. The size metric Number of Methods does not show up in the graph indicating that neither the TD indicators nor the interest indicators are very strongly associated with size.

5.1 RQ 1: Which techniques tend to report problems in the same sets of classes?

The results shown in Fig. 3 lead to our first finding in response to RQ1: Different TD techniques point to different classes and therefore to different problems.

The only arc in Fig. 3 that relates two different types of TD identification approaches is Intensive Coupling and FindBugs high priority issues. A method exhibits Intensive Coupling if it “calls too many methods from a few unrelated classes” (Lanza and Marinescu 2006). FindBugs high priority issues are those issues thought to have higher probability to detect serious problems in the code. The direction of the association indicates that classes with many high priority issues have methods affected by Intensive Coupling. A possible explanation for this relation is that both of these indicators point, more than any others, to generally poorly designed code.

Looking at the associations inside the boundaries of the techniques, we observe a characteristic of all FindBugs issues in the graph: significant relationships are only revealed between priority and type categories, which are not independent indicators and are constructed by the FindBugs authors.² Therefore, this relationship is not a surprising result. A follow-up analysis revealed that 81–87 % (depending on version) of all classes contain FindBugs issues of only one single category.

Shifting the focus of the results analysis to the code smells group, we observe three associations between particular code smells: Brain Class → Brain Method (2.0), Brain Method → Brain Class (1.15), and Dispersed Coupling → Intensive Coupling (1.23). The first relationship is stronger in the direction → Brain Method, and it indicates that Brain classes are more prone to contain Brain methods. This observation can be explained by the way Brain Class code smells are detected using Marinescu’s detection strategy (Lanza and Marinescu 2006; Marinescu 2004): the Brain Class detection requires that the inspected class contains at least one Brain Method. Therefore, the conditional probability as defined in Sect. 4 of P(Brain Method|Brain Class) is always 1.0. The second relationship between

² Each bug pattern is assigned a priority and category by the FindBugs authors. Some categories are biased toward single priorities: for example, *correctness* is considered more often to be of *high* priority.

code smells is Dispersed Coupling → Intensive Coupling (1.20). While the latter indicates that a class has methods that invoke *many* functions of a *few* other classes, the former shows classes having methods invoking functions of many other classes. Their association demonstrates that classes in Hadoop having the former of the coupling smells also have the latter smell, which intensifies the problem of coupling. No other relationship within different code smells was found in this analysis.

We also observe that Modularity violations are not strongly related to any other indicator. This confirms and validates one of the findings reported in Wong et al. (2011), who found that 40 % of Modularity violations in Hadoop are not defined as code smells and are not detectable using existing approaches.

To conclude, the 4 TD detection approaches (Modularity violation, code smells, grime, and ASA issues) have only very little overlap and are therefore pointing to different problems. Within the broad approaches, relations are stronger (as one would expect). However, the data also show that many code smells and some ASA issue types are not inter-related (i.e., the ones not showing in Fig. 3) indicating that even at a lower-level indicators point to different problem classes.

5.2 RQ2: Which TD indicators correlate with the interest indicators defect- and change-proneness?

Turning to RQ2, our major finding concerning defect-proneness is summarized as follows: **The Dispersed Coupling code smell and Modularity violations are located in the classes that are more defect prone.**

For each TD indicator (Column 1), Table 2 reports the average score obtained in matrix M9 for the association between the indicator and defect-proneness (columns 2–4) and change-proneness (column 5). TD indicators are listed in the same order as Table 1, but those with average score less than or equal to 0.3 in all associations with interest indicators are not shown.

Figure 3 shows that no single FindBugs indicator has a very strong relationship (>1) with bug fixes. However, when investigating less correlated indicators, we find the strongest FindBugs indicator to be Multithread Correctness having a borderline value of 1.0 (Table 2). This category is very specific but ties very well into the studied application; Hadoop has to deal with both distributed data storage and computations. Previous work (Vetro' et al. 2010, 2011; Wagner et al. 2005) reported that only a small percentage of FindBugs issues are actually related to bug fixes. This is supported by our results and a follow-up analysis: Multithread Correctness issues make up only 5.3 % of the total of FindBugs issues found in Hadoop.

Another strong relationship with bug fixes (in two approaches, between and fixed) involves the Dispersed Coupling code smell. In a related work (D'Ambros et al. 2010), Dispersed Coupling was highly correlated with bug fixes only when the prevalence of this smell increased during the evolution of the software. We observe a border value (1.0) also for one other code smell: the God class indicator has an average score of 1.0. Zazworka et al. (2011) reported in their previous work that in an industrial system God classes contained up to 13 times more defects than non-God classes.

The last indicator strongly related to bug fixes (between) is Modularity violations, which are located in the same classes where the more bug fixes are found (but not in all of them).

Although defect prediction is not a goal of this work, it is useful to look at measures of precision and recall to further describe the relationships we have found. We used the two

TD indicators most strongly related to bug fixes and the two border value indicators to predict, in each version, classes with at least one bug fix (strategy between). We observe in Table 3 high precision and low recall values. Each of the four indicators points out a small subset of defect-prone classes very well. When using all four indicators together recall can be raised to 0.33 by trading off some precision.

The second part of RQ2 was concerned with change-proneness. The following summarizes this result. **Modularity violations point to change-prone classes.**

Change-prone classes might indicate maintenance problems (e.g., classes that have to be changed unusually often are candidates for refactoring). As outlined in Sect. 4, we labeled on average the top 25 % most frequently changed classes as “change prone.”

The results indicate that Modularity violations are strongly related to change likelihood. Table 2 shows that the highest average scores are Modularity violations (1.38), Dispersed Coupling (0.92), and God classes (0.85). This fits expectations since all of the three approaches claim to identify maintenance problems. Modularity violations and Dispersed Coupling point to classes that have collaboration disharmonies. The God class code smell identifies classes that implement multiple responsibilities and should be refactored (e.g., split up into multiple classes).

Further, the relation between defect- and change-proneness shows that these issues are interconnected. Explanations for the phenomena can be that maintenance problems lead to less correct code, or that many quick-and-dirty bug fixes lead to less maintainable code.

Finally, we point out that a large set of TD indicators (i.e., 9 out of 25) do not show significant associations with defect- or change-proneness. This proportion suggests that these indicators either point to different classes of quality issues (e.g., FindBugs type Performance) or to none at all. Therefore, these results can be further used to tailor TD indicators toward quality attributes of interest. If one is most interested in defect- and change-proneness issues in Hadoop (or similar software), we suggest analyzing for Dispersed Coupling and Modularity violations.

6 Threats to validity

We list the threats to the validity and generalizability of the study following the structure proposed by Wohlin et al. (2000), who identify four categories: construct, internal, conclusion, and external threats.

A first conclusion threat concerns the impact of thresholds T1, T2, and T3 applied on Step 2 (Sect. 4.3.2) on the results. We documented the choice of thresholds based on values discussed and adopted in the literature, and we adopted higher values to decrease the level of uncertainty. The collection and aggregation of different statistical measures also lower the risk associated with this threat.

A further statistical point of discussion is the loss of measurement resolution in data transformation from ranges $[0..N]$ to the range $[0,1]$. The transformation was required to

Table 3 Precision and recall when predicting defect-prone classes

	Multithread correctness	Dispersed coupling	God class	Modularity violations	All 4
Avg. prec.	0.78	0.81	0.96	0.85	0.77
Avg. recall	0.13	0.12	0.11	0.21	0.33

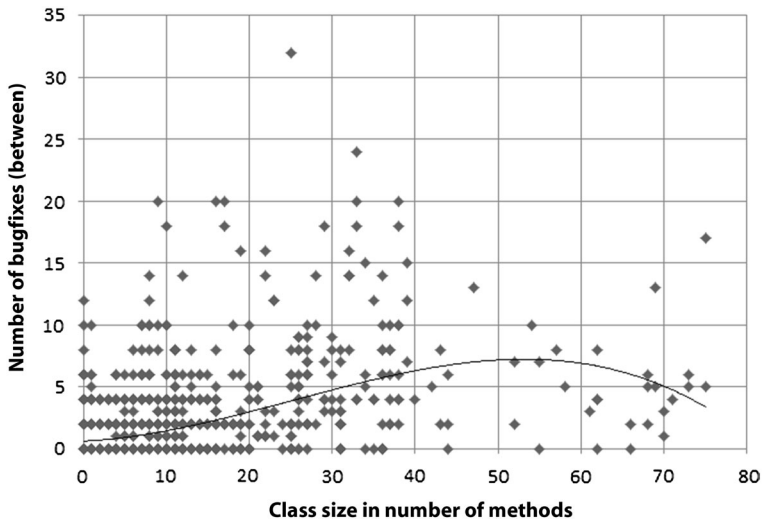


Fig. 4 Correlation plot for size versus defect-proneness

compute Cohen’s Kappa and the conditional probability. To limit this threat, we investigated distributions carefully to find reasonable thresholds (e.g., for Number of Methods we decide on a threshold of 11 for a top 25 % cutoff). Moreover, the choice of at least one occurrence as criterion for transformation was driven by a preliminary analysis of distribution that revealed that TD and interest indicators were equal to zero on average in 90 % of the classes.

Another conclusion threat is derived by the decision to not normalize measures by size. Our choice was based on past experiences (Zazworka et al. 2011) and from the analysis of the results of the current study. Figure 4 shows a correlation plot between size and number of defect fixes for each file over all versions. The 3rd order polynomial trend line shows that the correlation is not simply linear, for example, a class twice as large is not twice as defect prone. This analysis suggests that a linear normalization by Number of Methods is not required.

Moreover, the choice of this size metric rather than other size metrics is also a threat (construct). We examined whether the Total Number of Methods correlates with other size metrics in the different Hadoop versions. We obtained almost perfect correlations with Total Number of Statements, Total Number of Lines of Code (that include comments), and Total Number of Files: 0.9958, 0.9950, and 0.9711, respectively. In addition to that, we checked, for each version, the correlation between the Number of Methods and two other metrics, that is, the Number of Lines of Code and the Number of Statements in Java classes. We also obtained at this granularity very high correlation, respectively, 0.8975 and 0.8780. We conclude that using Number of Methods as measure of class size is equivalent to lines of code and did not affect results.

A further construct threat is the selection of outer classes, ignoring nested classes. At this point in time, the tools used did not allow the collection of all metrics for nested classes. Therefore, the validity scope of our results is limited to outer classes only. Future work will be devoted to include nested classes in the analysis.

We believe that the findings of this work apply only in the context analyzed (external threat). They may apply in similar applications, but we are not aware of any other published results that can be compared to ours. However, although results cannot be generalized, they contribute to begin composing the TD landscape.

7 Conclusion

The main findings of this study are as follows:

- Different TD techniques point to different classes and therefore to different problems. There are very few overlaps between the results reported by these techniques.
- Dispersed Coupling, God classes, Modularity violations, and Multithread Correctness issues are located in classes with higher defect-proneness.

Modularity violations are strongly associated with change-proneness.

Our results indicate that the issues raised by the different code analysis techniques are in different software classes. Moreover, only a subset of the problematic issue types identified by these techniques is shown to be more defective or change prone. This is consistent with the result of earlier work where these techniques were applied independently (Zazworka et al. 2011; Olbrich et al. 2010; Vetro' et al. 2010, 2011).

These findings contribute to building an initial picture of the TD landscape, where TD techniques are loosely overlapping and only a subset of them is strongly associated with software components' defect- and change-proneness.

This initial picture will contribute to future research efforts concerned about continuously monitoring and managing TD in software projects, a larger subject that is out of the scope of this study. But this study constitutes an important step in addressing the TD problem holistically: not all TD should be considered sufficiently harmful to warrant repayment, in particular if the cost of repair (paying the principal) outweighs the interest payments in long term. Thus, not all TD is bad and not all TD needs to be avoided. The results of this study build a stepping-stone for further trade off analysis studies by providing insights into how tools can help to point to TD that is worth being managed.

7.1 Implications for practice

The results indicate that, in practice, multiple TD indicators should be used instead of only one of the investigated tools. As a recommendation to practitioners, these initial results evidently show that different tools point to different problems in a code base. The use of a single tool or single indicator (e.g., a single code smells) will only in rare cases point to all important TD issues in a project. As a result, development teams need to make intentional decisions about which of the TD indicators are of most relevance to them, based on the quality goals of their project, as suggested in (Shull 2011). For example, is maintainability a priority for the team or is the system expected to be short lived? Is code performance important? Different answers to questions like these would lead to different choices for a TD detection strategy.

In the current state of research, we cannot yet give a more complete recommendation on which indicators are best for signaling specific quality shortcomings; however, our results give some preliminary advice on which indicators to start with when looking for TD related to defects and maintenance bottlenecks, namely: Modularity violations, God class, Dispersed Coupling, and Multithread Correctness issues.

7.2 Implications for research

Since results indicate that there might not be a project independent one-size-fits-all tool to detect TD, but rather a necessary tailoring process to the right subset of indicators required, future research should be concerned with investigating and showing connections between TD techniques, types of technical debt, effect and tailoring toward project-specific software quality characteristics. Future work should also investigate other TD indicators when they become available to broaden the landscape.

As more specific advice for future research directions, we recommend extending the interest indicators toward a broader range of software quality aspects, beyond defect- and change-proneness as investigated here. Further, we recommend extending this type of quantitative study with qualitative insights, for example, from practitioners that investigate whether the studied approaches point to the most important kinds of technical debt.

References

- Altman, D. G. (1990) *Practical Statistics for Medical Research (Statistics texts)*, 1st ed. Chapman & Hall/CRC, Nov. 1990. [Online]. Available: <http://www.worldcat.org/isbn/-0412276305>.
- Ayewah, N., Pugh, W. (2010) The google findbugs fixit, In *Proceedings of the 19th International Symposium on Software Testing and Analysis* (pp. 241–252), ser. ISSSTA'10. New York, NY, USA: ACM. [Online]. Available: [10.1145/-1831708.1831738](https://doi.org/10.1145/-1831708.1831738).
- Basili, V. R., & Weiss, D. M. (1984). A methodology for collecting valid software engineering data. *Software Engineering, IEEE Transactions on, SE-10*(6), 728–738.
- Izurrieta C., Bieman, J. (2008) Testing consequences of grime buildup in object oriented design patterns, In *Software Testing, Verification, and Validation, 2008 1st International Conference on*, April 2008, pp. 171–179.
- Bieman, J.M., Straw, G., Wang, H., Munger, P.W., Alexander, R.T. (2003) Design patterns and change proneness: An examination of five evolving systems, *Software Metrics Symposium, 2003. Proceedings. Ninth International*, pp. 40–49, 3–5 Sept. 2003.
- Boogerd, C., Moonen, L. (2009) Evaluating the relation between coding standard violations and faultswithin and across software versions, In *Mining Software Repositories, 2009. MSR'09. 6th IEEE International Working Conference on*, May 2009, pp. 41–50.
- Brown, N., Cai, Y., Guo, Y., Kazman, R., Kim, M., Kruchten, P., Lim, E., MacCormack, A., Nord, R., Ozkaya, I., Sangwan, R., Seaman, C., Sullivan, K., Zazworka, N. (2010) Managing technical debt in software-reliant systems. In *Proceedings of the FSE/SDP workshop on Future of Software Engineering Research* (pp. 47–52), ser. FoSER'10. New York, NY, USA: ACM, [Online]. Available: [10.1145/-1882362.1882373](https://doi.org/10.1145/-1882362.1882373).
- Brown, W. J., Malveau, R. C., Mowbray, T. J. (1998) *AntiPatterns: Refactoring software, architectures, and projects in crisis*. Wiley, Mar. 1998. [Online]. Available: <http://www.worldcat.org/isbn/-0471197130>.
- CAST, (2010) Cast worldwide application software quality study: Summary of key findings, Tech. Rep.
- Cohen, J. (1988) *Statistical power analysis for the behavioral sciences: Jacob Cohen.*, 2nd ed. Lawrence Erlbaum, Jan. 1988. [Online]. Available: <http://www.worldcat.org/isbn/-0805802835>.
- Cunningham W (1992) The wycash portfolio management system, In *Addendum to the Proceedings on Object-Oriented Programming Systems, Languages, and Applications (Addendum)*, ser. OOPSLA'92 (pp. 29–30). New York, NY: ACM. [Online]. Available: [10.1145/-157709.157715](https://doi.org/10.1145/-157709.157715).
- D'Ambros, M., Bacchelli, A., Lanza, M. (2010) On the impact of design flaws on software defects, In *Quality Software (QSIC), 2010 10th International Conference on*, July 2010 (pp. 23–31).
- El Emam, K., Wieczorek, I. (1998) The repeatability of code defect classifications, In *Software Reliability Engineering, 1998. Proceedings. The Ninth International Symposium on*, nov 1998 (pp. 322–333).
- Evans, J. (1996) *Straightforward Statistics for the Behavioral Sciences*. Brooks/Cole Pub. Co., 1996. [Online]. Available: <http://books.google.com/books?id=8Ca2AAAAIAAJ>.
- Fenton, N., Neil, M., Marsh, W., Hearty, P., Marquez, D., Krause, P., Mishra, R. (2007) Predicting software defects in varying development lifecycles using bayesian nets, *Information and Software Technology*,

- 49:(1): 32–43, Most Cited Journal Articles in Software Engineering—2000. [Online]. Available: <http://www.sciencedirect.com/-science/-article/-pii/-S0950584906001194>.
- Fleiss, J. L. (1981) *Statistical Methods for Rates and Proportions*, 2nd ed., Wiley series in probability and mathematical statistics. New York: Wiley.
- Fowler, M., Beck, K., Brant, J., Opdyke, W., & Roberts, D. (1999). *Refactoring: Improving the design of existing code* (1st ed.). Jul: Addison-Wesley Professional.
- Gat, I., Heintz, J.D., (2011) From assessment to reduction: How cutter consortium helps rein in millions of dollars in technical debt, In *Proceedings of the 2nd Workshop on Managing Technical Debt* (pp. 24–26), ser. MTD'11. New York, NY, USA: ACM. [Online]. Available: [10.1145/-1985362.1985368](https://doi.org/10.1145/-1985362.1985368).
- Guéhéneuc, Y.-G., Albin-Amiot, H. (2001) Using design patterns and constraints to automate the detection and correction of inter-class design defects, In *Proceedings of the 39th International Conference and Exhibition on Technology of Object-Oriented Languages and Systems (TOOLS39)* (p 296), ser. TOOLS'01. Washington, DC, USA: IEEE Computer Society, 2001, [Online]. Available: <http://dl.acm.org/-citation.cfm?id=882501.884740>.
- Hovemeyer, D., Pugh, W. (2004) Finding bugs is easy, *SIGPLAN Not.*, 39: 92–106. [Online]. Available: [10.1145/-1052883.1052895](https://doi.org/10.1145/-1052883.1052895).
- Izurieta, C., Bieman, J. (2007) How software designs decay: A pilot study of pattern evolution, In *Empirical Software Engineering and Measurement, 2007. ESEM 2007. First International Symposium on*, sept. 2007 (pp. 449–451).
- Izurieta, C., Bieman, J. (2012) A multiple case study of design pattern decay, grime, and rot in evolving software systems, *Springer Software Quality Journal*. February 2012, [10.1007/s11219-012-9175-x](https://doi.org/10.1007/s11219-012-9175-x).
- Khomh, F., Di Penta, M., Gueheneuc, Y.-G. (2009) An Exploratory study of the impact of code smells on software change-proneness, *Reverse Engineering, 2009. WCRE '09. 16th Working Conference on*, pp.75–84, 13–16 Oct. 2009.
- Kim, S., Ernst, M. D. (2007) Prioritizing warning categories by analyzing software history, In *Proceedings of the Fourth International Workshop on Mining Software Repositories* (p. 27), ser. MSR'07. Washington, DC, USA: IEEE Computer Society, [Online]. Available: [10.1109/-MSR.2007.26](https://doi.org/10.1109/-MSR.2007.26).
- Kim, S., Ernst, M. D. (2007) Which warnings should i fix first? In *Proceedings of the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering* (pp. 45–54), ser. ESEC-FSE'07. New York, NY, USA: ACM, [Online]. Available: [10.1145/-1287624.1287633](https://doi.org/10.1145/-1287624.1287633).
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174.
- Lanza, M., & Marinescu, R. (2006). *Object-oriented metrics in practice*. Berlin: Springer-Verlag.
- Marinescu, R. (2004). Detection strategies: Metrics-based rules for detecting design flaws. *Software Maintenance, IEEE International Conference on*, 0, 350–359.
- Muthanna, S., Kontogiannis, K., Ponnambalam, K. Stacey, B. (2000) A maintainability model for industrial software systems using design level metrics, In *Reverse Engineering, 2000. Proceedings. Seventh Working Conference on*, 2000, pp. 248–256.
- Nagappan, N., Ball, T. (2005) Static analysis tools as early indicators of pre-release defect density, In *Software Engineering, 2005. ICSE 2005. Proceedings. 27th International Conference on*, may 2005.
- Nagappan, N., Ball, T., Zeller, A. (2006) Mining metrics to predict component failures, In *Proceedings of the 28th International Conference on Software Engineering* (pp. 452–461), ser. ICSE'06. New York, NY, USA: ACM, [Online]. Available: [10.1145/-1134285.1134349](https://doi.org/10.1145/-1134285.1134349).
- Nugroho, A., Visser, J., Kuipers, T. (2011) An empirical model of technical debt and interest, In *Proceedings of the 2nd Workshop on Managing Technical Debt* (pp. 1–8), ser. MTD'11. New York, NY, USA: ACM, [Online]. Available: [10.1145/-1985362.1985364](https://doi.org/10.1145/-1985362.1985364).
- Olbrich, S. M., Cruzes, D. S., Sjöberg, D. I. K. (2010) Are all code smells harmful? a study of god classes and brain classes in the evolution of three open source systems, In *Proceedings of the 2010 IEEE International Conference on Software Maintenance* (pp. 1–10), ser. ICSM'10. Washington, DC, USA: IEEE Computer Society. [Online]. Available: [10.1109/-ICSM.2010.5609564](https://doi.org/10.1109/-ICSM.2010.5609564).
- Park, H.-M., Jung, H.-W. (2003) Evaluating interrater agreement with intraclass correlation coefficient in spice-based software process assessment, In *Quality Software, 2003. Proceedings of Third International Conference on*, Nov. 2003 (pp. 308–314).
- Riaz, M., Mendes, E., Tempero, E. (2009) A systematic review of software maintainability prediction and metrics, In *Empirical Software Engineering and Measurement, 2009. ESEM 2009. 3rd International Symposium on*, Oct. 2009 (pp. 367–377).
- Schumacher, J., Zazworka, N., Shull, F., Seaman, C., Shaw, M. (2010) Building empirical support for automated code smell detection, In *Proceedings of the 2010 ACM-IEEE International Symposium on*

- Empirical Software Engineering and Measurement* (pp. 8:1–8:10), ser. ESEM'10. New York, NY, USA: ACM. [Online]. Available: [10.1145/1852786.1852797](https://doi.org/10.1145/1852786.1852797).
- Shull, F. (2011). Perfectionists in a world of finite resources. *IEEE Software*, 28(2), 4–6.
- Vetro', A., Morisio, M., Torchiano, M. (2011) An empirical validation of findbugs issues related to defects, *IET Seminar Digests*, 2011(1):144–153. [Online]. Available: <http://link.aip.org/link/-abstract-IEEESEM/v2011/i1/p144-s1>.
- Vetro', A., Torchiano, M., Morisio, M. (2010) Assessing the precision of findbugs by mining java projects developed at a university, In *Mining Software Repositories (MSR), 2010 7th IEEE Working Conference on*, may 2010, pp. 110–113.
- Wagner, S., Jürjens, J., Koller, C., Trischberger, P. (2005) Comparing bug finding tools with reviews and tests, In *Proceedings of the International Conference on Testing of Communications Systems* (pp. 40–55).
- Wohlin, C., Runeson, P., Höst, M., Ohlsson, M. C., Regnell, B., & Wesslén, A. (2000). *Experimentation in Software Engineering: An Introduction*. Norwell, MA, USA: Kluwer Academic Publishers.
- Wong, S., Cai, Y., Kim, M., Dalton, M. (2011) Detecting software modularity violations, In *Proceedings of 33th International Conference on Software Engineering* (pp. 411–420), May 2011.
- Zazworka, N., Shaw, M. A., Shull, F., Seaman, C. (2011) Investigating the impact of design debt on software quality, In *Proceeding of the 2nd Working on Managing Technical Debt* (pp. 17–23), ser. MTD'11. New York, NY, USA: ACM, [Online]. Available: [10.1145/1985362.1985366](https://doi.org/10.1145/1985362.1985366).

Author Biographies



Nico Zazworka is a software project manager at Elsevier Information Systems. His main research interest is to investigate the applicability and performance of software development processes and to identify when and why developers stray away from them, and how tools can support developers better. He has previously worked at Fraunhofer to study this question. Besides, he is interested in research that aims at identifying and visualizing code smells, code decay, and other technical debt items through software repositories. He received his Ph.D. and M.S. in Computer Science from University of Maryland, College Park in 2010 and 2009, respectively, and his B.S. in Computer Science from University of Applied Sciences in Mannheim, Germany.



Antonio Vetro is Ph.D. candidate at Politecnico di Torino. The main goal of his doctoral activities is to evaluate how automatic static analysis can be used to improve specific attributes of software quality, related to standard quality models such as ISO/IEC 25010 and SQALE. His recent interests are focused on maintainability and performance issues in the general area of technical debt identification. The methodology he uses is that one of the empirical software engineering, and he is specialized on statistical techniques to analyze process and product data. He received a B.Sc. in Management Engineering and a M.Sc. in Computer Engineering from Politecnico di Torino. He has been PhD visitor at the Center for Experimental Software Engineering at Fraunhofer USA (College Park, MD).



Clemente Izurieta is an assistant professor of computer science at Montana State University. Dr. Izurieta has over 16 years of experience developing and architecting commercial software systems for Hewlett Packard Co. and Intel. His work involves empirical research and the evolution of software designs, as well as studies in ecological modeling and cyber infrastructures.



Sunny Wong is a software architect at Siemens Healthcare, Health Services. He received his Ph.D. in Computer Science from Drexel University. His research interests include software architecture and design modeling, mining of software engineering data, and techniques for automated modularity analysis of design structures. He is a member of the ACM and IEEE.



Yuanfang Cai is an Assistant Professor at Drexel University. She received her M.S. and Ph.D degrees in 2002 and 2006, respectively, from the University of Virginia. Dr. Cai's research areas include software evolution, software modularity, software economics, and socio-technical congruence. Her recent work focuses on assessing software design quality based on the synergy of software artifacts, evolution history, and the associated team structure.



Carolyn Seaman is an Associate Professor in Information Systems at UMBC and a Research Fellow at the Fraunhofer USA Center for Experimental Software Engineering. Her research interests include software maintenance, software metrics, empirical software engineering, and qualitative research methods. She received her Ph.D in computer science from the University of Maryland, her MS from Georgia Tech, and her BA from the College of Wooster (Ohio). She can be reached at the Department of Information Systems, UMBC, 1000 Hilltop Circle, Baltimore, MD, USA 21250.



Forrest Shull is a division director at the Fraunhofer Center for Experimental Software Engineering in Maryland, a nonprofit research and tech transfer organization, where he leads the Measurement and Knowledge Management Division. He has been a lead researcher on projects for NASA's Office of Safety and Mission Assurance, the NASA Safety Center, the US Department of Defense, DARPA, the National Science Foundation, and companies such as Motorola and Fujitsu Labs of America. He is an Adjunct Professor at the University of Maryland College Park and Editor in Chief of IEEE Software.